Oracle® Database Appliance

Frequently Asked Questions

Release 19.29
G42413-01
November 2025

Oracle Database Appliance Frequently Asked Questions, Release 19.29

G42413-01

Copyright © 2022, 2025, Oracle and/or its affiliates.

Primary Author: Aparna Kamath

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

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

If this is software, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

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

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

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

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

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

Contents

Copyright © 2022, 2025, Oracle and/or its affiliates.

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

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 customer access to and use of Oracle support services will be pursuant to the terms and conditions specified in their Oracle order for the applicable services.

Related Documents

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

Copyright © 2022, 2025, Oracle and/or its affiliates.

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.
italic	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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

1

Oracle Database Appliance Frequently Asked Questions

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

- Patching Oracle Database Appliance
  FAQs on patching Oracle Database Appliance.
- TDE Keys on Oracle Key Vault for TDE-Enabled Databases
  Understand general FAQs on storing TDE keys on Oracle Key Vault for TDE-enabled databases on Oracle Database Appliance.
- 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.
- Multiple Databases in DB Systems
  FAQs on creating multiple database in 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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Patching Oracle Database Appliance

FAQs on patching Oracle Database Appliance.

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, you can patch your deployment from Oracle Database Appliance release 19.25, 19.26, 19.27, and 19.28.

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.

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 X11 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 19.29, when patching from 19.28 to 19.29. 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 1-1 Time Taken for Patching Oracle Database Appliance X11 Bare Metal Systems

Component Being Patched	Time taken for patching X11-HA	Time taken for patching X11-L
Patching DCS stack (running the commands `odacli update-dcsadmin`, `odacli update-dcscomponents`)	approximately 9-12 minutes to complete both commands	approximately 5-8 minutes to complete both commands
Server components prechecks (running the command `odacli create-prepatchreport -sc`)	approximately 3-5 minutes	approximately 2-4 minutes
Server patching - updating operating system, Oracle ILOM, and firmware (running the command `odacli update-server`)	approximately 15-20 minutes for each node	approximately 7-12 minutes
Create the precheck report for Oracle Grid Infrastructure	approximately 2-5 minutes for each node	approximately 2-5 minutes

Copyright © 2022, 2025, Oracle and/or its affiliates.

Table 1-1 (Cont.) Time Taken for Patching Oracle Database Appliance X11 Bare Metal Systems

Component Being Patched	Time taken for patching X11-HA	Time taken for patching X11-L
Oracle Grid Infrastructure patching	approximately 10-15 minutes for each node	approximately 5-8 minutes
Oracle Database prechecks for a single database in an Oracle home (running the command `odacli create-prepatchreport --dbhome`)	approximately 7-10 minutes	approximately 2-5 minutes
Patching Oracle Database 19c for a single database in an Oracle home (running the command `odacli update-dbhome`)	approximately 9-12 minutes	approximately 5-8 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 19.29, when patching from 19.28 to 19.29.

Table 1-2 Time Taken for Patching Oracle Database Appliance DB Systems

Component Being Patched	Time taken for patching DB System on X11-HA	Time taken for patching DB System on X11-L
Patching DCS stack (running the commands `odacli update-dcsadmin`, `odacli update-dcscomponents`)	approximately 10-12 minutes to complete both commands	approximately 4-8 minutes to complete both commands
Update server components	approximately 2-4 minutes	approximately 1-3 minutes
Server patching - updating operating system, Oracle Grid Infrastructure, Oracle ILOM, and firmware (running the command `odacli update-server`)	approximately 5-8 minutes for both nodes	approximately 3-6 minutes
Oracle Grid Infrastructure prechecks	approximately 3-5 minutes for both nodes	approximately 2-4 minutes
Oracle Grid Infrastructure patching	approximately 5-8 minutes for both nodes	approximately 3-6 minutes
Oracle Database prechecks for a single database in an Oracle home (running the command `odacli create-prepatchreport --dbhome`)	approximately 5-8 minutes	approximately 3-7 minutes
Patching Oracle Database 19c for a single database in an Oracle home (running the command `odacli update-dbhome`)	approximately 7-10 minutes	approximately 10-15 minutes

Copyright © 2022, 2025, Oracle and/or its affiliates.

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` command, and hence you must patch the DCS admin with the `odacli update-dcsadmin` command first. Run the `odacli update-dcscomponents` command next.

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, Highly Available Metadata Infrastructure (Oracle HAMI), and MySQL components.

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

The command `odacli update-gihome` updates Oracle Grid Infrastructure.

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

Copyright © 2022, 2025, Oracle and/or its affiliates.

The command `odacli update-database` updates a database to a specific database home with later version, and within the same major release.

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 currenly 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 currenly 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 precheck report, generated with the command `odacli create-prepatchreport -sc` 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?

For various server component patching, analyse the failure and apply corresponding fixes manually and then re-run the `odacli update-server` command. Running the `odacli update-server` command again after you fix the root cause is generally successful. In other cases, check with Oracle Support for further assistance.

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` and `odacli update-dcscomponents` 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

Copyright © 2022, 2025, Oracle and/or its affiliates.

not present in the Oracle Database Appliance patch bundled 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`).

Copyright © 2022, 2025, Oracle and/or its affiliates.

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 Oracle Database Appliance release 19.11, Oracle Databse Appliance 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 releases later than 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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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

TDE Keys on Oracle Key Vault for TDE-Enabled Databases

Understand general FAQs on storing TDE keys on Oracle Key Vault for TDE-enabled databases on Oracle Database Appliance.

What is Oracle Key Vault server configuration?

Oracle Key Vault server configuration is an Oracle Database Appliance entity which represents the metadata about the Oracle Key Vault server such as the IP address, host name, the user name, and a brief description about the Oracle Key Vault server. The Oracle Key Vault user password must be provided while creating the Oracle Key Vault server configuration object. The client auto-login wallets are generated with the specified passwords.

What are the four different ways of creating a TDE-enabled database on Oracle Database Appliance which uses Oracle Key Vault to store TDE keys?

Oracle Database Appliance supports the following two ways of creating TDE enabled database that uses Oracle Key Vault to store TDE keys:

- **Create TDE-enabled databases using Oracle Key Vault with endpoints in Oracle Key Vault:** The virtual wallet and the endpoints are not created by Oracle Database Appliance tooling but created by the user, on Oracle Key Vault server. Oracle Database Appliance tooling only expects the user to specify the okvclient.jar file corresponding to the created endpoint during database creation.
- **Create TDE-enabled databases with user credentials:** In this case, the virtual wallet and the endpoints are created by Oracle Database Appliance tooling, on Oracle Key Vault server. The user must create an Oracle Key Vault server configuration corresponding to the required Oracle Key Vault server and specify it during database creation.
- **Create and Restore TDE-enabled databases manually and register it with Oracle Database Appliance using Oracle Key Vault server configuration object:** The database can be either created or restored manually from backups such that it uses Oracle Key Vault to store the TDE keys and then it can be registered using Oracle Key Vault server configuration object. While registering the database using Oracle Key Vault server configuration, the current TDE password must be provided, so that Oracle Database Appliance tooling can change it to a new random password. This new random password is managed by Oracle Database Appliance tooling and is not known to the user.

Copyright © 2022, 2025, Oracle and/or its affiliates.

- **Create and Restore TDE-enabled databases manually and register it with Oracle Database Appliance without using Oracle Key Vault server configuration object:** The database can be either created or restored manually from backups such that it uses Oracle Key Vault to store the TDE keys and then it can be registered without using Oracle Key Vault server configuration object. Registering the database without using Oracle Key Vault server configuration does not change the current TDE password.

How do I get to know the keystore type of my TDE database?

The `keystoreType` parameter value in the output of `odacli describe-database -n` *`dbname`* `-j` command displays the type of the keystore used by the TDE database.

How do I identify whether a TDE-enabled database is using software keystore or Oracle Key Vault for TDE configuration?

If the `keystoreType` attribute of the database has value `software`, then it means software keystore was used to configure TDE. Else, if the same attribute has OKV as its value, then it means Oracle Key Vault was used to configure TDE.

When should I choose the option of creating TDE-enabled databases using Oracle Key Vault with endpoints in Oracle Key Vault?

If you do not want to use Oracle Key Vault user credentials in Oracle Database Appliance, then you can use the option of creation of TDE-enabled databases that use pre-created virtual wallet and endpoints.

When should I choose the option of creating TDE-enabled databases using Oracle Key Vault with credentials of the user on Oracle Key Vault?

If you want Oracle Database Appliance to interact with Oracle Key Vault directly using the credentials of the Oracle Key Vault user to create virtual wallets and endpoints, then use this option.

Which user can create Oracle Key Vault server configuration?

In multi-user access or multi-user access-passwordless enabled systems, the odaadmin user and users with the ODA-OKVCONFIGADMIN role can create the Oracle Key Vault server configuration. The Oracle Key Vault server configuration can then be shared with the required DB user, that is, user with the ODA-DB role, with only odaadmin so that DB user can create TDE database using the shared Oracle Key Vault server configuration object. This ensures that the Oracle Key Vault server credentials used in creating the Oracle Key Vault server configuration object is not shared with the DB user. However, in case of multi-user access enabled systems, the Oracle Key Vault server configuration is created by the same user who creates the database.

When the database is created using Oracle Key Vault server configuration, Oracle Database Appliance tooling generates a random password and stores it in the wallet on the Oracle Key Vault server. The format of the name of the wallet is *db_unique_name*_pass_on_*oda-cluster-name*. This wallet is not deleted when the database is deleted.

Can a TDE database that uses software keystore be converted to use Oracle Key Vault?

Yes. See the *Oracle Database Appliance Deployment and User Guide* for your hardware model to view the manual steps to migrate from software keystore to Oracle Key Vault. Migration using ODACLI is not supported.

Copyright © 2022, 2025, Oracle and/or its affiliates.

What is the Oracle Key Vault Server port number that must be opened for Oracle Database Appliance to communicate with it?

Oracle Database Appliance relies on port number 5695 to communicate with Oracle Key Vault server. Ensure that port number 5695 is opened on the Oracle Key Vault server so that Oracle Database Appliance can communicate with it.

What is the version of Oracle Key Vault server that Oracle Database Appliance recommends?

Oracle Database Appliance recommends the version of the Oracle Key Vault server to be 21.10.

What are the database versions that support Oracle Key Vault feature on Oracle Database Appliance?

The Oracle Key Vault feature on Oracle Database Appliance is supported with Oracle Database 19c, on both bare metal and DB systems.

What are the database lifecycle management operations available on Oracle Database Appliance, but not supported for TDE databases that use Oracle Key Vault to store TDE keys?

Database lifecycle management operations such as cloning, upgrading, and registering databases are currently not supported for TDE databases that use Oracle Key Vault to store TDE keys.

Should I provide TDE password while creating the database that uses Oracle Key Vault to store TDE keys?

No, the TDE password is not required. The TDE password is randomly generated by Oracle Database Appliance tooling. The TDE password must be provided while creating the database with endpoints and wallet manually created on Oracle Key Vault.

Once the user with ODA-DB role is granted access to the Oracle Key Vault server configuration object, can they delete the Oracle Key Vault server configuration object?

No, the Oracle Key Vault server configuration object can only be deleted by the user who created it or the odaadmin user. The user with ODA-DB role can only use the shared Oracle Key Vault server configuration object and create TDE database but not delete the shared Oracle Key Vault server configuration object. Also, the database that used the Oracle Key Vault server configuration object during creation, must be deleted before deleting the Oracle Key Vault server configuration object.

Does Oracle Database Appliance tooling support backup and recovery of TDE wallet of a database which uses Oracle Key Vault to store TDE keys?

No, Oracle Database Appliance tooling does not support backup and recovery of TDE wallet since the TDE wallet is present outside the Oracle Database Appliance system, in an Oracle Key Vault server. The OKV ADMIN of the Oracle Key Vault server have to manage the backup and recovery of TDE wallet. When you select Oracle Key Vault server, you have the option of TDE wallet management outside the client, in an Oracle Key Vault server.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Does Oracle Database Appliance tooling identify whether the given Oracle Key Vault credentials belong to a user who has the least privilege of `Create Endpoint` on Oracle Key Vault server ?

No, Oracle Database Appliance tooling does not identify whether the given Oracle Key Vault credentials belong to a user who has the least privilege of `Create Endpoint` on Oracle Key Vault server. Ensure that you use the credentials of the least privileged user. This least privilege of `Create Endpoint` is recommended so that the given credentials can only be used to onboard the database on to the Oracle Key Vault server and no other operations are performed on the Oracle Key Vault server.

Why do I need to create a copy of the TDE wallet before restoring a TDE database with Oracle Key Vault keystore?

Oracle Database Appliance tooling does not support backup, recovery, and restore of the TDE wallet if the wallet is stored in the Oracle Key Vault server because the wallet is present outside the appliance. However, if you are restoring the database as a standby in an Oracle Data Guard configuration, then you must specify the wallet of the primary database.

Is NTP configuration required on Oracle Database Appliance to create a database with TDE configuration using Oracle Key Vault?

It is recommended to set up NTP on Oracle Database Appliance because the database creation or restore operation may fail if the clock on Oracle Database Appliance and the Oracle Key Vault server are not synchronized. Hence, you must ensure that the time is consistent across the servers or use NTP on Oracle Database Appliance. While using an NTP server, ensure that the server is reachable from the appliance and that the chrony settings in the `/etc/chrony.conf` file are the same on Oracle Database Appliance and the Oracle Key Vault server.

Why does re-key of TDE wallet fail with `ORA-28353: failed to open wallet` error?

The error code ORA-28353 could be due to incorrect TDE password. Retry the operation with the correct TDE password to resolve the issue. For more information about this error, see the *Oracle Database Error Messages Guide* at https://docs.oracle.com/en/database/oracle/oracle-database/19/errmg/ORA-24280.html.

Why does creation of a database with TDE configuration using Oracle Key Vault as keystore fail with an error?

You may encounter the following error:

```
Internal error encountered: PL/SQL procedure successfully completed.begin
*
ERROR at line 1:
ORA-00600: internal error code, arguments: [kcbtse_populate_tbskey_1], [The
request operation was denied.], [], [], [], [], [], [], [], [], []
ORA-06512: at line 2.
```

Creation of the database may fail with the above error, when setting the TDE master encryption key. This error is caused when the clock on the appliance is not synchronized with the clock on the Oracle Key Vault server, causing a lag between the time on Oracle Database Appliance and the Oracle Key Vault server. Due to the lag, the time at which the command to set the TDE master encryption key for the database is run does not match with the time on the Oracle Key Vault server, causing the operation to fail. To synchronize the clocks, set up NTP on Oracle Database Appliance. While setting up NTP, ensure that the NTP server is reachable

Copyright © 2022, 2025, Oracle and/or its affiliates.

from Oracle Database Appliance, and that the chrony settings are the same on Oracle Database Appliance and the Oracle Key Vault server.

Run the following commands to set up NTP manually:

```
systemctl stop chronyd


cat /etc/chrony.conf
server Enter NTP server IP address iburst
driftfile /var/lib/chrony/drift
makestep 1.0 -1
rtcsync
logdir /var/log/chrony


systemctl start chronyd


systemctl enable chronyd
```

Why does creation of a database with TDE configuration using Oracle Key Vault as keystore fail with an error?

You may encounter the following error:

```
{ "result" : "Failure", "message" : "Error occurred during install of Oracle
Key Vault endpoint software. Check log files for more information. Please
cleanup all of the files and directories created by the failed installation
attempt before re-install" }.
```

Creation of the database may fail with the above error, when setting the TDE master encryption key . This error is caused when the clock on the appliance is not synchronized with the clock on the Oracle Key Vault server, causing a lag between the time on Oracle Database Appliance and the Oracle Key Vault server. Due to the lag, the certificate generated for the Oracle Key Vault endpoint is not valid when the command to install the endpoint software is run on Oracle Database Appliance, causing the operation to fail. To synchronize the clocks, set up NTP on Oracle Database Appliance.

Run the following commands to set up NTP manually:

```
systemctl stop chronyd


cat /etc/chrony.conf
server Enter NTP server IP address iburst
driftfile /var/lib/chrony/drift
makestep 1.0 -1
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
rtcsync
logdir /var/log/chrony


systemctl start chronyd


systemctl enable chronyd
```

How is TDE password handled when the database is created with endpoints in Oracle Key Vault and when the database is created with user credentials?

When the database is created with endpoints in Oracle Key Vault, Oracle Database Appliance requires you to provide the TDE password. This password is not known to Oracle Database Appliance and the same password must be provided while performing operations such as rekey and change of TDE password. When the database is created with user credentials, the TDE password is randomly generated by Oracle Database Appliance tooling and is not known to the user. This password is stored in a wallet on the Oracle Key Vault server. You need not specify any TDE password while performing operations such as rekey and change of TDE password.

Registering TDE enabled database that uses Oracle Key Vault to store TDE keys using Oracle Key Vault server configuration is failing with the `ORA-28353: failed to open wallet` error.

This error may occur when an incorrect TDE password is specified. Ensure that you specify the correct TDE password when registering the database and then retry the operation.

A TDE enabled database that uses Oracle Key Vault to store TDE keys was first registered using Oracle Key Vault server configuration. Later, when the same database is re-registered using the same Oracle Key Vault server configuration and the same TDE password, it fails with the `DCS-12737:Failure to change the TDE password for the database wallet '/etc/OKV/db_unique_name/okv/ssl'` error.

When the Oracle Key Vault TDE database is registered using Oracle Key Vault server configuration, its current TDE password is changed to a new random password by the Oracle Database Appliance tooling. This new random password must be specified while re-registering the database using Oracle Key Vault server configuration. The random password generated by Oracle Database Appliance tooling is stored in a wallet created on the Oracle Key Vault server. Following are the steps to retrieve the stored random password. After the password is retrieved, the re-register operation using Oracle Key Vault server configuration can be retried using the retrieved password. For non-multi user access enabled environment, switch to the `oracle` user. For multi-user access enabled and multi-user passwordless environment, switch to the Oracle Key Vault user.

Re-enroll the password store endpoint whose name will be of the format *db_name*_pass_on_*cluster_name*. This endpoint uses the password store wallet where the random password is stored, as the default wallet.

```
 /etc/OKV/okv_server_config_name/bin/okv admin endpoint re-enroll --endpoint
db_name_pass_on_cluster_name
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

Download the `okvclient.jar` file corresponding to the endpoint *db_name*_pass_on_*cluster_name* to some temporary location, for example, `/tmp`.

```
/etc/OKV/<okv_server_config_name>/bin/okv admin endpoint download --endpoint
db_name_pass_on_cluster_name --location /tmp
```

Go to the temporary location where the `okvclient.jar` file is downloaded and install the endpoint. When prompted for the password, press Enter.

```
cd /tmp/db_name_pass_on_cluster_name

    /opt/oracle/dcs/java/1.8.0_xxx/bin/java -jar okvclient.jar
    Detected JAVA_HOME: /opt/oracle/dcs/java/1.8.0_xxx
    Enter new Key Vault endpoint password (enter for auto-login):
    The endpoint software for Oracle Key Vault installed successfully.
    Deleted the file : /tmp/db_name_pass_on_cluster_name/okvclient.jar
```

Run the following command by replacing the appropriate value to retrieve the UUID corresponding to the stored TDE password.

```
/etc/OKV/okv_server_config_name/bin/okv managed-object object locate --
okv_client_config /tmp/db_name_pass_on_cluster_name/conf/okvclient.ora --
output_format "TEXT" --state "ACTIVE" --custom-attribute '[{"name": "x-
REFRESH_STATE", "value": "0"}, {"name": "x-key", "value": "db_resource_id-
db_unique_name"}]'
```

Retrieve the TDE password corresponding to the UUID.

```
/etc/OKV/okv_server_config_name/bin/okv managed-object secret get --
output_format "TEXT" --okv_client_config /tmp/db_na
```

When should the endpoints of a database be updated?

The endpoints of a database must be updated whenever the Oracle Key Vault server where the corresponding TDE wallet is present is updated to newer version and whenever the certificates on the Oracle Key Vault server is updated.

Can the endpoints of selective databases among the list of Oracle Key Vault TDE databases created in the system be updated?

Yes, only those databases whose endpoints must be updated can be mentioned in Oracle Key Vault client JSON file, which is used as input in Update Endpoints request.

Why does the overall status of the update endpoints job says 'Success' even though Update endpoints tasks has failed for some databases?

To ensure that the update endpoints job is run for all specified databases, even if there is failure for some databases in between the job run, the overall status of the job is marked as 'Success'.

Update endpoint job of some databases have passed and some have failed. What should I do to update the endpoints of those failed database?

Remove entries corresponding to successful databases in the input Oracle Key Vault Client JSON file and keep the entries corresponding to failed databases only and then retry the

Copyright © 2022, 2025, Oracle and/or its affiliates.

operations. Before retrying, make sure the endpoints are re-enrolled and the new 'Oracle Key Vaultclient.jar' corresponding to the re-enrolled endpoints are specified in the Oracle Key Vault client JSON file.

What are Oracle Key Vault PKCS library and Oracle Key Vault REST Client library?

The Oracle Key Vault PKCS library (liborapkcs.so) present at '/opt/oracle/extapi/64/hsm/oracle/1.0.0/' location is to communicate with the Oracle Key Vault Server. Oracle Key Vault REST Client library (Oracle Key Vaultrestcli.jar) present at '/etc/Oracle Key Vault/<Oracle Key Vault_server_config_name>/lib' location is used to make REST API calls to the Oracle Key Vault Server.

When should the Oracle Key Vault PKCS library and Oracle Key Vault REST Client library be updated?

The Oracle Key Vault PKCS library and Oracle Key Vault REST Client library must be updated only when the corresponding Oracle Key Vault Server is updated to a later version.

There are no Oracle Key Vault Server Config objects created in the system. Is it necessary to update the Oracle Key Vault REST Client library?

If there are no Oracle Key Vault Server Config objects created in the system, then the update Oracle Key Vault endpoints job will skip the task of updating the Oracle Key Vault REST Client library.

Why is database restart required while updating the Oracle Key Vault PKCS library?

Since the Oracle Key Vault PKCS library is shared by all Oracle Key Vault TDE databases present in the system, a database restart is required after updating the Oracle Key Vault PKCS library.

Is it mandatory to update all endpoints available in the system before updating the Oracle Key Vault PKCS library?

Yes, it is mandatory to update all endpoints available in the system before updating the Oracle Key Vault PKCS library.

When should I register a database using Oracle Key Vault server configuration?

If the TDE password must be handled by Oracle Database Appliance tooling itself and not by the user, then the database can be registered using Oracle Key Vault server configuration. The current TDE password will be changed to a new random password by Oracle Database Appliance tooling internally.

When should I register a database without using Oracle Key Vault server configuration?

If the TDE password need not be handled by Oracle Database Appliance tooling and user would be handling the TDE password themselves, then the database can be registered without using Oracle Key Vault server configuration. The TDE password will not be changed during the register operation.

An Oracle Key Vault TDE database is already registered with Oracle Key Vault server configuration. Can it be re-registered with or without Oracle Key Vault server configuration?

Once the database is registered with Oracle Key Vault server configuration, its TDE password will be changed to a new random password which will be maintained by Oracle Database Appliance tooling and not known to the user. Since current TDE password is unknown, such databases cannot be Re-Registered, with or without Oracle Key Vault server configuration.

Copyright © 2022, 2025, Oracle and/or its affiliates.

An Oracle Key Vault TDE database is already registered without Oracle Key Vault server configuration. Can it be re-registered again with or without Oracle Key Vault server configuration?

Yes, an Oracle Key Vault TDE database which was registered without Oracle Key Vault server configuration can be re-registered with or without Oracle Key Vault server configuration. However when registered with Oracle Key Vault server configuration, the current TDE password will be changed to a new random password which will be maintained by Oracle Database Appliance tooling and not known to the user.

Is high availability enabled for single-instance Oracle Key Vault TDE database being registered on Oracle Database Appliance high availability system?

No. By default, high availability is not enabled for single-instance Oracle Key Vault TDE database being registered on Oracle Database Appliance high availability system.

How can I identify whether the database was earlier registered with or without Oracle Key Vault server configuration object?

If the Oracle Key `VaultServerConfigName` database attribute is set to the name of the Oracle Key Vault server configuration object that was used during database register operation, then the database was registered with Oracle Key Vault server configuration object. If the Oracle Key `VaultServerConfigName` database attribute is set to `NONE`, then the database was registered without Oracle Key Vault server configuration object.

Before registering the Oracle Key Vault TDE database, can the endpoint software files be copied from the existing location to the required /etc/Oracle Key Vault/ <db_unique_name>/Oracle Key Vault location?

No, the endpoint software files must be installed again at the location /etc/Oracle Key Vault/ <db_unique_name>/Oracle Key Vault and not copied, before registering.

After registering an Oracle Key Vault TDE database with the Oracle Key Vault server configuration object, an external store autologin wallet (cwallet.sso) is created in the /etc/Oracle Key Vault/*db_unqiue_name*/tde_seps location by Oracle Database Appliance tooling. Why is it created? Can it be deleted?

The external store autologin wallet is created while registering an Oracle Key Vault TDE database with Oracle Key Vault server configuration object, which is similar to creating the Oracle Key Vault TDE database with Oracle Key Vault server object. It is created because Oracle Database Appliance tooling would be using the clause 'identified by EXTERNAL STORE' and not the actual TDE password identified by `tde_password` in running Administer Key Management commands. This External Store Autologin wallet should not be deleted since its deletion would cause the TDE operations such as rekey through Oracle Database Appliance tooling to fail.

Related Topics

- Troubleshooting Oracle Database Appliance

Copyright © 2022, 2025, Oracle and/or its affiliates.

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 19.29?

No, SNMP cannot be configured in Oracle Database Appliance release 19.29 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
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

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 19.29?

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 and later?

In Oracle Database Appliance release 19.20.0.1, Oracle ASR is configured with SNMP protocol for fault alerts. In release 19.21 and later, 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`

Copyright © 2022, 2025, Oracle and/or its affiliates.

`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

Copyright © 2022, 2025, Oracle and/or its affiliates.

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
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

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.SA
XParserFactoryImpl -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

Copyright © 2022, 2025, Oracle and/or its affiliates.

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 1-3

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	`dcs-agent.json`	`dcs-agent.yml`	`dcs-agent-logback.xml`
DCS admin	`dcs-admin.json`	`dcs-admin.yml`	`dcs-admin-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`

Copyright © 2022, 2025, Oracle and/or its affiliates.

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>
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

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>
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

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>
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

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.
- `acrmap` 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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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 ifnormation, 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

Copyright © 2022, 2025, Oracle and/or its affiliates.

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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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-`odadmin` 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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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/reinstate?

It is recommended to use ODACLI for switchover/failover/reinstate 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?

Oracle Database Appliance tooling supports multiple 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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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?

Prior to Oracle Database Appliance release 19.23, the DB system supported one database per DB system. Starting with Oracle Database Appliance release 19.23, a DB system can run multiple databases.

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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

How many cores does a DB system utilize?

The DB system shape decides the cores a DB system can utilize. For example, a DB system with dbsX shape uses X CPU cores, with default memory set to 8 times of CPU core X. You can customize the memory used by the DB system by specifying the memory size at the time of DB system creation, either in the JSON file or in the BUI. Starting Oracle Database Appliance release 19.29, you can scale up and scale down the CPU cores for a DB system using the `odacli modify-dbsystem --core` command.

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 pages, and the DB system is configured using huge pages. So the total memory used by the DB system must be less than the huge pages allocated at the host. If there are databases and application VMs running on the host, then you must also consider them when you plan for CPU and memory resource for DB systems. It is recommended that you reserve at least 2 CPU cores and 24 GB regular memory for the bare metal system host, that is, hypervisor. Check /proc/meminfo for memory allocation details. Use the `odacli modify-osconfigurations` command to update the Huge Pages.

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. Starting with Oracle Database Appliance release 19.24, you can use the `odacli`

Copyright © 2022, 2025, Oracle and/or its affiliates.

`modify-dbsystem` command to attach a vdisk to a DB system. Then, use the vdisk to manually create new file system mount point in the DB system.

The DB system fails to be created due to lack of memory error. But the operating system 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.

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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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 19.29, you can deploy DB system running Oracle Database Appliance release 19.29. However, you can choose earlier releases of Oracle Grid Infrastructure and Oracle Database, for example, you can deploy Oracle Grid Infrastructure 19.28 and Oracle Database 19.28 on a DB system running Oracle Database Appliance release 19.29. 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. Starting with Oracle Database Appliance release 19.29, you can decrease the DB system memory using the `odacli modify-dbsystem --memory` command. This option does not change the configuration of the database. Make sure configuration of the database is updated first before decreasing the memory of the DB system; otherwise, databases may fail to start after you decrease 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?

Prior to Oracle Database Appliance release 19.23, modifying 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. Starting ith Oracle Database Appliance release 19.23, one DB system supports running multiple databases. Changing the DB system shape using the `odacli modify-dbsystem -s` command does not change the databases in the DB system anymore.

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

On Oracle Database Appliance releases earlier than 19.29, you can use `odacli modify-dbsystem -s` to revert the changes. For example, if you updated the DB system from the

Copyright © 2022, 2025, Oracle and/or its affiliates.

default dbs4/32G to dbs4/48G, and you want to revert it to dbs4/32G, you can modify the shape from dbs4 to dbs2, and then from dbs2 to dbs4 again. Starting with Oracle Database Appliance release 19.29, decreasing memory of the DB system is supported using the `odacli modify-dbsystem -m` command. Make sure you reduce the memory used by the database first, before decreasing the DB system memory to avoid database startup failure.

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.

Multiple Databases in DB Systems

FAQs on creating multiple database in Oracle Database Appliance DB systems.

What is the multi-database feature in DB system?

Prior to Oracle Database Appliance release 19.23, one DB system supported one database. Starting with Oracle Database Appliance release 19.23, a DB system can run multiple databases and multiple database homes.

I updated my bare metal system to Oracle Database Appliance release 19.23, what happens to the existing DB systems? How do I know if a DB system has multi-database feature enabled?

After the bare metal system is updated to Oracle Database Appliance release 19.23, the new DB system created has multi-database support by default. For existing DB systems, after the DB systems are updated to Oracle Database Appliance release 19.23, either by patching to 19.23 with the `odacli update-server` command or upgrading to 19.23 with the `odacli upgrade-dbsystem` command, multi-database support is enabled. If an existing DB system is not updated to 19.23, the DB system continues to support only one database. To check if the DB system has multi-database enabled, you can view the shape with the `odacli list-dbsystems` command. If the shape is dbsX, then multi-database feature is enabled. If the shape is odbX, then the DB system supports only one database.

How do I choose the DB system shapes to support multiple databases? How many database can be created in the multi-database enabled DB system?

Prior to Oracle Database Appliance release 19.23, one DB system supported one database. Starting with Oracle Database Appliance release 19.23, DB system supports running multiple databases. Selecting the shape of the DB system depends on the sizing requirements of the databases running in the DB system. Similarly, the number of databases supported in the DB system depends on the resources such as CPU, memory of DB system. DB system must be sized properly based on the sizing requirement of databases running in the DB system. Generally, the shapes of databases together running in a DB system must be no greater than the shape of the DB system. For example, if databases odbA, odbB, ..., odbN run in the DB system with shape dbsX, then A+B+...N <= X. You can change the DB system shapes to meet the new sizing requirement of the databases in the DB system.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Why must I manually change the database shapes when modifying DB system shapes in 19.23 and later?

Prior to Oracle Database Appliance release 19.23, one DB system supported one database and the database shape was the same as the DB system shape. Hence, modifying the DB system shape would automatically change the shape of the database. Starting with Oracle Database Appliance release 19.23, multi-database feature is enabled by default on the DB system. When you modify the DB system shape with the `odacli modify-DB system --shape` command, the shape of the DB system is modifed; the command does not change the shapes of the databases inside the DB system any more. You must modify database shapes accordingly based on the new shape of the DB system. For DB system shape scale up, increase the database shapes after DB system shape is scaled up. For DB system shape scale down, scale down database shapes before DB system shape is scaled down; otherwise databases could fail to start because of reduced DB system memory.

After the DB system shape is scaled down, the databases in the DB system no longer starts, why? What is the solution?

After DB system shape is scaled down, the CPU and memory are also reduced, so are the huge pages of the system. When the memory of the system cannot meet the memory requirements of the database, especially if the huge pages of the system cannot meet the SGA of the database, then databases fail to start. To fix the issue, scale up the DB system shape to the original shape, scale down the databases shape, and then scale down the shape of the DB system.

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.

Does Oracle Database Appliance support creation of Oracle Database on application VMs using the vdisk as storage?

Yes. Oracle Database Appliance supports creation of single-instance database using vdisks on application VMs, but you cannot create Oracle RAC databases using vdisks as shared storage on application VMs.

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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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?

After imaging the appliance, about half of the Oracle Database Appliance host memory is allocated to Huge pages pool by default. Huge pages are consumed by the SGA of the databases and DB systems. 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. If Oracle Database Appliance is dedicated to run application VMs, that is, there are no databases on the bare metal system or DB systems, then it is recommended that you reset the HugePages to zero using the odacli modify-osconfigurations -nhp command to release the memory configured as hugeHugePagespage back to the regular memory pool.

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

Copyright © 2022, 2025, Oracle and/or its affiliates.

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, cpupcores 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 `dbs` 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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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 `dbs` type is available. The `dbs` 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-dbsystem` to limit the DB system CPU within the CPUs inside this CPU pool. Multiple DB systems can share one `dbs` CPU pool and can oversubscribe to the `dbs` CPU pool. When DB system is created without specifying the `dbs` 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 `dbs` CPU pool is the DB system CPU pool is manually created through ODACLI commands. The `dbs` CPU pool is assigned to DB systems at creation time or after creation using the command `odacli modify-dbsystem`. The `dbs` 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 `dbs` type is available. The `dbs` 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-dbsystem` to limit the DB system CPU within the CPUs inside this CPU pool. Multiple DB systems can share one `dbs` CPU pool and can oversubscribe to the `dbs` 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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

Can DB systems share the same CPU pool?

Yes, different DB systems can share the same `dbs` 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).

Copyright © 2022, 2025, Oracle and/or its affiliates.

How do I change from SFP28 Copper to SPF28 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

Copyright © 2022, 2025, Oracle and/or its affiliates.

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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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 configure, 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:

Copyright © 2022, 2025, Oracle and/or its affiliates.

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:1vn21ugu1saj1v9i1v941sar1ugw1vo0",
       "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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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. Whey 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 backups 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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

- 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`.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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 Applaince 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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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
```


Copyright © 2022, 2025, Oracle and/or its affiliates.

```
                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`.

Copyright © 2022, 2025, Oracle and/or its affiliates.

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.
/
```


Copyright © 2022, 2025, Oracle and/or its affiliates.