

Oracle® Database Appliance

Frequently Asked Questions



Release 19.22

F92148-01

March 2024

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

ORACLE®

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.

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>.

Contents

Preface

Audience	iv
Documentation Accessibility	iv
Related Documents	v
Conventions	v

1 Oracle Database Appliance Frequently Asked Questions

Data Preserving Reprovisioning	1-2
Data Preserving Reprovisioning: General	1-2
Data Preserving Reprovisioning Process	1-7
Data Preserving Reprovisioning: Prechecks	1-10
Data Preserving Reprovisioning: DB systems and KVM	1-11
Data Preserving Reprovisioning: Databases	1-13
Data Preserving Reprovisioning: Backups	1-14
Data Preserving Reprovisioning: Oracle Data Guard	1-15
Data Preserving Reprovisioning: Oracle ASR	1-16
Data Preserving Reprovisioning: Multi-User Access	1-17
Patching Oracle Database Appliance	1-18
Oracle Auto Service Request	1-24
Updating Components using the odacli update-registry Command	1-27
Micronaut on Oracle Database Appliance	1-28
Adaptive Classification and Redaction (ACR) on Oracle Database Appliance	1-34
Multi-User Access on Oracle Database Appliance	1-37
Integrated Oracle Data Guard	1-40
Oracle Database Appliance DB Systems	1-41
Application KVMs on Oracle Database Appliance	1-45
CPU Pools	1-49
Oracle Database Appliance Networks	1-50
Browser User Interface	1-54
Backup and Recovery	1-56
ODACLI Commands	1-61

Preface

Oracle Database Appliance is an optimized, prebuilt database system that is easy to deploy, operate, and manage. By integrating hardware and software, Oracle Database Appliance eliminates the complexities of nonintegrated, manually assembled solutions. Oracle Database Appliance reduces the installation and software deployment times from weeks or months to just a few hours while preventing configuration and setup errors that often result in suboptimal, hard-to-manage database environments.

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

Audience

This guide is intended for anyone who configures, maintains, or uses Oracle Database Appliance:

- System administrators
- Network administrators
- Database administrators
- Application administrators and users

This book does not include information about Oracle Database architecture, tools, management, or application development that is covered in the main body of Oracle Documentation, unless the information provided is specific to Oracle Database Appliance. Users of Oracle Database Appliance software are expected to have the same skills as users of any other Linux-based Oracle Database installations.

Documentation Accessibility

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

Access to Oracle Support

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

Related Documents

For more information about Oracle Database Appliance, go to <http://www.oracle.com/goto/oda/docs> and click the appropriate release.

For more information about using Oracle Database, go to <http://docs.oracle.com/database/> and select the database release from the menu.

For more information about Oracle Integrated Lights Out Manager 3.2, see https://docs.oracle.com/cd/E37444_01/.

For more details about other Oracle products that are mentioned in Oracle Database Appliance documentation, see the Oracle Documentation home page at <http://docs.oracle.com>.

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 the text.
<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.
# prompt	The pound (#) prompt indicates a command that is run as the root user.

1

Oracle Database Appliance Frequently Asked Questions

Use these FAQs to address your queries about setting up, deploying, patching, and managing Oracle Database Appliance.

- [Data Preserving Reprovisioning](#)
Use these FAQs to address your queries about upgrading Oracle Database Appliance using Data Preserving Reprovisioning.
- [Patching Oracle Database Appliance](#)
FAQs on patching Oracle Database Appliance X10.
- [Oracle Auto Service Request](#)
Understand general FAQs on Oracle Auto Service Request (Oracle ASR).
- [Updating Components using the `odacli update-registry` Command](#)
FAQs about the `odacli update-registry` command.
- [Micronaut on Oracle Database Appliance](#)
FAQs on Micronaut on Oracle Database Appliance.
- [Adaptive Classification and Redaction \(ACR\) on Oracle Database Appliance](#)
FAQs on ACR on Oracle Database Appliance.
- [Multi-User Access on Oracle Database Appliance](#)
FAQs on Oracle Database Appliance Multi-User Access.
- [Integrated Oracle Data Guard](#)
FAQs on Integrated Oracle Data Guard.
- [Oracle Database Appliance DB Systems](#)
FAQs on Oracle Database Appliance DB Systems.
- [Application KVMs on Oracle Database Appliance](#)
FAQs on Application KVMs on Oracle Database Appliance.
- [CPU Pools](#)
FAQs on CPU pools on Oracle Database Appliance.
- [Oracle Database Appliance Networks](#)
FAQs on Oracle Database Appliance networks.
- [Browser User Interface](#)
FAQs on Browser User Interface (BUI).
- [Backup and Recovery](#)
FAQs on backup and recovery on Oracle Database Appliance.
- [ODACLI Commands](#)
Understand FAQs on ODACLI commands.

Data Preserving Re provisioning

Use these FAQs to address your queries about upgrading Oracle Database Appliance using Data Preserving Re provisioning.

- [Data Preserving Re provisioning: General](#)
Understand general FAQs on Data Preserving Re provisioning.
- [Data Preserving Re provisioning Process](#)
Understand FAQs on the Data Preserving Re provisioning process.
- [Data Preserving Re provisioning: Prechecks](#)
Understand FAQs on prechecks in Data Preserving Re provisioning.
- [Data Preserving Re provisioning: DB systems and KVM](#)
Understand FAQs on updating DB systems and KVM in the Data Preserving Re provisioning process.
- [Data Preserving Re provisioning: Databases](#)
Understand FAQs on updating databases in the Data Preserving Re provisioning process.
- [Data Preserving Re provisioning: Backups](#)
Understand FAQs about backups during the Data Preserving Re provisioning process.
- [Data Preserving Re provisioning: Oracle Data Guard](#)
Understand FAQs on Oracle Data Guard updates during the Data Preserving Re provisioning process.
- [Data Preserving Re provisioning: Oracle ASR](#)
Understand FAQs on Oracle ASR updates during the Data Preserving Re provisioning process.
- [Data Preserving Re provisioning: Multi-User Access](#)
Understand FAQs on multi-user access in the Data Preserving Re provisioning process.

Data Preserving Re provisioning: General

Understand general FAQs on Data Preserving Re provisioning.

My Oracle Database Appliance deployment is running Oracle Database Appliance release 18.8. Can I directly upgrade my system to Oracle Linux 8 and Oracle Database Appliance release 19.22?

No. You must upgrade the appliance, that is on a release earlier than 18.8 to either Oracle Database Appliance release 19.18 or later first, using the steps described in the Upgrading Oracle Database Appliance Using Data Preserving Re provisioning chapter in the *Oracle Database Appliance Deployment and User's Guide* in the Oracle Database Appliance release 19.20 documentation library. After you upgrade the system to Oracle Database release 19.18 or later which runs Oracle Linux 7, you can upgrade to Oracle Linux 8.

What is the typical time required for an upgrade using Data Preserving Re provisioning?

The total outage time depends on several factors such as system configuration and the number of resources such as database and DB systems created. On an Oracle Database Appliance X9-2-HA with 24 SSD disks and 32 cores enabled and an X9-2M single-node system, with three Oracle Database 19c databases on the bare metal system and two Oracle Database Appliance release 19.18 DB systems, the time to upgrade the bare metal system and DB system from 19.18 to 19.22 is about 4 hours for high-availability system and approximately 2 hours 30 minutes for the single node system.

Why is Oracle Database Appliance upgrading to Oracle Linux 8 using Data Preserving Re provisioning instead of doing an in-place upgrade?

The Linux operating system upgrade utility, Leapp, is a framework to perform in-place upgrade of Oracle Linux 7 systems to Oracle Linux 8. Leapp runs a set of preupgrade checks on the system to be upgraded before running the upgrade. These preupgrade checks impose specific requirements on the system. For example, Leapp expects access to YUM repository to update to the latest Oracle Linux 7 release before performing the actual upgrade, clearing of locked packages, disabling of secure boot and FIPS mode, kernel on Unbreakable Enterprise Kernel 6, unmounting NFS exports, and several others. Oracle Database Appliance, being a secured engineered system, employs design choices that would require extending the Leapp software to integrate with the Oracle Database Appliance software stack. For example, Oracle Database Appliance uses a very specific set of RPMs locks RPM package versions, does not assume access to YUM repositories, has FIPS support enabled, is not on Unbreakable Enterprise Kernel 6, and exports Oracle Database Appliance repository to DB systems through NFS. In addition, each Oracle Database Appliance system may have a different set of Leapp issues (also called inhibitors) which makes it difficult to support a consistent upgrade experience for all users. Conversely, with the clean install approach that Data Preserving Re provisioning employs, you experience a uniform upgrade while continuing to support the appliance principle of Oracle Database Appliance.

For information about LEAP, see the topic *Upgrading Systems With Leapp* at <https://docs.oracle.com/en/operating-systems/oracle-linux/8/leapp/leapp-AboutLeapp.html>

What are minimum supported versions for the databases before upgrading Oracle Database Appliance bare metal system and DB system to Oracle Linux 8?

The databases must be at the following minimum versions at the minimum to upgrade to Oracle Linux 8:

- **19.18.0.0.230117 or later:** Databases can be restarted in Oracle Linux 8 environment.
- **21.8.0.0.221018 (DB system only):** DB systems with 21c databases can be restarted on Oracle Linux 8 bare metal Oracle Linux 7. DB systems with 21c databases can be upgraded to Oracle Linux 8 DB systems and retain 21c databases.

My deployment is OVM and I am on Oracle Database Appliance release 19.13. Can I upgrade my system to Oracle Linux 8 using the Data Preserving Re provisioning process?

No. You must upgrade the appliance to either Oracle Database Appliance release 19.18 or later first, using the steps described in the *Migrating from Virtualized Platform to KVM on Oracle Database Appliance* chapter in the *Oracle Database Appliance Deployment and User's Guide* in the Oracle Database Appliance release 19.20 documentation library. After you upgrade the system to Oracle Database release 19.18 or later which has DCS software and Oracle Linux 7, you can upgrade to Oracle Linux 8.

Can I use ODABR when upgrading my system from Oracle Linux 7 to Oracle Linux 8?

See *My Oracle Support Note* ODA (Oracle Database Appliance): ODABR a System Backup/Restore Utility at <https://support.oracle.com/rs?type=doc&id=2466177.1>. The note states that ODABR may not be useful for systems upgrading from Oracle Linux 7 to Oracle Linux 8. The primary use of ODABR is to create LVM snapshots and backup data blocks that are modified while patching the system. Data Preserving Reprovisioning orchestrates an operating system upgrade of the Oracle Database Appliance system from Oracle Linux 7 to Oracle Linux 8 and involves a reimage of the system with an Oracle Linux 8 ISO image. During the reimage, the LVM is reinitialized and the snapshots on the boot drive are lost. The database homes and database files on the shared storage are not lost due to this reimage.

How much time does upgrading Oracle Database Appliance bare metal systems using Data Preserving Reprovisioning take?

The time taken for upgrading Oracle Database Appliance X9-2-HA is as follows:

Table 1-1 Time Taken for Patching Oracle Database Appliance Bare Metal Systems

Component being updated	Time taken for the update (h:mm:ss)
Infrastructure update of bare metal systems	<ul style="list-style-type: none"> Update DCS admin: 0:00:12 Update DCS components: 0:07:04 Update DCS Agent: 0:02:24
Infrastructure update of DB systems	<ul style="list-style-type: none"> Update DCS admin: 0:00:20 Update DCS components: 0:04:40 Update DCS Agent: 0:08:00
Pre-upgrade Report	0:04:16
Detach node operation	0:27:14
Re-image with Oracle Linux 8 ISO image, configure first network, update repository with server zip, and update repository with Oracle Grid Infrastructure clone	1:30:00
Restore of Oracle Grid Infrastructure, databases, Oracle KVM, and DB systems	<ul style="list-style-type: none"> Oracle Grid Infrastructure restore: 0:42:45 Databases restore: 0:05:20 KVM restore: 0:00:24 DB system restore: 0:06:55
Upgrade of DB system	<ul style="list-style-type: none"> Create pre-upgrade report: 0:01:37 Upgrade DB system: 0:47:46
Total time taken	approximately 04:00:30

The databases in my deployment are not on a supported release. How do I proceed to upgrade my bare metal system to Oracle Linux 8?

If the databases are at one of the following Release Update (RU) or later, it is recommended that you upgrade the databases to any of the following database versions before you run the `odacli detach-node` command.

- 18.14.0.0.210420
- 12.2.0.1.220118

- 12.1.0.2.220719
- 11.2.0.4.210119

What customizations are not restored when upgrading using the Data Preserving Reprovisioning process?

Customizations made to user profiles, kernel parameters, additional RPMs installed, or any additional services created on the databases are not restored during an upgrade using Data Preserving Reprovisioning.

I have additional operating system RPMs installed on my system. Are these restored as part of Data Preserving Reprovisioning?

No. Since the Data Preserving Reprovisioning process involves reimaging using the latest Oracle Database Appliance ISO image, any previous customizations in operating system RPMs, or any customizations in the settings or configurations of services in the operating system are lost. You must uninstall these RPMs and then continue with the next step in the upgrade process. You can reinstall these custom RPMs as required, after the upgrade.

I have configured VLANs on my system. Will these VLANs be configured on the upgraded system?

Yes.

Do I need to take backups of my databases before running the `odacli detach-node` command?

The database files are left intact on Oracle ASM disk groups and are used for bringing up the databases after successful reprovisioning. You must, however, take database backups as a safety measure.

I have a 2-JBOD configuration. Is Data Preserving Reprovisioning supported with this configuration?

Yes.

I have additional options (for example, ASO) enabled on my source databases. Are these options preserved?

If the source database home is on the root file system, then the options that are part of the Oracle software are not preserved. As part of the `odacli restore-node -d` command process, a new database home is created with the default option. If the source database homes are on Oracle ACFS file systems, then the options are preserved because the database home are not deleted.

Which are the supported Oracle Database Appliance releases as the source versions for Data Preserving Reprovisioning?

The source versions must be at the minimum on Oracle Database Appliance release 19.17. The DCS software and Oracle Grid Infrastructure must be on Oracle Database Appliance release 19.17 on the bare metal system. The same requirement exists for DB systems running on the bare metal systems. Note that Oracle Database Appliance release 19.20.0.1 is not a valid source version, since it already runs on Oracle Linux 8.

What happens to my database audit data after the reimage process in Data Preserving Re provisioning?

The audit data of databases can be either in database tables or in operating system files or in `SYSLLOG`. Only the data in database tables is preserved. You must save the audit data that is present on the root file system before initiating reimage of the system.

Can I run `cleanup.pl` after detaching the node on source ODA version?

No. Do not run **`cleanup.pl`** at any time during the upgrade process. This utility erases the Oracle ASM disk groups and then you cannot restore the databases. If there is a failure in running the command `odacli restore-node -g`, then you can use **`cleanup.pl`** to clean up the system while preserving Oracle ASM disk groups. Run **`cleanup.pl`** in DPR mode and do not use the `-nodpr` option.

Are all files present in database homes preserved and restored?

Local database homes are deleted during the reimage. Any files on these database homes (such as TDE wallets) are not preserved.

Can database homes be created on the root file system after re provisioning to Oracle Database Appliance release 19.21?

No, database homes can only be created on Oracle ACFS.

How much space is required for database homes and database clones on Oracle ACFS?

For database home on Oracle ACFS volumes, the database home is preserved and all files in these homes are preserved. For database clones on Oracle ACFS, a total of 150 GB is required. The space required for database homes is 15 GB multiplied by the number of database homes to be restored. The space availability is also validated at the source by the `odacli create-preupgradereport prechecks`.

Should I unpack Oracle Grid Infrastructure and Oracle Database clones together?

No, it is recommended that only the Oracle Grid Infrastructure clone is unpacked before running the command `odacli restore-node -g`. After the command `odacli restore-node -g` runs successfully, the Oracle Grid Infrastructure clone is moved to Oracle ACFS which has 150 GB of space.

Do the nodes restart during any of the operations?

The command `odacli restore-node -g` does not restart the nodes. When you run the `odacli update-server` command after the reimage, the node restarts.

Why does cleanup with `-erasedata` or `-f` fail after re provisioning using Data Preserving Re provisioning?

On Data Preserving Re provisioning environments, the default cleanup mode is the Data Preserving Re provisioning mode. This mode cleans up the appliance without erasing Oracle ASM disks to allow the command `odacli restore-node -g` to be reattempted. This mode is enabled after the command `odacli update-repository`

completes the step to generate the server archives. This mode does not support usage of `-erasedata` or `-f` options.

Override this behavior and force regular cleanup with the `-nodpr` option as follows:

```
cleanup.pl -nodpr -erasedata / -f
```

When is the Data Preserving Re provisioning tab displayed in the BUI ?

The Data Preserving Re provisioning tab is displayed on the left-hand side navigation list in the BUI only after the DCS infrastructure is upgraded to Oracle Database Appliance release 19.21.

Why are some tabs disabled in the Data Preserving Re provisioning screens in the BUI?

The Data Preserving Re provisioning process comprises of the following five tabs:

1. Re-provision
2. Restore Databases
3. Restore VM Instances
4. Restore DB Systems
5. Upgrade DB Systems

The Re-provision tab is always enabled, but tabs 2 to 5 are enabled only after the appliance is reimaged to Oracle Database Appliance release 19.21.

Can I use BUI to reimage the appliance?

No, you must reimage the appliance using the Oracle ILOM console. After you detach the nodes, the BUI prompts you to save a copy of the server archive file to a location outside of the appliance, and then proceed with the reimage.

Are there any steps to be completed, if I have LACP enabled on the network interface?

You must run the `odacli configure-firstnet` command to enable LACP after reimaging the appliance.

Related Topics

- Troubleshooting Oracle Database Appliance

Data Preserving Re provisioning Process

Understand FAQs on the Data Preserving Re provisioning process.

What are the server archive files referred in the Data Preserving Re provisioning process?

The Data Preserving Re provisioning process uses a method where the system is reimaged with the ISO image of the latest Oracle Database Appliance release. The metadata about the resources that were created, such as databases, application KVMs, DB systems, Oracle Data Guard, Oracle ASR configuration, and others are first saved in an archive file. This metadata is used for re provisioning the system after the reimage step. The saved metadata files are archived into a zip file with the name `serverarchive_hostname.zip`, where *hostname* is the

name of the Oracle Database Appliance system being re provisioned. This saved metadata file is called server archive file. You must copy this file outside the appliance before the reimage step in the Data Preserving Re provisioning process.

What if I lose the server archive files?

The command `odacli restore-node` cannot run without the server archive files. The server archive files are generated when running the `odacli detach-node` command and contain metadata that is required to re provision the system. The server archive file is generated and stored at `/opt/oracle/oak/restore/out/`. The `/opt` file system is deleted during the re-image step of the Data Preserving Re provisioning process. The server archive file must be saved before starting the reimage step. Failing to do so will result in not being able to re provision the system and Oracle Database Appliance will not be able to restart the resources such as database, DB systems, Oracle Data Guard, and others.

What are the files present in the server archive file?

Server archive files encapsulate metadata for Oracle Database Appliance resources, including databases, KVMs, DB systems, Oracle Data Guard, multi-user access, and Oracle ASR configuration. Additionally, these files encompass Oracle ASM and Oracle ACFS metadata, including volumes and file systems, as well as root file system files for both the database and Oracle Grid Infrastructure. It also includes other crucial files necessary for the seamless restoration of the system.

How many zip files are generated as server archive files?

One server archive file is generated in the `/opt/oracle/oak/restore/out` directory. A sha256 checksum file is also generated for the server archive file. Example:

```
# ls -l /opt/oracle/oak/restore/out/*.zip
-rw-r--r-- 1 root root 1454460 Dec 6 22:49
/opt/oracle/oak/restore/out/serverarchive_oda412.zip
```

How do I verify if the server archive file was copied successfully out of the appliance?

After copying the server archive file outside the appliance using the `scp` command, verify the file sha256sum checksum with the generated checksum file. The SHA256 checksum of outputs from both command must match. Example:

```
$ sha256sum serverarchive_oda6m003.zip
be7a517283e1aac89cb124fe68bc8e02b991a50ee336332af2e9fb06f753dd6f
serverarchive_oda6m003.zip
$ cat serverarchive_oda6m003.zip.sha256
be7a517283e1aac89cb124fe68bc8e02b991a50ee336332af2e9fb06f753dd6f
```

Are server archive files generated for DB system upgrade also?

Yes. The server archive zip file is generated during the detach phase of the upgrade-dbsystem process and then copied to the bare metal system. After creation of Oracle Linux 8 DB system, this file is copied to the DB system, and the database is restored. For each DB system being upgraded, the storage location of its server archive zip on the bare metal system is based on its name and is, therefore, unique.

I encountered an error when running the `odacli restore-node` command. Can I modify the saved metadata?

No. Contact *My Oracle Support* for assistance in correcting any erroneous metadata.

What are the prerequisites for Oracle Grid Infrastructure restore?

Ensure system Oracle ILOM and firmware are updated. Also, update the repository with the saved server archive file and Oracle Grid Infrastructure clone file before starting the Oracle Grid Infrastructure restore process.

Can I use an earlier release of Oracle Grid Infrastructure clone file, for example 19.20 after reimaging with 19.21 ISO image, before running the `odacli restore-node -g` command?

No, the Oracle Grid Infrastructure clone file must be of the same version as the DCS software.

Is a JSON input file required while running the `odacli restore-node -g` command, similar to the `odacli create-appliance` command?

No, a JSON input file is not required while running the `odacli restore-node -g` command. The Data Preserving Reprovisioning process collects the metadata during the detach of the system and uses the same metadata to reprovision the system. Ensure that you update the repository with the saved server archive file before running the `odacli restore-node -g` command.

Oracle Grid Infrastructure restore failed. How can I proceed?

To address the issue, review the error correlation report and the DCS agent logs. Address the reported problems and then run the `cleanup.pl` script on all nodes without using the `-f`, `-erasedata`, or `-nodpr` options. Subsequently, retry the Oracle Grid Infrastructure restore process.

Can I run the `odacli restore-node -d` command if there was a failure in an earlier run?

Yes. Before a rerun, the cause of the failure needs to be determined and fixed. Then rerun the command.

Are there any changes to be aware of after the system is reimaged with the Oracle Linux 8 ISO image?

If your appliance was reimaged with an Oracle Database Appliance release 19.12 or later ISO image, and you are now on a supported Oracle Database Appliance release, then you will not see any difference in the partition sizes. If your system was originally reimaged with an earlier release ISO image, then the partition sizes for `/u01` and `/opt` are smaller after your system is reimaged with the Oracle Linux 8 ISO image.

```
/u01 40GB 100GB
/opt 30GB 60GB
```

Related Topics

- [Troubleshooting Oracle Database Appliance](#)

Data Preserving Re provisioning: Prechecks

Understand FAQs on prechecks in Data Preserving Re provisioning.

Why should I run `odacli update-dcsadmin`, `odacli update-dcscomponents`, and `odacli update-dcsagent` commands before starting the Data Preserving Re provisioning process?

New commands and pre-upgrade checks required for Data Preserving Re provisioning are available only in the DCS software 19.21. Only after the DCS software stack is updated successfully on the bare metal system and on each DB system using the `odacli update-dcsadmin`, `odacli update-dcscomponents`, and `odacli update-dcsagent` commands, will the system be ready to run the required mandatory prechecks and initiate the detach-node process.

I ran the `odacli update-dcsadmin`, `odacli update-dcscomponents`, and `odacli update-dcsagent` commands successfully, but Zookeeper and MySQL versions have not changed. Is this expected?

Yes, this is expected. Zookeeper is no longer a part of the Oracle Database Appliance Server Patch and hence, it will not be updated. MySQL is also not updated since the reimage step erases MySQL and reinstalls it as part of the reimage.

Should I run the `odacli update-server` and `odacli update-storage` commands also before starting Data Preserving Re provisioning process by running the `odacli detach-node` command?

No. Updating Oracle Grid Infrastructure or the storage firmware is not required, if the system meets the minimum software requirements. The Oracle Grid Infrastructure software released with 19.21 or later is supported on Oracle Linux 8. However, you must run these commands after the reimage step in the Data Preserving Re provisioning process.

What is the best practice for running Data Preserving Re provisioning pre-upgrade checks? Can they be run outside the maintenance window?

It is recommended to have the pre-upgrade checks run as early as possible (before the start of the maintenance window). It is necessary to update the DCS software to 19.21 or later, to run the prechecks. Running the prechecks will not cause any outages to Oracle Grid Infrastructure and Oracle databases. The DCS infrastructure can be on a later version in the interim period, that is, until the start of the maintenance window.

Will the `odaadmin` password change after the re provisioning step?

Yes, the server archive files do not store passwords. Hence, a new password will be requested during the Oracle Grid Infrastructure re provisioning if the `odaadmin` user does not exist. You can specify the password that was set for the `odaadmin` user before the reimage step.

Is running the `odacli create-preupgradereport` command mandatory for upgrading Oracle Database Appliance to Oracle Linux 8?

Yes, the `odacli create-preupgradereport` command runs all the necessary prechecks to evaluate the system readiness for Oracle Linux 8 upgrade. If any errors

are found, then address these errors before you begin the detach-node process to ensure that the upgrade is smooth.

What are the different types of status results for prechecks that run as part of the `odacli create-preupgradereport` command?

Prechecks can have one of four statuses: SUCCESS, FAILED, ALERT, and ABORTED.

- **SUCCESS:** Indicates that the check ran without issues, requiring no user intervention.
- **FAILED:** Implies that users must inspect the system and address reported errors.
- **ALERT:** Indicates that users must take specific manual actions, as these cannot be handled by the Data Preserving Re provisioning process.
- **ABORTED:** Indicates that the check was not run due to an internal error, and the user should retry the command to rerun all aborted prechecks.

Can I run the `odacli detach-node` command even if the preupgrade report indicates failure?

You must review all the errors reported by the report before proceeding to run the `odacli detach-node` command. The errors pointed out can result in failure later if they are not resolved. To ensure smooth upgrade process, you must attend to all the errors reported.

There are alert messages in the report. They indicate actions that you must run to avoid loss of some files. Pay attention to these alerts and complete the suggested action.

Is running `odacli create-preupgradereport` mandatory for upgrading DB systems to Oracle Linux 8?

No. However, it is recommended that you run the `odacli create-preupgradereport -dbs dbsystem_name` command prior to upgrading a DB system. The errors displayed in the report can result in failure later if they are not resolved.

Can I run `odacli create-preupgradereport` for DB systems when the operating system is still Oracle Linux 7?

Yes. However, the DCS agent version of the DB system and the bare metal system must be release 19.21 or later.

Related Topics

- Troubleshooting Oracle Database Appliance

Data Preserving Re provisioning: DB systems and KVM

Understand FAQs on updating DB systems and KVM in the Data Preserving Re provisioning process.

Is DB system restored as an Oracle Linux 7 or Oracle Linux 8 environment?

After completing re provisioning of the bare metal system with Oracle Linux 8, the command `odacli restore-node -dbs` restarts all DB systems. The DB systems will run Oracle Linux 7. You must run the command `odacli upgrade-dbsystem` to upgrade the DB system from Oracle Linux 7 to Oracle Linux 8 and to upgrade the Oracle Grid Infrastructure to release 19.21 or 21.8.

Is there an order to follow when running the commands, `odacli restore-node -d`, `odacli restore-node -kvm`, and `odacli restore-node -dbs`?

Use the `odacli restore-node -d` command to restore database deployments in the bare metal system. If there is a shared CPU pool or a custom vnetwork associated with at least one DB system, then run the `odacli restore-node -kvm` command first before running the `odacli restore-node -db` command.

What are the resources restored during the restore KVM process?

CPU pools (except BM cpupools, which are restored during the restore database task), VM storages, vdisks, vnetworks, and VMs are restored during the restore KVM process.

Can I run the `odacli restore-kvm` command again?

Yes, if the restore kvm job fails, then run the command again. The restore task will skip the resources that were previously restored.

What state (online/offline) will the VMs be in after the restore operation?

VMs will be in the same state that they were before the detach-node operation.

Can I upgrade the operating system of the DB system but not change the Oracle Grid Infrastructure inside the DB system?

No. The Oracle Grid Infrastructure version is upgraded to the latest supported Oracle Grid Infrastructure version associated with the DB system image. For example, for Oracle Database Appliance release 19.21, the Oracle Grid Infrastructure supported versions are 19.21.0.0.231017 and 21.8.0.0.221018 respectively for 19c and 21c DB systems.

Can I upgrade different DB systems in parallel?

No. You must upgrade the DB systems in a sequential manner.

Can I run the command `odacli upgrade-dbsystem` again, if the command fails to complete?

Yes. If the command fails, then check the output of `odacli describe-job -i upgrade-dbsystem_job_id -ecr`. Take the appropriate action to resolve the reported error. Run the command again.

Is the database inside a DB system also updated when the DB system is upgraded?

No. The database continues to be of the same version as it was before the detach-node operation.

Related Topics

- Troubleshooting Oracle Database Appliance

Data Preserving Re provisioning: Databases

Understand FAQs on updating databases in the Data Preserving Re provisioning process.

Are there any database homes on the local file system after the Data Preserving Re provisioning process completes?

No. Database homes on the local file system are deleted when the system is reimaged.

Are there any database customizations that I need to be aware of when using Data Preserving Re provisioning?

Any customizations made to the database home, for example, SQL*Plus login file are lost during the Data Preserving Re provisioning process. It is the users responsibility to backup your customizations before the node is reimaged. These changes will then need to be manually applied after the `odacli restore-node -d job` completes successfully.

Will the TDE wallets, that is, Password Protected Wallet (ewallet.p12) and Autologin Wallet (cwallet.sso) be lost after the Data Preserving Re provisioning process completes ?

No, since TDE wallets are placed in DATA location and the Data Preserving Re provisioning process does not erase the contents of the DATA location, the TDE wallets are not lost after the Data Preserving Re provisioning process completes.

Is it required to reset the database initialization parameters `WALLET_ROOT` and `TDE_CONFIGURATION` for TDE databases after the Data Preserving Re provisioning process completes?

No. Since the database initialization parameters `WALLET_ROOT` and `TDE_CONFIGURATION` are not deleted after the Data Preserving Re provisioning process completes, resetting these parameter values after DPR completes is not required..

Why do I need to update the repository with database clones again?

Database homes existing on the local disk, before the reimage, are wiped out. These need to be created once again on the shared database home file system. You will be prompted to update the repository, if the clones are not available.

Are my customizations made to the database home also restored?

No. You must save any file that may have been customized, before the reimage. The customization must be manually applied after the `odacli restore-node -d job` completes successfully.

What is the type of database for which the `precheck Validate TDE wallet presence` is applicable?

The `precheck Validate TDE wallet presence` is applicable to TDE-enabled database whose `TDE wallet management` attribute is set to the value `EXTERNAL`.

If TDE wallet management is `EXTERNAL`, will the wallet be saved and restored by the Data Preserving Re provisioning process?

No, TDE wallets are not saved and restored by the Data Preserving Re provisioning process. The TDE wallets (`ewallet.p12` and `cwallet.sso`) must be copied by the user to the `dbhome/`

`admin/db_uniquename/wallets` location and restored to the earlier location after the Data Preserving Re provisioning process completes successfully.

The precheck TDE wallet presence check expects both password-protected TDE wallet (`ewallet.p12`) and Autologin TDE wallet (`cwallet.sso`) to be present at the `dbhome/admin/db_uniquename/wallets` location. Is it necessary to copy both the TDE wallets?

Yes, both password-protected TDE and Autologin wallets (`ewallet.p12` and `cwallet.sso`) must be copied to the `dbhome/admin/db_uniquename/wallets` location.

Should the TDE parameters such as `SQLNET.ENCRYPTION_WALLET_LOCATION` or `WALLET_ROOT` that were used to configure TDE be updated to point to the copied TDE wallets at the `dbhome/admin/db_uniquename/wallets` location?

No, the TDE parameters need not be updated. Once the node is restored, both the password-protected TDE and Autologin wallets (`ewallet.p12` and `cwallet.sso`) can be copied to the same location that was present before node restore. None of the TDE parameters are lost during node restore. Copy the TDE wallets to the `dbhome/admin/db_uniquename/wallets` location as a precautionary measure so that TDE wallets are not lost accidentally during node restore.

After the Data Preserving Re provisioning process completes and the TDE wallets are copied to their previous location, can the TDE wallets at the `dbhome/admin/db_uniquename/wallets` location be deleted?

Yes, the TDE wallets at the `dbhome/admin/db_uniquename/wallets` location can be deleted after the Data Preserving Re provisioning process completes, and the TDE wallets are copied to their original location.

Is it recommended to change the TDE wallet management of the database from 'EXTERNAL' to 'ODA' after upgrading to Oracle Database Appliance release 19.21? If yes, how can I change TDE wallet management?

Yes, it is recommended to change the TDE wallet management of the database from `EXTERNAL` to `ODA` after upgrading to Oracle Database Appliance release 19.21. When you make this change, Oracle Database Appliance then handles all TDE wallet management activities such as backup, restore, re-key and change of password.

Use the `--change-tdewallet-mgmt,-ctm` option in the `odacli modify-database` command to set TDE wallet management.

Related Topics

- Troubleshooting Oracle Database Appliance

Data Preserving Re provisioning: Backups

Understand FAQs about backups during the Data Preserving Re provisioning process.

Some database backup reports that were present before the reimage step are missing. Is this expected?

Yes, after restoring the node, the obsolete backup reports corresponding to Regular backups and Longterm backups are deleted. Hence some backup reports would not be present after restoring the databases. Look for the following log entries in the `/opt/`

oracle/dcs/log/dcs-agent.log file which specifies the actual number of Regular and Longterm Backup reports that were deleted.

c.o.d.a.r.s.b.DbBkupOperations: Number of long term backup reports deleted:
c.o.d.a.r.s.b.DbBkupOperations: Number of regular backup reports deleted:

Are TDE backup reports also deleted similar to the database backup reports?

No, TDE backup reports are not deleted. For each database, there is only one TDE backup report corresponding to each backup destination where the backup was taken, for example, if a TDE-enabled database backup was taken on NFS and Object Store backup destination, then there is one TDE backup report corresponding to NFS destination and one TDE backup report corresponding to Object Store destination.

After restoring the node, will I see the backup configurations with TDE backup location restored?

Yes, backup configurations with TDE backup location are restored to the same state as before the node restore process.

Related Topics

- Troubleshooting Oracle Database Appliance

Data Preserving Re provisioning: Oracle Data Guard

Understand FAQs on Oracle Data Guard updates during the Data Preserving Re provisioning process.

When does the Validate Dataguard status check run during the Data Preserving Re provisioning process?

This precheck is run only when there are databases with Oracle Data Guard configuration in your deployment. There is a separate precheck for every Oracle Data Guard configuration. This check verifies that Oracle Data Guard is configured correctly. If there are any warnings or errors, it displays the warnings or errors under Message column.

Does the precheck display multiple errors or warnings for Oracle Data Guard configuration?

Yes, it displays multiple errors or warnings for the Oracle Data Guard configuration. You just fix these errors and warnings before proceeding to the detach-node step.

Is there a separate metadata file for Oracle Data Guard configuration?

Yes, there is a separate metadata file for each Oracle Data Guard configuration which is later used to restore the Oracle Data Guard configuration.

Does restoration of Oracle Data Guard also restore Oracle Data Guard services?

Yes, Oracle Data Guard related services are restored after the Data Preserving Re provisioning process completes successfully.

Related Topics

- Troubleshooting Oracle Database Appliance

Data Preserving Re provisioning: Oracle ASR

Understand FAQs on Oracle ASR updates during the Data Preserving Re provisioning process.

Does Data Preserving Re provisioning restore Oracle ASR configuration?

Yes, the `odacli restore-node -g` step in the Data Preserving Re provisioning process restores Oracle ASR configuration.

I did not have Oracle ASR configured in Oracle Database Appliance release 19.20. Will running the `odacli restore-node -g` step configure Oracle ASR?

No, the `odacli restore-node -g` step restores Oracle ASR only if it was configured before running the `odacli detach-node` command.

I forgot the password for Oracle ASR user name. Can I skip Oracle ASR configuration during the `odacli restore-node -g` step?

Yes, specify the `--skip-asr/-sa` option in the `odacli restore-node -g` command to skip the restoration of Oracle ASR configuration.

I gave a wrong password for Oracle ASR user name when prompted by the `odacli restore-node -g` step. What do I do now?

Oracle ASR is not restored if a wrong password is provided when prompted in the `odacli restore-node -g` step. Hence, after the `odacli restore-node -g job` completes, you must configure Oracle ASR manually with the correct password using the `odacli configure-asr` command to avoid loss of service.

My Oracle Database Appliance assets are registered with an external Oracle ASR Manager. How should I upgrade my system?

Before upgrading your Oracle Database Appliance with assets registered with an external Oracle ASR Manager, that is, an external Oracle ASR configuration type, you must first upgrade your appliance where Oracle ASR Manager runs. After the upgrade, before you run the `odacli restore-node -g` step on the current appliance, ensure that you run the `odacli update-repository` with the Oracle ASR configuration zip file generated when you run the `odacli export-asrconfig` command on the Oracle ASR manager appliance.

I have external assets registered with Oracle ASR Manager. Are they registered when Oracle ASR configuration is restored in the `odacli restore-node -g` step?

No, external assets are not registered when Oracle ASR Manager is restored by the `odacli restore-node -g` step. For example, if Oracle ASR Manager is running on `oda1` and the assets of `oda2` are registered with Oracle ASR Manager on `oda1`, the Data Preserving Re provisioning process does not register the assets of `oda2` when restoring Oracle ASR configuration on `oda1`. To register these assets again, you must upgrade `oda2` as well using the Data Preserving Re provisioning process after `oda1` is upgraded.

Can I change my Oracle ASR configuration to external after completing the Data Preserving Re provisioning process?

Yes, Oracle ASR configuration type can be changed to external after completing the Data Preserving Re provisioning process. Delete the current restored Oracle ASR configuration using the `odacli delete-asr` command, and configure Oracle ASR configuration as external using the `odacli configure-asr` command.

Related Topics

- Troubleshooting Oracle Database Appliance

Data Preserving Re provisioning: Multi-User Access

Understand FAQs on multi-user access in the Data Preserving Re provisioning process.

Is it possible to convert a non-multi user access environment to a multi-user access deployment during Data Preserving Re provisioning?

No. During Data Preserving Re provisioning process, the Oracle Database Appliance deployment is restored to exactly the same state and configuration as was the case prior to the detach process. Hence, a non-multi user access environment continues to function in a similar manner even after the Data Preserving Re provisioning process is complete. Similarly, a multi user access-enabled environment is restored with the same configuration.

There is an error about User ID (UID) or Group ID (GID) conflict while performing restore node for Oracle Grid Infrastructure. How do I resolve this issue?

In a multi user access-enabled environment, all users must be restored with the same UID/GID as the ones prior to detach. Hence, post reimage of the system, if there are any conflicts, the restoration of Oracle Grid Infrastructure may fail. Manually resolve this conflict and then run the command again.

Do I need to remember old passwords for users to run ODA CLI commands after the Data Preserving Re-provisioning process?

For system users such as `odaadmin`, `oracle` and `grid`, you are prompted to enter a new password of your choice. These are preactivated accounts, and you can log into the appliance with the new password, either through ODA CLI commands or through the BUI. For custom users, Oracle Database Appliance accounts are created as locked-out accounts. You can connect to the appliance using SSH with your previous password. When you run the first ODA CLI command, you are prompted to unlock your account. After you unlock your account and reset to the new password in the process, you can log into the BUI or run ODA CLI commands using the new password.

Related Topics

- Troubleshooting Oracle Database Appliance

Patching Oracle Database Appliance

FAQs on patching Oracle Database Appliance X10.

What are the patching paths supported on Oracle Database Appliance?

Oracle Database Appliance supports patching from release N-4 to release N, where N is latest Oracle Database Appliance release.

For this release, on Oracle Database Appliance X9-2 and earlier, you can upgrade from 19.18, 19.19, 19.20, or 19.21 using Data Preserving Re provisioning. For Oracle Database Appliance X10, you can patch your deployment from Oracle Database Appliance release 19.20.0.1 and 19.21 to 19.22.

Oracle recommends that your appliance is on an Oracle Database Appliance release that is N-4 or later. If you want to patch from releases earlier than N-4 to N, then ensure that you try it on a test system first.

Are databases patched to Oracle Database release 19.22 as part of the Data Preserving Re provisioning process?

No. Data Preserving Re provisioning only re provisions the system and restarts all resources such as databases and DB systems on Oracle Linux 8. You must then run database patching commands such as `odacli update-database` after the re provisioning steps are complete.

What are the database upgrade paths supported on Oracle Database Appliance?

Oracle Database Appliance supports all upgrade paths supported by Oracle Database.

How much time does patching Oracle Database Appliance X10 bare metal systems take?

Oracle Database Appliance patching involves multiple steps, as documented in the *Oracle Database Appliance Deployment and User's Guide* for your hardware model. Listed below are times for bare metal system patching on Oracle Database Appliance release 19.22, when patching from 19.21 to 19.22. Note that patching times vary according to hardware, source and target version software, and configuration details such as the number of databases and database homes to be patched, patching of storage and firmware, and so on. Account for software download and running the command `odacli update-repository` times, which depend on your network bandwidth and the target system. If you do not run prechecks in advance, when estimating the maintenance window, you must also account for the time to run prechecks and corrective actions, if any, based on any precheck failures. The number of databases affects the time to patch a single database home, and patching multiple database homes takes proportionally longer.

Table 1-2 Time Taken for Patching Oracle Database Appliance X10 Bare Metal Systems

Component Being Patched	Time taken for patching X10-HA	Time taken for patching X10-S, X10-L
Patching DCS stack (running the commands <code>odacli update-dcsadmin</code> , <code>odacli update-dcscomponents</code> , <code>odacli update-dcsagent</code>)	approximately 9-12 minutes to complete all three commands	approximately 6-8 minutes to complete all three commands
Server prechecks (running the command <code>odacli create-prepatchreport -s</code>)	approximately 10-12 minutes	approximately 5-8 minutes
Server patching - updating operating system, Oracle Grid Infrastructure, Oracle ILOM, and firmware (running the command <code>odacli update-server</code>)	approximately 35-45 minutes for each node	approximately 30-40 minutes
Oracle Database prechecks for a single database in an Oracle home (running the command <code>odacli create-prepatchreport --dbhome</code>)	approximately 20-25 minutes	approximately 6-10 minutes
Patching Oracle Database 19c for a single database in an Oracle home (running the command <code>odacli update-dbhome</code>)	approximately 10-20 minutes	approximately 6-10 minutes

How much time does patching Oracle Database Appliance DB systems take?

Oracle Database Appliance patching involves multiple steps, as documented in the *Oracle Database Appliance Deployment and User's Guide* for your hardware model. Listed below are times for bare metal system patching on Oracle Database Appliance release 19.22, when patching from 19.21 to 19.22.

Table 1-3 Time Taken for Patching Oracle Database Appliance DB Systems

Component Being Patched	Time taken for patching X8-2-HA	Time taken for patching X8-2M
Patching DCS stack (running the commands <code>odacli update-dcsadmin</code> , <code>odacli update-dcscomponents</code> , <code>odacli update-dcsagent</code>)	approximately 10-15 minutes to complete all three commands	approximately 4-7 minutes to complete all three commands
Server prechecks (running the command <code>odacli create-prepatchreport -s</code>)	approximately 17-25 minutes	approximately 9-13 minutes
Server patching - updating operating system, Oracle Grid Infrastructure, Oracle ILOM, and firmware (running the command <code>odacli update-server</code>)	approximately 25-40 minutes for both nodes	approximately 30-40 minutes

Table 1-3 (Cont.) Time Taken for Patching Oracle Database Appliance DB Systems

Component Being Patched	Time taken for patching X8-2-HA	Time taken for patching X8-2M
Oracle Database prechecks for a single database in an Oracle home (running the command <code>odacli create-prepatchreport --dbhome</code>)	approximately 20-25 minutes	approximately 10-15 minutes
Patching Oracle Database 19c for a single database in an Oracle home (running the command <code>odacli update-dbhome</code>)	approximately 20-25 minutes	approximately 10-15 minutes

When should patching prechecks be run?

To run the latest prechecks, ensure the DCS admin, DCS components, and the DCS agent are updated to the latest Oracle Database Appliance release to which you want to patch your appliance.

It is recommended that you run patching prechecks in advance, outside the patching or maintenance window, and take corrective actions for precheck failures proactively.

Should I always run Oracle Database Appliance prechecks before patching?

Yes. When you run prechecks with the command `odacli create-prepatchreport`, checks specific to Oracle Database Appliance and ORAchk are run. Prechecks must be run at least once before running the command `odacli update-server` or `odacli update-dbhome`, otherwise these commands fail with an error. It is strongly recommended that you run patching prechecks in advance, outside the patching or maintenance window, and take corrective actions for precheck failures proactively.

Can I run ORAchk directly on Oracle Database Appliance?

Yes. Refer to the *Troubleshooting Oracle Database Appliance* chapter in the *Oracle Database Appliance Deployment and User's Guide* for your hardware model.

Do Oracle Database Appliance patching steps remain the same over releases?

No, the patching steps may change when new features are introduced or for other reasons. Always refer to the documentation for an Oracle Database Appliance release before you patch your deployment.

Why must the order of patching commands be strictly maintained?

Oracle Database Appliance patching comprises running multiple steps. These steps encompass patching of the DCS stack, Server, Storage, and Database. The DCS stack, Server, and Storage constitute the infrastructure, which must be updated before updating the databases. With every release, the new version of orchestration software, DCS, is updated and made capable of running new prechecks and handling differences, if any, for the new Oracle Grid Infrastructure and database versions. New DCS software also contains support for any new features in a Oracle Database Appliance release. Hence, DCS stack must be updated first, followed by the rest of infrastructure components that fall under the Server bucket. Strict order must be

followed within the DCS stack too. The DCS admin orchestrates `odacli update-dcscomponents` and `odacli update-dcsagent` commands, and hence you must patch the DCS admin with the `odacli update-dcsadmin` command. Since the DCS Agent has a dependency on ZooKeeper, which is updated when running `odacli update-dcscomponents`, the DCS components must be updated next. And lastly, DCS Agent must be updated using `odacli update-dcsagent`.

What are the components updated when running each patching command on Oracle Database Appliance?

The command `odacli update-dcsadmin` updates the DCS admin component.

The command `odacli update-dcscomponents` updates the DCS agent, DCS CLI, DCS Controller, Apache ZooKeeper and MySQL components.

The command `odacli update-dcsagent` updates Oracle Autonomous Health Framework (AHF).

The command `odacli update-server` updates OAK, operating system, local disk's firmware, Oracle ILOM firmware, and Oracle Grid Infrastructure.

The command `odacli update-dbhome` updates the database home and the databases running out of that Oracle home.

The command `odacli update-storage` updates shared storage (such as controller, disks, and expander) firmware.

Is server update rolling on Oracle Database Appliance?

For certain Oracle Database Appliance releases, server update supports the `--local` option for the command `odacli update-server` to patch one node at a time and maintain availability of at least one database at all times. This option is currently available.

Is database update rolling on Oracle Database Appliance?

For certain Oracle Database Appliance releases, database update supports the `--local` option for the command `odacli update-dbhome` to patch one node at a time and maintain availability of at least one database at all times. This option is currently available.

How do I know whether local server patching is available in the current release?

Check the prepatch report. It indicates with an alert whether local patching is available.

How do I know whether local database patching is available in the current release?

Check the prepatch report for the database home to be patched. It indicates with an alert whether local patching is available.

Under what conditions are Oracle Database Appliance storage patching commands rolling on an high-availability system?

The command `odacli update-storage -r` can be run only when shared controller or expander update is not required. When shared controller and expander patching is done, the Oracle Grid Infrastructure stack needs to be brought down on both nodes and both nodes are rebooted at the same time, hence rolling patching is not supported. Oracle Database Appliance release 19.13 onwards, storage patching prechecks, generated with the command `odacli create-prepatchreport -s` includes a check to indicate whether rolling patching option `-r` can be used in the `odacli update-storage` command.

Can `odacli update-server` be rerun if it fails?

Failure in running the `odacli update-server` command usually requires intervention from Oracle Support to run manual steps to fix the issues, before you can rerun the `odacli update-server` command.

Can `odacli update-dbhome` be rerun if it fails?

Yes. If a failed `odacli update-dbhome` command leaves behind a newly-created database home, it can be deleted using the `odacli delete-dbhome` command.

Can `odacli update-dcscomponents`, `odacli update-dcscomponents`, and `odacli update-dcsagent` commands be rerun if they fail?

Yes, these commands can be rerun if they fail.

Is applying one-off patches for Oracle Grid Infrastructure and database supported on Oracle Database Appliance?

Yes. However, these patches may need to be manually rolled back when applying the next Oracle Database Appliance patch bundle. Rollback is required when the one-off patches are not present in the Oracle Database Appliance patch bundle being applied. In such cases, another one-off patch on the new base release must be requested and applied.

Can the one-off patches be applied in a rolling manner?

If an Oracle Grid Infrastructure patch is a rolling patch, it can be applied in rolling manner. Oracle Database patches cannot be applied in a rolling manner because database homes on Oracle Database Appliance are on a shared Oracle ACFS file system starting with Oracle Database Appliance release 19.11.

Can a newer RU be applied on database homes if RU is not available through an Oracle Database Appliance Patch Bundle yet?

Yes, it is possible, though the recommended way to patch database homes is to update them as part of the quarterly Oracle Database Appliance release. Refer to the topic *Applying Out-of-Cycle Database Patches* in the *Oracle Database Appliance Deployment and User's Guide* for your hardware model for a release.

Just as is supported for Oracle Database on Oracle Database Appliance, can a new Oracle Grid Infrastructure RU be applied out of cycle?

No.

Can I update the kernel to apply security fixes?

You can use Oracle Ksplice to update the kernel. Refer to the topic *Applying Additional Patches and Updates on Bare Metal Systems* in the *Oracle Database Appliance Deployment and User's Guide* for your hardware model for a release.

Should I always take ODABR snapshots before patching?

It is an Oracle Database Appliance best practice and is highly recommended that Oracle Database Appliance snapshots be taken before patching the bare metal system.

Should I delete ODABR snapshots after patching completes successfully?

ODABR snapshots consume space on the boot disk and could also slow down steady-state operations. Besides, they have little utility once patching completes successfully and the system is back in the steady state. Hence, it is recommended that ODABR snapshots be deleted after patching completes successfully.

What is the impact of installing additional operating system RPMs on Oracle Database Appliance patching?

During Oracle Database Appliance server patching, newer versions of RPMs that Oracle Database Appliance ships are installed. If the additional RPMs installed on your system have dependencies on the Oracle Database Appliance RPMs, attempts to install the newer versions during patching may create conflicts, resulting in a patching failure. In such cases, you may need to uninstall additional RPMs, complete the patching, and reinstall the RPMs after patching succeeds.

What is different about patching Oracle Database Appliance starting with Oracle Database Appliance release 19.11?

Starting with Oracle Database Appliance release 19.11, Oracle Database Appliance follows the out-of-place patching model for Oracle Grid Infrastructure and database home patching. A new software home is created when patching Oracle Grid Infrastructure (with the command `odacli update-server`) and database homes (with the command `odacli update-dbhome`).

Why is the Oracle Database Appliance Server Patch for ODACLI/DCS Stack file much smaller in Oracle Database Appliance release 19.11 and later?

Starting 19.11, ODA follows the out of place patching model, which uses clone files during patching to create a new software (GI or database) home. Since GI and database patches (RUs) are no longer shipped on ODA, the "Oracle Database Appliance Server Patch for ODACLI/DCS Stack" file is smaller.

Are the new database homes always created on an Oracle ACFS file system after patching to Oracle Database Appliance release 19.11?

Yes.

Is there an option to keep database homes on /u01 in Oracle Database Appliance release 19.11 and later?

No.

Can I remove the old Oracle Grid Infrastructure home after successful patching to a newer release?

Yes. Refer to Oracle Support Document 2537389.1: ODA Administration: How to Manually Remove a 12.1 Grid Home on the ODA After an Upgrade to 12.2 or Higher at: <https://support.oracle.com/rs?type=doc&id=2537389.1>.

How can I free up space before patching?

To remove old patches and database clones, run the command `odacli cleanup-patchrepo`. To remove log files, run the command `odacli create-logcleanjob`.

Can I delete the old database clones and patch components after patching?

Yes. Run the command `odacli cleanup-patchrepo`.

Does the OJVM patch need to be installed separately?

See the following My Oracle Support Notes for more information:

- Oracle Support Document 1929745.1: Oracle Recommended Patches -- "Oracle JavaVM Component Database PSU and Update" (OJVM PSU and OJVM Update) Patches at: <https://support.oracle.com/rs?type=doc&id=1929745.1>.
- Oracle Support Document 2802019.1: Transparent rolling updates of OJVM components - Oracle Database 21c at <https://support.oracle.com/rs?type=doc&id=2802019.1>.

How can I update Oracle ILOM and the BIOS manually in case of Oracle ILOM and the BIOS patching failures?

Refer to Oracle Support Document 1427885.1: ODA (Oracle Database Appliance): OAK Bundle Patch failing on ILOM/BIOS component apply) at <https://support.oracle.com/rs?type=doc&id=1427885.1>.

Does Oracle Database Appliance patch KVM and OVM guest VMs?

No. You must patch them separately.

Is the prepatch report generated for the source database home with the `odacli create-prepatchreport -d dbhomeId -v version` command valid for running the `odacli update-database` command?

Yes, if the source and destination database homes are the same when you create the prepatch report and when you patch the database. Otherwise, the patch request is rejected with the error message that the corresponding prepatch report does not exist.

Can I use the `odacli update-database` command to resume from the failure point in a previously failed `odacli update-dbhome` command job?

Yes, `odacli update-database` command job can resume from the failure point for both `odacli update-database` and `odacli update-dbhome` commands.

Related Topics

- Patching Oracle Database Appliance

Oracle Auto Service Request

Understand general FAQs on Oracle Auto Service Request (Oracle ASR).

Can I configure SNMP to send alerts to Oracle ASR in Oracle Database Appliance release 19.21?

No, SNMP cannot be configured in Oracle Database Appliance release 19.21 with Oracle ASR. Starting with 19.21, ASR is configured to receive fault alerts using an XML payload over the HTTPS network of Oracle ASR Manager.

Where are the XML payload templates for Oracle ASR present?

The XML payload templates are present in the `/opt/oracle/asr/asrXMLTemplates/` directory. These templates are used to create an XML payload which is populated with the required fault information and is sent to the Oracle ASR Manager for further processing. Hence, these templates must not be modified.

How is the HTTPS network of Oracle ASR Manager used?

Starting with Oracle Database Appliance release 19.21, the HTTPS network is used, instead of the HTTP network. Hence, to enable HTTPS connection between Oracle ASR Manager and Oracle Database Appliance, the HTTPS certificate of Oracle ASR manager is imported into the keystore of the DCS agent and the keystore of Java used by Oracle ASR, and the certificate of the DCS agent is imported into the keystore of Oracle ASR Manager.

How do I verify if Oracle ASR is configured correctly?

Oracle ASR configuration can be verified using the `odacli test-asr` command. This command tests the connection between Oracle ASR Manager, and the Oracle Database Appliance host, and Oracle ILOM. When you run the command, an email is sent to the registered email ID, if the test is successful. Following is the command output for a successful run:

Job details

```
-----
ID: 3a3eca9e-94c8-4081-be9f-2a65962f6acf
Description: Test ASR
Status: Success
Created: September 5, 2023 1:43:52 PM CST
Message:
```

Task Name	Start Time
End Time	Status
Test ASR	September 5, 2023 1:43:52 PM CST
September 5, 2023 1:43:52 PM CST	Success

How do I list the assets configured with Oracle ASR Manager?

List the assets configured with Oracle ASR Manager with the `/opt/asrmanager/bin/asr list_asset` command.

```
[root@node1 ~]# /opt/asrmanager/bin/asr list_asset
IP ADDRESS  HOST_NAME          SERIAL_NUMBER PARENT_SERIAL ASR ASR_STATUS
PROTOCOL SOURCE   LAST_HEARTBEAT      PRODUCT_NAME
-----
1605NM10JF scaoda6s003-ilom 1605NM10JF          Y Active
HTTP      ILOM_XML 2023-12-08 07:24:13.703 ORACLE SERVER X8-2
10.128.5.125 scaoda6s003      1605NM10JF          Y Active
HTTP      ESSW_XML NA              ORACLE SERVER X8-2
```

Where are the logs generated by Oracle ASR stored?

The logs generated by Oracle ASR are stored in the `/var/opt/asrmanager/log/` directory.

What is the difference between Oracle ASR of type internal and external ?

With Internal Oracle ASR, Oracle Database Appliance assets are registered with an Oracle ASR Manager installed and running on the same appliance. With external Oracle ASR, Oracle Database Appliance assets are registered with an Oracle ASR Manager that runs on a remote appliance.

How do I configure external Oracle ASR in Oracle Database Appliance release 19.21?

See the chapter *Configuring and Using Oracle Auto Service Request* in the *Oracle Database Appliance Deployment and User's Guide* for your hardware model.

Can I list assets if Oracle ASR is configured as external?

No, assets cannot be listed if Oracle ASR is configured as external on Oracle Database Appliance because Oracle ASR Manager is not installed on the same appliance.

Can I modify the current Oracle ASR configuration?

Yes, you can modify Oracle ASR configuration using the `odacli modify-asr` command for internal Oracle ASR configuration type. You cannot modify external Oracle ASR configuration with the same command because Oracle ASR is configured with an external Oracle ASR Manager. To modify external Oracle ASR configuration, you must delete the current Oracle ASR configuration with the `odacli delete-asr` command and configure it again with the `odacli configure-asr` command after updating the repository with the modified Oracle ASR configuration zip file. For the complete procedure, see the chapter *Configuring and Using Oracle Auto Service Request* in the *Oracle Database Appliance Deployment and User's Guide* for your hardware model.

What happens to the Oracle ASR configuration while patching Oracle Database Appliance from release 19.20.0.1 to 19.21?

In Oracle Database Appliance release 19.20.0.1, Oracle ASR is configured with SNMP protocol for fault alerts. In release 19.21, Oracle ASR is configured with an XML payload for fault alerts. Hence, if Oracle ASR was configured before patching, it must be reconfigured after updating the DCS components but before you run the `odacli update-server` command. This is required to ensure that Oracle ASR is configured with the new changes before the server is updated. To reconfigure Oracle ASR, delete it using the `odacli delete-asr` command and run the `odacli configure-asr` command with the same parameter values as before.

Why is the DCS agent restarted while deleting Oracle ASR?

While configuring Oracle ASR, the HTTPS certificate of Oracle ASR Manager is imported into the Java keystore of the Java installed on the appliance because the Java binaries are used by Oracle ASR Manager. While deleting Oracle ASR configuration, the imported certificate is removed from the Java keystore. Hence, after removing the certificate, DCS agent must be restarted to reload the keystore.

Related Topics

- Troubleshooting Oracle Database Appliance

Updating Components using the `odacli update-registry` Command

FAQs about the `odacli update-registry` command.

What are the components affected when specifying the `system` option with `--component,-n` in the `odacli update-registry` command?

The `system` components are `sysinstance`, `gihome`, and `dbnode`. The metadata related to these component is updated in appliance registry metadata when you specify the `system` option with `--component,-n` in the `odacli update-registry` command.

What are the components are affected when specifying the `all` option with `--component,-n` in the `odacli update-registry` command?

When you specify `all`, the components updated are `DB`, `dbstorage`, `dbhome`, and `asr`. The metadata related to these component is updated in appliance registry metadata when you specify the `all` option with `--component,-n` in the `odacli update-registry` command.

How does the `odacli update-registry` command work?

The `odacli update-registry` command attempts to discover values for different parameters corresponding to a component using comands such as `srvctl`, `crsctl`, `sqlplus`, or operating system commands. If the query succeeds and if the information is available, then the values are written in the appliance registry. Note that the `odacli update-registry` command does not configure or troubleshoot any component-related issues.

What are non-discoverable parameters?

Parameter values that cannot be discovered using commands such as `srvctl`, `crsctl`, `sqlplus`, or operating system commands are non-discoverable parameters. Refer to the *Oracle Database Appliance Deployment and User's Guide* for your hardware model.

When do I use the `odacli update-registry` command?

Use the `odacli update-registry` command to discover appliance registry metadata for components `dbhome` and `gihome` after out-of-cycle patching.

Which values does the `odacli update-registry` command modify?

Except for values that are specific to Oracle Database Appliance, all other values can be modified correctly in the metadata. Values which are specific to Oracle Database Appliance such as `dbshape`, `dbclass`, and `enableSEHA` remain the same in the metadata. If the registry metadata for parameters of any component is incorrect, and the component is not mentioned in the Non Discoverable Parameters section, then when you run the `odacli update-registry` command, the components are rediscovered and the metadata is updated. Refer to the *Oracle Database Appliance Deployment and User's Guide* for your hardware model for the values that are updated. You can update the registry metadata of an individual database or `dbstorage` using the `-u` option.

If the value specific to Oracle Database Appliance was null before running the `odacli update-registry` command, then it remains the same after the command is run. Contact Oracle Support in these cases.

What values does the `odacli update-registry` command modify when the database option is specified using the `-n DB` option?

Databases for which the metadata already exists in appliance registry can only be re-discovered using the `odacli update-registry` command. For discovering databases that are migrated from outside your Oracle Database Appliance deployment, use the `odacli register-database` command. For information about migrating databases, see the *Oracle Database Appliance Deployment and User's Guide* for your hardware model.

What values does the `odacli update-registry` command modify when the database storage option is specified using the `-b dbstorage` option?

The dbstorages for which the metadata already exists, can only be discovered using the `odacli update-registry` command. For discovering dbstorage that is migrated from outside your Oracle Database Appliance deployment, use the `odacli create-dbstorage` command. The `odacli create-dbstorage` command does not create another dbstorage, it discovers the values and stores the same in the metadata. For Oracle ACFS dbstorage, it creates the file system, if required.

Can I manually modify or delete the appliance registry metadata and then run the `odacli update-registry` command to discover related information?

Erasing entries from the metadata and then using the `odacli update-registry` command for discovery is not supported.

Related Topics

- `odacli update-registry`
- Migrating Databases

Micronaut on Oracle Database Appliance

FAQs on Micronaut on Oracle Database Appliance.

Are there any changes visible to the user in Oracle Database Appliance release 19.16, because of the replacement of Dropwizard with Micronaut?

Yes, Existing Dropwizard-based DCS configuration JSON file is replaced with Micronaut based DCS configuration YML file and logback XML file. Running the `ps` command to list DCS processes displays different JVM parameters as the `java` command used to invoke Micronaut based DCS Application varies. For example, the Dropwizard-based output is as follows:

```
java -Doracle.security.jps.config=/opt/oracle/dcs/agent/jps-config.xml
-jar /opt/oracle/dcs/bin/dcs-agent*.jar server /opt/oracle/dcs/conf/
dcs-agent.json
```

The Micronaut-based output is as follows:

```
java -Doracle.security.jps.config=/opt/oracle/dcs/agent/jps-config.xml
-Dlogback.configurationFile=/opt/oracle/dcs/conf/dcs-agent-logback.xml
-Dmicronaut.config.files=/opt/oracle/dcs/conf/dcs-agent.yml -
```

```
Dold.config.files=/opt/oracle/dcs/conf/dcs-agent.json -  
Djavax.xml.parsers.SAXParserFactory=com.sun.org.apache.xerces.internal.jaxp.S  
AXParserFactoryImpl -jar /opt/oracle/dcs/bin/dcs-agent*.jar
```

How do I know whether DCS application is running on Micronaut on Oracle Database Appliance?

Check the DCS Application logs to see whether the application uses Micronaut or Dropwizard. When the DCS application completes startup, the following Micronaut informational message is displayed in the DCS application log file:

```
-----  
2022-01-29 14:14:20,886 INFO [main] [] i.m.r.Micronaut: Startup completed in  
1226ms. Server Running: https://<host>:<port>  
-----  
Note: i.m.r.Micronaut means io.micronaut.runtime.Micronaut
```

From the above log message, it is clear that the DCS application is running on Micronaut instead of Dropwizard.

With Dropwizard based DCS application, the log message is as follows:

```
-----  
2022-01-27 03:08:49,398 INFO [main] [] o.e.j.s.AbstractConnector: Started  
application@67be1284{SSL, (ssl, http/1.1)}{0.0.0.0:<port>}  
-----  
Note : Here, o.e.j.s.AbstractConnector means  
org.eclipse.jetty.server.AbstractConnector
```

Are there any changes to the log messages format after moving to Micronaut?

No. The existing log message format is carried forward to Micronaut.

Is there a difference in logging functionality between Oracle Database Appliance release 19.16 and earlier releases?

Yes. Logging configuration file has changed. Earlier logging configuration file changes were a part of application specific JSON configuration file such as `dcs-agent.json` for DCS agent, `dcs-admin.json` for DCS admin, and `dcs-controller.yml` for DCS controller.

From Oracle Database Appliance release 19.16 and later, logging configuration changes are in a separate application specific logback XML configuration file, such as the `dcs-agent-logback.xml` for DCS agent, `dcs-admin-logback.xml` for DCS admin and `dcs-controller-logback.xml` for DCS controller. Note that the log messages pattern and log file names remain unchanged.

In Oracle Database Appliance release 19.16 and later, dynamic reloading of log configuration is also supported. Whenever the DCS application-specific logback XML file is modified, such as when changing the log level from INFO to DEBUG, then the new changes are auto-detected. The new logging configuration changes are automatically applied without requiring restart of the DCS application. Note that this feature is introduced in Micronaut-based Oracle Database Appliance release 19.16 DCS applications and is not available in Oracle Database Appliance release 19.15 or earlier DropWizard-based DCS applications.

Where are the configuration parameters set with Micronaut-based DCS applications?

Configuration file location remains the same but file names are changed in Oracle Database Appliance release 19.16 and later.

In Oracle Database Appliance release 19.15 or earlier, DCS application-specific JSON configuration files had three sections:

- Server configuration parameters such as server port, HTTP or HTTPS, MTLS enabled or not, truststore and keystore certificate details, and so on.
- Custom configuration parameters such as node ID, nodes list, emulation mode, and so on.
- Log configuration parameters such as log file name, log rotation details, log messages format, log levels, and so on.

In Oracle Database Appliance release 19.16 or later DCS applications, configuration files are split into two files:

- YML configuration file had two sections: Server configuration parameters such as server port, HTTP or HTTPS, MTLS enabled or not, truststore and keystore certificate details, and so on. Custom configuration parameters such as node ID, nodes list, emulation mode, and so on.
- Logback XML configuration contains logging specific configuration changes only: Log configuration parameters such as log file name, log rotation details, log messages format, log levels, and so on.

When patching DCS applications, custom settings in JSON file are migrated to YML and logback XML files accordingly.

Table 1-4

DCS application	DCS application-specific configuration file (Server configuration parameters + Log configuration parameters + Custom configuration parameters) Note that releases earlier than 19.14 use JSON, whereas 19.14 or later use YML for DCS controller	DCS application specific configuration file (Server configuration parameters + Custom configuration parameters). Note that release 19.16 or later use YML	DCS application specific logging configuration file (Log configuration parameters). Note that release 19.16 or later use application-specific logback.xml
DCS agent	dc-agent.json	dc-agent.yml	dc-agent-logback.xml
DCS admin	dc-admin.json	dc-admin.yml	dc-admin-logback.xml

Table 1-4 (Cont.)

DCS application	DCS application-specific configuration file (Server configuration parameters + Log configuration parameters + Custom configuration parameters) Note that releases earlier than 19.14 use JSON, whereas 19.14 or later use YML for DCS controller	DCS application specific configuration file (Server configuration parameters + Custom configuration parameters). Note that release 19.16 or later use YML	DCS application specific logging configuration file (Log configuration parameters). Note that release 19.16 or later use application-specific logback.xml
DCS controller	dcs-controller.json (release 19.13 or earlier) dcs-controller.yml (release 19.14 or later)	dcs-controller.yml	dcs-controller-logback.xml

How is the logging configuration file structure in Oracle Database Appliance release 19.16 and later?

DCS agent logging configuration file (dcs-agent-logback.xml):

```
<configuration scan="true" scanPeriod="60 seconds">
  <property name="DCS_LOG_DIR" value="/opt/oracle/dcs/log" />

  <appender name="ROLLING"
class="ch.qos.logback.core.rolling.RollingFileAppender">
  <file>${DCS_LOG_DIR}/dcs-agent.log</file>
  <rollingPolicy
class="ch.qos.logback.core.rolling.SizeAndTimeBasedRollingPolicy">
  <!-- rollover daily -->
  <fileNamePattern>${DCS_LOG_DIR}/dcs-agent-%d{yyyy-MM-dd}.%i.log</
fileNamePattern>
  <!-- each file should be at most 100MB, keep 10 days worth of history
-->
  <maxFileSize>100MB</maxFileSize>
  <maxHistory>10</maxHistory>
</rollingPolicy>
<encoder>
  <pattern>%date{"yyyy-MM-dd HH:mm:ss,SSS"} %level [%thread] [%mdc]
%logger{15}: %m%n%rEx</pattern>
</encoder>
</appender>
```

```

    <appender name="httpAccessLogAppender"
class="ch.qos.logback.core.rolling.RollingFileAppender">
    <file>${DCS_LOG_DIR}/dcs-agent-requests.log</file>
    <rollingPolicy
class="ch.qos.logback.core.rolling.TimeBasedRollingPolicy">
    <!-- rollover daily -->
    <fileNamePattern>${DCS_LOG_DIR}/dcs-agent-requests-%d{yyyy-MM-
dd}.log
    </fileNamePattern>
    <maxFileSize>100MB</maxFileSize>
    <maxHistory>10</maxHistory>
    </rollingPolicy>
    <encoder>
    <charset>UTF-8</charset>
    <pattern>%msg%n</pattern>
    </encoder>
    </appender>

    <root level="INFO">
    <appender-ref ref="ROLLING" />
    </root>

    <logger name="HTTP_ACCESS_LOGGER" additivity="false" level="info">
    <appender-ref ref="httpAccessLogAppender" />
    </logger>

    <logger name="io.micronaut" level="INFO"/>
    <logger name="com.oracle.dcs.agent" level="DEBUG"/>
</configuration>

```

DCS admin logging configuration file (dcs-admin-logback.xml):

```

<configuration scan="true" scanPeriod="60 seconds">
    <property name="DCS_LOG_DIR" value="/opt/oracle/dcs/log" />

    <appender name="ROLLING"
class="ch.qos.logback.core.rolling.RollingFileAppender">
    <file>${DCS_LOG_DIR}/dcs-admin.log</file>
    <rollingPolicy
class="ch.qos.logback.core.rolling.SizeAndTimeBasedRollingPolicy">
    <!-- rollover daily -->
    <fileNamePattern>${DCS_LOG_DIR}/dcs-admin-%d{yyyy-MM-
dd}.%i.log</fileNamePattern>
    <!-- each file should be at most 100MB, keep 10 days worth of
history -->
    <maxFileSize>100MB</maxFileSize>
    <maxHistory>10</maxHistory>
    </rollingPolicy>
    <encoder>
    <pattern>%date{"yyyy-MM-dd HH:mm:ss,SSS"} %level [%thread]
[%mdc] %logger{15}: %m%n%rEx</pattern>
    </encoder>
    </appender>

```

```

    <appender name="httpAccessLogAppender"
class="ch.qos.logback.core.rolling.RollingFileAppender">
    <file>${DCS_LOG_DIR}/dcs-admin-requests.log</file>
    <rollingPolicy
class="ch.qos.logback.core.rolling.TimeBasedRollingPolicy">
    <!-- rollover daily -->
    <fileNamePattern>${DCS_LOG_DIR}/dcs-admin-requests-%d{yyyy-MM-dd}.log
    </fileNamePattern>
    <maxFileSize>100MB</maxFileSize>
    <maxHistory>10</maxHistory>
    </rollingPolicy>
    <encoder>
    <charset>UTF-8</charset>
    <pattern>%msg%n</pattern>
    </encoder>
    </appender>

    <root level="INFO">
    <appender-ref ref="ROLLING" />
    </root>

    <logger name="HTTP_ACCESS_LOGGER" additivity="false" level="info">
    <appender-ref ref="httpAccessLogAppender" />
    </logger>

    <logger name="io.micronaut" level="INFO"/>
    <logger name="com.oracle.dcs.admin" level="DEBUG"/>
</configuration>

```

DCS controller logging configuration file (dcs-controller-logback.xml):

```

<configuration scan="true" scanPeriod="60 seconds">
    <property name="DCS_LOG_DIR" value="/opt/oracle/dcs/log" />

    <appender name="ROLLING"
class="ch.qos.logback.core.rolling.RollingFileAppender">
    <file>${DCS_LOG_DIR}/dcs-controller.log</file>
    <rollingPolicy
class="ch.qos.logback.core.rolling.SizeAndTimeBasedRollingPolicy">
    <!-- rollover daily -->
    <fileNamePattern>${DCS_LOG_DIR}/dcs-controller-%d{yyyy-MM-dd}.%i.log</
fileNamePattern>
    <!-- each file should be at most 100MB, keep 10 days worth of history
-->
    <maxFileSize>100MB</maxFileSize>
    <maxHistory>10</maxHistory>
    </rollingPolicy>
    <filter class="com.oracle.oda.controller.CORSLogFilter" />
    <encoder>
    <pattern>%date{"yyyy-MM-dd HH:mm:ss,SSS"} %level [%thread] [%mdc]
%logger{15}: %m%n%rEx</pattern>
    </encoder>
    </appender>

```

```
<appender name="httpAccessLogAppender"
class="ch.qos.logback.core.rolling.RollingFileAppender">
  <file>${DCS_LOG_DIR}/dcs-controller-requests.log</file>
  <rollingPolicy
class="ch.qos.logback.core.rolling.TimeBasedRollingPolicy">
  <!-- rollover daily -->
  <fileNamePattern>${DCS_LOG_DIR}/dcs-controller-requests-%d{yyyy-
MM-dd}.log
  </fileNamePattern>
  <maxFileSize>100MB</maxFileSize>
  <maxHistory>10</maxHistory>
</rollingPolicy>
  <encoder>
  <charset>UTF-8</charset>
  <pattern>%msg%n</pattern>
  </encoder>
</appender>

<root level="INFO">
  <appender-ref ref="ROLLING" />
</root>

<logger name="HTTP_ACCESS_LOGGER" additivity="false" level="info">
  <appender-ref ref="httpAccessLogAppender" />
</logger>

<logger name="io.micronaut" level="INFO"/>
<logger name="com.oracle.oda.controller" level="DEBUG"/>
</configuration>
```

Adaptive Classification and Redaction (ACR) on Oracle Database Appliance

FAQs on ACR on Oracle Database Appliance.

What type of files does ACR handle?

ACR handles various files in a TFA collection such as trace files, alert logs, AWR, ASH, and ADDM reports, and ExaChk and ORAChk reports and other system generated logs. Also, ACR handles different compressed files such as zip, gzip, tar, tar.gz, bzip2, and tar.xz.

What are the files not redacted by ACR?

ACR skips binaries, shared libraries, compiled files, core dumps, jar files even if they are part of a TFA collection. This ensures that the diagnosability of the issue is not affected.

Does ACR redact literal values in the SQL statements?

Yes, ACR redacts SQL literals in the SQL statements present in AWR, ASH, and ADDM reports.

Are there any additional space requirements to run ACR?

ACR runs in 2 modes: In-place and Out-of-place. For out-of-place redaction, extra space equal to the size of the input files is required because redacted files are created in a different directory. For in-place redaction, extra space equal to the size of biggest file in the collection is required. In this case, each input file is replaced by the redacted file. For multiple ACR processes that are spawned, additional space requirement is the sum of the sizes of N largest files among the input files, where N is the number of parallel ACR processes.

Does ACR redact TFA collections generated on both nodes in a high availability environment?

In a high availability environment, TFA creates separate collections on both the nodes. ACR runs separately on each node on these collections to create two redacted collections. After redaction, redacted collection from the remote node is copied to the local node, that is, the node which initiated the TFA command.

Does ACR generate any trace files?

Yes, ACR generates trace files and logs in the `oracle.ahf/` directory present at the location `/opt/oracle/dcs/oracle.ahf/`. The trace files are as follows:

- `acrctl_*.trc` files generated in `oracle.ahf/data/HOSTNAME/diag/acr/HOSTNAME/acrctl/trace/` directory. These traces files are generated during an ACR run capturing various steps during redaction.
- `acr_info.json` file generated in `oracle.ahf/common/acr/acr_repo/` directory. This file contains the list of sanitized version of all entity instances identified during an ACR run.
- `acr_stats_*.txt` files generated in `oracle.ahf/common/acr/acr_repo/` directory. These files capture various statistics during an ACR run such as size of collection being redacted, number of parallel ACR processes, redaction time, and largest file size in collection.
- `acrmmap` file generated in `oracle.ahf/common/acr/acr_repo/` directory. This file contains the mapping from the redacted entity instance to the original instance. It is used to return the original entity instances while running `rmap` command.

I have specified 'mask' as the redaction mode, but I still see files and directory names getting sanitized. Is this expected?

Yes, this is correct ACR behaviour. Irrespective of redaction mode, ACR sanitizes sensitive entity instances appearing as part of files and directory names. This is done because changing files and directory names to '***' will make the collection unsuitable for navigation.

I have specified 'sanitize' as the redaction mode, but I still see some parts of trace files getting converted to '***'. Is this expected?

Yes, this is expected. Irrespective of the redaction mode, ACR converts sensitive data in trace files as part of block and redo log dumps to '***'.

Does ACR preserve features of entity instances during redaction?

Yes, ACR preserves features of entity instances in following ways:

- ACR preserves entity word's shape, that is, capitalisation of letters are preserved, and any trailing numerical characters are preserved.

```
HRDB1 => ONVL1  
hrdb1 => onvl1  
Hrdb1 => Onvl1  
HRDB1_host2_svc3 => ONVL1_ygmi2_wrj3  
myhost007 => ppspao007  
myhost008 => ppspao008
```

- **IPv4:** Restricts numbers for each octet to less than 255. Does not redact special octets such as 255, 127, and 0.
- **MAC address and IPv6:** Restricts numbers to follow hexadecimal number pattern.

Can I redact an existing unredacted TFA collection or any other files not part of a TFA collection?

Create a zip file containing all unredacted files, for example, `input.zip` and run the following command:

```
odacli redact-acr -i input.zip -m sanitize
```

The command redacts the files. The redacted file is present at `ACR_REPO/acr_runs/TIMESTAMP/outdir/input.zip`.

How do I verify whether a TFA collection is redacted?

Any redacted TFA collection has an empty file file `TFA.txt`. Run the following command on a zipped TFA collection:

```
zipinfo TFA_COLLECTION | grep "ACR.txt"
```

If the above command returns an entry, then the collection is redacted.

When redaction is performed, multiple ACR processes are spawned on my system. Is that expected?

Yes, ACR takes advantage of multiple free cores available on the machine to spawn multiple processes to reduce the overall redaction time.

If there are 10 cores on my system, will ACR use all 10 cores?

Number of parallel processes spawned by ACR never crosses 20% of the total number of cores. So, there will only be 2 parallel ACR processes created in this case.

CPU usage of each ACR process is close to 100%, is that an issue?

No, it is not an issue. Thus is expected behavior because ACR is a CPU intensive task. Whenever an ACR process is scheduled on a CPU, it utilises the CPU fully. But, since ACR does not run at an elevated priority, it does not starve other processes on the system.

How are `odaadmcli` commands different from `odacli` ACR commands?

The `odaadmcli manage diagcollect` command provides functionality to collect diagnostics data as TFA collection. It can be redacted using the following option:

```
odaadmcli manage diagcollect [-dataMask|-dataSanitize]
```

If auto redaction is enabled as follows, the `odaadmcli` command will always mask or sanitize even if the option is not explicitly specified.

```
tfactl set redact=[sanitize | mask]
```

The above command is specific to each node. The command cannot be run on the BUI.

```
odacli enable-acr  
odacli disable-acr
```

These commands enable or disable ACR on both nodes in a high-availability environment and can be invoked from the BUI. The `odaadmcli manage diagcollect` and `odacli` commands use `tfactl` command internally to create and redact the collection.

How do I redact bind variables appearing in audit trails?

Bind variables appearing in audit trails can be redacted to '*' by enabling Transparent Sensitive Data Protection (TSDP). To achieve this, sensitive column protection has to be enabled on the column of the table which one considers as sensitive. For more information, refer to the *Oracle Database Security Guide*.

How do I set up the staging server for ACR outside of the production environment?

For information about setting up the staging server for ACR, see My Oracle Support Note 2882798.1 at <https://support.oracle.com/rs?type=doc&id=2882798.1>.

Multi-User Access on Oracle Database Appliance

FAQs on Oracle Database Appliance Multi-User Access.

What are the advantages of enabling multi-user access on Oracle Database Appliance?

Multi-user access provides enhanced capabilities for authentication, authorization, resource ownership, and access control. This helps in providing an efficient mechanism for role separation. The Oracle Database Appliance administrator can create users with specific roles, limiting these users to only the operations allowed by that role. This role separation also restricts one user from accessing resources of other users, thereby ensuring resource ownership. This isolation removes barriers to consolidation. For example, creation and management of departmental databases can be delegated to different owners while maintaining separation. See the chapter *Implementing Multi-User Access on Oracle Database Appliance* in the *Oracle Database Appliance Deployment and User's Guide* for your hardware model for more information.

Is multi-user access enabled by default on provisioning Oracle Database Appliance?

No, multi-user access is enabled when provisioning Oracle Database Appliance only if you choose to enable the option at the time of provisioning the appliance.

Can I enable multi-user access on my current Oracle Database Appliance deployment?

No. The option to enable multi-user access is available only for new deployments of Oracle Database Appliance. Existing Oracle Database Appliance systems provisioned before Oracle Database Appliance release 19.13 continue to function without multi-user access feature even after patching to Oracle Database Appliance release 19.13 or later.

Can I disable multi-user access after enabling it during provisioning on Oracle Database Appliance?

No. Once enabled, multi-user access cannot be disabled. It is recommended that you provision the feature on your staging systems first, and then deploy it on your production system.

If I enable multi-user access on Oracle Database Appliance, do I need to enter my password every time I run an ODACLI command?

No, once you login with your Oracle Database Appliance account credentials, you are prompted for a password only when you run an ODACLI command for the first time. On successful authentication, an auth token is generated and used for authentication for running subsequent ODACLI commands. For every ODACLI command, the authentication token is further refreshed for a time interval equal to the token expiration duration. This implies that if your system is not idle for more than the token expiration duration, you need to enter the password only once. The default value of token expiration duration is 120 minutes and can be configured at the time of provisioning of the appliance up to a maximum of 600 minutes. However, it is recommended that the Oracle Database Appliance administrator takes the security policies of the organization into account while modifying the token expiration duration from its default value.

With multi-user access enabled, I am logged into the appliance as the `root` user. Why do I need to enter my password every time in spite of authentication token support ?

The authentication token support for ODACLI session management is linked to a multi-user access user account. Since `root` is an operating system administrative user and not a multi-user access user, auth token based session management system is not supported when a user logs in as `root`. Therefore, they must provide an Oracle Database Appliance account user name and password to run any ODACLI command.

Note that `root` access must be used only by the Oracle Database Appliance administrator and only for tasks that require `root` privileges. In all other cases, all multi-user access users in the system including `odaadmin` must use the Oracle Database Appliance credentials assigned to them.

Can I create new roles and entitlements?

No. This release supports default roles and entitlements and there is no provision for the `odaadmin` user to create new roles and entitlements.

What are the configurable parameters for multi-user access?

You can configure the following system settings:

- **Token expiration duration in minutes:** The minimum value you can specify is 10 minutes, the maximum value is 600 minutes, and the default is 120 minutes.
- **Password expiration duration in days:** The minimum value you can specify is 30 days, the maximum value is 180 days, and the default is 90 days.
- **Maximum failed login attempts allowed:** The minimum value you can specify is 2, the maximum value is 5, and the default is 3.

I have enabled multi-user access but do not plan to create new users. Will all Oracle Database Appliance features work as before ?

Yes. You can login as the default user created with the role of `oracleUser` and `gridUser` during the provisioning of the appliance, for all your needs.

What happens if I forget my Oracle Database Appliance account password?

You can reset your Oracle Database Appliance account password after authorization from the Oracle Database Appliance administrator. You can reset the password using ODACLI or the BUI.

What happens if I enter wrong password multiple times ?

Your account is locked when your consecutive login attempts have failed, as defined in the **Maximum failed login attempts allowed** setting. You can unlock your account by resetting the password after getting the authorization from the Oracle Database Appliance administrator. You can reset the password using ODACLI or the BUI.

I am a non-odaadmin user with ODA-DB role. Should I create a separate database home for my use?

The recommended practice is to create your own database home and then create databases in these database homes so that you have exclusive and full control of your database without anyone else being able to access it. However, in exceptional circumstances, you can request the Oracle Database Appliance administrator to grant you shared access to another database home. Creating a database on a shared database home restricts the operations you can perform on the database.

Are there any restrictions on the number of users that can be created on multi-user access enabled systems?

No. There is no maximum limit on the number of users that you can create on your deployment. The actual number of users depends on the availability of hardware resources such as CPU cores, disk space and memory on the appliance.

Integrated Oracle Data Guard

FAQs on Integrated Oracle Data Guard.

Can I configure Oracle Data Guard between high-availability and single-node Oracle Database Appliance system?

Yes. Oracle Data Guard configuration between high-availability to high-availability, high-availability to single-node, single-node to single-node, and single-node to high-availability are supported.

If Oracle Data Guard is implemented using ODACLI, can DGMGRL commands be used for switchover/failover/reinststate?

It is recommended to use ODACLI for switchover/failover/reinststate to better synchronize DCS metadata.

What is the best approach to setup Oracle Data Guard with Fast-Start Fail-Over?

Oracle Data Guard Observer for Fast-Start Fail-Over (FSFO) can be manually configured into Oracle Database Appliance Data Guard configuration.

Are different DB shapes of primary and standby allowed for using Oracle Database Appliance tooling to configure Oracle Data Guard?

Starting with Oracle Database Appliance release 19.11, different DB shapes of primary and standby are allowed.

Can the Oracle Database Appliance tooling configure Oracle Data Guard for one primary database to multiple standby databases?

Currently, the Oracle Database Appliance tooling supports one primary database to one standby database configuration.

Can standby database be created in ways other than irestore?

The Oracle Database Appliance tooling of configuring Oracle Data Guard only supports standby database created using irestore from NFS mount or Oracle Cloud backup of the database.

Is it mandatory to configure Oracle Data Guard using the Oracle Database Appliance tooling?

Oracle recommends that you use Oracle Database Appliance tooling to configure Oracle Data Guard. However, the manual steps to configure Oracle Data Guard should still work.

Will the manually configured Data Guard be affected by upgrading Oracle Database Appliance software?

No.

What are the best practices for migrating databases from an Oracle Database Appliance X5-2 hardware model to newer Oracle Database Appliance hardware models?

Use the Integrated Data Guard on Oracle Database Appliance feature, available from release 19.8, to migrate databases from old hardware to new hardware. Refer to the Oracle Database Appliance Deployment and User's Guide for your hardware model.

What is the procedure to patch or upgrade in the Integrated Data Guard environment?

Refer to the Oracle Database Appliance Deployment and User's Guide for your hardware model for the procedure to upgrade or patch databases.

What are the Oracle Database Appliance hardware models supported for Oracle Database Appliance Integrated Data Guard?

Any Oracle Database Appliance hardware model that supports Oracle Database Appliance release 19.8 and later supports Oracle Database Appliance Integrated Data Guard. Refer to the *Prerequisites for Oracle Database Guard Configuration* topic in the Oracle Database Appliance Deployment and User's Guide for your hardware model.

Oracle Database Appliance DB Systems

FAQs on Oracle Database Appliance DB Systems.

Can I create a single DBVM on a high-availability system?

For single-node systems, each DB system consists of only one DBVM. For high-availability systems, each DB system has two DBVMs, one on each Oracle Database Appliance system. Creation of a DB system with only one VM on a high-availability system is not supported.

How many databases does a DB System support?

One database per DB system. The database can be a multitenant container database (CDB) or a non-CDB. If you want to deploy multiple databases on each DB system, it is recommended that you consolidate into a CDB.

Can you run both Oracle Database Standard Edition and Enterprise Edition databases in KVMs on a single Oracle Database Appliance system?

Yes.

What is the CPU cores granularity for DBVM, that is, what is the smallest virtual DB system?

The minimal CPU cores a DB System supports is 2 cores.

Is it supported to run both bare metal databases and KVM DB system on an Oracle Database Appliance system?

Yes. Both bare metal databases and KVM DB system are supported on a single Oracle Database Appliance system.

Is Oracle Database Standard Edition2 licensing of one socket is possible?

In Oracle Database Appliance releases 19.10 and 19.11, you cannot choose from which socket the CPU is allocated. But based on how CPU is allocated, it is possible to use `odacli create-cpupool` to temporarily occupy the CPU pool on one socket, then create the DB system that uses the CPU pool on the second socket. Starting from release 19.12, you can specify the socket from which the CPU is allocated with the `odacli create-cpupool --socket` option. This is intended to support SE2 licensing. For other use cases, do not specify the `--socket` option and allow the system to handle the default CPU allocation. Do not create the CPU pool using mixed options, that is, some CPU pool created with the `--socket` option, and some created without. This causes imbalanced CPU allocations across the socket and do not optimize CPUs.

How many cores does a DB system utilize?

The shape of the DB system is derived from the shape you select for the database. The CPU cores for the database system is set as one defined in the database shape. For example, if you select `odb4` as the database shape, in which the database uses 4 cores, then the database system is created with 4 cores.

Can I use Oracle ACFS inside the DB system?

Oracle ACFS is not supported in the DB system.

Does DB system support oversubscription?

Each DB system is associated with a CPU pool created internally. In release 19.10, the CPU pool associated with a DB system cannot be shared. Starting with release 19.11, you can create the DB system type of CPU pool and share the CPU pool with multiple DB systems. Oversubscribing the CPU pool is allowed among DB systems. Oversubscribing the CPU pool may impact the performance. You must evaluate the impact before implementing CPU pool oversubscription in production.

What are the supported database versions on DB systems?

You can check the supported database version for a DB system version using the command `odacli describe-dbsystem-image`.

Can I install a later version of the database on the DB system than the Oracle Grid Infrastructure version on the bare metal system?

Yes, starting with release 19.12, DB system supports Oracle Database 21c. You can install Oracle Database 21.3 database in the DB system while the bare metal system runs Oracle Grid Infrastructure 19.x. Run the command `odacli describe-dbsystem-image` to confirm the Oracle Database and Oracle Grid Infrastructure versions supported on the DB system.

What is the maximum number of DB systems that can be created on an Oracle Database Appliance system?

The maximum number of DB systems depend on the CPU and the memory. By default, half of the host memory is allocated to the huge page, and the DB system is configured using huge page. So the total memory used by the DB system must be less than the huge page allocated at the host. If there are databases and blackbox VMs running at the host, then you must also consider them when you plan for CPU and memory resource for DB systems.

What is the storage used by the DB system?

The DB system local file system includes the `/u01`, used for Oracle Grid Infrastructure and Oracle Database home, boot partition, the partition for the operating system, swap space, and the `/opt` partition for the Oracle Database Appliance software. The size of the DB system local file system is fixed at 200 GB per node. For each high-availability system, each DB system node has 200 GB allocated for local storage, and placed on Oracle ACFS on the bare metal system. The databases on the DB system use Oracle ASM disk groups for storage. There is no quota limit per database. The space usage is dependent on how Oracle ASM disk group is used.

The DB system fails to be created due to lack of memory error. But the free space command shows that there is enough memory. What is the cause of the failure?

The DB system uses huge page. By default, Oracle Database Appliance allocates half of the physical memory as huge page. The huge page is shared by the databases on the bare metal system and DB system. Check `/proc/meminfo` on the bare metal for the number of free huge page available for the DB system. You can increase the huge page configuration based on the usage of the Oracle Database Appliance system.

Can I create multiple DB systems at the same time?

You can create DB systems that use different disk groups for the VM storage at the same time. If they use the same disk group as the VM storage, they must be created in serial.

Can I use FLASH disk group for the DB system VM storage?

Starting with release 19.12, DB systems can be created on FLASH disk group.

Is the name used in DB system case-sensitive or case-insensitive?

The DB system name is case-insensitive. For example, the name `DBSystemname1` is considered the same as `dbsystemname1`.

How can I backup a DB system?

For the detailed procedure, see My Oracle Support Note 2784991.1: *Database System backup on Oracle Database Appliance Release 19.10 and later* at <https://support.oracle.com/rs?type=doc&id=2784991.1>.

Why do I need to specify SCAN IP and VIP if I create a single-instance database?

Currently, a DB system with two nodes is created by default on Oracle Database Appliance high-availability environment to support Oracle RAC One node and Standard Edition High Availability databases. Therefore, SCAN IP and VIP are required.

How to resize the root partitions and mount points on DB systems?

The size of the internal file systems on a DB system are optimized and cannot be resized.

Since only one database is supported per DB system, can I create the database as CDB and consolidate databases into different PDBs? Is there any licensing issue?

Only one database is supported per DB system. You can select the CDB option to create a CDB when creating the DB system. One PDB is created by default, and additional PDBs can be created manually by using SQL Plus. You can create a maximum of three PDBs without

additional licenses. For more details, see <https://blogs.oracle.com/database/post/oracle-database-19c-up-to-3-pdbs-per-cdb-without-licensing-multitenant>.

Can I setup the Integrated Data Guard between the bare metal database and DB system database?

Yes, Integrated Oracle Data Guard is supported for DB system in release 19.11 and later. You can set up Oracle Data Guard between bare metal database and DB system database.

Can a different public network be used for the DB system?

Yes. By default, if no public network is provided when creating the DB system, the DB system uses the pubnet from the bare metal as its public network. Starting from release 19.11, you can have different networks for the public network of the DB system. You can create a vnetwork (bridge or bridged-vlan) and select this vnetwork when creating the DB system.

Can I change the DB system memory size without changing number of the CPU?

Yes. You can change the DB system memory size.

Can I detach the default public network from the DB system?

You can select a vnetwork as the public network when creating the DB system. By default, the DB system uses the pubnet as the public network. You cannot detach this public network after the DB system is created because it used by Oracle Grid Infrastructure and the database. You can add a new public network to an existing DB system and detach this second public network with the `odacli modify-dbsystem` command.

How do I change the public network of a DB system?

Refer to the following *My Oracle Support* Notes:

- *Doc ID 2824390.1: How To Change DBSystem Nodes' Public IP / Network After Deployment on ODA S/M/L Release 19.10 and later* at <https://support.oracle.com/rs?type=doc&id=2824390.1>
- *Doc ID 2797758.1: How To Change DBSystem Nodes' Public IP / Network After Deployment on ODA HA Release 19.10 and later* at <https://support.oracle.com/rs?type=doc&id=2797758.1>

Can I create a DB system of an earlier release?

You can only create a DB system of the same release as the bare metal system on Oracle Database Appliance. For example, if the bare metal system is on Oracle Database Appliance release 19.13, you can deploy DB system running Oracle Database Appliance release 19.13. However, you can choose earlier releases of Oracle Grid Infrastructure and Oracle Database, for example, you can deploy Oracle Grid Infrastructure 19.12 and Oracle Database 19.12 on a DB system running Oracle Database Appliance release 19.13. Use the command `odacli describe-dbsystem-image` to check the supported DB system, Oracle Grid Infrastructure, and Oracle Database releases.

What is the *modify dbsystem memory* feature?

The *modify dbsystem memory* feature allows you to increase the DB system memory so that you can add more memory to the database (SGA or PGA). It is recommended that you select DB system shapes instead, optimizing the sizing of CPU and memory used by the database. If you want to increase database memory without changing the number of CPUs, then use the *modify dbsystem memory* feature.

Why is decreasing the DB system memory not supported? How can I reduce the DB system memory?

Database CPU and memory sizing ratios on Oracle Database Appliance are optimized for pre-defined database shapes. Reducing the database memory for a given CPU setting is against the sizing best practices recommended on Oracle Database Appliance. If reducing database memory is justified, then you can reduce the database memory without modifying the DB system memory. Change the DB system shape to reduce the DB system memory.

If I increase the DB system memory and the database memory and then change the DB system shape by using the command `odacli modify-dbsystem -s`, what happens to the database memory?

Changing the DB system shape changes the database inside the DB system accordingly. So, both the DB system and the database setting, including memory, are reset to match the new shape of the DB system.

I increased the *dbsystem* memory. How do I revert the changes?

You can use `odacli modify-dbsystem -s` to revert the changes. For example, if you updated the DB system from the default `odb4/32G` to `odb4/48G`. To revert it to `odb4/32G`, you can modify the shape from `odb4` to `odb2`, and then from `odb2` to `odb4` again.

How does Oracle Database Appliance manage the access control to the database on a DB system?

The database on a DB system accesses the Oracle ASM disk groups managed by Oracle ASM. Oracle ASM access control prevents a user of one DB system from accessing files owned by another DB system. For more information, refer to the *Managing Oracle ASM File Access Control for Disk Groups* topic in the *Oracle ASM Administrator's Guide* at <https://docs.oracle.com/en/database/oracle/oracle-database/19/ostmg/asm-access-control-diskgroups.html#GUID-89BB888A-BA1C-4095-BB6D-C3BE142CB371>. When you create a DB system on Oracle Database Appliance, this security feature is implemented by default.

Application KVMs on Oracle Database Appliance

FAQs on Application KVMs on Oracle Database Appliance.

The Browser User Interface (BUI) prompts to enter `VM Disk` and `Virtual Disk` when creating VM compute instance. What are `VM Disk` and `Virtual Disk`?

The `VM Disk` is the disk created internally where the VM operating system is installed and shows up as root partition inside the VM. When a VM is created, the `Virtual Disk` is a separate disk that can be used to provide additional space in the VM. The `Virtual Disk` details can be specified when creating the VM, or attached to the existing VM. It can then be formatted to a file system and mounted to a mount point, such as `/data`, inside the VM.

How do I backup and restore the application VMs?

See My Oracle Support note 2779329.1 *Backup of KVM guests on ODA 19.9 BM and later*.

Does Oracle Database Appliance application KVM support Oracle Solaris 11.4?

Yes. Oracle Solaris 11.4 is included in the list of supported operating systems for Oracle Linux KVM.

What are the options for My Oracle Support for Oracle Solaris 11.4 running on Oracle Database Appliance application KVM?

Support for Oracle Solaris 11.4 is included in Oracle Premier Support for Systems with Oracle Database Appliance.

What are the guidelines for running Oracle Solaris 11.4 on Oracle Database Appliance application KVM?

Oracle Database Appliance does not provide application KVM management. However, for any issues that arise, you can contact *My Oracle Support* and provide details of the error by reproducing it in a bare metal system environment. For more information, see *My Oracle Support Note: Oracle Solaris Support on Virtualization Platforms (Doc ID 1681652.1)* at <https://support.oracle.com/rs?type=doc&id=1681652.1>.

What is the maximum size of RAM for an application VM?

The maximum size depends on the available memory of the Oracle Database Appliance host. It can be as much as the available regular memory and does not include the huge page memory.

How do I attach a new network to the application VM?

When an application VM is created, a default network interface is created in the VM. This interface is attached to the default KVM bridge (virbr0) at the host. This interface is intended for convenient communication between the host and VM. This interface does not attach to any physical network interface on the host. Do not configure this interface to use for any external communication. You should create a vnetwork using the command `odacli create-vnetwork` and attach the vnetwork using the command `odacli modify-vm -avn` to the VM. This creates a second interface in the VM. Configure this network interface for your network requirement. The network interface name in the VM depends on the operating system.

How do I configure Highly-Available NFS (HANFS) on Oracle Database Appliance KVM system?

Oracle Database Appliance KVM uses a different architecture than the Oracle Database Appliance OVM system. There is no ODA_BASE in the Oracle Database Appliance KVM system. The Oracle Grid Infrastructure and Oracle ASM run on the Oracle Database Appliance host where the application VMs and DB systems run. If you want to use Oracle ACFS for storage, then configure HANFS using Oracle ACFS and mount it on the DB system and application VM. See the *Oracle ASM Administrator's Guide* at the location: <https://docs.oracle.com/en/database/oracle/oracle-database/19/ostmg/understand-acfs-admin.html>.

Also, see <https://www.oracle.com/technetwork/database/database-technologies/cloud-storage/acfs/learnmore/acfs-nas-max-wp-3618364.html> for configuring HANFS. Choose the VIP from the public network used by the DB system or the application VM.

What happens to the KVMs the DB system and the application VM when one of the Oracle Database Appliance high-availability node fails?

Autostart and failover can be set for application VMs as required to enable the failover. Oracle RAC One Node database or Standard Edition High Availability databases on DB systems can failover to the second node automatically.

In case of disaster, if reimaging is only option to recover the Oracle Database Appliance, what are all the backups required to restore all the KVMs on Oracle Database Appliance?

DB systems cannot be restored in such a case and they need to be recreated. The database can be restore from backup once the DB system is recreated. Remote DB backup is essential to recover the database. KVM guest backups are required to restore application VMs.

How can I download the additional RPMS for the application VMs and DB systems?

You manage the application VMs on Oracle Database Appliance, and hence the RPMS in the VM. The policy of updating or installing additional RPMS for DB system is the same as the process on bare metal system. The DB system contains version lock to maintain the integrity of the RPMS. Any RPM that violate the version locks should not be installed or updated. Use the Oracle Database Appliance patch bundle to patch your bare metal and DB systems.

How is memory allocated on Oracle Database Appliance?

About half of the Oracle Database Appliance host memory is allocated to huge page pool, which is consumed by the SGA of the database and DB system. The remaining memory is shared by the database and other memory usage, Oracle Clusterware such as Oracle Grid Infrastructure and Oracle ASM, application VMs, operating system, swap space, and other applications that you may have installed. Therefore, when sizing the database, DB system, or the application VMs, the huge page and the remaining regular memory must be carefully considered.

How can I create the VM using the KVM template published by Oracle?

See My Oracle Support note Doc ID 2992638.1: Create Apps VM using Oracle Linux template on ODA at

<https://support.oracle.com/rs?type=doc&id=2992638.1>

How do I configure the application VM network after attaching a vnetwork?

See the *Oracle Database Appliance Deployment and User's Guide* for your hardware model.

How do I create a Microsoft Windows application VM?

See My Oracle Support Note 2748946.1 at:

<https://support.oracle.com/rs?type=doc&id=2748946.1>

How do I migrate OVM guest machines from Oracle Database Appliance Virtualized Platform to KVM on bare metal systems?

Starting with Oracle Database Appliance release 19.15, you can use ODACLI commands to perform in-place migration from OVM to KVM on the same Oracle Database Appliance system.

To migrate OVM guest machine to KVM on a different Oracle Database Appliance system, see *My Oracle Support* Note 2773840.1: *Migrating OVM guests from ODA VP to KVM on ODA BM* at:

<https://support.oracle.com/rs?type=doc&id=2773840.1>

Which are the guest operating systems supported on Oracle Database Appliance application VMs?

Oracle Database Appliance application VMs support the same guest operating systems as Oracle KVM. For more information about supported operating systems, see the *Oracle Linux KVM User's Guide*.

Why does the command `odacli modify-vm --enable- numa` fail on my application with the error `Cannot get NUMA configuration for odd number of max cores?`

`--enable- numa` is only supported for application VMs with even number of cpu cores (one cpucore is 2 vcpus). When enabling NUMA for existing application VMs, cpucores are evenly distributed across two NUMA nodes, that's why it requires VM to have even number of cpucores in order to enable NUMA.

Is overcommitting of CPU and memory supported on application VMs?

Overcommitting of CPU and memory is possible on Oracle KVM. Oracle Database Appliance does not restrict overcommitting of CPU or memory for application VMs. In general, it is not recommended to overcommit CPU or memory. Overcommitting CPU can lead to slow performance of the VMs and the host. Overcommitting memory can cause the system to run out of memory (OOM), which may lead to the Linux kernel shutting down important system processes. If you decide to overcommit CPU or memory, ensure that sufficient testing is performed.

What is the maximum number of application VMs that can be created on Oracle Database Appliance?

The maximum number of application VMs that can be created on Oracle Database Appliance depends on the available CPU and regular memory and whether you want to overcommit CPU and memory. It is recommended that you reserve at least 2 CPU cores and 24 GB memory for the bare metal host (hypervisor). A portion of the Oracle Database Appliance memory is also allocated to HugePages and is not accounted in the regular memory available to the application VMs. Check `/proc/meminfo` for memory allocation details.

CPU Pools

FAQs on CPU pools on Oracle Database Appliance.

What are the types of CPU pools supported?

There are three types of CPU pools available. BM type, VM type and DBS (DB system) type. The `dbss` type is available in Oracle Database appliance release 19.11 and later. In general, CPU pool is used to limit CPU usage by the objects, but it is not designed to prevent other objects from using the CPUs within the CPU pool.

What is the BM CPU pool?

BM CPU pool is designed to limit the database on the Oracle Database Appliance to only run on the CPUs within the CPU pool. Once a CPU pool is assigned to a database, then the database can only run on the CPU set within the CPU pool. BM CPU pool can be attached to different databases. But one database cannot use more than one CPU pool.

What is VM CPU pool?

VM CPU pool is designed to limit the application VMs to only run on the CPUs within the CPU pool. Once a CPU pool is assigned to a VM, then the VM can only run on the CPU set within the CPU pool. VM CPU pool can be attached to different VMs.

What is DB system CPU pool?

With Oracle Database appliance release 19.11 and later, a new CPU pool type called DB system CPU pool, or `dbss` type is available. The `dbss` CPU pool cannot overlap with any other CPU pool (BM, or VM CPU pool). After it is created, it can be passed as a parameter to the command `odacli create-dbssystem` to limit the DB system CPU within the CPUs inside this CPU pool. Multiple DB systems can share one `dbss` CPU pool and can oversubscribe to the `dbss` CPU pool. When DB system is created without specifying the `dbss` CPU pool, an internal CPU pool is created and this CPU pool is exclusive to this DB system; it cannot be shared with any other DB system. The `dbss` CPU pool is the DB system CPU pool is manually created through ODA CLI commands. The `dbss` CPU pool is assigned to DB systems at creation time or after creation using the command `odacli modify-dbssystem`. The `dbss` CPU pool can be shared by different DB systems, but the size of the CPU pool should be at least equal or greater than the largest DB system shape it is associated to.

What is shared DB system CPU pool?

With Oracle Database appliance release 19.11 and later, a new CPU pool type called DB system CPU pool, or `dbss` type is available. The `dbss` CPU pool cannot overlap with any other CPU pool (BM, or VM CPU pool). After it is created, it can be passed as a parameter to the command `odacli create-dbssystem` to limit the DB system CPU within the CPUs inside this CPU pool. Multiple DB systems can share one `dbss` CPU pool and can oversubscribe to the `dbss` CPU pool.

What happens if application VMs are not using CPU pool?

If a database or VM is not using a CPU pool, then the database or VM can run on any CPU. The operating system assigns the available CPUs at run time, including the CPUs within a CPU pool assigned to other VMs.

Can a CPU pool overlap with another CPU pool by sharing some CPUs?

No, two CPU pools cannot overlap.

Can I create the `vm` CPU pool on a single node on a high-availability environment?

Currently, this is not supported. If a VM uses a `vm` CPU pool, and wants to failover to the other node, the CPU pool is pre-created.

I have DB system and application VMs running on the same host, how do I prevent application VMs from using the CPUs of the DB system?

You should create a CPU pool for the application VMs. Attaching the `vm` CPU pool to the application VMs will limit VMs to only use the CPUs in the CPU pool, and not use the CPUs of the DB systems.

Can application VMs share the same CPU pool?

Yes, different application VMs can share the same CPU pool. Oversubscription to the CPU pool is allowed, although not recommended because of potential performance impact.

Can DB systems share the same CPU pool?

Yes, different DB systems can share the same `db`s CPU pool. Oversubscription to the CPU pool is allowed, although not recommended because of potential performance impact.

Can DB systems share the same CPU pool with application VMs?

No.

Can I use all bare metal system CPU cores in the CPU pool?

In general, you must not use all CPU cores in the CPU pool. The bare metal system host needs CPU as well. Starting with Oracle Database Appliance release 19.12, by default, the first core of the socket is reserved for the bare metal system host and cannot be used by the VM or DB system CPU pools, and can be used by BM CPU pools. With Oracle Database Appliance release 19.13, you can use the `--use-reserved-cpucores` option to override this behavior in the `odacli remap-cpupools` command. For more information, see the *Oracle Database Appliance Deployment and User's Guide* for your hardware model.

Oracle Database Appliance Networks

FAQs on Oracle Database Appliance networks.

Can a mix of 10 Gb Base-T and 10 Gb SFP+ network adapters be used on Oracle Database Appliance X8-2-HA?

Yes. For the detailed configuration options and instructions, refer to the chapter *Managing Networks* in the *Oracle Database Appliance Deployment and User's Guide* for your hardware model.

What is the `odacli` command in this release to set the Broadcom NIC speed to either 10Gb or 25Gb?

For 10Gb or dual-mode transceivers for the Oracle Dual Port 25 Gb Ethernet Adapter Whitney+ network controller, you can use the `odacli configure-firstnet -o` command to set the speed of the network card manually.

What are the supported network cables for Oracle Dual Port 25 Gb Ethernet Adapter SFP Ports?

The adapter supports cables and transceivers that align to the SFP28 and SFP+ MSAs. These include Direct Attach Copper (DAC) cables for shorter lengths as well as discrete transceivers and fiber cables for longer lengths. For the details, refer to the section *Ethernet Transceivers and Cables Supported* in the *Oracle Dual Port 25 Gb Ethernet Adapter Product Notes*.

For Oracle Database Appliance X8-2-HA systems, what is the order in which to install additional network cards?

For Oracle Database Appliance X8-2-HA and X8-2M, the additional network cards can be installed in the order of slot 2 and 10. For Oracle Database Appliance X8-2S, it might be in the order of slot 8 and 10.

Can the Net0 (on-board 1Gbps port) in Oracle Database Appliance X7-2 be used?

Yes, it can be configured using ODACLI or the Browser User interface (BUI).

How do I change from SFP28 Copper to SFP28 Fiber in Oracle Database Appliance X7-2?

Shut down the server, install the appropriate transceivers, swap cables, and restart the server. There are no changes necessary in the host operating system.

How can I disable DCB (Data Center Bridging) and LLDP (Link Layer Discovery Protocol) on the Oracle Dual Port 25 Gb Ethernet Adapter?

In the BIOS setup, select **Oracle Dual 25Gb Ethernet Adapter** in the **Advanced** tab. In the Device Configuration menu, disable LLDP Nearest Bridge and DCB Protocol.

Does Oracle Database Appliance support IPv6?

Currently IPv6 is not supported.

Does Oracle Database Appliance support Link Aggregation Control Protocol (LACP) or trunk mode?

Starting with Oracle Database Appliance release 19.18, you can configure LACP when you run the `odacli configure-firstnet` while deploying the appliance. You can also use the `-o` option with the `odacli configure-firstnet` command to enable or disable LACP on a bonding interface after deployment.

Can I create network bonding across two physically separated NICs?

Yes. Starting with Oracle Database Appliance release 19.15, Oracle Database Appliance X9-2 supports creation of network bonding across two physically separated NICs is supported.

Can I configure Oracle Database Appliance to break the network bonding?

Breaking network bonding feature is available only for Oracle Database Appliance X7-2.

Are there any issues when configuring the same subnet on two different bonding interfaces?

There is no assurance to which interface will be used to transmit traffic. The wrong interface may be used or one of IP addresses may become non-reachable.

How can I resolve the Multi-Homed Routing issue?

Refer to the Oracle blog *Advanced Routing for Multi-Homed Hosts* at <https://blogs.oracle.com/networking/advance-routing-for-multi-homed-hosts>.

Is bonding or HAIP used for private interconnect on bare metal systems?

Starting with Oracle Database Appliance release 18.7, the bonding interface is used instead of the HAIP on two network interfaces.

How can I disable Network Manager in Oracle Database Appliance release 19.x?

By default, the "NM_CONTROLLED=No" is set on the `ifcfg` files provided by Oracle Database Appliance. If you have networking issues, contact Oracle Support to get instructions on how to address your issue.

What is the default firewall on Oracle Database Appliance releases 18.x and 19.x?

The iptables is default for Oracle Linux 6 and firewalld is the default for Oracle Linux 7. By default, both iptables and firewalld are not running on Oracle Database Appliance. When the iptables or firewalld need to run, ensure that your rules do not affect the cluster interconnect interfaces, protocols and network address space. If iptables are already running in pre-19.x systems, refer to the My Oracle Support note *Oracle Linux: Migrate OL6 iptables(8) Rules To OL7 firewalld(1)* (Doc ID 2518939.1) or *Oracle Linux 7: How to Replace firewalld with iptables* (Doc ID 2620399.1).

How do I change the public IP or network for an existing Oracle Database Appliance after deployment?

Refer to the My Oracle Support note *How To Change the Public IP / Network For an Existing ODA HA After Deployment (DCS)* (Doc ID 2638458.1) and *How To Change the Public IP / Network For an Existing ODA M/L After Deployment (DCS)* (Doc ID 2638331.1).

How to change the DNS and NTP settings after deployment?

Refer to the My Oracle Support note *How To Change DNS and NTP Settings On an ODA Post-install for DCS Stack 18.3 to 18.6* (Doc ID 2680182.1) and *ODA:Changing NTP server Does Not Change Output Of odacli describe-system* (Doc ID 2675635.1).

How can I update the NTP server information on Oracle Database Appliance registry after deployment?

Refer to the My Oracle Support note *ODA:Changing NTP server Does Not Change Output Of odacli describe-system* (Doc ID 2675635.1).

Which SNMP version is supported for Oracle ASR on Oracle Database Appliance?

The SNMP v3 is supported on the host.

Is NTP supported in Oracle Database Appliance release 19.x?

By default, NTP is not running in Oracle Linux 7. To start NTP, refer to the My Oracle Support note *Oracle Linux: NTP Does Not Start Automatically After Server Reboot on OL7 (Doc ID 2422378.1)*.

Is enabling the FTP daemon supported on Oracle Database Appliance?

It is not recommended to enable the FTP due to security reasons. If you enable the FTP daemon for some reason, you must ensure that security is considered and that installation and configuration is done properly.

How can I use Oracle Database Appliance as an NFS server?

Refer to the My Oracle Support Note *ODA (Oracle Database Appliance): How To export ACFS (cloudfs) using HANFS (Doc ID 1934030.1)* or *ODA (Oracle Database Appliance): How to export ACFS using Samba-CIFS to MS Windows (Doc ID 1634154.1)*.

Can I change the VLAN ID of the default public VLAN configured by odacli configure-firstnet?

It is not possible to change the VLAN ID without redeploying the appliance.

How do I use native VLAN?

To use the native VLAN, do not choose the `VLAN` feature in the `odacli configure-firstnet` command.

Is the policy-based routing supported in Oracle Database Appliance?

The policy-based routing could be configured on non-NM_CONTROLLED interfaces. When setting your policy, be aware that there are reserved subnets for the HAIP, HAVIP, private interconnect, and 169.254.0.0 on Oracle Database Appliance. When patching your appliance to the latest release, you must set these settings again.

Is the VRF (Virtual Routing and Forwarding) supported in Oracle Database Appliance?

The layer 2 tagged VLAN is supported. Regarding VRF in layer 3 settings, you must set these settings again when patching your appliance to the latest release.

How do I resolve the subnet conflict with Oracle Database Appliance private network 192.168.16/17.xxx?

The 192.168.16.0/24 and 192.168.17.0/24 subnet are reserved for private network communications.

Is it possible to separate the networks used by the two databases?

Starting with Oracle Database Appliance release 19.7, a network can be attached or detached with the database with the `odacli modify-database` command.

Can I modify the network script under `/etc/sysconfig/network-scripts` manually?

It is not recommended to modify the `ifcfg-*` network script file manually, because these changes may get lost when you upgrade to a new Oracle Database Appliance release. If you modify these files, then you may need to reconfigure these settings when you patch the appliance to the latest release.

Browser User Interface

FAQs on Browser User Interface (BUI).

Why do I get an error when I try to access the BUI at port 7093?

Oracle Database Appliance uses self-signed certificates. Depending on the browser and browser version, you may receive a warning or error that the certificate is invalid or not trusted because it is self-signed, or that the connection is not private. Ensure that you accept the self-signed certificate for the agent and Browser User Interface. For more information about configuring the self-signed certificates, see the *Oracle Database Appliance Deployment and User's Guide* for your hardware model.

Why does the browser show `Not Secure` warning when I access the BUI?

Oracle Database Appliance BUI uses self-signed certificates. For custom certificates, unless configured, a warning is displayed. For more information about configuring custom certificates, see the *Oracle Database Appliance Deployment and User's Guide* for your hardware model.

Can I change the port number for accessing the BUI?

The BUI runs on port 7093 by default. You can manually change the port number following these steps:

1. Update the entry corresponding to "port" in `/opt/oracle/dcs/conf/dcs-controller.json`:

```
{
  "server": {
    "applicationConnectors": [
      {
        "type": "https",
        "port": 7093,
        "keyStorePath": "/opt/oracle/dcs/conf/dcsKey.jks",
        "trustStorePath": "/opt/oracle/dcs/conf/cacerts",
        "keyStorePassword": "OBF:1vn2lugulsajlv9ilv94lsarlugwlvo0",
        "certAlias": "selfsigned",
```

2. Restart DCS controller after making the changes:

```
systemctl restart initdcscontroller
```

Bill of Materials in the BUI does not show up-to-date information. Why does this happen?

The Bill of Materials Information shown in the BUI is based on the last collection on Oracle Database Appliance. The collection can be triggered using the command `odacli describe-system -b`. The BUI does not trigger the collection. If you notice that the information shown in the BUI is not up-to-date, run the command `odacli describe-system -b` and refresh the BUI to see the latest collection. The collection also runs as part of a weekly automatic scheduler on Oracle Database Appliance.

What are the supported browsers for running the BUI?

Firefox, Chrome, and Safari. The BUI may not display correctly in an Internet Explorer browser.

Is the Oracle Database Appliance documentation in the BUI up-to-date with the release?

The Oracle Database Appliance documentation shown in the BUI contains the latest documentation for the release.

Does BUI support all of the functionality that is available in the CLI?

While the BUI does support most of the functionality, there may be some functionality that is available only using ODACLI commands.

Does BUI support Oracle Data Guard?

BUI does not support Oracle Data Guard currently. Use ODACLI commands to configure and manage Oracle Data Guard on Oracle Database Appliance.

Is Backup Policy same as Backup Config?

Yes. The BUI uses the term Backup Policy which corresponds to backupconfig used by the ODACLI commands. They both mean the same.

Can I update DCS admin and DCS agent using BUI?

The BUI does not provide support to update the DCS admin. The DCS agent is updated when updating the server in the BUI. When you select the option to update the server in the BUI, the DCS agent is first updated, if it is an earlier version.

Can the BUI be used to run `odacli configure-firstnet`?

No. In order to access the BUI, `odacli configure-firstnet` has to be run first. The BUI can then be accessed using the host name or public IP.

If a BUI page fails to load, are there any logs that can be checked?

Different browsers have different ways for accessing the console logs. In most of the browsers, you can right-click and select **Inspect or Inspect Element**. This opens a window where you can check the console logs and look for errors or exception traces.

Backup and Recovery

FAQs on backup and recovery on Oracle Database Appliance.

What are all the backup destinations supported in Oracle Database Appliance?

Oracle Database Appliance supports backups to Fast Recovery Area (Disk), Object Store and NFS-mounted location.

How do I take a database backup?

See the *Oracle Database Appliance Deployment and User's Guide* for your hardware model for the complete procedure.

What are the Oracle Database Appliance components which support backup and restore?

Oracle Database Appliance supports backup of components such as database and TDE wallet. When you take a backup of the database, a backup of the TDE wallet is also performed, if the database is TDE-enabled.

If the database is TDE-enabled, is it required that the backupconfig that is used to attach to that database must have a separate Database backup location and TDE wallet location?

Yes. According to security guidelines, the backup of the database and TDE wallet must not be in the same location.

How is the database backup location and TDE wallet location separated in the case of disk backupconfig?

While creating the disk backupconfig, explicit paths for database backup and TDE wallet backup is not required to be provided. Oracle Database Appliance internally backs up the TDE wallet to separate path within the Fast Recovery Area (FRA).

What type of keystore is used in Oracle Database Appliance to configure TDE?

Oracle Database Appliance uses Software Keystore to configure TDE

What are the different types of keystores created in Oracle Database Appliance?

Oracle Database Appliance creates **Password Protected Software Keystore** and **Autologin Software Keystore**. However, **Local Autologin Software Keystore** is not created on Oracle Database Appliance.

Does Oracle Database Appliance tooling support setting up Oracle Key Vault as a key management solution?

Not at this time.

Which are the database versions that Oracle Database Appliance supports TDE on?

Oracle Database Appliance supports TDE configuration for 18c and above databases.

What is the TDE parameter used in Oracle Database Appliance while configuring TDE?

Oracle Database Appliance uses the Oracle recommended `WALLET_ROOT` parameter to configure TDE.

Does Oracle Database Appliance support auto backups?

Yes, Oracle Database Appliance supports auto backup of database and archivelogs.

When are the auto backup schedulers created?

Backup schedulers are created while attaching a backupconfig to the database. On successfully attaching the backupconfig to the database, a database backup scheduler and archivelog backup scheduler are created.

What is the frequency of the database backup scheduler and archivelog backup scheduler?

Default frequency of database backup scheduler is once every 24 hours and that of archivelog backup scheduler is every 30 mins.

Can the autobackup scheduler be disabled?

Yes, the database backup scheduler and archivelog backup scheduler associated to the database can be disabled by running the `odacli update-scheduler` command.

Can the frequency of the schedulers be changed?

Yes, the frequency of the schedulers can be changed by updating the corresponding `cron` expression using the `odacli update-scheduler` command.

Can the database backup be disabled?

Yes, the database backups can be disabled by using the `odacli modify-database` command with the `-no-bkp` option. Once disabled, manual backup, using the `odacli create-backup` command and automatic backups using schedulers, cannot be performed.

What are the backup compression algorithms supported in Oracle Database Appliance?

Oracle Database Appliance supports BASIC, LOW, MEDIUM and HIGH compression algorithms.

How can the compression algorithm be specified in Oracle Database Appliance backup?

Compression algorithm can be specified while creating the Backupconfig object. However, by default, BASIC compression algorithm is used while creating Disk or NFS backupconfig and MEDIUM compression algorithm is used while creating Objectstore backupconfig.

Can the compression algorithm be updated?

Yes, use the `odacli update-backupconfig` command to update the compression algorithm.

What is the license requirement to use different types of compression algorithm in Oracle Database Appliance?

BASIC compression algorithm does not require any license, whereas LOW, MEDIUM and HIGH compression algorithm require Advanced Compression Option license. However, if the user has Objectstore subscription, that is, Objectstore credentials, then any compression algorithm can be used.

Does Oracle Database Appliance check for license compliance before using any compression algorithm?

Compliance with license is the end user's responsibility, and Oracle Database Appliance tooling does not verify it.

Are the Objectstore credentials validated while creating the Objectstoreswift object?

No. The Objectstore credential is not validated while creating the Objectstoreswift object.

When does the validation of the Objectstore credentials happen?

The validation of the Objectstore credential happens while creating the backupconfig which uses the Objectstore.

What all attributes are considered to validate the Objectstore credentials?

Attributes such as Username, Password, Endpoint URL, Tenancy, and Container are considered while validating the Objectstore credentials.

Can more than one Objectstoreswift object can be created using the same Objectstore credential?

Yes, any number of Objectstoreswift objects can be created with one Objectstore credential. However, the Objectstoreswift object names must be different. Oracle Database Appliance does not allow multiple Objectstoreswift objects with the same name.

Can multiple Objectstore backupconfigs be created using same Objectstoreswift object?

Yes, multiple Objectstore backupconfigs can be created using same Objectstoreswift object. The names of the Objectstore backupconfigs must be different.

Can TDE enabled database be registered in Oracle Database Appliance?

Yes, Oracle Database Appliance supports registration of TDE-enabled database from 19.12 onwards with limited support for TDE wallet lifecycle management.

What are the prerequisites to register TDE enabled database in Oracle Database Appliance?

Following are the prerequisites:

- Software keystore must be used to configure TDE.
- `SQLNET.ENCRYPTION_WALLET_LOCATION` must be used to configure TDE.

- For bare metal deployments, password-protected wallet (ewallet.p12) and autologin TDE wallet (cwallet.sso) must be available in the location `/opt/oracle/dcs/commonstore/wallets/tde/db_uniquename/`.
- If autologin TDE wallet (cwallet.sso) is not present but password-protected wallet (ewallet.p12) is present, then TDE password must be provided in the register database request.
- For Oracle Database Appliance DB systems deployments, TDE must be configured using the `WALLET_ROOT` parameter and TDE wallets must be at location `+DATA/db_uniquename/tde`.

What is the significance of TDE Wallet Management attribute in database object?

TDE Wallet Management attribute signifies the one responsible for TDE wallet lifecycle management and is available starting from Oracle Database Appliance release 19.12.

What value does TDE Wallet Management attribute in database object take?

The TDE Wallet Management attribute takes the values: ODA, which means Oracle Database Appliance manages the lifecycle management of TDE wallets, EXTERNAL, which means Oracle Database Appliance does not manage the LCM of TDE wallets, and the value NULL which means database is not TDE-enabled.

What are the operations that are not supported by Oracle Database Appliance when TDE Wallet Management is EXTERNAL?

The following operations are not supported by Oracle Database Appliance when TDE Wallet Management is EXTERNAL:

- Re-key of TDE wallet
- TDE wallet password change
- TDE wallet backup
- Restore of TDE wallet
- Cloning of TDE database
- Deletion of TDE wallet on deleting the database

Is it possible to change the TDE Wallet Management attribute of a database from EXTERNAL to ODA?

Yes, the TDE Wallet Management attribute of a database can be changed from EXTERNAL to ODA if the following conditions are met:

- TDE must be configured using Software Keystore.
- TDE wallet must be present in the path `/opt/oracle/dcs/commonstore/wallets/tde/db_unique_name`.
- The database version must be 18c or later.

Is it possible to change the TDE Wallet Management attribute of a database from ODA to EXTERNAL?

No, it is not possible to change the TDE Wallet Management attribute of a database from ODA to EXTERNAL.

How can the TDE Wallet Management of a database be changed from EXTERNAL to ODA?

The TDE Wallet Management attribute of a database be changed from EXTERNAL to ODA in the following ways:

- Using the Modify Database operation, for database versions 18c or later.
- While upgrading the database from Oracle Database release 12c to 18c or later.

The `odacli list-backupreports` and `odacli describe-backupreport` commands do not complete quickly. What could cause it and how can I manage these backup reports?

The slowness in listing backup reports or describing a backup report could be because there may be many backup reports accumulated in the metadata. To resolve this issue, run the `odacli delete-backup` command. For Oracle Database Appliance release 19.12 or earlier, the `odacli delete-backup` command only deletes the obsolete backups for a database. Starting with Oracle Database Appliance release 19.13, the `odacli delete-backup` command also deletea the backup reports associated with such obsolete database backups. For example: `odacli delete-backup -in myDBname`

Is RMAN backup password, which is used to encrypt backup, required for TDE-enabled databases?

No, if a database is TDE-enabled, RMAN backup password is not required while modifying the database with the Objectstore or NFS backup configuration or while restoring the database.

Can the database be restored using regular L1 backup?

Yes, the database can be restored from a regular L1 backup. Ensure that the corresponding regular L0 backup which is the base for the regular L1 backup is also present in the backup location. Note that if a regular L0 backup is taken to the disk backup location and its corresponding incremental regular L1 backup is taken to either NFS or Objectstore backup location, then this regular L1 backup can not be used to restore the database.

Should the NFS mount point be configured on both the nodes if the target system is a high-availability system?

Yes, Oracle Database Appliance requires that the nodes are able to access the NFS client location.

What are the NFS mount options that Oracle Database Appliance supports?

The NFS mount options that Oracle Database Appliance supports are (`rw, sync, no_root_squash`). From Oracle Database Appliance release 19.17, you do not need to specify the 'no_root_squash' NFS option.

Is it mandatory that the ownership of the NFS server location must be the same as the `DB_Home/bin/oracle` binary of the source database?

Yes, since the backup files created by the `DB_Home/bin/oracle` binary have the same permissions as the `DB_Home/bin/oracle` directory, Oracle Database Appliance

requires the ownership of the NFS server location to be same as `DB_Home/bin/oracle` of the source database.

Should the backup location for database and TDE wallet always be separate, if the source database is TDE enabled?

Yes, the backup location of the database and TDE wallet must be different. To use the same NFS client location, the absolute paths must be different inside the NFS client location.

Creation of NFS backupconfig fails with the DCS-10074:User oracle does not have permission to access /nfs_backup_client error.

Make sure the ID of the `oracle` user is the same on the system where the NFS client and NFS server location is configured.

Creation of NFS backupconfig fails with the DCS-10013:Input parameter 'BackupLocation' contains unacceptable value: /nfs_backup_client.Directory: /nfs_backup_clientdoes not exist. error.

Make sure that the NFS client location exists on both the nodes in case of high-availability systems. Also, make sure the `'no_root_squash'` NFS option is used while exporting the NFS server location to the NFS client location. For Oracle Database Appliance release 19.16 and earlier, you must specify the `no_root_squash` parameter. Starting with Oracle Database Appliance release 19.17, you need not specify the `no_root_squash` parameter.

How can a TDE-enabled database with TDE wallet management set to value `EXTERNAL` be created?

You can create a TDE-enabled database with TDE wallet management set to value `EXTERNAL` as follows:

- Registering a TDE-enabled database, which had used hardware keystore for TDE configuration.
- Registering a TDE-enabled database in Oracle Database Appliance 19.12 release, which had used software keystore for TDE configuration.

ODACLI Commands

Understand FAQs on ODACLI commands.

What is the difference between `odacli` and `odaadmcli` commands?

ODACLI commands are used to perform most of the lifecycle management operations for the appliance, network, databases, data guard, application VMs and DB systems. ODAADMCLI is used to perform hardware management and monitoring operations.

What is the difference between `oakcli` and `odacli` commands?

OAKCLI commands are deprecated from Oracle Database Appliance release 19.10 on bare metal deployments. OAKCLI commands can be used on Oracle VM (Xen) deployments for lifecycle management of virtual resources. You can use ODACLI to perform the corresponding oakcli functionality on bare metal systems. ODACLI is also enhanced in each Oracle Database Appliance release and enables management of KVM and DB systems on Oracle Database Appliance.

Can I run `odacli` or `odaadmcli` commands as a non-root user?

No, only `root` user or `root` `sudoers` can run ODACLI and ODAADMCLI commands.

When I run `odacli` commands, why is the error `DCS-10033:Service DCS_AGENT is down. displayed?`

ODACLI commands communicates with the DCS agent to perform most operations on the appliance. Ensure that the system service `initdcsagent` is running properly. The ODACLI commands also communicate with the DCS admin to perform some patching operations. Ensure that the system service `initdcsadmin` is running properly if you face any issue communicating with that service.

```
# systemctl status initdcsagent
* initdcsagent.service - Oracle dcs-agent startup
   Loaded: loaded (/etc/systemd/system/initdcsagent.service; enabled;
 vendor preset: disabled)
   Active: active (running) since (...
 Main PID: 1512 (sh)
    Tasks: 45
   Memory: 352.9M
 (...)
```

How can I get more information about `odacli` commands?

The *Oracle Database Appliance Deployment and User's Guide* for your hardware model contains detailed information about each ODACLI command. You can use the command line to view the help for each command.

To list all the actions and object categories use: `odacli -h`

To list all the operations against an object category, use: `odacli object -h`

```
# odacli database -h
Command usage:
  database:
    clone-database
    create-database
    delete-database
    describe-database
    irestore-database
    list-databases
    modify-database
    move-database
    recover-database
    register-database
    upgrade-database
```

To list all the operations against an action, use: `odacli action -h`

```
# odacli create -h
Command usage:
  create:
    create-appliance
```

```

create-auto-logclean-policy
create-backup
create-backupconfig
create-cpupool
create-database
create-dbhome
create-dbstorage
create-dbsystem
create-jobdata-retention-policy
create-logcleanjob
create-network
create-objectstoreswift
create-prepatchreport
create-purge-jobdata-job
create-vdisk
create-vm
create-vmstorage
create-vnetwork

```

To get the help of a specific command, use: `odacli command -h`

```

# odacli create-database -h
create-database
    Creates a new database

```

Syntax

```

create-database -n db_name [-nn associated_network_names]
[-bi backup_config_id] [-bn backup_config_name] [-bp] [-c]
[-cs character_set] [-cp cpu_pool] [-u database_unique_name]
[-de db_editon] [-f] [-rd db_redundancy] [-cl db_class] [-co]
[-dn db_domain_name] [-dh db_home_id] [-l db_language]
[-s db_shape] [-r db_storage] [-dt db_territory]
[-y {RAC|RACOne|SI}] [-no-sh] [-sh] [-t] [-fc] [-io]
[-lb level_zero_backup_day] [-ns national_character_set] [-no-c]
[-no-f] [-no-co] [-no-fc] [-d pdb_adminuser_name] [-p pdb_name]
[-th db_target_node_name] [-g db_target_node_number]
[-v db_version] [-j] [-h]

```

Parameters

Options

Description

```

-----
--dbname, -n (*)           Database
name

--associated-networks, -nn Associated network names (in
format                    networkname1, networkname2, ...)

(...)

```

Can I change the default port to communicate with DCS agent?

Yes, by default the DCS agent service uses port 7070 but you can edit the file `/opt/oracle/dcs/dcscli/dcscli.conf` to change the property `"AgentPort=port"`. Note: by default the DCS admin service uses port 7060, you can edit the same file to edit the property `"AdminPort=port"` to change that port.

Where are the logs for odacli command?

The logs are in the `/opt/oracle/dcs/dcscli/log` directory. Most of the operations log entries are available in the `dcscli.log`. VM and DB system operations log entries are available in `dcscli_kvm.log`.

How I can process the output of odacli command more easily?

Each ODACLI command has the `-j` option which provides a JSON formatted output that may be processed by any other process. For complex integration, use the Oracle Database Appliance Java SDK.

Where I could check the cause and action for each DCS error message I get after running an odacli command?

You can use the tool `dcserr` to see the cause and action for any DCS error message displayed when you run an ODACLI command:

```
# dcserr
10037, System_Not_Provisioned, "System is not yet Provisioned."
// *Cause: System was not yet Provisioned.
// *Action: Contact Oracle Support Services.
/
```