Oracle® Exadata

Exascale User's Guide



26.1
F17209-20
July 2026



Oracle Exadata Exascale User's Guide, 26.1

F17209-20

Copyright © 2022, 2026, Oracle and/or its affiliates.

Primary Author: Peter Fusek

Contributing Authors: Akshay Shah, Alex Blyth, Alex Tsukerman, Allan Graves, Arun Venkataraman, Barb Glover, Boris Erlikhman, Chandra Pabba, Christian Craft, Elisabeth Thibault, Evangelos Vlachos, Frank Kobylanski, Francois Marcoux, Gaurav Chadha, Gautam Bhatt, Jaime Figueroa, Joshua Smith, Kai Zhang, Krishnadev Telikicherla, Kuan-Ju Chen, Maruti Sharma, Min Yun Law, Nilesh Choudhury, Pedro Gonzalez, Ruggero Citton, Saeed Abedigozalabad, Seth Miller, Siddhartha Datta, Soma Prasad, Tony Dziedzic, Vijay Sridharan, Vyomkesh Tripathi

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

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

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

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

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

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

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

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

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

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.

Contents

Preface

Audience i
Related Documents i

1 Oracle Exadata Exascale Overview

1.1 What is Oracle Exadata Exascale? 1
1.2 Exascale Components and Concepts 2
1.2.1 Exascale Cluster 2
1.2.2 Exascale Services 3
1.2.3 Exascale Resources and Attributes 5
1.2.4 Storage Media Types 5
1.2.5 Pool Disks 6
1.2.6 Storage Pools 6
1.2.7 Vaults 7
1.2.8 Files 8
1.2.9 File Storage Attributes 8
1.2.10 Templates 9
1.2.11 Extended Attributes 10
1.2.12 Clone Files 10
1.2.13 Snapshot Files 11
1.2.14 Datasets 11
1.2.15 Exascale Users 12
1.2.16 Exascale User Credentials 13
1.2.17 Cluster Level Privileges 13
1.2.18 Access Control Lists 16
1.2.19 Vault and File Access Control 17
1.2.20 Trust Store 19
1.2.21 Resource Management 19
1.2.21.1 Inter-Vault Resource Management Using Exadata IORM 19
1.2.21.2 Intra-Vault Resource Management Using Exascale Resource Profiles 20
1.2.21.3 Intra-Database Resource Management Using DBRM 23
1.2.22 Exascale Block Store 23
1.2.22.1 Exascale Volumes 24

Copyright © 2022, 2026, Oracle and/or its affiliates.

1.2.22.2 Volume Attachments 24
1.2.22.3 Volume Snapshots 25
1.2.22.4 Volume Clones 25
1.2.22.5 Volume Backups 25
1.2.22.6 Volume Groups 26
1.2.22.7 Oracle Advanced Cluster File System (ACFS) on Exascale 26
1.2.22.8 Virtual Machine Images 27

2 Oracle Exadata Exascale System Administration

2.1 Configure Exascale 1
2.1.1 Configure Exascale Using the OEDA Web User Interface 1
2.1.1.1 OEDA Web User Interface Settings and Options for Exascale 2
2.1.1.2 Exascale Post Deployment Settings 4
2.2 Administer the Exascale Cluster 7
2.2.1 Administer Exascale Services 7
2.2.1.1 About Automatic Service Placement 8
2.2.1.2 List Services 8
2.2.1.3 Alter the Number of Service Instances 9
2.2.1.4 Alter the Placement of Services 10
2.2.1.5 Restart a Service 10
2.2.1.6 Stop a Service 11
2.2.1.7 Start a Service 13
2.2.1.8 Disable a Service 14
2.2.1.9 Administer Client Services 15
2.2.2 Modify the Exascale Cluster 16
2.2.3 Stop the Exascale Cluster 16
2.2.4 Add a Storage Server 16
2.2.5 Manage Certificates for Exascale Control Services 17
2.3 Administer Exascale Users 18
2.3.1 Create a User 19
2.3.2 List User Information 19
2.3.3 Modify a User 20
2.3.3.1 Modify a User Name 20
2.3.3.2 Modify User Privileges 20
2.3.3.3 Manage a User's Public Keys 21
2.3.4 Remove a User 21
2.3.5 Administer the Internal User Accounts 22
2.3.5.1 Administer the ADMIN User 22
2.3.5.2 Administer the Node Administration Users 22
2.4 Administer Exascale Pool Disks 22
2.4.1 Create Pool Disks 23

Copyright © 2022, 2026, Oracle and/or its affiliates.

2.4.2 List Pool Disks 23
2.4.3 Modify the State of a Pool Disk 24
2.4.4 Resize Pool Disks 24
2.4.5 Drop Pool Disks 25
2.5 Administer Exascale Storage Pools 25
2.5.1 Create a Storage Pool 25
2.5.2 List Storage Pools 26
2.5.3 Maintain Free Space in a Storage Pool to Protect Against Disk Failure 26
2.5.4 Modify a Storage Pool 28
2.5.5 Drop a Storage Pool 28
2.6 Administer the Storage Devices in an Exascale Storage Pool 29
2.6.1 Replacing a Failed Storage Device 29
2.6.2 Proactively Replacing a Storage Device 31
2.7 Administer Cluster Templates 33
2.7.1 Create a Cluster Template 33
2.7.2 List Cluster Templates 34
2.7.3 Modify a Cluster Template 34
2.7.4 Remove a Cluster Template 35

3 Oracle Exadata Exascale User-Specific Administration

3.1 Administer Exascale User Credentials 1
3.1.1 Create User Keys 1
3.1.2 Create a Wallet 2
3.1.3 Store a Private Key in a Wallet 2
3.1.4 Fetch the Trust Store 3
3.1.5 Store the URL for Exascale Cluster Services 3
3.1.6 Store the Exascale Control Services Endpoint 4
3.1.7 Distribute a Wallet 4
3.1.8 Change User Keys 5
3.2 Administer Exascale Vaults 6
3.2.1 Create a Vault 6
3.2.2 Provision a Vault 6
3.2.2.1 Provision a Vault Using Multiple Storage Pools For Each Media Type 7
3.2.2.2 Provision a Vault Using Only One Storage Pool For Each Media Type 8
3.2.3 List Vaults 9
3.2.4 Remove a Vault 10
3.3 Administer Exascale Files 10
3.3.1 Create a File 11
3.3.2 Copy a File 11
3.3.3 List Files 12
3.3.4 Cloning Files 12

Copyright © 2022, 2026, Oracle and/or its affiliates.

3.3.5 Snapshotting Files 14
3.3.6 Lock File Snapshots 16
3.3.7 List Clones and Snapshots of Exascale Files 18
3.3.8 Remove a File 20
3.4 Administer Exascale Datasets 20
3.4.1 List Datasets 20
3.5 Administer Access Control Lists 21
3.5.1 List ACLs 21
3.5.2 Modify an ACL 22
3.6 Administer Extended Attributes 24
3.6.1 List Extended Attributes 24
3.6.2 Set an Extended Attribute 25
3.6.3 Remove an Extended Attribute 25
3.7 Administer Vault Templates 26
3.7.1 Create a Vault Template 26
3.7.2 List Vault Templates 27
3.7.3 Modify a Vault Template 28
3.7.4 Remove a Vault Template 29
3.8 Administer Resource Profiles 29
3.8.1 Create a Resource Profile 29
3.8.2 List Resource Profiles 31
3.8.3 Modify a Resource Profile 32
3.8.4 Remove a Resource Profile 34
3.9 Administer Exascale Block Store Data Objects 34
3.9.1 Administer Exascale Volumes 34
3.9.1.1 Create a Volume 34
3.9.1.2 List Volumes 35
3.9.1.3 Modify a Volume 36
3.9.1.4 Remove a Volume 36
3.9.2 Administer Volume Attachments 37
3.9.2.1 Create an EDV Attachment 37
3.9.2.2 List EDV Attachments 39
3.9.2.3 Remove an EDV Attachment 40
3.9.3 Administer Volume Snapshots 40
3.9.3.1 Create a Volume Snapshot 41
3.9.3.2 Lock Volume Snapshots 41
3.9.3.3 List Volume Snapshots 43
3.9.3.4 Remove a Volume Snapshot 43
3.9.4 Administer Volume Clones 44
3.9.4.1 Create a Volume Clone 44
3.9.4.2 List Volume Clones 45
3.9.4.3 Modify a Volume Clone 46

Copyright © 2022, 2026, Oracle and/or its affiliates.

3.9.4.4 Remove a Volume Clone 46
3.9.5 Administer Volume Backups 47
3.9.5.1 Create a Volume Backup 47
3.9.5.2 List Volume Backups 48
3.9.5.3 Restore a Volume Backup 48
3.9.5.4 Monitor the Progress of a Volume Backup or Restore 49
3.9.5.5 Remove a Volume Backup 50
3.9.6 Administer Volume Groups 50
3.9.6.1 Create a Volume Group 50
3.9.6.2 List Volume Groups 51
3.9.6.3 Adding and Removing Volumes in a Volume Group 51
3.9.6.4 Modify a Volume Group 52
3.9.6.5 Create a Volume Group Snapshot 53
3.9.6.6 Remove a Volume Group 53
3.9.7 Administer an Advanced Cluster File System (ACFS) on Exascale 54
3.9.7.1 Create an ACFS File System on Exascale 54
3.9.7.2 List ACFS File Systems on Exascale 55
3.9.7.3 Modify an ACFS File System on Exascale 55
3.9.7.4 Remove an ACFS File System on Exascale 57

4 Using Exascale with Oracle Grid Infrastructure and Oracle Database

4.1 Using Oracle Grid Infrastructure with Exascale 1
4.2 Using Oracle Database with Exascale 2
4.3 About Exascale User Credentials for Oracle Grid Infrastructure and Oracle Database 2
4.4 Using Oracle Managed Files with Exascale 3
4.5 Customizing the File Storage Attributes for Oracle Database Files 5
4.6 Using Exascale Resource Profiles with Oracle Database 5
4.7 Using Oracle Database with Exascale Snapshots and Clones 6
4.7.1 Thin Cloning a Pluggable Database 6
4.7.2 Thin Cloning a Pluggable Database in a Different Container Database 7
4.7.3 Using a Carousel of Thinly Provisioned Pluggable Database Snapshots 8
4.7.4 Cloning a Container Database 9
4.7.5 Capacity Planning for Oracle Database with Exascale Snapshots and Thin Clones 12
4.8 Using Exascale with Oracle Database 19c and Oracle Database 21c 14
4.9 Copying a TDE Keystore Between a File System and Exascale Storage 15

5 Using Exascale Block Storage

5.1 Using Exascale Block Storage with EDV 1

Copyright © 2022, 2026, Oracle and/or its affiliates.

5.2 Using Resource Profiles with Exascale Volumes 5

6 Using ESCLI

6.1 Start and Use ESCLI 1
6.2 ESCLI Command Reference 2
6.2.1 ESCLI Command Help 3
6.2.2 Describing Resources and Attributes 3
6.2.3 Service and Cluster Management 6
6.2.3.1 cellcli 6
6.2.3.2 chcluster 7
6.2.3.3 chclusteralert 8
6.2.3.4 chservice 9
6.2.3.5 dbmcli 11
6.2.3.6 lscell 12
6.2.3.7 lscelldisk 13
6.2.3.8 lscluster 14
6.2.3.9 lsclusteralert 16
6.2.3.10 lscomputeserver 18
6.2.3.11 lsgriddisk 19
6.2.3.12 lsservice 20
6.2.3.13 rmclusteralert 22
6.2.4 Security and User Management 22
6.2.4.1 chacl 23
6.2.4.2 chuser 25
6.2.4.3 chwallet 26
6.2.4.4 lsacl 28
6.2.4.5 lskey 29
6.2.4.6 lsuser 30
6.2.4.7 lswallet 31
6.2.4.8 mkkey 33
6.2.4.9 mkuser 33
6.2.4.10 mkwallet 35
6.2.4.11 rmuser 35
6.2.5 Storage Pool and Pool Disk Management 36
6.2.5.1 chpooldisk 36
6.2.5.2 chstoragepool 37
6.2.5.3 lspooldisk 38
6.2.5.4 lsstoragepool 39
6.2.5.5 lsstoragepooloperation 40
6.2.5.6 mkstoragepool 42
6.2.5.7 rmstoragepool 43

Copyright © 2022, 2026, Oracle and/or its affiliates.

6.2.6 Vault Management 43
6.2.6.1 chvault 44
6.2.6.2 lsvault 46
6.2.6.3 mkvault 48
6.2.6.4 rmvault 50
6.2.7 Dataset Management 50
6.2.7.1 lsdataset 50
6.2.8 File Management 53
6.2.8.1 cd 54
6.2.8.2 chfile 55
6.2.8.3 clonefile 56
6.2.8.4 extentmap 58
6.2.8.5 getfile 60
6.2.8.6 ls 60
6.2.8.7 lsfile 63
6.2.8.8 lssnapshots 66
6.2.8.9 mkfile 69
6.2.8.10 putfile 72
6.2.8.11 rmfile 72
6.2.8.12 snapshotfile 73
6.2.9 Template Management 75
6.2.9.1 chtemplate 75
6.2.9.2 lstemplate 77
6.2.9.3 mktemplate 80
6.2.9.4 rmtemplate 82
6.2.10 Resource Profile Management 83
6.2.10.1 chresourceprofile 84
6.2.10.2 lsresourceprofile 86
6.2.10.3 mkresourceprofile 88
6.2.10.4 rmresourceprofile 90
6.2.11 Extended Attribute Management 91
6.2.11.1 chxattr 91
6.2.11.2 lsxattr 92
6.2.11.3 rmxattr 93
6.2.12 Block Store Management 93
6.2.12.1 acfsctl 95
6.2.12.2 chacfsfilesystem 96
6.2.12.3 chvolume 97
6.2.12.4 chvolumeattachment 98
6.2.12.5 chvolumebackup 98
6.2.12.6 chvolumegroup 99
6.2.12.7 chvolumesnapshot 101

Copyright © 2022, 2026, Oracle and/or its affiliates.

6.2.12.8 lsacfsfilesystem 101
6.2.12.9 lsinitiator 103
6.2.12.10 lsvolume 104
6.2.12.11 lsvolumeattachment 105
6.2.12.12 lsvolumebackup 107
6.2.12.13 lsvolumegroup 109
6.2.12.14 lsvolumesnapshot 110
6.2.12.15 mkacfsfilesystem 112
6.2.12.16 mkvolume 113
6.2.12.17 mkvolumeattachment 115
6.2.12.18 mkvolumebackup 117
6.2.12.19 mkvolumegroup 118
6.2.12.20 mkvolumegroupsnapshot 120
6.2.12.21 mkvolumesnapshot 120
6.2.12.22 rmacfsfilesystem 122
6.2.12.23 rmvolume 122
6.2.12.24 rmvolumeattachment 123
6.2.12.25 rmvolumebackup 123
6.2.12.26 rmvolumegroup 124
6.2.12.27 rmvolumesnapshot 125

7 Using XSH

7.1 Start and Use XSH 1
7.2 XSH Command Reference 3
7.2.1 cat 5
7.2.2 chacl 6
7.2.3 chresourceprofile 8
7.2.4 chtemplate 10
7.2.5 chxattr 11
7.2.6 clone 12
7.2.7 cp 14
7.2.8 dd 17
7.2.9 hexdump 19
7.2.10 ls 20
7.2.11 lsacl 22
7.2.12 lsblk 23
7.2.13 lsdataset 24
7.2.14 lsinitiator 26
7.2.15 lsresourceprofile 27
7.2.16 lstemplate 28
7.2.17 lsxattr 29

Copyright © 2022, 2026, Oracle and/or its affiliates.

7.2.18 man 30
7.2.19 mkresourceprofile 31
7.2.20 mktemplate 33
7.2.21 mv 34
7.2.22 rm 36
7.2.23 rmresourceprofile 37
7.2.24 rmtemplate 38
7.2.25 rmxattr 39
7.2.26 scrub 40
7.2.27 snap 41
7.2.28 strings 43
7.2.29 touch 44
7.2.30 version 46
7.2.31 volumeiops 46

8 Using Exascale Utility Programs

8.1 edvmkvol 1
8.2 edvrmvol 3

A Exascale-Specific Information in the Oracle Data Dictionary

A.1 Oracle Database Dictionary Views A-1
A.1.1 V$EXA_FILE A-1
A.1.2 V$EXA_TEMPLATE A-2
A.1.3 V$EXA_VAULT A-3

B Exascale Resources and Attributes

Copyright © 2022, 2026, Oracle and/or its affiliates.

Preface

This guide describes how to use and administer Oracle Exadata Exascale.

This Preface contains the following topics:

- Audience
- Related Documents

Audience

The audience for this book includes system administrators, database administrators, and storage administrators. This book is intended for database and storage administrators who configure, administer, manage, and consume storage services using Oracle Exadata Exascale.

To use this document, you should be familiar with basic Oracle Exadata and Oracle Database concepts and administrative procedures. Also, you might want to review documentation about related products, such as Oracle Advanced Cluster File System (ACFS), Oracle Grid Infrastructure, and Oracle Real Application Clusters (Oracle RAC). For a list of related books, refer to Related Documents.

Related Documents

For more information, refer to the following Oracle resources:

- *Oracle Exadata System Overview*
- *Oracle Exadata Installation and Configuration Guide*
- *Oracle Exadata System Software User's Guide*
- *Oracle Exadata Maintenance Guide*
- *Oracle AI Database Administrator’s Guide*
- *Oracle AI Database Concepts*
- *Oracle AI Database Reference*
- *Oracle AI Database SQL Language Reference*
- *Oracle Clusterware Administration and Deployment Guide*
- *Oracle Real Application Clusters Administration and Deployment Guide*

Copyright © 2022, 2026, Oracle and/or its affiliates.

1

Oracle Exadata Exascale Overview

- What is Oracle Exadata Exascale?
- Exascale Components and Concepts

1.1 What is Oracle Exadata Exascale?

Oracle Exadata Exascale further empowers Exadata to meet the most demanding corporate and cloud computing requirements by decoupling Oracle Database and Oracle Grid Infrastructure clusters from the underlying Exadata storage servers. With Exascale, you can centrally manage a large fleet of Exadata storage servers connected by the Exadata RDMA Network Fabric. This architecture allows multiple Oracle Grid Infrastructure clusters and databases to securely access a large pool of shared storage resources, providing:

- Secure storage sharing with strict data isolation, ensuring users and databases can access only the data for which they have the proper privileges
- Flexible, dynamic storage provisioning for many users and databases
- Increased storage utilization and efficiency, maximizing your return on investment
- Improved utilization of storage processing resources to boost overall performance

Furthermore, Exascale introduces advanced snapshot and cloning capabilities that are tightly integrated with Oracle Database. For example, Oracle Database provides native snapshot and cloning functionality for pluggable databases (PDBs) through the `CREATE PLUGGABLE DATABASE` and `ALTER PLUGGABLE DATABASE` SQL commands. When Oracle Database utilizes Exascale storage, the pluggable database snapshot and snapshot copy (cloning) functions automatically use native Exascale snapshots and clones, which are space-efficient file copies based directly on the underlying Oracle Database files, thereby eliminating the need for a test master database to support snapshots and clones on Exadata.

In addition to unparalleled support for Oracle Database, Exascale provides block storage services, which deliver sophisticated capabilities to create and manage arbitrary-sized raw block volumes based on Exascale storage.

While end users can create and use Exascale block volumes for numerous applications, Exadata also leverages Exascale block volumes internally to store Exadata database server virtual machine images. Placing virtual machine images in Exascale removes the dependency on local storage inside the Exadata compute nodes, which enables the creation of more virtual machines and provides the infrastructure to support seamless migration of virtual machines between different Exadata compute nodes.

Despite the fact that Exascale transforms Exadata storage, Exascale also preserves the proven strengths and benefits of Exadata:

- Scalability - including efficient support for hundreds of Exadata storage servers in an Exascale cluster
- High availability - based on a clustered architecture with built-in redundancy and dynamic fail-over of software services
- High performance - utilizing Exadata RDMA memory and Exadata Smart Flash Cache

Copyright © 2022, 2026, Oracle and/or its affiliates.

- Reliability - using proven Exadata storage server technologies
- Security - employing advanced security protocols and automatic encryption

Exascale runs on all Oracle Exadata system hardware with RoCE Network Fabric (X8M-2 or later), except systems with 8-socket database servers (X8M-8 and X9M-8). For full-featured native Oracle Database file storage in Exascale, you must use Oracle Database 23ai release 23.5.0 or later. You can also employ Exascale block volumes to support databases using older Oracle Database software releases back to Oracle Database 19c.

1.2 Exascale Components and Concepts

This section describes key Exascale components and concepts:

- Exascale Cluster
- Exascale Services
- Exascale Resources and Attributes
- Storage Media Types
- Pool Disks
- Storage Pools
- Vaults
- Files
- File Storage Attributes
- Templates
- Extended Attributes
- Clone Files
- Snapshot Files
- Datasets
- Exascale Users
- Exascale User Credentials
- Cluster Level Privileges
- Access Control Lists
- Vault and File Access Control
- Trust Store
- Resource Management
- Exascale Block Store

1.2.1 Exascale Cluster

An Exascale cluster consists of Exadata storage servers operating together as a unified system, delivering storage to multiple Oracle Grid Infrastructure clusters and databases.

Each Exascale deployment comprises exactly one Exascale cluster, and each Exascale cluster operates independently of all others. An Exadata storage server may belong to only one Exascale cluster.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Within the Exascale cluster, a set of software services manages the cluster and associated system resources, such as storage pools, vaults, volumes, and other resources.

Related Topics

- Administer the Exascale Cluster

1.2.2 Exascale Services

Exascale system components are implemented using a set of software services in a clustered architecture. The core Exascale storage services predominantly run on the Exadata storage servers.

Exascale contains the following storage services:

- **Cluster Services**

  Exascale cluster services, also known as Exascale Global Services (EGS), provide the core foundation for the Exascale system. EGS primarily manages the storage allocated to Exascale storage pools. It also manages storage cluster membership, provides security and identity services for storage servers and Exascale clients, and monitors the other Exascale services. Exascale cluster services use the Raft consensus algorithm.

  For high availability, every Exascale cluster contains five EGS instances.

  For Exascale clusters with five or more Exadata storage servers, one EGS instance runs on each of five storage servers.

  For Exascale configurations with fewer than five storage servers, one EGS instance runs on each Exadata storage server, and the remaining EGS instances run on the Exadata compute nodes to make up the required total of five. In a bare-metal configuration, EGS compute node instances run on the server operating system. In a configuration with compute nodes running on virtual machines (VMs), EGS compute node instances run in the hypervisor.

- **Control Services**

  Exascale control services, also known as Exascale RESTful Services (ERS), provide a management endpoint for Exascale management operations. All Exascale management operations come through ERS. But, no file I/O operations come through ERS.

  ERS service instances are deployed using front-end and back-end server processes. The front-end ERS processes provide a highly available client endpoint with load-balancing capabilities. The back-end ERS processes work with other software services to process requests and reply back to the client.

  Multiple ERS instances provide high availability and share the Exascale management workload. Typically, five ERS instances are distributed across the Exadata storage servers. However, for configurations with fewer than five storage servers, one ERS instance usually runs on each Exadata storage server.

  The Exascale Command Line (ESCLI) utility provides a simple command-line interface to perform Exascale monitoring and management functions. ESCLI commands are translated into ERS calls and run through ERS. ESCLI works in conjunction with the Exadata cell command-line interface (CellCLI) and does not replace it.

- **Exascale Vault Manager Services**

  Exascale vault manager, also known as Exascale Data Services (EDS), is the collective name for the software services that manage metadata for Exascale files and vaults:

  – **System Vault Manager**

Copyright © 2022, 2026, Oracle and/or its affiliates.

The system vault manager service (SYSEDS) serves and manages the metadata for Exascale vaults. This metadata includes vault-level access control lists (ACLs) and attributes.

SYSEDS is a lightweight process, so one instance can usually service the load for an entire Exascale cluster. However, to ensure high availability, five SYSEDS instances are typically distributed across the Exadata storage servers. For configurations with fewer than five storage servers, one SYSEDS instance usually runs on each Exadata storage server.

- **User Vault Manager**

  The user vault manager service (USREDS) serves and manages the metadata for files inside the Exascale vaults. This metadata includes file-level access control lists (ACLs) and attributes, along with metadata that defines clones and snapshots. All file control operations, such as open and close, are serviced by the user vault manager service.

  Multiple USREDS instances provide high availability and share the user workload. Typically, five USREDS instances are distributed across the Exadata storage servers. However, for configurations with fewer than five storage servers, one USREDS instance usually runs on each Exadata storage server.

- **Block Store Manager**

  The Block Store Manager service (BSM) serves and manages the metadata for Exascale block storage. All block store management operations are serviced by BSM. These block store management operations include creating a volume, creating a volume snapshot, and so on. It also coordinates the block store worker processes.

  To provide high availability, three BSM instances are typically distributed across the Exadata storage servers.

- **Block Store Worker**

  The Block Store Worker service (BSW) primarily services requests from block store clients and performs the resulting storage server I/O. It plays a role in clone and snapshot creation operations and is also responsible for performing volume backup and restore operations.

  To provide high availability and share the user workload, three BSW instances are typically distributed across the Exadata storage servers.

- **Instant Failure Detection**

  The Instant Failure Detection (IFD) service ensures high availability by detecting failures across the Exascale cluster as they occur, enabling rapid recovery of critical services and minimizing the impact of failures. IFD automatically runs as a lightweight, dedicated service on every storage server in the Exascale cluster, as well as on any Exadata compute node that runs an Exascale Global Services (EGS) instance or the Block Store Worker (BSW) service.

- **Exadata Cell Services**

  Exascale works in conjunction with, and relies on, the core Exadata cell services. Specifically, Exascale requires running instances of Cell Server (CELLSRV), Management Server (MS), and Restart Server (RS) on every storage server. Also, Exascale requires running instances of Management Server (MS) and Restart Server (RS) on every compute server that is associated with the Exascale cluster.

In addition to the Exascale storage services, the following client-side services provide specific support for various Exascale functions on the Exadata compute nodes:

- **Exascale Node Proxy**

Copyright © 2022, 2026, Oracle and/or its affiliates.

The Exascale Node Proxy (ESNP) service maintains information about the current state of the Exascale cluster, which it provides to local Oracle Grid Infrastructure and Oracle Database processes.

ESNP is a background process that runs on each Exadata compute node. In a bare-metal configuration, ESNP runs on the server operating system. In a configuration with compute nodes running on virtual machines (VMs), ESNP runs inside the guest VMs.

- **Exascale Direct Volume**

  Exascale Direct Volume (EDV) is the default and recommended volume attachment mechanism within the Exadata RDMA Network Fabric.

  The EDV service exposes Exascale volumes as EDV devices on Exadata compute nodes and services the I/O on each EDV device. EDV-managed storage can be used as raw block devices or to support various file systems, including Oracle Advanced Cluster File System (ACFS).

  The EDV service is required on all Exadata compute nodes where you want to use EDV devices. In a bare-metal configuration, the EDV service runs on the server operating system. In a configuration with compute nodes running on virtual machines (VMs), the EDV service runs in the KVM host to provide the hypervisor access to VM image files based on EDV-managed storage. The EDV service also runs in each guest VM to enable direct access to EDV-managed storage from inside the VM.

Related Topics

- Administer the Exascale Cluster

1.2.3 Exascale Resources and Attributes

The metadata describing Exascale system objects is organized into a collection of resources. Exascale resources define the primary objects that Exascale users interact with, such as files, vaults, storage pools, volumes, and others.

Each resource contains a specific set of attributes. For example, every file has numerous attributes, including the file name, file size, and the name of the vault containing the file. Many attributes are automatically populated by the system, such as the file creation time. Some attributes, such as the size of a file, may be set during resource creation, while others, like the access control list (ACL), can be modified after creation.

At any point, users can utilize the Exascale Command Line (ESCLI) utility to view information about any Exascale resource and its attributes. See also Describing Resources and Attributes.

1.2.4 Storage Media Types

Exascale storage is categorized based on its underlying media type. The supported media types are:

- `HC`: Identifies high capacity storage, using hard disk drives (HDDs) on high-performance Exadata storage servers.
- `EF`: Identifies extreme flash storage, using low-latency, high-throughput flash devices.
- `XT`: Identifies extended storage, using HDDs on Exadata storage servers optimized for cost-effective online storage of less frequently accessed data.

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> The minimum requirements for using XT storage in Exascale are:
>
> - Exadata Storage Server X11M Extended (XT) hardware
> - Oracle Exadata System Software release 25.2

1.2.5 Pool Disks

Exascale uses pool disks and storage pools to organize the physical storage provided by Exadata storage servers.

A pool disk is an Exascale-specific Exadata grid disk that reserves storage space for Exascale from an Exadata storage device.

Like any other Exadata grid disk, a pool disk is created by allocating space from an Exadata cell disk.

Based on its underlying physical characteristics, each pool disk is implicitly associated with a storage media type.

Related Topics

- Administer Exascale Pool Disks

1.2.6 Storage Pools

Exascale uses pool disks and storage pools to organize the physical storage provided by Exadata storage servers.

An Exascale storage pool is a collection of pool disks that provides persistent physical storage for Exascale vaults and files.

A working Exascale cluster requires at least one storage pool. Note that there is no need to define multiple storage pools for data separation because Exascale vaults provide strong data isolation.

Each storage pool is a collection of pool disks with a common media type (for example, HC). You cannot define a storage pool with pool disks having a mixture of media types (for example, HC and XT).

Starting with Oracle Exadata System Software release 26.1.0, you can create multiple storage pools for each media type within a single Exascale cluster. Previously, only one storage pool was permitted for each media type in an Exascale cluster.

A storage pool can contain pool disks that reside on different generations of Exadata storage server hardware, which allows for easy migration to new generations of Exadata hardware.

You can dynamically reconfigure a storage pool by changing the size of the pool disks or by adding or removing pool disks (or Exadata storage servers). However, pay attention to the following:

- For each storage pool, use pool disks that are spread across all of the available storage devices in each storage server.
- On each storage server, maintain consistent sizing for all of the pool disks belonging to a storage pool. However, pool disks on different storage servers can have different sizes.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Because a storage pool may contain a vast amount of data, Exascale automatically organizes it using internal constructs to optimize performance, availability, and manageability.

Within each storage pool, Exascale implicitly groups pool disks and their associated storage servers into disk partner groups and cell partner groups. This organization maximizes availability by limiting the impact of multiple simultaneous failures. Separate groups enable the system to handle simultaneous failures independently, and failures that affect only a small group are statistically less likely than failures across the entire system.

Each storage pool also contains internal failure constructs called storage pool rings. Storage pool rings ensure that data with different content types and redundancy requirements use separate disk partner groups. Through the use of separate storage pool rings and native integration with Oracle Database, Exascale automatically isolates Oracle Database data files from recovery files. This organization guarantees that if multiple failures affect a data file, its corresponding recovery files remain available.

You cannot control the organization of storage pool rings or partner groups. However, Exascale provides information about these constructs to support system monitoring and to help users understand the impact of any storage failure scenario.

Related Topics

- Administer Exascale Storage Pools

1.2.7 Vaults

An Exascale vault is a logical storage container that uses the physical resources provided by Exascale storage pools.

By default, a vault can use all the underlying storage pool resources. However, an Exascale administrator can limit the amount of space, I/O resources (I/Os per second, or IOPS), and cache resources associated with each vault.

Exascale vaults facilitate strict data separation, ensuring that data is isolated to specific users and separated from other data and users. A vault, and its contents, are invisible to users without the appropriate privileges. Without the correct entitlements, users of one vault cannot see another vault, even though data from both vaults may be striped across the same underlying storage pools.

To an end-user and Oracle Database, a vault appears like a top-level directory that contains files. Exascale uses the convention of beginning vault names with the at sign (@) character. For example, `@my-vault`. So, a fully qualified Exascale file path always starts with the at sign (@) and vault name, followed by the rest of the file path. For example, `@my-vault/my-example/my-path/my-file-name`.

Note that although a vault seems to use a hierarchical directory structure, it does not actually contain directories. Instead, each file has a path that can include virtual directories. For example, in the file path `@my-vault/my-db1/datafiles/my-datafile1`, `my-db1` and `datafiles` function as virtual directories within the full file path. These virtual directories implicitly exist only as long as files reference them, and they are removed automatically when no files use them. As a result, there are no commands to create and delete virtual directories within a vault, and there are no permissions or access controls assigned to them.

Related Topics

- Administer Exascale Vaults

Copyright © 2022, 2026, Oracle and/or its affiliates.

1.2.8 Files

Exascale is optimized to directly store and manage files associated with Oracle Database and Oracle Grid Infrastructure.

When Oracle Database or Oracle Grid Infrastructure creates a file inside a vault, Exascale automatically recognizes the file type and applies the appropriate file storage attributes using a template associated with that file type. Among other things, the template determines the storage pool where the file is stored.

You can also use the Exascale Command Line (ESCLI) mkfile command to create a file. By this means, you can explicitly set the file storage attributes, or you can influence the file storage attributes by specifying the file type or by using a specific template. You can then use the ESCLI putfile command to copy the contents of a file into Exascale.

To optimize storage efficiency, physical space is consumed only when data is written to the file. For example, when you use the ESCLI mkfile command to create a file, the specified file size is reserved in the vault, but the file does not use any physical storage at that point. Physical storage is only allocated when data is written to the file.

Technically, you can store any type of files inside an Exascale vault. However, Exascale is optimized to manage Oracle Database files, which are typically much larger than regular files. To optimally store and manage regular files, you can define a block volume on Exascale, attach the volume to a server, and then use the volume attachment to support a file system. See Exascale Block Store.

Related Topics

- Administer Exascale Files

1.2.9 File Storage Attributes

During file creation, every file in Exascale is associated with core file storage attributes, which govern how the file is stored and managed. The attributes are:

- `mediaType`: Specifies the physical media type used to store the file. Possible values are:
  - `HC`: Identifies high capacity storage, using hard disk drives (HDDs) on high-performance Exadata storage servers.
  - `EF`: Identifies extreme flash storage, using low-latency, high-throughput flash devices.
  - `XT`: Identifies extended storage, using HDDs on Exadata storage servers optimized for cost-effective online storage of less frequently accessed data.
- `storagePool`: Specifies the name of the storage pool that contains the file.
- `redundancy`: Specifies the number of data copies that are maintained. Currently, the only permitted value is:
  - `high`: Indicates three mirrored copies of the file data.
- `contentType`: Specifies the type of content in the file. Exascale internally uses this attribute to place file extents on physically separate devices in a manner that maximizes availability in the event of a failure. Possible values are:
  - `DATA`: Principally associated with user data.
  - `RECO`: Primarily for data used in backup and recovery operations.

Copyright © 2022, 2026, Oracle and/or its affiliates.

File storage attributes are assigned to all files when they are created. Typically, the file storage attributes are implicitly assigned using templates, which assign attribute settings based on the file type. File storage attributes can be assigned explicitly when a file is created by using the ESCLI mkfile command.

You cannot change the file storage attributes after the file is created.

1.2.10 Templates

During file creation, every file in Exascale is associated with attributes that govern how the file is stored and managed.

An Exascale template is a named collection of file storage attribute settings. For example:

```
name: DATAFILE
mediaType: HC
redundancy: high
contentType: DATA
```

With templates, you can set file storage attributes automatically and consistently. When creating a file, Exascale automatically uses the template that matches the file type by default. For example, when Oracle Database creates a data file, the `DATAFILE` template is automatically used by default.

Note that templates only govern the assignment of file storage attributes during file creation. You cannot change the file storage attributes after the file is created. Any change to a template applies only to files created after the change.

Within Exascale, templates are defined at two levels:

- **Cluster Templates** are defined at the cluster level.

  A cluster template contains settings for the file storage attributes: `mediaType`, `redundancy`, and `contentType`.

  A cluster template is used if no vault-level template exists to override it.

- **Vault Templates** are defined at the vault level.

  A vault template may also contain settings for the file storage attributes: `mediaType`, `redundancy`, and `contentType`.

  Additionally, starting with Oracle Exadata System Software release 26.1.0, a vault template can also include a `storagePool` setting to specify the name of a storage pool used to store files created with the template. If a template specifies both `mediaType` and `storagePool`, the values must be compatible. For example, you cannot specify the use of high capacity (HC) storage media (`mediaType=HC`) and also specify the name of a storage pool that uses extreme flash (EF) storage.

  A vault template defines vault-specific attributes that override the corresponding cluster template. For example, your cluster-level `DATAFILE` template may specify using HC storage media (`mediaType=HC`), but you may override that setting for a vault by having a vault-level `DATAFILE` template that specifies EF storage (`mediaType=EF`). Or, starting with Oracle Exadata System Software release 26.1.0, you could define a vault-level `DATAFILE` template that uses a specific storage pool rather than nominating a media type.

Additionally, **User Templates** are templates with user-specified names that do not correspond to a specific file type. User templates effectively override file-type templates. However, to use a user template, you must explicitly specify it during file creation. User templates can be defined

Copyright © 2022, 2026, Oracle and/or its affiliates.

at the cluster level or at the vault level. A vault-level user template will override a cluster-level user template having the same name.

Starting with Oracle Exadata System Software release 26.1.0, an Exascale cluster can support multiple storage pools for each media type. If a template does not specify a `storagePool` setting and multiple storage pools match the template's `mediaType` setting, Exascale automatically selects the storage pool to use.

Related Topics

- Administer Cluster Templates
- Administer Vault Templates

1.2.11 Extended Attributes

An extended attribute is non-standard metadata that is associated with an Exascale vault or file. An extended attribute consists of an extended attribute name and an associated value. The extended attribute name is a user-defined string value. The extended attribute value can be a string value or a binary value read from a file.

In addition to user-defined extended attributes, extended attributes reserved for internal use by Exascale services and tools have names starting with a dollar sign ($). You cannot create, modify, or delete these extended attributes.

Multiple extended attributes are allowed for each Exascale vault or file. For each file or vault, the total space available for all extended attributes is 16 kilobytes. This allowance includes the extended attribute names and values, along with a small allowance for internal metadata.

Related Topics

- Administer Extended Attributes

1.2.12 Clone Files

An Exascale clone is a thinly-provisioned point-in-time copy of a file that is readable and writable. The source file for a clone can be a regular file, a snapshot, or another clone. Exascale uses redirect-on-write techniques to create and maintain clones very quickly and space-efficiently.

The clone creation process can work on an individual file or a group of files. All clones created in the same operation are point-in-time consistent, and all files in a clone operation must be in the same vault.

A clone file inherits the ACL and file storage attribute settings from the original file. After creation, you can modify the clone ACL and any modifiable attributes independently from the original file. You can also modify, or even delete, the original file without affecting the clone.

Typically, you will automatically use Exascale clones through the Oracle Database pluggable database (PDB) snapshot and cloning capabilities. For example, the Oracle Database `CREATE PLUGGABLE DATABASE ... FROM ...` command internally uses Exascale clones when the `FROM` database is on Exascale.

You can also use the ESCLI clonefile command to manually create clones, and you can use lssnapshots to view the association between clones and their sources.

Exascale also contains separate cloning capabilities for Exascale block volumes. See Volume Clones.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Related Topics

- Administer Exascale Files
- Cloning Files

1.2.13 Snapshot Files

An Exascale snapshot is a thinly-provisioned read-only point-in-time copy of a file. The source file for a snapshot can be a regular file, a clone, or another snapshot. Exascale uses redirect-on-write techniques to create and maintain snapshots very quickly and space-efficiently.

The snapshot creation process can work on an individual file or a group of files. All snapshots created in the same operation are point-in-time consistent, and all files in a snapshot operation must be in the same vault.

Additionally, starting with Oracle Exadata System Software release 26.1.0, you can lock a snapshot file to prevent its deletion before a specified future date and time. You can also make the lock immutable, which prevents users from removing the lock or changing it to expire earlier.

A snapshot file inherits the ACL and file storage attribute settings from the original file. After creation, you can modify the snapshot ACL and any modifiable attributes independently from the original file. You can also modify, or even delete, the original file without affecting the snapshot.

You can use the ESCLI snapshotfile command to create snapshots, and you can use lssnapshots to view the association between snapshots and their sources.

While an Exascale snapshot is read-only, you can effectively manufacture a writable snapshot by creating a snapshot and then making a clone of the snapshot. In this case, the snapshot provides a durable point-in-time copy of the file, and the clone enables writing to a thinly-provisioned copy of the snapshot. At any point, you can effectively roll back to the snapshot creation time by removing and re-creating the clone.

Exascale also contains separate snapshotting capabilities for Exascale block volumes. See Volume Snapshots.

Related Topics

- Administer Exascale Files
- Snapshotting Files
- Lock File Snapshots

1.2.14 Datasets

An Exascale dataset is a logical grouping of files within a vault. The primary purpose of a dataset is to enable tracking and management of storage utilization for files associated with Oracle Database and Oracle Grid Infrastructure.

Exascale automatically creates and maintains system-defined datasets, which are conceptually organized in a hierarchical tree. Each dataset directly contains some number of files and indirectly contains child datasets that correspond with other entities in the hierarchy. The hierarchy of system-defined datasets contains the following levels:

1. Exascale vault
2. Oracle Grid Infrastructure (GI) cluster
3. Oracle multitenant container database (CDB)

Copyright © 2022, 2026, Oracle and/or its affiliates.

4. Oracle pluggable database (PDB)

Starting at the bottom, the files belonging to each PDB are contained in a separate dataset. The parent CDB dataset contains the files belonging to the CDB and all of the associated PDB datasets. All of the CDBs in a GI cluster are grouped in a GI-level dataset, which also contains GI-specific files such as the Oracle Cluster Registry (OCR) and voting files. At the top of the hierarchy, the vault-level dataset contains all of the GI-level datasets that consume storage in the vault. The vault-level dataset is also the default container for any files that don't belong to Oracle Database and Oracle Grid Infrastructure.

Each system-defined dataset is identified by a composite ID, which contains unique identifiers for the associated entities in the hierarchy. For example, each vault-level dataset is simply identified by the vault name (for example, `@my-vault`). But, a PDB dataset ID has the following format:

`@Vault-name:GI-cluster-ID:CDB-ID:PDB-ID`

Each system-defined dataset also has a name. The dataset name is also a composite of the relevant entities, but it contains human-readable names for the GI cluster, CDB and PDB components (instead of the system-generated unique identifiers). For example, a PDB dataset name has the following format:

`@Vault-name/GI-cluster-name:CDB-name.PDB-name`

Related Topics

- Administer Exascale Datasets

1.2.15 Exascale Users

Exascale has a system of user accounts enabling different users to perform actions and access data according to their assigned privileges. Though you can create a single Exascale user with privileges to do everything, a typical configuration contains cluster administration users and storage users:

- Cluster administration users are typically provisioned with privileges to administer the Exascale cluster. Cluster administration users typically administer the physical storage objects; namely storage servers, storage pools and pool disks. They also administer Exascale software services, vaults, cluster templates and user accounts. See Oracle Exadata Exascale System Administration.

  By default, each Exascale cluster contains one superuser account. The user identifier (ID) for the superuser account is `admin`. The `admin` user can implicitly perform any system operation and effectively holds all system privileges.

  While you can use the `admin` user to perform cluster administration tasks, Oracle recommends that you create your own cluster administration users with specific privileges. For example, rather than having one cluster administrator that does everything, you may choose to have dedicated user accounts for security administration, storage administration, and so on.

- Storage users are typically provisioned with privileges to use storage within Exascale vaults. Storage users often administer the vaults they use, and sometimes even create new vaults. Storage users typically manage their own files and the access control lists (ACLs) that govern file access. They also administer vault-level templates, extended file attributes, and their own user credentials. See Oracle Exadata Exascale User-Specific Administration.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Additionally, Exascale contains one node administration account for every node (storage server or compute node) that runs Exascale software services. Each node administration account inherits its user ID from the server hostname and each account contains the privileges required to run the Exascale software services on the node. Do not directly use or modify these accounts.

Related Topics

- Administer Exascale Users

1.2.16 Exascale User Credentials

Exascale user authentication uses public and private key pairs, with user credentials stored in a digital key store.

Each Exascale user is associated with a public key. To prove their identity and connect to Exascale, a user must supply the matching private key.

Exascale user credentials are contained in a digital key store, also known as a wallet. To use Exascale, a user's wallet must contain their private key. It must also contain a copy of the Exascale trust store. As a matter of convenience, a wallet can also store the default endpoint for Exascale control services.

To facilitate flexible key management, each Exascale user can be associated with up to three public and private key pairs. However, each wallet should contain only one Exascale user name and one private key.

For maximum security, an Exascale user should create their own public and private key pairs and manage their own wallet. This is recommended to ensure the integrity of the private keys, since the user should never share a private key, not even with the Exascale administrator.

Related Topics

- Administer Exascale User Credentials

1.2.17 Cluster Level Privileges

Cluster level privileges are high-level privileges that control various actions performed by Exascale users.

Cluster level privileges are assigned by using the ESCLI mkuser or chuser commands.

There are four types of Exascale cluster level privileges, and any user may hold privileges across multiple privilege types. The following list describes the privilege types and the available privileges:

- **Cluster Level Storage Privileges** are powerful privileges that are typically reserved for users that administer the Exascale cluster. A user can hold zero or one of the following cluster level storage privileges:
  - `cl_monitor`: Enables the receiving user to monitor the Exascale cluster by performing list operations using the Exascale command line interface (ESCLI).
  - `cl_operator`: Enables the receiving user to:
    - Monitor the Exascale cluster by performing list operations using ESCLI.
    - Manage pool disks (create, drop, online, offline).
    - Manage software services (list, startup, shutdown, restart, delete).
    - Manage the trust store.

Copyright © 2022, 2026, Oracle and/or its affiliates.

  - `cl_admin`: A set of system administrator privileges that includes all the `cl_monitor` and `cl_operator` privileges, along with all of the privileges from the other privilege types; namely:
    * All of the cluster level vault privileges specified in `vlt_manage`.
    * All the cluster level user privileges: `rest_vault_client`, `rest_volume_client`, `user_create`, `system_restore`, and `on_behalf_of`.
    * All the service privileges: `cellsrv`, `egs`, `ers`, `syseds`, `usreds`, `bsm`, and `bsw`.

    This privilege also enables the receiving user to:
    * Grant any privilege to any user.
    * Reset a key for any user.
    * Create, modify, and delete storage pools.
    * View extent map information.
    * Create, modify, and delete Exascale volumes and related resources, including volume attachments, volume groups, volume backups, volume snapshots, and Exascale-managed ACFS file systems.

- **Cluster Level Vault Privileges** are powerful data access privileges governing the actions that the receiving user is allowed to perform on all vaults and files. Typically, cluster level vault privileges are assigned to administration users that manage files in Exascale vaults.

  Cluster level vault privileges work in addition to access control lists (ACLs). To perform an action on a vault or file, a user requires the appropriate cluster level vault privilege or the appropriate ACL privilege. See Vault and File Access Control.

  A user can hold zero or one of the following cluster level vault privileges:

  - `vlt_inspect`: Enables the receiving user to create new vaults. The receiving user also gets complete control over files created in those vaults.

    This privilege is assigned to new users by default. However, wherever possible, you should consider removing this privilege and using specific ACL privileges instead.
  - `vlt_read`: Includes the `vlt_inspect` privileges and also enables the receiving user to list all existing vaults, display attributes for any vault, create files in any vault, list files in any vault, and display attributes for any file.
  - `vlt_use`: Includes the `vlt_read` privileges and also enables the receiving user to open any file for reading.
  - `vlt_manage`: Includes the `vlt_use` privileges and also enables the receiving user to open any file for read and write, alter any vault or file, and drop vaults and files.

- **Cluster Level User Privileges** govern the actions that the receiving user is allowed to perform on the Exascale cluster. A user can hold zero or more of the following cluster level privileges:

  - `rest_vault_client`: Enables the receiving user to run ESCLI commands that operate on vaults, files (including file clones and snapshots), and other associated resources. These resources include templates, extended attributes, datasets, access control lists (ACLs), and resource profiles.

    This privilege enables primary access to ESCLI commands that send RESTful requests to Exascale control services (ERS) to operate on vault and file-related resources. However, specific operations or visibility into resources may require additional privileges. For example, while a user with this privilege can run the ESCLI `ls` command, other cluster level privileges and ACLs ultimately govern the specific vaults and files that are displayed.

Copyright © 2022, 2026, Oracle and/or its affiliates.

This privilege is first introduced in Oracle Exadata System Software release 25.2.3 (October 2025).

  - `rest_volume_client`: Enables the receiving user to run ESCLI commands that operate on Exascale volumes and related resources, including volume attachments, volume groups, volume backups, volume snapshots, and Exascale-managed ACFS file systems.

    This privilege enables primary access to ESCLI commands that send RESTful requests to Exascale control services (ERS) to operate on volume-related resources. However, specific operations or visibility into resources may require additional privileges. For example, while a user with this privilege can run the ESCLI `lsvolume` command, volume ownership and other cluster level privileges ultimately govern the specific volumes that are displayed.

    This privilege is first introduced in Oracle Exadata System Software release 25.2.3 (October 2025).

  - `user_create`: Enables the receiving user to create new users in the cluster.

    You may choose to assign this privilege to dedicated personnel performing user administration.

  - `system_restore`: Enables the receiving user to restore an Exascale system.

    Do not assign this privilege to any of your Exascale users. This privilege and the accompanying operation are internal and are not used under normal circumstances.

  - `on_behalf_of`: A special privilege that enables the receiving user to send a request to Exascale control services (ERS) on behalf of another user.

    Do not assign this privilege to any of your Exascale users. Typically, this privilege is only assigned to the internal administration accounts that reside on each Exascale node.

- **Service Privileges** govern the Exascale software services that the receiving user is allowed to run. Typically, service privileges are only assigned to the internal node-specific administration accounts that reside on each Exascale node. A user can hold zero or more of the following service privileges:
  - `cellsrv`: Enables the receiving user to run the core Exadata cell services.
  - `egs`: Enables the receiving user to run Exascale cluster services (also known as Exascale Global Services).
  - `ers`: Enables the receiving user to run Exascale control services (also known as Exascale RESTful Services).
  - `syseds`: Enables the receiving user to run the system vault manager service.
  - `usreds`: Enables the receiving user to run the user vault manager service.
  - `bsm`: Enables the receiving user to run the block storage manager service.
  - `bsw`: Enables the receiving user to run the block storage worker service.
  - `ms`: Facilitates the transfer of telemetry information between Exascale RESTful Services (ERS) and the Exadata Management Server (MS).

Additionally, `no_privilege` is a special privilege that removes all privileges from the receiving user. When it is assigned to a user, `no_privilege` cannot be combined with any other privilege.

Related Topics

- Modify User Privileges

Copyright © 2022, 2026, Oracle and/or its affiliates.

1.2.18 Access Control Lists

Access control lists (ACLs) govern the operations that users can perform on Exascale vaults and files.

Each Exascale vault or file has an ACL. A vault ACL enables users to perform actions on the vault and on the files that it contains. A file ACL only controls the file that it is associated with.

The following table lists the ACL privileges and the actions that they enable users to perform:

ACL Privilege	In a vault ACL, the ACL privilege enables the user to:	In a file ACL, the ACL privilege enables the user to:
`inspect`	• Create a file in the vault. • View attributes of the vault, but not the vault contents.	• View attributes of the file, but not the file contents.
`read`	• View attributes of all files in the vault, but not their contents. • Perform all `inspect` actions.	• Read the file contents. • Perform all `inspect` actions.
`use`	• Read the contents of all files in the vault. • Perform all `inspect` and `read` actions.	• Read and write the file contents. • Alter attributes of the file. • Perform all `inspect` and `read` actions.
`manage`	• Read and write the contents of any file in the vault. • Alter the attributes and ACL for the vault and any file in the vault. • Drop the vault and any file in the vault. • Perform all `inspect`, `read` and `use` actions.	• Alter the file ACL. • Drop the file. • Perform all `inspect`, `read` and `use` actions.

Note that the same ACL privilege enables different actions in a vault ACL or a file ACL. For example, in a file ACL the `read` privilege enables the user to read the contents of the file. However, to read file contents using a vault ACL requires the `use` privilege.

Every ACL is a list of user IDs and privilege pairs. Depending on the user creation method, the user ID may be a system-generated value or a user-specified value. For example:

```
96a68014-5762-4579-86ee-29eb743decbd:manage;scott:use;sue:inspect;dd7c8e35-3c8
d-4441-a9b0-f58e959b84ba:read
```

A user is added to an ACL when they are assigned one of the ACL privileges. A user is removed from an ACL when they are assigned the `none` privilege. It is possible for a vault or file to have an empty list of user and privilege pairs, which is also known as a null ACL.

ACLs work in conjunction with cluster level privileges. To perform an action on a vault or file, a user requires appropriate ACL privilege or the appropriate cluster level vault privilege (`vlt_inspect`, `vlt_read`, `vlt_use`, or `vlt_manage`).

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> In addition to the appropriate cluster level vault privilege or ACL privilege, starting with Oracle Exadata System Software release 25.2.3 (October 2025):
>
> - To perform operations that create, modify, or delete Exascale vaults, files, and related resources, a user must also have the `rest_vault_client` privilege or the `cl_admin` privilege.
> - To list vault and file-related resources, a user must also have the `rest_vault_client`, `cl_monitor`, or `cl_operator` privilege.

See Vault and File Access Control.

Related Topics

- Administer Access Control Lists

1.2.19 Vault and File Access Control

Access control lists (ACLs) work together with cluster level privileges, in particular cluster level vault privileges, to control access to Exascale vaults and files. To perform an action on a vault or file, a user requires the appropriate ACL privilege or the appropriate cluster level vault privilege. Because Exascale has no formal concept of vault or file ownership, all operations are governed by the combination of cluster level privileges and ACLs.

The following table lists the minimum cluster level vault privilege, vault ACL privilege, or file ACL privilege that is required to perform various operations on Exascale vaults and files. Where relevant, associated ESCLI commands are listed along with each operation.

Operation	Required Cluster Level Vault Privilege	Required Vault ACL Privilege	Required File ACL Privilege
Create vault (mkvault)	`vlt_inspect`	Not applicable.	Not applicable.
List vaults (ls)	`vlt_read`	`inspect`	Not applicable.
List files in a vault (ls)	`vlt_read`	`read`	Not applicable.
Drop vault (rmvault)	`vlt_manage`	`manage`	Not applicable.
View vault attributes (lsacl, lsxattr, lstemplate)	`vlt_read`	`inspect`	Not applicable.
Alter vault attributes (chxattr, mktemplate, rmtemplate, rmxattr)	`vlt_manage`	`manage`	Not applicable.
Alter vault ACL (chacl)	`vlt_manage`	`manage`	Not applicable.
Create file (mkfile)	`vlt_read`	`inspect`	Not applicable.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Operation	Required Cluster Level Vault Privilege	Required Vault ACL Privilege	Required File ACL Privilege
Drop file (rmfile)	`vlt_manage`	`manage`	`manage`
Read and write file contents (putfile)	`vlt_manage`	`manage`	`use`
Read file contents (getfile)	`vlt_use`	`use`	`read`
View file attributes (lsacl, lsxattr)	`vlt_read`	`read`	`inspect`
Alter file attributes (chxattr, rmxattr)	`vlt_manage`	`manage`	`use`
Alter file ACL (chacl)	`vlt_manage`	`manage`	`manage`

To perform an operation, a user requires at least one of the privileges that is listed beside the operation. For example, to open a file for read-only access the requesting user must have at least one of the following:

- The `vlt_use` cluster level vault privilege.
- The `use` vault ACL privilege for the vault containing the file.
- The `read` file ACL privilege for the file being opened.

Note

- In addition to the appropriate cluster level vault privilege or ACL privilege, starting with Oracle Exadata System Software release 25.2.3 (October 2025):
  - To perform operations that create, modify, or delete Exascale vaults, files, and related resources, a user must also have the `rest_vault_client` privilege or the `cl_admin` privilege.
  - To list vault and file-related resources, a user must also have the `rest_vault_client`, `cl_monitor`, or `cl_operator` privilege.
- To create a snapshot or a clone, the user requires the privileges for the 'read file contents' operation to read the source file, and they also require the privileges for the 'create file' operation to create a file for the snapshot or clone. After creation, operations on snapshots and clones require the same privileges as for any other file.
- Exascale ensures that users can manage the vaults and files that they create. During vault creation, if the creating user does not have the `vlt_manage` cluster level vault privilege, then Exascale adds the creating user to the vault ACL with the `manage` privilege. During file creation, if the creating user does not have the `vlt_manage` cluster level vault privilege and the user does not have the `manage` privilege in the vault ACL, then Exascale adds the creating user to the file ACL with the `manage` privilege.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Related Topics

- Modify User Privileges
- Administer Access Control Lists

1.2.20 Trust Store

The Exascale trust store is a group of digital certificates that facilitates trusted communication between servers running Exascale cluster services (EGS). The digital certificates ensure the identity of the EGS servers, which enables clients to trust the legitimacy of the Exascale cluster.

During the initial configuration of Exascale, the trust store is automatically created and distributed across the cluster.

The trust store is modified when EGS servers are added to, or deleted from, the Exascale cluster.

While most operations involving the trust store occur automatically, some manual operations are occasionally required. For example, the trust store must be manually fetched into a new user wallet.

Related Topics

- Exascale User Credentials
- Fetch the Trust Store

1.2.21 Resource Management

Resource management governs how various storage resources are consumed. Exascale resource management works in conjunction with existing Exadata I/O Resource Management (IORM) and Oracle Database Resource Manager (DBRM) capabilities.

Exascale resource management is implemented in a hierarchical manner:

- Firstly, inter-vault resource management governs how resources are allocated to different vaults using Exadata IORM.
- The resulting resource allocation for each vault is further divided among the different databases and block volumes using intra-vault resource management with Exascale resource profiles.
- Finally, the resulting resource allocation for each database is divided among PDBs and consumer groups using intra-database resource management with DBRM definitions.

The following topics describe each of these areas in greater detail.

- Inter-Vault Resource Management Using Exadata IORM
- Intra-Vault Resource Management Using Exascale Resource Profiles
- Intra-Database Resource Management Using DBRM

1.2.21.1 Inter-Vault Resource Management Using Exadata IORM

Inter-vault resource management defines how resources are allocated to different vaults. Inter-vault resource management utilizes Exadata IORM in conjunction with a vault plan, which controls resource allocation across multiple Exascale vaults.

Copyright © 2022, 2026, Oracle and/or its affiliates.

By default, each vault can access all of the resources in the Exascale cluster. However, an Exascale administrator can set various vault-specific resource provisioning attributes, which may limit the space, I/O bandwidth (IOPS), and caching resources available to the vault.

The system automatically derives the vault plan from all vault-level resource provisioning attribute settings and propagates it to every storage server in the Exascale cluster. The vault plan also updates automatically when any resource provisioning attributes change. No administrator intervention is required.

On a system that uses Exascale and traditional Exadata storage based on Oracle Automatic Storage Management (Oracle ASM), the vault plan coexists with other parts of the IORM plan. For example, an Exadata administrator can still define a cluster plan (`clusterplan`) to manage resource allocation across multiple Oracle ASM clusters. By default, unless the IORM plan specifies otherwise, resources are shared equally between Exascale and Oracle ASM.

1.2.21.2 Intra-Vault Resource Management Using Exascale Resource Profiles

Intra-vault resource management defines how resources are shared by the Exascale clients (Oracle databases and block store volumes) that use a vault. Within Exascale, intra-vault resource management is governed by resource profiles.

By default, without any resource profiles in the system, every Exascale client has access to all of the resources in their associated vault. Furthermore, I/O resources are shared equally when the system is under load.

To enable more granular I/O resource management and prioritize resource utilization, you can associate each Exascale client with a resource profile. Resource profiles enable consistent resource management across numerous clients by defining resource settings once and applying them to all associated clients. In most deployments, a small number of profiles define a resource hierarchy, and each client is assigned to a single profile.

Each resource profile contains the following resource limits and settings:

- I/O Bandwidth (IOPS) — For each Exascale media type (HC, EF, and XT), you can define a limit value and a share value. While I/O bandwidth utilization is less than the vault capacity, the limit controls each client. However, when I/O bandwidth utilization reaches capacity, resource allocation across all clients is further governed by the share values.

  The limit value specifies the absolute limit of I/O bandwidth for each client associated with the resource profile. The value represents a fraction of the available I/O bandwidth. If not specified, the default value is 100% of the available I/O bandwidth (effectively unlimited).

  The share value defines the proportional share of I/O bandwidth available to each client associated with the resource profile. Each client's share is relative to all other client shares. A higher share value implies higher priority. The range of valid share values is 1-100. If not specified, the default share value is 1.

  For example, consider a system with 2 clients, where client A has 2 shares and client B has 1 share, resulting in a total of 3 shares across the system. In this case, client A gets 66.67% of the I/O bandwidth (2 out of 3 shares), and client B gets 33.33% (1 out of 3 shares). Next, consider adding client C with 7 shares. After the addition of client C, there are a total of 10 shares across the system. Now, client A gets 20% of the I/O bandwidth (2 out of 10 shares), which is still twice as much as client B (10%, 1 out of 10 shares), but client C gets 70% of the I/O bandwidth (7 out of 10 shares).
- Flash Cache and Exadata RDMA Memory (XRMEM) Cache — For each type of cache, you can enable or disable use of the cache by clients associated with the resource profile. Then, for each enabled cache type, you can specify a minimum usage value and maximum size.

Copyright © 2022, 2026, Oracle and/or its affiliates.

The minimum value guarantees a portion of the cache for each client associated with the resource profile. The value nominally represents a fraction of the cache. However, if the sum of all minimums (across all clients and resource profiles) exceeds 100% of the cache size, then all values are scaled down proportionally to ensure that minimum guarantees can be honored. If not specified, the default minimum value is 0 (no guaranteed minimum).

The maximum size specifies the absolute limit of cache space available to each client associated with the resource profile. The value represents a fraction of the cache. If not specified, the default value is 100% of the cache size (effectively unlimited).

- Flash Log — You can enable or disable use of the flash log accelerator for clients associated with the resource profile. If not specified, the default setting enables the use of flash log.

You can define any number of resource profiles, but you can only associate an Exascale client with one resource profile at a time.

You can also create a system-reserved resource profile named `$UNASSIGNED`. All Exascale clients not explicitly associated with a resource profile are automatically governed by the `$UNASSIGNED` profile. The `$UNASSIGNED` resource profile has only two modifiable attributes that define the maximum fractions of flash cache and XRMEM cache space assigned to the profile. All other attributes of the `$UNASSIGNED` resource profile use the previously described default values.

All Exascale clients governed by the `$UNASSIGNED` profile share the specified cache resources. The behavior differs from regular resource profiles, where each application of the resource profile defines the resource allocation for one associated client.

If you do not create the `$UNASSIGNED` resource profile, the system reserves 5% of each cache for unassigned Exascale clients. These clients can also use any remaining unassigned flash cache and XRMEM cache space.

Resource over-provisioning across multiple clients and resource profiles is allowed. For over-provisioned resources, the system automatically adjusts the resource shares to maintain the relative proportions for each client.

For example, consider a system with the following resource profiles:

- GOLD resource profile:
  - EF IOPS: share=11, limit=51%
  - HC IOPS: share=40, limit=52%
  - XT IOPS: share=3, limit=unlimited
  - Flash Cache: enabled=TRUE, minimum=5%, maximum=10%
  - XRMEM Cache: enabled=TRUE, minimum=4%, maximum=8%
  - Flash Log: enabled=TRUE
- SILVER resource profile:
  - EF IOPS: share=6, limit=25%
  - HC IOPS: share=20, limit=25%
  - XT IOPS: share=2, limit=unlimited
  - Flash Cache: enabled=TRUE, minimum=3%, maximum=6%
  - XRMEM Cache: enabled=TRUE, minimum=2%, maximum=4%
  - Flash Log: enabled=TRUE

Copyright © 2022, 2026, Oracle and/or its affiliates.

- BRONZE resource profile:
  - EF IOPS: share=3, limit=9%
  - HC IOPS: share=10, limit=15%
  - XT IOPS: share=1, limit=10%
  - Flash Cache: enabled=TRUE, minimum=2%, maximum=5%
  - XRMEM Cache: enabled=TRUE, minimum=1%, maximum=2%
  - Flash Log: enabled=FALSE
- `$UNASSIGNED` resource profile:
  - Flash Cache: maximum=10%
  - XRMEM Cache: maximum=5%

Also, imagine that the system hosts 8 databases, which are associated with the resource profiles as follows:

- DB1 and DB2 are associated with the GOLD resource profile.
- DB3 is associated with the SILVER resource profile.
- DB4 and DB5 are associated with the BRONZE resource profile.
- DB6, DB7, and DB8 are not explicitly associated with any resource profile. Therefore, they are implicitly associated with the `$UNASSIGNED` resource profile.

Now consider how the resource profiles are used to manage a specific resource. For example, consider I/O bandwidth on high capacity storage media (HC IOPS):

- While the resource utilization is less than 100% of the vault capacity, each database is capped using the applicable limit value.

  For example:

Database	Resource Profile	HC IOPS Limit
DB1	GOLD	52%
DB2	GOLD	52%
DB3	SILVER	25%
DB4	BRONZE	15%
DB5	BRONZE	15%
DB6	`$UNASSIGNED`	100% (default value, effectively unlimited)
DB7	`$UNASSIGNED`	100% (default value, effectively unlimited)
DB8	`$UNASSIGNED`	100% (default value, effectively unlimited)

  If the limits keep utilization within the capacity of the vault, then no further intervention is required.
- When the resource utilization reaches the capacity of the vault, share-based allocation is used. In this case, each database receives a portion equal to its share value divided by the sum of all shares in the system.

  For example:

Copyright © 2022, 2026, Oracle and/or its affiliates.

Database	Resource Profile	HC IOPS Share Value	HC IOPS Allocation
DB1	GOLD	40	35.52% (40 out of 123 shares)
DB2	GOLD	40	35.52% (40 out of 123 shares)
DB3	SILVER	20	16.26% (20 out of 123 shares)
DB4	BRONZE	10	8.13% (10 out of 123 shares)
DB5	BRONZE	10	8.13% (10 out of 123 shares)
DB6	`$UNASSIGNED`	1 (default value)	0.81% (1 out of 123 shares)
DB7	`$UNASSIGNED`	1 (default value)	0.81% (1 out of 123 shares)
DB8	`$UNASSIGNED`	1 (default value)	0.81% (1 out of 123 shares)

Related Topics

- Administer Resource Profiles
- Using Exascale Resource Profiles with Oracle Database
- Using Resource Profiles with Exascale Volumes

1.2.21.3 Intra-Database Resource Management Using DBRM

Intra-database resource management defines how resources are shared within an Oracle database, which may include multiple pluggable databases (PDBs). Intra-database resource management uses existing Oracle Database Resource Manager (DBRM) definitions, which are transparently propagated to Exadata storage servers.

1.2.22 Exascale Block Store

The Exascale block store provides capabilities to create and manage arbitrary-sized raw block volumes based on Exascale storage. Each Exascale volume can be used as an Exascale Direct Volume (EDV) attachment.

The Exascale block store features:

- High availability - based on a clustered architecture with built-in redundancy and dynamic fail-over of software services
- High performance - utilizing Exadata RDMA memory and flash cache
- Reliability - using proven Exadata storage server technologies
- Security - employing standard security protocols and automatic volume encryption
- Rich functionality - including instantaneous volume snapshots and backups

The following topics introduce various Exascale block store concepts and features:

- Exascale Volumes
- Volume Attachments

Copyright © 2022, 2026, Oracle and/or its affiliates.

- Volume Snapshots
- Volume Clones
- Volume Backups
- Volume Groups
- Oracle Advanced Cluster File System (ACFS) on Exascale
- Virtual Machine Images

1.2.22.1 Exascale Volumes

An Exascale block volume is an arbitrary-sized allocation of storage space, which can be used as an Exascale Direct Volume (EDV) attachment.

Internally, each volume is an Exascale file with special properties that identify it as block storage space. Like other Exascale files, physical space for the volume is only materialized when data is written to the volume.

Each volume is created in a user-specified vault. You can also optionally specify the physical media type used to store the volume.

Each volume can also be associated with a series of optional attributes, which define the detailed characteristics of the volume, such as the number of data copies (mirrors) that are maintained or the system resources that a volume is allowed to consume.

To use a volume, you must create a volume attachment.

Unlike regular Exascale files, volumes are not subject to vault privileges and access control lists (ACLs). Consequently, volume administration tasks (attachment, detachment, modification, and so on) must be performed by a volume owner with the `rest_volume_client` privilege or an Exascale cluster administrator (having the `cl_admin` privilege).

Related Topics

- Administer Exascale Volumes

1.2.22.2 Volume Attachments

To use a block volume, you must create a volume attachment.

An Exascale Direct Volume (EDV) attachment creates an association between the volume and an EDV device file on the Exadata compute node or Oracle Grid Infrastructure (GI) cluster hosting the attachment.

If you create a cluster-wide attachment, the EDV device file is created on every node in the GI cluster. If you create a node-specific attachment, the corresponding EDV device is only created on that node.

After attachment, the volume can be used as a raw block storage device or to support a file system. To implement Oracle Advanced Cluster File System (ACFS) on Exascale block volume storage, you must use an EDV attachment.

The I/O for each attachment is serviced by an EDV server process. Automatic data encryption and decryption occurs within the EDV process.

In addition to general file storage, an EDV attachment may be used to store data files for Oracle Database 19c and Oracle Database 21c. However, an EDV attachment cannot be used for a bootable volume.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Exascale supports the simultaneous use of multiple attachments for each volume, which can be used to support various types of clustered file systems.

Related Topics

- Administer Volume Attachments

1.2.22.3 Volume Snapshots

A volume snapshot is a thinly-provisioned read-only point-in-time copy of a volume.

Volume snapshots have the following uses:

- You can create attachments to a volume snapshot and use it as a read-only volume.
- You can use a volume snapshot as the source for a volume clone.
- You can use a volume snapshot as the source for a volume backup.

Additionally, starting with Oracle Exadata System Software release 26.1.0, you can lock a volume snapshot to prevent its deletion before a specified future date and time. You can also make the lock immutable, which prevents users from removing the lock or changing it to expire earlier.

Related Topics

- Administer Volume Snapshots
- Lock Volume Snapshots

1.2.22.4 Volume Clones

A volume clone is a thinly-provisioned read-write point-in-time copy of a volume snapshot.

A volume clone is functionally equivalent to a standard (non-cloned) volume.

Volume clones have the following uses:

- You can create attachments to a volume clone and use it like any other writable volume.
- You can create volume snapshots that are based on the volume clone.

Related Topics

- Administer Volume Clones

1.2.22.5 Volume Backups

A volume backup is a backup of an Exascale volume snapshot, which provides a consistent point-in-time copy of the volume.

The backup is stored in an Exascale vault, and the backup destination vault may differ from the vault containing the volume being backed up. Furthermore, Exascale automatically places file extents for volume backups on physically separate devices from the volume being backed up. This organization guarantees that if multiple failures affect a volume, its corresponding backups remain available.

You can restore a volume backup into a new volume. When you restore a volume backup, the restored volume is separate from the volume backup and the original backup source volume. There is no ongoing association between a restored volume and the volume backup or the original backup source volume.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Related Topics

- Administer Volume Backups

1.2.22.6 Volume Groups

A volume group is a specific collection of Exascale volumes.

Volume groups can be used to:

- Easily identify volumes with a common purpose or inherent relationship.

  For example, you can use a volume group to collect all the volumes underpinning the image files associated with an individual virtual machine (VM). Also, you can optionally name the volume group, making it easier to understand its purpose and distinguish it from numerous other groups.

- Collectively manage I/O resources consumed by all volumes in the group.

  A volume group can optionally be configured as a resource-sharing volume group, using either aggregate or specified resource sharing.

  Under aggregate resource sharing, the volume group automatically shares the total I/Os per second (IOPS) bandwidth provisioned to all volumes in the group. In this case, an individual volume can exceed its volume-level limit by sharing unused IOPS provisioned to other volumes in the group. But the entire group is always governed by the aggregated limit.

  With specified resource sharing, a specific IOPS bandwidth limit is associated with the volume group. In this case, the group is governed only by the specified limit.

  In either case, a resource-sharing volume group also works in conjunction with I/O resource management (IORM) directives defined at other levels, such as the Exascale vault level.

- Efficiently create consistent snapshots of all volumes in a volume group.

  In one operation, you can create a set of volume snapshots with point-in-time consistency, which are easily identifiable by a unique batch ID.

Note the following details about volume groups:

- A volume group can contain only writable volumes, including volume clones. Volume snapshots cannot be part of a volume group.
- All volumes belonging to a volume group must reside in the same Exascale vault.
- A volume must always have at least one owner in common with each volume group it belongs to.
- A volume can be a member of up to five different volume groups. However, at any time, a volume can belong to only one resource-sharing volume group.

Related Topics

- Administer Volume Groups

1.2.22.7 Oracle Advanced Cluster File System (ACFS) on Exascale

Exascale contains integrated support for Oracle Advanced Cluster File System (ACFS) on Exascale block storage using Exascale Direct Volumes (EDV).

By using this facility, an Exascale administrator can quickly and easily create and manage ACFS file systems on Exascale block storage.

Copyright © 2022, 2026, Oracle and/or its affiliates.

For example, using one simple command an Exascale administrator can:

1. Create an ACFS file system on the specified EDV block storage volume.
2. Register the ACFS details with the Oracle Grid Infrastructure (GI) cluster
3. Mount the file system on all nodes associated with the GI cluster.

Related Topics

- Administer an Advanced Cluster File System (ACFS) on Exascale
- About Oracle ACFS

1.2.22.8 Virtual Machine Images

Exadata systems configured with Exascale can use Exascale Direct Volumes (EDV) to store virtual machine (VM) guest image files for Exadata database server VMs.

This capability removes previous constraints caused by the limited local storage space on each Exadata database server. It can be used regardless of whether the databases running inside each VM use native Exascale storage or Oracle ASM.

Furthermore, VM images on Exascale are easily accessible from any VM host. This effectively decouples VM guests and hosts, providing the infrastructure to enable quick and easy migration of a VM guest to another host.

Decoupling VM guests and hosts also opens up additional possibilities for:

- Balancing database server VM workloads by separating busy VMs onto different hosts.
- Reducing the impact of scheduled VM host maintenance by proactively moving VMs to another host.
- Reducing the impact of unscheduled VM host downtime by reactively moving VMs to another host.

Copyright © 2022, 2026, Oracle and/or its affiliates.

2

Oracle Exadata Exascale System Administration

This chapter covers administration tasks that define and manage an Exascale system.

Administration activities for Oracle Exadata Exascale are logically divided into two groups. System administration activities define and manage the Exascale system components that are shared across the system. User-specific administration activities define and manage Exascale storage objects and metadata that belong to a specific Exascale user.

This chapter contains the following system administration topics:

- Configure Exascale
- Administer the Exascale Cluster
- Administer Exascale Users
- Administer Exascale Pool Disks
- Administer Exascale Storage Pools
- Administer the Storage Devices in an Exascale Storage Pool
- Administer Cluster Templates

2.1 Configure Exascale

- Configure Exascale Using the OEDA Web User Interface

2.1.1 Configure Exascale Using the OEDA Web User Interface

You can use the Oracle Exadata Deployment Assistant (OEDA) Web user interface to initially configure Exascale on a fresh Exadata deployment, or to add Exascale alongside an existing Exadata deployment.

To initially configure Exascale on a fresh Exadata deployment:

1. Use the OEDA Web user interface to create a new Exadata XML configuration file (`es.xml`) that includes configuration details for an Exascale cluster and Exascale-enabled Oracle Database and Oracle Grid Infrastructure installations.

   See OEDA Web User Interface Settings and Options for Exascale.

2. Use the OEDA deployment utility (`install.sh`) to configure the fresh Exadata deployment as specified in the Exadata XML configuration file (`es.xml`).

To add Exascale alongside an existing Exadata deployment:

1. Ensure that the existing Exadata system configuration has free space in the storage servers to accommodate Exascale. The total amount of required free space is governed by your projected use of Exascale.

   If required, reconfigure the existing ASM disk groups and Exadata grid disks to ensure that every Exadata cell disk contains the same amount of space available for Exascale to use.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Also, ensure that any adjustments to the storage configuration are reflected in the Exadata XML configuration file (`es.xml`).

2. Start the OEDA Web user interface and import the Exadata XML configuration file (`es.xml`) from your existing Exadata deployment.
3. Use the OEDA Web user interface to modify the Exadata configuration as follows:

   See also OEDA Web User Interface Settings and Options for Exascale.

   a. Navigate to the Exascale page and click the check box to **Enable Exascale**.
   b. Use the Exascale page to define an Exascale cluster.
   c. Use the Define Clusters page to add new Exascale-enabled VM clusters alongside your existing ASM-based VM clusters.
   d. Use the Create Database page to add new databases that use Exascale storage.
   e. Generate the updated Exadata XML configuration files.
4. Use the OEDA deployment utility (`install.sh`) in conjunction with the cluster-specific XML configuration files for the new Exascale-enabled VM clusters.

- OEDA Web User Interface Settings and Options for Exascale
- Exascale Post Deployment Settings

2.1.1.1 OEDA Web User Interface Settings and Options for Exascale

The following list describes the Exascale-specific settings and options in the OEDA Web user interface.

- **Exascale page**

  OEDA includes the Exascale page, which you use to define the Exascale cluster.

  To commence configuration, you must click the check box to **Enable Exascale**. Otherwise, the following configuration options are not available.

  After you click the check box to **Enable Exascale**, you can use the Exascale page to specify the following settings and options:

  – **Cluster Name**: Specify the name used to identify the Exascale cluster.
  – **ERS IP Address**: Specify an IP address for Exascale control services (also known as Exascale RESTful Services or ERS).

    The specified IP address should reside in the Exadata administration network (also known as the management network).

    The specified IP address is associated with a highly-available virtual IP (VIP) network interface, which provides a consistent network end point for Exascale control services. The VIP is hosted by one of the Exadata storage servers, which also runs front-end ERS processes. If the storage server or ERS instance becomes unavailable, then the VIP moves to another server containing front-end ERS processes.
  – **ERS Host Name**: Specify a host name for Exascale control services (also known as Exascale RESTful Services or ERS).

    The specified host name is associated with the **ERS IP Address**.
  – **Available nodes/Selected nodes**: Specify the Exadata storage servers that become members of the Exascale cluster.

Copyright © 2022, 2026, Oracle and/or its affiliates.

The specified cells must all be of the same Exadata storage server type. For example, all high capacity (HC).

- **Storage Pool Name**: Specify the name for the storage pool.

  You can only define one storage pool for each Exascale cluster using the OEDA Web user interface.

  > **Note**
  >
  > The media type for the storage pool matches the underlying Exadata storage server type. For example, if the Exadata system uses high capacity (HC) storage cells, then the storage pool media type will be `HC`.

- **Storage Pool Size (%/GB/TB)**: Specify the amount of storage space to provision in the storage pool.

  Specify a percentage of the available space or an amount of storage space in gigabytes (GB) or terabytes (TB).

On the Exascale page, the **Advanced** button displays a window containing:

- **Disable ERS IP Address**: This option configures Exascale without a highly-available VIP network interface that provides a consistent network end point for Exascale control services. If you select this option, then **ERS IP Address** and **ERS Host Name** are disabled.

  > **Note**
  >
  > This option enables Exascale configurations on existing Exadata systems that have no spare IP addresses to support the ERS VIP.
  >
  > Using this option is generally not recommended for customer Exascale deployments!

- **Enable ERS Network ID**: This option adds a field to specify the ERS Network ID value associated with the Exascale cluster. The ERS Network ID is used to maintain the highly-available virtual IP (VIP) network interface that provides a consistent network end point for Exascale control services (ERS).

  > **Note**
  >
  > This option is not required when deploying only one Exascale cluster in the subnet supporting the ERS IP address.
  >
  > However, when multiple Exascale clusters have ERS IP addresses that share the same subnet, each Exascale cluster requires a unique ERS Network ID.
  >
  > By default, OEDA attempts to allocate a unique ERS Network ID to each Exascale cluster. However, this option enables users to provide a specific ERS Network ID value.

- **Define Clusters page**

  The Define Clusters page configures the Oracle Grid Infrastructure (GI) clusters on the Exadata Database Machine.

Copyright © 2022, 2026, Oracle and/or its affiliates.

To enable Exascale storage on a GI cluster, you must use the **Exascale clusters** list to select an Exascale cluster that you defined previously on the Exascale page. If you do not select an Exascale cluster, then Oracle Automatic Storage Management (ASM) is used for the GI cluster and all of the databases that it contains.

Enabling Exascale storage on a GI cluster implicitly creates an Exascale vault, which is used to store the Oracle Cluster Registry (OCR) and voting files. Consequently, you must also specify:

- **Vault Name**: Specify a name to identify the Exascale vault.
- **Vault Size**: Specify the amount of storage to provision for the vault.

Additionally, if the system is configured to use Virtual Machines (VMs) for the database servers, you are presented with the option to **Use Exascale for VM file storage**. If you select the check box to enable this option, Exascale volumes are used to store the guest image files for the Exadata database server VMs.

Using Exascale volumes to store VM images removes previous constraints imposed by the limited amount of local storage space available on each physical Exadata database server. Furthermore, VM images on Exascale volumes are easily accessible from any VM host. This capability effectively decouples VM guests and hosts, providing the infrastructure to enable quick and easy migration of a VM guest to another host.

All other settings and options on the Define Clusters page are unaffected by Exascale.

> **Note**
>
> In each GI cluster, all databases must use one database storage model. Databases in the same GI cluster can use either native Exascale storage or Oracle ASM, but not both. Hybrid Exascale and Oracle ASM configurations are supported on the same Exadata system by using separate GI clusters.
>
> Furthermore, on an Exascale-enabled system, Exascale volumes can store guest image files for Exadata database server VMs associated with any GI cluster, regardless of whether the databases running inside the VMs use native Exascale storage or Oracle ASM.

- **Create Database page**

  The Create Database page is used to create Oracle databases on the Exadata Database Machine.

  To enable Exascale storage for a database, you can:

  - Use the **Available vaults for this database** list to select a previously defined Exascale vault.
  - Click **Add Vault**, and specify the **Vault Name** and **Vault Size**.

  All other settings and options on the Create Database page are unaffected by Exascale.

2.1.1.2 Exascale Post Deployment Settings

This section describes the Exascale objects and settings left after deployment of an Exascale configuration created with the OEDA Web user interface.

- Exascale Storage Configuration
- Exascale User Accounts and Wallets

Copyright © 2022, 2026, Oracle and/or its affiliates.

2.1.1.2.1 Exascale Storage Configuration

This topic describes the Exascale storage objects created during system deployment.

During system deployment, one Exascale storage pool is created for each Exascale cluster defined on the OEDA Web UI Exascale page:

- The storage pool uses the **Storage Pool Name** and sizing details (**Storage Pool Size** and **Storage Pool Size Type**) specified on the Exascale page.
- The storage pool spans the Exadata storage servers specified in the **Selected Cells** list.
- The storage pool media type is derived from the type of Exadata storage servers in use. For example, an `HC` storage pool is created on high capacity (HC) Exadata storage servers.
- To support the storage pool, an Exascale pool disk is created on each primary storage device in the Exadata storage servers. Each pool disk is uniformly sized as a fraction of the overall storage pool size specified on the Exascale page.

Additionally, Exascale vaults are created according to definitions specified in the OEDA Web UI on the Define Clusters page and the Create Database page.

To use Exascale storage, each compute cluster (database server cluster or VM cluster) in the Define Clusters page is associated with an Exascale cluster. An Exascale cluster can support multiple compute clusters. However, each compute cluster requires a separate vault to store the Oracle Grid Infrastructure shared clusterware files (Oracle Cluster Registry and voting disks). Consequently, one vault is created for each Exascale-enabled compute cluster.

Then, for each database that you define on an Exascale-enabled compute cluster, you can use the **Available Vaults** list on the Create Database page to place the Oracle Database files (data files, control files, log files, and so on) on a previously defined vault. Or, you can use the **Add new Vault** button to define a new vault to house the Oracle Database files.

During deployment, each vault that is created uses the **Vault Name** specified in the OEDA Web UI. The corresponding **Vault Size** is applied to the vault space provisioning attribute associated with the underlying storage pool media type. For example, if the vault is created on an Exascale cluster containing an `HC` storage pool, the `spaceProvHC` vault attribute is set to **Vault Size**.

2.1.1.2.2 Exascale User Accounts and Wallets

Exascale has a system of user accounts enabling different users to perform actions and access data according to their assigned privileges. This topic describes the Exascale user accounts and associated key stores (wallets) created during system deployment.

After system deployment, the number of Exascale user accounts, and their settings, depends on the configuration details specified on the OEDA Web UI User(s) and Groups page and the clusters defined on the Define Clusters page.

The following occurs for each virtual machine (VM) cluster or bare-metal database server cluster on the Define Clusters page:

- If the cluster is associated with a default (non-role-separated) user configuration containing one Oracle OS user account, then one matching Exascale user account is created.

  In this case, the Exascale user manages the Oracle Grid Infrastructure shared clusterware files (Oracle Cluster Registry and voting disks) and the Oracle Database files (data files, control files, log files, and so on) for all databases associated with the cluster. Consequently, the Exascale user is the owner and manager of all Exascale vaults associated with the cluster.

Copyright © 2022, 2026, Oracle and/or its affiliates.

The Exascale user identifier (ID) is set to a concatenation of:

- The **User Name** (typically `oracle`) as specified on the User(s) and Groups page.
- The value of the **Cluster Name** as specified on the Define Clusters page.

For example, if the **User Name** is `oracle` and the **Cluster Name** is `Escluster1`, then the Exascale user ID is `oracleEscluster1`.

The Exascale user account uses a system-generated public and private key pair for authentication. The Exascale user account definition contains the public key, and each compute node (database server or VM) contains a system-generated key store (wallet), which contains the corresponding private key. The wallet is created at `/etc/oracle/cell/network-config/eswallet`.

- If the cluster is associated with a role-separated user configuration containing one Grid OS user account and one Oracle OS user account, then two matching Exascale user accounts are created.

  The Exascale user that matches the Grid OS user manages the Oracle Grid Infrastructure shared clusterware files (Oracle Cluster Registry and voting disks). Also known as the grid Exascale user, this Exascale user account has the following characteristics:

  - The user ID is set to a concatenation of:
    - The **User Name** for the Grid OS user (typically `grid`) as specified on the User(s) and Groups page.
    - The value of the **Cluster Name** as specified on the Define Clusters page.

    For example, if the **User Name** is `grid` and the **Cluster Name** is `Escluster2`, then the Exascale user ID is `gridEscluster2`.
  - The Exascale user account uses a system-generated public and private key pair for authentication. The Exascale user account definition contains the public key, and each compute node (database server or VM) contains a system-generated key store (wallet), which contains the corresponding private key. The wallet for the grid Exascale user account is created at `/etc/oracle/cell/network-config/eswallet`.
  - The Exascale grid user is the owner and manager of the Exascale vault specified on the Define Clusters page.

  The Exascale user that matches the Oracle OS user manages the Oracle Database files (data files, control files, log files, and so on) for all databases associated with the cluster. Also known as the oracle Exascale user, this Exascale user account has the following characteristics:

  - The user ID is set to a concatenation of:
    - The **User Name** for the Oracle OS user (typically `oracle`) as specified on the User(s) and Groups page.
    - The value of the **Cluster Name** as specified on the Define Clusters page.

    For example, if the **User Name** is `oracle` and the **Cluster Name** is `Escluster2`, then the Exascale user ID is `oracleEscluster2`.
  - The Exascale user account uses a system-generated public and private key pair for authentication. The Exascale user account definition contains the public key, and each compute node (database server or VM) contains a system-generated key store (wallet), which contains the corresponding private key. The wallet for the Exascale user account that corresponds with the Oracle OS user is created at `$ORACLE_BASE/admin/eswallet` (typically `/u01/app/oracle/admin/eswallet`).

Copyright © 2022, 2026, Oracle and/or its affiliates.

  - The Exascale oracle user is the owner and manager of the database-specific Exascale vaults specified on the Create Database page.

    However, if any database uses the vault that contains the Oracle Grid Infrastructure shared clusterware files, then the Exascale oracle user is added as a manager of that vault.

In addition to the aforementioned Exascale user accounts, each Exascale cluster contains:

- One superuser account. The user ID for the superuser account is `admin`. The `admin` user can implicitly perform any system operation and effectively holds all system privileges.

  During system deployment, the `admin` user wallet is created on every storage server at `/opt/oracle/cell/cellsrv/deploy/config/security/admwallet` and every wallet contains the same system-generated private key.
- One node administration account for every node (storage server or compute node) that runs Exascale software services. Each node administration account inherits its user ID from the server hostname and each account contains the privileges required to run the Exascale software services on the node. Do not directly use or modify these accounts.

  During system deployment, the node administration user wallet is created on every node, and every wallet contains a system-generated private key. On a storage server, the wallet is located at `/opt/oracle/cell/cellsrv/deploy/config/eswallet`. On a compute node, the wallet is located at `/opt/oracle/dbserver/dbms/deploy/config/eswallet`.

2.2 Administer the Exascale Cluster

- Administer Exascale Services
- Modify the Exascale Cluster
- Stop the Exascale Cluster
- Add a Storage Server
- Manage Certificates for Exascale Control Services

2.2.1 Administer Exascale Services

- About Automatic Service Placement
- List Services
- Alter the Number of Service Instances
- Alter the Placement of Services
- Restart a Service
- Stop a Service
- Start a Service
- Disable a Service
- Administer Client Services

Related Topics

- About Exascale Services

Copyright © 2022, 2026, Oracle and/or its affiliates.

2.2.1.1 About Automatic Service Placement

Exascale system components are implemented using a series of clustered software services. This topic describes the default automatic placement of the Exascale storage services.

By default, Exascale uses automatic service placement to ensure that the required Exascale storage services are available across the Exascale cluster. If Exascale detects that a storage service is not running, the service instance is automatically restarted on another node.

The following outlines the default automatic service placement methodology for Exascale storage services:

- **Cluster Services (EGS)**

  Every Exascale cluster contains five EGS instances.

  For Exascale clusters with five or more Exadata storage servers, one EGS instance runs on each of five storage servers.

  For Exascale configurations with fewer than five storage servers, one EGS instance runs on each Exadata storage server, and the remaining EGS instances run on the Exadata compute nodes to make up the required total of five. In a bare-metal configuration, EGS compute node instances run on the server operating system. In a configuration with compute nodes running on virtual machines (VMs), EGS compute node instances run in the hypervisor.

- **Control Services (ERS)**

  For Exascale clusters with five or more Exadata storage servers, a total of five ERS instances are distributed across the Exadata storage servers. The first three ERS instances contain front-end and back-end processes, and the final two instances only contain back-end processes.

  For Exascale configurations with fewer than five storage servers, one ERS instance runs on each storage server. The first three ERS instances contain front-end and back-end processes, and the remaining instances only contain back-end processes.

- **System Vault Manager (SYSEDS) and User Vault Manager (USREDS)**

  For Exascale clusters with five or more Exadata storage servers, a total of five SYSEDS instances and five USREDS instances are distributed across the Exadata storage servers.

  For Exascale configurations with fewer than five storage servers, one SYSEDS instance and one USREDS instance runs on each storage server.

- **Block Store Manager (BSM) and Block Store Worker (BSW)**

  By default, a total of three BSM instances and three BSW instances are distributed across the Exadata storage servers.

In all cases, the Exascale storage service instances are automatically distributed across the available storage servers with the aim of evenly distributing the overall workload.

In addition to the Exascale storage services governed by automatic service placement, every storage server that is associated with the Exascale cluster also contains the Instance Failure Detection (IFD) service and the core Exadata cell services: Cell Server (CELLSRV), Management Server (MS), and Restart Server (RS).

2.2.1.2 List Services

Exascale system components are implemented using a series of software services.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- To display information about Exascale storage services running across the Exascale cluster, including Exascale cluster services (EGS), control services (ERS), vault manager services (SYSEDS and USREDS), and block store services (BSM and BSW), use the ESCLI lsservice command. For example:

  ```
  @> lsservice --detail
  ```

  Using the various command options available with the lsservice command, you can view output with different levels of detail and specify the services you want to display.

- To display information about client-side Exascale services, including Exascale Node Proxy (ESNP) and Exascale Direct Volume (EDV), use the DBMCLI `LIST DBSERVER` command, as follows:

  ```
  DBMCLI> list dbserver detail
  ```

  The command output displays status information about all of the client-side services running on the Exadata compute node. You must run the command separately on each compute node.

2.2.1.3 Alter the Number of Service Instances

Exascale system components are implemented using a series of software services.

By default, each Exascale cluster is configured with a number of storage service instances that is usually sufficient for most workloads. However, as an Exascale administrator, you can alter the number of service instances that are deployed.

To alter the number of storage service instances deployed in the Exascale cluster, use the ESCLI chcluster command as follows:

```
@> chcluster --attributes attributeName=value
```

In the command, specify one of the following as the *attributeName*:

- `numControlServicesFrontends`: The number of Control Service (ERS) instances that contain front-end and back-end processes.
- `numControlServicesBackends`: The number of Control Service (ERS) instances that contain only back-end processes.
- `numSystemVaultManagers`: The number of System Vault Manager (SYSEDS) instances.
- `numUserVaultManagers`: The number of User Vault Manager (USREDS) instances.
- `numVolumeManagers`: The number of Block Store Manager (BSM) instances.
- `numStorageVolumeWorkers`: The number of Block Store Worker (BSW) instances.

> **Note**
>
> You cannot alter the number of Exascale cluster services (EGS). EGS always runs on five nodes in the Exascale cluster.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Related Topics

- About Automatic Service Placement

2.2.1.4 Alter the Placement of Services

Exascale system components are implemented using a series of software services.

By default, the storage service instances are distributed across the Exascale cluster with the aim of evenly distributing the overall workload. However, as an Exascale administrator, you can alter the placement of the services to optimize the workload balance.

To alter the placement of Exascale storage services, you can start a service on a node where it is not running and then disable the same service on another node where it is running.

For example, you can use the following ESCLI command sequence to alter the placement of the block store worker service (BSW) by starting the service on `cell09` and disabling it on `cell03`.

```
@> chservice --attributes name=bsw_cell09 --start
@> chservice --attributes name=bsw_cell03 --disable
```

> **Note**
>
> You cannot alter the placement of Exascale cluster services (EGS). EGS always runs on five nodes in the Exascale cluster.

Related Topics

- About Automatic Service Placement
- Start a Service
- Disable a Service

2.2.1.5 Restart a Service

Exascale system components are implemented using a series of software services.

Typically, the Exascale cluster automatically contains the required Exascale storage services. However, as an Exascale administrator, you may occasionally need to restart a software service instance.

To restart an Exascale storage service, use the ESCLI chservice command with the `--restart` option.

If you issue a command to restart a stopped or disabled service, then the command starts the service and completes successfully.

- To restart an Exascale storage service on an Exadata storage server, the command syntax is:

  ```
  @> chservice --attributes name=service_cell[,frontend={true|false}] --
  restart
  ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

To identify the service instance you want to restart on an Exadata storage server, specify the `name` attribute using the format: `name=service_cell`.

In the attribute value:

– *service*: Specifies the service that you want to act on, which is one of the following:
  * `cellsrv`: Specifies the Exadata Cell Server (CELLSRV).
  * `egs`: Specifies the Exascale cluster service, also known as Exascale global service (EGS).
  * `ers`: Specifies the Exascale control service, also known as Exascale RESTful service (ERS).
  * `syseds`: Specifies the system vault manager service.
  * `usreds`: Specifies the user vault manager service.
  * `bsm`: Specifies the block store manager (BSM) service.
  * `bsw`: Specifies the block store worker (BSW) service.
– *cell*: Specifies the name of the storage server (`cell`) that hosts the service.

  The cell must be identified by its name. Use the lscell command to find all cell names.

Additionally, when you restart ERS only, you can optionally set the `frontend` attribute to control the front-end server processes on the ERS instance. To enable the front-end server processes on the ERS instance, specify `frontend=true`. To disable the front-end server processes on the ERS instance, specify `frontend=false`. The `frontend` attribute only applies to ERS.

- To restart an Exascale storage service on an Exadata compute server, the command syntax is:

```
@> chservice --attributes name=service_compute --restart
```

To identify the service instance you want to restart on an Exadata compute server, specify the `name` attribute using the format: `name=service_compute`.

In the attribute value:

– *service*: Specifies the service that you want to act on, which is one of the following:
  * `egs`: Specifies the Exascale cluster service, also known as Exascale global service (EGS).
  * `bsw`: Specifies the block store worker service.
– *compute*: Specifies the name of the compute server (`dbserver`) that hosts the service.

  Depending on the required service location, specify either the name of a specific compute node virtual machine (VM) or the Exadata compute node host domain (KVM host or Dom0). Use the lscomputeserver command to display information about compute servers associated with the Exascale cluster.

2.2.1.6 Stop a Service

Exascale system components are implemented using a series of software services.

Typically, the Exascale cluster automatically contains the required Exascale storage services. However, as an Exascale administrator, you may occasionally need to stop a software service instance.

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> By default, Exascale uses automatic service placement to ensure that the required storage services are available across the Exascale cluster. Consequently, if you manually stop a software service instance (as described below), a replacement service instance is automatically started on another node (if possible).

To stop an Exascale storage service, use the ESCLI chservice command with the `--stop` option.

If you issue a command to stop a service that is not running, then the command completes successfully but it has no effect.

- To stop an Exascale storage service on an Exadata storage server, the command syntax is:

  ```
  @> chservice --attributes name=service_cell --stop
  ```

  To identify the service instance you want to stop on an Exadata storage server, specify the `name` attribute using the format: `name=service_cell`.

  In the attribute value:

  – *service*: Specifies the service that you want to act on, which is one of the following:

    * `cellsrv`: Specifies the Exadata Cell Server (CELLSRV).
    * `egs`: Specifies the Exascale cluster service, also known as Exascale global service (EGS).
    * `ers`: Specifies the Exascale control service, also known as Exascale RESTful service (ERS).
    * `syseds`: Specifies the system vault manager service.
    * `usreds`: Specifies the user vault manager service.
    * `bsm`: Specifies the block store manager (BSM) service.
    * `bsw`: Specifies the block store worker (BSW) service.

  – *cell*: Specifies the name of the storage server (`cell`) that hosts the service.

    The cell must be identified by its name. Use the lscell command to find all cell names.

- To stop an Exascale storage service on an Exadata compute server, the command syntax is:

  ```
  @> chservice --attributes name=service_compute --stop
  ```

  To identify the service instance you want to stop on an Exadata compute server, specify the `name` attribute using the format: `name=service_compute`.

  In the attribute value:

  – *service*: Specifies the service that you want to act on, which is one of the following:

    * `egs`: Specifies the Exascale cluster service, also known as Exascale global service (EGS).
    * `bsw`: Specifies the block store worker service.

Copyright © 2022, 2026, Oracle and/or its affiliates.

  - *compute*: Specifies the name of the compute server (`dbserver`) that hosts the service.

    Depending on the required service location, specify either the name of a specific compute node virtual machine (VM) or the Exadata compute node host domain (KVM host or Dom0). Use the lscomputeserver command to display information about compute servers associated with the Exascale cluster.

2.2.1.7 Start a Service

Exascale system components are implemented using a series of software services.

Typically, the Exascale cluster automatically contains the required Exascale storage services. However, as an Exascale administrator, you may occasionally need to start a software service instance.

To start an Exascale storage service, use the ESCLI chservice command with the `--start` option.

If you issue a command to start an already running service, then the command completes successfully but it has no effect.

- To start an Exascale storage service on an Exadata storage server, the command syntax is:

  ```
  @> chservice --attributes name=service_cell --start
  ```

  To identify the service instance you want to start on an Exadata storage server, specify the `name` attribute using the format: `name=service_cell`.

  In the attribute value:

  - *service*: Specifies the service that you want to act on, which is one of the following:
    - `cellsrv`: Specifies the Exadata Cell Server (CELLSRV).
    - `egs`: Specifies the Exascale cluster service, also known as Exascale global service (EGS).
    - `ers`: Specifies the Exascale control service, also known as Exascale RESTful service (ERS).
    - `syseds`: Specifies the system vault manager service.
    - `usreds`: Specifies the user vault manager service.
    - `bsm`: Specifies the block store manager (BSM) service.
    - `bsw`: Specifies the block store worker (BSW) service.
  - *cell*: Specifies the name of the storage server (`cell`) that hosts the service.

    The cell must be identified by its name. Use the lscell command to find all cell names.
- To start an Exascale storage service on an Exadata compute server, the command syntax is:

  ```
  @> chservice --attributes name=service_compute --start
  ```

  To identify the service instance you want to start on an Exadata compute server, specify the `name` attribute using the format: `name=service_compute`.

  In the attribute value:

Copyright © 2022, 2026, Oracle and/or its affiliates.

- *service*: Specifies the service that you want to act on, which is one of the following:
  - `egs`: Specifies the Exascale cluster service, also known as Exascale global service (EGS).
  - `bsw`: Specifies the block store worker service.
- *compute*: Specifies the name of the compute server (`dbserver`) that hosts the service.

  Depending on the required service location, specify either the name of a specific compute node virtual machine (VM) or the Exadata compute node host domain (KVM host or Dom0). Use the lscomputeserver command to display information about compute servers associated with the Exascale cluster.

2.2.1.8 Disable a Service

Exascale system components are implemented using a series of software services.

Typically, the Exascale cluster automatically contains the required Exascale storage services. However, as an Exascale administrator, you may occasionally need to disable a software service instance.

To disable an Exascale storage service, use the ESCLI chservice command with the `--disable` option.

A disabled service instance is stopped and prevented from automatically restarting on the associated server. To enable a disabled service instance, you must start it manually.

- To disable an Exascale storage service on an Exadata storage server, the command syntax is:

  ```
  @> chservice --attributes name=service_cell --disable
  ```

  To identify the service instance you want to disable on an Exadata storage server, specify the `name` attribute using the format: `name=service_cell`.

  In the attribute value:

  - *service*: Specifies the service that you want to act on, which is one of the following:
    - `egs`: Specifies the Exascale cluster service, also known as Exascale global service (EGS).
    - `ers`: Specifies the Exascale control service, also known as Exascale RESTful service (ERS).
    - `syseds`: Specifies the system vault manager service.
    - `usreds`: Specifies the user vault manager service.
    - `bsm`: Specifies the block store manager (BSM) service.
    - `bsw`: Specifies the block store worker (BSW) service.
  - *cell*: Specifies the name of the storage server (`cell`) that hosts the service.

    The cell must be identified by its name. Use the lscell command to find all cell names.
- To disable an Exascale storage service on an Exadata compute server, the command syntax is:

  ```
  @> chservice --attributes name=service_compute --disable
  ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

To identify the service instance you want to disable on an Exadata compute server, specify the `name` attribute using the format: `name=service_compute`.

In the attribute value:

- *service*: Specifies the service that you want to act on, which is one of the following:
  - `egs`: Specifies the Exascale cluster service, also known as Exascale global service (EGS).
  - `bsw`: Specifies the block store worker service.
- *compute*: Specifies the name of the compute server (`dbserver`) that hosts the service.

  Depending on the required service location, specify either the name of a specific compute node virtual machine (VM) or the Exadata compute node host domain (KVM host or Dom0). Use the lscomputeserver command to display information about compute servers associated with the Exascale cluster.

You can also add the `--force` option to forcibly disable the service, even if the server hosting the service is unavailable.

2.2.1.9 Administer Client Services

In addition to the Exascale storage services, further client-side services provide support for specific Exascale functions on the Exadata compute nodes.

By default, the following client-side service instances are configured on each Exascale compute node:

- **Exascale Node Proxy (ESNP)**

  In a bare-metal configuration, ESNP runs on the server operating system. In a configuration with compute nodes running on virtual machines (VMs), ESNP runs inside the guest VMs.

- **Exascale Direct Volume (EDV)**

  This service is required on all Exadata compute nodes where you want to use EDV devices. In a bare-metal configuration, the EDV service runs on the server operating system. In a configuration with compute nodes running on virtual machines (VMs), the EDV service runs in the KVM host to provide the hypervisor access to VM image files based on EDV-managed storage. The EDV service also runs in each guest VM to enable direct access to EDV-managed storage from inside the VM.

To administer the client-side services, use the `ALTER DBSERVER` command in the Oracle Exadata Database Machine Command-Line Interface (DBMCLI).

The relevant command syntax is:

```
DBMCLI> ALTER DBSERVER { STARTUP | SHUTDOWN | RESTART } SERVICES { ESNP |
EDV }
```

To display the status of the client-side services, use the DBMCLI `LIST DBSERVER` command.

Related Topics

- List Services

Copyright © 2022, 2026, Oracle and/or its affiliates.

2.2.2 Modify the Exascale Cluster

Exascale system components are implemented using a series of clustered software services. You can control various aspects of the Exascale cluster by modifying the cluster attributes.

To modify attributes of the Exascale cluster, use the ESCLI chcluster command and specify one or more attribute settings. For example:

```
@> chcluster --attributes alterServiceTimeoutInSecs=60
```

2.2.3 Stop the Exascale Cluster

Exascale system components are implemented using a series of clustered software services.

To stop all of the Exascale software services in the Exascale cluster, use the ESCLI chcluster command with the `--shutdown` option.

However, to shut down the Exascale cluster, note that you must:

- Run ESCLI as the Exascale `admin` user or as another user with the `cl_admin` cluster level storage privilege.
- Connect ESCLI directly to an online Exascale control services (ERS) back-end server process. To do this you must start ESCLI and specify the `--ctrl` option as follows:

  ```
  $ escli --wallet admin-wallet-location --ctrl ERS-server-IP:8080
  ```

  To find the IP address for an online ERS server (*ERS-server-IP*), start a regular ESCLI session and use the following command:

  ```
  @> lsservice --filter serviceType=controlServices,status=ONLINE --
  attributes url
  ```

  The command displays the private IP addresses associated with all of the online ERS servers. You can use any of the reported IP addresses.

For example:

```
$ escli --wallet admin-wallet-location --ctrl ERS-server-IP:8080 chcluster --
shutdown
```

2.2.4 Add a Storage Server

Adding an Exadata storage server to an existing Exascale cluster enables you to:

- Add storage space to the Exascale storage pools.
- Run Exascale cluster services on the Exadata storage server.

To add a storage server to an existing cluster use the Oracle Exadata Deployment Assistant (OEDA) command line interface (OEDACLI):

Copyright © 2022, 2026, Oracle and/or its affiliates.

1. Load the OEDA XML system configuration file.

   ```
   oedacli> LOAD FILE NAME=es.xml
   ```

2. Clone an existing storage server to add a storage server to the Exascale cluster definition.

   For example:

   ```
   oedacli> CLONE NEWCELL SRCNAME=cell07.example.com
   TGTNAME=cell10.example.com
   oedacli> SET ADMINNET NAME=cell10.example.com IP=192.0.2.21
   NETMASK=255.255.248.0 GATEWAY=192.0.2.1
   oedacli> SET PRIVNET NAME1=cell10-priv1.example.com IP1=192.168.0.27
   NAME2=cell10-priv2.example.com IP2=192.168.0.28
   oedacli> SET ILOMNET NAME=cell10-ilom.example.com IP=198.51.100.110
   oedacli> SET RACK NUM=1, ULOC=8
   oedacli> SAVE ACTION FORCE
   ```

   In the example:

   - The name of the existing storage server being cloned is `cell07.example.com`.
   - The name of the new storage server being added is `cell10.example.com`.
   - The host names, network addresses, and rack details are all examples and you must substitute your own valid values.

3. Save the OEDA XML system configuration file.

   ```
   oedacli> SAVE FILE NAME=es.xml
   ```

4. Alter the cluster containing the storage servers and deploy the changes.

   For example:

   ```
   oedacli> ALTER CLUSTER ADDCELLS='cell10.example.com' WHERE CLUSTERNUMBER=1
   oedacli> SAVE ACTION
   oedacli> MERGE ACTIONS
   oedacli> DEPLOY ACTIONS
   ```

   In the example, the example cluster number is *1* and you must substitute your own valid value.

2.2.5 Manage Certificates for Exascale Control Services

Exascale control services, also known as Exascale RESTful Services (ERS), provide a management endpoint for Exascale management operations. Certificates enable secure communications for ERS processes over HTTPS.

ERS service instances are deployed using front-end and back-end server processes:

- The front-end ERS processes provide a highly available client endpoint with load-balancing capabilities. A corresponding front-end ERS certificate enables secure communications between ERS and external clients, such as the Exascale command line interface (ESCLI).
- The back-end ERS processes work with other software services to process requests and reply back to the client. To provide high-availability and share the workload, multiple back-end ERS instances are distributed across the Exadata storage servers.

Copyright © 2022, 2026, Oracle and/or its affiliates.

By default, ERS uses self-signed security certificates. Optionally, you can specify the certificates, either your own self-signed certificates or trusted certificates issued by a Certificate Authority (CA). Trusted certificates are generally recommended for front-end ERS servers because they are directly visible to external clients.

To specify the ERS certificates, you must use the ExaCLI utility. For information about configuring and using the ExaCLI utility, see Using the ExaCLI utility.

- The front-end ERS processes use a certificate that is bound to the virtual IP (VIP) address and hostname associated with the ERS endpoint. To specify a front-end ERS security certificate, you must use ExaCLI to modify the storage server `ESNODE` object and set the following attributes:
  - `securityPubKey` - Specifies a URL for the certificate public key file.
  - `securityPrivKey` - Specifies a URL for the certificate private key file.
  - `securityPrivKeyPW` - Specifies the password to use if the private key file is encrypted.

    If you specify `securityPrivKeyPW=*`, the user is prompted for the password.

  The keys must be supplied as PEM-encoded files. Each URL can use the `http`, `https`, or `file` access scheme.

  For example:

  ```
  $ exacli -l celladministrator -c cellhost  -
          -e 'ALTER ESNODE securityPubKey="file:///root/security/
  newkey.pem.crt", -
              securityPrivKey="file:///root/security/newkey-private.pem", -
              securityPrivKeyPW=*'

  Password: *********
  ...
  ```

  The same certificate must be deployed to all storage servers capable of hosting front-end ERS processes.
- The back-end ERS processes use the same certificates as the Exadata Management Server (MS). For details about how to specify a certificate for MS and back-end ERS processes, see Specifying a Security Certificate.

2.3 Administer Exascale Users

- Create a User
- List User Information
- Modify a User
- Remove a User
- Administer the Internal User Accounts

Related Topics

- About Exascale Users
- About Exascale User Credentials

Copyright © 2022, 2026, Oracle and/or its affiliates.

2.3.1 Create a User

By using different users, you can enforce data separation within Exascale.

Prior to performing the following procedure, the prospective new Exascale user should create their own public/private key pair and supply the public key. See Create User Keys. This is recommended to ensure the integrity of the private key, since the user should never share the private key, not even with the Exascale administrator.

To provision a new Exascale user:

1. Create the Exascale user.

   Use the ESCLI mkuser command and specify the name of the new Exascale user and a unique user ID.

   For example:

   ```
   @> mkuser my-user-name --attributes id=my-user-ID
   User created with ID: my-user-ID
   ```

   If you do not specify the user ID, then the user is assigned a system-generated unique user ID. For example:

   ```
   @> mkuser my-user-name
   User created with ID: 96a68014-5762-4579-86ee-29eb743decbd
   ```

   > **Note**
   >
   > The user ID is fixed at user creation time. There is no way to modify the user ID afterward.

   Take note of the user ID for the newly created user. You must supply the user ID value to the actual user so that they can configure their credential store (wallet).

2. Associate the new Exascale user with the user's supplied public key.

   Use the ESCLI chuser command and specify:

   - The ID for the user that is being modified in the `chuser` command.
   - The location of the file that contains the user's public key in PEM format.

   For example:

   ```
   @> chuser 96a68014-5762-4579-86ee-29eb743decbd --public-key-file1 pub.pem
   ```

2.3.2 List User Information

This topic describes how to display information about Exascale users.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- To display a complete list of Exascale user names and corresponding user IDs, use the ESCLI lsuser command with no options. For example:

  ```
  @> lsuser
  ```

- To display detailed information about a specific Exascale user, use the lsuser command with the `--detail` option and specify the ID for the user that is being queried. For example:

  ```
  @> lsuser my-user-ID --detail
  ```

  Using the command options available with the lsuser command, you can specify the user attributes you want to display and display the output in a long, tabular form.

2.3.3 Modify a User

You can modify attributes for an existing Exascale user by using the ESCLI chuser command. You can:

- Modify a User Name
- Modify User Privileges
- Manage a User's Public Keys

2.3.3.1 Modify a User Name

To change the name of a user, use the ESCLI chuser command with the `--name` option and specify:

- The unique user ID for the user that is being modified. You can use the lsuser command to find the ID for each user.
- The new name to apply to the user.

For example:

```
@> chuser my-user-ID --attributes name=my-user-name
```

2.3.3.2 Modify User Privileges

To change a user's privileges, use the ESCLI chuser command with the `--privilege` option and specify:

- The unique user ID for the user that is being modified. You can use the lsuser command to find the ID for each user.
- The list of privileges to apply to the user.

For example:

```
@> chuser my-user-ID --privilege vlt_read|egs
```

The `--privilege` option specifies a list of one or more privileges of the form *`privilege-1`*`|`*`privilege-2`*`|...`.

A privilege is one of the following: `no_privilege`, `cellsrv`, `egs`, `ers`, `syseds`, `usreds`, `bsm`, `bsw`, `ms`, `vlt_manage`, `vlt_use`, `vlt_read`, `vlt_inspect`, `cl_admin`, `cl_operator`, `cl_monitor`,

Copyright © 2022, 2026, Oracle and/or its affiliates.

`on_behalf_of`, `user_create`, `system_restore`, `rest_vault_client`, `rest_volume_client`. For descriptions, see Cluster Level Privileges.

`no_privilege` cannot be combined with any other privilege; otherwise, an error is returned.

Vault top-level privileges (`vlt_manage`, `vlt_use`, `vlt_read`, and `vlt_inspect`) are mutually exclusive. If any two or more are combined, an error is returned.

Cluster privileges (`cl_admin`, `cl_operator`, and `cl_monitor`) are mutually exclusive. If any two or more are combined, an error is returned.

2.3.3.3 Manage a User's Public Keys

This topic describes how to manage the public keys that are associated with an existing Exascale user.

- To associate a public key with a user, use the ESCLI chuser command with the `--public-key-file1`, `--public-key-file2`, or `--public-key-file3` option and specify:
  - The unique user ID for the user that is being modified. You can use the lsuser command to find the ID for each user.
  - The location of a file that contains the user's public key in PEM format.

  For example:

  ```
  @> chuser my-user-ID --public-key-file1 pub.pem
  ```

  By using the `--public-key-file1`, `--public-key-file2`, and `--public-key-file3` options, Exascale provides three public key 'slots' for each Exascale user.

  If you supply a public key for a previously unused slot, then the specified key is inserted into the slot. Otherwise, the existing key in the slot is overwritten with the specified public key.
- To remove a public key from a user, use the chuser command with the `--rm-public-key1`, `--rm-public-key2`, or `--rm-public-key3` option and specify the unique user ID for the user that is being modified.

  For example:

  ```
  @> chuser my-user-ID --rm-public-key1
  ```

  By using the `--rm-public-key1`, `--rm-public-key2`, and `--rm-public-key3` options, you can selectively remove the public key in any of the three public key 'slots' for the Exascale user.

2.3.4 Remove a User

To remove an existing Exascale user, use the ESCLI rmuser command and specify the ID for the user that is being removed. For example:

```
@> rmuser my-user-ID
```

You can use the lsuser command to find the ID for each user.

Copyright © 2022, 2026, Oracle and/or its affiliates.

2.3.5 Administer the Internal User Accounts

- Administer the ADMIN User
- Administer the Node Administration Users

2.3.5.1 Administer the ADMIN User

By default, each Exascale cluster contains one superuser account. The user identifier (ID) for the superuser account is `admin`. The `admin` user can implicitly perform any system operation and effectively holds all system privileges.

During system deployment, the `admin` user wallet is created on every storage server at `/opt/oracle/cell/cellsrv/deploy/config/security/admwallet` and every wallet contains the same system-generated private key.

If you choose to use the `admin` user for ongoing system administration, then you must manage access to the `admin` user wallet.

Alternatively, Oracle recommends the following approach:

1. Use the `admin` user to create your own dedicated Exascale administrator account or accounts.
2. Extract the `admin` user private key from the wallet and store it in a secure off-site key store. You can extract the private key from a wallet by using the ESCLI lswallet command.
3. Remove all copies of the `admin` user wallet.

By using this approach, you effectively disable the `admin` user and you must recreate the wallet if you require future `admin` access.

2.3.5.2 Administer the Node Administration Users

Exascale contains one node administration account for every node (storage server or compute node) that runs Exascale software services. Each node administration account inherits its user identifier (ID) from the server hostname and each account contains the privileges required to run the Exascale software services on the node. Do not directly use or modify these accounts.

During system deployment, the node administration user wallet is created on every node, and every wallet contains a system-generated private key. On a storage server, the wallet is located at `/opt/oracle/cell/cellsrv/deploy/config/eswallet`. On a compute node, the wallet is located at `/opt/oracle/dbserver/dbms/deploy/config/eswallet`.

By default, each node is configured with appropriate operating system privileges to limit wallet access.

Additionally, Oracle recommends that you extract the private key from each wallet and store it in a secure off-site key store. By doing this, you have a backup of the private key, which you can use if you need to rebuild the wallet. You can extract the private key from a wallet by using the ESCLI lswallet command.

2.4 Administer Exascale Pool Disks

- Create Pool Disks
- List Pool Disks

Copyright © 2022, 2026, Oracle and/or its affiliates.

- Modify the State of a Pool Disk
- Resize Pool Disks
- Drop Pool Disks

Related Topics

- About Pool Disks

2.4.1 Create Pool Disks

A pool disk is a portion of an Exadata storage device that is designated for use by Exascale. A pool disk is created by allocating space from an Exadata cell disk.

To create pool disks, use the CellCLI `CREATE GRIDDISK` command and specify the name of the storage pool that you want to associate with the newly created pool disks.

For example:

```
CellCLI> create griddisk all harddisk prefix=pool1, storagepool=my-pool1
```

You must specify the `storagepool` attribute to create Exascale pool disks. Otherwise, you will create Exadata grid disks, which Exascale cannot use.

When you specify the `storagepool` attribute, the Exadata grid disks are marked as Exascale pool disks, and the storage pool name is written into the pool disk metadata.

The `create griddisk` command does not immediately add the pool disks to the specified storage pool, and at this point, the storage pool does not need to exist. The pool disks are only available for data storage after they are included in the storage pool, either during storage pool creation or after the storage pool is modified by using the chstoragepool command with the `--reconfig` option.

Related Topics

- Create a Storage Pool
- Modify a Storage Pool

2.4.2 List Pool Disks

A pool disk is a portion of an Exadata storage device that is designated for use by Exascale.

To display information about pool disks, use the ESCLI lspooldisk command. For example:

```
@> lspooldisk
```

By default, without any optional arguments, the command displays basic information about all of the pool disks associated with the Exascale cluster.

Using the command options available with the ESCLI lspooldisk command, you can:

- Specify the pool disks that you want to display.
- Filter the output according to your specified conditions.
- Sort the output using specific attributes.
- Choose to display output with different levels of detail.

Copyright © 2022, 2026, Oracle and/or its affiliates.

2.4.3 Modify the State of a Pool Disk

This topic describes how to modify a pool disk and change its state; either online or offline.

A pool disk is a portion of an Exadata storage device that is designated for use by Exascale. You can modify a pool disk by setting its state to online or offline.

To modify the state of a pool disk, use the ESCLI chpooldisk command and specify:

- The identifier for the pool disk that you want to modify.

  You can use the ESCLI lspooldisk command to display the pool disk identifier for any Exascale pool disk.
- The desired pool disk state. Either `--offline` or `--online`.

For example:

```
@> chpooldisk 15 --offline
```

2.4.4 Resize Pool Disks

A pool disk is a portion of an Exadata storage device that is designated for use by Exascale.

When you resize pool disks that are allocated to a storage pool, you indirectly change the size of the storage pool. Depending on the situation, the change may trigger a rebalance operation within the storage pool.

To resize pool disks:

1. Modify the size attribute for the pool disks.

   Use the CellCLI `ALTER GRIDDISK` command and specify the pool disks that you want to resize along with a setting for the size attribute. For example:

   ```
   CellCLI> alter griddisk size=2000G where storagepool=my-pool1
   ```

   You can increase or reduce the size of existing pool disks.

   Increasing the size of a pool disk is an almost instantaneous operation. But, you cannot increase the size of a pool disk beyond the available free space in the underlying cell disk.

   When reducing the size of a pool disk, existing allocations beyond the new pool disk boundary move into available locations inside the pool disk. The time required to complete the operation depends on the amount of data being moved. The operation fails if you attempt to reduce the size to the point where the pool disk cannot accommodate the existing data.
2. Reconfigure the storage pool to accommodate the resized pool disks.

   Use the ESCLI chstoragepool command. For example:

   ```
   @> chstoragepool my-pool1 --reconfig
   ```

   If required, reconfiguring the storage pool triggers a rebalance operation.

Copyright © 2022, 2026, Oracle and/or its affiliates.

2.4.5 Drop Pool Disks

A pool disk is a portion of an Exadata storage device that is designated for use by Exascale.

Exascale only supports dropping of pool disks for the purpose of shrinking a storage pool by removing all of the pool disks provided by a storage server. You cannot drop individual pool disks.

To drop the pool disks associated with a storage server:

1. Prepare the pool disks for removal.

   Use the CellCLI `DROP GRIDDISK` command with the `prepare` option. For example:

   ```
   CellCLI> drop griddisk all harddisk prefix=pool1 prepare
   ```

   This command marks the pool disks for removal and sets the pool disk status to `waiting for storage pool reconfig`.

   > **Note**
   >
   > If required, you can undo this step before you reconfigure the storage pool by using the CellCLI `DROP GRIDDISK` command with the `cancel` option. For example:
   >
   > ```
   > CellCLI> drop griddisk all harddisk prefix=pool1 cancel
   > ```

2. Reconfigure the storage pool.

   Use the ESCLI chstoragepool command with the `--reconfig` option. For example:

   ```
   @> chstoragepool my-pool1 --reconfig
   ```

   This command permanently removes the pool disks and rebalances the storage pool.

2.5 Administer Exascale Storage Pools

- Create a Storage Pool
- List Storage Pools
- Maintain Free Space in a Storage Pool to Protect Against Disk Failure
- Modify a Storage Pool
- Drop a Storage Pool

Related Topics

- About Storage Pools

2.5.1 Create a Storage Pool

This topic describes how to create a new Exascale storage pool.

Copyright © 2022, 2026, Oracle and/or its affiliates.

An Exascale storage pool is a collection of pool disks that provides persistent physical storage for Exascale vaults and files. Each storage pool is a collection of pool disks using only one type of storage media (for example, HC). You cannot define a storage pool with pool disks having a mixture of media types (for example, HC and EF).

Before you create a storage pool, you need some pool disks that are already associated with the storage pool name.

To create a storage pool, use the ESCLI mkstoragepool command and specify the name for the new storage pool.

The essential command syntax is:

```
@> mkstoragepool storage-pool-name
```

Additionally, you can specify the `--nowait` option to instruct the command to return immediately instead of waiting until the storage pool is created.

Related Topics

- Create Pool Disks

2.5.2 List Storage Pools

An Exascale storage pool is a collection of pool disks that provides the persistent physical storage for Exascale vaults and files.

To display information about storage pools, use the ESCLI lsstoragepool command. For example:

```
@> lsstoragepool
```

By default, without any optional arguments, the command displays basic information about all of the storage pools associated with the Exascale cluster.

Using the command options available with the lsstoragepool command, you can:

- Specify the storage pools that you want to display.
- Filter the output according to your specified conditions.
- Sort the output using specific attributes.
- Choose to display output with different levels of detail.

2.5.3 Maintain Free Space in a Storage Pool to Protect Against Disk Failure

An Exascale storage pool is a collection of pool disks that provides the persistent physical storage for Exascale vaults and files. If the storage device hosting one pool disk fails, then the storage pool can continue to function with reduced redundancy. However, the storage pool may go offline if multiple pool disks fail at the same time.

To maintain maximum protection against storage failures, you should maintain sufficient free space in each storage pool to facilitate an automatic data rebalance operation after the failure of a main storage device.

The amount of free space in a storage pool is the difference between the following storage pool attributes:

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `spaceRaw`: Specifies the total amount of storage space associated with the storage pool.
- `spaceUsed`: Specifies the total amount of storage space allocated (used) in the storage pool.

You can use the ESCLI lsstoragepool command to examine these attributes. For example, the following command shows the `spaceRaw` and `spaceUsed` values associated with the storage pool named `my-pool1`:

```
@> lsstoragepool my-pool1 --attributes spaceRaw,spaceUsed
```

The amount of free space required to rebalance a storage pool after losing a storage device largely depends on the number and type of storage servers associated with the storage pool. Another consideration is the redundancy level used to protect against data loss. However, the following recommendations assume that the storage pool only contains high redundancy (triple mirrored) files.

- When the storage pool occupies only three storage servers, high redundancy (triple mirrored) data must be reconstructed on the same server. In this case, the required free space to reconstruct data after losing a storage device is linked to the number of physical storage devices in the storage server.

  The following formula describes the free space requirement to reconstruct data after losing one main storage device for a storage pool using only three storage servers. The free space requirement is expressed as a percentage of the total amount of storage space associated with the storage pool. In the formula, *DeviceCount* represents the number of main storage devices contained in each storage server.

  ```
  Free Space Percentage = ceiling(100/DeviceCount)+3
  ```

  For example, on a high capacity (HC) storage server with 12 disks, one disk represents 8.33% (`100/12`) of the total storage capacity. So, for a storage pool using three 12-disk HC servers, you should maintain 12% (`ceiling(8.33)+3`) of the storage pool as free space to reconstruct data after losing a storage device.

  Based on the preceding formula, the following table summarizes the amount of free space that should be maintained to successfully rebalance a storage pool spanning only three storage servers after losing one main storage device.

Storage Server Type	Free Space Requirement
High capacity (HC) or extended (XT) storage server with 12 disks	12% of the total space in the storage pool.
HC storage server with 6 disks	20% of the total space in the storage pool.
Extreme flash (EF) storage server with 8 performance-optimized flash devices	16% of the total space in the storage pool
EF storage server with 4 capacity-optimized flash devices	28% of the total space in the storage pool

- When a storage pool occupies four or more storage servers, the lost data can be reconstructed on any partner storage server. In this case, the amount of free space that should be maintained to successfully rebalance a storage pool after losing one main storage device is whichever is greater out of the following:
  - 3% of the total space in the storage pool.
  - Three times the size of the lost pool disk(s).

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> On an EF storage server with 4 capacity-optimized flash devices, each device contains two pool disks.

2.5.4 Modify a Storage Pool

An Exascale storage pool is a collection of pool disks that provides the persistent physical storage for Exascale vaults and files.

Fundamentally, you can modify a storage pool in two ways:

- You can modify the storage configuration associated with the storage pool.
- You can modify other attributes that are associated with the storage pool.

To modify a storage pool, use the ESCLI chstoragepool command:

- To modify the storage configuration associated with the storage pool, use the chstoragepool command in conjunction with the following options:
  - `--reconfig`: Reconfigures the storage pool and makes permanent any changes to the number or size of the underlying pool disks. Reconfiguration may trigger a rebalance operation to ensure that data is spread evenly across the pool disks in the storage pool. The time required for reconfiguration depends on the amount of data movement required to rebalance the storage pool.
  - `--force`: Optionally forces a reconfiguration operation even if the system detects no apparent change to the underlying pool disks.

  For example:

  ```
  @> chstoragepool storage-pool-name --reconfig --force
  ```

  Implicitly, you also modify the storage configuration associated with a storage pool when you perform any of the following tasks:

  - Add a Storage Server.
  - Create Pool Disks and allocate them to an existing storage pool.
  - Resize Pool Disks.
  - Drop Pool Disks.
- To modify the attributes that are associated with the storage pool, use the chstoragepool command, and specify the attribute that you want to modify. For example:

  ```
  @> chstoragepool storage-pool-name --attributes diskOfflineTimerInMins=60
  ```

  Use the `describe chstoragepool` command to view details about the storage pool attributes you can modify with `chstoragepool`.

2.5.5 Drop a Storage Pool

An Exascale storage pool is a collection of pool disks that provides the persistent physical storage for Exascale vaults and files.

Copyright © 2022, 2026, Oracle and/or its affiliates.

To remove a storage pool, use the ESCLI rmstoragepool command and specify the storage pool that you want to remove. For example:

```
@> rmstoragepool storage-pool-name
```

2.6 Administer the Storage Devices in an Exascale Storage Pool

Use the following procedures to administer the storage devices in an Exascale storage pool. If your system is configured to use Exascale and Oracle ASM (also known as a hybrid storage configuration), use these procedures in conjunction with the directions at Maintaining Oracle Exadata Storage Servers (in the *Oracle Exadata Maintenance Guide*).

- Replacing a Failed Storage Device
  Failure of a storage device can affect performance and data redundancy. Consequently, a failed storage device should be replaced as soon as possible.
- Proactively Replacing a Storage Device
  You may need to proactively replace a storage device that has not failed.

Related Topics

- About Pool Disks
- About Storage Pools

2.6.1 Replacing a Failed Storage Device

Failure of a storage device can affect performance and data redundancy. Consequently, a failed storage device should be replaced as soon as possible.

A storage device is considered to have failed in the following circumstances:

- A hardware or firmware fault causes the device to stop functioning.
- The device enters a predictive failure state.

  In this case, the device is still usable, but there is an indication that the device may soon stop functioning. For example, a hard disk drive (HDD) could be short of spare sectors, or a flash device could be approaching wear limits.
- Exadata software confines the device, and the device fails the post-confinement checks.

  Exadata automatically confines a storage device after detecting a significant performance problem or functional anomaly. After confinement, Exadata attempts to resolve the issue and recheck the device. However, if the post-confinement checks fail, the device is considered to have failed.

If a storage device fails, Exadata automatically drops all of the grid disks contained on the storage device. If any grid disk is used as an Exascale pool disk, Exascale automatically removes it and rebalances the storage pool to restore data redundancy.

Exadata also generates an alert when a storage device fails. The alert message includes specific instructions for replacing the device. If alert notifications are configured on the storage server, then the alert notification is automatically sent using email and SNMP.

After the failed storage device is replaced, Exadata automatically creates the cell and grid disks on the new device. If any grid disk is configured as an Exascale pool disk, it is automatically added to the storage pool, and the storage pool is rebalanced.

The following steps outline the procedure for replacing a failed storage device:

Copyright © 2022, 2026, Oracle and/or its affiliates.

1. Confirm the location of the failed device.

   Use the following CellCLI `LIST PHYSICALDISK` command:

   ```
   CellCLI> list physicaldisk where status!=normal detail
   ```

   In the output, the `slotNumber` value describes the physical location of the device.

   You can also classify the failure type by examining the `status` value:

   - `failed` or `failed - dropped for replacement` - Indicates that the device stopped functioning due to a hardware or firmware failure.
   - `warning - predictive failure` - Indicates that the device entered a predictive failure state.
   - `warning - poor performance` - Indicates that Exadata software confined the device, and the device failed the post-confinement checks.

   For example, the following output indicates that the hard disk drive (HDD) in slot 5 stopped functioning due to a hardware or firmware failure:

   ```
   CellCLI> list physicaldisk where status!=normal detail
            name:                   0:5
            deviceName:             /dev/sdi
            diskType:               HardDisk
            enclosureDeviceId:      0
            luns:                   0_5
            makeModel:              "WDC W7222A520ORA022T"
            physicalFirmware:       A7B0
            physicalInsertTime:     2023-07-07T17:20:44-07:00
            physicalInterface:      sas
            physicalSerial:         70SP8E
            physicalSize:           20.009765625T
            slotNumber:             5
            status:                 failed
   ```

2. Ensure that the storage server Do Not Service LED is not lit.
3. Ensure that the failed device is ready for removal.
   - If the failed device is a flash drive or HDD located in one of the hot-swappable drive bays in the front of the server, ensure that the blue OK to Remove LED on the device is lit before removing the device.
   - If the failed device is a hot-swappable flash card contained inside the server, ensure that the power LED on the flash card is not lit before removing the device. Starting with Exadata Storage Server X7-2, all storage server models contain hot-swappable flash cards.
4. Remove the failed storage device and install the replacement.

   See the associated server hardware guide for additional details about physical hardware replacement.
5. Wait for the server to recognize the replaced device.

   When you physically replace a hot-swappable storage device, it may take a few minutes for the server to recognize the new device.
6. Confirm the status of the replacement device.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Use the CellCLI `LIST PHYSICALDISK` command to confirm that the status of the replacement device is `normal`.

For example:

```
CellCLI> list physicaldisk 0:5 detail
         name:                   0:5
         deviceName:             /dev/sdi
         diskType:               HardDisk
         enclosureDeviceId:      0
         luns:                   0_5
         makeModel:              "WDC W7222A520ORA022T"
         physicalFirmware:       A7B0
         physicalInsertTime:     2023-09-01T12:00:25-07:00
         physicalInterface:      sas
         physicalSerial:         75X8RD
         physicalSize:           20.009765625T
         slotNumber:             5
         status:                 normal
```

7. Monitor the storage pool rebalance operation.

   As part of reintegrating any Exascale pool disk on the replacement storage device, the affected storage pool undergoes a rebalance operation.

   Use the ESCLI lsstoragepooloperation command to monitor the storage pool rebalance operation.

8. Confirm that Exascale is using the replacement device.

   Use the ESCLI lspooldisk command and examine the `status` attribute.

   Initially, as the replacement device comes online, the pool disk status is briefly set to `BEING ADDED`. However, the status value transitions to `ONLINE` as Exascale reintegrates the replacement device.

2.6.2 Proactively Replacing a Storage Device

You may need to proactively replace a storage device that has not failed.

For example, you may choose to replace a device after Exadata detects a temporary performance anomaly (temporary confinement), or you may need to replace a functioning device for any other reason.

The following steps outline the procedure to proactively replace a storage device:

1. Prepare the storage device for replacement.

   Use the CellCLI `ALTER PHYSICALDISK` command to prepare the system for replacement of the specified storage device. You can choose whether or not to maintain redundancy throughout the device replacement process.

   - If you maintain redundancy throughout the device replacement process, Exascale moves data off any affected pool disk and rebalances the associated storage pool. This option is generally recommended and provides the best protection from any other failure that could occur while replacing the device. However, it comes with the cost of moving the affected data and rebalancing the storage pool.

     To choose this option, add the `MAINTAIN REDUNDANCY` clause to the `ALTER PHYSICALDISK ... DROP FOR REPLACEMENT` command. For example, you can use the

Copyright © 2022, 2026, Oracle and/or its affiliates.

following command to prepare storage device `0:5` for replacement while maintaining redundancy:

```
CellCLI> alter physicaldisk 0:5 drop for replacement maintain redundancy
```

- If you choose not to maintain redundancy throughout the device replacement process, Exascale only performs a quick redundancy check before the device is taken offline. The redundancy check ensures that another copy of data is available for any affected pool disk. This option allows for quicker device replacement using reduced redundancy. However, it comes with an increased risk of data loss or a storage pool outage if another failure occurs while replacing the device. Consequently, this option is only recommended where data loss can be tolerated (for example, a non-critical test system) or where the increased risk associated with reduced redundancy is mitigated by some other means (for example, the data is replicated using Oracle Data Guard).

  To choose this option, use the `ALTER PHYSICALDISK ... DROP FOR REPLACEMENT` command without the `MAINTAIN REDUNDANCY` clause. For example, you can use the following command to prepare storage device `0:5` for replacement without maintaining redundancy:

  ```
  CellCLI> alter physicaldisk 0:5 drop for replacement
  ```

2. Ensure that the storage server Do Not Service LED is not lit.
3. Ensure that the device is ready for removal.
   - If the failed device is a hard disk drive (HDD) or flash drive located in one of the hot-swappable drive bays in the front of the server, ensure that the blue OK to Remove LED on the device is lit before removing the device.
   - If the failed device is a hot-swappable flash card contained inside the server, ensure that the power LED on the flash card is not lit before removing the device. Starting with Exadata Storage Server X7-2, all storage server models contain hot-swappable flash cards.
4. Remove the failed storage device and install the replacement.

   See the associated server hardware guide for additional details about physical hardware replacement.
5. Wait for the server to recognize the replaced device.

   When you physically replace a hot-swappable storage device, it may take a few minutes for the server to recognize the new device.
6. Confirm the status of the replacement device.

   Use the CellCLI `LIST PHYSICALDISK` command to confirm that the status of the replacement device is `normal`.

   For example:

   ```
   CellCLI> list physicaldisk 0:5 detail
            name:                   0:5
            deviceName:             /dev/sdi
            diskType:               HardDisk
            enclosureDeviceId:      0
            luns:                   0_5
            makeModel:              "WDC W7222A520ORA022T"
            physicalFirmware:       A7B0
            physicalInsertTime:     2023-09-01T12:00:25-07:00
   ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
physicalInterface:      sas
physicalSerial:         75X8RD
physicalSize:           20.009765625T
slotNumber:             5
status:                 normal
```

7. Monitor the storage pool rebalance operation.

   After the storage device is replaced, the Exadata cell disks and grid disks that existed on the original device are re-created on the new device. If any grid disk is designated as an Exascale pool disk, it is added to the storage pool and the storage pool is automatically rebalanced.

   Use the ESCLI lsstoragepooloperation command to monitor the storage pool rebalance operation.

8. Confirm that Exascale is using the replacement device.

   Use the ESCLI lspooldisk command and examine the `status` attribute.

   Initially, as the replacement device comes online, the pool disk status is briefly set to `BEING ADDED`. However, the status value transitions to `ONLINE` as Exascale reintegrates the replacement device.

2.7 Administer Cluster Templates

- Create a Cluster Template
- List Cluster Templates
- Modify a Cluster Template
- Remove a Cluster Template

Related Topics

- About Exascale Templates

2.7.1 Create a Cluster Template

This topic describes how to create a user-defined cluster template.

A cluster template is a named collection of file storage attribute settings defined at the cluster level. By using cluster templates, you can manage file storage attributes automatically and consistently across the Exascale cluster.

Exascale is automatically preconfigured with cluster templates for each system-defined file type; for example, `DATAFILE`, `ONLINELOG`, `CONTROLFILE`, and so on. Therefore, you can only create new user templates at the cluster level.

To create a user-defined cluster template, use the ESCLI mktemplate command with the `--cluster` and `--name` options and specify the following file storage attribute settings:

- `--media-type`: Specifies the physical media type used to store files created with the template. Possible values are:
  - `HC`: Identifies high capacity storage, using hard disk drives (HDDs) on high-performance Exadata storage servers.
  - `EF`: Identifies extreme flash storage, using low-latency, high-throughput flash devices.

Copyright © 2022, 2026, Oracle and/or its affiliates.

  - `XT`: Identifies extended storage, using HDDs on Exadata storage servers optimized for cost-effective online storage of less frequently accessed data.
- `--redundancy`: Specifies the number of data copies that are maintained for files created with the template. Currently, the only permitted value is:
  - `high`: Indicates three mirrored copies of the file data.
- `--content-type`: Specifies the type of content in files created with the template. Exascale internally uses this attribute to place file extents on physically separate devices in a manner that maximizes availability in the event of a failure. Possible values are:
  - `DATA`: Principally associated with user data.
  - `RECO`: Primarily for data used in backup and recovery operations.

The command syntax is:

```
@> mktemplate --cluster --name template-name --content-type content-type --
media-type media-type --redundancy redundancy
```

2.7.2 List Cluster Templates

A cluster template is a named collection of file storage attribute settings defined at the cluster level. By using cluster templates, you can manage file storage attributes automatically and consistently across the Exascale cluster.

To display information about Exascale cluster templates, use the ESCLI lstemplate command with the `--cluster` option. For example:

```
@> lstemplate --cluster
```

This simple form of the command displays basic information about all of the cluster templates.

Using the command options available with the lstemplate command, you can specify the templates you want to display and display output with different levels of detail.

2.7.3 Modify a Cluster Template

A cluster template is a named collection of file storage attribute settings defined at the cluster level. By using cluster templates, you can manage file storage attributes automatically and consistently across the Exascale cluster.

- To modify an existing cluster template that is associated with a file type, use the ESCLI chtemplate command with the `--cluster` and `--file-type` options and specify the file storage attribute settings:
  - `--media-type`: Specifies the physical media type used to store files created with the template. Possible values are:
    - `HC`: Identifies high capacity storage, using hard disk drives (HDDs) on high-performance Exadata storage servers.
    - `EF`: Identifies extreme flash storage, using low-latency, high-throughput flash devices.
    - `XT`: Identifies extended storage, using HDDs on Exadata storage servers optimized for cost-effective online storage of less frequently accessed data.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `--redundancy`: Specifies the number of data copies that are maintained for files created with the template. Currently, the only permitted value is:
  - `high`: Indicates three mirrored copies of the file data.
- `--content-type`: Specifies the type of content in files created with the template. Exascale internally uses this attribute to place file extents on physically separate devices in a manner that maximizes availability in the event of a failure. Possible values are:
  - `DATA`: Principally associated with user data.
  - `RECO`: Primarily for data used in backup and recovery operations.

The command syntax is:

```
@> chtemplate --cluster --file-type file-type --content-type content-type
--media-type media-type --redundancy redundancy
```

- To modify a user-defined cluster template, use the chtemplate command with the `--cluster` and `--name` options and specify the file storage attribute settings. The command syntax is:

```
@> chtemplate --cluster --name template-name --content-type content-type --
media-type media-type --redundancy redundancy
```

2.7.4 Remove a Cluster Template

A cluster template is a named collection of file storage attribute settings defined at the cluster level. By using cluster templates, you can manage file storage attributes automatically and consistently across the Exascale cluster.

You can only remove user templates at the cluster level. You cannot remove a cluster template that is associated with a file type.

To remove a user-defined cluster template, use the ESCLI rmtemplate command with the `--cluster` and `--name` options. The command syntax is:

```
@> rmtemplate --cluster --name template-name
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

3

Oracle Exadata Exascale User-Specific Administration

This chapter covers administration tasks that define and manage Exascale storage objects and metadata that belong to a specific Exascale user.

Administration activities for Oracle Exadata Exascale are logically divided into two groups. System administration activities define and manage the Exascale system components that are shared across the system. User-specific administration activities define and manage Exascale storage objects and metadata that belong to a specific Exascale user.

This chapter contains the following user-specific administration topics:

- Administer Exascale User Credentials
- Administer Exascale Vaults
- Administer Exascale Files
- Administer Exascale Datasets
- Administer Access Control Lists
- Administer Extended Attributes
- Administer Vault Templates
- Administer Resource Profiles
- Administer Exascale Block Store Data Objects

3.1 Administer Exascale User Credentials

- Create User Keys
- Create a Wallet
- Store a Private Key in a Wallet
- Fetch the Trust Store
- Store the URL for Exascale Cluster Services
- Store the Exascale Control Services Endpoint
- Distribute a Wallet
- Change User Keys

Related Topics

- About Exascale User Credentials

3.1.1 Create User Keys

This topic describes how to create a public/private key pair for an Exascale user.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Exascale authentication uses a system of public and private key pairs. Each Exascale user is associated with a public key. To prove their identity, an Exascale user must supply the matching private key.

For maximum security, an Exascale user should create their own public/private key pair. This is recommended to ensure the integrity of the private key, since the user should never share the private key, not even with the Exascale administrator.

For use with Exascale, the public and private keys must be in PEM format.

To create a key pair, you can use standard utilities, such as `openssl`, or you can use the ESCLI mkkey command. For example:

```
@> mkkey --private-key-file priv.pem --public-key-file pub.pem
```

After you create a key pair, you will need to send the public key to the Exascale administrator so that they can associate it with your Exascale user account.

Never share the private key, not even with the Exascale administrator.

3.1.2 Create a Wallet

Exascale user credentials are contained in a digital key store, also known as a wallet.

For maximum security, an Exascale user should create and manage their own wallet.

Before you can associate a wallet with an Exascale user, you must know the user identifier for the Exascale user. Consult with your Exascale administrator to determine the required value.

To create a wallet for an Exascale user:

1. Create an empty wallet.

   Use the ESCLI mkwallet command and specify the directory location for the new wallet. For example:

   ```
   @> mkwallet --wallet /home/user/user.wallet
   ```

2. Associate the wallet with the Exascale user.

   Use the ESCLI chwallet command and specify:

   - The user ID for the user that is being associated with the wallet.
   - The wallet directory location.

   For example:

   ```
   @> chwallet --wallet /home/user/user.wallet --attributes user=my-user-ID
   ```

3.1.3 Store a Private Key in a Wallet

Exascale authentication uses a system of public and private key pairs. Each Exascale user is associated with a public key. To prove their identity, an Exascale user must supply the matching private key, which is stored inside a digital wallet.

Before you can store a private key in your wallet, you must already have the key and the wallet.

Copyright © 2022, 2026, Oracle and/or its affiliates.

To store a private key in a wallet, use the ESCLI chwallet command and specify:

- The wallet location.
- The location of the file that contains the user's private key in PEM format.

For example:

```
@> chwallet --wallet /home/user/user.wallet --private-key-file priv.pem
```

If the wallet did not contain a private key before the chwallet command, then the specified private key is inserted into the wallet. Otherwise, the existing private key is overwritten with the specified private key.

3.1.4 Fetch the Trust Store

The Exascale trust store is a group of digital certificates that facilitates trusted communication between servers in the Exascale cluster. To use Exascale, users need to copy the trust store into their wallet.

To copy the trust store into your wallet, use the ESCLI chwallet command and specify the wallet location. For example:

```
@> chwallet --wallet /home/user/user.wallet --fetch-trust-store
```

3.1.5 Store the URL for Exascale Cluster Services

If you want to use a wallet to access data files on Exascale storage directly, the wallet must contain the URL for Exascale cluster services (EGS).

This requirement does not apply if the wallet does not directly access data files on Exascale storage. For example, a wallet need not contain the URL for Exascale cluster services (EGS) if it is used only for Exascale administration operations through the ESCLI utility.

You must already have a wallet before you can store the URL for Exascale cluster services (EGS) in it.

The URL for Exascale cluster services (EGS) is contained in the `exaRootUrl` attribute associated with the Exascale cluster. You can also get the required value by copying the `exaRootUrl` value from another previously configured wallet.

To store the URL for Exascale cluster services (EGS) in your wallet, use the ESCLI chwallet command and specify:

- The wallet location.
- The URL for Exascale cluster services (EGS).

For example:

```
@> chwallet --wallet /home/user/user.wallet --attributes
exaRootUrl="egs=egsexc4:192.0.2.217:5045 egs=egsexc4:192.0.2.218:5045
egs=egsexc4:192.0.2.219:5045"
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

3.1.6 Store the Exascale Control Services Endpoint

As a matter of convenience, you can store the default endpoint for Exascale control services in your wallet. Unless another value is explicitly provided in the command line, the ESCLI utility automatically uses the Exascale control services endpoint that is contained in your wallet.

You must already have a wallet before you can store the Exascale control services endpoint in it.

To store the Exascale control services endpoint in your wallet, use the ESCLI chwallet command and specify:

- The wallet location.
- The server and port that host the Exascale control services endpoint.

For example:

```
@> chwallet --wallet /home/user/user.wallet --attributes
restEndPoint=exa01:61926
```

3.1.7 Distribute a Wallet

You may need to move or copy your Exascale user wallet to connect using various different Exascale clients.

- To move or copy an Exascale user wallet within the same platform (for example, from one Exadata Linux server to another), you can use an archiving utility, such as `zip`. When you do this, you must ensure that you copy the entire contents of the wallet location.

  For example:

  ```
  $ zip -r wallet.zip /home/user/user.wallet
  ```

  In the previous example, the wallet located at `/home/user/user.wallet` is zipped into an archive named `wallet.zip`. You can then distribute the archive (`wallet.zip`) to the required client and unzip it into the required location.

  For example:

  ```
  $ unzip wallet.zip -d /path/to/wallet
  ```

- To copy an Exascale user wallet across different platforms (for example, from an Exadata Linux server to a non-Linux Exascale client system), you must create a new wallet on the desired host and populate it with the required user credentials. To move an Exascale user wallet across platforms, you must create a copy of the wallet and then delete the original.

  If you have the required user credentials, then you can use the procedures described in:

  – Create a Wallet
  – Store a Private Key in a Wallet
  – Fetch the Trust Store
  – Store the Exascale Control Services Endpoint

  If you do not have the required user credentials, then you can get them from the existing wallet by using the ESCLI lswallet command.

Copyright © 2022, 2026, Oracle and/or its affiliates.

For example, the following command displays the Exascale user identifier and the Exascale control services endpoint that are stored in the wallet at `/home/user/user.wallet`:

```
@> lswallet --wallet /home/user/user.wallet --attributes user,restEndPoint
```

In the following example, the private key is extracted from the wallet and stored in a regular PEM file located at `/home/user/priv.pem`:

```
@> lswallet --wallet /home/user/user.wallet --private-key-file /home/user/
priv.pem
```

Regardless of the method used to duplicate the wallet, to use an Exascale user wallet in conjunction with Oracle Database or Oracle Grid Infrastructure, you must copy the wallet into the correct location on each database cluster node.

Related Topics

- About Exascale User Credentials for Oracle Grid Infrastructure and Oracle Database

3.1.8 Change User Keys

This topic describes how to change a public/private key pair for an Exascale user.

Exascale authentication uses a system of public and private key pairs. Each Exascale user is associated with a public key. To prove their identity, an Exascale user must supply the matching private key, which is stored inside a digital wallet.

If your private key is ever compromised, you will need to generate a new key pair and change to using it. It is also good practice to periodically change, or rotate, your keys.

To change a public/private key pair:

1. Generate a new key pair:

   To create a key pair, you can use standard utilities, such as `openssl`, or you can use the ESCLI mkkey command. For example:

   ```
   @> mkkey --private-key-file new-priv.pem --public-key-file new-pub.pem
   ```

2. Associate the new public key with the Exascale user:

   Use the ESCLI chuser command and specify:

   - The unique user ID for the user that is being modified in the `chuser` command. By default, an Exascale user can change their own public key.
   - The location of the file that contains the user's new public key in PEM format.

   For example:

   ```
   @> chuser my-user-ID --public-key-file1 new-pub.pem
   ```

   If you want to ensure that a previously associated public key is no longer usable, then use the public key slot (`--public-key-file1`, `--public-key-file2`, or `--public-key-file3`) that contains the key that you want to overwrite.

3. Store the new private key in the user's wallet:

Copyright © 2022, 2026, Oracle and/or its affiliates.

Repeat this step for every copy of the user's wallet, which may reside on different Exascale servers.

Use the ESCLI chwallet command and specify:

- The wallet location.
- The location of the file that contains the user's new private key in PEM format.

For example:

```
@> chwallet --wallet /home/user/user.wallet --private-key-file new-priv.pem
```

When prompted, specify that you want to overwrite the existing private key in the wallet.

3.2 Administer Exascale Vaults

- Create a Vault
- Provision a Vault
- List Vaults
- Remove a Vault

Related Topics

- About Exascale Vaults

3.2.1 Create a Vault

An Exascale vault is a logical file container.

To create a vault, use the ESCLI mkvault command. For example:

```
@> mkvault @my-vault-name
```

As shown in the example, the simplest form of the command specifies only the name of the new vault. In this case, the resulting vault has unlimited access to all of the resources in the Exascale cluster.

Alternatively, you can provision a vault with specific resources. For details, see Provision a Vault.

3.2.2 Provision a Vault

An Exascale vault is a logical file container. You can provision a vault with specific storage resources.

Starting with Oracle Exadata System Software release 26.1.0, you can create multiple storage pools for each media type within a single Exascale cluster. Previously, only one storage pool was permitted for each media type in an Exascale cluster.

Consequently, there are two approaches for provisioning a vault with specific storage resources:

- Starting with Oracle Exadata System Software release 26.1.0, you can provision a vault with a specific set of storage resources from each storage pool.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- If your Exascale cluster is configured with Oracle Exadata System Software before release 26.1.0, you can specify one set of resource provisioning attributes, covering all media types, that is linked directly to the vault.

  This approach remains in later Oracle Exadata System Software releases for backward compatibility. You can transition to using multiple storage pools for each media type by deleting the original resource provisioning definitions linked directly to the vault, and then creating new definitions that provision a specific set of storage resources from each storage pool associated with the vault.

Choose the approach that suits your system configuration:

- Provision a Vault Using Multiple Storage Pools For Each Media Type
- Provision a Vault Using Only One Storage Pool For Each Media Type

3.2.2.1 Provision a Vault Using Multiple Storage Pools For Each Media Type

Starting with Oracle Exadata System Software release 26.1.0, you can provision a vault with specific storage resources from multiple storage pools.

To provision the required storage resources, you must create a vault and then separately provision the resources from the desired storage pools using the ESCLI chvault command with the `--provision` option. The general command syntax is:

```
@> chvault [@]vault --provision storagePool=storagepool-name[,provisioning-attribute=value]...
```

In the command, you must specify:

- The vault name (*vault*).
- The name of the storage pool (`storagePool=`*`storagepool-name`*) that you want to provision resources from.

Optionally, you can also specify any combination of the following provisioning attributes:

- `spaceProv`: Provisions the vault with the specified amount of storage space.
- `iopsProv`: Provisions the vault with the specified number of IOPS.
- `flashCacheProv`: Provisions the vault with the specified amount of flash cache space.
- `xrmemCacheProv`: Provisions the vault with the specified amount of Exadata RDMA Memory Cache (XRMEM cache) space.
- `flashLogProv`: Boolean value (`true` or `false`) indicating whether the vault is provisioned with access to Exadata Smart Flash Log.
- `xrmemLogProv`: Boolean value (`true` or `false`) indicating whether the vault is provisioned with access to XRMEM log (if available).

You can use the same general command syntax to provision a vault with storage resources, modify existing provisioning arrangements, and remove (deprovision) storage resources from a vault.

- To provision a vault with resources from multiple storage pools, you must use a separate chvault command for each storage pool.

Copyright © 2022, 2026, Oracle and/or its affiliates.

For example, the following commands create a new vault (`@my-data`) and provision it with 10 TB of storage space from the storage pool named `my-pool1` and 20 TB of space from `my-pool2`.

```
@> mkvault @my-data
@> chvault @my-data --provision storagePool=my-pool1,spaceProv=10T
@> chvault @my-data --provision storagePool=my-pool2,spaceProv=20T
```

- To modify existing provisioning arrangements, you can use the chvault command to set or modify specific provisioning attributes.

  For example, the following command builds on the previous example by provisioning the specified vault (`@my-data`) with 10 GB of flash cache space from `my-pool1`. The command also changes the amount of provisioned storage space to 30 TB.

  ```
  @> chvault @my-data --provision storagePool=my-
  pool1,flashCacheProv=10G,spaceProv=30T
  ```

- To remove provisioned storage pool resources from a vault, you must specify zero values for both `spaceProv` and `iopsProv`.

  For example, the following command removes all storage resources provided by `my-pool2` from the specified vault (`@my-data`).

  ```
  @> chvault @my-data --provision storagePool=my-pool2,spaceProv=0,iopsProv=0
  ```

  Removing (deprovisioning) storage pool resources is only permitted if the vault contains no data files in the specified storage pool.

For further details, see the chvault command reference.

3.2.2.2 Provision a Vault Using Only One Storage Pool For Each Media Type

If your Exascale cluster is configured with Oracle Exadata System Software before release 26.1.0, you can specify one set of resource provisioning attributes, covering all media types, that is linked directly to the vault.

This approach remains in later Oracle Exadata System Software releases for backward compatibility.

Following is the list of resource provisioning attributes that you can specify:

- `spaceProvEF`, `spaceProvHC`, `spaceProvXT`: Provisions the vault with the specified amount of extreme flash (EF), high capacity (HC), or extended (XT) storage space.
- `iopsProvEF`, `iopsProvHC`, `iopsProvXT`: Provisions the vault with the specified number of IOPS from extreme flash (EF), high capacity (HC), or extended (XT) storage.
- `flashCacheProv`: Provisions the vault with the specified amount of flash cache space.
- `flashLogProv`: Boolean value (`true` or `false`) indicating whether the vault is provisioned with access to Exadata Smart Flash Log.
- `xrmemCacheProv`: Provisions the vault with the specified amount of Exadata RDMA Memory Cache (XRMEM cache) space.

  Use this attribute only on systems with `HC` or `EF` storage, but not both. For systems with `HC` and `EF` storage, use the following media-specific XRMEM cache provisioning attributes.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `xrmemCacheProvEF`: Provisions the vault with the specified amount of XRMEM cache space associated with `EF` media.
- `xrmemCacheProvHC`: Provisions the vault with the specified amount of XRMEM cache space associated with `HC` media.

You can specify the resource provisioning attributes when you create a vault or separately after vault creation. You can also modify or remove existing attribute settings.

- When you create a vault, using the ESCLI mkvault command, you can specify one or more resource provisioning attributes that are linked directly to the vault.

  For example, the following command creates a new vault that is provisioned with 5 TB of high capacity storage (HC) space and 100 GB of flash cache space:

  ```
  @> mkvault @my-data --attributes spaceProvHC=5T,flashCacheProv=100G
  ```

  > **Note**
  >
  > If you specify a value for any provisioning attributes using the mkvault command, then the vault is limited to use the provisioned resources from the corresponding storage pools and the other storage pools are effectively disabled.
  >
  > For further details, see the mkvault command reference.

- You can also specify resource provisioning attributes that are linked directly to the vault when you modify a vault using the ESCLI chvault command.

  For example, the following command modifies the specified vault by setting the amount of provisioned HC storage space to 10 TB.

  ```
  @> chvault @my-data --attributes spaceProvHC=10T
  ```

  > **Note**
  >
  > The chvault command only changes the specified attributes. All other settings are unaffected.
  >
  > For further details, see the chvault command reference.

- You can effectively remove the resource provisioning settings that are linked directly to the vault by resetting the attributes to their default values. To restore the default values, set any modified attribute with a space or IOPS value to `unlimited`, and reset `flashLogProv` to `true` (if it is set otherwise).

  For example:

  ```
  @> chvault @my-data --attributes
  spaceProvHC=unlimited,flashCacheProv=unlimited
  ```

3.2.3 List Vaults

An Exascale vault is a logical file container.

Copyright © 2022, 2026, Oracle and/or its affiliates.

To display information about vaults, use the ESCLI ls command while the root directory is the current working directory in the ESCLI session. For example:

```
@> ls
```

By default, without any optional arguments, the command displays basic information about all of the vaults in the Exascale cluster.

Using the command options available with the ls command, you can:

- Specify the vault(s) that you want to display.
- Filter the output according to your specified conditions.
- Sort the output using specific attributes.
- Choose to display output with different levels of detail.

3.2.4 Remove a Vault

An Exascale vault is a logical file container.

To delete a vault, use the ESCLI rmvault command. For example:

```
@> rmvault @my-data
```

As shown in the example, the simplest form of the command deletes the specified vault (`@my-data`).

You can add the `--nowait` option to return immediately and run the operation in the background instead of waiting for the command to finish.

You can add the `--force` option to delete the vault and all of the files that it contains. Take extreme care with this option!

If you attempt to delete a vault that is not empty without the `--force` option, then the operation fails and an error is displayed.

3.3 Administer Exascale Files

- Create a File
- Copy a File
- List Files
- Cloning Files
- Snapshotting Files
- Lock File Snapshots
- List Clones and Snapshots of Exascale Files
- Remove a File

Related Topics

- About Exascale Files

Copyright © 2022, 2026, Oracle and/or its affiliates.

3.3.1 Create a File

This topic describes how to create a file in Exascale using ESCLI.

Typically, you implicitly create files in Exascale as a byproduct of performing some action with an Exascale client, such as Oracle Database. For example, when you create a database or add a file to a tablespace.

However, you can also explicitly create a file inside an Exascale vault by using ESCLI.

To create a file using ESCLI, use the mkfile command. For example:

```
@> mkfile @my-vault-name/my-file-name
```

As shown in the example, the simplest form of the command creates an empty file in the specified vault location.

Using other command options available with the mkfile command, you can:

- Specify the file size.

  For example:

  ```
  @> mkfile @my-data/my-file --attributes size=1G
  ```

- Specify the file type.

  For example:

  ```
  @> mkfile @my-data/my-data-file --attributes fileType=DATAFILE
  ```

- Specify a template to apply.

  For example:

  ```
  @> mkfile @my-data/my-data-file2 --template DATAFILE
  ```

- Explicitly specify the file storage attributes.

  For example:

  ```
  @> mkfile @my-data/my-data-file3 --attributes storagePool=my-pool1
  ```

Related Topics

- About Exascale File Storage Attributes
- About Exascale Templates
- Using Oracle Database with Exascale

3.3.2 Copy a File

You can copy a file from your local filesystem into Exascale, or you can copy a file from Exascale into your local filesystem.

Before you copy a file into Exascale, you might want to first create an empty copy of the file so that you can specify the file storage attributes.

Copyright © 2022, 2026, Oracle and/or its affiliates.

To copy a file between the local filesystem and Exascale, use the XSH cp command:

- In the following example, XSH copies the local file located at `/home/my-file` into an Exascale file at `@my-vault-name/my-file-copy`.

  ```
  $ xsh cp /home/my-file @my-vault-name/my-file-copy
  ```

- In the following example, XSH copies the Exascale file located at `@my-vault-name/my-file` to `/home/my-file-copy` on the system running XSH.

  ```
  @> xsh cp @my-vault-name/my-file /home/my-file-copy
  ```

You can also use the ESCLI putfile and getfile commands to copy files to and from Exascale. However, the XSH cp command supports the transfer of much larger files and provides superior transfer performance.

3.3.3 List Files

To display information about files inside an Exascale vault, use the ESCLI ls command and specify the vault name followed by the slash character (/). For example:

```
@> ls @my-vault-name/
```

As shown in the example, a simple form of the command displays basic information about all of the files inside the specified vault.

Alternatively, you can change the working directory in the ESCLI session by using the cd command, and then run the ls command within the context of the specified vault. For example:

```
@> cd @my-vault-name
@my-vault-name/> ls
```

Using the command options available with the ls command, you can also:

- Specify the files that you want to display.
- Filter the output according to your specified conditions.
- Sort the output using specific attributes.
- Choose to display output with different levels of detail.

3.3.4 Cloning Files

An Exascale clone is a thinly-provisioned point-in-time copy of a file that is readable and writable. The source file for a clone can be a regular file, a snapshot, or another clone. Exascale uses redirect-on-write techniques to create and maintain clones very quickly and space-efficiently.

To clone a file or group of files, use the ESCLI clonefile command.

All clones created in the same operation are point-in-time consistent, and all files in a clone operation must be in the same vault.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- To clone a single file, use the clonefile command and specify the source file and the clone destination. For example:

  ```
  @> clonefile @my-vault-name/my-file-name @my-vault-name/my-clone-file
  ```

  In the previous example, the clone (`@my-vault-name/my-clone-file`) is created in the same location as the original file (`@my-vault-name/my-file-name`). In the following example, the clone is created under `@my-vault-name/clones`.

  ```
  @> clonefile @my-vault-name/my-file-name @my-vault-name/clones/my-other-clone-file
  ```

- To clone a group of files, use the clonefile command and specify the source files and the clone destinations by including the wildcard character (*). For example:

  ```
  @> clonefile @my-vault-name/my-files* @my-vault-name/my-clone-files*
  ```

  In the previous example, the clones are created in the same location as the original files. In the following example, the clones are created in a different location.

  ```
  @> clonefile @my-vault-name/my-files* @my-vault-name/more-clones/more-clone-files*
  ```

- To clone multiple files or groups of files, use the clonefile command and specify multiple source and destination pairs. For example:

  ```
  @> clonefile @my-vault-name/file1 @my-vault-name/file1-clone @my-vault-name/my-files* @my-vault-name/my-clone-files*
  ```

  In the previous example, an individual file (`@my-vault-name/file1`) and a group of files (`@my-vault-name/my-files*`) are cloned in the same operation, and the clones (`@my-vault-name/file1-clone` and `@my-vault-name/my-clone-files*`) are created in the same location as the original files.

  In the following example, the clones are created in multiple different locations.

  ```
  @> clonefile @my-vault-name/file1 @my-vault-name/location1/file1-clone @my-vault-name/my-files* @my-vault-name/location2/my-clone-files*
  ```

  When you specify multiple source and destination pairs, the source file specifications are considered in order, and only the first match is used.

  For example, the following command creates clones under `@my-vault-name/clone/a` for files matching `@my-vault-name/a*`, while the clones for the other files are created under `@my-vault-name/clone/other`.

  ```
  @> clonefile @my-vault-name/a* @my-vault-name/clone/a/a* @my-vault-name/* @my-vault-name/clone/other/*
  ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

However, the following command creates all of the clones under `@my-vault-name/clone/other` because all of the files match `@my-vault-name/*`. So, in this case, the second argument pair is never used.

```
@> clonefile @my-vault-name/* @my-vault-name/clone/other/* @my-vault-name/a* @my-vault-name/clone/a/a*
```

- To clone some files while excluding others, use the clonefile command with one of the following exclusion methods:
  - You can use the `--exclude` option to specify files that are excluded from the clone operation.

    For example, the following command clones files in `@my-vault-name`, except for files matching `@my-vault-name/a*` or `@my-vault-name/b*`:

    ```
    @> clonefile @my-vault-name/* @my-vault-name/clone/* --exclude @my-vault-name/a* --exclude @my-vault-name/b*
    ```

  - If you specify a null (empty) string as a clone destination value, no clones are created for the corresponding source files. The null destination and corresponding source is considered in order along with all other source and destination pairs, and only the first match is used.

    For example, the following command also clones files in `@my-vault-name`, except for those matching `@my-vault-name/a*` or `@my-vault-name/b*`.

    ```
    @> clonefile @my-vault-name/a* "" @my-vault-name/b* "" @my-vault-name/* @my-vault-name/clone/*
    ```

After creation, clones are managed using the same commands and procedures that are used for other files. For example, you must use the rmfile command to delete a clone.

Related Topics

- About Exascale File Clones

3.3.5 Snapshotting Files

An Exascale snapshot is a thinly-provisioned read-only point-in-time copy of a file. The source file for a snapshot can be a regular file, a file clone, or another snapshot. Exascale uses redirect-on-write techniques to create and maintain snapshots very quickly and space-efficiently.

To snapshot a file or group of files, use the ESCLI snapshotfile command.

All snapshots created in the same operation are point-in-time consistent, and all files in a snapshot operation must be in the same vault.

- To snapshot a single file, use the snapshotfile command and specify the source file and the snapshot destination. For example:

  ```
  @> snapshotfile @my-vault-name/my-file-name @my-vault-name/my-snap-file
  ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

In the previous example, the snapshot (`@my-vault-name/my-snap-file`) is created in the same location as the original file (`@my-vault-name/my-file-name`). In the following example, the snapshot is created under `@my-vault-name/snapshots`.

```
@> snapshotfile @my-vault-name/my-file-name @my-vault-name/snapshots/my-other-snap-file
```

- To snapshot a group of files, use the snapshotfile command and specify the source files and the snapshot destinations by including the wildcard character (*). For example:

  ```
  @> snapshotfile @my-vault-name/my-files* @my-vault-name/my-snap-files*
  ```

  In the previous example, the snapshots are created in the same location as the original files. In the following example, the snapshots are created in a different location.

  ```
  @> snapshotfile @my-vault-name/my-files* @my-vault-name/more-snapshots/more-snap-files*
  ```

- To snapshot multiple files or groups of files, use the snapshotfile command and specify multiple source and destination pairs. For example:

  ```
  @> snapshotfile @my-vault-name/file1 @my-vault-name/file1-snapshot @my-vault-name/my-files* @my-vault-name/my-snap-files*
  ```

  In the previous example, the snapshot operation processes an individual file (`@my-vault-name/file1`) and a group of files (`@my-vault-name/my-files*`), and the snapshots (`@my-vault-name/file1-snapshot` and `@my-vault-name/my-snap-files*`) are created in the same location as the original files.

  In the following example, the snapshots are created in multiple different locations:

  ```
  @> snapshotfile @my-vault-name/file1 @my-vault-name/location1/file1-snapshot @my-vault-name/my-files* @my-vault-name/location2/my-snap-files*
  ```

  When you specify multiple source and destination pairs, the source file specifications are considered in order, and only the first match is used.

  For example, the following command creates snapshots under `@my-vault-name/snap/a` for files matching `@my-vault-name/a*`, while the snapshots for the other files are created under `@my-vault-name/snap/other`.

  ```
  @> snapshotfile @my-vault-name/a* @my-vault-name/snap/a/a* @my-vault-name/* @my-vault-name/snap/other/*
  ```

  However, the following command creates all of the snapshots under `@my-vault-name/snap/other` because all of the files match `@my-vault-name/*`. So, in this case, the second argument pair is never used.

  ```
  @> snapshotfile @my-vault-name/* @my-vault-name/snap/other/* @my-vault-name/a* @my-vault-name/snap/a/a*
  ```

- To snapshot some files while excluding others, use the snapshotfile command with one of the following exclusion methods:

Copyright © 2022, 2026, Oracle and/or its affiliates.

- You can use the `--exclude` option to specify files that are excluded from the snapshot operation.

  For example, the following command snapshots files in `@my-vault-name`, except for files matching `@my-vault-name/a*` or `@my-vault-name/b*`:

  ```
  @> snapshotfile @my-vault-name/* @my-vault-name/snap/* --exclude @my-vault-name/a* --exclude @my-vault-name/b*
  ```

- If you specify a null (empty) string as a snapshot destination value, no snapshots are created for the corresponding source files. The null destination and corresponding source is considered in order along with all other source and destination pairs, and only the first match is used.

  For example, the following command also snapshots files in `@my-vault-name`, except for those matching `@my-vault-name/a*` or `@my-vault-name/b*`.

  ```
  @> snapshotfile @my-vault-name/a* "" @my-vault-name/b* "" @my-vault-name/* @my-vault-name/snap/*
  ```

After creation, snapshots are generally managed using the same commands and procedures that are used for other files. For example, you can use the rmfile command to delete a snapshot.

Related Topics

- About Exascale File Snapshots

3.3.6 Lock File Snapshots

An Exascale snapshot is a thinly-provisioned read-only point-in-time copy of a file. The source file for a snapshot can be a regular file, a file clone, or another snapshot. Exascale uses redirect-on-write techniques to create and maintain snapshots very quickly and space-efficiently.

Starting with Oracle Exadata System Software release 26.1.0, you can lock a snapshot to prevent its deletion before a specified future date and time. You can also make the lock immutable, which prevents users from removing the lock or changing it to expire earlier.

To set or change a lock on a snapshot file, use the chfile command and set the following attributes:

- `snapLockExpiryTime`: Specifies the snapshot lock expiry time. The value must be an ISO 8601-formatted future date and time.
- `snapLockImmutable`: Boolean value (`true` or `false`) indicating whether the lock is immutable:
  - A mutable (variable) lock (`snapLockImmutable=false`) prevents the snapshot from being deleted before the snapshot lock expiry time (`snapLockExpiryTime`).
  - An immutable lock (`snapLockImmutable=true`) prevents the snapshot from being deleted before the snapshot lock expiry time (`snapLockExpiryTime`), and also prevents users from removing the lock or changing it to expire earlier.

For example:

```
@> chfile @my-vault-name/my-snap-file --attributes
snapLockExpiryTime=2050-06-30T13:00:00,snapLockImmutable=false
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

In the example, the snapshot (`@my-vault-name/my-snap-file`) is locked with a variable lock until 1PM on June 30, 2050.

After locking a snapshot file:

- You can unlock a snapshot with a variable lock (`snapLockImmutable=false`) by setting the `snapLockExpiryTime` attribute to `null`.

  For example:

  ```
  @> chfile @my-vault-name/my-snap-file --attributes snapLockExpiryTime=null
  ```

  > **Note**
  >
  > You cannot unlock a snapshot with an immutable lock (`snapLockImmutable=true`).

- You can convert a variable snapshot lock into an immutable lock by setting the `snapLockImmutable` attribute to `true`. You must also re-specify the snapshot lock expiry time (`snapLockExpiryTime`).

  For example:

  ```
  @> chfile @my-vault-name/my-snap-file --attributes
  snapLockExpiryTime=2050-06-30T13:00:00,snapLockImmutable=true
  ```

- You can extend an immutable lock by setting the `snapLockExpiryTime` attribute to a later date and time. In the command, you must also include `snapLockImmutable=true`.

  For example:

  ```
  @> chfile @my-vault-name/my-snap-file --attributes
  snapLockExpiryTime=2051-12-31T19:00:00,snapLockImmutable=true
  ```

To view the status of a locked snapshot, use the ls command and examine the `snapLockExpiryTime`, `snapLockImmutable`, and `snapLockStatus` attributes.

For example:

```
@> ls @my-vault-name/my-snap-file --detail
id                                  @my-vault-name/my-snap-file
name                                my-snap-file
contentType                         DATA
createTime                          2025-10-29T05:23:05+00:00
createdBy                           exa01
mediaType                           HC
redundancy                          high
size                                20.0000M
snapLockExpiryTime                  2051-12-31T19:00:00+00:00
snapLockImmutable                   true
snapLockStatus                      locked
vaultName                           my-vault-name
```

The `snapLockExpiryTime`, `snapLockImmutable`, and `snapLockStatus` attributes display only for snapshots with a `snapLockExpiryTime` setting.

Copyright © 2022, 2026, Oracle and/or its affiliates.

The `snapLockStatus` attribute is a derived property that indicates the current status of the snapshot lock. The value is `locked` when the lock is active, or `expired` when the lock expiry time has passed.

> **Note**
>
> - The `snapLockExpiryTime`, `snapLockImmutable`, and `snapLockStatus` attributes apply only to snapshot files and are not available for regular (non-snapshot) files.
> - At any time, an Exascale cluster administrator who has the `cl_admin` privilege can modify an immutable lock or delete a locked snapshot.

Related Topics

- About Exascale File Snapshots

3.3.7 List Clones and Snapshots of Exascale Files

To display information about related Exascale files, clones, and snapshots, use the ESCLI lssnapshots command. This command has two principal modes of operation:

- To display information about the snapshots that are based on a specific file, use the lssnapshots command and specify the file that is the subject of the command. The specified file can be a snapshot, clone, or regular file. For example:

  ```
  @> lssnapshots @my-vault-name/file1
  name
  @my-vault-name/file1
  @my-vault-name/snap2-of-file1
  @my-vault-name/snap1-of-file1
  ```

  As shown in the example, a simple form of the command displays the snapshots that are directly based on the specified file.

  Using the command options available with the lssnapshots command, you can also:

  - Filter the output according to your specified conditions.
  - Sort the output using specific attributes.
  - Choose to display output with different levels of detail.
- To display information about related files, clones, and snapshots, use the lssnapshots command and specify a file name in conjunction with the `--all` option or the `--tree` option.

  Both options display information about all files, clones, and snapshots belonging to the snapshot tree that contains the specified file. The specified file can be a snapshot, clone, or regular file at any level in the snapshot tree. No other command options are permitted in conjunction with the `--all` option or the `--tree` option.

  For example, consider a group of related files that is based on the following command sequence:

  ```
  @my-vault-name/> putfile some-local-file file1
  Putting file some-local-file to @my-vault-name/file1
  Success.
  ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
@my-vault-name/> snapshotfile file1 snap1-of-file1
Success.

@my-vault-name/> clonefile file1 clone1
Success.

@my-vault-name/> snapshotfile file1 snap2-of-file1
Success.

@my-vault-name/> snapshotfile clone1 snap1-of-clone1
Success.
```

The following example shows how to display related files using the `--all` command option. The output is presented in a series of lists. The first list contains the files and clones in the snapshot tree. The remaining output lists snapshots that are grouped by their source file.

```
@> lssnapshots @my-vault-name/clone1 --all

---FILE/CLONES:---

2022-05-02 06:14:41 UTC @my-vault-name/file1
2022-05-02 06:15:15 UTC @my-vault-name/clone1

---SNAPSHOTS:---

@my-vault-name/file1
  2022-05-02 06:15:06 UTC @my-vault-name/snap1-of-file1
  2022-05-02 06:15:15 UTC @my-vault-name/?8000_0000_000a:00000001
  2022-05-02 06:15:24 UTC @my-vault-name/snap2-of-file1

@my-vault-name/clone1
  2022-05-02 06:15:45 UTC @my-vault-name/snap1-of-clone1
```

Note that the command output includes an internal snapshot (`@my-vault-name/?8000_0000_000a:00000001`), which was implicitly created as part of the cloning operation to create `@my-vault-name/clone1`.

The following example shows how to display related files in a snapshot tree using the `--tree` command option.

```
@> lssnapshots @my-vault-name/clone1 --tree

+---- [1] @my-vault-name/snap1-of-file1
+---- [2] @my-vault-name/?8000_0000_000a:00000001
|
|`---+---- [4] @my-vault-name/snap1-of-clone1
|    +---- [6] @my-vault-name/clone1
|
+---- [3] @my-vault-name/snap2-of-file1
+---- [5] @my-vault-name/file1
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

3.3.8 Remove a File

To delete a file in an Exascale vault, use the ESCLI rmfile command. For example:

```
@> rmfile @my-vault-name/my-file-name
```

As shown in the example, the simplest form of the command deletes the specified file.

You can add the `--force` option to delete the file even if it is in use.

You can add the `--nowait` option to return immediately and run the operation in the background instead of waiting for the command to finish.

You can add the `--with-snapshots` option to also delete any associated snapshots that are not locked.

3.4 Administer Exascale Datasets

Exascale automatically creates and maintains system-defined datasets, which you can use to track storage utilization for groups of related files, such as all of the files belonging to a specific database.

- List Datasets

Related Topics

- About Exascale Datasets

3.4.1 List Datasets

To display information about an Exascale dataset, use the ESCLI lsdataset command.

- To display a simple list of all datasets in all vaults, use the lsdataset command with no arguments or options. For example:

  ```
  @> lsdataset
  ```

- To display detailed information about a specific dataset, use the lsdataset command with the `--detail` option and specify the dataset identifier.

  Each system-defined dataset has a unique composite identifier, which contains unique identifiers for the associated entities. The dataset identifier has one of the following formats:

  - *`@Vault-name`*: Identifies the vault-level dataset for a specific named vault.
  - *`@Vault-name:GI-cluster-ID`*: Identifies the dataset for a specific Oracle Grid Infrastructure (GI) cluster, which consumes storage space in the specified vault.
  - *`@Vault-name:GI-cluster-ID:CDB-ID`*: Identifies the dataset for a specific Oracle multitenant container database (CDB) that belongs to the specified GI cluster.
  - *`@Vault-name:GI-cluster-ID:CDB-ID:PDB-ID`*: Identifies the dataset for an Oracle pluggable database (PDB) that is associated with the specified CDB, GI cluster, and vault.

Copyright © 2022, 2026, Oracle and/or its affiliates.

For example, the following command displays detailed information about a specific PDB dataset.

```
@> lsdataset --detail @my-
data:a5b4997a027d6f91ffd9729702ff6ec5:1427076301:2164757665
```

- To display information about the files in a specific dataset, use the lsdataset command with the `--files` option and specify the dataset identifier.

  For example, the following command lists the files that belong to the specified GI-level dataset.

  ```
  @> lsdataset --files @my-data:a5b4997a027d6f91ffd9729702ff6ec5
  ```

  If you add the `--recursive` option, then the command output also includes the files in all of the descendant datasets.

  For example, the following command shows the files belonging to the specified GI cluster and all of the associated CDB and PDB datasets.

  ```
  @> lsdataset --files --recursive @my-data:a5b4997a027d6f91ffd9729702ff6ec5
  ```

- When specifying a dataset identifier, the asterisk (*) can be used for wildcard searches.

  For example, the following command lists all CDB datasets associated with the specified GI cluster and vault.

  ```
  @> lsdataset @my-data:a5b4997a027d6f91ffd9729702ff6ec5:*
  ```

Using other command options available with the lsdataset command, you can also:

- Filter the output according to your specified conditions.
- Sort the output using specific attributes.
- Choose the dataset attributes that are displayed.

3.5 Administer Access Control Lists

- List ACLs
- Modify an ACL

Related Topics

- About Exascale Access Control Lists

3.5.1 List ACLs

Access control lists (ACLs) govern the operations that users can perform on Exascale vaults and files.

To display information about ACLs, use the ESCLI lsacl command:

- To display the ACLs for all of the vaults, use the lsacl command with no arguments while the root directory is the current working directory in the ESCLI session. For example:

  ```
  @> lsacl
  ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

- To display the ACLs for all files in a vault, use the lsacl command and specify the vault name along with a wildcard (*) for the files. For example:

  ```
  @> lsacl @my-vault-name/*
  ```

  Alternatively, you can use the lsacl command with no arguments while the current working directory is in the desired vault. For example:

  ```
  @> cd @my-vault-name
  @my-vault-name/> lsacl
  ```

- To display the ACLs for a specific vault or file, use the lsacl command and specify the desired vault or file. For example:

  ```
  @> lsacl @my-vault-name/some-file
  ```

- To display the ACLs for multiple vaults and files, use the lsacl command and specify multiple vault and file specification strings. For example:

  ```
  @> lsacl @my-vault-name @my-vault-name/* @another-vault-name @another-vault-name/some-file
  ```

3.5.2 Modify an ACL

To modify an access control list (ACL) for an Exascale vault or file, use the ESCLI chacl command and specify:

1. The vault or file for which you are modifying the ACL.
2. The ACL string, which defines the modification to the ACL.

   The ACL string has the following format:

   ```
   [+]user-ID1:acl-permission[+][;user-ID2:acl-permission[+]]...
   ```

   In the ACL string:

   - An optional plus (+) operator at the beginning of the ACL string indicates that the specified ACL string merges into the existing ACL. In this case, the changes only impact the users in the specified ACL string, and all other users in the existing ACL retain their permissions. Without this optional plus (+) operator, the ACL is completely overwritten by the specified ACL string, enabling permissions only for the users in the specified ACL string.
   - *user-IDn*: Specifies an Exascale user ID.

     Depending on the user creation method, the user ID may be a system-generated value (for example, `96a68014-5762-4579-86ee-29eb743decbd`) or a user-specified value (for example, `scott`).
   - *acl-permission*: Specifies an ACL permission setting, which can be one of the following:
     - `I | inspect`: Specifies that the user can view attributes of the file or vault, but not its contents.

Copyright © 2022, 2026, Oracle and/or its affiliates.

  - `R | read`: Specifies that the user can read contents of the file, or list files in the vault. Also confers the `inspect` permission.
  - `U | use`: Specifies that the user can write to the file, or use the vault. Also confers all preceding permissions.
  - `M | manage`: Specifies that the user can manage the file or vault. Also confers all preceding permissions.
  - `0 | none`: Specifies that the user is removed from the ACL and loses all permissions. This setting can be used only in conjunction with the initial plus (+) operator to remove a user from an existing ACL.

  Starting with Oracle Exadata System Software release 26.1.0, an optional plus (+) operator following an *acl-permission* value guarantees that the user has at least that privilege level. For example, `R+` ensures the user has at least the `R` (`read`) privilege. If the user already has a higher privilege, such as `M` (`manage`), their privileges remain unchanged. You can only use this option when updating an ACL and you cannot use this option with the `0` (`none`) privilege.

Following are various use-cases for modifying an ACL, along with accompanying examples:

- To set the ACL for a vault, use the chacl command and specify the vault and the new ACL string.

  For example, to make `jenny` the sole manager of `@vault1` use:

  ```
  @> chacl @vault1 jenny:manage
  ```

- To set the ACL for a file, use the chacl command and specify the file and the new ACL string.

  For example, to make `jenny` the manager and `jill` a reader of `@vault1/file1` use:

  ```
  @> chacl @vault1/file1 jenny:manage;jill:read
  ```

- To modify an existing ACL, use the chacl command, specify the vault or file name, and specify an ACL string beginning with a + operator.

  For example, the following command enables `peter` as a user of `@vault1`:

  ```
  @> chacl @vault1 +peter:use
  ```

  In this example, if `peter` is not listed in the existing ACL, an entry for `peter` is added to the ACL. Otherwise, if the existing ACL already includes an entry for `peter`, the existing entry is modified. In either case, all other users in the existing ACL retain their permissions.

- To modify an existing ACL to ensure at least a specified privilege level, use the chacl command, specify the vault or file name, and specify an ACL string beginning with a + operator that also contains ACL privileges with trailing + operators.

  For example, the following command ensures that both `peter` and `jason` have at least the `read` privilege on `@vault1`:

  ```
  @> chacl @vault1 +peter:read+;jason:read+
  ```

  In this example, if `jason` previously had `inspect` privilege and `peter` had `use` privilege, the command would upgrade `jason` to `read`, while `peter` would retain the more powerful `use` privilege.

Copyright © 2022, 2026, Oracle and/or its affiliates.

If either `peter` or `jason` were not previously in the ACL, then they would be added with the `read` privilege. In any case, all other users in the existing ACL would retain their permissions.

The option to specify a minimum privilege level using a trailing + operator is only available starting with Oracle Exadata System Software release 26.1.0.

- To remove a user from an ACL, use the chacl command, specify the vault or file name, and specify an ACL string that begins with + and contains `0` or `none` as the ACL privilege.

  For example, to remove `jill` from the ACL for `@vault1/file1` use:

  ```
  @> chacl @vault1/file1 +jill:none
  ```

3.6 Administer Extended Attributes

- List Extended Attributes
- Set an Extended Attribute
- Remove an Extended Attribute

Related Topics

- About Exascale Extended Attributes

3.6.1 List Extended Attributes

An extended attribute is non-standard user-defined metadata that is associated with an Exascale vault or file.

To display information about extended attributes, use the ESCLI lsxattr command:

- To display all of the extended attributes for a specific vault or file, use the lsxattr command and specify the desired vault or file name. For example:

  ```
  @> lsxattr @my-vault-name/my-file
  ```

- To display a specific extended attribute for a vault or file, use the lsxattr command and specify:
  - The name of the vault or file that is associated with the extended attribute.
  - The extended attribute name.

  For example:

  ```
  @> lsxattr @my-vault-name/my-file --name my-attribute
  ```

- To write a binary extended attribute value to a local file, use the lsxattr command and specify:
  - The name of the vault or file that is associated with the extended attribute.
  - The extended attribute name.
  - The name of the file where you want to write the value of the extended attribute.

Copyright © 2022, 2026, Oracle and/or its affiliates.

For example:

```
@> lsxattr @my-vault-name/my-file --name my-binary-attribute --bin-value-
file /home/my-output-file
```

3.6.2 Set an Extended Attribute

A user-defined extended attribute is non-standard metadata that is associated with an Exascale vault or file.

You can set new extended attributes for a vault or a file, or you can overwrite the value for existing user-defined extended attributes.

> **Note**
>
> Extended attributes with names that start with a dollar sign ($) are reserved for internal use by Exascale services and tools. You cannot create, modify, or delete these extended attributes.

To set a user-defined extended attribute and associate it with an Exascale vault or file, use the ESCLI chxattr command:

- To set an extended attribute with a string value, use the chxattr command and specify:
  - The name of the vault or file that is associated with the extended attribute.
  - The extended attribute name.
  - The extended attribute value.

  For example:

  ```
  @> chxattr @my-vault-name/my-file --name my-attribute --value my-attribute-
  value
  ```

- To set an extended attribute with a binary value, use the chxattr command and specify:
  - The name of the vault or file that is associated with the extended attribute.
  - The extended attribute name.
  - The name of the file that contains the binary extended attribute value.

  For example:

  ```
  @> chxattr @my-vault-name/my-file --name my-binary-attribute --bin-value-
  file /home/my-binary-file
  ```

3.6.3 Remove an Extended Attribute

A user-defined extended attribute is non-standard metadata that is associated with an Exascale vault or file.

To delete an extended attribute that is associated with an Exascale vault or file, use the ESCLI rmxattr command and specify:

- The name of the vault or file that is associated with the extended attribute.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- The extended attribute name.

For example:

```
@> rmxattr @my-vault-name/my-file --name my-attribute
```

> **Note**
>
> Extended attributes with names that start with a dollar sign ($) are reserved for internal use by Exascale services and tools. You cannot create, modify, or delete these extended attributes.

3.7 Administer Vault Templates

- Create a Vault Template
- List Vault Templates
- Modify a Vault Template
- Remove a Vault Template

Related Topics

- About Exascale Templates

3.7.1 Create a Vault Template

A vault template is a named collection of file storage attribute settings defined at the vault level and used when creating new files in the vault. By using vault templates, you can manage file storage attributes automatically and consistently across the vault. A vault template overrides the settings in the corresponding cluster template.

- To create a vault template that is associated with a file type, use the ESCLI mktemplate command with the `--vault` and `--file-type` options and specify the file storage attribute settings.

  Use the following command options to specify the file storage attributes:

  - `--media-type`: Specifies the physical media type used to store files created with the template. Possible values are:
    - `HC`: Identifies high capacity storage, using hard disk drives (HDDs) on high-performance Exadata storage servers.
    - `EF`: Identifies extreme flash storage, using low-latency, high-throughput flash devices.
    - `XT`: Identifies extended storage, using HDDs on Exadata storage servers optimized for cost-effective online storage of less frequently accessed data.
  - `--redundancy`: Specifies the number of data copies that are maintained for files created with the template. Currently, the only permitted value is:
    - `high`: Indicates three mirrored copies of the file data.
  - `--content-type`: Specifies the type of content in files created with the template. Exascale internally uses this attribute to place file extents on physically separate

Copyright © 2022, 2026, Oracle and/or its affiliates.

devices in a manner that maximizes availability in the event of a failure. Possible values are:

* `DATA`: Principally associated with user data.
* `RECO`: Primarily for data used in backup and recovery operations.

For example:

```
@> mktemplate --vault @my-vault-name --file-type DATAFILE --content-type
DATA --media-type HC --redundancy high
```

Additionally, starting with Oracle Exadata System Software release 26.1.0, you can use the `--storage-pool` option to specify the name of a storage pool used to store files created with the template. This option is available only for vault templates.

For example:

```
@> mktemplate --vault @my-vault-name --file-type DATAFILE --content-type
DATA --storage-pool my-pool1 --redundancy high
```

When a template includes a storage pool specification, it does not require a media type setting. If a storage pool and media type are both specified, the settings must be compatible. For example, a template cannot specify the use of high capacity (HC) storage media and also specify the name of a storage pool that uses extreme flash (EF) storage.

- To create a user-defined vault template, use the mktemplate command with the `--vault` and `--name` options and specify the file storage attribute settings.

  To set the file storage attributes, use the same command options as for creating a vault template for a file type, as described above. For example:

  ```
  @> mktemplate --vault @my-vault-name --name my-data-template --content-
  type DATA --media-type HC --redundancy high
  ```

3.7.2 List Vault Templates

A vault template is a named collection of file storage attribute settings defined at the vault level and used when creating new files in the vault. By using vault templates, you can manage file storage attributes automatically and consistently across the vault. A vault template overrides the settings in the corresponding cluster template.

To display information about Exascale vault templates, use the ESCLI lstemplate command with the `--vault` option. For example:

```
@> lstemplate --vault vault-name
```

This simple form of the command displays basic information about all templates that are associated with the specified vault (*vault-name*), including the vault-specific templates and any cluster-level templates not overridden by vault-specific templates.

Using the command options available with the lstemplate command, you can specify the templates you want to display and display output with different levels of detail.

Copyright © 2022, 2026, Oracle and/or its affiliates.

3.7.3 Modify a Vault Template

A vault template is a named collection of file storage attribute settings defined at the vault level and used when creating new files in the vault. By using vault templates, you can manage file storage attributes automatically and consistently across the vault. A vault template overrides the settings in the corresponding cluster template.

- To modify an existing vault template that is associated with a file type, use the ESCLI chtemplate command with the `--vault` and `--file-type` options and specify the file storage attribute settings.

  Use the following command options to specify the file storage attributes:

  – `--media-type`: Specifies the physical media type used to store files created with the template. Possible values are:

    * `HC`: Identifies high capacity storage, using hard disk drives (HDDs) on high-performance Exadata storage servers.
    * `EF`: Identifies extreme flash storage, using low-latency, high-throughput flash devices.
    * `XT`: Identifies extended storage, using HDDs on Exadata storage servers optimized for cost-effective online storage of less frequently accessed data.

  – `--redundancy`: Specifies the number of data copies that are maintained for files created with the template. Currently, the only permitted value is:

    * `high`: Indicates three mirrored copies of the file data.

  – `--content-type`: Specifies the type of content in files created with the template. Exascale internally uses this attribute to place file extents on physically separate devices in a manner that maximizes availability in the event of a failure. Possible values are:

    * `DATA`: Principally associated with user data.
    * `RECO`: Primarily for data used in backup and recovery operations.

  For example:

  ```
  @> chtemplate --vault @my-vault-name --file-type DATAFILE --content-type
  DATA --media-type EF --redundancy high
  ```

  Additionally, starting with Oracle Exadata System Software release 26.1.0, you can use the `--storage-pool` option to specify the name of a storage pool used to store files created with the template. This option is available only for vault templates.

  For example:

  ```
  @> chtemplate --vault @my-vault-name --file-type DATAFILE --content-type
  DATA --storage-pool my-ef-pool1 --redundancy high
  ```

  When a template includes a storage pool specification, it does not require a media type setting. If a storage pool and media type are both specified, the settings must be compatible. For example, a template cannot specify the use of high capacity (HC) storage media and also specify the name of a storage pool that uses extreme flash (EF) storage.

- To modify a user-defined vault template, use the chtemplate command with the `--vault` and `--name` options and specify the file storage attribute settings.

Copyright © 2022, 2026, Oracle and/or its affiliates.

To set the file storage attributes, use the same command options as for modifying a vault template for a file type, as described above. For example:

```
@> chtemplate --vault @my-vault-name --name my-data-template --content-
type DATA --media-type EF --redundancy high
```

3.7.4 Remove a Vault Template

A vault template is a named collection of file storage attribute settings defined at the vault level and used when creating new files in the vault. By using vault templates, you can manage file storage attributes automatically and consistently across the vault. A vault template overrides the settings in the corresponding cluster template.

To remove a vault template, use the ESCLI rmtemplate command with the `--vault` option as follows:

- To remove a template that is associated with a file type, use the `--file-type` option. For example:

  ```
  @> rmtemplate --vault vault-name --file-type file-type
  ```

- To remove a user-defined vault template, use the `--name` option. For example:

  ```
  @> rmtemplate --vault vault-name --name template-name
  ```

3.8 Administer Resource Profiles

- Create a Resource Profile
- List Resource Profiles
- Modify a Resource Profile
- Remove a Resource Profile

Related Topics

- About Exascale Resource Management
- Using Exascale Resource Profiles with Oracle Database
- Using Resource Profiles with Exascale Volumes

3.8.1 Create a Resource Profile

Resource profiles govern how resources are managed within a vault. Each resource profile defines a set of resource limits and settings, which can be associated with various Exascale clients (Oracle databases and block store volumes).

To create a regular resource profile, use the ESCLI mkresourceprofile command and specify:

- The name of the resource profile.
- The name of the vault that the resource profile governs.
- An attribute list that defines the resource limits and settings. The following list outlines all of the available attributes:

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `iopsShareEF`, `iopsShareHC`, `iopsShareXT`: Specifies the relative share of I/O bandwidth (IOPS) from extreme flash (EF), high capacity (HC), or extended (XT) storage that is available to each client associated with the resource profile. Each client's share is relative to all other client shares associated with the corresponding media type. A higher share value implies higher priority. The range of valid values is `1-100`, and the default value is `1`.
- `iopsLimitEF`, `iopsLimitHC`, `iopsLimitXT`: Specifies the upper limit of the I/O bandwidth (IOPS) from extreme flash (EF), high capacity (HC), or extended (XT) storage that is available to each client associated with the resource profile.

  Starting with Oracle Exadata System Software release 26.1.0, each value is a percentage with up to two decimal places. The range of valid values is `0.01-100`, and the default value is `100` (effectively unlimited).

  Before Oracle Exadata System Software release 26.1.0, each value represents a fraction out of 10000. The range of valid values is `1-10000`, and the default value is `10000` (effectively unlimited).
- `enableFlashCache`, `enableXrmemCache`, `enableFlashLog`: Enables or disables use of the flash cache, XRMEM cache, or flash log for clients associated with the resource profile. Each value is Boolean, and the default is `true` (enabled).
- `flashCacheMin`, `xrmemCacheMin`: Specifies the guaranteed minimum fraction of flash cache or XRMEM cache space available to each client associated with the resource profile.

  Starting with Oracle Exadata System Software release 26.1.0, each value is nominally a percentage with up to two decimal places. The range of valid values is `0-100`, and the default value is `0` (no set minimum).

  Before Oracle Exadata System Software release 26.1.0, each value represents a fraction out of 10000. The range of valid values is `0-10000`, and the default value is `0` (no set minimum).

  For each cache type, if the sum of all values across all resource profiles exceeds the maximum, then all the values are proportionally scaled down. Each value applies only when the corresponding cache is enabled in the resource profile (for example, `enableFlashCache=true`).
- `flashCacheSize` / `flashCacheMax`, `xrmemCacheSize` / `xrmemCacheMax`: Specifies the maximum fraction of flash cache or XRMEM cache space available to each client associated with the resource profile. Starting with Oracle Exadata System Software release 26.1.0, the attribute names are `flashCacheSize` and `xrmemCacheSize`. Previously, they were `flashCacheMax` and `xrmemCacheMax`.

  Starting with Oracle Exadata System Software release 26.1.0, each value is nominally a percentage with up to two decimal places. The range of valid values is `0-100`, and the default value is `100`.

  Before Oracle Exadata System Software release 26.1.0, each value represents a fraction out of 10000. The range of valid values is `0-10000`, and the default value is `10000`.

  Each value applies only when the corresponding cache is enabled in the resource profile (for example, `enableFlashCache=true`).

For example:

```
@> mkresourceprofile my-vault-name/my-resource-profile-name --attributes
iopsShareHC=40,iopsLimitHC=6000,enableXrmemCache=false
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

In addition to regular user-defined resource profiles, you can also define a system-reserved resource profile named `$UNASSIGNED`. All Exascale clients not explicitly associated with a resource profile are automatically governed by the `$UNASSIGNED` profile. The `$UNASSIGNED` resource profile contains only two modifiable attributes:

- `flashCacheSize` / `flashCacheMax`: Specifies the maximum fraction of flash cache shared by clients associated with the `$UNASSIGNED` profile.

  Starting with Oracle Exadata System Software release 26.1.0, the attribute name is `flashCacheSize`. The value represents a percentage with up to two decimal places. The range of valid values is `0-100` and the default value is `100`.

  Previously, the attribute name was `flashCacheMax`. The range of valid values was `0-10000` and the default value was `10000`.

- `xrmemCacheSize` / `xrmemCacheMax`: Specifies the maximum fraction of XRMEM cache shared by clients associated with the `$UNASSIGNED` profile.

  Starting with Oracle Exadata System Software release 26.1.0, the attribute name is `xrmemCacheSize`. The value represents a percentage with up to two decimal places. The range of valid values is `0-100` and the default value is `100`.

  Previously, the attribute name was `xrmemCacheMax`. The range of valid values was `0-10000` and the default value was `10000`.

For example:

```
@> mkresourceprofile my-vault-name/$UNASSIGNED --attributes
flashCacheSize=30,xrmemCacheSize=20
```

All Exascale clients governed by the `$UNASSIGNED` profile share the corresponding cache resources. The behavior differs from regular resource profiles, where each application of the resource profile defines the resource allocation for one associated client.

If the `$UNASSIGNED` resource profile does not exist, the system automatically reserves 5% of each cache for unassigned Exascale clients, but unassigned clients share all of the unassigned cache space if more space is available.

3.8.2 List Resource Profiles

Resource profiles govern how resources are managed within a vault. Each resource profile defines a set of resource limits and settings, which can be associated with various Exascale clients (Oracle databases and block store volumes).

To display information about resource profiles, use the ESCLI lsresourceprofile command.

For example, the following command displays basic information about all resource profiles associated with the specified vault (`my-vault-name`):

```
@> lsresourceprofile --filter vault=my-vault-name
```

And the following command displays detailed information about the specified resource profile (`my-vault-name/my-resource-profile-name`):

```
@> lsresourceprofile my-vault-name/my-resource-profile-name --detail
```

Using the command options available with the lsresourceprofile command, you can also:

Copyright © 2022, 2026, Oracle and/or its affiliates.

- Filter the output according to your specified conditions.
- Sort the output using specific attributes.
- Choose to display output with different levels of detail.

3.8.3 Modify a Resource Profile

Resource profiles govern how resources are managed within a vault. Each resource profile defines a set of resource limits and settings, which can be associated with various Exascale clients (Oracle databases and block store volumes).

To modify a resource profile, use the ESCLI chresourceprofile command and specify:

- The name of the resource profile.
- The name of the vault that the resource profile governs.
- An attribute list that defines the resource limits and settings that you want to modify. Unspecified attributes are not modified. The following list outlines all of the available attributes:
  - `iopsShareEF`, `iopsShareHC`, `iopsShareXT`: Specifies the relative share of I/O bandwidth (IOPS) from extreme flash (EF), high capacity (HC), or extended (XT) storage that is available to each client associated with the resource profile. Each client's share is relative to all other client shares associated with the corresponding media type. A higher share value implies higher priority. The range of valid values is `1-100`, and the default value is `1`.
  - `iopsLimitEF`, `iopsLimitHC`, `iopsLimitXT`: Specifies the upper limit of the I/O bandwidth (IOPS) from extreme flash (EF), high capacity (HC), or extended (XT) storage that is available to each client associated with the resource profile.

    Starting with Oracle Exadata System Software release 26.1.0, each value is a percentage with up to two decimal places. The range of valid values is `0.01-100`, and the default value is `100` (effectively unlimited).

    Before Oracle Exadata System Software release 26.1.0, each value represents a fraction out of 10000. The range of valid values is `1-10000`, and the default value is `10000` (effectively unlimited).
  - `enableFlashCache`, `enableXrmemCache`, `enableFlashLog`: Enables or disables use of the flash cache, XRMEM cache, or flash log for clients associated with the resource profile. Each value is Boolean, and the default is `true` (enabled).
  - `flashCacheMin`, `xrmemCacheMin`: Specifies the guaranteed minimum fraction of flash cache or XRMEM cache space available to each client associated with the resource profile.

    Starting with Oracle Exadata System Software release 26.1.0, each value is nominally a percentage with up to two decimal places. The range of valid values is `0-100`, and the default value is `0` (no set minimum).

    Before Oracle Exadata System Software release 26.1.0, each value represents a fraction out of 10000. The range of valid values is `0-10000`, and the default value is `0` (no set minimum).

    For each cache type, if the sum of all values across all resource profiles exceeds the maximum, then all the values are proportionally scaled down. Each value applies only when the corresponding cache is enabled in the resource profile (for example, `enableFlashCache=true`).

Copyright © 2022, 2026, Oracle and/or its affiliates.

  - `flashCacheSize` / `flashCacheMax`, `xrmemCacheSize` / `xrmemCacheMax`: Specifies the maximum fraction of flash cache or XRMEM cache space available to each client associated with the resource profile. Starting with Oracle Exadata System Software release 26.1.0, the attribute names are `flashCacheSize` and `xrmemCacheSize`. Previously, they were `flashCacheMax` and `xrmemCacheMax`.

    Starting with Oracle Exadata System Software release 26.1.0, each value is nominally a percentage with up to two decimal places. The range of valid values is `0-100`, and the default value is `100`.

    Before Oracle Exadata System Software release 26.1.0, each value represents a fraction out of 10000. The range of valid values is `0-10000`, and the default value is `10000`.

    Each value applies only when the corresponding cache is enabled in the resource profile (for example, `enableFlashCache=true`).

For example:

```
@> chresourceprofile my-vault-name/my-resource-profile-name --attributes
iopsShareHC=30,enableXrmemCache=true
```

In addition to regular user-defined resource profiles, you can also define a system-reserved resource profile named `$UNASSIGNED`. All Exascale clients not explicitly associated with a resource profile are automatically governed by the `$UNASSIGNED` profile. The `$UNASSIGNED` resource profile contains only two modifiable attributes:

- `flashCacheSize` / `flashCacheMax`: Specifies the maximum fraction of flash cache shared by clients associated with the `$UNASSIGNED` profile.

  Starting with Oracle Exadata System Software release 26.1.0, the attribute name is `flashCacheSize`. The value represents a percentage with up to two decimal places. The range of valid values is `0-100` and the default value is `100`.

  Previously, the attribute name was `flashCacheMax`. The range of valid values was `0-10000` and the default value was `10000`.

- `xrmemCacheSize` / `xrmemCacheMax`: Specifies the maximum fraction of XRMEM cache shared by clients associated with the `$UNASSIGNED` profile.

  Starting with Oracle Exadata System Software release 26.1.0, the attribute name is `xrmemCacheSize`. The value represents a percentage with up to two decimal places. The range of valid values is `0-100` and the default value is `100`.

  Previously, the attribute name was `xrmemCacheMax`. The range of valid values was `0-10000` and the default value was `10000`.

For example:

```
@> chresourceprofile my-vault-name/$UNASSIGNED --attributes
flashCacheSize=40,xrmemCacheSize=10
```

All Exascale clients governed by the `$UNASSIGNED` profile share the corresponding cache resources. The behavior differs from regular resource profiles, where each application of the resource profile defines the resource allocation for one associated client.

You must create the `$UNASSIGNED` resource profile before you can modify it. If the `$UNASSIGNED` resource profile does not exist, the system automatically reserves 5% of each cache for unassigned Exascale clients, but unassigned clients share all of the unassigned cache space if more space is available.

Copyright © 2022, 2026, Oracle and/or its affiliates.

3.8.4 Remove a Resource Profile

Resource profiles govern how resources are managed within a vault. Each resource profile defines a set of resource limits and settings, which can be associated with various Exascale clients (Oracle databases and block store volumes).

To delete a resource profile, use the ESCLI rmresourceprofile command and specify the resource profile name and the associated vault name.

For example:

```
@> rmresourceprofile my-vault-name/my-resource-profile-name
```

3.9 Administer Exascale Block Store Data Objects

- Administer Exascale Volumes
- Administer Volume Attachments
- Administer Volume Snapshots
- Administer Volume Clones
- Administer Volume Backups
- Administer Volume Groups
- Administer an Advanced Cluster File System (ACFS) on Exascale

Related Topics

- About Exascale Block Store

3.9.1 Administer Exascale Volumes

- Create a Volume
- List Volumes
- Modify a Volume
- Remove a Volume

Related Topics

- About Exascale Volumes

3.9.1.1 Create a Volume

An Exascale block volume is an arbitrary-sized allocation of storage space, which can be used as an Exascale Direct Volume (EDV) attachment.

To create an Exascale block volume, use the ESCLI mkvolume command. For example:

```
@> mkvolume  1TB --attributes vault=my-vault-name,name=my-volume-name
```

As shown in the example, a typical form of the command specifies:

- The amount of storage space that is allocated to the volume.

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> If you intend to use the volume to support Oracle ACFS, note that ACFS requires a minimum volume size of 512 MB.

- The vault where the volume is stored.
- The name for the volume, which makes it easier for you to identify it later.

You can also specify a series of optional attributes, which define the detailed characteristics of the volume or constrain the system resources that the volume is allowed to consume. For instance, you can specify the physical media type used to store the volume. By default, volumes are stored on hard disk drives (`HC` media type).

The following example shows creating a volume that is 100 GB in size and located in the vault named `my-vols`. The example also includes the following specific attribute settings:

- The volume name is set to `vol9`.
- The volume is created on extreme flash (EF) storage.
- The volume is provisioned with 1000 IOPS (IOs per second).

```
@> mkvolume 100g --attributes vault=my-
vols,name=vol9,mediaType=EF,iopsProvisioned=1000
```

After you create a volume, you must create a volume attachment to use the volume. An unattached volume is not available to any clients and is useless in isolation.

> **Note**
>
> For simple cases, you can streamline the creation of an Exascale volume and associated Exascale Direct Volume (EDV) attachment by using the edvmkvol command.

3.9.1.2 List Volumes

An Exascale block volume is an arbitrary-sized allocation of storage space, which can be used as an Exascale Direct Volume (EDV) attachment.

To display information about existing volumes, use the ESCLI lsvolume command. For example:

```
@> lsvolume
```

By default, without any optional arguments, the command displays basic information about all of the volumes associated with the Exascale cluster.

Using the command options available with the lsvolume command, you can:

- Specify the volumes that you want to display.
- Filter the output according to your specified conditions.
- Sort the output using specific attributes.
- Choose to display output with different levels of detail.

Copyright © 2022, 2026, Oracle and/or its affiliates.

3.9.1.3 Modify a Volume

An Exascale block volume is an arbitrary-sized allocation of storage space, which can be used as an Exascale Direct Volume (EDV) attachment.

To modify an existing Exascale block volume, use the ESCLI chvolume command and specify:

- The volume identifier.

  You can use the lsvolume command to find the identifier for each volume.

- The attributes that you want to change.

  You can use the `describe chvolume` command to see details about the volume attributes you can change. See also Describing Resources and Attributes.

For example, the following command sets the size of the specified volume to 200 GB:

```
@> chvolume vol0002_50e52177583f4be4bad68ac20b65001e --attributes size=200g
```

A volume can have up to two owners. To modify volume ownership, you can add or remove one volume owner at a time. To add a volume owner, specify `owners=+`*`user-ID`* as an attribute setting. To remove a volume owner, use `owners=-`*`user-ID`*.

For example:

- The following command adds `peter` as an owner of the specified volume:

  ```
  @> chvolume vol0002_50e52177583f4be4bad68ac20b65001e --attributes
  owners=+peter
  ```

- The following command removes `dave` as an owner of the specified volume:

  ```
  @> chvolume vol0002_50e52177583f4be4bad68ac20b65001e --attributes owners=-
  dave
  ```

3.9.1.4 Remove a Volume

An Exascale block volume is an arbitrary-sized allocation of storage space, which can be used as an Exascale Direct Volume (EDV) attachment.

Removing a volume irreversibly deletes its contents. You cannot remove a volume with current attachments. You must remove all attachments to the volume before the volume can be removed.

To remove an existing volume, use the ESCLI rmvolume command and specify the identifier for the volume that is being removed. For example:

```
@> rmvolume vol0002_50e52177583f4be4bad68ac20b65001e
```

You can use the lsvolume command to find the identifier for each volume.

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> You can use the edvrmvol command to simplify the removal of an Exascale Direct Volume (EDV) attachment and associated Exascale volume that were created using the edvmkvol command.

Related Topics

- Remove an EDV Attachment

3.9.2 Administer Volume Attachments

A volume attachment enables the use of an Exascale volume. An Exascale Direct Volume (EDV) attachment can be used as a raw block storage device or to support a file system.

- Create an EDV Attachment
- List EDV Attachments
- Remove an EDV Attachment

Related Topics

- About Exascale Volume Attachments

3.9.2.1 Create an EDV Attachment

To use a volume, you must create a volume attachment. An Exascale Direct Volume (EDV) attachment can be used as a raw block storage device or to support a file system. To implement Oracle Advanced Cluster File System (ACFS) on Exascale block volume storage, you must use an EDV attachment.

> **Note**
>
> During initial system deployment with Oracle Exadata Deployment Assistant (OEDA), the Exascale Direct Volume (EDV) service is configured on each Exadata compute node (bare-metal or VM) and runs with the permissions of the Exascale user that manages the Oracle Grid Infrastructure (GI) cluster. To create an EDV attachment, you must use the Exascale user linked with the EDV service.
>
> If the GI cluster uses a non-role-separated user configuration with one Oracle OS user account, then the associated Exascale user is related to the EDV service. If the GI cluster uses a role-separated configuration with a Grid OS user account and an Oracle OS user account, then the EDV service is linked to the Exascale user associated with the Grid OS account.
>
> To find the Exascale user linked with the EDV service, use the ESCLI lsinitiator command with the `--detail` option and examine the `user` attribute.

To create an EDV attachment, connect to ESCLI as the Exascale user linked with the EDV service. Then, run the mkvolumeattachment command and specify:

- The volume identifier. You can use the lsvolume command to find the identifier for each volume.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- The device name to use in conjunction with the attachment. This is a user-supplied name, which is applied to the device file that is associated with the attachment. After attachment, the corresponding device file is located under `/dev/exc/`.
- The cluster identifier (`giClusterId`) or specific EDV initiator identifier (`initiator`) where the attachment resides. You can use the lsinitiator command to find cluster and EDV initiator identifiers.

  If you create a cluster-wide attachment, the EDV device file is created on every node in the Oracle Grid Infrastructure (GI) cluster. If you create a node-specific attachment, the corresponding EDV device file is only created on the node associated with the specified EDV initiator.

For example, the following command creates a node-specific EDV attachment:

```
@> mkvolumeattachment vol0001_8e4fbaff261440b493e0a5e5e6808e66 my-vol --
attributes initiator=1fd9b363-079a-84e2-1fd9-b363079a84e2
```

In the example:

- The volume identifier is `vol0001_8e4fbaff261440b493e0a5e5e6808e66`.
- The device name to use in conjunction with the attachment is `my-vol`, and the corresponding device file is located at `/dev/exc/my-vol`.
- The EDV initiator identifier is `1fd9b363-079a-84e2-1fd9-b363079a84e2`. This value identifies a specific cluster node (server) that hosts the EDV attachment.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Note

- Each EDV attachment also has a kernel device file at `/dev/exc-devN`, where *N* is the minor number of the device. The kernel device name is contained as an attribute of the EDV attachment and is visible using the ESCLI lsvolumeattachment command.

  Note that most Linux tools, such as `iostat`, display the kernel device file at `/dev/exc-devN`, while Exascale commands use the user-named device file (under `/dev/exc/`).

  The relationship between the kernel device file and the user-named device file is also recorded in the udev database and is visible using the following Linux command:

  ```
  # udevadm info device-file
  ```

  In the `udevadm` command, for the *device-file* value, you can specify either the kernel device file (`/dev/exc-devN`) or the user-named device file (under `/dev/exc/`).

- By default, read and write access to EDV device files is only available to the `root` operating system user and members of the `disk` group. Depending on your use case, you may need to modify the permissions on the EDV device files before using them.

  For example, to make the EDV device file at `/dev/exc/my-vol` readable and writable by the `oracle` user and `dba` group, you could configure it using a udev rule similar to the following:

  ```
  # cat /etc/udev/rules.d/57-edv-user.rules
  KERNEL=="exc-*", ENV{EXC_ALIAS}=="my-vol", OWNER="oracle",
  GROUP="dba", MODE="0660"
  ```

- To facilitate the management of udev rules related to EDV devices, each EDV client node is configured with a template udev rules file at `/etc/udev/rules.d/57-edv-user.rules`, which you can modify to fulfill your requirements. To maintain existing udev rules, `/etc/udev/rules.d/57-edv-user.rules` is preserved whenever the EDV client software is updated.
- Each EDV client node can support a maximum of 3000 attachments at the same time. This limit includes the total of all cluster attachments involving the server, as well as local attachments specific to the server.
- For simple cases, you can streamline the creation of an Exascale volume and associated Exascale Direct Volume (EDV) attachment by using the edvmkvol command.

3.9.2.2 List EDV Attachments

An Exascale Direct Volume (EDV) attachment can be used as a raw block storage device or to support a file system, such as Oracle Advanced Cluster File System (ACFS).

Copyright © 2022, 2026, Oracle and/or its affiliates.

To display information about EDV attachments, use the ESCLI lsvolumeattachment command. For example:

```
@> lsvolumeattachment
```

By default, without any optional arguments, the command displays basic information about all of the EDV attachments associated with the Exascale cluster.

Using the command options available with the lsvolumeattachment command, you can:

- Specify the attachments that you want to display.
- Filter the output according to your specified conditions.
- Sort the output using specific attributes.
- Choose to display output with different levels of detail.

3.9.2.3 Remove an EDV Attachment

An Exascale Direct Volume (EDV) attachment can be used as a raw block storage device or to support a file system, such as Oracle Advanced Cluster File System (ACFS).

Removing an EDV attachment removes only the ability for clients to access the volume. It does not modify or delete any of the volume contents.

To remove an existing EDV attachment, use the ESCLI rmvolumeattachment command and specify the EDV attachment identifier. For example:

```
@> rmvolumeattachment att0001_90565644befb4c4abab1558b81b940eb
```

In the example, the EDV attachment identifier is `att0001_90565644befb4c4abab1558b81b940eb`. You can use the lsvolumeattachment command to find the identifier for each volume attachment.

> **Note**
>
> You can use the edvrmvol command to simplify the removal of an Exascale Direct Volume (EDV) attachment and associated Exascale volume that were created using the edvmkvol command.

3.9.3 Administer Volume Snapshots

- Create a Volume Snapshot
- Lock Volume Snapshots
- List Volume Snapshots
- Remove a Volume Snapshot

Related Topics

- About Exascale Volume Snapshots

Copyright © 2022, 2026, Oracle and/or its affiliates.

3.9.3.1 Create a Volume Snapshot

A volume snapshot is a thinly-provisioned read-only point-in-time copy of a volume.

To create a volume snapshot, use the ESCLI mkvolumesnapshot command. For example:

```
@> mkvolumesnapshot vol0002_390535212ea7448299bcc7bc7f7653b2
```

As shown in the preceding example, the simplest form of the command specifies the identifier for the volume that you want to copy. You can use the lsvolume command to find the identifier for each volume.

You can also use the `--attributes` option to specify various attributes for the volume snapshot. For example:

```
@> mkvolumesnapshot vol0001_2ee25b4c77fa423a87a25bd2160154de --attributes
name=snap1
```

> **Note**
>
> The volume snapshot operation fails if the specified volume supports Exascale Direct Volume (EDV) attachments and Oracle Advanced Cluster File System (ACFS) is currently mounted. To create a volume snapshot in this situation, you must either:
>
> - Unmount the affected file system on all servers before you run the command to create the volume snapshot. Then, mount the file system again afterward.
> - Alternatively, you can use the Oracle ACFS command-line utility (`acfsutil`) to synchronize the file system and create the volume snapshot. On an EDV client node, run the `acfsutil volsnap create` command as the `root` user and specify the ACFS mountpoint or EDV device.
>
>   For example:
>
>   ```
>   # acfsutil volsnap create /mnt/my-acfs
>   ```

After you create a volume snapshot, you must create an attachment to read the contents. By itself, an unattached volume snapshot is not available to any clients.

3.9.3.2 Lock Volume Snapshots

A volume snapshot is a thinly-provisioned read-only point-in-time copy of a volume.

Starting with Oracle Exadata System Software release 26.1.0, you can lock a volume snapshot to prevent its deletion before a specified future date and time. You can also make the lock immutable, which prevents users from removing the lock or changing it to expire earlier.

To set or change a lock on a volume snapshot, use the chvolumesnapshot command and set the following attributes:

- `lockExpiryTime`: Specifies the lock expiry time. The value must be an ISO 8601-formatted future date and time.
- `lockImmutable`: Boolean value (`true` or `false`) indicating whether the lock is immutable:

Copyright © 2022, 2026, Oracle and/or its affiliates.

  - A mutable (variable) lock (`lockImmutable=false`) prevents the volume snapshot from being deleted before the lock expiry time (`lockExpiryTime`).
  - An immutable lock (`lockImmutable=true`) prevents the volume snapshot from being deleted before the lock expiry time (`lockExpiryTime`), and also prevents users from removing the lock or changing it to expire earlier.

For example:

```
@> chvolumesnapshot vol0001_snap0001_50e52177583f4be4bad68ac20b65001e --
attributes lockExpiryTime=2050-06-30T13:00:00,lockImmutable=false
```

In the example, the volume snapshot (`vol0001_snap0001_50e52177583f4be4bad68ac20b65001e`) is locked with a variable lock until 1PM on June 30, 2050.

After locking a volume snapshot:

- You can unlock a volume snapshot with a variable lock (`lockImmutable=false`) by setting the `lockExpiryTime` attribute to `null`.

  For example:

  ```
  @> chvolumesnapshot vol0001_snap0001_50e52177583f4be4bad68ac20b65001e --
  attributes lockExpiryTime=null
  ```

  > **Note**
  >
  > You cannot unlock a volume snapshot with an immutable lock (`lockImmutable=true`).

- You can convert a variable lock into an immutable lock by setting the `lockImmutable` attribute to `true`. You must also re-specify the lock expiry time (`lockExpiryTime`).

  For example:

  ```
  @> chvolumesnapshot vol0001_snap0001_50e52177583f4be4bad68ac20b65001e --
  attributes lockExpiryTime=2050-06-30T13:00:00,lockImmutable=true
  ```

- You can extend an immutable lock by setting the `lockExpiryTime` attribute to a later date and time. In the command, you must also include `lockImmutable=true`.

  For example:

  ```
  @> chvolumesnapshot vol0001_snap0001_50e52177583f4be4bad68ac20b65001e --
  attributes lockExpiryTime=2051-12-31T19:00:00,lockImmutable=true
  ```

To view the status of a locked volume snapshot, use the lsvolumesnapshot command and examine the `lockExpiryTime`, `lockImmutable`, and `lockStatus` attributes.

For example:

```
@> lsvolumesnapshot vol0001_snap0001_50e52177583f4be4bad68ac20b65001e --
attributes lockExpiryTime,lockImmutable,lockStatus
lockExpiryTime                  2051-12-31T19:00:00+00:00
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
lockImmutable                   true
lockStatus                      locked
```

The `lockExpiryTime`, `lockImmutable`, and `lockStatus` attributes display only for volume snapshots with a `lockExpiryTime` setting.

The `lockStatus` attribute is a derived property that indicates the current status of the lock. The value is `locked` when the lock is active, or `expired` when the lock expiry time has passed.

> **Note**
>
> At any time, an Exascale cluster administrator who has the `cl_admin` privilege can modify an immutable lock or delete a locked volume snapshot.

Related Topics

- About Volume Snapshots

3.9.3.3 List Volume Snapshots

A volume snapshot is a thinly-provisioned read-only point-in-time copy of a volume.

To display information about existing volume snapshots, use the ESCLI lsvolumesnapshot command. For example:

```
@> lsvolumesnapshot
```

By default, without any optional arguments, the command displays basic information about all of the volume snapshots associated with the Exascale cluster.

Using the command options available with the lsvolumesnapshot command, you can:

- Specify the volume snapshots that you want to display.
- Filter the output according to your specified conditions.
- Sort the output using specific attributes.
- Choose to display output with different levels of detail.

3.9.3.4 Remove a Volume Snapshot

A volume snapshot is a thinly-provisioned read-only point-in-time copy of a volume.

Removing a volume snapshot irreversibly deletes it. Exascale provides no means to recover a deleted volume snapshot.

You cannot remove a volume snapshot with current attachments. You must remove all attachments to the volume snapshot before the snapshot can be removed.

To remove an existing volume snapshot, use the ESCLI rmvolumesnapshot command and specify the identifier for the volume snapshot that is being removed. For example:

```
@> rmvolumesnapshot vol0002_snap0001_3c0e9866d66345afbb316472f0bce825
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

You can use the lsvolumesnapshot command to find the identifier for each volume snapshot.

Related Topics

- Remove an EDV Attachment

3.9.4 Administer Volume Clones

- Create a Volume Clone
- List Volume Clones
- Modify a Volume Clone
- Remove a Volume Clone

Related Topics

- About Exascale Volume Clones

3.9.4.1 Create a Volume Clone

An Exascale volume clone is a thinly-provisioned read-write point-in-time copy of a volume snapshot.

You can create an Exascale volume clone by cloning a volume snapshot or by directly cloning an existing volume.

- Create a Volume Clone from a Volume Snapshot
- Create a Volume Clone Directly from an Existing Volume

3.9.4.1.1 Create a Volume Clone from a Volume Snapshot

To create an Exascale volume clone from an Exascale volume snapshot, use the ESCLI mkvolume command and specify the attributes for the volume clone. The command syntax is:

```
@> mkvolume --attributes volumeSnapshot=parent-snapshot-ID[,name=clone-name]
[,iopsProvisioned=integer-value][,iopsInherited={true|false}]
```

In the command:

- `volumeSnapshot=`*`parent-snapshot-ID`*: Identifies the volume snapshot that you want to clone. You can use the lsvolumesnapshot command to find the identifier for each volume snapshot.
- `name=`*`clone-name`*: Optionally specifies the name of the volume clone, which makes it easier for you to identify later. If not specified, a system-generated name is assigned.
- `iopsProvisioned=`*`integer-value`*: Optionally specifies the I/O bandwidth provisioned for the volume clone. The I/O bandwidth is expressed as the number of I/Os per second (IOPS).
- `iopsInherited={true|false}`: Optionally specifies whether the volume clone inherits I/O bandwidth from the nearest ancestor in the volume hierarchy with provisioned (not inherited) I/O bandwidth.

For more details, see the mkvolume command reference.

After you create a volume clone, you must create an attachment to access the contents. By itself, an unattached volume clone is not available to any clients.

Copyright © 2022, 2026, Oracle and/or its affiliates.

3.9.4.1.2 Create a Volume Clone Directly from an Existing Volume

> **Note**
>
> To create an Exascale volume clone directly from an existing Exascale volume, you must use Oracle Exadata System Software release 25.1.0 or later.

To create a volume clone directly from an existing volume, use the ESCLI mkvolume command and specify the attributes for the volume clone. The command syntax is:

```
@> mkvolume --attributes volumeSource=parent-volume-ID[,name=clone-name]
[,iopsProvisioned=integer-value][,iopsInherited={true|false}]
```

In the command:

- `volumeSource=`*`parent-volume-ID`*: Identifies the existing volume that you want to clone. You can use the lsvolume command to find the identifier for each existing volume.
- `name=`*`clone-name`*: Optionally specifies the name of the volume clone, which makes it easier for you to identify later. If not specified, a system-generated name is assigned.
- `iopsProvisioned=`*`integer-value`*: Optionally specifies the I/O bandwidth provisioned for the volume clone. The I/O bandwidth is expressed as the number of I/Os per second (IOPS).
- `iopsInherited={true|false}`: Optionally specifies whether the volume clone inherits I/O bandwidth from the nearest ancestor in the volume hierarchy with provisioned (not inherited) I/O bandwidth.

For more details, see the mkvolume command reference.

After you create a volume clone, you must create an attachment to access the contents. By itself, an unattached volume clone is not available to any clients.

3.9.4.2 List Volume Clones

An Exascale volume clone is a thinly-provisioned read-write point-in-time copy of a volume snapshot.

To display information about existing volumes, including volume clones, use the ESCLI lsvolume command.

To focus on snapshot-based volume clones, look for volumes that contain a value for the `volumeSnapshot` attribute. For example:

```
@> lsvolume --detail --filter volumeSnapshot!=""
```

Without any other arguments or command options, the previous command displays detailed information about all of the snapshot-based volume clones associated with the Exascale cluster.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Likewise, to focus on clones that are based directly on other volumes, look for volumes that contain a value for the `volumeSource` attribute. For example:

```
@> lsvolume --detail --filter volumeSource!=""
```

Using the other command options available with the lsvolume command, you can:

- Display information about specific volumes or volume clones.
- Filter the output according to your specified conditions.
- Sort the output using specific attributes.
- Choose to display output with different levels of detail.

3.9.4.3 Modify a Volume Clone

An Exascale volume clone is a thinly-provisioned read-write point-in-time copy of a volume snapshot.

To modify an existing Exascale volume clone, use the ESCLI chvolume command and specify:

- The volume identifier for the volume clone. You can use the lsvolume command to find the volume identifier for each volume clone.
- The attributes that you want to change.

For a volume clone, you can modify the following attributes:

- `name=`*`clone-name`*: Modifies the name of the volume clone.
- `iopsProvisioned=`*`integer-value`*: Modifies the I/O bandwidth provisioned for the volume clone. The I/O bandwidth is expressed as the number of I/Os per second (IOPS).
- `iopsInherited={true|false}`: Specifies whether the volume clone inherits I/O bandwidth from the nearest ancestor in the volume hierarchy with provisioned I/O bandwidth (not inherited).

  If you change `iopsInherited` from `true` to `false`, then the volume clone is governed by the `iopsProvisioned` setting. If `iopsProvisioned` is not set to a value, then the default value is `unlimited`.

  If you change `iopsInherited` from `false` to `true`, then the volume clone inherits I/O bandwidth from its nearest ancestor and any previous `iopsProvisioned` setting is nullified.

For example, the following command changes the name of the specified volume clone:

```
@> chvolume vol0002_50e52177583f4be4bad68ac20b65001e --attributes name=my-new-
name
```

3.9.4.4 Remove a Volume Clone

An Exascale volume clone is a thinly-provisioned read-write point-in-time copy of a volume snapshot.

Removing a volume clone irreversibly deletes its unique contents but does not affect the underlying volume snapshot. You cannot remove a volume clone with current attachments. You must remove all attachments to the volume clone before the clone can be removed.

Copyright © 2022, 2026, Oracle and/or its affiliates.

To remove an existing volume clone, use the ESCLI rmvolume command and specify the volume identifier for the volume clone that is being removed. For example:

```
@> rmvolume vol0002_50e52177583f4be4bad68ac20b65001e
```

You can use the lsvolume command to find the volume identifier for each volume clone.

Related Topics

- Remove an EDV Attachment

3.9.5 Administer Volume Backups

- Create a Volume Backup
- List Volume Backups
- Restore a Volume Backup
- Monitor the Progress of a Volume Backup or Restore
- Remove a Volume Backup

Related Topics

- About Exascale Volume Backups

3.9.5.1 Create a Volume Backup

A volume backup is a backup of an Exascale volume snapshot, which provides a consistent point-in-time copy of the volume.

You can use the ESCLI mkvolumebackup command to create an Exascale volume backup, which may be stored as a file copy in an Exascale vault.

To create a volume backup, use the mkvolumebackup command and specify (at least):

- The identifier for the volume snapshot that you want to back up.

  You can use the lsvolumesnapshot command to find the identifier for the volume snapshot you want to back up.
- The destination of the backup file as a location in an Exascale vault.

  If the destination path ends with a forward slash (/), a file name is automatically created using the following pattern: `vol_<volume-ID>_bkp_<backup-ID>`.

For example, the following command creates a backup of the specified volume snapshot to `@my-vault1/bkup/`.

```
@> mkvolumebackup --attributes
volumeSnapshot=vol0002_snap0001_3c0e9866d66345afbb316472f0bce825,destination=@
my-vault1/bkup/
```

You can also specify optional attributes for a volume backup, including a user-defined name that you can use to easily identify the backup and Exascale file storage attributes that you want to associate with the backup file.

Copyright © 2022, 2026, Oracle and/or its affiliates.

For example, the following command creates a backup of the specified volume snapshot to `@my-vault1/bkup2/`. In this case, the backup file content type is explicitly set to `RECO` and the backup name is set to `my-backup2`.

```
@> mkvolumebackup --attributes
volumeSnapshot=vol0002_snap0001_3c0e9866d66345afbb316472f0bce825,destination=@
my-vault1/bkup2/,contentType=RECO,name=my-backup2
```

The mkvolumebackup command completes at the beginning of the backup operation. See also Monitor the Progress of a Volume Backup or Restore.

3.9.5.2 List Volume Backups

A volume backup is a backup of an Exascale volume snapshot, which provides a consistent point-in-time copy of the volume.

To display information about existing Exascale volume backups, use the ESCLI lsvolumebackup command. For example:

```
@> lsvolumebackup
```

By default, without any optional arguments, the command displays basic information about all of the volume backups associated with the Exascale cluster.

Using the command options available with the lsvolumebackup command, you can:

- Specify the volume backups that you want to display.
- Filter the output according to your specified conditions.
- Sort the output using specific attributes.
- Choose to display output with different levels of detail.

3.9.5.3 Restore a Volume Backup

A volume backup is a backup of an Exascale volume snapshot, which provides a consistent point-in-time copy of the volume.

You can restore a volume backup into a new volume. When you restore a volume backup, the restored volume is separate from the volume backup and the original backup source volume. There is no ongoing association between a restored volume and the volume backup or the original backup source volume.

To restore a volume backup, use the ESCLI mkvolume command and specify:

- The volume backup identifier to use as the source for the new volume. You can use the lsvolumebackup command to find the identifier for each volume backup.
- The vault that hosts the new volume.

For example:

```
@> mkvolume --attributes
volumeBackup=vol0002_bkp0001_7476692b2499488794951dfe59606e1b,vault=my-vault
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

In the example, the volume backup identifier is `vol0002_bkp0001_7476692b2499488794951dfe59606e1b` and the vault that hosts the restored volume is `my-vault`.

The command completes at the beginning of the restore operation, which may take some time to complete. See Monitor the Progress of a Volume Backup or Restore.

3.9.5.4 Monitor the Progress of a Volume Backup or Restore

A volume backup is a backup of an Exascale volume snapshot, which provides a consistent point-in-time copy of the volume.

The operation to create or restore a volume backup may take some time to complete. You can monitor the ongoing progress by viewing the completion percentage.

- To display the completion percentage for a volume backup operation, use the ESCLI lsvolumebackup command to display information about the desired volume backup operation and examine the `state` and `progress` attributes. For example:

  ```
  @> lsvolumebackup vol0002_bkp0001_7476692b2499488794951dfe59606e1b --
  attributes state,progress --detail
  state           CREATING
  progress        27%

  @>
  ```

  In the example, `vol0002_bkp0001_7476692b2499488794951dfe59606e1b` is the volume backup identifier for the backup operation being monitored.

  You can also use the lsvolumebackup command to find the identifier for each volume backup.

- To display the completion percentage for a volume restore operation, use the ESCLI lsvolume command to display information about the new volume that is being created as part of the volume restore operation and examine the `state` and `progress` attributes. For example:

  ```
  @> lsvolume vol0003_4f71bc664cdb47e3928083505819aafa --attributes
  state,progress --detail
  state           RESTORING
  progress        59%

  @>
  ```

  In the example, the identifier for the new volume being created by the restore operation is `vol0003_4f71bc664cdb47e3928083505819aafa`.

  You can also use the lsvolume command to find the identifier for each volume being created as part of the volume restore operation.

During a backup operation, the `state` attribute contains the value `CREATING`. During a restore operation, the `state` attribute contains the value `RESTORING`. During both operations, the `progress` attribute contains the current completion percentage.

Upon successful completion, the `state` changes to `AVAILABLE`, and the `progress` attribute contains no value. If the operation fails, then the `state` attribute indicates the failure, and the `progress` attribute reflects the point of failure for a while before it is reset.

Copyright © 2022, 2026, Oracle and/or its affiliates.

3.9.5.5 Remove a Volume Backup

A volume backup is a backup of an Exascale volume snapshot, which provides a consistent point-in-time copy of the volume.

To remove an existing volume backup, use the ESCLI rmvolumebackup command and specify the backup identifier. For example:

```
@> rmvolumebackup vol0002_bkp0001_7476692b2499488794951dfe59606e1b
```

You can use the lsvolumebackup command to see details about volume backups, including the identifier for each volume backup.

3.9.6 Administer Volume Groups

- Create a Volume Group
- List Volume Groups
- Adding and Removing Volumes in a Volume Group
- Modify a Volume Group
- Create a Volume Group Snapshot
- Remove a Volume Group

Related Topics

- About Exascale Volume Groups

3.9.6.1 Create a Volume Group

A volume group is a specific collection of Exascale volumes.

To create an Exascale volume group, use the ESCLI mkvolumegroup command. For example:

```
@> mkvolumegroup vol-group1
```

As shown in the example, the simplest form of the command must specify at least the volume group name (`vol-group1`). In this case, the volume group is not enabled for resource sharing, because the `resourceSharing` attribute is set to `none` by default.

A volume group can optionally be configured as a resource-sharing volume group, using either aggregate or specified resource sharing.

- Under aggregate resource sharing, the volume group automatically shares the total I/Os per second (IOPS) bandwidth provisioned to all volumes in the group. In this case, an individual volume can exceed its volume-level limit by sharing unused IOPS provisioned to other volumes in the group. But the entire group is always governed by the aggregated limit.

  For example:

  ```
  @> mkvolumegroup vol-group2 --attributes resourceSharing=aggregate
  ```

- With specified resource sharing, a specific IOPS bandwidth limit is associated with the volume group. In this case, the group is governed only by the specified limit.

Copyright © 2022, 2026, Oracle and/or its affiliates.

When a volume group employs specified resource sharing (`resourceSharing=specified`), the `iopsProvisioned` attribute value must be greater than zero.

For example:

```
@> mkvolumegroup vol-group3 --attributes
resourceSharing=specified,iopsProvisioned=1000
```

A new volume group is initially empty and you must separately modify each volume that you want to associate with the group. See Adding and Removing Volumes in a Volume Group.

3.9.6.2 List Volume Groups

A volume group is a specific collection of Exascale volumes.

To display information about existing Exascale volume groups, use the ESCLI lsvolumegroup command. For example:

```
@> lsvolumegroup
```

By default, without any optional arguments, the command displays basic information about all of the volume groups associated with the Exascale cluster.

Using the command options available with the lsvolumegroup command, you can:

- Specify the volume groups that you want to display.
- Filter the output according to your specified conditions.
- Sort the output using specific attributes.
- Choose to display output with different levels of detail.

3.9.6.3 Adding and Removing Volumes in a Volume Group

A volume group is a specific collection of Exascale volumes. A new volume group is initially empty and you must separately modify each volume that you want to associate with the group.

To associate an Exascale block volume (or volume clone) with a volume group, use the ESCLI chvolume command and:

- Specify the identifier for the volume you are acting on.

  You can use the lsvolume command to find the identifier for each volume.
- Set the `volumeGroup` attribute to identify the volume group that you want the volume to be part of.

  You can use the lsvolumegroup command to find the identifier for each volume group.

For example, the following command adds the volume with ID `vol0001_96ffd78f28fa475fa4e31cf58c5a00db` to the volume group with ID `volgrp0001_c2a3cec3a7e7413181fa1730093b3110`:

```
@> chvolume vol0001_96ffd78f28fa475fa4e31cf58c5a00db --attributes
volumeGroup=volgrp0001_c2a3cec3a7e7413181fa1730093b3110
```

If the volume being changed was previously associated with another volume group, then setting the `volumeGroup` attribute effectively moves the volume from one group to another.

Copyright © 2022, 2026, Oracle and/or its affiliates.

You can also remove a volume from all volume groups by setting the `volumeGroup` attribute to an empty string. For example:

```
@> chvolume vol0001_96ffd78f28fa475fa4e31cf58c5a00db --attributes
volumeGroup=''
```

Note the following details about volume groups. Exascale denies any modification to a volume group or associated volumes if it violates any of the specified conditions.

- A volume group can contain only writable volumes, including volume clones. Volume snapshots cannot be part of a volume group.
- All volumes belonging to a volume group must reside in the same Exascale vault.
- A volume must always have at least one owner in common with each volume group it belongs to.
- A volume can be a member of up to five different volume groups. However, at any time, a volume can belong to only one resource-sharing volume group.

3.9.6.4 Modify a Volume Group

A volume group is a specific collection of Exascale volumes.

To modify an existing Exascale volume group, use the ESCLI chvolumegroup command and specify:

- The volume group identifier.

  You can use the lsvolumegroup command to find the identifier for each volume group.
- The attributes that you want to change.

  You can use the `describe chvolumegroup` command to see details about the volume group attributes you can change. See also Describing Resources and Attributes.

Also note:

- When resource sharing is disabled on a volume group (`resourceSharing=none`), the I/O bandwidth provisioned to the volume group must be zero (`iopsProvisioned=0`).
- When a volume group uses aggregate resource sharing (`resourceSharing=aggregate`), the I/O bandwidth provisioned to the volume group must be zero (`iopsProvisioned=0`).
- When a volume group employs specified resource sharing (`resourceSharing=specified`), the `iopsProvisioned` attribute value must be greater than zero.
- A volume group can have up to two owners. To modify volume group ownership, you can add or remove one owner at a time.

  To add an owner, specify `owners=+`*`user-ID`* as an attribute setting. To remove an owner, use `owners=-`*`user-ID`*.

  A volume must always have at least one owner in common with each volume group it belongs to. Any modification violating this requirement is not permitted.

For example:

Copyright © 2022, 2026, Oracle and/or its affiliates.

- The following command sets the I/O bandwidth provisioning to 3000 I/Os per second (IOPS) for the specified volume group:

  ```
  @> chvolumegroup volgrp0001_08a99eb7e8ff41499dee41de75af3e62 --attributes
  iopsProvisioned=3000
  ```

- The following command adds `peter` as an owner of the specified volume group:

  ```
  @> chvolumegroup volgrp0001_08a99eb7e8ff41499dee41de75af3e62 --attributes
  owners=+peter
  ```

- The following command removes `dave` as an owner of the specified volume group:

  ```
  @> chvolumegroup volgrp0001_08a99eb7e8ff41499dee41de75af3e62 --attributes
  owners=-dave
  ```

To modify the volumes contained in the volume group, you must separately modify each volume. See Adding and Removing Volumes in a Volume Group.

3.9.6.5 Create a Volume Group Snapshot

A volume group is a specific collection of Exascale volumes.

A volume group enables you to easily create consistent snapshots of all volumes in the group. In one operation, you can create a set of volume snapshots with point-in-time consistency, which are easily identifiable by a unique batch ID.

To create a snapshot of a volume group, use the ESCLI mkvolumegroupsnapshot command and specify the identifier of the volume group. For example:

```
@> mkvolumegroupsnapshot volgrp0001_8b599d1725244432a1f47407b4e9493e
```

You can use the lsvolumegroup command to find the identifier for each volume group.

The output from mkvolumegroupsnapshot contains a unique volume group snapshot ID that includes the timestamp of the operation. For example, `1:1c61dd0d85914da3bae34e31a37c2f16_20250624T191208`.

You can use the ID to view all of the volume snapshots as a group using the lsvolumesnapshot command. For example, the following command displays information about all the volume snapshots associated with the specified volume group snapshot ID.

```
@> lsvolumesnapshot --filter
volumeGroupSnapshot=1:1c61dd0d85914da3bae34e31a37c2f16_20250624T191208
```

Related Topics

- Create a Volume Snapshot

3.9.6.6 Remove a Volume Group

A volume group is a specific collection of Exascale volumes.

Before you can remove a volume group, you must first remove all volumes from it. See Adding and Removing Volumes in a Volume Group.

Copyright © 2022, 2026, Oracle and/or its affiliates.

To remove an existing volume group, use the ESCLI rmvolumegroup command and specify the volume group identifier. For example:

```
@> rmvolumegroup volgrp0001_08a99eb7e8ff41499dee41de75af3e62
```

You can use the lsvolumegroup command to find the identifier for each volume group.

3.9.7 Administer an Advanced Cluster File System (ACFS) on Exascale

- Create an ACFS File System on Exascale
- List ACFS File Systems on Exascale
- Modify an ACFS File System on Exascale
- Remove an ACFS File System on Exascale

Related Topics

- About Oracle Advanced Cluster File System (ACFS) on Exascale
- About Oracle ACFS

3.9.7.1 Create an ACFS File System on Exascale

Exascale contains integrated support for Oracle Advanced Cluster File System (ACFS) on Exascale block storage using Exascale Direct Volumes (EDV).

Before you create an Exascale-managed ACFS file system, you must:

- Create the Exascale volume to store the file system.

  The specified volume must be associated with a cluster-wide EDV attachment, otherwise the operation cannot proceed. Using a volume with a node-specific EDV attachment is not permitted.
- Ensure that ACFS is configured appropriately on the target system.

  For example, you can create a file system using ACFS encryption by specifying the `encryptionEnabled`, `encryptionAlgorithm`, and `encryptionKeyLength` attributes. However, the operation fails if ACFS encryption is not initialized on the target system. To initialize ACFS encryption, the system administrator must run the following command before the file system is created:

  ```
  # acfsutil encr init
  ```

To create an Exascale-managed ACFS file system, use the ESCLI mkacfsfilesystem command and specify:

- The identifier of the Exascale volume that you want to use to store the file system. You can use the lsvolume command to find the identifier for each Exascale volume.
- The mount point location where you want to mount the file system.
- Optional attributes that define additional characteristics of the file system.

For example:

```
@> mkacfsfilesystem vol0001_bbd6fb4c75e2411b9bf366fe702eabaf /mnt/acfs1 --
attributes mountReadOnly=true
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

In the example:

- The volume identifier is `vol0001_bbd6fb4c75e2411b9bf366fe702eabaf`.
- The file system mount point is `/mnt/acfs1`.
- The file system is mounted in read-only mode by specifying the optional attribute setting: `mountReadOnly=true`.

Following the operation, the ACFS file system is mounted on every node in the Oracle Grid Infrastructure (GI) cluster and the ACFS details are registered with the GI cluster. After registration, the file system is automatically mounted (or remounted) by Oracle Clusterware as required.

3.9.7.2 List ACFS File Systems on Exascale

Exascale contains integrated support for Oracle Advanced Cluster File System (ACFS) on Exascale block storage using Exascale Direct Volumes (EDV).

To display information about existing Exascale-managed ACFS file systems, use the ESCLI lsacfsfilesystem command. For example:

```
@> lsacfsfilesystem
```

By default, without any optional arguments, the command displays basic information about all of the Exascale-managed ACFS file systems associated with the Exascale cluster.

Using the command options available with the lsacfsfilesystem command, you can:

- Specify the file systems that you want to display.
- Filter the output according to your specified conditions.
- Sort the output using specific attributes.
- Choose to display output with different levels of detail.

3.9.7.3 Modify an ACFS File System on Exascale

Exascale contains integrated support for Oracle Advanced Cluster File System (ACFS) on Exascale block storage using Exascale Direct Volumes (EDV).

To modify an existing Exascale-managed ACFS file system:

1. Use the ESCLI acfsctl command to deregister the file system.

   If you deregister an Exascale-managed ACFS file system, the file system is dismounted and the mount point directory is removed on every node in the Oracle Grid Infrastructure (GI) cluster. The operation also removes the ACFS file system registration in the GI cluster. However, the file system contents remain in the underlying Exascale volume.

   To deregister an Exascale-managed ACFS file system, use the `acfsctl deregister` command and specify the file system identifier. For example:

   ```
   @> acfsctl deregister acfs0001_2fb06f13cddd4a8d8d636d1f794046cb
   ```

   You can use the lsacfsfilesystem command to find the identifier for each Exascale-managed ACFS file system.

   You can also bypass errors in the deregistration process by adding the `--force` option.

Copyright © 2022, 2026, Oracle and/or its affiliates.

2. After deregistering the file system, you can optionally remove and re-create the EDV attachment. See Administer Volume Attachments.

   Re-creating the EDV attachment enables you to move the attachment to a different GI cluster, which effectively moves the file system when it is registered again.

3. Use the ESCLI acfsctl command to register the file system.

   During file system registration, you must specify:

   - The identifier of the Exascale volume that stores the file system. You can use the lsacfsfilesystem command to find the volume identifier associated with each Exascale-managed ACFS file system.

     The specified volume must be associated with an existing cluster-wide EDV attachment and contain an existing ACFS file system; otherwise, the command fails. Registration of a volume with a node-specific EDV attachment is not permitted.

     Furthermore, the ACFS file system is mounted on every node in the Oracle Grid Infrastructure (GI) cluster and the ACFS details are registered with the GI cluster. After registration, the file system is automatically mounted (or remounted) by Oracle Clusterware as required.

   - The mount point location where you want to mount the file system.

     The specified mount point location may differ from the location used in previous registrations.

   During file system registration, you may specify optional attributes governing how the file system is mounted. The permitted attributes are:

   - `mountLeafOwner`, `mountLeafGroup`, `mountLeafMode`: Specify the ownership and permission attributes applied to the leaf directory in the file system mount path.
   - `mountNoRootSUID`: Controls whether non-root users are permitted to run root-owned binaries with the set user ID (setuid) permission.
     - `true`: Only the root user can run root-owned setuid binaries.
     - `false`: Non-root users can also run root-owned setuid binaries.

     If not specified, the default value is `false`.
   - `mountReadOnly`: Specifies whether the file system is mounted with read-only access.
     - `true`: Mounts the file system with read-only access.
     - `false`: Mounts the file system with read-write access.

     If not specified, the default value is `false`.

   You cannot set any other file system attributes during registration.

   You can use the `describe acfsfilesystems` to see further details about all Exascale-managed ACFS file system attributes. See also Describing Resources and Attributes.

   For example:

   ```
   @> acfsctl register vol0001_bbd6fb4c75e2411b9bf366fe702eabaf /mnt/acfs1 --
   attributes mountReadOnly=true
   ```

   In the example:

   - The volume identifier is `vol0001_bbd6fb4c75e2411b9bf366fe702eabaf`.
   - The file system mount point is `/mnt/acfs1`.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- The file system is mounted in read-only mode by specifying the optional attribute setting: `mountReadOnly=true`.

3.9.7.4 Remove an ACFS File System on Exascale

Exascale contains integrated support for Oracle Advanced Cluster File System (ACFS) on Exascale block storage using Exascale Direct Volumes (EDV).

To delete and remove an existing Exascale-managed ACFS file system, use the ESCLI rmacfsfilesystem command and specify the identifier for the file system that is being removed. For example:

```
@> rmacfsfilesystem acfs0001_2fb06f13cddd4a8d8d636d1f794046cb
```

You can use the lsacfsfilesystem command to find the identifier for each Exascale-managed ACFS file system.

The ACFS file system is dismounted and the mount point directory is removed on every node in the Oracle Grid Infrastructure (GI) cluster. The command also removes the ACFS file system registration in the GI cluster. Furthermore, the file system is removed from the underlying volume, effectively deleting the files in the file system. The command does not alter the underlying volume or volume attachment in any other way.

Copyright © 2022, 2026, Oracle and/or its affiliates.

4

Using Exascale with Oracle Grid Infrastructure and Oracle Database

This chapter covers requirements and tasks for using Exascale storage to support Oracle Grid Infrastructure and Oracle Database.

This chapter contains the following topics:

- Using Oracle Grid Infrastructure with Exascale
- Using Oracle Database with Exascale
- About Exascale User Credentials for Oracle Grid Infrastructure and Oracle Database
- Using Oracle Managed Files with Exascale
- Customizing the File Storage Attributes for Oracle Database Files
- Using Exascale Resource Profiles with Oracle Database
- Using Oracle Database with Exascale Snapshots and Clones
- Using Exascale with Oracle Database 19c and Oracle Database 21c
- Copying a TDE Keystore Between a File System and Exascale Storage

4.1 Using Oracle Grid Infrastructure with Exascale

This topic outlines how Exascale storage natively supports Oracle Grid Infrastructure.

Oracle Grid Infrastructure (GI) includes two critical components that manage GI cluster configuration and node membership: Oracle Cluster Registry (OCR), and voting files. Starting with Oracle Grid Infrastructure 23ai release 23.5.0, you can use Exascale to natively store these critical files.

Following is a high-level overview of the procedure for configuring Oracle Grid Infrastructure in conjunction with Exascale. This procedure is automated for you when you use Oracle Exadata Deployment Assistant (OEDA) to create a VM cluster that uses Exascale storage.

1. Provision an Exascale user to manage the OCR and voting files.

   The same Exascale user must be used to access the OCR and voting files on all nodes in the GI cluster.

2. Create an Exascale vault for the OCR and voting files.

   You can dedicate a vault to store the OCR and voting files for a GI cluster, or you can use the vault to also store other files.

   To avoid potential confusion and compromised operations, you should not store the OCR and voting files for multiple GI clusters in one vault.

3. Install the Exascale user wallet into the Oracle Grid Infrastructure software on every node in the GI cluster.

4. Configure Oracle Grid Infrastructure to use the Exascale vault as the location for the OCR and voting files.

Copyright © 2022, 2026, Oracle and/or its affiliates.

After configuration is completed, the OCR and voting files are located in the Exascale vault and can be accessed and managed like any Exascale file. Referencing the OCR and voting files on Exascale is essentially the same as using Oracle Automatic Storage Management (ASM). The fundamental difference is that you use an Exascale vault instead of an ASM disk group.

By default, the OCR and voting files are created with high redundancy (triple mirroring).

Related Topics

- Using Exascale with Oracle Database 19c and Oracle Database 21c

4.2 Using Oracle Database with Exascale

This topic outlines how Exascale storage natively supports Oracle Database.

Starting with Oracle Database 23ai release 23.5.0, you can use Exascale to natively store Oracle Database files such as data files, control files, log files, and so on.

Following is a high-level overview of the procedure for using Oracle Database in conjunction with Exascale. This procedure is automated for you when you use Oracle Exadata Deployment Assistant (OEDA) to create an Oracle database that uses Exascale storage.

1. Provision an Exascale user to manage the Oracle Database files.

   The same Exascale user must be used to access the database files on all nodes in the database cluster.

2. Create an Exascale vault for the Oracle Database files.

   You can dedicate a vault to store the files for one database, or you can use a vault to store the files for many databases.

3. Install the Exascale user wallet on every node in the database cluster.

4. Use the Exascale vault as the location for Oracle Database files.

   Referencing Oracle Database files on Exascale is essentially the same as using Oracle Automatic Storage Management (ASM). The fundamental difference is that you use an Exascale vault instead of an ASM disk group.

   In the following fully-qualified file name example, `@my-vault` specifies the vault that contains the data file.

   ```
   SQL> CREATE TABLESPACE my-tbs1 DATAFILE '@my-vault/my-dbfiles/data/my-
   tbs1.dbf' SIZE 10G;
   ```

   For more details and options regarding Oracle Database files on Exascale, see Using Oracle Managed Files with Exascale and Customizing the File Storage Attributes for Oracle Database Files.

Related Topics

- Using Exascale with Oracle Database 19c and Oracle Database 21c

4.3 About Exascale User Credentials for Oracle Grid Infrastructure and Oracle Database

This topic describes how to identify the Exascale user credentials used in conjunction with Oracle Grid Infrastructure and Oracle Database software.

Copyright © 2022, 2026, Oracle and/or its affiliates.

You can use Exascale to store the Oracle Grid Infrastructure shared clusterware files (Oracle Cluster Registry and voting disks) and Oracle Database files such as data files, control files, log files, and so on.

To natively use Exascale storage, Oracle Grid Infrastructure and Oracle Database must have access to an Exascale user account. The account access is provided using credentials located in a credential store file known as a wallet.

By default, Oracle Exadata Deployment Assistant (OEDA) automatically creates the required Exascale user accounts and wallets when you use it to deploy databases and VM clusters that use Exascale storage. See Exascale User Accounts and Wallets.

Otherwise, to create and administer an Exascale user, you can use the procedures described in Administer Exascale Users and Administer Exascale User Credentials.

To access Exascale storage, Oracle Grid Infrastructure and Oracle Database use the first Exascale user wallet available in the following search path:

1. `$ORACLE_BASE/admin/eswallet`
2. `/etc/oracle/cell/network-config/eswallet`

If the system uses the same operating system (OS) software owner (typically `oracle`) for both Oracle Grid Infrastructure and Oracle Database, then one Exascale user is typically used throughout the Oracle Grid Infrastructure cluster.

A role-separated software installation contains a separate Oracle Grid Infrastructure OS owner (typically `grid`) and Oracle Database OS owner (typically `oracle`). In this case, separate Exascale users are typically associated with the Oracle Database OS owner and the Oracle Grid Infrastructure OS owner.

4.4 Using Oracle Managed Files with Exascale

To store Oracle Database files on Exascale, Oracle recommends using the Oracle Managed Files (OMF) feature of Oracle Database.

Using OMF, you can simply define Oracle Database instance parameters (init.ora parameters) that specify the Exascale vault as the target destination. For example:

```
db_create_file_dest=@my-vault
db_create_online_log_dest_1=@my-vault
db_recovery_file_dest=@my-vault
```

You can also use OMF when specifying an Oracle Database file name in a SQL command. For example:

```
SQL> CREATE TABLESPACE my-tbs DATAFILE '@my-vault' SIZE 10G;
```

If desired, you can specify OMF file locations that contain a more specific path inside an Exascale vault. For example:

```
db_create_file_dest=@my-vault/mylocation/
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> If you specify an OMF file location with a path inside an Exascale vault, you must include a forward slash character (/) at the end of the value.
>
> This requirement also applies when specifying a path inside an Exascale vault in a SQL command or to an Oracle Database client (such as RMAN). For example:
>
> ```
> RMAN> set archivelog destination to '@my-vault/my-logs/';
> ```

Using OMF, Oracle Database files are created and managed automatically using the following file naming convention:

OMF-dest/[*cluster-details*/]*db-unique-name*/[*pdb-guid-name*/]*file-type*/[*date-stamp*/]*name*.OMF.*random*

In the file naming convention:

- *OMF-dest*: Depending on the context, *OMF-dest* specifies the OMF target value specified directly in the SQL command or the value of the relevant OMF instance parameter. For example, if no value is explicitly specified in the SQL command, the value of `db_create_file_dest` is used for a data file, while the value of `db_recovery_file_dest` is used for an archived redo log file.
- *vault*: Specifies the Exascale vault name.
- *cluster-details*: Identifies the Oracle Grid Infrastructure (GI) cluster that contains the Oracle database. The value uses the format: *clustername-clusterGUID*.

  To find the *clusterGUID* for a GI cluster, you can use the following command:

  ```
  # crsctl get css clusterguid
  ```

  If the Oracle database is not associated with a GI cluster, then this element is not included in the file path.
- *db-unique-name*: Identifies the Oracle database by using the *db_unique_name* database parameter value.
- *pdb-guid-name*: Identifies the pluggable database (PDB). If the Oracle database is not a multitenant container database (CDB), then this element is not included in the file path.
- *file-type*: Identifies the Oracle Database file type; for example, `DATAFILE`, `ONLINELOG`, `CONTROLFILE`, and so on.
- *date-stamp*: Specifies the file creation date. This element is only included for specific Oracle Database file types where files are automatically generated over time, such as archive log files and automatically generated backup files.
- *name*: Specifies the OMF file name. The specific value depends on the Oracle Database file type. For example, a data file name contains the relevant tablespace name, while an archive log file name identifies the log thread and sequence number.
- *random*: Specifies a random alpha-numeric string to guarantee uniqueness.

Copyright © 2022, 2026, Oracle and/or its affiliates.

4.5 Customizing the File Storage Attributes for Oracle Database Files

Every Exascale file is associated with file storage attributes that govern how the file is stored and managed.

By default, when a file is created, Exascale automatically applies the attributes from the template that matches the file type. For example, when Oracle Database creates a data file, the `DATAFILE` template is automatically used by default. You cannot change the file storage attributes after the file is created.

However, you can override the default template by specifying the template name in an Oracle Managed Files (OMF) specification or fully-qualified file name. In the following examples, *mytemplate* is used instead of the default template:

- OMF parameter example:

  ```
  db_create_file_dest=@my-vault(my-template)
  ```

- OMF data file example:

  ```
  SQL> CREATE TABLESPACE my-tbs2 DATAFILE '@my-vault(mytemplate)' SIZE 10G;
  ```

- Fully-qualified data file name example:

  ```
  SQL> CREATE TABLESPACE my-tbs3 DATAFILE '@my-vault(my-template)/my-dbfiles/
  data/my-tbs3.dbf' SIZE 10G;
  ```

By specifying a non-default template, you can customize the Exascale file storage attributes for specific databases or files within a database. For example, if the vault is provisioned with high capacity (HC) and extreme flash (EF) storage, you can choose EF storage instead of the default (HC) by using a template with `mediaType=EF`.

Related Topics

- About Exascale File Storage Attributes
- About Exascale Templates

4.6 Using Exascale Resource Profiles with Oracle Database

This topic describes how to manage Exascale resources by associating a database with a resource profile.

By default, without any resource profiles in the system, every Oracle database has unlimited access to Exascale vault I/O resources and multiple databases share vault I/O resources equally when the system is under load.

To enable more granular I/O resource management and prioritize resource utilization, you can associate each Oracle database with an Exascale resource profile, which defines a collection of resource limits and settings. Resource profiles enable consistent resource management across numerous clients by defining resource settings once and applying them to all associated databases. In most deployments, a small number of profiles define a resource hierarchy, and each database is assigned to a single profile.

Copyright © 2022, 2026, Oracle and/or its affiliates.

To associate an Oracle database with an Exascale resource profile, you must specify the resource profile name as the value for the `db_performance_profile` Oracle Database instance parameter (init.ora parameter). For example:

```
SQL> ALTER SYSTEM SET db_performance_profile=my-resource-profile SCOPE=spfile;
```

When the database starts, Exascale notes the `db_performance_profile` setting and automatically applies the corresponding resource profile.

Related Topics

- About Intra-Vault Resource Management Using Exascale Resource Profiles
- Administer Resource Profiles
- Using Resource Profiles with Exascale Volumes

4.7 Using Oracle Database with Exascale Snapshots and Clones

Exascale introduces advanced space-efficient snapshot and cloning capabilities that are tightly integrated with Oracle Database and eliminate the requirement for a test master database to support snapshots and clones on Exadata. The following topics describe typical use cases for Oracle Database in conjunction with Exascale snapshots and clones.

- Thin Cloning a Pluggable Database
- Thin Cloning a Pluggable Database in a Different Container Database
- Using a Carousel of Thinly Provisioned Pluggable Database Snapshots
- Cloning a Container Database
- Capacity Planning for Oracle Database with Exascale Snapshots and Thin Clones

4.7.1 Thin Cloning a Pluggable Database

Within Oracle Database, you can use the `CREATE PLUGGABLE DATABASE` SQL command to create a writable snapshot copy (clone) of an existing pluggable database (PDB). For example, the following command creates a clone of `pdb1` named `pdb1c`.

```
SQL> CREATE PLUGGABLE DATABASE pdb1c
     FROM pdb1 SNAPSHOT COPY;
```

When Oracle Database uses Exascale storage, the PDB snapshot copy operation automatically uses native Exascale cloning functionality to thinly provision the underlying Oracle Database files.

The thin cloned PDB data files reside in the same Exascale vault as the underlying source PDB data files, and Exascale uses redirect-on-write techniques to create and maintain the cloned data files. As a result, the cloning operation is instantaneous, regardless of the size of the underlying data files. Also, the clone is extremely space-efficient. A cloned data file initially consumes no additional space, and physical storage space is only allocated when new data is written.

Within the Oracle Database, the cloned files function like regular data files and appear in the standard Oracle Database dictionary views, such as `DBA_DATA_FILES`. Additionally, you can examine various Exascale-specific file details in the V$EXA_FILE Oracle Database dictionary view.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Within Exascale, you can use the ESCLI lssnapshots command to view the association between the cloned data files and their underlying source files.

> **Note**
>
> If you create a thinly provisioned PDB clone on an Oracle Data Guard primary database, the corresponding standby PDB data files contain a complete (not thinly provisioned) copy of the data, even if the standby database uses Exascale storage. To avoid this, you can prevent the PDB clone from propagating to the standby database by including `STANDBYS=NONE` in the `CREATE PLUGGABLE DATABASE` SQL command.

Related Topics

- Cloning a Local PDB

4.7.2 Thin Cloning a Pluggable Database in a Different Container Database

In conjunction with a database link, you can use the `CREATE PLUGGABLE DATABASE` SQL command to create a writable snapshot copy (clone) of an existing pluggable database (PDB) that resides in a different container database (CDB). For example, the following command creates a clone of `pdb1` named `pdb1c`.

```
SQL> CREATE PLUGGABLE DATABASE pdb1c
     FROM pdb1@CDB1-link SNAPSHOT COPY;
```

In this example, the cloned PDB (`pdb1c`) resides in the CDB processing the SQL command. The original (source) PDB (`pdb1`) resides in the CDB referenced in the database link (`CDB1-link`).

When both CDBs use the same Exascale vault, the PDB snapshot copy operation automatically uses native Exascale cloning functionality to thinly provision the underlying Oracle Database files.

The thin cloned PDB data files reside in the same Exascale vault as the underlying source PDB data files, and Exascale uses redirect-on-write techniques to create and maintain the cloned data files. As a result, the cloning operation is instantaneous, regardless of the size of the underlying data files. Also, the clone is extremely space-efficient. A cloned data file initially consumes no additional space, and physical storage space is only allocated when new data is written.

Within the Oracle Database, the cloned files function like regular data files and appear in the standard Oracle Database dictionary views, such as `DBA_DATA_FILES`. Additionally, you can examine various Exascale-specific file details in the V$EXA_FILE Oracle Database dictionary view.

Within Exascale, you can use the ESCLI lssnapshots command to view the association between the cloned data files and their underlying source files.

You can also use the technique introduced here to clone a PDB residing on an Oracle Data Guard standby database. This approach enables space-efficient test and development databases separate from the primary database.

Copyright © 2022, 2026, Oracle and/or its affiliates.

To clone a PDB from a standby database, you must briefly stop the redo apply process on the standby database while creating the clone. For example

```
DGMGRL> edit database STANDBY set state='apply-off';


SQL> CREATE PLUGGABLE DATABASE pdb1test
     FROM pdb1@STANDBY-link SNAPSHOT COPY;


DGMGRL> edit database STANDBY set state='apply-on';
```

In this example, the cloned PDB (`pdb1test`) resides in the CDB processing the SQL command. The original (source) PDB (`pdb1`) resides in the standby database referenced in the database link (`STANDBY-link`). As in the earlier case, to use Exascale thin clones, the CDB processing the SQL command must use the same Exascale vault that houses the underlying source PDB data files.

> **Note**
>
> Regardless of whether the source PDB is an Oracle Data Guard primary or standby database, if you create a thinly provisioned PDB clone on a primary database, the corresponding standby PDB data files contain a complete (not thinly provisioned) copy of the data, even if the standby database uses Exascale storage. To avoid this, you can prevent the PDB clone from propagating to the standby database by including `STANDBYS=NONE` in the `CREATE PLUGGABLE DATABASE` SQL command.

Related Topics

- Cloning a Remote PDB

4.7.3 Using a Carousel of Thinly Provisioned Pluggable Database Snapshots

Within Oracle Database, you can create a library of pluggable database (PDB) snapshots, known as a PDB snapshot carousel. A PDB snapshot is a point-in-time copy of a PDB. The source PDB can be open read-only or read/write while the snapshot is created. A PDB snapshot carousel is useful for maintaining a library of recent PDB copies for point-in-time recovery and cloning.

For example, the following command creates a snapshot carousel based on `pdb1` where a new snapshot is taken every 2 hours.

```
SQL> ALTER PLUGGABLE DATABASE pdb1
     SNAPSHOT MODE EVERY 2 hours;
```

When Oracle Database uses Exascale storage, the PDB snapshot operation automatically uses native Exascale functionality to snapshot the underlying Oracle Database files. A PDB snapshot carousel based on Exascale storage can contain up to 4096 PDB snapshots. Once the carousel reaches its limit, each new snapshot automatically replaces the oldest snapshot.

Copyright © 2022, 2026, Oracle and/or its affiliates.

The snapshots reside in the same Exascale vault as the underlying source data files, and Exascale uses redirect-on-write techniques to create and maintain the thinly provisioned snapshot data files. As a result, the snapshot operation is instantaneous, regardless of the size of the underlying data files. Also, each snapshot consumes no additional space. Space is only consumed when new data is written to the source PDB or to a writable clone based on a snapshot.

Within the Oracle Database, the snapshot data files appear in the standard Oracle Database dictionary views, such as `DBA_PDB_SNAPSHOTFILE`. Additionally, you can examine various Exascale-specific file details in the V$EXA_FILE Oracle Database dictionary view.

Within Exascale, you can use the ESCLI lssnapshots command to view the association between the snapshots and their underlying source files.

> **Note**
>
> If you create a thinly provisioned PDB snapshot on an Oracle Data Guard primary database, the snapshot metadata propagates to the standby. However, the snapshot is not accessible at the standby site, even after an Oracle Data Guard role transition (switchover or failover).

To utilize a snapshot from the PDB snapshot carousel as a writable PDB, you must create a clone (snapshot copy) based on the snapshot. For example:

```
SQL> CREATE PLUGGABLE DATABASE pdbCloneFromSnap FROM pdb1
     USING SNAPSHOT pdb1SnapshotN SNAPSHOT COPY;
```

In the example, if `pdb1SnapshotN` identifies a thinly provisioned PDB snapshot on Exascale storage, the new PDB (`pdbCloneFromSnap`) is also thinly provisioned.

Note that creating a clone (snapshot copy) based on the PDB snapshot requires access to archived redo log files from the time of the snapshot to ensure consistency across the PDB data files.

Related Topics

- Thin Cloning a Pluggable Database
- Administering a PDB Snapshot Carousel

4.7.4 Cloning a Container Database

Oracle Database contains integrated snapshot and cloning functionality for Oracle multitenant pluggable databases (PDBs), which is primarily available through the `CREATE PLUGGABLE DATABASE` and `ALTER PLUGGABLE DATABASE` SQL commands. When Oracle Database uses Exascale storage, the PDB snapshot and snapshot copy (cloning) functions automatically use native Exascale snapshots and clones to create space-efficient copies of the required files.

However, cloning an entire container database (CDB) requires more than just cloning the database files. The clone requires a unique database name and separate configuration settings for the database instances, requiring coordination of various operations apart from the original database.

To facilitate efficient CDB cloning, Oracle provides a generic database cloning utility called `gDBClone`. Historically, `gDBClone` provides support for various methods of CDB duplication. With

Copyright © 2022, 2026, Oracle and/or its affiliates.

the introduction of Exascale, the gDBClone utility is enhanced to facilitate efficient CDB cloning in conjunction with Exascale storage.

Creating a CDB clone using the gDBClone utility and Exascale storage is:

- Easy. The utility requires only a few inputs.
- Reliable. The entire procedure is automated, reducing the risk of errors.
- Flexible. Using gDBClone, you can clone:
  - Single-instance and Oracle RAC databases.
  - Unencrypted databases and databases using Oracle Transparent Data Encryption (TDE).
  - Stand-alone databases, Oracle Data Guard primary databases, and Oracle Data Guard standby databases.

To use the gDBClone utility, you must first ensure that the source CDB is running in ARCHIVELOG mode.

Then, to create a CDB clone on Exascale storage, use the gDBClone utility with the following syntax and options:

```
# /opt/gDBClone/gDBClone.bin { clone | snap } -sdbname source-db-name -
tdbname target-db-name
     [ -tdbhome target-db-home ]
     [[[ -sga_max_size size-MB ] [ -sga_target size-MB ]] | -pfile file-name ]
     [ -sdbport source-listener-port ]
     [ -racmod { 0 | 1 | 2 }]
     [ -syspwf sys-password-file ]
     [ -walletpwf wallet-password-file ]
```

In the command:

- clone: Instructs gDBClone to create a CDB clone using a complete (not thinly provisioned) database copy. This option is recommended for copying a database into Exascale storage from elsewhere.
- snap: After the source database resides in an Exascale vault, the snap option instructs gDBClone to create a thinly provisioned space-efficient CDB clone in the same Exascale vault.

  Using this option, Exascale uses redirect-on-write techniques to efficiently create and maintain data file clones using minimal storage space. As a result, you can clone an entire CDB in minutes, with the underlying Exascale file cloning operation occurring instantaneously, regardless of the source database size.
- -sdbname *source-db-name*: Specifies the db_unique_name of the source CDB being cloned.
- -tdbname *target-db-name*: Specifies the db_unique_name to use for the new thinly provisioned target CDB.
- -tdbhome *target-db-home*: Optionally specifies the ORACLE_HOME location for the target CDB. If not specified, the target CDB automatically uses the same Oracle Database binaries as the source CDB.
- -sga_max_size *size-MB*: Optionally specifies the maximum size (in megabytes) of the target database system global area (SGA). If not specified, the target CDB automatically uses the same size as the source database.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `-sga_target size-MB`: Optionally specifies the nominal total size (in megabytes) of the target database system global area (SGA). If not specified, the target CDB automatically uses the same size as the source database.
- `-pfile file-name`: Optionally specifies a file that contains one or more supported parameter settings for use by the target database. You can use this option as an alternative to specifying `-sga_max_size` and `-sga_target` on the command line.

  For the list of supported parameters, see gDBClone Powerful Database Clone/Snapshot Management Tool (My Oracle Support Doc ID KB145187).
- `-sdbport source-listener-port`: Optionally specifies the Oracle Net listener port number for the source CDB. If not specified, the default value is 1521.
- `-racmod { 0 | 1 | 2 }`: Optionally specifies the target CDB configuration using one of the following values:
  - `0` - Single instance. This is the `gDBClone` default value.
  - `1` - Oracle RAC One Node.
  - `2` - Oracle RAC. This is the typical setting for databases on Oracle Exadata Exascale.
- `-syspwf sys-password-file`: Optionally specifies a file containing an encrypted version of the source database SYS user password, previously created using `gDBClone` with the `syspwf` option.

  If this option is not specified, `gDBClone` checks the default location (`/opt/gDBClone/SYSpasswd_file`) for a usable password file, and if that is not successful, `gDBClone` prompts to supply the password during command execution.
- `-walletpwf wallet-password-file`: Optionally specifies a file containing an encrypted version of the source database keystore (wallet) password. This option must be specified if the source database uses Oracle Transparent Data Encryption (TDE).

For example, the following command shows the simple case of thin cloning the `SOURCE` CDB to create the `THIN` CDB.

```
# /opt/gDBClone/gDBClone.bin snap -sdbname SOURCE -tdbname THIN -racmod 2
```

In the previous example:

- The `SOURCE` CDB was previously copied into an Exascale vault, and the `THIN` CDB resides in the same vault.
- The `SOURCE` CDB is an Oracle RAC database and the `THIN` CDB is also configured as an Oracle RAC database.
- Both CDBs share the same Oracle binaries and the `THIN` CDB inherits Oracle Database parameter settings from the `SOURCE` CDB.
- The `gDBClone` utility looks for an encrypted file containing the `SOURCE` CDB SYS user password at `/opt/gDBClone/SYSpasswd_file`. If the file doesn't exist or doesn't contain the required password, `gDBClone` prompts for the password at runtime.
- The `SOURCE` CDB is not encrypted using Oracle Transparent Data Encryption (TDE). If the source database contains TDE-encrypted tablespaces, then additional steps are required to configure the TDE keystore (wallet) before creation of the CDB thin clone.

Copyright © 2022, 2026, Oracle and/or its affiliates.

The `gDBClone` utility provides many additional functions. For example, you can view the databases created with `gDBClone` and display the relationships between them by running:

```
# /opt/gDBClone/gDBClone.bin listdbs -tree
```

You can also remove a database quickly and easily by using the `deldb` option. For example:

```
# /opt/gDBClone/gDBClone.bin deldb -tdbname THIN
```

For a complete list of functions and further details, see gDBClone Powerful Database Clone/ Snapshot Management Tool (My Oracle Support Doc ID KB145187).

4.7.5 Capacity Planning for Oracle Database with Exascale Snapshots and Thin Clones

Exascale snapshots and thin clones provide fast, space-efficient copies of Oracle Database files. However, thin provisioning does not eliminate the need for capacity planning. A writable thin clone consumes additional storage as it diverges from the source database. If many clones are created and then modified, the resulting storage consumption can be much higher than the initial clone footprint suggests.

To reduce the risk that clone changes consume all available space in a shared Exascale vault, Exascale uses reserved clone space for updateable thin clones. Reserved clone space provides capacity for changes made to the clone. As additional previously unchanged blocks are modified in the clone, Exascale can increase the reserved space until enough space exists to materialize the full clone. Excluding additional data added to the database, a clone does not consume more storage than the original source database.

As a baseline for capacity planning, Oracle recommends reserving 20% to 25% of the size of the parent data files for each updateable clone. This reserve accounts for changes to clones, parents of snapshots, and associated metadata. It should be sufficient when approximately 5% to 10% of the cloned data changes. Workloads that modify a larger portion of the cloned data require additional capacity.

For example, if a parent database contains 100 TB of data files and you plan to maintain four updateable thin clones, reserve approximately 80 TB to 100 TB of additional storage for clone-related changes and metadata:

```
100 TB parent data files x 4 updateable clones x 20% to 25% = 80 TB to 100 TB
```

This reserve is in addition to the free space required for normal database growth, recovery files, storage pool maintenance, and protection against storage-device failures.

Use the following table to estimate snapshot and clone storage requirements for different amounts of changed data. The values are capacity-planning estimates and should not be interpreted as exact measurements for every workload.

Snapshot or Clone Type	Amount of Changed Data	Estimated Size of Snapshot or Clone	Savings Compared with a Full Copy
Updateable thin clone	0% to 6%	10% to 20% of the parent	80% to 90%
Updateable thin clone	6% to 10%	20% to 25% of the parent	75% to 80%

Copyright © 2022, 2026, Oracle and/or its affiliates.

Snapshot or Clone Type	Amount of Changed Data	Estimated Size of Snapshot or Clone	Savings Compared with a Full Copy
Updateable thin clone	10% to 20%	25% to 40% of the parent	60% to 75%
Updateable thin clone	20% to 40%	40% to 70% of the parent	30% to 60%
Updateable thin clone	40% to 60%	70% to 100% of the parent	0% to 30%
Updateable thin clone	60% to 100%	100% of the parent	0%
Snapshot	Not applicable	1% of the parent	99%
Snapshot or clone of a read-only file	Not applicable	0% of the parent	100%

Note that clone growth depends on the distribution of changed blocks, not only on the volume of write activity. Updating a block that was already modified in the same clone does not require additional reserved clone space. A workload that repeatedly updates the same working set can therefore consume less clone space than a workload that modifies new blocks across a larger portion of the database.

The final rows in the table show that snapshots are highly space efficient, and that files marked read-only in Exascale storage do not allocate additional reserved clone space when they are snapshotted or cloned. If write operations are later enabled on a cloned file that was originally marked read-only, normal clone space reservation behavior resumes.

When reviewing storage consumption, consider the number of updateable clones, the size of the parent data files, and the expected percentage of changed data in each clone. Use snapshot and clone lineage information to identify parent files that support multiple snapshots or clones. The ESCLI lssnapshots command can help identify snapshot and clone relationships.

```
@> lssnapshots @my-vault-name/my-clone-file --all
```

You can also use the ESCLI lsstoragepool command to review storage pool free space as part of the same planning process.

```
@> lsstoragepool my-pool1 --attributes spaceRaw,spaceUsed
```

Reserved clone space is managed automatically by Exascale. This topic provides planning guidance; it does not describe a user-visible setting for configuring clone space reservation.

Related Topics

- About Exascale File Clones
- About Exascale File Snapshots
- Cloning Files
- Snapshotting Files
- List Clones and Snapshots of Exascale Files
- Maintain Free Space in a Storage Pool to Protect Against Disk Failure

Copyright © 2022, 2026, Oracle and/or its affiliates.

4.8 Using Exascale with Oracle Database 19c and Oracle Database 21c

This topic outlines options for using Exascale storage to support Oracle Database 19c and Oracle Database 21c.

Starting with Oracle Database 23ai release 23.5.0, you can use Exascale to natively store and manage the Oracle Grid Infrastructure shared files (OCR and voting files), and Oracle Database files such as data files, control files, log files, and so on. Using Exascale in conjunction with natively supported Oracle software releases enables full functionality for optimal system performance.

However, you can use Exascale storage to support other Oracle Database software releases back to Oracle Database 19c. The following outlines the available options:

- **Use Exascale Direct Volume (EDV).**

  This option uses Oracle Grid Infrastructure release 23.5.0 (or later) and native Exascale vault storage for the Oracle Grid Infrastructure shared files.

  You can use EDV in either of the following ways to provide storage for databases that use Oracle Database software versions earlier than 23.5.0.

  – You can use EDV attachments as ASM grid disks to define an ASM disk group based on Exascale volumes. The resulting disk group can store your Oracle Database files the same way as any other ASM disk group.

  – You can use Oracle Advanced Cluster File System (ACFS) to create a cluster-aware file system based on EDV storage. The resulting file system can store your Oracle Database files and other files, such as Oracle Database binaries, log files, and so on.

  > **Note**
  >
  > By using either of the above approaches, the associated Oracle databases cannot use Exadata Storage Server offloading capabilities, such as Exadata Smart Scan, to offload database processing to the storage servers.

- **Use a Hybrid Configuration with Exascale and Oracle ASM.**

  In a hybrid configuration, part of the available Exadata storage is allocated to Exascale, while the other part is allocated to Oracle ASM. Under this arrangement, you can use Exascale natively for supported Oracle Grid Infrastructure and Oracle Database releases. And, you can use Oracle ASM natively for older Oracle Grid Infrastructure and Oracle Database releases.

  By using this option, all of your Oracle databases can take full advantage of Exadata Smart Storage features, such as Exadata Smart Scan.

Copyright © 2022, 2026, Oracle and/or its affiliates.

> **Note**
>
> Within a single Oracle Grid Infrastructure cluster, all databases must use one database storage model. Databases in the same Oracle Grid Infrastructure cluster can use either native Exascale storage or Oracle ASM, but not both. Hybrid Exascale and Oracle ASM configurations are supported on the same Exadata system by using separate GI clusters.

Related Topics

- Using Exascale Block Storage with EDV

4.9 Copying a TDE Keystore Between a File System and Exascale Storage

You cannot use the XSH `cp` command or the ESCLI `putfile` or `getfile` commands to successfully copy an Oracle Transparent Database Encryption (TDE) keystore (wallet) file in either direction between a file system and Exascale storage. If you copy a TDE keystore between a file system and an Exascale vault using XSH or ESCLI, Oracle Database cannot open the resulting file.

The examples in the following procedure show how to copy a TDE keystore from a file system to an Exascale vault. By appropriately modifying the keystore locations, you can also use the same approach to copy a TDE keystore from Exascale storage to a file system location.

1. Configure Oracle Database to use the desired target location for the TDE keystore.

   Set the `WALLET_ROOT` database parameter to the desired target keystore location. If you want to copy the TDE keystore into Exascale, specify a location in an Exascale vault.

   For example:

   ```
   SQL> ALTER SYSTEM SET WALLET_ROOT=@my-vault/my-db/wallet SCOPE=SPFILE;
   ```

2. Restart the database to use the new keystore location.
3. Create a new (empty) keystore for the database.

   For example:

   ```
   SQL> ADMINISTER KEY MANAGEMENT CREATE KEYSTORE IDENTIFIED BY
   'newKeystorePassword';
   ```

4. Note the location of the newly created keystore.

   For example:

   ```
   SQL> SELECT wrl_parameter FROM v$encryption_wallet;

   WRL_PARAMETER
   --------------------------------------------------------------------------------
   -----
   @my-vault/my-db/wallet/tde/
   ```

5. Merge the contents of the source keystore into the newly created (empty) keystore.

Copyright © 2022, 2026, Oracle and/or its affiliates.

In the following example, the source keystore is on file system storage (at `/u01/app/oracle/admin/my-db/wallet/`) and the newly created keystore is on Exascale storage at the location identified in the previous step.

```
SQL> ADMINISTER KEY MANAGEMENT MERGE KEYSTORE '/u01/app/oracle/admin/my-db/
wallet/'
         IDENTIFIED BY 'originalKeystorePassword'
       INTO EXISTING KEYSTORE '@my-vault/my-db/wallet/tde/'
         IDENTIFIED BY 'newKeystorePassword'
       WITH BACKUP;
```

6. Open the new keystore.

   For example:

```
SQL> ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN IDENTIFIED BY
'newKeystorePassword';
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

5

Using Exascale Block Storage

This chapter outlines how to use Exascale block storage with Exascale Direct Volumes (EDV).

This chapter contains the following topics:

- Using Exascale Block Storage with EDV
- Using Resource Profiles with Exascale Volumes

5.1 Using Exascale Block Storage with EDV

This topic outlines the procedure for using Exascale block storage with Exascale Direct Volumes (EDV).

The Exascale block store provides capabilities to create and manage arbitrary-sized raw block volumes based on Exascale storage.

EDV is recommended for use-cases where clients want to use storage volumes inside the Exadata RDMA Network Fabric. For example:

- You can use EDV to support Oracle Advanced Cluster File System (ACFS) or local Linux file systems (such as XFS, EXT4, and so on) on the Exadata compute nodes.
- You can use EDV raw devices to support earlier Oracle Database versions, not natively supported on Exascale. For example, you can create an Oracle ASM disk group based on EDV raw devices and use the disk group to contain the Oracle Database data files. Or, you can create a file system on an EDV volume and use it to contain the Oracle Database software and data files.

Before you can use Exascale block storage with EDV, note the following:

- The block store manager (BSM) and block store worker (BSW) services must be running in the Exascale cluster.

  To display information about Exascale storage services running across the Exascale cluster, use the ESCLI lsservice command. For example:

  ```
  @> lsservice --detail
  ```

- The Exascale Direct Volume (EDV) service must be running on each Exadata compute node that you want to host an EDV attachment. If you are planning to use a cluster-wide attachment, then the EDV service must be running on every node in the Oracle Grid Infrastructure (GI) cluster.

  To display information about client-side Exascale services, including Exascale Node Proxy (ESNP) and Exascale Direct Volume (EDV), use the DBMCLI `LIST DBSERVER` command, as follows:

  ```
  DBMCLI> list dbserver detail
  ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

The command output displays status information for all of the client-side Exascale services running on the Exadata compute node. You must run the command separately on each compute node.

- During initial system deployment with Oracle Exadata Deployment Assistant (OEDA), the Exascale Direct Volume (EDV) service is configured on each Exadata compute node (bare-metal or VM) and runs with the permissions of the Exascale user that manages the Oracle Grid Infrastructure (GI) cluster. To create an EDV attachment, you must use the Exascale user linked with the EDV service.

  If the GI cluster uses a non-role-separated user configuration with one Oracle OS user account, then the associated Exascale user is related to the EDV service. If the GI cluster uses a role-separated configuration with a Grid OS user account and an Oracle OS user account, then the EDV service is linked to the Exascale user associated with the Grid OS account.

  To find the Exascale user linked with the EDV service, use the ESCLI lsinitiator command with the `--detail` option and examine the `user` attribute.

To begin, you can create an Exascale block volume and EDV attachment.

For example:

```
@>mkvault my-vault
Vault @my-vault created.

@>mkvolume 1g --attributes vault=my-vault,name=edv1
Created volume with id vol0002_7eb2b5cc5d1a47f09abed0fa83514a36

@>lsvolume vol0002_7eb2b5cc5d1a47f09abed0fa83514a36 --detail
id                                  vol0002_7eb2b5cc5d1a47f09abed0fa83514a36
name                                edv1
bandwidthProvisioned                unlimited
contentType                         DATA
creationTime                        2025-06-06T05:06:46+00:00
filePath                            @my-vault/
vol.e2173df2021a4368bbf986483f8332b8
iopsProvisioned                     unlimited
mediaType                           HC
numAttachments                      0
owners                              exa01
redundancy                          high
size                                1G
state                               AVAILABLE
vault                               @my-vault

@>lsinitiator --detail
id                                  e7e0db8c-9a2a-0279-e7e0-db8c9a2a0279
edvBaseDriverVersion                25.2.0.0.0.250602
edvEffectiveDriverVersion           25.2.0.0.0.250602
giClusterId                         deadbeef-badc-0fee-dead-beefbadc0fee
giClusterName                       edvTestCluster
hostName                            exa01
lastHeartbeat                       2025-06-06T05:08:20+00:00
registerTime                        2025-06-03T07:41:10+00:00
state                               ONLINE
user                                exa01
version                             25.2.0.0.0.250602
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
@>mkvolumeattachment vol0002_7eb2b5cc5d1a47f09abed0fa83514a36 my-edv1 --
attributes giClusterId=deadbeef-badc-0fee-dead-beefbadc0fee
Created edv attachment with id att0001_9deed3d3f7f944838f5767151d9f06de

@>lsvolumeattachment att0001_9deed3d3f7f944838f5767151d9f06de --detail
id                                  att0001_9deed3d3f7f944838f5767151d9f06de
attachTime                          2025-06-06T05:11:20+00:00
deviceName                          my-edv1
devicePath                          /dev/exc/my-edv1
giClusterId                         deadbeef-badc-0fee-dead-beefbadc0fee
giClusterName                       edvTestCluster
hostName
initiator
kernelDeviceName                    exc-dev1
logicalSectorSize                   512
volume                              vol0002_7eb2b5cc5d1a47f09abed0fa83514a36
volumeSnapshot

@>
```

The EDV attachment creates an association between the volume and an EDV device file, which resides on the Exadata compute nodes hosting the attachment. A node hosting an EDV attachment is also known as an EDV initiator. If you create a cluster-wide attachment, then the EDV device file is created on every node in the Oracle Grid Infrastructure (GI) cluster. If you create a node-specific attachment, then the corresponding EDV device is only created on that node.

In the example, the EDV attachment is a cluster-wide attachment, and the EDV device name is `my-edv1`, so the corresponding device file is located at `/dev/exc/my-edv1` on each cluster node. The volume identifier (`vol0002_7eb2b5cc5d1a47f09abed0fa83514a36`) was reported to the user during volume creation. Volume identifiers can also be discovered using the ESCLI lsvolume command. The GI cluster identifier (`deadbeef-badc-0fee-dead-beefbadc0fee`) was found by using the ESCLI lsinitiator command.

Alternatively, for a node-specific attachment you must specify the node-specific EDV initiator identifier instead of the GI cluster identifier. For example:

```
@>lsinitiator --detail
id                                  e7e0db8c-9a2a-0279-e7e0-db8c9a2a0279
edvBaseDriverVersion                25.2.0.0.0.250602
edvEffectiveDriverVersion           25.2.0.0.0.250602
giClusterId                         deadbeef-badc-0fee-dead-beefbadc0fee
giClusterName                       edvTestCluster
hostName                            exa01
lastHeartbeat                       2025-06-06T05:21:18+00:00
registerTime                        2025-06-03T07:41:10+00:00
state                               ONLINE
user                                exa01
version                             25.2.0.0.0.250602

@>mkvolumeattachment vol0002_7eb2b5cc5d1a47f09abed0fa83514a36 my-edv1 --
attributes initiator=e7e0db8c-9a2a-0279-e7e0-db8c9a2a0279
Created edv attachment with id att0001_4118794cebae49d6836139b34df08b6b
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
@>lsvolumeattachment att0001_4118794cebae49d6836139b34df08b6b --detail
id                                  att0001_4118794cebae49d6836139b34df08b6b
attachTime                          2025-06-06T05:22:47+00:00
deviceName                          my-edv1
devicePath                          /dev/exc/my-edv1
giClusterId
giClusterName
hostName                            exa01
initiator                           e7e0db8c-9a2a-0279-e7e0-db8c9a2a0279
kernelDeviceName                    exc-dev1
logicalSectorSize                   512
volume                              vol0002_7eb2b5cc5d1a47f09abed0fa83514a36
volumeSnapshot

@>
```

After attachment, the EDV device can be used on the Exadata compute node or GI cluster hosting the attachment.

For example, you could use the following command sequence to create and mount an ACFS file system using the EDV device defined in the previous examples (`/dev/exc/my-edv1`):

```
# # Confirmation of the EDV device at /dev/exc/my-edv1.
# lsblk /dev/exc/my-edv1
NAME     MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
exc-dev1 251:1    0   1G  0 disk
# # Note the default ownership and permissions for the EDV device.
# ls -l /dev/exc/my-edv1
brw-rw---- 2 root disk 251, 1 Jun  6 05:11 /dev/exc/my-edv1
# # Now use the device to support ACFS.
# mkfs -t acfs /dev/exc/my-edv1
mkfs.acfs: version                   = 23.0.0.0.0
mkfs.acfs: ACFS compatibility        = 23.0.0.0.0
mkfs.acfs: on-disk version           = 53.0
mkfs.acfs: volume                    = /dev/exc/my-edv1
mkfs.acfs: volume size               = 1073741824  (   1.00 GB )
mkfs.acfs: file system size          = 1073741824  (   1.00 GB )
mkfs.acfs: Format complete.
# mkdir /mnt/my-acfs1
# mount /dev/exc/my-edv1 /mnt/my-acfs1
# df /mnt/my-acfs1
Filesystem       1K-blocks   Used Available Use% Mounted on
/dev/exc/my-edv1   1048576 318180    730396  31% /mnt/my-acfs1
#
# # Mount the file system on other cluster nodes as required.
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Note

- Each EDV attachment also has a kernel device file at `/dev/exc-devN`, where *N* is the minor number of the device. The kernel device name is contained as an attribute of the EDV attachment and is visible using the ESCLI lsvolumeattachment command.

  Note that most Linux tools, such as `iostat`, display the kernel device file at `/dev/exc-devN`, while Exascale commands use the user-named device file (under `/dev/exc/`).

  The relationship between the kernel device file and the user-named device file is also recorded in the udev database and is visible using the following Linux command:

  ```
  # udevadm info device-file
  ```

  In the `udevadm` command, for the *device-file* value, you can specify either the kernel device file (`/dev/exc-devN`) or the user-named device file (under `/dev/exc/`).

- By default, read and write access to EDV device files is only available to the `root` operating system user and members of the `disk` group. Depending on your use case, you may need to modify the permissions on the EDV device files before using them.

  For example, to make the EDV device file at `/dev/exc/my-vol` readable and writable by the `oracle` user and `dba` group, you could configure it using a udev rule similar to the following:

  ```
  # cat /etc/udev/rules.d/57-edv-user.rules
  KERNEL=="exc-*", ENV{EXC_ALIAS}=="my-vol", OWNER="oracle",
  GROUP="dba", MODE="0660"
  ```

- To facilitate the management of udev rules related to EDV devices, each EDV client node is configured with a template udev rules file at `/etc/udev/rules.d/57-edv-user.rules`, which you can modify to fulfill your requirements. To maintain existing udev rules, `/etc/udev/rules.d/57-edv-user.rules` is preserved whenever the EDV client software is updated.
- Each EDV client node can support a maximum of 3000 attachments at the same time. This limit includes the total of all cluster attachments involving the server, as well as local attachments specific to the server.

An ACFS file system on EDV can also be exported to clients outside the Exadata RDMA Network Fabric by using ACFS HANFS. For more information about ACFS, see Oracle Advanced Cluster File System Administrator's Guide.

EDV devices can also be used as raw devices or in conjunction with other Linux file systems, such as XFS, EXT4, and so on.

5.2 Using Resource Profiles with Exascale Volumes

This topic describes how to manage Exascale resources by associating resource profiles with Exascale volumes and volume groups.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Before Oracle Exadata System Software release 26.1.0, resource management for Exascale volumes and volume groups was limited to provisioning a number of I/Os per second (IOPS) for each volume or volume group using the `iopsProvisioned` attribute. Using this arrangement, access to other resources (such as cache space) was proportional to the `iopsProvisioned` setting.

Starting with Oracle Exadata System Software release 26.1.0, you also associate resource profiles with Exascale volumes and volume groups. This change aligns resource management across all clients that share an Exascale vault and enables consistent management of many Oracle databases, Exascale volumes, and volume groups.

To associate a resource profile with an Exascale volume or volume group, you must specify the resource profile name in the `resourceProfile` attribute. You can specify the resource profile setting when you create a new volume or volume group. For example:

```
@> mkvolume 100g --attributes vault=my-vols,resourceProfile=gold-resource-
profile
```

```
@> mkvolumegroup vol-group1 --attributes resourceProfile=gold-resource-profile
```

You can also change the resource profile setting when you alter an existing volume or volume group. For example:

```
@> chvolume vol0125_79e1c93ccfd246be8fe14a8092f38991 --attributes
resourceProfile=silver-resource-profile
```

```
@> chvolumegroup volgrp0025_ce7c459718d94a2ea8eaa066adf43db7 --attributes
resourceProfile=silver-resource-profile
```

The following requirements and restrictions also apply:

- Each volume or volume group can be associated with only one resource profile at a time.
- A volume or volume group cannot be associated with a resource profile and at the same time have I/O bandwidth provisioned using the `iopsProvisioned` attribute.

  For existing volumes and volumes groups you must set the provisioned I/O bandwidth to zero (`iopsProvisioned=0`) before you can specify the `resourceProfile` attribute.
- If a volume is part of a volume group, the group's resource sharing policy (specified by the `resourceSharing` attribute setting) governs the behavior of resource management:
  - If the group is not enabled for resource sharing (`resourceSharing=none`), resource management is governed by the resource profile settings associated with each individual volume.
  - If the group uses specified resource sharing (`resourceSharing=specified`), resource management applies the resource profile associated with the volume group.
  - The aggregate group resource sharing policy (`resourceSharing=aggregate`) is not permitted in conjunction with resource profiles. An error is returned if you attempt to create an incompatible configuration.
- Unless explicitly specified otherwise, when a volume clone is created, the `resourceProfile` setting is copied from the parent volume. After creation, the clone is an independent entity and the `resourceProfile` setting can be changed independently.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- A volume snapshot inherits the `resourceProfile` setting from its parent volume. In this case, no change is permitted to the `resourceProfile` setting.

Related Topics

- About Intra-Vault Resource Management Using Exascale Resource Profiles
- Administer Resource Profiles
- Using Exascale Resource Profiles with Oracle Database

Copyright © 2022, 2026, Oracle and/or its affiliates.

6
Using ESCLI

This chapter describes how to use the Exascale command line interface (ESCLI):

- Start and Use ESCLI
- ESCLI Command Reference

6.1 Start and Use ESCLI

This topic describes how to start and use the Exascale command line interface (ESCLI).

ESCLI is a command-line administration tool that is located on Exadata compute nodes and storage servers. You can use ESCLI to perform Exascale monitoring and management functions. ESCLI works in conjunction with the Exadata cell command-line interface (CellCLI) and does not replace it.

To use ESCLI you must have access to an Exascale user account and the digital key store (wallet) for the Exascale user account.

To start ESCLI, use the following command line syntax:

```
$ escli [ --wallet wallet-location ] [ --ctrl server[:port] ] [ ESCLI-
command ]
```

In the command line:

- `--wallet` *`wallet-location`*: Identifies the wallet that you want to use for authentication to Exascale. If you do not specify the wallet location, then ESCLI checks the local server and uses the first Exascale user wallet available in the following search path:
  1. `$OSSCONF/eswallet`
  2. `$ORACLE_BASE/admin/eswallet`
  3. `/etc/oracle/cell/network-config/eswallet`
- `--ctrl` *`server`*`[:`*`port`*`]`: Specifies the endpoint address for connection to Exascale control services, also known as Exascale RESTful Services (ERS). If the ERS address is not specified on the command line, then ESCLI connects to the ERS address stored in the user's wallet.
- *ESCLI-command*: Specifies an ESCLI command to run immediately. For example:

```
$ escli --wallet /home/user/user.wallet ls -l @my-data
Total 5
  10.0M   05 Jan 12:59 x
  19.5k   05 Jan 13:08 y
   5.0k   05 Jan 13:09 z1
  10.0M   05 Jan 13:23 z2
  20.0G   05 Jan 13:14 z3
$
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

If an ESCLI command is not specified on the `escli` command line, ESCLI starts an interactive session and presents a command prompt. During an interactive session, you can run a series of ESCLI commands. The interactive session ends with the `exit` command. For example:

```
$ escli --wallet /home/user/user.wallet
@> ls
my-data
vault2

@> cd @my-data

@my-data/> ls -l
Total 5
  10.0M   05 Jan 12:59 x
  19.5k   05 Jan 13:08 y
   5.0k   05 Jan 13:09 z1
  10.0M   05 Jan 13:23 z2
  20.0G   05 Jan 13:14 z3

@my-data/> exit
$
```

During an interactive session, ESCLI supports command auto-completion by using the <tab> key. When you press the <tab> key, ESCLI attempts to complete the current command token. If multiple possibilities exist, ESCLI provides a list of alternatives. For example:

```
@> lsc<tab>
lscell                    (List all cells in the cluster)
lscelldisk             (Lists cell disks in this cluster)
lscluster                      (List cluster information)
lscomputeserver (List all compute servers in the cluster)
```

Command auto-completion also works within a command to suggest or complete command options. For example:

```
@> lscell --<tab>
--attributes                         (Lists the specified attributes)
--count                          (Maximum number of results to report)
--detail                    (Lists all attributes in a detailed form)
--filter (Used to specify conditions for filtering the list output)
--sort     (Used to sort the output using the specified attributes)
```

6.2 ESCLI Command Reference

This section contains references for the Exascale command line interface (ESCLI) commands. Apart from command help and object description, which are universal, the commands are grouped by functional area:

- ESCLI Command Help
- Describing Resources and Attributes
- Service and Cluster Management
- Security and User Management

Copyright © 2022, 2026, Oracle and/or its affiliates.

- Storage Pool and Pool Disk Management
- Vault Management
- Dataset Management
- File Management
- Template Management
- Resource Profile Management
- Extended Attribute Management
- Block Store Management

6.2.1 ESCLI Command Help

Purpose

The ESCLI `help` command displays help information for Exascale command line interface (ESCLI) commands.

Syntax

```
help ( command | list | all )
```

Command Options

The options for the `help` command are:

- *command*: Specifies the ESCLI command for which you want to print help information.
- `list`: Lists all of the available ESCLI commands.
- `all`: Prints help information for all of the available ESCLI commands.

Examples

Example 6-1 Print Help Information for a Specific Command

The following example prints help information for the `ls` command.

```
@> help ls
```

6.2.2 Describing Resources and Attributes

Purpose

The `describe` command displays information about Exascale resources and their attributes.

Syntax

```
describe { --all | resource-spec } [ --detail ]
```

Command Options

The options for the `describe` command are:

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `--all`: Directs the command to display information for all Exascale resources. This option displays a large amount of information, but it is a good starting point if you are unfamiliar with Exascale resources and their attributes.
- *resource-spec*: Directs the command to display information about specific Exascale resources and their attributes as follows:
  - The *resource-spec* value is not case-sensitive.

    Consequently, the following commands are functionally equivalent:

    ```
    @> describe vaults
    ```

    ```
    @> describe VAULTS
    ```

    ```
    @> describe vAUltS
    ```

  - You can specify resource names in singular or plural form.

    Consequently, the following commands are functionally equivalent:

    ```
    @> describe vault
    ```

    ```
    @> describe vaults
    ```

  - If you specify a complete resource name (in singular or plural form), the output contains all attributes of the specified resource.
  - If you specify a resource name with a wildcard (*), the output contains all resources that match the pattern.

    For example, the following command describes all resources starting with the letter 'v' (vaults, volumes, and so on):

    ```
    @> describe v*
    ```

  - If you specify an attribute name (using the notation: `<resource>.<attribute>`), the output describes the specified attribute.

    For example, the following command describes the attribute that contains the vault name.

    ```
    @> describe vaults.name
    ```

  - If you specify an attribute name with a wildcard (*), the output contains all attributes that match the pattern.

    For example, the following command describes all attributes starting with 'n' in all resources starting with 'v':

    ```
    @> describe v*.n*
    ```

  - If *resource-spec* begins with 'mk', the output contains only the attributes that are initializable during resource creation (`initializable != No`).

Copyright © 2022, 2026, Oracle and/or its affiliates.

For example, the following command describes all attributes that are initializable during file creation:

```
@> describe mkfiles
```

You can further refine the output using the following options:

* `--mandatory`: Describes only the mandatory attributes that must be initialized during resource creation (`initializable = Mandatory`).

  For example:

  ```
  @> describe mkfiles --mandatory
  ```

* `--optional`: Describes only the optional attributes that may be initialized during resource creation (`initializable = Optional`).

  For example:

  ```
  @> describe mkfiles --optional
  ```

– If *resource-spec* begins with 'ch', the output contains only the attributes that are modifiable for an existing resource (`modifiable = true`).

  For example, the following command describes the attributes that may be modified for existing files:

  ```
  @> describe chfiles
  ```

- `--detail`: Optionally formats the command output so that values are displayed on separate lines, instead of using the default tabular output.

Usage Notes

For each resource, the output includes the resource name and a list of associated attributes. For each attribute, the following information is provided:

- `name`: The attribute name.
- `type`: The data type for the attribute value (for example: `String`, `Integer`, `Date`, `Boolean`, and so on).
- `description`: A brief description of the attribute, including valid values and default settings when applicable.
- `initializable`: Specifies whether the attribute value can or must be set during resource creation. Possible values are:
  – `No`: The attribute value cannot be set during resource creation. The value is set automatically by the system or inherits a default value.
  – `Mandatory`: The attribute value must be specified during resource creation.
  – `Optional`: The attribute value may be specified during resource creation. If not specified, the attribute is typically assigned a default value.
- `modifiable`: Indicates whether the attribute value can be modified after the resource is created (either `true` or `false`).

Copyright © 2022, 2026, Oracle and/or its affiliates.

6.2.3 Service and Cluster Management

This section contains references for the Exascale command line interface (ESCLI) commands that are associated with service and cluster management:

- cellcli
  Run a CellCLI command on a specified Exadata storage server.
- chcluster
  Shut down the Exascale cluster or modify its attributes.
- chclusteralert
  Change the `examinedBy` attribute for a cluster alert.
- chservice
  Start, stop, restart, or disable a software service.
- dbmcli
  Execute a DBMCLI command on a compute server.
- lscell
  List storage cells.
- lscelldisk
  List cell disks.
- lscluster
  List cluster information.
- lsclusteralert
  List information about cluster alerts.
- lscomputeserver
  List compute servers.
- lsgriddisk
  List grid disks.
- lsservice
  List software services.
- rmclusteralert
  Delete a cluster alert.

6.2.3.1 cellcli

Run a CellCLI command on a specified Exadata storage server.

Syntax

```
cellcli --cell cell-name [ --xml ] [ -e ] CellCLI-command
```

Command Options

The options for the `cellcli` command are:

- *cell-name*: Specifies the name of the cell that is the subject of the operation.

  The cell must be identified by its name. Use the lscell command to find all cell names.
- `--xml`: Generates XML-formatted command output.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `-e CellCLI-command`: Specifies the CellCLI command. The `-e` argument is optional.

Usage Notes

To run the `cellcli` command in ESCLI, the Exascale user must have the `cl_admin` privilege. On the Exadata storage server, the command runs with the privilege of the `celladmin` user.

Examples

Example 6-2 Run a CellCLI command

In the following examples, the `list cell` command is run on `CELL1`. Both commands perform the same function. However, the second command specifies the optional `-e` argument.

```
@> cellcli --cell CELL1 list cell


@> cellcli --cell CELL1 -e list cell
```

Example 6-3 Run a CellCLI command and generate XML-formatted output

In this example, the `list cell` command is run on `CELL1` and the command output is rendered using XML.

```
@> cellcli --cell CELL1 --xml list cell
```

6.2.3.2 chcluster

Shut down the Exascale cluster or modify its attributes.

Syntax

```
chcluster [ --attributes attribute=value[,attribute=value]... ]
          [ --shutdown ]
```

Command Options

The options for the `chcluster` command are:

- `--attributes`: Specifies one or more attribute settings for the cluster.

  Use the `describe chcluster` command to see details about the cluster attributes you can change. See also Describing Resources and Attributes.

- `--shutdown`: Shuts down the Exascale cluster.

  To shut down the Exascale cluster, note that you must:

  – Run ESCLI as the Exascale `admin` user or as another user with the `cl_admin` cluster level storage privilege.

  – Connect ESCLI directly to an online Exascale control services (ERS) back-end server process. To do this you must start ESCLI and specify the `--ctrl` option as follows:

    ```
    $ escli --wallet admin-wallet-location --ctrl ERS-server-IP:8080
    ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

To find the IP address for an online ERS server (*ERS-server-IP*), start a regular ESCLI session and use the following command:

```
@> lsservice --filter serviceType=controlServices,status=ONLINE --
attributes url
```

The command displays the private IP addresses associated with all of the online ERS servers. You can use any of the reported IP addresses.

Examples

Example 6-4 Modify the Cluster

This example shows setting the cluster attribute that specifies the Exascale control services (ERS) REST endpoint.

```
@> chcluster --attributes restEndPoint="198.51.100.1/24"
```

Example 6-5 Shut Down the Cluster

This example shows how to shut down the Exascale cluster using the specified wallet and Exascale control services (ERS) IP address.

```
$ escli --wallet /opt/oracle/cell/cellsrv/deploy/config/security/admwallet --
ctrl 198.51.100.1:8080 chcluster --shutdown
```

6.2.3.3 chclusteralert

Change the `examinedBy` attribute for a cluster alert.

Purpose

The `chclusteralert` command allows you to change the `examinedBy` attribute for an Exascale cluster alert.

Syntax

`chclusteralert` *`cluster-alert-name`* `[ --attributes ] examinedBy=`*`user-name`*

Command Options

The options for the `chclusteralert` command are:

- *cluster-alert-name*: Identifies the cluster alert that is the subject of the operation. You can use the lsclusteralert command to find the name for each cluster alert.
- *user-name*: Specifies the name of the user that has examined the cluster alert.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Examples

Example 6-6 Mark a Cluster Alert as Examined

The following example shows how to mark the specified cluster alert as examined by the `admin` user.

```
@> chclusteralert StoragePool:3d9c3c5f-c4411d4d-80ddfd82-
c90d84a5_StoragePoolRing:6291456_2026-01-15T00:05:35+00:00 --attributes
examinedBy=admin
```

6.2.3.4 chservice

Start, stop, restart, or disable a software service.

Purpose

The `chservice` command allows you to start, stop, restart, or disable an Exascale software service. A `chservice` command can only perform one operation on one service instance.

Syntax

```
chservice --attributes name=service_cell[,frontend={true|false}]
          { --start | --stop | --restart | --disable [ --force ]}


chservice --attributes name=service_compute
          { --start | --stop | --restart | --disable [ --force ]}
```

Command Options

The options for the `chservice` command are:

- `--attributes`: Primarily identifies the server and service that is the subject of the operation. This option is also used to specify optional attribute settings.
  - `name`: Identifies the server and service that is the subject of the operation.

    For a service on an Exadata storage server, specify the `name` attribute using the format: `name=service_cell`. In the attribute value:
    - *service*: Specifies the service that you want to act on, which is one of the following:
      - `cellsrv`: Specifies the Exadata Cell Server (CELLSRV).
      - `egs`: Specifies the Exascale cluster service, also known as Exascale global service (EGS).
      - `ers`: Specifies the Exascale control service, also known as Exascale RESTful service (ERS).
      - `syseds`: Specifies the system vault manager service.
      - `usreds`: Specifies the user vault manager service.
      - `bsm`: Specifies the block store manager (BSM) service.
      - `bsw`: Specifies the block store worker (BSW) service.
    - *cell*: Specifies the name of the storage server (`cell`) that hosts the service.

Copyright © 2022, 2026, Oracle and/or its affiliates.

The cell must be identified by its name. Use the lscell command to find all cell names.

For a service on an Exadata compute server, specify the `name` attribute using the format: `name=service_compute`. In the attribute value:

* *service*: Specifies the service that you want to act on, which is one of the following:
  * `egs`: Specifies the Exascale cluster service, also known as Exascale global service (EGS).
  * `bsw`: Specifies the block store worker service.
* *compute*: Specifies the name of the compute server (`dbserver`) that hosts the service.

  Depending on the required service location, specify either the name of a specific compute node virtual machine (VM) or the Exadata compute node host domain (KVM host or Dom0). Use the lscomputeserver command to display information about compute servers associated with the Exascale cluster.

– `frontend`: Optionally specifies whether to start ERS front-end server processes. This option is only valid when starting or restarting an ERS instance.

  To start the front-end server processes on the ERS instance, specify `frontend=true`. To not start the front-end server processes on the ERS instance, specify `frontend=false`.

- You must also specify one of the following operations:
  – `--start`: Starts a service. The service is also enabled if it was previously disabled.
  – `--stop`: Stops a service.
  – `--restart`: Restarts a service. The service is also enabled if it was previously disabled.
  – `--disable`: Disables a service. A disabled service is stopped and prevented from automatically restarting on the associated server.

    You can also add the `--force` option to forcibly disable the service, even if the server hosting the service is unavailable. This option only applies when disabling a service.

Examples

Example 6-7 Start the System Vault Manager Service (SYSEDS)

In this example, the system vault manager service is started on `CELL1`.

```
@> chservice --attributes name=syseds_CELL1 --start
```

Example 6-8 Restart the Exascale Control Service (ERS)

In this example, ERS is restarted on `CELL2`, along with the front-end ERS server processes.

```
@> chservice --attributes name=ers_CELL2,frontend=true --restart
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

6.2.3.5 dbmcli

Execute a DBMCLI command on a compute server.

Purpose

The `dbmcli` command enables you to run a DBMCLI command on a specified compute server.

Syntax

```
dbmcli --compute compute-name [ --xml ] [ -e ] dbmcli-command
```

Command Options

The options for the `dbmcli` command are:

- *compute-name*: Specifies the name of the compute server that is the subject of the operation.

  The compute server must be identified by its name. Use the lscomputeserver command to find all compute server names.
- `--xml`: Generates XML-formatted command output.
- `-e` *dbmcli-command*: Specifies the DBMCLI command. The `-e` argument is optional.

Usage Notes

To run the `dbmcli` command in ESCLI, the Exascale user must have the `cl_admin` privilege. On the Exadata compute server, the command runs with the privilege of the `dbmadmin` user.

Examples

Example 6-9 Run a DBMCLI command

In the following examples, the `list dbserver` command is run on `DB01`. Both commands perform the same function. However, the second command specifies the optional `-e` argument.

```
@> dbmcli --compute DB01 list dbserver
```

```
@> dbmcli --compute DB01 -e list dbserver
```

Example 6-10 Run a DBMCLI command and generate XML-formatted output

In this example, the `list dbserver` command is run on `DB01` and the command output is rendered using XML.

```
@> dbmcli --compute DB01 --xml list dbserver
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

6.2.3.6 lscell

List storage cells.

Purpose

The `lscell` command displays information about Exadata storage servers in the Exascale cluster.

Syntax

```
lscell [ -l ] [ --detail ] [ --attributes attribute[,attribute]... ]
       [ --filter filter[,filter]... ]
       [ --sort [-]attribute[,[-]attribute]... ]
       [ --count value ]
```

Command Options

The options for the `lscell` command are:

- `-l`: Displays output in a long, tabular format.
- `--detail`: Displays detailed output with additional attributes.
- `--attributes`: Identifies specific attributes to display in the output.

  To see a list of all available attributes and their descriptions, use the `describe cells` command. See also Describing Resources and Attributes.
- `--filter`: Specifies conditions for including items in the output.
- `--sort`: Sorts the output using the specified attributes.
- `--count`: Specifies the maximum number of items to display in the output.

Usage Notes

Note the following information when using this command:

- Filter conditions are specified as: <*attribute*><*operator*><*value*>.

  The allowed operators are =, !=, >=, <=, >, and <.

  Multiple comma-separated filter conditions are combined using AND logic.

  Dates can be specified using the following formats:

  - `yyyy-MM-dd''T''HH:mm:ss`
  - `yyyy-MM-dd` (Time is assumed to be 00:00 AM)
  - `HH:mm:ss` (Date is assumed to be today)

  A date can also be followed by a timezone specification.

  Sizes can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.
- Sorting attributes are specified as: `[-]`<*attribute*>. Multiple sort attributes are comma-delimited. The default sort order is ascending. For descending sort order, prefix the attribute name with `-`.

Copyright © 2022, 2026, Oracle and/or its affiliates.

For example, use the following to primarily sort by `name` in descending order, and use *otherAttribute* in ascending order to further sort entries with the same `name`:

```
--sort -name,otherAttribute
```

Examples

Example 6-11 List Storage Servers

The following example shows how to list essential information about all of the storage servers in the Exascale cluster.

```
@> lscell
```

Example 6-12 List Storage Servers

The following example shows how to list detailed information about all of the storage servers in the Exascale cluster.

```
@> lscell --detail
```

6.2.3.7 lscelldisk

List cell disks.

Purpose

The `lscelldisk` command displays information about Exadata storage server cell disks in the Exascale cluster.

Syntax

```
lscelldisk [ celldisk [ celldisk ]... ] [ -l ] [ --detail ]
           [ --attributes attribute[,attribute]... ]
           [ --filter filter[,filter]... ]
           [ --sort [-]attribute[,[-]attribute]... ]
           [ --count value ]
```

Command Options

The options for the `lscelldisk` command are:

- *celldisk*: Identifies an Exadata cell disk that you want to list information about. If not specified, the command displays information about all cell disks.
- `-l`: Displays output in a long, tabular format.
- `--detail`: Displays detailed output with additional attributes.
- `--attributes`: Identifies specific attributes to display in the output.

  To see a list of all available attributes and their descriptions, use the `describe celldisks` command. See also Describing Resources and Attributes.
- `--filter`: Specifies conditions for including items in the output.
- `--sort`: Sorts the output using the specified attributes.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `--count`: Specifies the maximum number of items to display in the output.

Usage Notes

Note the following information when using this command:

- Filter conditions are specified as: <*attribute*><*operator*><*value*>.

  The allowed operators are =, !=, >=, <=, >, and <.

  Multiple comma-separated filter conditions are combined using AND logic.

  Dates can be specified using the following formats:

  – `yyyy-MM-dd''T''HH:mm:ss`
  – `yyyy-MM-dd` (Time is assumed to be 00:00 AM)
  – `HH:mm:ss` (Date is assumed to be today)

  A date can also be followed by a timezone specification.

  Sizes can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.

- Sorting attributes are specified as: `[-]`<*attribute*>. Multiple sort attributes are comma-delimited. The default sort order is ascending. For descending sort order, prefix the attribute name with `-`.

  For example, use the following to primarily sort by `name` in descending order, and use *otherAttribute* in ascending order to further sort entries with the same `name`:

  ```
  --sort -name,otherAttribute
  ```

Examples

Example 6-13 List Cell Disk Information

The following example shows how to list detailed information for all cell disks in the Exascale cluster.

```
@> lscelldisk --detail
```

Example 6-14 List Information for Specific Cell Disks

The following example shows how to list information about specific cell disks named `DISK1`.

```
@> lscelldisk --filter name=DISK1
```

6.2.3.8 lscluster

List cluster information.

Purpose

The `lscluster` command displays information about the Exascale cluster.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Syntax

```
lscluster [ -l ] [ --detail ] [ --attributes attribute[,attribute]... ]
          [ --filter filter[,filter]... ]
          [ --sort [-]attribute[,[-]attribute]... ]
          [ --count value ] [ --backup ]
```

Command Options

The options for the `lscluster` command are:

- `-l`: Displays output in a long, tabular format.
- `--detail`: Displays detailed output with additional attributes.
- `--attributes`: Identifies specific attributes to display in the output.

  To see a list of all available attributes and their descriptions, use the `describe clusters` command. See also Describing Resources and Attributes.
- `--filter`: Specifies conditions for including items in the output.
- `--sort`: Sorts the output using the specified attributes.
- `--count`: Specifies the maximum number of items to display in the output.
- `--backup`: Limits the output to list only volume backup attributes.

Usage Notes

Note the following information when using this command:

- Filter conditions are specified as: <*attribute*><*operator*><*value*>.

  The allowed operators are =, !=, >=, <=, >, and <.

  Multiple comma-separated filter conditions are combined using AND logic.

  Dates can be specified using the following formats:

  – `yyyy-MM-dd''T''HH:mm:ss`
  – `yyyy-MM-dd` (Time is assumed to be 00:00 AM)
  – `HH:mm:ss` (Date is assumed to be today)

  A date can also be followed by a timezone specification.

  Sizes can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.
- Sorting attributes are specified as: `[-]`<*attribute*>. Multiple sort attributes are comma-delimited. The default sort order is ascending. For descending sort order, prefix the attribute name with `-`.

  For example, use the following to primarily sort by `name` in descending order, and use *otherAttribute* in ascending order to further sort entries with the same `name`:

  ```
  --sort -name,otherAttribute
  ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Examples

Example 6-15 List Cluster Information

The following example shows how to list essential information about the Exascale cluster.

```
@> lscluster
```

Example 6-16 List Cluster Information

The following example shows how to list detailed information about the Exascale cluster.

```
@> lscluster --detail
```

Example 6-17 List Cluster Information

The following example shows how to list detailed information about volume backup attributes that are associated with the Exascale cluster.

```
@> lscluster --detail --backup
```

6.2.3.9 lsclusteralert

List information about cluster alerts.

Purpose

The `lsclusteralert` command displays information about Exascale cluster alerts.

Syntax

```
lsclusteralert [ name [ name ]... | --openalerts]
               [ -l ] [ --detail ] [ --attributes attribute[,attribute]... ]
               [ --filter filter[,filter]... ]
               [ --sort [-]attribute[,[-]attribute]... ]
               [ --count value ]
```

Command Options

The options for the `lsclusteralert` command are:

- *name*: Identifies an Exascale cluster alert that you want to list information about.
- `--openalerts`: Restricts the command to only display information about open Exascale cluster alerts.
- `-l`: Displays output in a long, tabular format.
- `--detail`: Displays detailed output with additional attributes.
- `--attributes`: Identifies specific attributes to display in the output.

  To see a list of all available attributes and their descriptions, use the `describe clusteralerts` command. See also Describing Resources and Attributes.
- `--filter`: Specifies conditions for including items in the output.
- `--sort`: Sorts the output using the specified attributes.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `--count`: Specifies the maximum number of items to display in the output.

Usage Notes

Note the following information when using this command:

- If no cluster alert name is specified and the `--openalerts` option is not specified, the command displays information about all Exascale cluster alerts.
- Filter conditions are specified as: <*attribute*><*operator*><*value*>.

  The allowed operators are =, !=, >=, <=, >, and <.

  Multiple comma-separated filter conditions are combined using AND logic.

  Dates can be specified using the following formats:

  – `yyyy-MM-dd''T''HH:mm:ss`
  – `yyyy-MM-dd` (Time is assumed to be 00:00 AM)
  – `HH:mm:ss` (Date is assumed to be today)

  A date can also be followed by a timezone specification.

  Sizes can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.
- Sorting attributes are specified as: `[-]`<*attribute*>. Multiple sort attributes are comma-delimited. The default sort order is ascending. For descending sort order, prefix the attribute name with `-`.

  For example, use the following to primarily sort by `name` in descending order, and use *otherAttribute* in ascending order to further sort entries with the same `name`:

  ```
  --sort -name,otherAttribute
  ```

Examples

Example 6-18 List All Cluster Alerts

The following example shows how to list essential information about all Exascale cluster alerts.

```
@> lsclusteralert
```

Example 6-19 List All Open Cluster Alerts

The following example shows how to list essential information about all open Exascale cluster alerts.

```
@> lsclusteralert --openalerts
```

Example 6-20 List Information for a Specific Cluster Alert

The following example shows how to list detailed information about a specific cluster alert.

```
@> lsclusteralert StoragePool:3d9c3c5f-c4411d4d-80ddfd82-
c90d84a5_StoragePoolRing:6291456_2026-01-15T00:05:35+00:00 --detail
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

6.2.3.10 lscomputeserver

List compute servers.

Purpose

The `lscomputeserver` command displays information about compute servers associated with the Exascale cluster.

Syntax

```
lscomputeserver [ -l ] [ --detail ] [ --attributes attribute[,attribute]... ]
                [ --filter filter[,filter]... ]
                [ --sort [-]attribute[,[-]attribute]... ]
                [ --count value ]
```

Command Options

The options for the `lscomputeserver` command are:

- `-l`: Displays output in a long, tabular format.
- `--detail`: Displays detailed output with additional attributes.
- `--attributes`: Identifies specific attributes to display in the output.

  To see a list of all available attributes and their descriptions, use the `describe computes` command. See also Describing Resources and Attributes.
- `--filter`: Specifies conditions for including items in the output.
- `--sort`: Sorts the output using the specified attributes.
- `--count`: Specifies the maximum number of items to display in the output.

Usage Notes

Note the following information when using this command:

- Filter conditions are specified as: <*attribute*><*operator*><*value*>.

  The allowed operators are =, !=, >=, <=, >, and <.

  Multiple comma-separated filter conditions are combined using AND logic.

  Dates can be specified using the following formats:

  – `yyyy-MM-dd''T''HH:mm:ss`
  – `yyyy-MM-dd` (Time is assumed to be 00:00 AM)
  – `HH:mm:ss` (Date is assumed to be today)

  A date can also be followed by a timezone specification.

  Sizes can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.
- Sorting attributes are specified as: `[-]`<*attribute*>. Multiple sort attributes are comma-delimited. The default sort order is ascending. For descending sort order, prefix the attribute name with `-`.

Copyright © 2022, 2026, Oracle and/or its affiliates.

For example, use the following to primarily sort by `name` in descending order, and use *otherAttribute* in ascending order to further sort entries with the same `name`:

```
--sort -name,otherAttribute
```

Examples

Example 6-21 List Compute Servers

The following example shows how to list essential information about all of the compute servers associated with the Exascale cluster.

```
@> lscomputeserver
```

Example 6-22 List Compute Servers

The following example shows how to list detailed information about all of the compute servers associated with the Exascale cluster.

```
@> lscomputeserver --detail
```

6.2.3.11 lsgriddisk

List grid disks.

Purpose

The `lsgriddisk` command displays information about Exadata storage server grid disks in the Exascale cluster.

Syntax

```
lsgriddisk [ grid-disk [ grid-disk ]... ] [ -l ] [ --detail ]
           [ --attributes attribute[,attribute]... ]
           [ --filter filter[,filter]... ]
           [ --sort [-]attribute[,[-]attribute]... ]
           [ --count value ]
```

Command Options

The options for the `lsgriddisk` command are:

- *grid-disk*: Identifies an Exadata grid disk that you want to list information about. If not specified, the command displays information about all grid disks.
- `-l`: Displays output in a long, tabular format.
- `--detail`: Displays detailed output with additional attributes.
- `--attributes`: Identifies specific attributes to display in the output.

  To see a list of all available attributes and their descriptions, use the `describe griddisks` command. See also Describing Resources and Attributes.
- `--filter`: Specifies conditions for including items in the output.
- `--sort`: Sorts the output using the specified attributes.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `--count`: Specifies the maximum number of items to display in the output.

Usage Notes

Note the following information when using this command:

- Filter conditions are specified as: <*attribute*><*operator*><*value*>.

  The allowed operators are =, !=, >=, <=, >, and <.

  Multiple comma-separated filter conditions are combined using AND logic.

  Dates can be specified using the following formats:

  – `yyyy-MM-dd''T''HH:mm:ss`
  – `yyyy-MM-dd` (Time is assumed to be 00:00 AM)
  – `HH:mm:ss` (Date is assumed to be today)

  A date can also be followed by a timezone specification.

  Sizes can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.

- Sorting attributes are specified as: `[-]`<*attribute*>. Multiple sort attributes are comma-delimited. The default sort order is ascending. For descending sort order, prefix the attribute name with `-`.

  For example, use the following to primarily sort by `name` in descending order, and use *otherAttribute* in ascending order to further sort entries with the same `name`:

  ```
  --sort -name,otherAttribute
  ```

Examples

Example 6-23 List Grid Disk Information

The following example shows how to list detailed information for all grid disks in the Exascale cluster.

```
@> lsgriddisk --detail
```

Example 6-24 List Information for Specific Grid Disks

The following example shows how to list information about specific grid disks named `DISK2`.

```
@> lsgriddisk --filter name=DISK2
```

6.2.3.12 lsservice

List software services.

Purpose

The `lsservice` command displays information about software services in the Exascale cluster.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Syntax

```
lsservice [ -l ] [ --detail ] [ --attributes attribute[,attribute]... ]
          [ --filter filter[,filter]... ]
          [ --sort [-]attribute[,[-]attribute]... ]
          [ --count value ]
```

Command Options

The options for the `lsservice` command are:

- `-l`: Displays output in a long, tabular format.
- `--detail`: Displays detailed output with additional attributes.
- `--attributes`: Identifies specific attributes to display in the output.

  To see a list of all available attributes and their descriptions, use the `describe services` command. See also Describing Resources and Attributes.
- `--filter`: Specifies conditions for including items in the output.
- `--sort`: Sorts the output using the specified attributes.
- `--count`: Specifies the maximum number of items to display in the output.

Usage Notes

Note the following information when using this command:

- Filter conditions are specified as: <*attribute*><*operator*><*value*>.

  The allowed operators are =, !=, >=, <=, >, and <.

  Multiple comma-separated filter conditions are combined using AND logic.

  Dates can be specified using the following formats:

  – `yyyy-MM-dd''T''HH:mm:ss`
  – `yyyy-MM-dd` (Time is assumed to be 00:00 AM)
  – `HH:mm:ss` (Date is assumed to be today)

  A date can also be followed by a timezone specification.

  Sizes can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.
- Sorting attributes are specified as: `[-]`<*attribute*>. Multiple sort attributes are comma-delimited. The default sort order is ascending. For descending sort order, prefix the attribute name with `-`.

  For example, use the following to primarily sort by `name` in descending order, and use *otherAttribute* in ascending order to further sort entries with the same `name`:

  ```
  --sort -name,otherAttribute
  ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Examples

Example 6-25 List Service Information

The following example shows how to list essential information about all of the services in the Exascale cluster.

```
@> lsservice
```

Example 6-26 List System Vault Manager Service Information

The following example shows how to list detailed information about system vault manager services in the Exascale cluster.

```
@> lsservice --detail --filter serviceType=systemVaultManagers
```

6.2.3.13 rmclusteralert

Delete a cluster alert.

Purpose

The `rmclusteralert` command deletes an Exascale cluster alert matching the specified cluster alert name.

Syntax

```
rmclusteralert cluster-alert-name
```

Command Options

The options for the `rmclusteralert` command are:

- *cluster-alert-name*: Specifies the name of the cluster alert that is being deleted. You can use the lsclusteralert command to find the name for each cluster alert.

Examples

Example 6-27 Delete a Cluster Alert

This example shows deleting a cluster alert with the specified name.

```
@> rmclusteralert StoragePool:3d9c3c5f-c4411d4d-80ddfd82-
c90d84a5_StoragePoolRing:6291456_2026-01-15T00:05:35+00:00
```

6.2.4 Security and User Management

This section contains references for the Exascale command line interface (ESCLI) commands that are associated with security and user management:

- chacl
  Change an access control list (ACL).
- chuser
  Change attributes of a user, or upload API keys.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- chwallet
  Change authentication-related data in a local wallet.
- lsacl
  List ACLs for files and vaults.
- lskey
  Display key information.
- lsuser
  List user information.
- lswallet
  Display information stored in a wallet.
- mkkey
  Generate and write keys in PEM format.
- mkuser
  Create a new user.
- mkwallet
  Create a local wallet.
- rmuser
  Delete a user.

6.2.4.1 chacl

Change an access control list (ACL).

Purpose

The `chacl` command allows you to change the ACL for an Exascale vault or file.

Syntax

```
chacl  { file-name | vault-name } acl-string
```

Command Options

The options for the `chacl` command are:

- `{ file-name | vault-name }`: Specifies the name of the file or vault that is the subject of the operation. Only one file or vault name is permitted (no wildcards).
- *acl-string*: Specifies an ACL string having the following format.

  ```
  [+]user-ID1:acl-permission[+][;user-ID2:acl-permission[+]]...
  ```

  In the ACL string:

  - An optional plus (+) operator at the beginning of the ACL string indicates that the specified ACL string merges into the existing ACL. In this case, the changes only impact the users in the specified ACL string, and all other users in the existing ACL retain their permissions. Without this optional plus (+) operator, the ACL is completely overwritten by the specified ACL string, enabling permissions only for the users in the specified ACL string.
  - *user-IDn*: Specifies an Exascale user ID.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Depending on the user creation method, the user ID may be a system-generated value (for example, `96a68014-5762-4579-86ee-29eb743decbd`) or a user-specified value (for example, `scott`).

- *acl-permission*: Specifies an ACL permission setting, which can be one of the following:
  - `I | inspect`: Specifies that the user can view attributes of the file or vault, but not its contents.
  - `R | read`: Specifies that the user can read contents of the file, or list files in the vault. Also confers the `inspect` permission.
  - `U | use`: Specifies that the user can write to the file, or use the vault. Also confers all preceding permissions.
  - `M | manage`: Specifies that the user can manage the file or vault. Also confers all preceding permissions.
  - `0 | none`: Specifies that the user is removed from the ACL and loses all permissions. This setting can be used only in conjunction with the initial plus (+) operator to remove a user from an existing ACL.

  Starting with Oracle Exadata System Software release 26.1.0, an optional plus (+) operator following an *acl-permission* value guarantees that the user has at least that privilege level. For example, `R+` ensures the user has at least the `R` (`read`) privilege. If the user already has a higher privilege, such as `M` (`manage`), their privileges remain unchanged. You can only use this option when updating an ACL and you cannot use this option with the `0` (`none`) privilege.

Examples

Example 6-28 Replace a File ACL

In this example, the ACL string for the file `@my-vault/my-file` is replaced with the new ACL string. Under the new ACL, `scott` is permitted to read and inspect the file. Any other users included in the previous ACL effectively lose all previously held permissions.

```
@> chacl @my-vault/my-file scott:R
```

Example 6-29 Change a File ACL

In this example, the plus sign (+) at the beginning of the ACL string indicates that the specified ACL string is merged into the existing file ACL. In this case, any pre-existing permissions for `jason` are overwritten, and `jason` is now permitted to inspect, read, write, and manage the file. Also, `scott` is removed from the ACL and loses all previously held permissions. No other ACL permissions are changed.

```
@> chacl @my-vault/my-file +jason:manage;scott:none
```

Example 6-30 Replace a File ACL using an ACL String that Specifies Multiple Users

In this example, the ACL string for the file is replaced with the new ACL string that specifies permissions for multiple users. Under the new ACL, `scott` can inspect the file, and `jason` can read and inspect the file. Any other users included in the previous ACL effectively lose all previously held permissions.

```
@> chacl @my-vault/my-file scott:inspect;jason:read
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

6.2.4.2 chuser

Change attributes of a user, or upload API keys.

Purpose

The `chuser` command allows you to change an attribute for an existing Exascale user, or manage API keys for a user.

Syntax

```
chuser user-ID [ --attributes attribute=value[,attribute=value]... ]
       [ --privilege privileges ] [ --public-key-file1 public-key ]
       [ --public-key-file2 public-key ] [ --public-key-file3 public-key ]
       [ --rm-public-key1 ] [ --rm-public-key2 ] [ --rm-public-key3 ]
```

Command Options

The options for the `chuser` command are:

- *user-ID*: Identifies the unique user ID for the user that is the subject of the operation. You can use the lsuser command to find the ID for each user.
- `--attributes`: Optionally specifies values for attributes of the user.

  Use the `describe chuser` command to see details about the user attributes you can change. See also Describing Resources and Attributes.
- `--privilege`: Changes the user's privileges. The *privileges* value is a list of one or more privileges of the form `[+|-]privilege-1|privilege-2|...`.

  A privilege is one of the following: `no_privilege`, `cellsrv`, `egs`, `ers`, `syseds`, `usreds`, `bsm`, `bsw`, `ms`, `vlt_manage`, `vlt_use`, `vlt_read`, `vlt_inspect`, `cl_admin`, `cl_operator`, `cl_monitor`, `on_behalf_of`, `user_create`, `system_restore`, `rest_vault_client`, `rest_volume_client`. For descriptions, see Cluster Level Privileges.

  `no_privilege` cannot be combined with any other privilege; otherwise, an error is returned.

  Vault top-level privileges (`vlt_manage`, `vlt_use`, `vlt_read`, and `vlt_inspect`) are mutually exclusive. If any two or more are combined, an error is returned.

  Cluster privileges (`cl_admin`, `cl_operator`, and `cl_monitor`) are mutually exclusive. If any two or more are combined, an error is returned.

  An optional plus sign (+) at the beginning of the privilege string indicates that the specified privileges are added to the user's existing privileges. An optional minus sign (-) at the beginning of the privilege string indicates that the specified privileges are removed from the user's existing privileges. Otherwise, the user's existing privileges are overwritten by the specified privileges.
- `--public-key-file1`: Uploads a public key in slot 1.
- `--public-key-file2`: Uploads a public key in slot 2.
- `--public-key-file3`: Uploads a public key in slot 3.
- `--rm-public-key1`: Deletes the public key in slot 1.
- `--rm-public-key2`: Deletes the public key in slot 2.
- `--rm-public-key3`: Deletes the public key in slot 3.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Examples

Example 6-31 Change the Name of a User

In this example, the name of the specified user is changed to `SCOTTY`.

```
@> chuser 59f6dce4-5687-4728-8751-acfe99089be2 --attributes name=SCOTTY
```

Example 6-32 Assign the Vault Top-Level Read Privilege to a User

In this example, the vault top-level read privilege is assigned to the specified user.

```
@> chuser 59f6dce4-5687-4728-8751-acfe99089be2 --privilege vlt_read
```

Example 6-33 Upload a Public Key for a User

In this example, the public key contained in the file `new-pub.pem` is uploaded into public key slot 1 for the specified user.

```
@> chuser 59f6dce4-5687-4728-8751-acfe99089be2 --public-key-file1 new-pub.pem
```

Example 6-34 Delete a Public Key Associated With a User

In this example, the key in public key slot 1 is deleted for the specified user.

```
@> chuser 59f6dce4-5687-4728-8751-acfe99089be2 --rm-public-key1
```

6.2.4.3 chwallet

Change authentication-related data in a local wallet.

Purpose

The `chwallet` command changes ESCLI authentication-related data like the user identifier and private key in a local wallet.

Syntax

```
chwallet --wallet wallet-location
         [ --attributes attribute=value[,attribute=value]... ]
         [ --private-key-file pem-file [ --force ]]
         [ --fetch-trust-store ]
         [ --trusted-cert-file trust-file-name ]
         [ --clear-old-trusted-certs ]
         [ --private-key-remove ]
```

Command Options

The options for the `chwallet` command are:

- `--wallet`: Identifies the wallet that is the subject of the operation. *wallet-location* must be specified as a directory location or SSO file location. The wallet must already exist. A wallet can be created using the ESCLI mkwallet command.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `--attributes`: Optionally specifies values for attributes of the wallet.

  To see various attributes of a wallet, use the `lswallet` command with the `--detail` option.
- `--private-key-file`: Specifies a regular file location that contains a private key in PEM format.
- `--force`: Optionally replaces the private key in the wallet without confirmation. This option works in conjunction with the `--private-key-file` option.
- `--fetch-trust-store`: Gets the trusted certificates from the Exascale storage cluster and stores them in the wallet.
- `--trusted-cert-file`: Reads trusted certificates from the specified file and stores them in the wallet. You can use this option to read a PEM format file containing one trusted certificate. Or, you can use a file containing multiple trusted certificates, which is a concatenation of multiple PEM format files, each containing one trusted certificate.
- `--clear-old-trusted-certs`: Removes all previous trusted certificates from the wallet.

  This option is useful for removing old and expired trusted certificates.

  > **Note**
  >
  > This option also removes manually added trusted certificates and trusted certificates not issued by the trust store.
- `--private-key-remove`: Removes the private key from the wallet.

Usage Notes

Before running the `chwallet` command to modify the Exascale user associated with a wallet, shut down all clients using the wallet (including Exascale services, Oracle Database processes, Oracle Grid Infrastructure processes, and so on). If you modify the user associated with an active wallet, the associated client connections end abnormally.

This requirement applies only to modifying the Exascale user associated with a wallet. Other wallet operations, including key rotation and certificate updates, may occur while a wallet is in use.

Examples

Example 6-35 Set a Private Key and User Identifier in a Wallet

The following example shows setting the private key and user identifier inside the wallet at `/home/user/user.wallet`.

```
@> chwallet --wallet /home/user/user.wallet --private-key-file /home/user/
privatekey.pem --attributes user=0b9b8510-f88e-4f7b-ac57-10943c73dbe8
```

Example 6-36 Set the URL for the Exascale cluster services (EGS) in a Wallet

The following example shows setting the URL for the Exascale cluster services (EGS) inside the wallet at `/home/user/user.wallet`.

```
@> chwallet --wallet /home/user/user.wallet --attributes
exaRootUrl="egs=egsexc4:192.0.2.217:5045 egs=egsexc4:192.0.2.218:5045
egs=egsexc4:192.0.2.219:5045"
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Example 6-37 Fetch Trust Store Certificates to a Wallet

The following example gets the trusted certificates from the Exascale storage cluster and stores them in the specified wallet. The command also removes the certificates that previously resided in the wallet.

```
@> chwallet --wallet /home/user/user.wallet --fetch-trust-store --clear-old-
trusted-certs
```

Example 6-38 Load Trust Store Certificates to a Wallet

The following examples read the trusted certificates from the specified PEM files and store them in the specified wallet. Each command also removes the certificates that previously resided in the wallet.

Both examples are functionally equivalent assuming that `/tmp/concatenated-trust.pem` is a concatenation of `/tmp/trust1.pem`, `/tmp/trust2.pem`, and `/tmp/trust3.pem`.

```
@> chwallet --wallet /home/user/user.wallet --trusted-cert-file /tmp/
trust1.pem --trusted-cert-file /tmp/trust2.pem --trusted-cert-file /tmp/
trust3.pem --clear-old-trusted-certs
```

```
@> chwallet --wallet /root/eswallet/ --trusted-cert-file /tmp/concatenated-
trust.pem --clear-old-trusted-certs
```

Related Topics

- mkwallet
  Create a local wallet.

6.2.4.4 lsacl

List ACLs for files and vaults.

Purpose

The `lsacl` command allows you to list access control lists (ACLs) for Exascale files and vaults.

Syntax

```
lsacl  [ file-name | vault-name ]...
```

Usage Notes

Note the following information when using this command:

- *file-name* and *vault-name* specify the names of files or vaults that are the subject of the operation. You can specify multiple files or vaults for which you want to display ACLs. The asterisk (*) can be used for wildcard searches.
- If you do not specify any *file-name* or *vault-name*, then the output depends on the current working directory in the ESCLI session. If the current working directory is the root directory, then the command returns the ACLs for all vaults. If the current working directory is not the root directory, then the command returns the ACLs for all files under the current working directory.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Examples

Example 6-39 List ACLs for all vaults

This example lists the ACLs for all vaults.

```
@> lsacl
```

Example 6-40 List ACLs for all files in a vault

This example lists the ACLs for all files in the `@my-data` vault.

```
@> cd @my-data
@my-data/> lsacl
```

Example 6-41 List ACLs for a Vault and Files

This example lists the ACLs for the `@my-data` vault and the files that it contains.

```
@> lsacl @my-data @my-data/*
```

6.2.4.5 lskey

Display key information.

Purpose

The `lskey` command displays information about the specified key file.

Syntax

```
lskey pem-file [ --attributes attribute[,attribute]... ] [ --detail ] [ -l ]
```

Command Options

The options for the `lskey` command are:

- *pem-file*: Identifies the key file that you want to display information about. The *pem-file* must specify a regular file location for a key file in PEM format.
- `--attributes`: Identifies specific attributes to display in the output.

  You can use this option to focus on a subset of attributes as an alternative to the `--detail` option.
- `--detail`: Displays detailed output with additional attributes.
- `-l`: Displays output in a long, tabular format.

Examples

Example 6-42 List Detailed Key Information

The following example displays detailed information from the key file named `pub.pem`.

```
@> lskey pub.pem --detail
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Example 6-43 List Specific Key Information

The following example displays specific attributes (`id`, `privateKeyFile`, and `publicKeyFingerprint`) from the key file named `priv.pem`.

```
@> lskey priv.pem --attributes id,privateKeyFile,publicKeyFingerprint
```

6.2.4.6 lsuser

List user information.

Purpose

The `lsuser` command displays information about Exascale users.

Syntax

```
lsuser [ user-ID ] [ --attributes attribute[,attribute]... ]
       [ --filter filter[,filter]... ]
       [ --sort [-]attribute[,[-]attribute]... ]
       [ --count value ]
       [ --detail ] [ -l ]
```

Command Options

The options for the `lsuser` command are:

- *user-ID*: Identifies the user that you want to list information about. If no user is identified, the command displays information about all users.
- `--attributes`: Identifies specific attributes to display in the output.

  You can use this option to focus on a subset of attributes as an alternative to the `--detail` option.
- `--filter`: Specifies conditions for including items in the output.
- `--sort`: Sorts the output using the specified attributes.
- `--count`: Specifies the maximum number of items to display in the output.
- `--detail`: Displays detailed output with additional attributes.
- `-l`: Displays output in a long, tabular format.

Usage Notes

Note the following information when using this command:

- The `--filter`, `--sort`, and `--count` options are available starting with Oracle Exadata System Software release 26.1.0.
- Filter conditions are specified as: <*attribute*><*operator*><*value*>.

  The allowed operators are =, !=, >=, <=, >, and <.

  Multiple comma-separated filter conditions are combined using AND logic.

  Dates can be specified using the following formats:

  - `yyyy-MM-dd''T''HH:mm:ss`
  - `yyyy-MM-dd` (Time is assumed to be 00:00 AM)

Copyright © 2022, 2026, Oracle and/or its affiliates.

  - `HH:mm:ss` (Date is assumed to be today)

  A date can also be followed by a timezone specification.

  Sizes can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.

- Sorting attributes are specified as: `[-]<attribute>`. Multiple sort attributes are comma-delimited. The default sort order is ascending. For descending sort order, prefix the attribute name with `-`.

  For example, use the following to primarily sort by `name` in descending order, and use *otherAttribute* in ascending order to further sort entries with the same `name`:

  ```
  --sort -name,otherAttribute
  ```

Examples

Example 6-44 List Information for a Specific User

The following example shows how to list detailed information about the user ID: `SCOTT0123`.

```
@> lsuser SCOTT0123 --detail
```

6.2.4.7 lswallet

Display information stored in a wallet.

Purpose

The `lswallet` command displays ESCLI authentication data that is stored in a wallet.

Syntax

```
lswallet [ --wallet wallet-location ] [ --detail ]
         [ --private-key-file pem-file ]
         [ --public-key-file pem-file ]
         [ --public-key-file8 pem-file ]
         [ --attributes attribute[,attribute]... ]
         [ --trusted-cert-file cert-file ]
```

Command Options

The options for the `lswallet` command are:

- `--wallet`: Identifies the wallet that you want to display information about. The *wallet-location* must be specified as a directory location or SSO file location. If not specified, the command displays information about the wallet used to authenticate the current ESCLI session.
- `--detail`: Displays detailed output with additional attributes, including the user ID and fingerprint of the public key from the wallet.
- `--private-key-file`: Extracts the private key from the wallet and stores it in the specified PEM file.
- `--public-key-file`: Extracts the public key from the wallet and stores it in the specified PEM file.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `--public-key-file8`: Extracts the public key from the wallet in PKCS8 format and stores it in the specified PEM file.
- `--attributes`: Identifies specific attributes to display in the output.

  You can use this option to focus on a subset of attributes as an alternative to the `--detail` option.
- `--trusted-cert-file`: Extracts the trusted certificates from the wallet and saves them in PEM format files.

  Files are created with numeric suffixes `1`, `2`, `3`, and so on. For example, if you specify `--trusted-cert-file cert_file`, then the certificates are stored in the files `cert_file1`, `cert_file2`, `cert_file3`, and so on.

Examples

Example 6-45 List Wallet Information

The following example displays essential information from the wallet used to authenticate the current ESCLI session.

```
@> lswallet
```

Example 6-46 List Detailed Wallet Information

The following example displays detailed information from the wallet used to authenticate the current ESCLI session.

```
@> lswallet --detail
```

Example 6-47 Extract a Private Key

The following example extracts the private key from the wallet used to authenticate the current ESCLI session. The private key is stored in `/home/user/privatekey.pem`

```
@> lswallet --private-key-file /home/user/privatekey.pem
```

Example 6-48 Extract a Public Key

The following example extracts the public key from the wallet used to authenticate the current ESCLI session. The public key is stored in `/home/user/publickey.pem`

```
@> lswallet --public-key-file /home/user/publickey.pem
```

Example 6-49 Extract Certificates to Files

The following example extracts the trusted certificates from the wallet used to authenticate the current ESCLI session. The certificates are stored in regular files `/home/user/cert1`, `/home/user/cert2`, `/home/user/cert3`, and so on.

```
@> lswallet --trusted-cert-file /home/user/cert
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

6.2.4.8 mkkey

Generate and write keys in PEM format.

Purpose

The `mkkey` command generates and writes a private key to a regular file in PEM format. You can also generate and write the corresponding public key to a regular file in PEM format.

Syntax

```
mkkey [ --private-key-file pem-file ] [ --public-key-file pem-file ]
      [ --attributes attribute=value[,attribute=value]... ]
```

Command Options

The options for the `mkkey` command are:

- `--private-key-file`: Specifies a regular file location for the private key.
- `--public-key-file`: Optionally specifies a regular file location for the corresponding public key.
- `--attributes`: Optionally specifies attributes to change. This option provides an alternative method for specifying the key file names.

Examples

Example 6-50 Generate a Key Pair

The following examples generate a key pair and store the private and public keys in `/home/user/privatekey.pem` and `/home/user/publickey.pem`, respectively.

```
@> mkkey --private-key-file /home/user/privatekey.pem --public-key-file /home/
user/publickey.pem
```

```
@> mkkey --attributes privateKeyFile=/home/user/privatekey.pem,publicKeyFile=/
home/user/publickey.pem
```

6.2.4.9 mkuser

Create a new user.

Purpose

The `mkuser` command allows you to create a new Exascale user with the specified user name.

Syntax

```
mkuser user-name [ --attributes attribute=value[,attribute=value]... ]
```

Command Options

The options for the `mkuser` command are:

Copyright © 2022, 2026, Oracle and/or its affiliates.

- *user-name*: Specifies the name of the user. The user name must start with an alphanumeric character, and can only contain alphanumeric characters, hyphens (-) and periods (.).
- `--attributes`: Optionally specifies attributes to set.

  Use the `describe mkuser` command to see details about the user attributes you can set. See also Describing Resources and Attributes.

Usage Notes

Note the following information when using this command:

- You can set the user identifier by specifying the `id` attribute.

  If specified, the user identifier must start with an alphanumeric character, and can only contain alphanumeric characters, hyphens (-) and periods (.).

  If none is specified, then a unique user ID is generated automatically and returned by the command. The user ID value is required when managing users using chuser, rmuser, and lsuser.

  The user ID is fixed at user creation time. There is no way to modify the user ID afterward.
- You can set the user's privileges by specifying the `privilege` attribute.

  The `privilege` value is a list of one or more privileges of the form *`privilege-1|privilege-2|...`*.

  A privilege is one of the following: `no_privilege`, `cellsrv`, `egs`, `ers`, `syseds`, `usreds`, `bsm`, `bsw`, `ms`, `vlt_manage`, `vlt_use`, `vlt_read`, `vlt_inspect`, `cl_admin`, `cl_operator`, `cl_monitor`, `on_behalf_of`, `user_create`, `system_restore`, `rest_vault_client`, `rest_volume_client`. For descriptions, see Cluster Level Privileges.

  `no_privilege` cannot be combined with any other privilege; otherwise, an error is returned.

  Vault top-level privileges (`vlt_manage`, `vlt_use`, `vlt_read`, and `vlt_inspect`) are mutually exclusive. If any two or more are combined, an error is returned.

  Cluster privileges (`cl_admin`, `cl_operator`, and `cl_monitor`) are mutually exclusive. If any two or more are combined, an error is returned.

  If the `privilege` attribute is not specified, then the user is created and assigned only the `vlt_inspect` privilege by default.

Examples

Example 6-51 Create a User

This example shows creating a user named `CHERIE`.

```
@> mkuser CHERIE
```

Example 6-52 Create a User with a Specific User Identifier

This example shows creating a user named `SCOTT` using `SCOTT0123` as the user identifier.

```
@> mkuser SCOTT --attributes id=SCOTT0123
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Example 6-53 Create a User with Specific Privileges

This example shows creating a user named `PETER` with a specific set of privileges.

```
@> mkuser PETER --attributes privilege=vlt_read|ers
```

6.2.4.10 mkwallet

Create a local wallet.

Purpose

The `mkwallet` command creates a new Oracle wallet at the specified location.

Syntax

```
mkwallet --wallet wallet-location
```

Command Options

The options for the `mkwallet` command are:

- `--wallet`: Specifies the location where the new wallet is created. *wallet-location* must be specified as a directory location or SSO file location.

Examples

Example 6-54 Create a Wallet

The following example shows creating a wallet at `/home/user/user.wallet`.

```
@> mkwallet --wallet /home/user/user.wallet
```

6.2.4.11 rmuser

Delete a user.

Purpose

The `rmuser` command deletes an existing Exascale user.

Syntax

```
rmuser user-ID
```

Command Options

The options for the `rmuser` command are:

- *user-ID*: Specifies the unique user ID for the user that is being deleted. You can use the lsuser command to find the ID for each user.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Examples

Example 6-55 Delete a User

This example shows deleting the user with the ID `0b9b8510-f88e-4f7b-ac57-10943c73dbe8`.

```
@> rmuser 0b9b8510-f88e-4f7b-ac57-10943c73dbe8
```

6.2.5 Storage Pool and Pool Disk Management

This section contains references for the Exascale command line interface (ESCLI) commands that are associated with storage pool and pool disk management:

- chpooldisk
  Modify a pool disk.
- chstoragepool
  Modify a storage pool.
- lspooldisk
  List pool disks.
- lsstoragepool
  List storage pools.
- lsstoragepooloperation
  List ongoing storage pool operations.
- mkstoragepool
  Create a storage pool.
- rmstoragepool
  Delete a storage pool.

6.2.5.1 chpooldisk

Modify a pool disk.

Purpose

The `chpooldisk` command enables you to modify an Exascale pool disk.

Syntax

```
chpooldisk pooldisk-ID { --offline | --online }
```

Command Options

The options for the `chpooldisk` command are:

- *pooldisk-ID*: Identifies the Exadata pool disk that you want to modify. You can use the lspooldisk command to find the ID for each pool disk.
- `--offline`: Takes the pool disk offline.
- `--online`: Brings the pool disk online.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Examples

Example 6-56 Bring a Pool Disk Online

The following example shows how to bring a pool disk online.

```
@> chpooldisk 10 --online
```

6.2.5.2 chstoragepool

Modify a storage pool.

Purpose

The `chstoragepool` command enables you to modify an Exascale storage pool.

Syntax

```
chstoragepool storagepool-name
              [ --attributes attribute=value[,attribute=value]... ]
              [ --reconfig [ --force ]]
```

Command Options

The options for the `chstoragepool` command are:

- *storagepool-name*: Identifies an Exascale storage pool that you want to modify.
- `--attributes`: Optionally specifies one or more attribute settings for the storage pool.

  Use the `describe chstoragepool` command to see details about the storage pool attributes you can change. See also Describing Resources and Attributes.
- `--reconfig`: Reconfigures the storage pool and makes permanent any changes to the number or size of the underlying pool disks. Reconfiguration may trigger a rebalance operation to ensure that data is spread evenly across the pool disks in the storage pool. The time required for reconfiguration depends on the amount of data movement required to rebalance the storage pool.
- `--force`: Optionally forces a reconfiguration operation even if the system detects no apparent change to the underlying pool disks.

Examples

Example 6-57 Set the Disk Offline Timer for a Storage Pool

The following example shows how to set the disk offline timer attribute associated with the Exascale storage pool named `my-pool1`.

```
@> chstoragepool my-pool1 --attributes diskOfflineTimerInMins=60
```

Example 6-58 Rediscover Pool Disks and Reconfigure the Storage Pool

The following example shows how to forcibly re-discover the pool disks and reconfigure the specified storage pool.

```
@> chstoragepool my-pool1 --reconfig --force
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

6.2.5.3 lspooldisk

List pool disks.

Purpose

The `lspooldisk` command displays information about Exadata storage server pool disks in the Exascale cluster.

Syntax

```
lspooldisk [ pooldisk-ID [ pooldisk-ID ]... ] [ -l ] [ --detail ]
           [ --attributes attribute[,attribute]... ]
           [ --filter filter[,filter]... ]
           [ --sort [-]attribute[,[-]attribute]... ]
           [ --count value ]
```

Command Options

The options for the `lspooldisk` command are:

- *pooldisk-ID*: Identifies an Exadata pool disk that you want to list information about. If not specified, the command displays information about all pool disks.
- `-l`: Displays output in a long, tabular format.
- `--detail`: Displays detailed output with additional attributes.
- `--attributes`: Identifies specific attributes to display in the output.

  To see a list of all available attributes and their descriptions, use the `describe pooldisks` command. See also Describing Resources and Attributes.
- `--filter`: Specifies conditions for including items in the output.
- `--sort`: Sorts the output using the specified attributes.
- `--count`: Specifies the maximum number of items to display in the output.

Usage Notes

Note the following information when using this command:

- Filter conditions are specified as: <*attribute*><*operator*><*value*>.

  The allowed operators are =, !=, >=, <=, >, and <.

  Multiple comma-separated filter conditions are combined using AND logic.

  Dates can be specified using the following formats:

  – `yyyy-MM-dd''T''HH:mm:ss`

  – `yyyy-MM-dd` (Time is assumed to be 00:00 AM)

  – `HH:mm:ss` (Date is assumed to be today)

  A date can also be followed by a timezone specification.

  Sizes can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- Sorting attributes are specified as: `[-]<attribute>`. Multiple sort attributes are comma-delimited. The default sort order is ascending. For descending sort order, prefix the attribute name with `-`.

  For example, use the following to primarily sort by `name` in descending order, and use *otherAttribute* in ascending order to further sort entries with the same `name`:

  ```
  --sort -name,otherAttribute
  ```

Examples

Example 6-59 List Pool Disk Information

The following example shows how to list detailed information for all pool disks in the Exascale cluster.

```
@> lspooldisk --detail
```

Example 6-60 List Information for Specific Pool Disks

The following example shows how to list information about specific pool disks named `DISK3`.

```
@> lspooldisk --filter name=DISK3
```

6.2.5.4 lsstoragepool

List storage pools.

Purpose

The `lsstoragepool` command displays information about Exascale storage pools.

Syntax

```
lsstoragepool [ storagepool-name [ storagepool-name ]... ]
              [ -l ] [ --detail ]
              [ --attributes attribute[,attribute]... ]
              [ --filter filter[,filter]... ]
              [ --sort [-]attribute[,[-]attribute]... ]
              [ --count value ]
```

Command Options

The options for the `lsstoragepool` command are:

- *storagepool-name*: Identifies an Exascale storage pool that you want to list information about. If not specified, the command displays information about all storage pools.
- `-l`: Displays output in a long, tabular format.
- `--detail`: Displays detailed output with additional attributes.
- `--attributes`: Identifies specific attributes to display in the output.

  To see a list of all available attributes and their descriptions, use the `describe storagepools` command. See also Describing Resources and Attributes.
- `--filter`: Specifies conditions for including items in the output.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `--sort`: Sorts the output using the specified attributes.
- `--count`: Specifies the maximum number of items to display in the output.

Usage Notes

Note the following information when using this command:

- Filter conditions are specified as: <*attribute*><*operator*><*value*>.

  The allowed operators are =, !=, >=, <=, >, and <.

  Multiple comma-separated filter conditions are combined using AND logic.

  Dates can be specified using the following formats:

  – `yyyy-MM-dd''T''HH:mm:ss`
  – `yyyy-MM-dd` (Time is assumed to be 00:00 AM)
  – `HH:mm:ss` (Date is assumed to be today)

  A date can also be followed by a timezone specification.

  Sizes can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.

- Sorting attributes are specified as: `[-]`<*attribute*>. Multiple sort attributes are comma-delimited. The default sort order is ascending. For descending sort order, prefix the attribute name with `-`.

  For example, use the following to primarily sort by `name` in descending order, and use *otherAttribute* in ascending order to further sort entries with the same `name`:

  ```
  --sort -name,otherAttribute
  ```

Examples

Example 6-61 List Storage Pool Information

The following example shows how to list information for all storage pools in the Exascale cluster.

```
@> lsstoragepool
```

Example 6-62 List Information for a Specific Storage Pool

The following example shows how to list detailed information about the storage pool named `my-pool1`.

```
@> lsstoragepool my-pool1 --detail
```

6.2.5.5 lsstoragepooloperation

List ongoing storage pool operations.

Purpose

The `lsstoragepooloperation` command displays information about storage pool operations that are currently active in the Exascale cluster.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Syntax

```
lsstoragepooloperation [ operation-ID [ operation-ID ]... ]
                       [ -l ] [ --detail ]
                       [ --attributes attribute[,attribute]... ]
                       [ --filter filter[,filter]... ]
                       [ --sort [-]attribute[,[-]attribute]... ]
                       [ --count value ]
```

Command Options

The options for the `lsstoragepooloperation` command are:

- *operation-ID*: Identifies an Exascale storage pool operation that you want to list information about. If not specified, the command displays information about all active operations.
- `-l`: Displays output in a long, tabular format.
- `--detail`: Displays detailed output with additional attributes.
- `--attributes`: Identifies specific attributes to display in the output.

  To see a list of all available attributes and their descriptions, use the `describe storagepooloperations` command. See also Describing Resources and Attributes.
- `--filter`: Specifies conditions for including items in the output.
- `--sort`: Sorts the output using the specified attributes.
- `--count`: Specifies the maximum number of items to display in the output.

Usage Notes

Note the following information when using this command:

- Filter conditions are specified as: <*attribute*><*operator*><*value*>.

  The allowed operators are =, !=, >=, <=, >, and <.

  Multiple comma-separated filter conditions are combined using AND logic.

  Dates can be specified using the following formats:

  – `yyyy-MM-dd''T''HH:mm:ss`
  – `yyyy-MM-dd` (Time is assumed to be 00:00 AM)
  – `HH:mm:ss` (Date is assumed to be today)

  A date can also be followed by a timezone specification.

  Sizes can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.
- Sorting attributes are specified as: `[-]`<*attribute*>. Multiple sort attributes are comma-delimited. The default sort order is ascending. For descending sort order, prefix the attribute name with `-`.

  For example, use the following to primarily sort by `name` in descending order, and use *otherAttribute* in ascending order to further sort entries with the same `name`:

  ```
  --sort -name,otherAttribute
  ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Examples

Example 6-63 List Storage Pool Operations

The following example shows how to list information for all ongoing storage pool operations in the Exascale cluster.

```
@> lsstoragepooloperation
```

Example 6-64 List Information for a Specific Operation

The following example shows how to list detailed information about a specific operation.

```
@> lsstoragepooloperation 3192092 --detail
```

6.2.5.6 mkstoragepool

Create a storage pool.

Purpose

The `mkstoragepool` command creates a pool of Exascale storage.

Syntax

```
mkstoragepool storagepool-name
              [ --attributes attribute=value[,attribute=value]... ]
              [ --nowait ]
```

Command Options

The options for the `mkstoragepool` command are:

- `--attributes`: Optionally specifies one or more attribute settings for the storage pool.

  Use the `describe mkstoragepool` command to see details about the storage pool attributes you can set. See also Describing Resources and Attributes.

- `--nowait`: Specifies that the command returns immediately while the operation completes in the background. Without this option, the command does not return until the storage pool is created.

Usage Notes

Before you create a storage pool, you need some pool disks that are already associated with the storage pool name.

Examples

Example 6-65 Create a Storage Pool

The following example shows creating a storage pool named `my-pool1`.

```
@> mkstoragepool my-pool1
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

6.2.5.7 rmstoragepool

Delete a storage pool.

Purpose

The `rmstoragepool` command enables you to delete an Exascale storage pool.

Syntax

```
rmstoragepool storagepool-name
```

Command Options

The options for the `rmstoragepool` command are:

- *storagepool-name*: Identifies the Exascale storage pool that you want to delete.

Usage Notes

You can use the `rmstoragepool` command in cases where you want to:

- Remove storage hardware from the Exascale cluster.
- Migrate data to another storage pool.

Examples

Example 6-66 Delete a Storage Pool

The following example shows how to delete a storage pool named `my-pool1`.

```
@> rmstoragepool my-pool1
```

6.2.6 Vault Management

This section contains references for the Exascale command line interface (ESCLI) commands that are associated with vault management:

- chvault
  Change the attributes of a vault and provision storage pool resources.
- lsvault
  List vaults.
- mkvault
  Create a vault.
- rmvault
  Delete a vault.

Copyright © 2022, 2026, Oracle and/or its affiliates.

6.2.6.1 chvault

Change the attributes of a vault and provision storage pool resources.

Purpose

The `chvault` command modifies the attributes of a vault and enables provisioning of storage pool resources.

Syntax

```
chvault [@]vault --attributes attribute=value[,attribute=value]...
```

```
chvault [@]vault --provision storagePool=storagepool-name[,provisioning-attribute=value]...
```

Command Options

The options for the `chvault` command are:

- *vault*: Identifies the name of the vault that you want to change.
- `--attributes`: Specifies attributes to change.

  Use the `describe chvault` command to see details about the vault attributes you can change. See also Describing Resources and Attributes.
- `--provision`: Provisions the specified vault with resources from a specific storage pool.

  To provision a vault with resources from multiple storage pools, you must use a separate command for each storage pool.

  This option is available starting with Oracle Exadata System Software release 26.1.0, and may not be used in conjunction with any other options.

Usage Notes

Note the following information when using this command:

- The command only changes the specified attributes. Other settings are unaffected.
- Starting with Oracle Exadata System Software release 26.1.0, you can create multiple storage pools for each media type within a single Exascale cluster and you can provision a vault with a specific set of storage resources from each storage pool.

  The following list describes these attributes, which may be specified using the `--provision` option. See also Provision a Vault Using Multiple Storage Pools For Each Media Type.

  - `spaceProv`: Provisions the vault with the specified amount of storage space.
  - `iopsProv`: Provisions the vault with the specified number of IOPS.
  - `flashCacheProv`: Provisions the vault with the specified amount of flash cache space.
  - `xrmemCacheProv`: Provisions the vault with the specified amount of Exadata RDMA Memory Cache (XRMEM cache) space.
  - `flashLogProv`: Boolean value (`true` or `false`) indicating whether the vault is provisioned with access to Exadata Smart Flash Log.

Copyright © 2022, 2026, Oracle and/or its affiliates.

  - `xrmemLogProv`: Boolean value (`true` or `false`) indicating whether the vault is provisioned with access to XRMEM log (if available).

- If your Exascale cluster is configured with Oracle Exadata System Software before release 26.1.0, you can specify one set of resource provisioning attributes, covering all media types, that is linked directly to the vault.

  The following list describes these attributes, which may be changed using the `--attributes` option. See also Provision a Vault Using Only One Storage Pool For Each Media Type.

  - `spaceProvEF`, `spaceProvHC`, `spaceProvXT`: Provisions the vault with the specified amount of extreme flash (EF), high capacity (HC), or extended (XT) storage space.
  - `iopsProvEF`, `iopsProvHC`, `iopsProvXT`: Provisions the vault with the specified number of IOPS from extreme flash (EF), high capacity (HC), or extended (XT) storage.
  - `flashCacheProv`: Provisions the vault with the specified amount of flash cache space.
  - `flashLogProv`: Boolean value (`true` or `false`) indicating whether the vault is provisioned with access to Exadata Smart Flash Log.
  - `xrmemCacheProv`: Provisions the vault with the specified amount of Exadata RDMA Memory Cache (XRMEM cache) space.

    Use this attribute only on systems with `HC` or `EF` storage, but not both. For systems with `HC` and `EF` storage, use the following media-specific XRMEM cache provisioning attributes.
  - `xrmemCacheProvEF`: Provisions the vault with the specified amount of XRMEM cache space associated with `EF` media.
  - `xrmemCacheProvHC`: Provisions the vault with the specified amount of XRMEM cache space associated with `HC` media.

- To specify an amount of space, you can:
  - Use an integer value representing the number of bytes. The maximum value is `4503599626321920`.
  - Specify a space value using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.
  - Use `unlimited` to specify an unlimited amount of space.
- To specify a number of IOPS, you can:
  - Use an integer value. The maximum value is `4294967295`.
  - Use `unlimited` to specify an unlimited number of IOPS.
- Specifying `0` (zero) as the value for a cache provisioning attribute (`flashCacheProv`, `xrmemCacheProv`, `xrmemCacheProvEF`, or `xrmemCacheProvHC`) disables use of the corresponding cache for the vault.

Related Topics

- Provision a Vault

Copyright © 2022, 2026, Oracle and/or its affiliates.

6.2.6.2 lsvault

List vaults.

Purpose

The `lsvault` command displays information about Exascale vaults.

Syntax

```
lsvault [ [@]name [ [@]name ]... ] [ -l ] [ --detail ]
   [ --attributes attribute[,attribute]... ]
   [ --filter filter[,filter]... ]
   [ --sort [-]attribute[,[-]attribute]... ]
   [ --sum attribute ] [ --count value ] [ -t ]


lsvault [@]vault --provision
```

Command Options

The options for the `lsvault` command are:

- *name*: A vault name that you want to display information for.

  In this context, the asterisk (*) can be used for wildcard searches.

  If no vault name is specified, then information is displayed about all vaults in the current level of the ESCLI hierarchy.
- `-l`: Displays output in a long, tabular format.
- `--detail`: Displays detailed output with additional attributes.
- `--attributes`: Identifies specific attributes to display in the output.

  You can use this option to focus on a subset of attributes as an alternative to the `--detail` option.
- `--filter`: Specifies conditions for including items in the output.
- `--sort`: Sorts the output using the specified attributes.
- `--sum`: Calculates the sum of the specified numeric attribute.
- `--count`: Specifies the maximum number of items to display in the output.
- `-t`: Sorts vaults by ascending creation time.
- `--provision`: Displays information about storage pool resources provisioned to the specified vault.

  This option is available starting with Oracle Exadata System Software release 26.1.0, and may not be used in conjunction with any other options.

Usage Notes

Note the following information when using this command:

- Filter conditions are specified as: <*attribute*><*operator*><*value*>.

  The allowed operators are =, !=, >=, <=, >, and <.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Multiple comma-separated filter conditions are combined using AND logic.

Dates can be specified using the following formats:

- `yyyy-MM-dd''T''HH:mm:ss`
- `yyyy-MM-dd` (Time is assumed to be 00:00 AM)
- `HH:mm:ss` (Date is assumed to be today)

A date can also be followed by a timezone specification.

Sizes can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.

- Sorting attributes are specified as: `[-]<attribute>`. Multiple sort attributes are comma-delimited. The default sort order is ascending. For descending sort order, prefix the attribute name with `-`.

  For example, use the following to primarily sort by `name` in descending order, and use *otherAttribute* in ascending order to further sort entries with the same `name`:

  ```
  --sort -name,otherAttribute
  ```

- You can use the `--sum` option only by itself or together with the `--filter` option.

  The `--sum` option is available starting with Oracle Exadata System Software release 26.1.0.

Examples

Example 6-67 Display All Vaults

You can use the `lsvault` command at the root level to get a listing of all vaults.

```
@> lsvault
my-data
vault2
```

Example 6-68 Display Vault Information in a Long Format

When you use the `-l` option, the output includes additional information about each vault.

```
@> lsvault -l
createTime                spaceUsedEF spaceUsedHC name
2024-11-24T01:41:41+00:00     0.0000     10.0000G my-data
2024-11-25T05:47:11+00:00     0.0000    100.0000M vault2
```

Example 6-69 Display Vault Provisioning Information

Use the `--provision` option to show vault provisioning information for multiple storage pools.

```
@> lsvault @my-data --provision

id          spaceProv  spaceUsed  iopsProv   flashCacheProv  xrmemCacheProv
mediaType  flashLogProv  xrmemLogProv
my-pool1    1.0000T    10.0000G   32000      unlimited       10.0000G
HC         enabled       disabled
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
my-pool2    unlimited  0.0000     unlimited  unlimited       unlimited
HC         enabled       disabled
```

6.2.6.3 mkvault

Create a vault.

Purpose

The `mkvault` command creates an Exascale vault.

Syntax

```
mkvault [@]vault [ --attributes attribute=value[,attribute=value]... ]
```

Command Options

The options for the `mkvault` command are:

- *vault*: Specifies the name of the vault.
- `--attributes`: Optionally specifies attributes to set.

  Use the `describe mkvault` command to see details about the vault attributes you can set. See also Describing Resources and Attributes.

Usage Notes

Note the following information when using this command:

- Starting with Oracle Exadata System Software release 26.1.0, you can create multiple storage pools for each media type within a single Exascale cluster and you can provision a vault with a specific set of storage resources from each storage pool.

  To use this capability, do not specify any resource provisioning using the `mkvault` command. Rather, create the vault using the `mkvault` command, and then complete provisioning by using a separate chvault command for each storage pool that you want to associate with the vault.

  See also Provision a Vault Using Multiple Storage Pools For Each Media Type.
- If your Exascale cluster is configured with Oracle Exadata System Software before release 26.1.0, you can specify one set of resource provisioning attributes, covering all media types, that is linked directly to the vault.

  The following list describes these attributes, which may be set using the `--attributes` option. See also Provision a Vault Using Only One Storage Pool For Each Media Type.

  - `spaceProvEF`, `spaceProvHC`, `spaceProvXT`: Provisions the vault with the specified amount of extreme flash (EF), high capacity (HC), or extended (XT) storage space.
  - `iopsProvEF`, `iopsProvHC`, `iopsProvXT`: Provisions the vault with the specified number of IOPS from extreme flash (EF), high capacity (HC), or extended (XT) storage.
  - `flashCacheProv`: Provisions the vault with the specified amount of flash cache space.
  - `flashLogProv`: Boolean value (`true` or `false`) indicating whether the vault is provisioned with access to Exadata Smart Flash Log.
  - `xrmemCacheProv`: Provisions the vault with the specified amount of Exadata RDMA Memory Cache (XRMEM cache) space.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Use this attribute only on systems with `HC` or `EF` storage, but not both. For systems with `HC` and `EF` storage, use the following media-specific XRMEM cache provisioning attributes.

- `xrmemCacheProvEF`: Provisions the vault with the specified amount of XRMEM cache space associated with `EF` media.
- `xrmemCacheProvHC`: Provisions the vault with the specified amount of XRMEM cache space associated with `HC` media.

The following notes also apply when specifying resource provisioning attributes using the `--attributes` option:

- The default value for all resource provisioning attributes is effectively unlimited. Consequently, a vault that is created without any resource provisioning attribute settings has access to all of the resources in the Exascale cluster.
- If you specify a value for any provisioning attributes, then the vault is limited to use the provisioned resources from the corresponding storage pools and the other storage pools are effectively disabled.

  If you specify a space provisioning attribute (`spaceProvEF`, `spaceProvHC`, or `spaceProvXT`) without a corresponding IOPS provisioning attribute (`iopsProvEF`, `iopsProvHC`, or `iopsProvXT`), then the vault can consume the specified space and use unlimited IOPS in the associated storage pool.

  If you specify an IOPS provisioning attribute without a corresponding space provisioning attribute, then the vault can consume unlimited space in the associated storage pool but the I/O bandwidth is limited by the IOPS provisioning attribute.
- Prior to consumption, all provisioned resources are logically provisioned. For example, the `spaceProvHC` attribute specifies the maximum amount of `HC` storage space that the vault can consume. However, physical space is only consumed when data is written to files in the vault.
- To specify an amount of space, you can:
  * Use an integer value representing the number of bytes. The maximum value is `4503599626321920`.
  * Specify a space value using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.
  * Use `unlimited` to specify an unlimited amount of space.
- To specify a number of IOPS, you can:
  * Use an integer value. The maximum value is `4294967295`.
  * Use `unlimited` to specify an unlimited number of IOPS.
- Specifying `0` (zero) as the value for a cache provisioning attribute (`flashCacheProv`, `xrmemCacheProv`, `xrmemCacheProvEF`, or `xrmemCacheProvHC`) disables use of the corresponding cache for the vault.

Examples

Example 6-70 Create a Vault with Unlimited Storage

The following example shows creating a vault named `my-data` that has unlimited access to all of the resources in the Exascale cluster.

```
@> mkvault @my-data
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

6.2.6.4 rmvault

Delete a vault.

Purpose

The `rmvault` command deletes a vault.

Syntax

```
rmvault [@]vault [ --force ] [ --nowait ]
```

Command Options

The options for the `rmvault` command are:

- *vault*: Identifies the vault that is being deleted.
- `--force`: Also deletes all files in the vault.
- `--nowait`: Instead of waiting for the command to finish, the command returns immediately while the operation completes in the background.

Examples

Example 6-71 Delete a Vault

The following example shows deleting the vault named `my-data`.

```
@> rmvault @my-data
```

6.2.7 Dataset Management

This section contains references for the Exascale command line interface (ESCLI) commands that are associated with dataset management:

- lsdataset
  List datasets.

6.2.7.1 lsdataset

List datasets.

Purpose

The `lsdataset` command displays information about Exascale datasets. It can also list the files contained in a specific dataset.

Syntax

```
lsdataset [ dataset-ID ] [ -l ] [ --detail ]
          [ --attributes attribute[,attribute]... ]
          [ --filter filter[,filter]... ]
          [ --sort [-]attribute[,[-]attribute]... ]
          [ --count value ] [ --files [ --recursive ]]
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Command Options

The options for the `lsdataset` command are:

- *dataset-ID*: The identifier for a dataset that you want to display information about.

  Each system-defined dataset has a unique composite identifier, which contains unique identifiers for the associated entities. The dataset identifier has one of the following formats:

  - *`@Vault-name`*: Identifies the vault-level dataset for a specific named vault.
  - *`@Vault-name:GI-cluster-ID`*: Identifies the dataset for a specific Oracle Grid Infrastructure (GI) cluster, which consumes storage space in the specified vault.
  - *`@Vault-name:GI-cluster-ID:CDB-ID`*: Identifies the dataset for a specific Oracle multitenant container database (CDB) that belongs to the specified GI cluster.
  - *`@Vault-name:GI-cluster-ID:CDB-ID:PDB-ID`*: Identifies the dataset for an Oracle pluggable database (PDB) that is associated with the specified CDB, GI cluster, and vault.

  The asterisk (*) can be used for wildcard searches. For example, *`@Vault-name:GI-cluster-ID:*`* displays information about all CDB datasets associated with the specified GI cluster and vault.

  If *dataset-ID* is not specified, the command displays information about all datasets in all vaults.
- `-l`: Displays output in a long, tabular format.
- `--detail`: Displays detailed output with additional attributes.
- `--attributes`: Identifies specific attributes to display in the output.

  To see a list of all available attributes and their descriptions, use the `describe datasets` command. See also Describing Resources and Attributes.
- `--filter`: Specifies conditions for including items in the output.
- `--sort`: Sorts the output using the specified attributes.
- `--count`: Specifies the maximum number of items to display in the output.
- `--files`: Lists files that are in the specified dataset.
- `--recursive`: Also lists files that are in descendant datasets.

Usage Notes

Note the following information when using this command:

- Filter conditions are specified as: <*`attribute`*><*`operator`*><*`value`*>.

  The allowed operators are =, !=, >=, <=, >, and <.

  Multiple comma-separated filter conditions are combined using AND logic.

  Dates can be specified using the following formats:

  - `yyyy-MM-dd''T''HH:mm:ss`
  - `yyyy-MM-dd` (Time is assumed to be 00:00 AM)
  - `HH:mm:ss` (Date is assumed to be today)

  A date can also be followed by a timezone specification.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Sizes can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.

- Sorting attributes are specified as: `[-]<attribute>`. Multiple sort attributes are comma-delimited. The default sort order is ascending. For descending sort order, prefix the attribute name with `-`.

  For example, use the following to primarily sort by `name` in descending order, and use *otherAttribute* in ascending order to further sort entries with the same `name`:

  ```
  --sort -name,otherAttribute
  ```

- You must specify a *dataset-ID* in conjunction with the `--files` option.

Examples

Example 6-72 Display All Datasets

You can use the `lsdataset` command with no other options to display a list of all datasets across all vaults.

```
@> lsdataset
id                                                                                  name
@my-data                                                                            @my-data
@my-data:a5b4997a027d6f91ffd9729702ff6ec5                                           @my-data/
cluster1
@my-data:a5b4997a027d6f91ffd9729702ff6ec5:1427076301                                @my-data/
cluster1:mydb
@my-data:a5b4997a027d6f91ffd9729702ff6ec5:1427076301:1402749181                     @my-data/
cluster1:mydb.PDB$SEED
@my-data:a5b4997a027d6f91ffd9729702ff6ec5:1427076301:2132037342                     @my-data/
cluster1:mydb.CDB$ROOT
@my-data:a5b4997a027d6f91ffd9729702ff6ec5:1427076301:2164757665                     @my-data/
cluster1:mydb.MYDB_PDB1
@vault2                                                                             @vault2
...
```

Example 6-73 Display Additional Dataset Attributes

```
@> lsdataset --attributes id,name,fileUsageHC,datasetUsageHC
id
name                                fileUsageHC  datasetUsageHC
@my-data                                                                            @my-
data                                 0.0000         3.4961G
@my-data:a5b4997a027d6f91ffd9729702ff6ec5                                           @my-data/
cluster1                      294.5339M         3.4961G
@my-data:a5b4997a027d6f91ffd9729702ff6ec5:1427076301                                @my-data/
cluster1:mydb                  25.2656M         3.2085G
@my-data:a5b4997a027d6f91ffd9729702ff6ec5:1427076301:1402749181                     @my-data/
cluster1:mydb.PDB$SEED       650.6719M       650.6719M
@my-data:a5b4997a027d6f91ffd9729702ff6ec5:1427076301:2132037342                     @my-data/
cluster1:mydb.CDB$ROOT         1.9056G         1.9056G
@my-data:a5b4997a027d6f91ffd9729702ff6ec5:1427076301:2164757665                     @my-data/
cluster1:mydb.MYDB_PDB1      658.1719M       658.1719M
@vault2
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
@vault2                                    0.0000       110.8750M
...
```

Example 6-74 Display Files in a Dataset

```
@> lsdataset --files @my-
data:a5b4997a027d6f91ffd9729702ff6ec5:1427076301:2164757665
DATAFILE/mydb_pdb1_db.f
DATAFILE/mydb_pdb1_tmp.f
DATAFILE/mydb_pdb1_xdb.f
DATAFILE/mydb_pdb1_ax.f
```

Example 6-75 Using the Recursive Display Option

The following examples illustrate using the `--recursive` option.

Without the `--recursive` option, the following example shows the files belonging directly to the specified GI cluster dataset.

```
@> lsdataset --files @my-data:a5b4997a027d6f91ffd9729702ff6ec5
cluster1.ocr
cluster1/vfile1
```

With the `--recursive` option, the following example shows the files belonging to the specified GI cluster and all of the associated CDB and PDB datasets.

```
@> lsdataset --files --recursive @my-data:a5b4997a027d6f91ffd9729702ff6ec5
DATAFILE/mydb_pdb0_ax.f
dbs/xspfile.ora
DATAFILE/mydb_pdb0.xml
DATAFILE/mydb_pdb1_db.f
DATAFILE/t_ax1.f
DATAFILE/t_db1.f
CLUSTER1-A5B4997A027D6F91FFD9729702FF6EC5/MYDB/TEMPFILE/TEMP.OMF.2012709C
DATAFILE/t_xdb1.f
DATAFILE/mydb_pdb1_tmp.f
dbs/t_cf1.f
DATAFILE/mydb_pdb0_xdb.f
DATAFILE/mydb_pdb0_tmp.f
DATAFILE/t_undo1.f
cluster1.ocr
DATAFILE/mydb_pdb1_xdb.f
CLUSTER1-A5B4997A027D6F91FFD9729702FF6EC5/MYDB/PASSWORD/pwdMYDB.4ABC8604
cluster1/vfile1
DATAFILE/mydb_pdb0_db.f
DATAFILE/mydb_pdb1_ax.f
...
```

6.2.8 File Management

This section contains references for the Exascale command line interface (ESCLI) commands that are associated with file management:

Copyright © 2022, 2026, Oracle and/or its affiliates.

- cd
  Change the current directory.
- chfile
  Change file attributes.
- clonefile
  Clone one or more files.
- extentmap
  Display storage extent information.
- getfile
  Copy a file from Exascale to the local file system.
- ls
  List files and vaults.
- lsfile
  List files.
- lssnapshots
  List snapshots associated with a file.
- mkfile
  Create a file.
- putfile
  Copy a file from the local file system to Exascale.
- rmfile
  Delete a file.
- snapshotfile
  Create snapshots.

6.2.8.1 cd

Change the current directory.

Purpose

The `cd` command allows you to change the current directory for file operations. File operations are performed relative to the specified directory if absolute paths are not used.

Syntax

```
cd directory
```

Command Options

directory: Specifies the directory to change to. You can specify an absolute directory path or a directory path that is relative to the current directory.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Examples

Example 6-76 Change Directories using an Absolute Directory Path

You can change to a specified absolute directory path, such as `@my-vault/dir`.

```
@> cd @my-vault/dir
@my-vault/dir/>
```

Example 6-77 Change Directories using a Relative Directory Path

You can change directories using a specified relative directory path, such as `../dir2`.

```
@my-vault/dir/> cd ../dir2
@my-vault/dir2/>
```

6.2.8.2 chfile

Change file attributes.

Purpose

The `chfile` command changes attributes of a file.

Syntax

```
chfile file-name --attributes attribute=value[,attribute=value]...
```

Command Options

The options for the `chfile` command are:

- *file-name*: Specifies the name of the file you are changing.
- `--attributes`: Specifies attributes to change.

  Use the `describe chfile` command to see details about the file attributes you can change. See also Describing Resources and Attributes.

Examples

Example 6-78 Change a File

The following example shows changing the size of the file named `my-file` to 100 KB.

```
@my-data/> chfile my-file --attributes size=100k
File altered.

@my-data/> ls -l
Total 1
100.0k  29 Jun 21:58   my-file
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

6.2.8.3 clonefile

Clone one or more files.

Purpose

The `clonefile` command allows you to create a writable clone of a file or group of files.

Syntax

```
clonefile source target [ source target ]... [ --exclude exclude-spec ]...
```

Command Options

The options for the `clonefile` command are:

- *source*: Specifies the name of the file or files that are the clone source.
- *target*: Specifies the name of the file or files that are the clone destination.
- `--exclude`: Specifies the name of the file or files that are excluded from the clone operation.

Usage Notes

Note the following information when using this command:

- You can use a wildcard (*) in the *source* to specify multiple source files, in which case the corresponding *target* must also contain a matching wildcard.
- All files in a clone operation must be in the same vault; otherwise, an error is returned.
- Multiple *source* and *target* pairs are permitted. In this case, the source file specifications are considered in order, and only the first match is used.
- Files specified using the `--exclude` option are excluded from the clone operation. Such exclusions apply across all of the *source* and *target* pairs.
- If you specify a null (empty) string as a *target* value, no clones are created for the corresponding *source*. Specifying a null *target* is an alternative method of excluding files from the clone operation.

  The *source* and null *target* is considered in order along with all other *source* and *target* pairs, and only the first match is used.
- All clones created in the same operation are point-in-time consistent.

Examples

Example 6-79 Clone a File

The following example shows cloning a file. The source file is `@my-data/file1`, and the clone is `@my-data/clone-file1`.

```
@> clonefile @my-data/file1 @my-data/clone-file1
```

Example 6-80 Clone a Group of Files

The following example shows cloning a group of files using a wildcard. Assuming the existence of files named `@my-data/file1`, `@my-data/file2`, and so on, the example creates a clone of

Copyright © 2022, 2026, Oracle and/or its affiliates.

`@my-data/file1` named `@my-data/clone-dir/file1`, `@my-data/file2` named `@my-data/clone-dir/file2`, and so on.

```
@> clonefile @my-data/file* @my-data/clone-dir/file*
```

Example 6-81 Clone Multiple File Groups

The following example shows cloning multiple groups of files. The example clones `@my-data/a*` and `@my-data/b*`, with the resulting clones located under `@my-data/clone/`.

```
@> clonefile @my-data/a* @my-data/clone/a* @my-data/b* @my-data/clone/b*
```

Example 6-82 Ordering Significance

These examples show the significance of the order in which the *source* and *target* pairs are specified.

In the first command, the clones for files matching `@my-data/a*` are created in `@my-data/clone/a`, and the clones for the other files are created in `@my-data/clone/other`.

```
@> clonefile @my-data/a* @my-data/clone/a/a* @my-data/* @my-data/clone/other/*
```

In the following command, all of the clones are created in `@my-data/clone/other` because all of the files match `@my-data/*`. In this case, the second *source* and *target* pair is never used.

```
@> clonefile @my-data/* @my-data/clone/other/* @my-data/a* @my-data/clone/a/a*
```

Example 6-83 Clone a Group of Files with Exclusions

The following examples show cloning files in `@my-data`, except for those matching `@my-data/a*` or `@my-data/b*`.

```
@> clonefile @my-data/* @my-data/clone/* --exclude @my-data/a* --exclude @my-data/b*
```

```
@> clonefile @my-data/a* "" @my-data/b* "" @my-data/* @my-data/clone/*
```

Example 6-84 Invalid Commands

The following examples show invalid commands. The first and second commands are invalid because multiple vaults are referenced. The final command is invalid because the clone destination does not contain a wildcard to match the source specification.

```
@> clonefile @my-data/a* @my-data-backup/a*
```

```
@> clonefile @my-data/a* @my-data/clone/a* @vault2/a* @vault2/clone/a*
```

```
@> clonefile @my-data/withwc* @my-data/clone/withoutwc
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

6.2.8.4 extentmap

Display storage extent information.

Purpose

The `extentmap` command enables Exascale cluster administrators to display storage extent information for Exascale files and vaults.

Syntax

```
extentmap [ --name name --file-offset offset [ --mirror mirror ]]
          [ --extId extent --target target ]
```

Command Options

The options for the `extentmap` command are:

- `--name`: A file or vault name for which you want to display extent information.

  Vault names are preceded with the `@` symbol.
- `--file-offset`: Specifies the offset location for which to display information.

  The value is expressed in bytes and can be a single number or a range. Use this option when you specify a file name for translation.
- `--mirror`: Optionally specifies the mirror number for which to display information.

  For high (3-way) redundancy, the mirror number can be `0`, `1`, or `2`.

  You can use this option when you specify a file name for translation.
- `--extId`: An extent identifier for which you want to display information.

  You can use this option to translate an extent identifier to a vault or file name.
- `--target`: Specifies the translation target; either `vaults` or `files`.

  You can use this option when you specify an extent identifier for translation.

Usage Notes

This command requires the cluster administrator privilege.

Examples

Example 6-85 Display Extent Information for a Vault

The following example shows how to display extent information for the specified vault.

```
@> extentmap --name @my-data
Vault Extent ID: c000_0000_003a:0000_0000_0000:00000000:00000000
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Example 6-86 Display Extent Information for a Specific File Offset

The following example shows how to display extent information for a specified byte offset in a specified file.

```
@> extentmap --name @my-data/file1 --file-offset 123
extId: c000_0000_003a:8000_0000_000a:00000000:00000000
rpmId: 3222802720
extentEndOffset: 8388608
mirrorHealth: OK
translationPath:
o/192.0.2.1;192.0.2.1;192.0.2.1;192.0.2.1:41681/cell1_CD_exdisk6_EDSCELL1
mirrorNumber: 0

extId: c000_0000_003a:8000_0000_000a:00000000:00000000
rpmId: 3222802720
extentEndOffset: 8388608
mirrorHealth: OK
translationPath:
o/192.0.2.3;192.0.2.3;192.0.2.3;192.0.2.3:41681/cell3_CD_exdisk1_EDSCELL3
mirrorNumber: 1

extId: c000_0000_003a:8000_0000_000a:00000000:00000000
rpmId: 3222802720
extentEndOffset: 8388608
mirrorHealth: OK
translationPath:
o/192.0.2.6;192.0.2.6;192.0.2.6;192.0.2.6:41681/cell6_CD_exdisk9_EDSCELL6
mirrorNumber: 2
```

Example 6-87 Display Extent Information for a Specific Mirror

The following example shows how to display extent information for a specific extent mirror, using a specified byte offset in a specified file.

```
@> extentmap --name @my-data/file1 --file-offset 123 --mirror 1
extId: c000_0000_003a:8000_0000_000a:00000000:00000000
rpmId: 3222802720
extentEndOffset: 8388608
mirrorHealth: OK
translationPath:
o/192.0.2.3;192.0.2.3;192.0.2.3;192.0.2.3:41681/cell3_CD_exdisk1_EDSCELL3
mirrorNumber: 1
```

Example 6-88 Display Vault Name for a Specific Extent

The following example shows how to display the vault name for a specific extent.

```
@> extentmap --extId c000_0000_003a:8000_0000_000a:00000000:00000000 --target
vaults
vaultName: my-data
SysEdsURL: 192.0.2.3;192.0.2.3;192.0.2.3;192.0.2.3:41681
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Example 6-89 Display File Name for a Specific Extent

The following example shows how to display the file name for a specific extent.

```
@> extentmap --extId c000_0000_003a:8000_0000_000a:00000000:00000000 --target
files
vaultName: my-data
fileName: file1
UserEdsURL: 192.0.2.3;192.0.2.3;192.0.2.3;192.0.2.3:41681
```

6.2.8.5 getfile

Copy a file from Exascale to the local file system.

Purpose

The `getfile` command copies a file from an Exascale vault to the local file system.

Syntax

```
getfile exa-file local-file
```

Command Options

The options for the `getfile` command are:

- *exa-file*: Specifies a file in an Exascale vault.
- *local-file*: Specifies a regular file location in the local file system.

Examples

Example 6-90 Get a File

The following example copies the Exascale file at `@my-data/file1` to `/tmp/file1.copy` on the local file system.

```
@> getfile @my-data/file1 /tmp/file1.copy
```

6.2.8.6 ls

List files and vaults.

Purpose

The `ls` command displays information about Exascale files and vaults.

Syntax

```
ls [ name [ name ]... ] [ -l ] [ --detail ]
   [ --attributes attribute[,attribute]... ]
   [ --filter filter[,filter]... ]
   [ --sort [-]attribute[,[-]attribute]... ]
   [ --sum attribute ] [ --count value ]
   [ -t ] [ --include-being-deleted ]
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Command Options

The options for the `ls` command are:

- *name*: A file or vault name that you want to display information for.

  Vault names can be preceded with the `@` symbol. The asterisk (`*`) can be used for wildcard searches.

  If no value is specified, then information is displayed about all files or vaults in the current level of the ESCLI file hierarchy.
- `-l`: Displays output in a long, tabular format.
- `--detail`: Displays detailed output with additional attributes.
- `--attributes`: Identifies specific attributes to display in the output.

  You can use this option to focus on a subset of attributes as an alternative to the `--detail` option.
- `--filter`: Specifies conditions for including items in the output.
- `--sort`: Sorts the output using the specified attributes.
- `--sum`: Calculates the sum of the specified numeric attribute.
- `--count`: Specifies the maximum number of items to display in the output.
- `-t`: Sorts files or vaults by ascending creation time.
- `--include-being-deleted`: Lists all files, including deleted files in the process of releasing their allocated storage.

Usage Notes

Note the following information when using this command:

- Filter conditions are specified as: <*attribute*><*operator*><*value*>.

  Attributes can be file attributes, or vault attributes with the `vault.` prefix.

  The allowed operators are =, !=, >=, <=, >, and <.

  Multiple comma-separated filter conditions are combined using AND logic.

  Dates can be specified using the following formats:

  - `yyyy-MM-dd''T''HH:mm:ss`
  - `yyyy-MM-dd` (Time is assumed to be 00:00 AM)
  - `HH:mm:ss` (Date is assumed to be today)

  A date can also be followed by a timezone specification.

  File sizes can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.

  For example, the following filter only includes files created after 1 PM today and with a size greater than 10 megabytes:

  ```
  --filter createTime>13:00:00,size>10M
  ```

- Sorting attributes are specified as: `[-]`<*attribute*>. Multiple sort attributes are comma-delimited. The default sort order is ascending. For descending sort order, prefix the attribute name with `-`.

Copyright © 2022, 2026, Oracle and/or its affiliates.

For example, use the following to primarily sort by `name` in descending order, and use *otherAttribute* in ascending order to further sort entries with the same `name`:

```
--sort -name,otherAttribute
```

- You can use the `--sum` option only by itself or together with the `--filter` option.

  The `--sum` option is available starting with Oracle Exadata System Software release 26.1.0.

Examples

Example 6-91 Display All Vaults

If you use the `ls` command at the root level, you get a listing of all vaults.

```
@> ls
my-data
vault2
```

Example 6-92 Display Files in a Vault

You can specify the name of a vault followed by the slash character (`/`) to list all the files within the vault.

```
@> ls @my-data/
x
y
z1
z2
z3
```

Example 6-93 Display File Information in a Long Format

If you use the `-l` option, the output includes additional information about each file or vault.

```
@my-data/> ls -l
Total 5
  10.0M   05 Jan 12:59 x
  19.5k   05 Jan 13:08 y
   5.0k   05 Jan 13:09 z1
  10.0M   05 Jan 13:23 z2
  20.0G   05 Jan 13:14 z3
```

Example 6-94 Sorting the Output of the ls Command

The following example shows the files in the vault `my-data`, with the most recent appearing first.

```
@my-data/> ls --sort -createTime
z2
z3
z1
y
x
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Example 6-95 Listing Files Using a Wildcard Search

You can use the * symbol as a wildcard when specifying the files to list.

```
@my-data> ls z*
z1
z2
z3
```

Example 6-96 Display Detailed File Information

If you use the `--detail` option, the output includes detailed information about the files and vaults. This example shows the detailed output for the file `x` in the vault `my-data`.

```
@> ls @my-data/x --detail
name                    x
size                    0
createTime              2019-01-05 12:59:57 GMT
fileType                8
vault.name              my-data
vault.createTime        2018-11-08 22:41:07 GMT
```

Example 6-97 Display Specific Attributes for Files or Vaults

The following example shows the size, name, and vault name attributes for the files in the `my-data` vault.

```
@my-data/> ls -l --attributes size,name,vault.name
Total 5
      0 x my-data
  19.5k y my-data
   5.0k z1 my-data
  10.0M z2 my-data
  20.0G z3 my-data
```

6.2.8.7 lsfile

List files.

Purpose

The `lsfile` command displays information about Exascale files.

Syntax

```
lsfile [ name [ name ]... ] [ -l ] [ --detail ]
   [ --attributes attribute[,attribute]... ]
   [ --filter filter[,filter]... ]
   [ --sort [-]attribute[,[-]attribute]... ]
   [ --sum attribute ] [ --count value ]
   [ -t ] [ --include-being-deleted ]
```

Command Options

The options for the `lsfile` command are:

Copyright © 2022, 2026, Oracle and/or its affiliates.

- *name*: A file name that you want to display information for.

  Vault names can be preceded with the `@` symbol. The asterisk (`*`) can be used for wildcard searches.

  If no value is specified, then information is displayed about all files in the current level of the ESCLI file hierarchy.
- `-l`: Displays output in a long, tabular format.
- `--detail`: Displays detailed output with additional attributes.
- `--attributes`: Identifies specific attributes to display in the output.

  You can use this option to focus on a subset of attributes as an alternative to the `--detail` option.
- `--filter`: Specifies conditions for including items in the output.
- `--sort`: Sorts the output using the specified attributes.
- `--sum`: Calculates the sum of the specified numeric attribute.
- `--count`: Specifies the maximum number of items to display in the output.
- `-t`: Sorts files by ascending creation time.
- `--include-being-deleted`: Lists all files, including deleted files in the process of releasing their allocated storage.

Usage Notes

Note the following information when using this command:

- Filter conditions are specified as: <*attribute*><*operator*><*value*>.

  Attributes can be file attributes, or vault attributes with the `vault.` prefix.

  The allowed operators are =, !=, >=, <=, >, and <.

  Multiple comma-separated filter conditions are combined using AND logic.

  Dates can be specified using the following formats:

  - `yyyy-MM-dd''T''HH:mm:ss`
  - `yyyy-MM-dd` (Time is assumed to be 00:00 AM)
  - `HH:mm:ss` (Date is assumed to be today)

  A date can also be followed by a timezone specification.

  File sizes can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.

  For example, the following filter only includes files created after 1 PM today and with a size greater than 10 megabytes:

  ```
  --filter createTime>13:00:00,size>10M
  ```

- Sorting attributes are specified as: `[-]`<*attribute*>. Multiple sort attributes are comma-delimited. The default sort order is ascending. For descending sort order, prefix the attribute name with `-`.

Copyright © 2022, 2026, Oracle and/or its affiliates.

For example, use the following to primarily sort by `name` in descending order, and use *otherAttribute* in ascending order to further sort entries with the same `name`:

```
--sort -name,otherAttribute
```

- You can use the `--sum` option only by itself or together with the `--filter` option.

  The `--sum` option is available starting with Oracle Exadata System Software release 26.1.0.

Examples

Example 6-98 Display Files in a Vault

You can specify the name of a vault followed by the slash character (`/`) to list all the files within the vault.

```
@> lsfile @my-data/
x
y
z1
z2
z3
```

Example 6-99 Display File Information in a Long Format

If you use the `-l` option, the output includes additional information about each file.

```
@my-data/> lsfile -l
Total 5
  10.0M   05 Jan 12:59 x
  19.5k   05 Jan 13:08 y
   5.0k   05 Jan 13:09 z1
  10.0M   05 Jan 13:23 z2
  20.0G   05 Jan 13:14 z3
```

Example 6-100 Sorting the Output of the lsfile Command

The following example shows the files in the vault `my-data`, with the most recent appearing first.

```
@my-data/> lsfile --sort -createTime
z2
z3
z1
y
x
```

Example 6-101 Listing Files Using a Wildcard Search

You can use the * symbol as a wildcard when specifying the files to list.

```
@my-data> lsfile z*
z1
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
z2
z3
```

Example 6-102 Display Detailed File Information

If you use the `--detail` option, the output includes detailed information about the files. This example shows the detailed output for the file `x` in the vault `my-data`.

```
@> lsfile @my-data/x --detail
name                    x
size                    0
createTime              2019-01-05 12:59:57 GMT
fileType                8
vault.name              my-data
vault.createTime        2018-11-08 22:41:07 GMT
```

Example 6-103 Display Specific Attributes for Files

The following example shows the size, name, and vault name attributes for the files in the `my-data` vault.

```
@my-data/> lsfile -l --attributes size,name,vault.name
Total 5
      0 x my-data
  19.5k y my-data
   5.0k z1 my-data
  10.0M z2 my-data
  20.0G z3 my-data
```

6.2.8.8 lssnapshots

List snapshots associated with a file.

Purpose

The `lssnapshots` command displays information about snapshots that are associated with the specified file.

Syntax

```
lssnapshots file-name [ --all | --tree ] [ -l ] [ --detail ]
            [ --attributes attribute[,attribute]... ]
            [ --filter filter[,filter]... ] [ --count value ]
```

Command Options

The options for the `lssnapshots` command are:

- *file-name*: Identifies the file that is the subject of the command. The specified file can be a snapshot, clone, or regular file.
- `--all`: Lists information about all files, clones, and snapshots belonging to the snapshot tree that contains the specified file. The output is presented in a series of lists. The first list contains the files and clones in the snapshot tree. The remaining output lists snapshots that are grouped by their source file.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `--tree`: Lists information about all files, clones, and snapshots belonging to the snapshot tree that contains the specified file. The output is arranged in a graphical tree.
- `-l`: Displays output in a long, tabular format.
- `--detail`: Displays detailed output with additional attributes.
- `--attributes`: Identifies specific attributes to display in the output.

  You can use this option to focus on a subset of attributes as an alternative to the `--detail` option.
- `--filter`: Specifies conditions for including items in the output.
- `--sort`: Sorts the output using the specified attributes.
- `--count`: Specifies the maximum number of items to display in the output.

Usage Notes

Note the following information when using this command:

- If you do not specify the command options `--all` or `--tree`, then the command displays information about immediate snapshots of the specified source file.
- You cannot use any other command options in conjunction with `--all` or `--tree`.
- Filters are specified as: <*attribute*><*operator*><*value*>.

  Attributes can be file attributes, or vault attributes with the `vault.` prefix. Multiple filter conditions are delimited by commas.

  The allowed operators are =, >=, <=, >, and <.

  Dates can be specified using the following formats:

  – `yyyy-MM-dd''T''HH:mm:ss`
  – `yyyy-MM-dd` (Time is assumed to be 00:00 AM)
  – `HH:mm:ss` (Date is assumed to be today)

  Any of these formats can be followed by a timezone specification.

  File sizes can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.

  For example, to filter the results to show only snapshots created after noon today, with a size greater than 10 megabytes, you could use the following:

  ```
  --filter creationTime>12:00:00,size>10M
  ```

- Output is ordered according to the snapshot creation time, starting with the most recent.

Examples

Example 6-104 List Basic Snapshot Information

The following example shows how to list information about all of the snapshots that are based on the specified file.

```
@> lssnapshots @my-data/my-file
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Example 6-105 List Detailed Snapshot Information

The following example shows how to list detailed information about all of the snapshots that are based on the specified file.

```
@> lssnapshots @my-data/my-file --detail
```

Example 6-106 List Snapshot Information with Specific Attributes

The following example shows how to list specific attributes about all of the snapshots that are based on the specified file.

```
@> lssnapshots @my-data/my-file --detail --attributes
name,vault.name,createTime,vault.createTime
```

Example 6-107 Filter the List of Snapshots

The following example shows how to specify a filter that constrains the command output. In the example, the output is limited to the snapshots that were created after the specified date.

```
@> lssnapshots @my-data/my-file --filter createTime>2020-01-01
```

Example 6-108 List the Latest Snapshots

The following example shows how to specify a count that constrains the command output. In the example, the output is limited to the 5 most recently created snapshots.

```
@> lssnapshots @my-data/my-file --count 5
```

Example 6-109 Display a Snapshot Tree

The command output in this example assumes a group of related files created using the following command sequence:

```
@my-data/> putfile somelocalfile file1
Putting file somelocalfile to @my-data/file1
Success.

@my-data/> snapshotfile file1 snap1_of_file1
Success.

@my-data/> clonefile file1 clone1
Success.

@my-data/> snapshotfile file1 snap2_of_file1
Success.

@my-data/> snapshotfile clone1 snap1_of_clone1
Success.
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

The following example shows how to display all files in a snapshot tree using the `--tree` command option.

```
@> lssnapshots @my-data/clone1 --tree

+---- [1] @my-data/snap1_of_file1
+---- [2] @my-data/?8000_0000_000a:00000001
|
|`---+---- [4] @my-data/snap1_of_clone1
|    +---- [6] @my-data/clone1
|
+---- [3] @my-data/snap2_of_file1
+---- [5] @my-data/file1
```

Note that the command output includes an internal snapshot (`@my-data/?8000_0000_000a:00000001`), which was implicitly created as part of the cloning operation to create `@my-data/clone1`.

Example 6-110 List all Files in a Snapshot Tree

The following example shows how to display all files in a snapshot tree using the `--all` command option. This example assumes the same scenario as the previous example.

```
@> lssnapshots @my-data/clone1 --all

---FILE/CLONES:---

2022-05-02 06:14:41 UTC @my-data/file1
2022-05-02 06:15:15 UTC @my-data/clone1

---SNAPSHOTS:---

@my-data/file1
  2022-05-02 06:15:06 UTC @my-data/snap1_of_file1
  2022-05-02 06:15:15 UTC @my-data/?8000_0000_000a:00000001
  2022-05-02 06:15:24 UTC @my-data/snap2_of_file1

@my-data/clone1
  2022-05-02 06:15:45 UTC @my-data/snap1_of_clone1
```

6.2.8.9 mkfile

Create a file.

Purpose

The `mkfile` command creates a file within a vault.

Syntax

```
mkfile file-name [ --template template ]
       [ --attributes attribute=value[,attribute=value]... ]
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Command Options

The options for the `mkfile` command are:

- *file-name*: Specifies the name of the file you are creating. The file name cannot contain wildcard characters. Only one file name can be specified.
- `--template`: Creates the file using the specified template.
- `--attributes`: Optionally specifies attributes to change.

  Use the `describe mkfile` command to see details about the file attributes you can set. See also Describing Resources and Attributes.

Usage Notes

Note the following information when using this command:

- You can set the file size by specifying the `size` attribute.

  File sizes can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.

  Regardless of the file size setting, storage space is physically materialized only when the file contents is written. So, after the `mkfile` command, the file initially consumes no physical storage space. The file only consumes physical storage space when data is written to it.
- You can explicitly set the file storage attributes:
  - `mediaType`: Specifies the physical media type used to store the file. Possible values are:
    - `HC`: Identifies high capacity storage, using hard disk drives (HDDs) on high-performance Exadata storage servers.
    - `EF`: Identifies extreme flash storage, using low-latency, high-throughput flash devices.
    - `XT`: Identifies extended storage, using HDDs on Exadata storage servers optimized for cost-effective online storage of less frequently accessed data.
  - `redundancy`: Specifies the number of data copies that are maintained. Currently, the only permitted value is:
    - `high`: Indicates three mirrored copies of the file data.
  - `contentType`: Specifies the type of content in the file. Exascale internally uses this attribute to place file extents on physically separate devices in a manner that maximizes availability in the event of a failure. Possible values are:
    - `DATA`: Principally associated with user data.
    - `RECO`: Primarily for data used in backup and recovery operations.

  Additionally, starting with Oracle Exadata System Software release 26.1.0, you can set the `storagePool` attribute to specify the name of a storage pool used to store the file.

  If you explicitly set any of the file storage attributes, the specified value overrides the setting in the default template. If you specify both `mediaType` and `storagePool`, the values must be compatible. For example, you cannot specify the use of high capacity (HC) storage media (`mediaType=HC`) and also specify the name of a storage pool that uses extreme flash (EF) storage.
- You can set the Exascale file type by specifying the `fileType` attribute.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Every Exascale file type is associated with an Exascale template. So, you can view a complete list of the Exascale file types by using the ESCLI lstemplate command. For example:

```
@> lstemplate --cluster
```

```
@> lstemplate --vault vault2
```

Examples

Example 6-111 Create a File in a Vault

The following example shows creating a file named `my-file` in a vault named `my-data`.

```
@my-data/> mkfile my-file --attributes size=20k
File created.

@my-data/> ls -l
  20.0k   08 Jan 15:53 my-file
```

Example 6-112 Create a File with a Specified File Type

The following example shows creating a file named `my-data-file` in a vault named `my-data` using a specified file type (`fileType=DATAFILE`).

```
@> mkfile @my-data/my-data-file --attributes fileType=DATAFILE
```

Example 6-113 Create a File with a Specified File Template

The following example shows creating a file named `my-data-file2` in a vault named `my-data` using a specified file template (`--template DATAFILE`).

```
@> mkfile @my-data/my-data-file2 --template DATAFILE
```

Example 6-114 Create a File with Explicit File Storage Attributes

The following example shows creating a file named `my-data-file3` in a vault named `my-data` using specified file storage attributes (`storagePool=my-pool1,contentType=DATA`).

```
@> mkfile @my-data/my-data-file3 --attributes storagePool=my-
pool1,contentType=DATA
```

Related Topics

- About Exascale File Storage Attributes
- lstemplate
  List file templates.

Copyright © 2022, 2026, Oracle and/or its affiliates.

6.2.8.10 putfile

Copy a file from the local file system to Exascale.

Purpose

The `putfile` command copies a file from the local file system to an Exascale vault.

Syntax

```
putfile local-file exa-file
```

Command Options

The options for the `putfile` command are:

- *local-file*: Specifies a regular file in the local file system.
- *exa-file*: Specifies a file location in an Exascale vault.

Examples

Example 6-115 Put a File

The following example copies the local file at `/tmp/file1` to `@my-data/file1.copy` on Exascale storage.

```
@> putfile /tmp/file1 @my-data/file1.copy
```

6.2.8.11 rmfile

Delete a file.

Purpose

The `rmfile` command deletes a file in an Exascale vault.

Syntax

```
rmfile file-name [ --force ] [ --nowait ] [ --with-snapshots ]
```

Command Options

The options for the `rmfile` command are:

- *file-name*: Identifies the file to be deleted.
- `--force`: Deletes the specified file even if it is in use.
- `--nowait`: Instead of waiting for the command to finish, the command returns immediately while the operation completes in the background.
- `--with-snapshots`: Also deletes unlocked snapshots linked to the specified file. Snapshots are locked if the `snapLockExpiryTime` attribute is set to a future date and time. Locked snapshots are skipped and remain unchanged.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Examples

Example 6-116 Delete a File

The following example deletes the file at `@my-data/file1`.

```
@> rmfile @my-data/file1
```

6.2.8.12 snapshotfile

Create snapshots.

Purpose

The `snapshotfile` command creates a snapshot of a file or group of files.

Syntax

```
snapshotfile source1 target1 [ sourceN targetN ]...
             [ --exclude exclude-spec ]...
```

Command Options

The options for the `snapshotfile` command are:

- *source1-N*: Specifies the name of the file or files that are the snapshot source.
- *target1-N*: Specifies the name of the file or files that are the snapshot destination.
- `--exclude`: Specifies the name of the file or files that are excluded from the snapshot operation.

Usage Notes

Note the following information when using this command:

- You can use a wildcard (*) in the *source* to specify multiple source files, in which case the corresponding *target* must also contain a matching wildcard.
- All files in a snapshot operation must be in the same vault; otherwise, an error is returned.
- Multiple *source* and *target* pairs are permitted. In this case, the source file specifications are considered in order, and only the first match is used.
- Files specified using the `--exclude` option are excluded from the snapshot operation. Such exclusions apply across all of the *source* and *target* pairs.
- If you specify a null (empty) string as a *target* value, no snapshots are created for the corresponding *source*. Specifying a null *target* is an alternative method of excluding files from the snapshot operation.

  The *source* and null *target* is considered in order along with all other *source* and *target* pairs, and only the first match is used.
- All snapshots created in the same operation are point-in-time consistent.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Examples

Example 6-117 Snapshot a File

This example shows creating a snapshot. The source file is `@my-data/file1`, and the snapshot is `@my-data/snap-file1`.

```
@> snapshotfile @my-data/file1 @my-data/snap-file1
```

Example 6-118 Snapshot a Group of Files

This example shows creating snapshots for a group of files using a wildcard. Assuming the existence of files named `@my-data/file1`, `@my-data/file2`, and so on, the example creates a snapshot of `@my-data/file1` named `@my-data/snap-dir/file1`, `@my-data/file2` named `@my-data/snap-dir/file2`, and so on.

```
@> snapshotfile @my-data/file* @my-data/snap-dir/file*
```

Example 6-119 Snapshot Multiple File Groups

This example shows creating snapshots for multiple groups of files using wildcards. The source file groups are `@my-data/a*` and `@my-data/b*`, with the resulting snapshots located under `@my-data/snap/`.

```
@> snapshotfile @my-data/a* @my-data/snap/a* @my-data/b* @my-data/snap/b*
```

Example 6-120 Ordering Significance

These examples show the significance of the order in which the *source* and *target* pairs are specified.

In the first command, the snapshots for files matching `@my-data/a*` are created in `@my-data/snap/a`, and the snapshots for the other files are created in `@my-data/snap/other`.

```
@> snapshotfile @my-data/a* @my-data/snap/a/a* @my-data/* @my-data/snap/
other/*
```

In the following command, all of the snapshots are created in `@my-data/snap/other` because all of the files match `@my-data/*`. In this case, the second *source* and *target* pair is never used.

```
@> snapshotfile @my-data/* @my-data/snap/other/* @my-data/a* @my-data/
snap/a/a*
```

Example 6-121 Snapshot a Group of Files with Exclusions

The following examples show how to snapshot files in `@my-data`, except for those matching with `@my-data/a*` or `@my-data/b*`.

```
@> snapshotfile @my-data/* @my-data/snap/* --exclude @my-data/a* --exclude
@my-data/b*
```

```
@> snapshotfile @my-data/a* "" @my-data/b* "" @my-data/* @my-data/snap/*
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Example 6-122 Invalid Commands

The following examples show invalid commands. The first and second commands are invalid because multiple vaults are referenced. The final command is invalid because the snapshot destination does not contain a wildcard to match the source specification.

```
@> snapshotfile @my-data/a* @my-data-backup/a*


@> snapshotfile @my-data/a* @my-data/snap/a* @vault2/a* @vault2/snap/a*


@> snapshotfile @my-data/withwc* @my-data/snap/withoutwc
```

6.2.9 Template Management

This section contains references for the Exascale command line interface (ESCLI) commands that are associated with template management:

- chtemplate
  Change an attribute for an existing template.
- lstemplate
  List file templates.
- mktemplate
  Create a file template.
- rmtemplate
  Delete a file template.

6.2.9.1 chtemplate

Change an attribute for an existing template.

Purpose

The `chtemplate` command changes an attribute for an existing Exascale file template.

Syntax

```
chtemplate [ --vault vault | --cluster ]
           { --file-type file-type | --name template-name }
           [ --media-type media-type ][ --storage-pool storage-pool-name ]
           [ --redundancy redundancy ]
           [ --content-type content-type ]
```

Command Options

Specify one or more of the following `chtemplate` command options to identify the template that is the subject of the operation:

- `--vault`: Alters a template associated with the specified vault.
- `--cluster`: Alters a template associated with the cluster.
- `--file-type`: Alters a template associated with the specified file type.
- `--name`: Alters a user-defined template having the specified name.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Use the following command options to specify the file storage attributes:

- `--media-type`: Specifies the physical media type used to store files created with the template. Possible values are:
  - `HC`: Identifies high capacity storage, using hard disk drives (HDDs) on high-performance Exadata storage servers.
  - `EF`: Identifies extreme flash storage, using low-latency, high-throughput flash devices.
  - `XT`: Identifies extended storage, using HDDs on Exadata storage servers optimized for cost-effective online storage of less frequently accessed data.
- `--storage-pool`: Specifies the name of a storage pool used to store files created with the template.

  This option is available starting with Oracle Exadata System Software release 26.1.0, and is only permitted in conjunction with a vault template.
- `--redundancy`: Specifies the number of data copies that are maintained for files created with the template. Currently, the only permitted value is:
  - `high`: Indicates three mirrored copies of the file data.
- `--content-type`: Specifies the type of content in files created with the template. Exascale internally uses this attribute to place file extents on physically separate devices in a manner that maximizes availability in the event of a failure. Possible values are:
  - `DATA`: Principally associated with user data.
  - `RECO`: Primarily for data used in backup and recovery operations.

Usage Notes

Note the following information when using this command:

- When a template includes a storage pool specification, it does not require a media type setting. If a storage pool and media type are both specified, the settings must be compatible. For example, a template cannot specify the use of high capacity (HC) storage media and also specify the name of a storage pool that uses extreme flash (EF) storage.

Examples

Example 6-123 Change a User-Defined Cluster-Wide Template

In this example, the content type is set to `DATA` for the user-defined template named `T1`. The template is associated with the cluster, as neither `--vault` nor `--cluster` is specified, and the current working directory in the ESCLI session is the root directory.

```
@> chtemplate --name T1 --content-type DATA
```

Example 6-124 Change a User-Defined Vault-Specific Template

In this example, the media type is set to `HC` and the redundancy is set to `high` for the user-defined template named `T2`. The template is associated with the vault named `vault2`, as neither `--vault` nor `--cluster` is specified, and the current working directory in the ESCLI session is inside `vault2`.

```
@vault2> chtemplate --name T2 --media-type HC --redundancy high
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Example 6-125 Change a User-Defined Vault-Specific Template

In this example, the media type, content type, and redundancy are all set for the user-defined template named `T3` that is associated with the vault named `vault2`.

```
@> chtemplate --vault vault2 --name T3 --content-type DATA --media-type HC --
redundancy high
```

Example 6-126 Change a User-Defined Cluster-Wide Template

In this example, the media type, content type, and redundancy are all set for the cluster-wide user-defined template named `T4`. The `--cluster` option overrides the fact that the current working directory in the ESCLI session is inside `vault2`.

```
@vault2> chtemplate --cluster --name T4 --content-type DATA --media-type HC --
redundancy high
```

Example 6-127 Change a Template for a Specific Vault and File Type

This example changes the template associated with the `DATAFILE` file type and the vault named `vault2`.

```
@> chtemplate --vault vault2 --file-type DATAFILE --content-type DATA --media-
type HC --redundancy high
```

Example 6-128 Change a Cluster-Wide Template for a Specific File Type

This example changes the cluster-wide template associated with the `DATAFILE` file type.

```
@> chtemplate --cluster --file-type DATAFILE --content-type DATA --media-type
HC --redundancy high
```

Related Topics

- About Exascale File Storage Attributes

6.2.9.2 lstemplate

List file templates.

Purpose

The `lstemplate` command displays information about Exascale file templates that are associated with the vault or the Exascale cluster.

Syntax

```
lstemplate [ template-name ]
           [ --vault vault [ --vault-level-only ] | --cluster ]
           [ -l ] [ --detail ] [ --default ]
           [ --attributes attribute[,attribute]... ]
           [ --filter filter[,filter]... ]
           [ --sort [-]attribute[,[-]attribute]... ]
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Command Options

The options for the `lstemplate` command are:

- *template-name*: Optionally limits the output to templates that match the specified template name.
- `--vault`: Lists templates associated with the specified vault.
- `--vault-level-only`: List only the vault-specific templates associated with the specified vault.
- `--cluster`: Lists templates associated with the cluster.
- `-l`: Displays output in a long, tabular format.
- `--detail`: Displays detailed output with additional attributes.
- `--default`: List default configuration details for system-defined templates (`templateType=systemType` or `templateType=oracleFileType`).
- `--attributes`: Identifies specific attributes to display in the output.

  You can use this option to focus on a subset of attributes as an alternative to the `--detail` option.
- `--filter`: Specifies conditions for including items in the output.
- `--sort`: Sorts the output using the specified attributes.

Usage Notes

Note the following information when using this command:

- If `--vault` is specified without `--vault-level-only`, then the output contains all templates affecting the specified vault, including the vault-specific templates and any cluster-level templates not overridden by vault-specific templates.
- If neither `--vault` nor `--cluster` is specified, and the user entered a vault in the ESCLI session, then the output contains all templates affecting the specified vault, including the vault-specific templates and any cluster-level templates not overridden by vault-specific templates.
- If neither `--vault` nor `--cluster` is specified, and the user hasn't entered a vault in the ESCLI session (ESCLI prompt shows `@>`), then templates associated with the cluster are listed.
- Filter conditions are specified as: <*attribute*><*operator*><*value*>.

  The allowed operators are =, !=, >=, <=, >, and <.

  Multiple comma-separated filter conditions are combined using AND logic.
- Sorting attributes are specified as: `[-]`<*attribute*>. Multiple sort attributes are comma-delimited. The default sort order is ascending. For descending sort order, prefix the attribute name with `-`.

  For example, use the following to primarily sort by `name` in descending order, and use *otherAttribute* in ascending order to further sort entries with the same `name`:

  ```
  --sort -name,otherAttribute
  ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Examples

Example 6-129 List Cluster-Wide Templates

These examples list all templates associated with the cluster.

```
@> lstemplate
```

```
@vault2> lstemplate --cluster
```

Example 6-130 List Effective Vault-Specific Templates

These examples list the effective templates associated with the vault `vault2`. The output includes templates associated with the vault, and templates associated with the cluster that are not overridden by vault-specific templates.

```
@vault2> lstemplate
```

```
@> lstemplate --vault vault2
```

Example 6-131 List Only Templates Affecting a Vault

This example displays only the vault-specific templates affecting files in the vault `vault2`.

```
@> lstemplate --vault vault2 --vault-level-only
```

Example 6-132 List a Specific Cluster-Wide Template

This example shows detailed information about the cluster-wide template named `DATAFILE`.

```
@> lstemplate --detail DATAFILE
```

Example 6-133 List Default Configuration Information for a Specific Cluster-Wide Template

This example displays default configuration information for the cluster-wide template named `DATAFILE`.

```
@> lstemplate --default DATAFILE
```

Example 6-134 List Specific Cluster-Wide Templates

This example displays detailed information about cluster-wide templates whose names end with `FILE` and whose content type is `DATA`. The results are sorted by name in ascending order, then by media type in descending order.

```
@> lstemplate --detail *FILE --sort name,-mediaType --filter contentType=DATA
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

6.2.9.3 mktemplate

Create a file template.

Purpose

The `mktemplate` command creates an Exascale file template.

Syntax

```
mktemplate [ --vault vault | --cluster ]
           { --file-type file-type | --name template-name }
           [ --media-type media-type ][ --storage-pool storage-pool-name ]
           --redundancy redundancy
           --content-type content-type
```

Command Options

Specify the following `mktemplate` command options to identify the template and its scope:

- `--vault`: Create a template associated with the specified vault.
- `--cluster`: Create a template associated with the cluster.
- `--file-type`: Create a template associated with the specified file type. If you specify this option, the template name is the same as the file type.
- `--name`: Create a user-defined template having the specified name.

Specify the following File Storage Attributes, which are associated with the template:

- `--media-type`: Specifies the physical media type used to store files created with the template. Possible values are:
  - `HC`: Identifies high capacity storage, using hard disk drives (HDDs) on high-performance Exadata storage servers.
  - `EF`: Identifies extreme flash storage, using low-latency, high-throughput flash devices.
  - `XT`: Identifies extended storage, using HDDs on Exadata storage servers optimized for cost-effective online storage of less frequently accessed data.
- `--storage-pool`: Specifies the name of a storage pool used to store files created with the template.

  This option is available starting with Oracle Exadata System Software release 26.1.0, and is only permitted in conjunction with a vault template.
- `--redundancy`: Specifies the number of data copies that are maintained for files created with the template. Currently, the only permitted value is:
  - `high`: Indicates three mirrored copies of the file data.
- `--content-type`: Specifies the type of content in files created with the template. Exascale internally uses this attribute to place file extents on physically separate devices in a manner that maximizes availability in the event of a failure. Possible values are:
  - `DATA`: Principally associated with user data.
  - `RECO`: Primarily for data used in backup and recovery operations.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Usage Notes

Note the following information when using this command:

- You can define the scope of the template by specifying the `--vault` option or the `--cluster` option. You cannot specify both options simultaneously.

  If neither option is specified, then the scope of the template is inferred from the current working directory in the ESCLI session. If the current working directory in the ESCLI session is the root directory, then the template is associated with the cluster. Otherwise, the template is associated with the vault that is referenced in the current working directory.
- You must define the template as a user template by specifying the `--name` option, or you must associate the template with a file type by using the `--file-type` option. You cannot specify both options simultaneously.
- When a template includes a storage pool specification, it does not require a media type setting. If a storage pool and media type are both specified, the settings must be compatible. For example, a template cannot specify the use of high capacity (HC) storage media and also specify the name of a storage pool that uses extreme flash (EF) storage.
- You can view a list of Exascale file types by using the ESCLI lstemplate command. For example:

```
@> lstemplate --cluster
```

```
@> lstemplate --vault vault2 --all
```

Examples

Example 6-135 Create a User-Defined Cluster-Wide Template

This example shows creating a user-defined template named `T1`. The template is associated with the cluster, as neither `--vault` nor `--cluster` is specified, and the current working directory in the ESCLI session is the root directory.

```
@> mktemplate --name T1 --content-type DATA --media-type HC --redundancy high
```

Example 6-136 Create a User-Defined Vault-Specific Template

This example shows creating a user-defined template named `T2`. The template is associated with the vault named `vault2`, as neither `--vault` nor `--cluster` is specified, and the current working directory in the ESCLI session is inside `vault2`.

```
@vault2> mktemplate --name T2 --content-type DATA --media-type HC --
redundancy high
```

Example 6-137 Create a User-Defined Vault-Specific Template

This example shows creating a user-defined template named `T3` that is associated with the vault named `vault2`.

```
@> mktemplate --vault vault2 --name T3 --content-type DATA --media-type HC --
redundancy high
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Example 6-138 Create a User-Defined Cluster-Wide Template

This example shows creating a user-defined cluster-wide template named `T4`. The `--cluster` option overrides the fact that the current working directory in the ESCLI session is inside `vault2`.

```
@vault2> mktemplate --cluster --name T4 --content-type DATA --media-type HC --
redundancy high
```

Example 6-139 Create a Template for a Specific Vault and File Type

This example shows creating a template associated with the `DATAFILE` file type and the vault named `vault2`.

```
@> mktemplate --vault vault2 --file-type DATAFILE --content-type DATA --media-
type HC --redundancy high
```

Example 6-140 Create a Cluster-Wide Template for a Specific File Type

This example shows creating a cluster-wide template associated with the `DATAFILE` file type.

```
@> mktemplate --cluster --file-type DATAFILE --content-type DATA --media-type
HC --redundancy high
```

Related Topics

- About Exascale File Storage Attributes
- lstemplate
  List file templates.

6.2.9.4 rmtemplate

Delete a file template.

Purpose

The `rmtemplate` command deletes an Exascale file template that is associated with the vault or the cluster.

Syntax

```
rmtemplate [ --vault vault | --cluster ]
           { --file-type file-type | --name template-name }
```

Command Options

Specify one or more of the following `rmtemplate` command options to identify the template to delete:

- `--vault`: Specifies a template associated with the specified vault.
- `--cluster`: Specifies a template associated with the cluster.
- `--file-type`: Specifies a template associated with the specified file type.
- `--name`: Specifies a user-defined template having the specified name.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Usage Notes

Note the following information when using this command:

- File type templates associated with the cluster cannot be deleted.

Examples

Example 6-141 Delete a User-Defined Cluster-Wide Template

This example shows deleting a user-defined template named `T1`. The template is associated with the cluster, as neither `--vault` nor `--cluster` is specified, and the current working directory in the ESCLI session is the root directory.

```
@> rmtemplate --name T1
```

Example 6-142 Delete a User-Defined Vault-Specific Template

This example shows deleting a user-defined template named `T2`. The template is associated with the vault named `vault2`, as neither `--vault` nor `--cluster` is specified, and the current working directory in the ESCLI session is inside `vault2`.

```
@vault2> rmtemplate --name T2
```

Example 6-143 Delete a User-Defined Vault-Specific Template

This example shows deleting a user-defined template named `T3` that is associated with the vault named `vault2`.

```
@> rmtemplate --name T3 --vault vault2
```

Example 6-144 Delete a User-Defined Cluster-Wide Template

This example shows deleting a user-defined cluster-wide template named `T4`. The `--cluster` option overrides the fact that the current working directory in the ESCLI session is inside `vault2`.

```
@vault2> mktemplate --name T4 --cluster
```

Example 6-145 Delete a Template for a Specific Vault and File Type

This example shows deleting a template associated with the `DATAFILE` file type and the vault named `vault2`.

```
@> rmtemplate --file-type DATAFILE --vault vault2
```

6.2.10 Resource Profile Management

This section contains references for the Exascale command line interface (ESCLI) commands that are associated with resource profile management:

- chresourceprofile
  Change a resource profile associated with a vault.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- lsresourceprofile
  List resource profiles.
- mkresourceprofile
  Create a resource profile.
- rmresourceprofile
  Delete a resource profile.

6.2.10.1 chresourceprofile

Change a resource profile associated with a vault.

Purpose

The `chresourceprofile` command changes the attribute settings for an existing resource profile that is associated with an Exascale vault.

Syntax

```
chresourceprofile [vault-name/]resource-profile-name
                  --attributes attribute=value[,attribute=value]...
```

Command Options

The options for the `chresourceprofile` command are:

- *resource-profile-name*: Identifies the resource profile being changed.
- *vault-name*: Identifies the vault that the resource profile is associated with.

  If not specified, then the current working directory is used to identify the vault.
- `--attributes`: Specifies attributes to change. The following list outlines all of the available attributes:
  - `iopsShareEF`, `iopsShareHC`, `iopsShareXT`: Specifies the relative share of I/O bandwidth (IOPS) from extreme flash (EF), high capacity (HC), or extended (XT) storage that is available to each client associated with the resource profile. Each client's share is relative to all other client shares associated with the corresponding media type. A higher share value implies higher priority. The range of valid values is `1-100`, and the default value is `1`.
  - `iopsLimitEF`, `iopsLimitHC`, `iopsLimitXT`: Specifies the upper limit of the I/O bandwidth (IOPS) from extreme flash (EF), high capacity (HC), or extended (XT) storage that is available to each client associated with the resource profile.

    Starting with Oracle Exadata System Software release 26.1.0, each value is a percentage with up to two decimal places. The range of valid values is `0.01-100`, and the default value is `100` (effectively unlimited).

    Before Oracle Exadata System Software release 26.1.0, each value represents a fraction out of 10000. The range of valid values is `1-10000`, and the default value is `10000` (effectively unlimited).
  - `enableFlashCache`, `enableXrmemCache`, `enableFlashLog`: Enables or disables use of the flash cache, XRMEM cache, or flash log for clients associated with the resource profile. Each value is Boolean, and the default is `true` (enabled).

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `flashCacheMin`, `xrmemCacheMin`: Specifies the guaranteed minimum fraction of flash cache or XRMEM cache space available to each client associated with the resource profile.

  Starting with Oracle Exadata System Software release 26.1.0, each value is nominally a percentage with up to two decimal places. The range of valid values is `0-100`, and the default value is `0` (no set minimum).

  Before Oracle Exadata System Software release 26.1.0, each value represents a fraction out of 10000. The range of valid values is `0-10000`, and the default value is `0` (no set minimum).

  For each cache type, if the sum of all values across all resource profiles exceeds the maximum, then all the values are proportionally scaled down. Each value applies only when the corresponding cache is enabled in the resource profile (for example, `enableFlashCache=true`).

- `flashCacheSize` / `flashCacheMax`, `xrmemCacheSize` / `xrmemCacheMax`: Specifies the maximum fraction of flash cache or XRMEM cache space available to each client associated with the resource profile. Starting with Oracle Exadata System Software release 26.1.0, the attribute names are `flashCacheSize` and `xrmemCacheSize`. Previously, they were `flashCacheMax` and `xrmemCacheMax`.

  Starting with Oracle Exadata System Software release 26.1.0, each value is nominally a percentage with up to two decimal places. The range of valid values is `0-100`, and the default value is `100`.

  Before Oracle Exadata System Software release 26.1.0, each value represents a fraction out of 10000. The range of valid values is `0-10000`, and the default value is `10000`.

  Each value applies only when the corresponding cache is enabled in the resource profile (for example, `enableFlashCache=true`).

Usage Notes

In addition to regular user-defined resource profiles, you can also define a system-reserved resource profile named `$UNASSIGNED`. All Exascale clients not explicitly associated with a resource profile are automatically governed by the `$UNASSIGNED` profile. The `$UNASSIGNED` resource profile contains only two modifiable attributes:

- `flashCacheSize` / `flashCacheMax`: Specifies the maximum fraction of flash cache shared by clients associated with the `$UNASSIGNED` profile.

  Starting with Oracle Exadata System Software release 26.1.0, the attribute name is `flashCacheSize`. The value represents a percentage with up to two decimal places. The range of valid values is `0-100` and the default value is `100`.

  Previously, the attribute name was `flashCacheMax`. The range of valid values was `0-10000` and the default value was `10000`.

- `xrmemCacheSize` / `xrmemCacheMax`: Specifies the maximum fraction of XRMEM cache shared by clients associated with the `$UNASSIGNED` profile.

  Starting with Oracle Exadata System Software release 26.1.0, the attribute name is `xrmemCacheSize`. The value represents a percentage with up to two decimal places. The range of valid values is `0-100` and the default value is `100`.

  Previously, the attribute name was `xrmemCacheMax`. The range of valid values was `0-10000` and the default value was `10000`.

Copyright © 2022, 2026, Oracle and/or its affiliates.

All Exascale clients governed by the `$UNASSIGNED` profile share the corresponding cache resources. The behavior differs from regular resource profiles, where each application of the resource profile defines the resource allocation for one associated client.

You must create the `$UNASSIGNED` resource profile before you can modify it. If the `$UNASSIGNED` resource profile does not exist, the system automatically reserves 5% of each cache for unassigned Exascale clients, but unassigned clients share all of the unassigned cache space if more space is available.

Examples

Example 6-146 Change a Resource Profile

In this example, the `iopsShareHC` attribute is set to `11` for the `silver` resource profile in conjunction with the `vault2` vault.

```
@> chresourceprofile vault2/silver --attributes iopsShareHC=11
```

Example 6-147 Change a Resource Profile

In this example, the `enableFlashLog` attribute is set to `false` for the `silver` resource profile in conjunction with the `vault2` vault. In this example, the vault association is not specified in the command but is derived from the current working directory in the ESCLI session.

```
@vault2> chresourceprofile silver --attributes enableFlashLog=false
```

Related Topics

- About Exascale Resource Profiles

6.2.10.2 lsresourceprofile

List resource profiles.

Purpose

The `lsresourceprofile` command displays information about Exascale resource profiles that are associated with a vault.

Syntax

```
lsresourceprofile [ [vault-name/]resource-profile-name
                        [,[vault-name/]resource-profile-name ]... ]
                  [ -l ] [ --detail ]
                  [ --attributes attribute[,attribute]... ]
                  [ --filter filter[,filter]... ]
                  [ --sort [-]attribute[,[-]attribute]... ]
```

Command Options

The options for the `lsresourceprofile` command are:

- *resource-profile-name*: Identifies a specific resource profile name.
- *vault-name*: Identifies a specific vault.

  If not specified, then the current working directory is used to specify the vault.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `-l`: Displays output in a long, tabular format.
- `--detail`: Displays detailed output with additional attributes.
- `--attributes`: Identifies specific attributes to display in the output.

  To see a list of all available attributes and their descriptions, use the `describe resourceprofiles` command. See also Describing Resources and Attributes.
- `--filter`: Specifies conditions for including items in the output.
- `--sort`: Sorts the output using the specified attributes.

Usage Notes

Note the following information when using this command:

- Filter conditions are specified as: <*attribute*><*operator*><*value*>.

  The allowed operators are =, !=, >=, <=, >, and <.

  Multiple comma-separated filter conditions are combined using AND logic.
- Sorting attributes are specified as: `[-]`<*attribute*>. Multiple sort attributes are comma-delimited. The default sort order is ascending. For descending sort order, prefix the attribute name with `-`.

  For example, use the following to primarily sort by `name` in descending order, and use *otherAttribute* in ascending order to further sort entries with the same `name`:

  ```
  --sort -name,otherAttribute
  ```

Examples

Example 6-148 List Resource Profiles

This example lists the resource profiles associated with the vault named `vault2`.

```
@> lsresourceprofile --filter vault=vault2
```

Example 6-149 List Resource Profiles

This example lists the resource profiles associated with the vault named `vault2`. In this example, the vault association is not specified in the command but is derived from the current working directory in the ESCLI session.

```
@vault2> lsresourceprofile
```

Example 6-150 List a Specific Resource Profile

This example shows two ways to display detailed information about the resource profile named `silver` that is associated with the vault named `vault2`.

```
@> lsresourceprofile vault2/silver --detail
```

```
@> lsresourceprofile silver --filter vault=vault2 --detail
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

6.2.10.3 mkresourceprofile

Create a resource profile.

Purpose

The `mkresourceprofile` command creates a resource profile.

Syntax

```
mkresourceprofile [vault-name/]resource-profile-name
                  [ --attributes attribute=value[,attribute=value]... ]
```

Command Options

The options for the `mkresourceprofile` command are:

- *resource-profile-name*: Specifies the name for the resource profile being created.
- *vault-name*: Specifies the vault that the resource profile is associated with.

  If not specified, then the current working directory is used to specify the vault.
- `--attributes`: Optionally specifies values for attributes of the resource profile. The following list outlines all of the available attributes:
  - `iopsShareEF`, `iopsShareHC`, `iopsShareXT`: Specifies the relative share of I/O bandwidth (IOPS) from extreme flash (EF), high capacity (HC), or extended (XT) storage that is available to each client associated with the resource profile. Each client's share is relative to all other client shares associated with the corresponding media type. A higher share value implies higher priority. The range of valid values is `1-100`, and the default value is `1`.
  - `iopsLimitEF`, `iopsLimitHC`, `iopsLimitXT`: Specifies the upper limit of the I/O bandwidth (IOPS) from extreme flash (EF), high capacity (HC), or extended (XT) storage that is available to each client associated with the resource profile.

    Starting with Oracle Exadata System Software release 26.1.0, each value is a percentage with up to two decimal places. The range of valid values is `0.01-100`, and the default value is `100` (effectively unlimited).

    Before Oracle Exadata System Software release 26.1.0, each value represents a fraction out of 10000. The range of valid values is `1-10000`, and the default value is `10000` (effectively unlimited).
  - `enableFlashCache`, `enableXrmemCache`, `enableFlashLog`: Enables or disables use of the flash cache, XRMEM cache, or flash log for clients associated with the resource profile. Each value is Boolean, and the default is `true` (enabled).
  - `flashCacheMin`, `xrmemCacheMin`: Specifies the guaranteed minimum fraction of flash cache or XRMEM cache space available to each client associated with the resource profile.

    Starting with Oracle Exadata System Software release 26.1.0, each value is nominally a percentage with up to two decimal places. The range of valid values is `0-100`, and the default value is `0` (no set minimum).

    Before Oracle Exadata System Software release 26.1.0, each value represents a fraction out of 10000. The range of valid values is `0-10000`, and the default value is `0` (no set minimum).

Copyright © 2022, 2026, Oracle and/or its affiliates.

For each cache type, if the sum of all values across all resource profiles exceeds the maximum, then all the values are proportionally scaled down. Each value applies only when the corresponding cache is enabled in the resource profile (for example, `enableFlashCache=true`).

- `flashCacheSize` / `flashCacheMax`, `xrmemCacheSize` / `xrmemCacheMax`: Specifies the maximum fraction of flash cache or XRMEM cache space available to each client associated with the resource profile. Starting with Oracle Exadata System Software release 26.1.0, the attribute names are `flashCacheSize` and `xrmemCacheSize`. Previously, they were `flashCacheMax` and `xrmemCacheMax`.

  Starting with Oracle Exadata System Software release 26.1.0, each value is nominally a percentage with up to two decimal places. The range of valid values is `0-100`, and the default value is `100`.

  Before Oracle Exadata System Software release 26.1.0, each value represents a fraction out of 10000. The range of valid values is `0-10000`, and the default value is `10000`.

  Each value applies only when the corresponding cache is enabled in the resource profile (for example, `enableFlashCache=true`).

Usage Notes

In addition to regular user-defined resource profiles, you can also define a system-reserved resource profile named `$UNASSIGNED`. All Exascale clients not explicitly associated with a resource profile are automatically governed by the `$UNASSIGNED` profile. The `$UNASSIGNED` resource profile contains only two modifiable attributes:

- `flashCacheSize` / `flashCacheMax`: Specifies the maximum fraction of flash cache shared by clients associated with the `$UNASSIGNED` profile.

  Starting with Oracle Exadata System Software release 26.1.0, the attribute name is `flashCacheSize`. The value represents a percentage with up to two decimal places. The range of valid values is `0-100` and the default value is `100`.

  Previously, the attribute name was `flashCacheMax`. The range of valid values was `0-10000` and the default value was `10000`.

- `xrmemCacheSize` / `xrmemCacheMax`: Specifies the maximum fraction of XRMEM cache shared by clients associated with the `$UNASSIGNED` profile.

  Starting with Oracle Exadata System Software release 26.1.0, the attribute name is `xrmemCacheSize`. The value represents a percentage with up to two decimal places. The range of valid values is `0-100` and the default value is `100`.

  Previously, the attribute name was `xrmemCacheMax`. The range of valid values was `0-10000` and the default value was `10000`.

All Exascale clients governed by the `$UNASSIGNED` profile share the corresponding cache resources. The behavior differs from regular resource profiles, where each application of the resource profile defines the resource allocation for one associated client.

If the `$UNASSIGNED` resource profile does not exist, the system automatically reserves 5% of each cache for unassigned Exascale clients, but unassigned clients share all of the unassigned cache space if more space is available.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Examples

Example 6-151 Create a Resource Profile with Default Attribute Values

This example shows two ways to create a resource profile with default attribute values, which is named `silver` and is associated with the vault named `vault2`.

```
@> mkresourceprofile vault2/silver
```

```
@> mkresourceprofile silver --attributes vault=vault2
```

Example 6-152 Create a Resource Profile with Default Attribute Values

This example shows how to create a resource profile with default attribute values, which is named `bronze` and is associated with the vault named `vault2`. In this example, the vault association is not specified in the command but is derived from the current working directory in the ESCLI session.

```
@vault2> mkresourceprofile bronze
```

Example 6-153 Create a Resource Profile with Specific Attribute Values

This example shows how to create a resource profile that contains specific attribute settings.

```
@> mkresourceprofile vault2/gold --attributes
iopsShareHC=11,enableFlashLog=false
```

Related Topics

- About Exascale Resource Profiles

6.2.10.4 rmresourceprofile

Delete a resource profile.

Purpose

The `rmresourceprofile` command deletes a resource profile that is associated with an Exascale vault.

Syntax

```
rmresourceprofile [vault-name/]resource-profile-name
```

Command Options

The options for the `rmresourceprofile` command are:

- *resource-profile-name*: Identifies the name for the resource profile to delete.
- *vault-name*: Identifies the vault that the resource profile is associated with.

  If not specified, then the current working directory is used to identify the vault.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Examples

Example 6-154 Delete a Resource Profile

This example shows how to delete a resource profile named `silver` that is associated with the vault named `vault2`.

```
@> rmresourceprofile vault2/silver
```

Example 6-155 Delete a Resource Profile

This example shows how to delete a resource profile named `bronze` that is associated with the vault named `vault2`. In this example, the vault association is not specified in the command but is derived from the current working directory in the ESCLI session.

```
@vault2> rmresourceprofile bronze
```

6.2.11 Extended Attribute Management

This section contains references for the Exascale command line interface (ESCLI) commands that are associated with extended attribute management:

- chxattr
  Set an extended attribute for a file or vault.
- lsxattr
  List an extended attribute for a file or vault.
- rmxattr
  Delete an extended attribute for a file or vault.

6.2.11.1 chxattr

Set an extended attribute for a file or vault.

Purpose

The `chxattr` command allows you to set an extended attribute for an Exascale file or vault.

Syntax

```
chxattr { file-name | vault-name } --name attribute-name
        { --value attribute-value | --bin-value-file file-name }
```

Command Options

The options for the `chxattr` command are:

- `{ file-name | vault-name }`: Specifies the name of the file or vault that is the subject of the operation.
- `--name`: Specifies the name of the extended attribute.
- `--value`: Specifies the value of the extended attribute.
- `--bin-value-file`: Specifies the regular file from which to read the value of the extended attribute. The value is assumed to be binary.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Examples

Example 6-156 Set an Extended Attribute for a Vault

You can add an extended attribute to a vault, such as `@my-data`. The name of the attribute is `custom-attr`, and its value is `my-value`.

```
@> chxattr @my-data --name custom-attr --value my-value
```

Example 6-157 Set an Extended Attribute for a File

You can add an extended attribute to a file, such as `@my-data/file1`. The name of the attribute is `custom-attr`, and its value is `my-value`.

```
@> chxattr @my-data/file1 --name custom-attr --value myv-alue
```

Example 6-158 Set a Binary Extended Attribute for a File

You can add a binary extended attribute to a file, such as `@my-data/file1`. The name of the attribute is `binary-attr`, and its value is the contents of the file at `/home/user/bin-val-in.dat`.

```
@> chxattr @my-data/file1 --name binary-attr --bin-value-file /home/user/bin-val-in.dat
```

6.2.11.2 lsxattr

List an extended attribute for a file or vault.

Purpose

The `lsxattr` command displays information about non-standard attributes created for Exascale files and vaults.

Syntax

```
lsxattr { file-name | vault-name } --name attribute-name
        [ --bin-value-file file-name ]
```

Command Options

The options for the `lsxattr` command are:

- `{ file-name | vault-name }`: Specifies the name of the file or vault for which you want to list the specified extended attribute value.
- `--name`: Specifies the name of the extended attribute.
- `--bin-value-file`: Specifies the file where you want to write the value of the extended attribute.

Usage Notes

If you do not specify the `--bin-value-file` option, binary values are output in Base64 format.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Examples

Example 6-159 List an Extended Attribute

You can use the following commands to list the `custom-attr` extended attribute.

```
@> lsxattr @my-data --name custom-attr


@> lsxattr @my-data/file1 --name custom-attr
```

Example 6-160 List a Binary Extended Attribute

You can use the following commands to list the `binary-attr` extended attribute and write the binary value to a regular file at `/home/user/bin-val-out.dat`.

```
@> lsxattr @my-data/file1 --name binary-attr --bin-value-file /home/user/bin-
val-out.dat
```

6.2.11.3 rmxattr

Delete an extended attribute for a file or vault.

Purpose

The `rmxattr` command deletes non-standard attributes created for Exascale files and vaults.

Syntax

```
rmxattr { file-name | vault-name } --name attribute-name
```

Command Options

The options for the `rmxattr` command are:

- `{ file-name | vault-name }`: Specifies the name of the file or vault for which you want to delete the specified extended attribute value.
- `--name`: Specifies the name of the extended attribute.

Examples

Example 6-161 Delete an Extended Attribute

You can use the following commands to delete the `custom-attr` extended attribute.

```
@> rmxattr @my-data --name custom-attr


@> rmxattr @my-data/file1 --name custom-attr
```

6.2.12 Block Store Management

This section contains references for the Exascale command line interface (ESCLI) commands that are associated with block store management:

Copyright © 2022, 2026, Oracle and/or its affiliates.

- acfsctl
  Control registration for Exascale-managed ACFS file systems.
- chacfsfilesystem
  Change attributes of an Exascale-managed ACFS file system.
- chvolume
  Change attributes of a volume.
- chvolumeattachment
  Change attributes of a volume attachment.
- chvolumebackup
  Change attributes of a volume backup.
- chvolumegroup
  Change attributes of a volume group.
- chvolumesnapshot
  Change attributes of a volume snapshot.
- lsacfsfilesystem
  List Exascale-managed ACFS file systems.
- lsinitiator
  List EDV initiator information.
- lsvolume
  List volumes.
- lsvolumeattachment
  List volume attachments.
- lsvolumebackup
  List volume backups.
- lsvolumegroup
  List volume groups.
- lsvolumesnapshot
  List volume snapshots.
- mkacfsfilesystem
  Create an Exascale-managed ACFS file system.
- mkvolume
  Create a volume or volume clone.
- mkvolumeattachment
  Create an attachment for an Exascale volume.
- mkvolumebackup
  Create a backup of a volume snapshot.
- mkvolumegroup
  Create a volume group.
- mkvolumegroupsnapshot
  Create snapshots of all volumes in the specified volume group.
- mkvolumesnapshot
  Create a snapshot of a volume.
- rmacfsfilesystem
  Delete an Exascale-managed ACFS file system.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- rmvolume
  Delete volumes.
- rmvolumeattachment
  Delete a volume attachment.
- rmvolumebackup
  Delete a volume backup.
- rmvolumegroup
  Delete volume groups.
- rmvolumesnapshot
  Delete volume snapshots.

6.2.12.1 acfsctl

Control registration for Exascale-managed ACFS file systems.

Purpose

The `acfsctl` command controls registration for Exascale-managed implementations of Oracle Advanced Cluster File System (ACFS).

Syntax

```
acfsctl register volume-ID mount-path
        [ --attributes attribute=value[,attribute=value]... ]


acfsctl deregister acfs-ID [ --force ]
```

Command Options

The options for the `acfsctl register` command are:

- *volume-ID*: Identifies the volume that hosts the file system. You can use the lsacfsfilesystem command to find the volume identifier associated with each Exascale-managed ACFS file system.
- *mount-path*: Specifies the file system mount path.
- `--attributes`: Optionally specifies attributes governing how the file system is mounted. The permitted attributes are:
  - `mountLeafOwner`, `mountLeafGroup`, `mountLeafMode`: Specify the ownership and permission attributes applied to the leaf directory in the file system mount path.
  - `mountNoRootSUID`: Controls whether non-root users are permitted to run root-owned binaries with the set user ID (setuid) permission.
    - `true`: Only the root user can run root-owned setuid binaries.
    - `false`: Non-root users can also run root-owned setuid binaries.

    If not specified, the default value is `false`.
  - `mountReadOnly`: Specifies whether the file system is mounted with read-only access.
    - `true`: Mounts the file system with read-only access.
    - `false`: Mounts the file system with read-write access.

Copyright © 2022, 2026, Oracle and/or its affiliates.

If not specified, the default value is `false`.

You cannot set any other file system attributes during registration.

Use the `describe acfsfilesystems` to see further details about all Exascale-managed ACFS file system attributes. See also Describing Resources and Attributes.

The options for the `acfsctl deregister` command are:

- *acfs-ID*: Identifies the ACFS file system being deregistered. You can use the lsacfsfilesystem command to find the identifier for each Exascale-managed ACFS file system.
- `--force`: Optionally ignores any errors and forces deregistration of the file system, even if it is in use.

Usage Notes

Note the following information when using the `acfsctl register` command:

- The specified volume (*volume-ID*) must be associated with an existing cluster-wide EDV attachment and contain an existing ACFS file system; otherwise, the command fails. Registration of a volume with a node-specific EDV attachment is not permitted.
- The ACFS file system is mounted on every node in the Oracle Grid Infrastructure (GI) cluster and the ACFS details are registered with the GI cluster. After registration, the file system is automatically mounted (or remounted) by Oracle Clusterware as required.

Note the following when using the `acfsctl deregister` command:

- The ACFS file system is dismounted and the mount point directory is removed on every node in the Oracle Grid Infrastructure (GI) cluster. The command also removes the ACFS file system registration in the GI cluster.

6.2.12.2 chacfsfilesystem

Change attributes of an Exascale-managed ACFS file system.

Purpose

The `chacfsfilesystem` command allows you to modify the attributes of an Exascale-managed implementation of Oracle Advanced Cluster File System (ACFS).

Syntax

```
chacfsfilesystem acfs-ID [ --attributes attribute=value[,attribute=value]... ]
```

Command Options

The options for the `chacfsfilesystem` command are:

- *acfs-ID*: Identifies the ACFS file system being changed. You can use the lsacfsfilesystem command to find the identifier for each ACFS file system.
- `--attributes`: Optionally specifies attributes to change.

  Use the `describe chacfsfilesystem` command to see details about the ACFS file system attributes you can change. See also Describing Resources and Attributes.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Examples

Example 6-162 Change the Quarantine State of an ACFS File System

This example shows changing the quarantine state of the specified ACFS file system to `false`.

```
@> chacfsfilesystem acfs0003_50e52177583f4be4bad68ac20b65001e --attributes
quarantined=false
```

6.2.12.3 chvolume

Change attributes of a volume.

Purpose

The `chvolume` command allows you to modify the attributes of an Exascale volume.

Syntax

```
chvolume volume-ID [ --attributes attribute=value[,attribute=value]... ]
```

Command Options

The options for the `chvolume` command are:

- *volume-ID*: Identifies the volume being changed. You can use the lsvolume command to find the identifier for each volume.
- `--attributes`: Optionally specifies attributes to change.

  Use the `describe chvolume` command to see details about the volume attributes you can change. See also Describing Resources and Attributes.

Usage Notes

- The `chvolume` command only proceeds if the specified volume is not attached.
- A volume can have up to two owners. To modify volume ownership, you can add or remove one volume owner at a time.

  To add a volume owner, specify `owners=+user-ID` as an attribute setting. To remove a volume owner, use `owners=-user-ID`.

  For example, see Example 6-164.

Examples

Example 6-163 Change the Volume Size

This example shows changing the size of the volume with the volume ID `vol0002_50e52177583f4be4bad68ac20b65001e`.

```
@> chvolume vol0002_50e52177583f4be4bad68ac20b65001e --attributes size=200m
```

Example 6-164 Change the Volume Owners

The following examples show different ways to change the volume owners.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- The following command adds `peter` as an owner of the specified volume:

  ```
  @> chvolume vol0002_50e52177583f4be4bad68ac20b65001e --attributes
  owners=+peter
  ```

- The following command removes `dave` as an owner of the specified volume:

  ```
  @> chvolume vol0002_50e52177583f4be4bad68ac20b65001e --attributes owners=-
  dave
  ```

6.2.12.4 chvolumeattachment

Change attributes of a volume attachment.

Purpose

The `chvolumeattachment` command allows you to modify the attributes of an Exascale volume attachment.

Syntax

```
chvolumeattachment attachment-ID
                   [ --attributes attribute=value[,attribute=value]... ]
```

Command Options

The options for the `chvolumeattachment` command are:

- *attachment-ID*: Identifies the volume attachment being changed. You can use the lsvolumeattachment command to find the identifier for each volume attachment.
- `--attributes`: Optionally specifies attributes to change.

  Use the `describe ch*volumeattachment` command to see details about the volume attachment attributes you can change. See also Describing Resources and Attributes.

Examples

Example 6-165 Change the Quarantine State of a Volume Attachment

This example shows changing the quarantine state of the specified volume attachment to `false`.

```
@> chvolumeattachment att0002_50e52177583f4be4bad68ac20b65001e --attributes
quarantined=false
```

6.2.12.5 chvolumebackup

Change attributes of a volume backup.

Purpose

The `chvolumebackup` command allows you to modify the attributes of an Exascale volume backup.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Syntax

```
chvolumebackup volume-backup-ID
               [ --attributes attribute=value[,attribute=value]... ]
```

Command Options

The options for the `chvolumebackup` command are:

- *volume-backup-ID*: Identifies the volume backup being changed. You can use the lsvolumebackup command to find the identifier for each volume backup.
- `--attributes`: Optionally specifies attributes to change.

  Use the `describe chvolumebackup` command to see details about the volume backup attributes you can change. See also Describing Resources and Attributes.

Examples

Example 6-166 Change the Name of a Volume Backup

This example shows changing the name of the specified volume backup to `my-volume-backup`.

```
@> chvolumebackup vol0002_bkp0001_50e52177583f4be4bad68ac20b65001e --
attributes name=my-volume-backup
```

6.2.12.6 chvolumegroup

Change attributes of a volume group.

Purpose

The `chvolumegroup` command allows you to modify the attributes of an Exascale volume group.

Syntax

```
chvolumegroup volume-group-ID
              [ --attributes attribute=value[,attribute=value]... ]
```

Command Options

The options for the `chvolumegroup` command are:

- *volume-group-ID*: Identifies the volume group being changed. You can use the lsvolumegroup command to find the identifier for each volume group.
- `--attributes`: Optionally specifies attributes to change.

  The following attributes can be changed:

  - `name`: Specifies the name of the volume group.
  - `resourceSharing`: Optionally specifies how provisioned resources are shared by volumes in the group. Permitted values are:

    * `aggregate`: Configures the volume group to share the I/Os per second (IOPS) bandwidth provisioned to each volume in the group. In this case, a volume can

Copyright © 2022, 2026, Oracle and/or its affiliates.

exceed its provisioned limit by sharing unused IOPS provisioned to other volumes in the group.

For example, consider a group with two volumes that are provisioned with 1000 IOPS each. In this example, aggregate resource sharing enables one volume to use 2000 IOPS when the other volume is idle.

* `specified`: Configures the volume group to share the I/Os per second (IOPS) bandwidth limit specified by the `iopsProvisioned` attribute.
* `none`: Specifies that the volume group is not a resource-sharing volume group. In this case, each volume is only subject to resource limits imposed at other levels, such as the individual volume level or the Exascale vault level.

– `iopsProvisioned`: Optionally specifies the I/O bandwidth limit provisioned to the volume group, expressed in I/Os per second (IOPS).

– `owners`: Modifies the volume group owners.

A volume group can have up to two owners. To modify volume group ownership, you can add or remove one owner at a time.

To add an owner, specify `owners=+`*`user-ID`* as an attribute setting. To remove an owner, use `owners=-`*`user-ID`*.

For example, see Example 6-167.

Usage Notes

- When resource sharing is disabled on a volume group (`resourceSharing=none`), the I/O bandwidth provisioned to the volume group must be zero (`iopsProvisioned=0`).
- When a volume group uses aggregate resource sharing (`resourceSharing=aggregate`), the I/O bandwidth provisioned to the volume group must be zero (`iopsProvisioned=0`).
- When a volume group employs specified resource sharing (`resourceSharing=specified`), the `iopsProvisioned` attribute value must be greater than zero.
- A volume must always have at least one owner in common with each volume group it belongs to. Any modification violating this requirement is not permitted.

Examples

Example 6-167 Change the Volume Group Owners

The following examples show different ways to change the volume group owners.

- The following command adds `peter` as an owner of the specified volume group:

```
@> chvolumegroup volgrp0001_08a99eb7e8ff41499dee41de75af3e62 --attributes
owners=+peter
```

- The following command removes `dave` as an owner of the specified volume group:

```
@> chvolumegroup volgrp0001_08a99eb7e8ff41499dee41de75af3e62 --attributes
owners=-dave
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

6.2.12.7 chvolumesnapshot

Change attributes of a volume snapshot.

Purpose

The `chvolumesnapshot` command allows you to modify the attributes of an Exascale volume snapshot.

Syntax

```
chvolumesnapshot volume-snapshot-ID
                 [ --attributes attribute=value[,attribute=value]... ]
```

Command Options

The options for the `chvolumesnapshot` command are:

- *volume-snapshot-ID*: Identifies the volume snapshot being changed. You can use the lsvolumesnapshot command to find the identifier for each volume snapshot.
- `--attributes`: Optionally specifies attributes to change.

  Use the `describe chvolumesnapshot` command to see details about the volume snapshot attributes you can change. See also Describing Resources and Attributes.

Examples

Example 6-168 Change the Name of a Volume Snapshot

This example shows changing the name of the specified volume snapshot to `vol-snap1`.

```
@> chvolumesnapshot vol0001_snap0001_50e52177583f4be4bad68ac20b65001e --
attributes name=vol-snap1
```

6.2.12.8 lsacfsfilesystem

List Exascale-managed ACFS file systems.

Purpose

The `lsacfsfilesystem` command displays information about Exascale-managed implementations of Oracle Advanced Cluster File System (ACFS).

Syntax

```
lsacfsfilesystem [ acfs-ID [ acfs-ID ]... ] [ -l ] [ --detail ]
                 [ --attributes attribute[,attribute]... ]
                 [ --filter filter[,filter]... ]
                 [ --sort [-]attribute[,[-]attribute]... ]
                 [ --count value ]
```

Command Options

The options for the `lsacfsfilesystem` command are:

Copyright © 2022, 2026, Oracle and/or its affiliates.

- *acfs-ID*: Identifies the ACFS file system that you want to list information about. If not specified, the command displays information about all file systems.
- `-l`: Displays output in a long, tabular format.
- `--detail`: Displays detailed output with additional attributes.
- `--attributes`: Identifies specific attributes to display in the output.

  To see a list of all available attributes and their descriptions, use the `describe acfsfilesystems` command. See also Describing Resources and Attributes.
- `--filter`: Specifies conditions for including items in the output.
- `--sort`: Sorts the output using the specified attributes.
- `--count`: Specifies the maximum number of items to display in the output.

Usage Notes

Note the following information when using this command:

- Filter conditions are specified as: <*attribute*><*operator*><*value*>.

  The allowed operators are =, !=, >=, <=, >, and <.

  Multiple comma-separated filter conditions are combined using AND logic.

  Dates can be specified using the following formats:

  – `yyyy-MM-dd''T''HH:mm:ss`
  – `yyyy-MM-dd` (Time is assumed to be 00:00 AM)
  – `HH:mm:ss` (Date is assumed to be today)

  A date can also be followed by a timezone specification.

  Sizes can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.
- Sorting attributes are specified as: `[-]`<*attribute*>. Multiple sort attributes are comma-delimited. The default sort order is ascending. For descending sort order, prefix the attribute name with `-`.

  For example, use the following to primarily sort by `name` in descending order, and use *otherAttribute* in ascending order to further sort entries with the same `name`:

  ```
  --sort -name,otherAttribute
  ```

Examples

Example 6-169 List an Exascale-Managed ACFS File System

The following example shows how to list detailed information for an Exascale-managed ACFS file system by using an ACFS ID. In the example, the ACFS ID is `acfs0001_2fb06f13cddd4a8d8d636d1f794046cb`.

```
@> lsacfsfilesystem acfs0001_2fb06f13cddd4a8d8d636d1f794046cb --detail
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Example 6-170 List Specific Exascale-Managed ACFS File System Attributes

The following example shows how to list specific attributes of all Exascale-managed ACFS file systems. In the example, the attributes being listed are `name` and `mountPath`.

```
@> lsacfsfilesystem --attributes name,mountPath
```

6.2.12.9 lsinitiator

List EDV initiator information.

Purpose

The `lsinitiator` command displays information about Exascale Direct Volume (EDV) initiators.

Syntax

```
lsinitiator [ edv-initiator-ID ]
            [ --attributes attribute[,attribute]... ]
            [ --filter filter[,filter]... ]
            [ --detail ] [ -l ]
```

Command Options

The options for the `lsinitiator` command are:

- *edv-initiator-ID*: Identifies the EDV initiator that you want to list information about. If not specified, the command displays information about all EDV initiators.
- `--attributes`: Identifies specific attributes to display in the output.

  You can use this option to focus on a subset of attributes as an alternative to the `--detail` option.
- `--filter`: Specifies conditions for including items in the output.
- `--detail`: Displays detailed output with additional attributes.
- `-l`: Displays output in a long, tabular format.

Usage Notes

Note the following information when using this command:

- Filter conditions are specified as: <*attribute*><*operator*><*value*>.

  The allowed operators are =, !=, >=, <=, >, and <.

  Multiple comma-separated filter conditions are combined using AND logic.

  Dates can be specified using the following formats:

  - `yyyy-MM-dd''T''HH:mm:ss`
  - `yyyy-MM-dd` (Time is assumed to be 00:00 AM)
  - `HH:mm:ss` (Date is assumed to be today)

  A date can also be followed by a timezone specification.

  Sizes can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Examples

Example 6-171 List EDV Initiator Information

The following example shows how to list all EDV initiators in the Exascale cluster.

```
@> lsinitiator
```

Example 6-172 List Information for a Specific EDV Initiator

The following example shows how to list detailed information about the EDV initiator with the ID: `3d8c1b1c-2fa8-fc6e-3d8c-1b1c2fa8fc6e`.

```
@> lsinitiator 3d8c1b1c-2fa8-fc6e-3d8c-1b1c2fa8fc6e
```

6.2.12.10 lsvolume

List volumes.

Purpose

The `lsvolume` command displays information about Exascale volumes.

Syntax

```
lsvolume [ volume-ID [ volume-ID ]... ] [ -l ] [ --detail ]
         [ --attributes attribute[,attribute]... ]
         [ --filter filter[,filter]... ]
         [ --sort [-]attribute[,[-]attribute]... ]
         [ --count value ]
```

Command Options

The options for the `lsvolume` command are:

- *volume-ID*: Identifies an Exascale volume that you want to list information about. If not specified, the command displays information about all volumes.
- `-l`: Displays output in a long, tabular format.
- `--detail`: Displays detailed output with additional attributes.
- `--attributes`: Identifies specific attributes to display in the output.

  To see a list of all available attributes and their descriptions, use the `describe volumes` command. See also Describing Resources and Attributes.
- `--filter`: Specifies conditions for including items in the output.
- `--sort`: Sorts the output using the specified attributes.
- `--count`: Specifies the maximum number of items to display in the output.

Usage Notes

Note the following information when using this command:

- Filter conditions are specified as: <*attribute*><*operator*><*value*>.

Copyright © 2022, 2026, Oracle and/or its affiliates.

The allowed operators are =, !=, >=, <=, >, and <.

Multiple comma-separated filter conditions are combined using AND logic.

Dates can be specified using the following formats:

- `yyyy-MM-dd''T''HH:mm:ss`
- `yyyy-MM-dd` (Time is assumed to be 00:00 AM)
- `HH:mm:ss` (Date is assumed to be today)

A date can also be followed by a timezone specification.

Sizes can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.

- Sorting attributes are specified as: `[-]<attribute>`. Multiple sort attributes are comma-delimited. The default sort order is ascending. For descending sort order, prefix the attribute name with `-`.

  For example, use the following to primarily sort by `name` in descending order, and use *otherAttribute* in ascending order to further sort entries with the same `name`:

  ```
  --sort -name,otherAttribute
  ```

Examples

Example 6-173 List Information for a Specific Volume

The following example shows how to list detailed information for volume ID: `vol0002_50e52177583f4be4bad68ac20b65001e`.

```
@> lsvolume vol0002_50e52177583f4be4bad68ac20b65001e --detail
```

Example 6-174 List Specific Volume Attributes

The following example shows how to list specific volume attributes for volume ID: `vol0002_50e52177583f4be4bad68ac20b65001e`.

```
@> lsvolume vol0002_50e52177583f4be4bad68ac20b65001e --attributes
id,size,iopsProvisioned
```

Example 6-175 List Volumes Matching a Filter

The following example shows how to list detailed volume information for volumes that are sized greater than 100 MB. The output is also sorted by volume size.

```
@> lsvolume --detail --filter size>100M --sort size
```

6.2.12.11 lsvolumeattachment

List volume attachments.

Purpose

The `lsvolumeattachment` command displays information about Exascale volume attachments.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Syntax

```
lsvolumeattachment [ attachment-ID [ attachment-ID ]... ] [ -l ] [ --detail ]
                   [ --attributes attribute[,attribute]... ]
                   [ --filter filter[,filter]... ]
                   [ --sort [-]attribute[,[-]attribute]... ]
                   [ --count value ]
```

Command Options

The options for the `lsvolumeattachment` command are:

- *attachment-ID*: Identifies the volume attachment that you want to list information about. If not specified, the command displays information about all volume attachments.
- `-l`: Displays output in a long, tabular format.
- `--detail`: Displays detailed output with additional attributes.
- `--attributes`: Identifies specific attributes to display in the output.

  To see a list of all available attributes and their descriptions, use the `describe volumeattachments` command. See also Describing Resources and Attributes.
- `--filter`: Specifies conditions for including items in the output.
- `--sort`: Sorts the output using the specified attributes.
- `--count`: Specifies the maximum number of items to display in the output.

Usage Notes

Note the following information when using this command:

- Filter conditions are specified as: <*attribute*><*operator*><*value*>.

  The allowed operators are =, !=, >=, <=, >, and <.

  Multiple comma-separated filter conditions are combined using AND logic.

  Dates can be specified using the following formats:

  – `yyyy-MM-dd''T''HH:mm:ss`
  – `yyyy-MM-dd` (Time is assumed to be 00:00 AM)
  – `HH:mm:ss` (Date is assumed to be today)

  A date can also be followed by a timezone specification.

  Sizes can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.
- Sorting attributes are specified as: `[-]`<*attribute*>. Multiple sort attributes are comma-delimited. The default sort order is ascending. For descending sort order, prefix the attribute name with `-`.

  For example, use the following to primarily sort by `name` in descending order, and use *otherAttribute* in ascending order to further sort entries with the same `name`:

  ```
  --sort -name,otherAttribute
  ```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Examples

Example 6-176 List an EDV Attachment Using an Attachment ID

The following example shows how to list attachment information for an Exascale Direct Volume by using an attachment ID. In the example, the attachment ID is `att0003_50e52177583f4be4bad68ac20b65001e`.

```
@> lsvolumeattachment att0003_50e52177583f4be4bad68ac20b65001e
```

Example 6-177 List Specific EDV Attachment Attributes

The following example shows how to list specific attributes of all EDV attachments. In the example, the attributes being listed are `id` and `devicePath`.

```
@> lsvolumeattachment --attributes id,devicePath
```

6.2.12.12 lsvolumebackup

List volume backups.

Purpose

The `lsvolumebackup` command displays information about Exascale volume backups.

Syntax

```
lsvolumebackup [ volume-backup-ID [ volume-backup-ID ]... ]
               [ -l ] [ --detail ]
               [ --attributes attribute[,attribute]... ]
               [ --filter filter[,filter]... ]
               [ --sort [-]attribute[,[-]attribute]... ]
               [ --count value ]
```

Command Options

The options for the `lsvolumebackup` command are:

- *volume-backup-ID*: Identifies an Exascale volume backup that you want to list information about. If not specified, the command displays information about all volume backups.
- `-l`: Displays output in a long, tabular format.
- `--detail`: Displays detailed output with additional attributes.
- `--attributes`: Identifies specific attributes to display in the output.

  To see a list of all available attributes and their descriptions, use the `describe volumebackups` command. See also Describing Resources and Attributes.
- `--filter`: Specifies conditions for including items in the output.
- `--sort`: Sorts the output using the specified attributes.
- `--count`: Specifies the maximum number of items to display in the output.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Usage Notes

Note the following information when using this command:

- Filter conditions are specified as: <*attribute*><*operator*><*value*>.

  The allowed operators are =, !=, >=, <=, >, and <.

  Multiple comma-separated filter conditions are combined using AND logic.

  Dates can be specified using the following formats:

  – `yyyy-MM-dd''T''HH:mm:ss`
  – `yyyy-MM-dd` (Time is assumed to be 00:00 AM)
  – `HH:mm:ss` (Date is assumed to be today)

  A date can also be followed by a timezone specification.

  Sizes can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.

- Sorting attributes are specified as: `[-]`<*attribute*>. Multiple sort attributes are comma-delimited. The default sort order is ascending. For descending sort order, prefix the attribute name with `-`.

  For example, use the following to primarily sort by `name` in descending order, and use *otherAttribute* in ascending order to further sort entries with the same `name`:

  ```
  --sort -name,otherAttribute
  ```

Examples

Example 6-178 List Information for Volume Backups

The following example shows how to list detailed information about all volume backups.

```
@> lsvolumebackup --detail
```

Example 6-179 List Information for Specific Volume Backups

The following example shows how to list detailed information about two volume backups with specific IDs.

```
@> lsvolumebackup vol0002_bkp0001_fedd311081ee490481b7e19bdb691999
vol0002_bkp0002_fedd311081ee490481b8f2acab692334 --detail
```

Example 6-180 List Information for Specific Volume Backups

The following example shows how to list information about volume backups with names starting with `my-volume-backup`.

```
@> lsvolumebackup --filter name=my-volume-backup*
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

6.2.12.13 lsvolumegroup

List volume groups.

Purpose

The `lsvolumegroup` command displays information about Exascale volume groups.

Syntax

```
lsvolumegroup [ volume-group-ID [ volume-group-ID ]... ] [ -l ] [ --detail ]
              [ --attributes attribute[,attribute]... ]
              [ --filter filter[,filter]... ]
              [ --sort [-]attribute[,[-]attribute]... ]
              [ --count value ]
```

Command Options

The options for the `lsvolumegroup` command are:

- *volume-group-ID*: Identifies an Exascale volume group that you want to list information about. If not specified, the command displays information about all volume groups.
- `-l`: Displays output in a long, tabular format.
- `--detail`: Displays detailed output with additional attributes.
- `--attributes`: Identifies specific attributes to display in the output.

  To see a list of all available attributes and their descriptions, use the `describe volumegroups` command. See also Describing Resources and Attributes.
- `--filter`: Specifies conditions for including items in the output.
- `--sort`: Sorts the output using the specified attributes.
- `--count`: Specifies the maximum number of items to display in the output.

Usage Notes

Note the following information when using this command:

- Filter conditions are specified as: <*attribute*><*operator*><*value*>.

  The allowed operators are =, !=, >=, <=, >, and <.

  Multiple comma-separated filter conditions are combined using AND logic.

  Dates can be specified using the following formats:

  – `yyyy-MM-dd''T''HH:mm:ss`

  – `yyyy-MM-dd` (Time is assumed to be 00:00 AM)

  – `HH:mm:ss` (Date is assumed to be today)

  A date can also be followed by a timezone specification.
- Sorting attributes are specified as: `[-]`<*attribute*>. Multiple sort attributes are comma-delimited. The default sort order is ascending. For descending sort order, prefix the attribute name with `-`.

Copyright © 2022, 2026, Oracle and/or its affiliates.

For example, use the following to primarily sort by `name` in descending order, and use *otherAttribute* in ascending order to further sort entries with the same `name`:

```
--sort -name,otherAttribute
```

Examples

Example 6-181 List Information for Volume Groups

The following example shows how to list detailed information about all volume groups.

```
@> lsvolumegroup --detail
```

Example 6-182 List Information for Specific Volume Groups

The following example shows how to list detailed information about two volume groups with specific IDs.

```
@> lsvolumegroup volgrp0001_08a99eb7e8ff41499dee41de75af3e62
volgrp0002_67e399099f3845eeb2310485a13d1aee --detail
```

Example 6-183 List Information for Specific Volume Groups

The following examples show how to list information about volume groups matching the specified filter conditions.

```
@> lsvolumegroup --filter name=my-group*
```

```
@> lsvolumegroup --filter iopsProvisioned>1000
```

6.2.12.14 lsvolumesnapshot

List volume snapshots.

Purpose

The `lsvolumesnapshot` command displays information about Exascale volume snapshots.

Syntax

```
lsvolumesnapshot [ volume-snap-ID [ volume-snap-ID ]... ] [ -l ] [ --detail ]
                 [ --attributes attribute[,attribute]... ]
                 [ --filter filter[,filter]... ]
                 [ --sort [-]attribute[,[-]attribute]... ]
                 [ --count value ]
```

Command Options

The options for the `lsvolumesnapshot` command are:

- *volume-snap-ID*: Identifies an Exascale volume snapshot that you want to list information about. If not specified, the command displays information about all volume snapshots.
- `-l`: Displays output in a long, tabular format.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `--detail`: Displays detailed output with additional attributes.
- `--attributes`: Identifies specific attributes to display in the output.

  To see a list of all available attributes and their descriptions, use the `describe volumesnapshots` command. See also Describing Resources and Attributes.
- `--filter`: Specifies conditions for including items in the output.
- `--sort`: Sorts the output using the specified attributes.
- `--count`: Specifies the maximum number of items to display in the output.

Usage Notes

Note the following information when using this command:

- Filter conditions are specified as: <*attribute*><*operator*><*value*>.

  The allowed operators are =, !=, >=, <=, >, and <.

  Multiple comma-separated filter conditions are combined using AND logic.

  Dates can be specified using the following formats:

  – `yyyy-MM-dd''T''HH:mm:ss`
  – `yyyy-MM-dd` (Time is assumed to be 00:00 AM)
  – `HH:mm:ss` (Date is assumed to be today)

  A date can also be followed by a timezone specification.

  Sizes can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.
- Sorting attributes are specified as: `[-]`<*attribute*>. Multiple sort attributes are comma-delimited. The default sort order is ascending. For descending sort order, prefix the attribute name with `-`.

  For example, use the following to primarily sort by `name` in descending order, and use *otherAttribute* in ascending order to further sort entries with the same `name`:

  ```
  --sort -name,otherAttribute
  ```

Examples

Example 6-184 List Information for Volume Snapshots

The following example shows how to list detailed information about all of the volume snapshots.

```
@> lsvolumesnapshot --detail
```

Example 6-185 List Information for a Specific Volume Snapshot

The following example shows how to list detailed information for volume snapshot ID: `vol0002_snap0001_50e52177583f4be4bad68ac20b65001e`.

```
@> lsvolumesnapshot vol0002_snap0001_50e52177583f4be4bad68ac20b65001e
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Example 6-186 List Specific Attributes for Volume Snapshots

The following example shows how to list specific attributes of all the volume snapshots.

```
@> lsvolumesnapshot --attributes id,volume,size
```

6.2.12.15 mkacfsfilesystem

Create an Exascale-managed ACFS file system.

Purpose

The `mkacfsfilesystem` command creates an Exascale-managed implementation of Oracle Advanced Cluster File System (ACFS) on an Exascale volume using Exascale Direct Volumes (EDV).

Syntax

```
mkacfsfilesystem volume-ID mount-path [ --force ]
                 [ --attributes attribute=value[,attribute=value]... ]
```

Command Options

The options for the `mkacfsfilesystem` command are:

- *volume-ID*: Identifies the volume that hosts the file system. You can use the lsvolume command to find the identifier for each volume.
- *mount-path*: Specifies the file system mount path.
- `--force`: Optionally instructs the command to ignore and overwrite any existing file system on the volume.
- `--attributes`: Optionally specifies attributes for the file system.

  Use the `describe mkacfsfilesystem` command to see details about the ACFS file system attributes you can set. See also Describing Resources and Attributes.

Usage Notes

Note the following information when using this command:

- The specified volume (*volume-ID*) must be associated with a cluster-wide EDV attachment, otherwise the command fails. Using a volume with a node-specific EDV attachment is not permitted.
- The ACFS file system is mounted on every node in the Oracle Grid Infrastructure (GI) cluster and the ACFS details are registered with the GI cluster. After registration, the file system is automatically mounted (or remounted) by Oracle Clusterware as required.
- Before using the `mkacfsfilesystem` command, ensure that ACFS is configured appropriately on the target system.

Copyright © 2022, 2026, Oracle and/or its affiliates.

For instance, you can create a file system using ACFS encryption by specifying the `encryptionEnabled`, `encryptionAlgorithm`, and `encryptionKeyLength` attributes in the `mkacfsfilesystem` command. For example:

```
@> mkacfsfilesystem vol0001_bbd6fb4c75e2411b9bf366fe702eabaf /mnt/acfs1
    --attributes
encryptionEnabled=true,encryptionAlgorithm=AES,encryptionKeyLength=192
```

However, the command fails if ACFS encryption is not initialized on the target system. To initialize ACFS encryption, the system administrator must run the following command before the file system is created:

```
# acfsutil encr init
```

6.2.12.16 mkvolume

Create a volume or volume clone.

Purpose

The `mkvolume` command allows you to create a new Exascale volume, clone a volume snapshot, or directly clone an existing volume.

Syntax

To create a new empty volume:

```
mkvolume size --attributes vault=vault[,attribute=value]...
```

To create a new volume by restoring a volume backup:

```
mkvolume --attributes volumeBackup=volume-backup-ID,vault=vault
```

To clone a volume snapshot:

```
mkvolume --attributes volumeSnapshot=parent-snapshot-ID[,name=clone-name]
[,iopsProvisioned=integer-value][,iopsInherited={true|false}]
```

To directly clone an existing volume:

```
mkvolume --attributes volumeSource=parent-volume-ID[,name=clone-name]
[,iopsProvisioned=integer-value][,iopsInherited={true|false}]
```

> **Note**
>
> To directly clone an existing volume, you must use Oracle Exadata System Software release 25.1.0 or later.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Command Options

The options for the `mkvolume` command are:

- *size*: Specifies the size of the new empty volume. The *size* can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.

  > **Note**
  >
  > If you intend to use the volume to support Oracle ACFS, note that ACFS requires a minimum volume size of 512 MB.

- `--attributes`: Optionally specifies attributes for the volume.

  Use the `describe mkvolume` command to see details about the volume attributes you can set. See also Describing Resources and Attributes.

Usage Notes

The following notes apply only when the `mkvolume` command is used to create a new volume by restoring a volume backup.

- The new volume contains the contents of the volume backup specified by `volumeBackup=`*`volume-backup-ID`*.
- After creation, the new volume is completely separate from the volume backup and the original backup source volume. There is no ongoing association between the new volume and the volume backup or its original source.

The following notes apply only when the `mkvolume` command is used to clone a volume.

- When cloning a volume, either directly or using a volume snapshot, the resulting volume is a thin clone.
- When cloning a volume snapshot, `volumeSnapshot=`*`parent-snapshot-ID`* identifies the volume snapshot that you want to clone.
- When directly cloning an existing volume, `volumeSource=`*`parent-volume-ID`* identifies the volume that you want to clone.
- When cloning a volume, either directly or using a volume snapshot, you can optionally specify the following attribute settings. No additional command options or attribute settings are required or permitted.
  - `name=`*`clone-name`*: Optionally specifies the name of the volume clone, which makes it easier for you to identify later. If not specified, a system-generated name is assigned.
  - `iopsProvisioned=`*`integer-value`*: Optionally specifies the I/O bandwidth provisioned for the volume clone. The I/O bandwidth is expressed as the number of I/Os per second (IOPS).
  - `iopsInherited={true|false}`: Optionally specifies whether the volume clone inherits I/O bandwidth from the nearest ancestor in the volume hierarchy with provisioned (not inherited) I/O bandwidth.
- When cloning a volume, either directly or using a volume snapshot, the following describes the relationship between the `iopsProvisioned` and `iopsInherited` attribute settings:
  - If `iopsProvisioned` is unspecified and `iopsInherited` is not set to `false` (`iopsInherited` is unspecified or `iopsInherited=true`), then:

Copyright © 2022, 2026, Oracle and/or its affiliates.

    - The volume clone inherits I/O bandwidth from its nearest ancestor.
    - The volume clone has the setting: `iopsInherited=true`.
  - If `iopsProvisioned` is unspecified and `iopsInherited=false`, then:
    - The volume clone is provisioned with unlimited `iopsProvisioned`.
    - The volume clone has the setting: `iopsInherited=false`.
  - If you specify a value for `iopsProvisioned` and `iopsInherited` is not set to `true` (`iopsInherited` is unspecified or `iopsInherited=false`), then:
    - The volume clone is governed by the specified `iopsProvisioned` value.
    - The volume clone has the setting: `iopsInherited=false`.
  - An error occurs if you specify a value for `iopsProvisioned` and `iopsInherited=true`. This is an invalid combination.

Examples

Example 6-187 Create a Volume

This example shows creating a volume that is 100 MB in size and located in the vault named `my-volumes`.

```
@> mkvolume 100m --attributes vault=my-volumes
```

Example 6-188 Create a Volume with Attribute Settings

This example shows creating a volume that is 200 MB in size and located in the vault named `my-volumes`. The example also includes the following specific attribute settings:

- The volume name is set to `vol2`.
- The volume is provisioned with 1000 IOPS (IOs per second).

```
@> mkvolume 200m --attributes vault=my-volumes,name=vol2,iopsProvisioned=1000
```

6.2.12.17 mkvolumeattachment

Create an attachment for an Exascale volume.

Purpose

The `mkvolumeattachment` command allows you to create an attachment for an Exascale volume.

Syntax

```
mkvolumeattachment volume-ID device-name
                   --attributes { giClusterId=cluster-ID |
                                  initiator=edv-initiator-ID }
```

Command Options

The options for the `mkvolumeattachment` command are:

Copyright © 2022, 2026, Oracle and/or its affiliates.

- *volume-ID*: Specifies the identifier for the Exascale volume for which you want to create an attachment. You can use the lsvolume command to find the identifier for each volume.
- *device-name*: For an Exascale Direct Volume (EDV) attachment, specifies the device name to use in conjunction with the attachment. This is a user-supplied name, which is applied to the device file that is associated with the attachment. After attachment, the corresponding device file is located under `/dev/exc/`.
- `--attributes`: Specifies different attributes depending on the attachment type.
  - For a cluster-wide EDV attachment only, the `giClusterId` attribute specifies the Oracle Grid Infrastructure (GI) cluster ID associated with the volume attachment. When you specify this attribute, the corresponding device file is created on all nodes in the Oracle GI cluster.
  - For a node-specific EDV attachment only, the `initiator` attribute specifies the EDV initiator ID associated with the volume attachment. When you specify this attribute, the corresponding device file is only created on the cluster node associated with the specified EDV initiator ID.

Usage Notes

The following notes apply to EDV attachments:

- During initial system deployment with Oracle Exadata Deployment Assistant (OEDA), the Exascale Direct Volume (EDV) service is configured on each Exadata compute node (bare-metal or VM) and runs with the permissions of the Exascale user that manages the Oracle Grid Infrastructure (GI) cluster. To create an EDV attachment, you must use the Exascale user linked with the EDV service.

  If the GI cluster uses a non-role-separated user configuration with one Oracle OS user account, then the associated Exascale user is related to the EDV service. If the GI cluster uses a role-separated configuration with a Grid OS user account and an Oracle OS user account, then the EDV service is linked to the Exascale user associated with the Grid OS account.

  To find the Exascale user linked with the EDV service, use the ESCLI lsinitiator command with the `--detail` option and examine the `user` attribute.
- Each EDV attachment also has a kernel device file at `/dev/exc-devN`, where *N* is the minor number of the device. The kernel device name is contained as an attribute of the EDV attachment and is visible using the ESCLI lsvolumeattachment command.

  Note that most Linux tools, such as `iostat`, display the kernel device file at `/dev/exc-devN`, while Exascale commands use the user-named device file (under `/dev/exc/`).

  The relationship between the kernel device file and the user-named device file is also recorded in the udev database and is visible using the following Linux command:

  ```
  # udevadm info device-file
  ```

  In the `udevadm` command, for the *device-file* value, you can specify either the kernel device file (`/dev/exc-devN`) or the user-named device file (under `/dev/exc/`).
- By default, read and write access to EDV device files is only available to the `root` operating system user and members of the `disk` group. Depending on your use case, you may need to modify the permissions on the EDV device files before using them.

Copyright © 2022, 2026, Oracle and/or its affiliates.

For example, to make the EDV device file at `/dev/exc/my-vol` readable and writable by the `oracle` user and `dba` group, you could configure it using a udev rule similar to the following:

```
# cat /etc/udev/rules.d/57-edv-user.rules
KERNEL=="exc-*", ENV{EXC_ALIAS}=="my-vol", OWNER="oracle", GROUP="dba",
MODE="0660"
```

- To facilitate the management of udev rules related to EDV devices, each EDV client node is configured with a template udev rules file at `/etc/udev/rules.d/57-edv-user.rules`, which you can modify to fulfill your requirements. To maintain existing udev rules, `/etc/udev/rules.d/57-edv-user.rules` is preserved whenever the EDV client software is updated.
- Each EDV client node can support a maximum of 3000 attachments at the same time. This limit includes the total of all cluster attachments involving the server, as well as local attachments specific to the server.

Examples

Example 6-189 Create a Cluster-Wide EDV Volume Attachment

You can create a cluster-wide EDV volume attachment by specifying a *volume-ID*, *device-name*, and *cluster-ID*.

In the following example, the *volume-ID* is `vol0007_50e52177583f4be4bad68ac20b65001e`, the *device-name* is `my-vol`, and the *cluster-ID* is `72071863FA3E7FCAFF9F42A96957F4C5`.

```
@> mkvolumeattachment vol0007_50e52177583f4be4bad68ac20b65001e my-vol --
attributes giClusterId=72071863-fa3e-7fca-ff9f-42a96957f4c5
```

Example 6-190 Create a Node-Specific EDV Volume Attachment

You can create a node-specific EDV volume attachment by specifying a *volume-ID*, *device-name*, and *edv-initiator-ID*.

In the following example, the *volume-ID* is `vol0007_50e52177583f4be4bad68ac20b65001e`, the *device-name* is `my-vol`, and the *edv-initiator-ID* is `b0b057aa-1f2c-0f48-b0b0-57aa1f2c0f48`.

```
@> mkvolumeattachment vol0007_50e52177583f4be4bad68ac20b65001e my-vol --
attributes initiator=b0b057aa-1f2c-0f48-b0b0-57aa1f2c0f48
```

6.2.12.18 mkvolumebackup

Create a backup of a volume snapshot.

Purpose

The `mkvolumebackup` command creates a backup of an Exascale volume snapshot.

Syntax

```
mkvolumebackup --attributes volumeSnapshot=volume-snapshot-
ID,destination=backup-destination[,attribute=value]...
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Command Options

The command options for the `mkvolumebackup` command are:

- `--attributes`: Specifies attribute settings for the volume backup:
  - `volumeSnapshot`: Identifies the volume snapshot to back up. This attribute must be specified. You can use the lsvolumesnapshot command to find the identifier for each volume snapshot.
  - `destination`: Specifies the backup destination as a location in an Exascale vault.

    If the destination path ends with a forward slash (`/`), a file name is automatically created using the following pattern: `vol_<volume-ID>_bkp_<backup-ID>`.
  - You can also specify optional attributes, including a user-defined name that you can use to easily identify the backup and Exascale file storage attributes that you want to associate with the backup file.

    Use the `describe mkvolumebackup` command to view details about all the volume attributes you can set with `mkvolumebackup`.

Examples

Example 6-191 Create a Volume Snapshot Backup

This example shows creating a backup for volume snapshot ID `vol0002_snap0001_50e52177583f4be4bad68ac20b65001e` with the backup being stored at `@my-vault1/bkup/`.

```
@> mkvolumebackup --attributes
volumeSnapshot=vol0002_snap0001_3c0e9866d66345afbb316472f0bce825,destination=@
my-vault1/bkup/
```

Example 6-192 Create a Volume Snapshot Backup with Optional Attributes

This example shows creating another backup for volume snapshot ID `vol0002_snap0001_50e52177583f4be4bad68ac20b65001e`. In this case, the backup destination is `@my-vault1/bkup2/`. Additionally, the backup file content type is explicitly set to `RECO` and the backup name is set to `my-backup2`.

```
@> mkvolumebackup --attributes
volumeSnapshot=vol0002_snap0001_3c0e9866d66345afbb316472f0bce825,destination=@
my-vault1/bkup2/,contentType=RECO,name=my-backup2
```

6.2.12.19 mkvolumegroup

Create a volume group.

Purpose

The `mkvolumegroup` command creates a volume group.

Syntax

```
mkvolumegroup [ volume-group-name ]
              [ --attributes attribute=value[,attribute=value]... ]
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Command Options

The command options for the `mkvolumegroup` command are:

- *volume-group-name*: Specifies the name of the volume group.

  If the name is not specified here, it must be specified as an attribute setting.

- `--attributes`: Specifies attribute settings for the volume group:
  - `name`: Specifies the name of the volume group.

    If the name is not specified as an attribute setting, it must be specified as the first argument of the `mkvolumegroup` command.
  - `resourceSharing`: Optionally specifies how provisioned resources are shared by volumes in the group. The default value is `none`. Permitted values are:
    * `aggregate`: Configures the volume group to share the I/Os per second (IOPS) bandwidth provisioned to each volume in the group. In this case, a volume can exceed its provisioned limit by sharing unused IOPS provisioned to other volumes in the group.

      For example, consider a group with two volumes that are provisioned with 1000 IOPS each. In this example, aggregate resource sharing enables one volume to use 2000 IOPS when the other volume is idle.
    * `specified`: Configures the volume group to share the I/Os per second (IOPS) bandwidth limit specified by the `iopsProvisioned` attribute.
    * `none`: Specifies that the volume group is not a resource-sharing volume group. In this case, each volume is only subject to resource limits imposed at other levels, such as the individual volume level or the Exascale vault level.
  - `iopsProvisioned`: Optionally specifies the I/O bandwidth limit provisioned to the volume group, expressed in I/Os per second (IOPS). The default value is `0`.

Usage Notes

- When resource sharing is disabled on a volume group (`resourceSharing=none`), the I/O bandwidth provisioned to the volume group must be zero (`iopsProvisioned=0`).
- When a volume group uses aggregate resource sharing (`resourceSharing=aggregate`), the I/O bandwidth provisioned to the volume group must be zero (`iopsProvisioned=0`).
- When a volume group employs specified resource sharing (`resourceSharing=specified`), the `iopsProvisioned` attribute value must be greater than zero.

Examples

Example 6-193 Create a Volume Group

These examples show how to create volume groups with the specified names and other attribute settings.

```
@> mkvolumegroup vol-group1 --attributes
resourceSharing=specified,iopsProvisioned=1000


@> mkvolumegroup --attributes name=vol-group2,resourceSharing=aggregate
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

6.2.12.20 mkvolumegroupsnapshot

Create snapshots of all volumes in the specified volume group.

Purpose

The `mkvolumegroupsnapshot` command creates a set of volume snapshots with point-in-time consistency for all volumes in the specified volume group.

Syntax

```
mkvolumegroupsnapshot volume-group-ID
```

Command Options

The options for the `mkvolumegroupsnapshot` command are:

- *volume-group-ID*: Identifies the volume group that you want to act on. You can use the lsvolumegroup command to find the identifier for each volume group.

Usage Notes

The volume snapshot operation fails if the specified volume supports Exascale Direct Volume (EDV) attachments and Oracle Advanced Cluster File System (ACFS) is currently mounted. To create a volume snapshot in this situation, you must either:

- Unmount the affected file system on all servers before you run the command to create the volume snapshot. Then, mount the file system again afterward.
- Alternatively, you can use the Oracle ACFS command-line utility (`acfsutil`) to synchronize the file system and create the volume snapshot. On an EDV client node, run the `acfsutil volsnap create` command as the `root` user and specify the ACFS mountpoint or EDV device.

  For example:

  ```
  # acfsutil volsnap create /mnt/my-acfs
  ```

Examples

Example 6-194 Snapshot Volumes in a Volume Group

This example shows creating a snapshot of volumes in the volume group with the ID `volgrp0001_8b599d1725244432a1f47407b4e9493e`.

```
@> mkvolumegroupsnapshot volgrp0001_8b599d1725244432a1f47407b4e9493e
```

6.2.12.21 mkvolumesnapshot

Create a snapshot of a volume.

Purpose

The `mkvolumesnapshot` command creates a snapshot of the specified Exascale volume.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Syntax

```
mkvolumesnapshot volume-ID
                 [ --attributes attribute=value[,attribute=value]... ]
```

Command Options

The options for the `mkvolumesnapshot` command are:

- *volume-ID*: Identifies the volume that you want to snapshot. You can use the lsvolume command to find the identifier for each volume.
- `--attributes`: Optionally specifies attributes for the volume snapshot.

  Use the `describe mkvolumesnapshot` command to see details about the volume snapshot attributes you can set. See also Describing Resources and Attributes.

Usage Notes

The volume snapshot operation fails if the specified volume supports Exascale Direct Volume (EDV) attachments and Oracle Advanced Cluster File System (ACFS) is currently mounted. To create a volume snapshot in this situation, you must either:

- Unmount the affected file system on all servers before you run the command to create the volume snapshot. Then, mount the file system again afterward.
- Alternatively, you can use the Oracle ACFS command-line utility (`acfsutil`) to synchronize the file system and create the volume snapshot. On an EDV client node, run the `acfsutil volsnap create` command as the `root` user and specify the ACFS mountpoint or EDV device.

  For example:

  ```
  # acfsutil volsnap create /mnt/my-acfs
  ```

Examples

Example 6-195 Create a Volume Snapshot

This example shows creating a snapshot of volume `vol0001_50e52177583f4be4bad68ac20b65001e`.

```
@> mkvolumesnapshot vol0001_50e52177583f4be4bad68ac20b65001e
```

Example 6-196 Create a Volume Snapshot with Attribute Settings

This example shows creating a snapshot of volume `vol0001_50e52177583f4be4bad68ac20b65001e` with specific attribute settings.

```
@> mkvolumesnapshot vol0001_50e52177583f4be4bad68ac20b65001e --attributes
name=snap1
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

6.2.12.22 rmacfsfilesystem

Delete an Exascale-managed ACFS file system.

Purpose

The `rmacfsfilesystem` command deletes an Exascale-managed implementation of Oracle Advanced Cluster File System (ACFS).

Syntax

```
rmacfsfilesystem acfs-ID
```

Command Options

The options for the `rmacfsfilesystem` command are:

- *acfs-ID*: Identifies the ACFS file system being deleted. You can use the lsacfsfilesystem command to find the identifier for each Exascale-managed ACFS file system.

Usage Notes

Note the following information when using this command:

- The ACFS file system is dismounted and the mount point directory is removed on every node in the Oracle Grid Infrastructure (GI) cluster. The command also removes the ACFS file system registration in the GI cluster. Furthermore, the file system is removed from the underlying volume, effectively deleting the files in the file system. The command does not alter the underlying volume or volume attachment in any other way.

6.2.12.23 rmvolume

Delete volumes.

Purpose

The `rmvolume` command deletes Exascale volumes.

Syntax

```
rmvolume volume-ID [ volume-ID ]... [ --keep-data ]
```

Command Options

The options for the `rmvolume` command are:

- *volume-ID*: Identifies the volume being deleted. You can use the lsvolume command to find the identifier for each volume.
- `--keep-data`: Optionally instructs the command to delete the metadata associated with the specified volume(s) while preserving the underlying volume data file(s) for potential future use.

Usage Notes

Note the following information when using this command:

Copyright © 2022, 2026, Oracle and/or its affiliates.

- You must detach the volume before it can be deleted.

Examples

Example 6-197 Delete a Volume

This example shows deleting the volume with the volume ID `vol0001_50e52177583f4be4bad68ac20b65001e`.

```
@> rmvolume vol0001_50e52177583f4be4bad68ac20b65001e
```

6.2.12.24 rmvolumeattachment

Delete a volume attachment.

Purpose

The `rmvolumeattachment` command allows you to delete an Exascale volume attachment.

Syntax

```
rmvolumeattachment attachment-ID [ --force ]
```

Command Options

The options for the `rmvolumeattachment` command are:

- *attachment-ID*: Identifies the volume attachment that is the subject of the operation. You can use the lsvolumeattachment command to find the identifier for each volume attachment.
- `--force`: Forcibly removes the volume attachment even if it is currently in use.

  Use this option only when the client is unavailable or is expected to become unavailable soon after the command (such as immediately before rebooting a client). After a forced detachment, the volume stays open on the initiator and any associated file system remains mounted, but all I/O operations will fail. Also, you cannot reattach the volume until it is closed, usually by unmounting the file system.

Examples

Example 6-198 Delete an EDV Attachment

You can delete an Exascale Direct Volume (EDV) attachment by specifying just the *attachment-ID*. In the following example, the *attachment-ID* is `att0003_50e52177583f4be4bad68ac20b65001e`.

```
@> rmvolumeattachment att0003_50e52177583f4be4bad68ac20b65001e
```

6.2.12.25 rmvolumebackup

Delete a volume backup.

Purpose

The `rmvolumebackup` command deletes a backup of an Exascale volume.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Syntax

```
rmvolumebackup volume-backup-ID [ --keep-data ]
```

Command Options

The options for the `rmvolumebackup` command are:

- *volume-backup-ID*: Identifies the volume backup that is being deleted. You can use the lsvolumebackup command to find the identifier for each volume backup.
- `--keep-data`: Optionally instructs the command to delete the metadata associated with the specified volume backups(s) while preserving the underlying volume data file(s) for potential future use.

Examples

Example 6-199 Delete a Volume Backup

This example shows deleting the volume backup having the ID `vol0002_bkp0001_50e52177583f4be4bad68ac20b65001e`.

```
@> rmvolumebackup vol0002_bkp0001_50e52177583f4be4bad68ac20b65001e
```

6.2.12.26 rmvolumegroup

Delete volume groups.

Purpose

The `rmvolumegroup` command deletes the specified Exascale volume groups.

Syntax

```
rmvolumegroup volume-group-ID [ volume-group-ID ]...
```

Command Options

The options for the `rmvolumegroup` command are:

- *volume-group-ID*: Identifies the volume group being deleted. You can use the lsvolumegroup command to find the identifier for each volume group.

Usage Notes

To delete a volume group, you must first remove all volumes from it. See Adding and Removing Volumes in a Volume Group.

Examples

Example 6-200 Delete a Volume Group

This example shows deleting the volume group having the ID `volgrp0001_c2a3cec3a7e7413181fa1730093b3110`.

```
@> rmvolumegroup volgrp0001_c2a3cec3a7e7413181fa1730093b3110
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

6.2.12.27 rmvolumesnapshot

Delete volume snapshots.

Purpose

The `rmvolumesnapshot` command allows you to delete one or more Exascale volume snapshots.

Syntax

```
rmvolumesnapshot volume-snap-ID [ volume-snap-ID ]... [ --keep-data ]
```

Command Options

The options for the `rmvolumesnapshot` command are:

- *volume-snap-ID*: Identifies the volume snapshot being deleted. You can use the lsvolumesnapshot command to find the identifier for each volume snapshot.
- `--keep-data`: Optionally instructs the command to delete the metadata associated with the specified volume snapshot(s) while preserving the underlying volume data file(s) for potential future use.

Usage Notes

You must detach the volume snapshot before it can be deleted.

Examples

Example 6-201 Delete a Volume Snapshot

This example shows deleting the volume snapshot having the ID `vol0002_snap0001_50e52177583f4be4bad68ac20b65001e`.

```
@> rmvolumesnapshot vol0002_snap0001_50e52177583f4be4bad68ac20b65001e
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

7

Using XSH

This chapter describes how to use the Exascale shell command line interface (XSH):

- Start and Use XSH
- XSH Command Reference

7.1 Start and Use XSH

This topic describes how to start and use the Exascale shell command line interface (XSH).

XSH is a command-line tool that is located on Exadata compute nodes and storage servers. You can use XSH to perform Linux-like commands on Exascale files and storage.

To use XSH you must have access to an Exascale user account and the digital key store (wallet) for the Exascale user account.

If you do not specify a wallet in the XSH command, XSH uses the first Exascale user wallet available in the following search path:

1. `$OSSCONF/eswallet`
2. `$ORACLE_BASE/admin/eswallet`
3. `/etc/oracle/cell/network-config/eswallet`

To run an XSH command, use the following command line syntax:

```
$ xsh [ XSH-command ]
```

In the command line, *XSH-command* specifies an XSH command to run. For example:

```
$ xsh ls
@my-data
$ xsh ls @my-data
@my-data/x  @my-data/y  @my-data/z1  @my-data/z2  @my-data/z3
$
```

If no command is specified, XSH displays a command list and basic help information. For example:

```
$ xsh
cat
clone
cp
dd
ls
man
hexdump
rm
scrub
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
snap
strings
template
touch
xattr
version
chacl
mv
lsacl
lsdataset
chtemplate
mkresourceprofile
lsresourceprofile
chresourceprofile
rmresourceprofile

Enter 'xsh man <command>' or 'xsh man -e <command>' for details
E.g., Enter 'xsh man dd' or 'xsh man -e dd'
Enter 'xsh man man' to see other options accepted by 'xsh man'
```

For specific online command help, use the XSH `man` command. For example:

```
$ xsh man cat
NAME
       cat - Dump the contents of the file to
             standard out

SYNOPSIS
       cat      [{ -w | --wallet } <wallet-location> ]
                [{ -T | --trace } [ <trace-level> ]]
                [--aio=<naio>] filename

DESCRIPTION
       Dump the contents of the file to standard out.

OPTIONS
       -w, --wallet
              Path to the wallet directory.

       -T, --trace
              Trace level, default level = 1.

       --aio=<naio>
              Number of async IOs, default=4.

EXAMPLES
cat --aio=8 @DATA/example.dbf
cat /tmp/example.txt -w $OSSCONF/eswallet
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

To facilitate troubleshooting, you can enable tracing by setting the `$XSH_TRACE_LEVEL` environment variable. You can set the trace level to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). For example:

```
# sh/bash/ksh
$ export $XSH_TRACE_LEVEL=3
```

A trace file is written for each XSH command issued while tracing is enabled using the `$XSH_TRACE_LEVEL` environment variable. The trace file is written to the first accessible location in the following list:

1. If the `$ADR_BASE` environment variable is set:

   `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

To stop and deactivate tracing, you can set the `$XSH_TRACE_LEVEL` environment variable to `0` (zero). For example:

```
# sh/bash/ksh
$ export $XSH_TRACE_LEVEL=0
```

Apart from using the `$XSH_TRACE_LEVEL` environment variable, you can enable tracing for a single command by including the `--trace` or `-T` option in the XSH command. For example:

```
$ xsh ls --trace 3 @my-data
```

7.2 XSH Command Reference

This section contains references for the Exascale shell command line interface (XSH) commands:

- cat
  Dump the contents of a file to standard output.
- chacl
  Modify the access control list (ACL) for an Exascale file.
- chresourceprofile
  Change a resource profile associated with a vault.
- chtemplate
  Change template attributes.
- chxattr
  Set an extended attribute for a file or vault.
- clone
  Clone Exascale files.
- cp
  Copy files residing in Exascale or a local file system.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- dd
  Convert and copy a file residing in Exascale or a local file system.
- hexdump
  Dump the contents of a file to standard output in hex format.
- ls
  List file/directory details.
- lsacl
  Display the access control list (ACL) for Exascale files and vaults.
- lsblk
  List volume details.
- lsdataset
  List Exascale datasets.
- lsinitiator
  List Exascale Direct Volume (EDV) initiator details.
- lsresourceprofile
  List resource profiles.
- lstemplate
  Display template attributes.
- lsxattr
  Display extended attributes for a file or vault.
- man
  Display online manual pages.
- mkresourceprofile
  Create a resource profile.
- mktemplate
  Create an Exascale template.
- mv
  Move Exascale files.
- rm
  Remove Exascale files.
- rmresourceprofile
  Delete a resource profile.
- rmtemplate
  Delete an Exascale template.
- rmxattr
  Delete an extended attribute for a file or vault.
- scrub
  On-demand Exascale file scrubbing for logical consistency.
- snap
  Snapshot Exascale files.
- strings
  Display strings of printable characters in an Exascale file.
- touch
  Create a file or modify attributes of an existing file.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- version
  Print XSH version information to standard output.
- volumeiops
  Create and manage volume groups and their allocated I/O bandwidth (IOPS).

7.2.1 cat

Dump the contents of a file to standard output.

Syntax

```
cat [ --aio=naio ] file-name
    [{ -w | --wallet } wallet-location ]
    [{ -T | --trace } trace-level ]
```

Command Options

The options for the `cat` command are:

- *file-name*: Specifies the name of the file being output.
- `--aio`: Specifies the number of asynchronous I/O channels used to process the command. Asynchronous I/O enables the command to initiate I/O operations and continue executing other tasks without waiting (blocking) for the I/Os to complete. The default value is 4, which generally provides an optimal balance of performance and resource utilization.
- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

Examples

Example 7-1 Dump the contents of an Exascale file

The following example dumps the contents of the file at `@my-data/my-file` to standard output.

```
$ xsh cat @my-data/my-file
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

7.2.2 chacl

Modify the access control list (ACL) for an Exascale file.

Syntax

```
chacl file-name acl-string
      [{ -w | --wallet } wallet-location ]
      [{ -T | --trace } trace-level ]
      [{ -j | --json } [ --compact ]]
```

Command Options

The options for the `chacl` command are:

- *`file-name`*: Specifies the name of the file that is the subject of the operation.
- *acl-string*: Specifies an ACL string having the following format.

  `[+]`*`user-ID1`*`:`*`acl-permission`*`[+][;`*`user-ID2`*`:`*`acl-permission`*`[+]]...`

  In the ACL string:

  – An optional plus (+) operator at the beginning of the ACL string indicates that the specified ACL string merges into the existing ACL. In this case, the changes only impact the users in the specified ACL string, and all other users in the existing ACL retain their permissions. Without this optional plus (+) operator, the ACL is completely overwritten by the specified ACL string, enabling permissions only for the users in the specified ACL string.

  – *user-IDn*: Specifies an Exascale user ID.

    Depending on the user creation method, the user ID may be a system-generated value (for example, `96a68014-5762-4579-86ee-29eb743decbd`) or a user-specified value (for example, `scott`).

  – *acl-permission*: Specifies an ACL permission setting, which can be one of the following:

    * `I | inspect`: Specifies that the user can view attributes of the file, but not its contents.
    * `R | read`: Specifies that the user can read contents of the file. Also confers the `inspect` permission.
    * `U | use`: Specifies that the user can write to the file. Also confers all preceding permissions.
    * `M | manage`: Specifies that the user can manage the file. Also confers all preceding permissions.
    * `0 | none`: Specifies that the user is removed from the ACL and loses all permissions. This setting can be used only in conjunction with the initial plus (+) operator to remove a user from an existing ACL.

    Starting with Oracle Exadata System Software release 26.1.0, an optional plus (+) operator following an *acl-permission* value guarantees that the user has at least that privilege level. For example, `R+` ensures the user has at least the `R` (`read`) privilege. If the user already has a higher privilege, such as `M` (`manage`), their privileges remain

Copyright © 2022, 2026, Oracle and/or its affiliates.

unchanged. You can only use this option when updating an ACL and you cannot use this option with the `0` (`none`) privilege.

- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

- `-j`, `--json`: Displays the output in JSON format.
- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

Examples

Example 7-2 Replace a File ACL

In this example, the ACL string for the file `@my-data/file1` is replaced with the new ACL string. Under the new ACL, `scott` is permitted to read and inspect the file. No other user can access this file unless permitted by the vault ACL.

```
$ xsh chacl @my-data/file1 scott:R
```

Example 7-3 Change a File ACL

In this example, the plus sign (+) at the beginning of the ACL string indicates that the specified ACL string is merged into the existing file ACL. In this case, any pre-existing permissions for `jason` are overwritten, and `jason` is permitted to inspect, read, write, and manage the file. No other user permissions are changed.

```
$ xsh chacl @my-data/file1 +jason:M
```

Example 7-4 Remove User Privileges from a File ACL

In this example, the plus sign (+) at the beginning of the ACL string indicates that the specified ACL string is merged into the existing file ACL. However, in this case, any pre-existing permissions for `scott` are removed. No other user permissions are changed.

```
$ xsh chacl @my-data/file1 +scott:none
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Example 7-5 Replace a File ACL using an ACL String that Specifies Multiple Users

In this example, the ACL string for the file is replaced with the new ACL string that specifies permissions for multiple users. Under the new ACL, `scott` can inspect the file, and `jason` can read and inspect the file. No other user can access this file unless permitted by the vault ACL.

```
$ xsh chacl @my-data/file1 scott:inspect;jason:read
```

Related Topics

- About Exascale Access Control Lists

7.2.3 chresourceprofile

Change a resource profile associated with a vault.

Syntax

```
chresourceprofile --vault vault-name --name resource-profile-name
                  [ --iops-share-ef iopsShareEF ]
                  [ --iops-limit-ef iopsLimitEF ]
                  [ --iops-share-hc iopsShareHC ]
                  [ --iops-limit-hc iopsLimitHC ]
                  [ --iops-share-xt iopsShareXT ]
                  [ --iops-limit-xt iopsLimitXT ]
                  [ --enable-flash-cache {true|false} ]
                  [ --flash-cache-min flashCacheMin ]
                  [ --flash-cache-size flashCacheSize ]
                  [ --enable-xrmem-cache {true|false} ]
                  [ --xrmem-cache-min xrmemCacheMin ]
                  [ --xrmem-cache-size xrmemCacheSize ]
                  [ --enable-flash-log {true|false} ]
                  [{ -w | --wallet } wallet-location ]
                  [{ -T | --trace } trace-level ]
                  [{ -j | --json } [ --compact ]]
```

Command Options

The options for the `chresourceprofile` command are:

- `--vault`: Identifies the vault that the resource profile is associated with.
- `--name`: Identifies the resource profile being changed.
- `--iops-share-ef`, `--iops-share-hc`, `--iops-share-xt`: Specifies the relative share of I/O bandwidth (IOPS) from extreme flash (EF), high capacity (HC), or extended (XT) storage that is available to each client associated with the resource profile. Each client's share is relative to all other client shares associated with the corresponding media type. A higher share value implies higher priority. The range of valid values is `1-100`, and the default value is 1.
- `--iops-limit-ef`, `--iops-limit-hc`, `--iops-limit-xt`: Specifies the upper limit of the I/O bandwidth (IOPS) from extreme flash (EF), high capacity (HC), or extended (XT) storage that is available to each client associated with the resource profile.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Starting with Oracle Exadata System Software release 26.1.0, each value is a percentage with up to two decimal places. The range of valid values is `0.01-100`, and the default value is `100` (effectively unlimited).

Before Oracle Exadata System Software release 26.1.0, each value represents a fraction out of 10000. The range of valid values is `1-10000`, and the default value is `10000` (effectively unlimited).

- `--enable-flash-cache`, `--enable-xrmem-cache`, `--enable-flash-log`: Enables or disables use of the flash cache, XRMEM cache, or flash log for clients associated with the resource profile. Each value is Boolean, and the default is `true` (enabled).
- `--flash-cache-min`, `--xrmem-cache-min`: Specifies the guaranteed minimum fraction of flash cache or XRMEM cache space available to each client associated with the resource profile.

  Starting with Oracle Exadata System Software release 26.1.0, each value is nominally a percentage with up to two decimal places. The range of valid values is `0-100`, and the default value is `0` (no set minimum).

  Before Oracle Exadata System Software release 26.1.0, each value represents a fraction out of 10000. The range of valid values is `0-10000`, and the default value is `0` (no set minimum).

  For each cache type, if the sum of all values across all resource profiles exceeds the maximum, then all the values are proportionally scaled down. Each value applies only when the corresponding cache is enabled in the resource profile (for example, `enableFlashCache=true`).
- `--flash-cache-size`, `--xrmem-cache-size`: Specifies the maximum fraction of flash cache or XRMEM cache space available to each client associated with the resource profile.

  Starting with Oracle Exadata System Software release 26.1.0, each value is nominally a percentage with up to two decimal places. The range of valid values is `0-100`, and the default value is `100`.

  Before Oracle Exadata System Software release 26.1.0, each value represents a fraction out of 10000. The range of valid values is `0-10000`, and the default value is `10000`.

  Each value applies only when the corresponding cache is enabled in the resource profile (for example, `enableFlashCache=true`).
- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`
- `-j`, `--json`: Displays the output in JSON format.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

Examples

Example 7-6 Change a Resource Profile

In this example, the `--iops-share-hc` option is set to `20` for the `silver` resource profile in conjunction with the `vault2` vault.

```
$ xsh chresourceprofile --vault vault2 --name silver --iops-share-hc 20
```

Related Topics

- About Exascale Resource Profiles

7.2.4 chtemplate

Change template attributes.

Syntax

```
chtemplate [ --vault vault ]
           { --file-type file-type | --name template-name }
           [ --media-type media-type | --storage-pool storage-pool-name]
           [ --redundancy redundancy ]
           [ --content-type content-type ]
           [{ -w | --wallet } wallet-location ]
           [{ -T | --trace } trace-level ]
           [{ -j | --json } [ --compact ]]
```

Command Options

The options for the `chtemplate` command are:

- `--vault`: Identifies the vault associated with the vault-level template being changed.
- `--file-type`: Alters the template associated with the specified file type.
- `--name`: Alters the user-defined template having the specified name.
- `--media-type`: Specifies the physical media type used to store files created with the template. Possible values are:
  - `HC`: Identifies high capacity storage, using hard disk drives (HDDs) on high-performance Exadata storage servers.
  - `EF`: Identifies extreme flash storage, using low-latency, high-throughput flash devices.
  - `XT`: Identifies extended storage, using HDDs on Exadata storage servers optimized for cost-effective online storage of less frequently accessed data.
- `--storage-pool`: Specifies the name of a storage pool used to store files created with the template.

  This option is available starting with Oracle Exadata System Software release 26.1.0, and is only permitted in conjunction with a vault template.
- `--redundancy`: Specifies the number of data copies that are maintained for files created with the template. Currently, the only permitted value is:

Copyright © 2022, 2026, Oracle and/or its affiliates.

  - `high`: Indicates three mirrored copies of the file data.
- `--content-type`: Specifies the type of content in files created with the template. Exascale internally uses this attribute to place file extents on physically separate devices in a manner that maximizes availability in the event of a failure. Possible values are:
  - `DATA`: Principally associated with user data.
  - `RECO`: Primarily for data used in backup and recovery operations.
- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`
- `-j`, `--json`: Displays the output in JSON format.
- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

Examples

Example 7-7 Modify an Exascale template

The following example sets the specified media type, redundancy, and content type settings in the user-defined template named `my-template`, which is associated with the `my-data` vault.

```
$ xsh chtemplate --vault my-data --name my-template --media-type EF --
redundancy high --content-type DATA
```

Related Topics

- About Exascale Templates

7.2.5 chxattr

Set an extended attribute for a file or vault.

Syntax

```
chxattr { file-name | vault-name }
        --name xname --value xval
        [{ -w | --wallet } wallet-location ]
        [{ -T | --trace } trace-level ]
        [{ -j | --json } [ --compact ]]
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Command Options

The options for the `chxattr` command are:

- `{ file-name | vault-name }`: Specifies the file or vault that is the subject of the command.
- `--name`: Specifies the name of the extended attribute that is the subject of the command.
- `--value`: Specifies the value of the extended attribute.
- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

- `-j`, `--json`: Displays the output in JSON format.
- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

Related Topics

- About Exascale Extended Attributes

7.2.6 clone

Clone Exascale files.

Syntax

```
clone source1 target1 [ sourceN targetN ]...
      [{ -w | --wallet } wallet-location ]
      [{ -T | --trace } trace-level ]
      [{ -j | --json } [ --compact ]]
```

Command Options

The options for the `clone` command are:

- *source1-N*: Specifies the source file(s) being cloned.
- *target1-N*: Specifies the location of the clone(s).
- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

- `-j`, `--json`: Displays the output in JSON format.
- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

Usage Notes

Note the following information when using this command:

- You can use a wildcard (%) in the *source* to specify multiple source files, in which case the corresponding *target* must also contain a matching wildcard.
- All files in a clone operation must be in the same vault.
- Multiple *source* and *target* pairs are permitted. In this case, the source file specifications are considered in order, and only the first match is used.
- All clones created in the same operation are point-in-time consistent.

Examples

Example 7-8 Clone a File

The following example shows cloning a file. The source file is `@my-data/file1`, and the clone is `@my-data/clone-file1`.

```
$ xsh clone @my-data/file1 @my-data/clone-file1
```

Example 7-9 Clone a Group of Files

The following example shows cloning a group of files using a wildcard. Assuming the existence of files named `@my-data/file1`, `@my-data/file2`, and so on, the example creates a clone of `@my-data/file1` named `@my-data/clone-dir/file1`, `@my-data/file2` named `@my-data/clone-dir/file2`, and so on.

```
$ xsh clone @my-data/file* @my-data/clone-dir/file*
```

Example 7-10 Clone Multiple File Groups

The following example shows cloning multiple groups of files. The example clones `@my-data/a*` and `@my-data/b*`, with the resulting clones located under `@my-data/clone/`.

```
$ xsh clone @my-data/a* @my-data/clone/a* @my-data/b* @my-data/clone/b*
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Example 7-11 Ordering Significance

These examples show the significance of the order in which the *source* and *target* pairs are specified.

In the first command, the clones for files matching `@my-data/a*` are created in `@my-data/clone/a`, and the clones for the other files are created in `@my-data/clone/other`.

```
$ xsh clone @my-data/a* @my-data/clone/a/a* @my-data/* @my-data/clone/other/*
```

In the following command, all of the clones are created in `@my-data/clone/other` because all of the files match `@my-data/*`. In this case, the second *source* and *target* pair is never used.

```
$ xsh clone @my-data/* @my-data/clone/other/* @my-data/a* @my-data/clone/a/a*
```

Example 7-12 Invalid Commands

The following examples show invalid commands. The first and second commands are invalid because multiple vaults are referenced. The final command is invalid because the clone destination does not contain a wildcard to match the source specification.

```
$ xsh clone @my-data/a* @my-data-backup/a*
```

```
$ xsh clone @my-data/a* @my-data/clone/a* @vault2/a* @vault2/clone/a*
```

```
$ xsh clone @my-data/withwc* @my-data/clone/withoutwc
```

Related Topics

- About Exascale File Clones

7.2.7 cp

Copy files residing in Exascale or a local file system.

Syntax

```
cp [ -f | --force ] [ --ftype=file-type ] [ --bs=block-size ]
   [ --template=template-name ] [ --content-type=content-type ]
   [ --media-type=media-type ] [ --redundancy=redundancy ]
   [[ -s | --snap ] | [ -c | --clone ]]
   [ --io-size=io-size ] [ --aio=naio ]
   [ --status=status-level] [ --sparse ]
   source1 target1 [ sourceN targetN ]...
   [{ -w | --wallet } wallet-location ]
   [{ -T | --trace } trace-level ]
   [{ -j | --json } [ --compact ]]
```

Command Options

The options for the `cp` command are:

- *source1-N*: Specifies the source file being copied.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- *target1-N*: Specifies the location of the file copy. The value cannot be a directory.
- `-f`, `--force`: Forces the target to be overwritten if it exists. This is the default and only behavior.
- `--ftype`: Optionally specifies the Oracle Database file type. The valid values are:
  - `ctrl` - control file
  - `data` - data file
  - `olog` - on-line log file
  - `alog` - archived log file
  - `temp` - temporary sort file
  - `init` - initialization parameter file
  - `pswd` - password file
  - `flog` - flashback log file
  - `ctrk` - change tracking file

  This option only applies when copying into Exascale storage.
- `--bs`: Optionally specifies the block size (in bytes) to use for copying the file(s).
- `--template`: Optionally specifies the name of the Exascale template to use when copying new files into Exascale storage.

  This option only applies when copying new files into Exascale storage.
- `--content-type`: Optionally specifies the content type setting for file creation on Exascale storage.

  If not specified, the value is determined by the file type and the associated template. If specified, this setting overrides the template setting. This option only applies when copying new files into Exascale storage.

  Permitted values are:
  - `DATA`: Principally associated with user data.
  - `RECO`: Primarily for data used in backup and recovery operations.
- `--media-type`: Optionally specifies the media type setting for file creation on Exascale storage.

  If not specified, the value is determined by the file type and the associated template. If specified, this setting overrides the template setting. This option only applies when copying new files into Exascale storage.

  Permitted values are:
  - `HC`: Identifies high capacity storage, using hard disk drives (HDDs) on high-performance Exadata storage servers.
  - `EF`: Identifies extreme flash storage, using low-latency, high-throughput flash devices.
  - `XT`: Identifies extended storage, using HDDs on Exadata storage servers optimized for cost-effective online storage of less frequently accessed data.
- `--redundancy`: Optionally specifies the redundancy setting for file creation on Exascale storage.

Copyright © 2022, 2026, Oracle and/or its affiliates.

If not specified, the value is determined by the file type and the associated template. If specified, this setting overrides the template setting. This option only applies when copying new files into Exascale storage.

Permitted values are:

  - `high`: Indicates three mirrored copies of the file data.

- `-s`, `--snap`: Creates a read-only snapshot of the specified source file.

  This option only applies when copying files in Exascale storage.

- `-c`, `--clone`: Creates a writable thin clone of the specified source file.

  This option only applies when copying files in Exascale storage.

- `--io-size`: Optionally specifies the file copy I/O batch size. The default value is 1M.
- `--aio`: Specifies the number of asynchronous I/O channels used to process the command. Asynchronous I/O enables the command to initiate I/O operations and continue executing other tasks without waiting (blocking) for the I/Os to complete. The default value is 4, which generally provides an optimal balance of performance and resource utilization.
- `--status`: Specifies the level of status information output during a long-running operation.

  Permitted values are:

  - `progress`: Periodically shows progress statistics.
  - `none`: Suppresses progress statistics and only shows applicable error messages. This is the default setting.

- `--sparse`: Performs the operation in a sparse-aware manner by skipping read and write I/Os for empty source file blocks (blocks containing all zeros) and generating sparse output files.

  This option only applies when copying files in Exascale storage.

- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

- `-j`, `--json`: Displays the output in JSON format.
- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

Usage Notes

Note the following information when using this command:

Copyright © 2022, 2026, Oracle and/or its affiliates.

- If both the source and target locations reside on Exascale storage, you can use a wildcard (%) in the source location to specify multiple source files, in which case the corresponding target location must also contain a matching wildcard. For example, see Example 7-15.

Examples

Example 7-13 Copy a File to Exascale

The following example copies the local file at `/tmp/file1` to `@my-data/file1.copy` on Exascale storage.

```
$ xsh cp /tmp/file1 @my-data/file1.copy
```

Example 7-14 Copy a File from Exascale

The following example copies the Exascale file at `@my-data/file2` to `/tmp/file2.copy` on the local file system.

```
$ xsh cp @my-data/file2 /tmp/file2.copy
```

Example 7-15 Copy Multiple Files on Exascale Storage

The following example uses the wildcard character (%) to copy multiple files from `@my-data/location1/prefix1%` to `@my-data/location2/anotherlocation/prefix2%`. Using the example command, a file named `@my-data/location1/prefix1mydata` would be copied to `@my-data/location2/anotherlocation/prefix2mydata`.

```
$ xsh cp @my-data/location1/prefix1% @my-data/location2/anotherlocation/
prefix2%
```

7.2.8 dd

Convert and copy a file residing in Exascale or a local file system.

Syntax

```
dd [ --if=name ] [ --of=name ] [ --bs=bytes ] [ --count=num ]
   [ --ibs=bytes ] [ --obs=bytes ] [ --seek=num ] [ --skip=num ]
   [ --imirror=num ] [ --omirror=num ] [ --mirror=num ]
   [ --truncate ] [ --aio=naio ] [ --block ]
   [ --status=status-level] [ --sparse ]
   [{ -w | --wallet } wallet-location ]
   [{ -T | --trace } trace-level ]
   [{ -j | --json } [ --compact ]]
```

Command Options

The options for the `dd` command are:

- `--if`: Specifies the name of the input file to use instead of standard input.
- `--of`: Specifies the name of the output file to use instead of standard output.
- `--bs`: Specifies the number of bytes to read and write in each block. The default value is `512`.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `--count`: Specifies the number of input blocks.
- `--ibs`: Specifies the number of bytes to read in each block. The default value is `512`.
- `--obs`: Specifies the number of bytes to write in each block. The default value is `512`.
- `--seek`: Specifies the number of blocks to skip at the start of the output.
- `--skip`: Specifies the number of blocks to skip at the start of the input.
- `--imirror`: Specifies the input mirror number. The first mirror is `0`, the second mirror is `1`, and so on. The default value is `255`, which means 'read any mirror.'
- `--omirror`: Specifies the output mirror number. The first mirror is `0`, the second mirror is `1`, and so on. The default value is `255`, which means 'write all mirrors.'
- `--mirror`: Specifies the input and output mirror number. The first mirror is `0`, the second mirror is `1`, and so on. The default value is `255`, which means 'all mirrors.'
- `--truncate` Resizes the output file.
- `--aio`: Specifies the number of asynchronous I/O channels used to process the command. Asynchronous I/O enables the command to initiate I/O operations and continue executing other tasks without waiting (blocking) for the I/Os to complete. The default value is 4, which generally provides an optimal balance of performance and resource utilization.
- `--block`: Specifies that the I/O offsets should be divisible by the relevant I/O block size.
- `--status`: Specifies the level of status information output during a long-running operation.

  Permitted values are:

  - `progress`: Periodically shows progress statistics.
  - `none`: Suppresses progress statistics and only shows applicable error messages. This is the default setting.
- `--sparse`: Performs the operation in a sparse-aware manner by skipping read and write I/Os for empty source file blocks (blocks containing all zeros) and generating sparse output files.

  This option only applies when copying files in Exascale storage.
- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`
- `-j`, `--json`: Displays the output in JSON format.
- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Examples

Example 7-16 Copy a Specific Portion of an Exascale File

This example demonstrates how to copy a specific section of an Exascale file to a local file. This type of command is valuable for extracting or examining a portion of a large file without copying the entire file.

Using a 16 KB (16,384 bytes) block size, the command skips the first 160 KB (10 blocks) of the Exascale file at `@my-data/dbs/t_cf1.f`, and then copies the next 32 KB (2 blocks) to the local file at `/tmp/16K.file`.

```
$ xsh dd --if=@my-data/dbs/t_cf1.f --of=/tmp/16K.file --bs=16384 --skip=10 --
count=2
```

7.2.9 hexdump

Dump the contents of a file to standard output in hex format.

Syntax

```
hexdump [ -j bytes | --skip-bytes=bytes ]
        [ -N bytes | --read-bytes=bytes ]
        [ --aio=naio ] [ -x ]
        file-name
        [{ -w | --wallet } wallet-location ]
        [{ -T | --trace } trace-level ]
```

Command Options

The options for the `hexdump` command are:

- *file-name*: Specifies the name of the file being output.
- `-j`, `--skip-bytes`: Optionally specifies the number of input bytes to skip.
- `-N`, `--read-bytes`: Optionally specifies the number of bytes to read and output.
- `--aio`: Specifies the number of asynchronous I/O channels used to process the command. Asynchronous I/O enables the command to initiate I/O operations and continue executing other tasks without waiting (blocking) for the I/Os to complete. The default value is 4, which generally provides an optimal balance of performance and resource utilization.
- `-x`: Displays the file contents in hex notation. This is the default behavior anyway.
- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

Copyright © 2022, 2026, Oracle and/or its affiliates.

2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

Examples

Example 7-17 Dump the contents of an Exascale file

The following example dumps the contents of the file at `@my-data/my-file` to standard output in hex format.

```
$ xsh hexdump @my-data/my-file
```

7.2.10 ls

List file/directory details.

Syntax

```
ls [ -h ] [ -l ] [ -r ] [ -t ] [ -S ]
   [ -a ] [ -1 ] [ -s ] [ -d | --detail ]
   [ name ]
   [{ -w | --wallet } wallet-location ]
   [{ -T | --trace } trace-level ]
   [{ -j | --json } [ -p | --storage-pool] [ --compact ]]
```

Command Options

The options for the `ls` command are:

- *name*: Specifies the file or vault that is the subject of the command.

  If no *name* is specified, the command lists all accessible vaults. If the *name* is a specific file, the command lists details for the specified file.

  If the specified *name* is a vault name and the `-s` option is not specified, the command lists all of the files in the vault.

  A wildcard (%) is permitted when specifying either a vault name or file name.
- `-h`: Displays output values in a human readable format (for example, `10M`, `22K`, `300G`, and so on). This option is used in conjunction with the `-l` option.
- `-l`: Displays a longer listing including file size and creation date.
- `-r`: Reverses the sort order.
- `-t`: Sort by creation time instead of alphabetically.
- `-S`: Sort by file size (largest first).
- `-a`: Displays hidden/special entries.
- `-1`: List one entry on each line of output.
- `-s`: Displays vault statistics instead of the vault contents.

  This option applies only when the *name* value specifies a vault name. The option is ignored when the *name* value specifies a file name.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `-d`, `--detail`: Displays a detailed listing of attributes.
- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

- `-j`, `--json`: Displays the output in JSON format.
- `-p`, `--storage-pool`: Displays detailed storage pool usage statistics. Only applicable with JSON output and when listing vaults.
- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

Examples

Example 7-18 List all Exascale vaults

The following example lists all vaults visible to the Exascale user associated with the XSH command invocation.

```
$ xsh ls
```

Example 7-19 List the contents of an Exascale vault

The following example lists the files contained in the `@my-data` vault.

```
$ xsh ls @my-data
```

Example 7-20 List the specific details about contents of an Exascale vault

The following example lists the files contained in the `@my-data` vault. The command displays a detailed (long) listing that includes hidden files. The output is sorted according to the file creation time and displayed in reverse order (oldest first).

```
# xsh ls -alrt @my-data
       583421952             Mar 10 14:57:34 2025 @my-data/
EXA1-82B1CA80638E4F86BF2E4B3EABC7B6E7/data.ocr
        20971520             Mar 10 14:57:41 2025 @my-data/
EXA1-82B1CA80638E4F86BF2E4B3EABC7B6E7/vfile1
            2048             Mar 10 15:01:53 2025 @my-data/
EXA1-82B1CA80638E4F86BF2E4B3EABC7B6E7/DB1/PASSWORD/pwdDB1.42A5F9E0
       175292416             Mar 10 15:02:00 2025 @my-data/
EXA1-82B1CA80638E4F86BF2E4B3EABC7B6E7/DB1/CONTROLFILE/Current.OMF.253F31EB
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
      4294967808             Mar 10 15:02:02 2025 @my-data/
EXA1-82B1CA80638E4F86BF2E4B3EABC7B6E7/DB1/ONLINELOG/group_1.OMF.52326749
      4294967808             Mar 10 15:02:04 2025 @my-data/
EXA1-82B1CA80638E4F86BF2E4B3EABC7B6E7/DB1/ONLINELOG/group_2.OMF.3628C85F
     17179877376             Mar 10 15:02:06 2025 @my-data/
EXA1-82B1CA80638E4F86BF2E4B3EABC7B6E7/DB1/DATAFILE/SYSTEM.OMF.05DAA9B3
      7964991488             Mar 10 15:02:08 2025 @my-data/
EXA1-82B1CA80638E4F86BF2E4B3EABC7B6E7/DB1/DATAFILE/SYSAUX.OMF.3EFADCA5
...
```

7.2.11 lsacl

Display the access control list (ACL) for Exascale files and vaults.

Syntax

```
lsacl [ file-name | vault-name ]...
      [{ -w | --wallet } wallet-location ]
      [{ -T | --trace } trace-level ]
      [{ -j | --json } [ --compact ]]
```

Command Options

The options for the `lsacl` command are:

- `[ file-name | vault-name ]...`: Specifies the names of the files or vaults that are the subjects of the operation.

  A wildcard (%) is permitted when specifying a file name.

- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

- `-j`, `--json`: Displays the output in JSON format.
- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Examples

Example 7-21 Display a File ACL

The following command shows the ACL string for the file named `@vault2/my-file1`.

```
$ xsh lsacl @vault2/my-file1
```

Example 7-22 Display File ACLs Using a Wildcard

The following command shows the ACL string for all files in `@vault2` having names starting with `my-file`.

```
$ xsh lsacl @vault2/my-file%
```

Related Topics

- About Exascale Access Control Lists

7.2.12 lsblk

List volume details.

Syntax

```
lsblk [ -h ] [ -c cluster-ID ]
      [[ -a | --all ] [ -d | --detail ] | [ -1 | --long ]]
      [ -p | --probe ] [ -1 ]
      [ volume-name ]...
      [{ -w | --wallet } wallet-location ]
      [{ -T | --trace } trace-level ]
      [{ -j | --json } [ -n | --noopen] [ --compact ]]
```

Command Options

The options for the `lsblk` command are:

- *volume-name*: Specifies the volume that is the subject of the command.

  If no *volume-name* is specified, the command lists all accessible volumes.

  A wildcard (`%`) is permitted when specifying a volume name.
- `-h`: Displays output values in a human readable format (for example, `10M`, `22K`, `300G`, and so on).
- `-c`: Display volume information associated with the specified Exascale cluster ID.
- `-a`, `--all`: Displays hidden/special entries.
- `-d`, `--detail`: Displays a detailed listing of attributes.
- `-l`, `--long`: Displays a longer listing in tabular form.
- `-p`, `--probe`: Displays volume state output.
- `-1`: List one entry on each line of output.
- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

- `-j`, `--json`: Displays the output in JSON format.
- `-n`, `--noopen`: Returns major, minor, device path, and cluster ID without opening the underlying devices. Only applicable with JSON output and in conjunction with the `-c` option.
- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

Examples

Example 7-23 List all Exascale volumes

The following example lists all volumes visible to the Exascale user associated with the XSH command invocation.

```
$ xsh lsblk
```

Example 7-24 List details about an Exascale volume

The following example displays detailed information about the `vol1` volume in JSON format.

```
$ xsh lsblk -d -j vol1
```

7.2.13 lsdataset

List Exascale datasets.

Syntax

```
lsdataset [ -d | --detail ] [ -h ] [ --vault vault ]
          [ -f | --files ] [ -r | --recursive ] [ -p | --storage-pool ]
          [ dataset-ID ]
          [{ -w | --wallet } wallet-location ]
          [{ -T | --trace } trace-level ]
          [{ -j | --json } [ --compact ]]
```

Command Options

The options for the `lsdataset` command are:

Copyright © 2022, 2026, Oracle and/or its affiliates.

- *dataset-ID*: The identifier for a dataset that you want to display information about.

  Each system-defined dataset has a unique composite identifier, which contains unique identifiers for the associated entities. The dataset identifier has one of the following formats:

  - *@Vault-name*: Identifies the vault-level dataset for a specific named vault.
  - *@Vault-name*:*GI-cluster-ID*: Identifies the dataset for a specific Oracle Grid Infrastructure (GI) cluster, which consumes storage space in the specified vault.
  - *@Vault-name*:*GI-cluster-ID*:*CDB-ID*: Identifies the dataset for a specific Oracle multitenant container database (CDB) that belongs to the specified GI cluster.
  - *@Vault-name*:*GI-cluster-ID*:*CDB-ID*:*PDB-ID*: Identifies the dataset for an Oracle pluggable database (PDB) that is associated with the specified CDB, GI cluster, and vault.

  The percent symbol (%) can be used for wildcard searches. For example, *@Vault-name*:*GI-cluster-ID*:% displays information about all CDB datasets associated with the specified GI cluster and vault.

  If *dataset-ID* is not specified, the command displays information about all datasets.
- `--vault`: Displays information for datasets in the specified vault.
- `-d`, `--detail`: Displays a detailed listing of attributes.
- `-h`: Displays output values in a human readable format (for example, `10M`, `22K`, `300G`, and so on).
- `-f`, `--files`: Displays files in the specified dataset.
- `-r`, `--recursive`: When listing participating files, also list files that are part of descendant datasets.
- `-p`, `--storage-pool`: Displays additional detailed storage pool space usage statistics.
- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`
- `-j`, `--json`: Displays the output in JSON format.
- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Examples

Example 7-25 List all Exascale datasets

The following example lists all datasets visible to the Exascale user associated with the XSH command invocation.

```
$ xsh lsdataset
```

Example 7-26 List datasets in an Exascale vault

The following examples list the datasets contained in the `@my-data` vault.

```
$ xsh lsdataset --vault @my-data
```

```
$ xsh lsdataset @my-data:%
```

Related Topics

- About Exascale Datasets

7.2.14 lsinitiator

List Exascale Direct Volume (EDV) initiator details.

Syntax

```
lsinitiator [ -c cluster-ID ]
      [ -d | --detail ] [ -1 ]
      [{ -w | --wallet } wallet-location ]
      [{ -T | --trace } trace-level ]
      [{ -j | --json } [ -n | --noopen] [ --compact ]]
```

Command Options

The options for the `lsinitiator` command are:

- `-c`: Display initiator information associated with the specified Exascale cluster ID.

  If not specified, the command displays initiator information associated with the current node.
- `-d`, `--detail`: Displays a detailed listing of attributes.
- `-l`: Displays a longer listing in tabular form.
- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

Copyright © 2022, 2026, Oracle and/or its affiliates.

1. If the `$ADR_BASE` environment variable is set:

   `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

- `-j`, `--json`: Displays the output in JSON format.
- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

7.2.15 lsresourceprofile

List resource profiles.

Syntax

```
lsresourceprofile --vault vault-name [ --name resource-profile-name ]
                  [ -l ] [ --detail ]
                  [{ -w | --wallet } wallet-location ]
                  [{ -T | --trace } trace-level ]
                  [{ -j | --json } [ --compact ]]
```

Command Options

The options for the `lsresourceprofile` command are:

- `--vault`: Identifies a specific vault.
- `--name`: Identifies a specific resource profile name.
- `-l`: Returns output in a long, tabular form.
- `--detail`: Lists all attributes in a detailed form.
- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

- `-j`, `--json`: Displays the output in JSON format.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

Examples

Example 7-27 List Resource Profiles

This example lists the resource profiles associated with the vault named `vault2`.

```
$ xsh lsresourceprofile --vault vault2
```

Example 7-28 List a Specific Resource Profile

This example shows how to display detailed information about the resource profile named `silver` that is associated with the vault named `vault2`.

```
$ xsh lsresourceprofile --vault vault2 --name silver --detail
```

Related Topics

- About Exascale Resource Profiles

7.2.16 lstemplate

Display template attributes.

Syntax

```
lstemplate template-name
           [ --vault vault [ --storage-pool ] | --cluster ]
           [{ -w | --wallet } wallet-location ]
           [{ -T | --trace } trace-level ]
           [{ -j | --json } [ --compact ]]
```

Command Options

The options for the `lstemplate` command are:

- *template-name*: Identifies the template that is the subject of the command.
- `--vault`: Indicates that the specified template is a vault-level template and identifies the associated vault.
- `--storage-pool`: Optionally includes the storage pool name in the output.
- `--cluster`: Indicates that the specified template is a cluster-level template.
- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

Copyright © 2022, 2026, Oracle and/or its affiliates.

2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

- `-j`, `--json`: Displays the output in JSON format.
- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

Usage Notes

Note the following information when using this command:

- If neither `--vault` nor `--cluster` is specified, the command displays information about the effective template.

Examples

Example 7-29 Display Attributes of an Exascale Template

The following example displays attributes of the user-defined template named `my-template`, which is associated with the `my-data` vault.

```
$ xsh lstemplate --vault my-data --name my-template
```

Related Topics

- About Exascale Templates

7.2.17 lsxattr

Display extended attributes for a file or vault.

Syntax

```
lsxattr { file-name | vault-name } [ --name xname ]
        [ -v | --verbose ] [ -u | --unlimited ]
        [{ -w | --wallet } wallet-location ]
        [{ -T | --trace } trace-level ]
        [{ -j | --json } [ --compact ]]
```

Command Options

The options for the `lsxattr` command are:

- `{ file-name | vault-name }`: Specifies the file or vault that is the subject of the command.
- `--name`: Specifies the name of the extended attribute that is being displayed.

  If not specified, the command displays all extended attributes associated with the specified file or vault.
- `-v`, `--verbose`: Displays extended attributes with size information.
- `-u`, `--unlimited`: Displays complete extended attribute values. If not specified, the output of extended attribute values is limited to 80 characters.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

- `-j`, `--json`: Displays the output in JSON format.
- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

Related Topics

- About Exascale Extended Attributes

7.2.18 man

Display online manual pages.

Syntax

```
man [ -a ] [ -c ] [ -e ] [ -f ] [ -o ] [ -x ] [ entry ]
```

Command Options

The options for the `man` command are:

- *entry*: Specifies the manual entry that you want to display. If none is specified, the command lists all manual pages.
- `-a`: Displays all details, including examples.
- `-c`: Displays the command name.
- `-e`, `-x`: Displays example information.
- `-f`: Displays details except for examples. This is the default behavior if *entry* is specified without any other options.
- `-o`: Displays brief (one line) information.

Examples

Example 7-30 Display command details

The following example displays details about the `dd` command.

```
$ xsh man dd
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

7.2.19 mkresourceprofile

Create a resource profile.

Syntax

```
mkresourceprofile --vault vault-name --name resource-profile-name
                  [ --iops-share-ef iopsShareEF ]
                  [ --iops-limit-ef iopsLimitEF ]
                  [ --iops-share-hc iopsShareHC ]
                  [ --iops-limit-hc iopsLimitHC ]
                  [ --iops-share-xt iopsShareXT ]
                  [ --iops-limit-xt iopsLimitXT ]
                  [ --enable-flash-cache {true|false} ]
                  [ --flash-cache-min flashCacheMin ]
                  [ --flash-cache-size flashCacheSize ]
                  [ --enable-xrmem-cache {true|false} ]
                  [ --xrmem-cache-min xrmemCacheMin ]
                  [ --xrmem-cache-size xrmemCacheSize ]
                  [ --enable-flash-log {true|false} ]
                  [{ -w | --wallet } wallet-location ]
                  [{ -T | --trace } trace-level ]
                  [{ -j | --json } [ --compact ]]
```

Command Options

The options for the `mkresourceprofile` command are:

- `--vault`: Identifies the vault that the resource profile is associated with.
- `--name`: Specifies the name of the resource profile being created.
- `--iops-share-ef`, `--iops-share-hc`, `--iops-share-xt`: Specifies the relative share of I/O bandwidth (IOPS) from extreme flash (EF), high capacity (HC), or extended (XT) storage that is available to each client associated with the resource profile. Each client's share is relative to all other client shares associated with the corresponding media type. A higher share value implies higher priority. The range of valid values is `1-100`, and the default value is 1.
- `--iops-limit-ef`, `--iops-limit-hc`, `--iops-limit-xt`: Specifies the upper limit of the I/O bandwidth (IOPS) from extreme flash (EF), high capacity (HC), or extended (XT) storage that is available to each client associated with the resource profile.

  Starting with Oracle Exadata System Software release 26.1.0, each value is a percentage with up to two decimal places. The range of valid values is `0.01-100`, and the default value is `100` (effectively unlimited).

  Before Oracle Exadata System Software release 26.1.0, each value represents a fraction out of 10000. The range of valid values is `1-10000`, and the default value is `10000` (effectively unlimited).
- `--enable-flash-cache`, `--enable-xrmem-cache`, `--enable-flash-log`: Enables or disables use of the flash cache, XRMEM cache, or flash log for clients associated with the resource profile. Each value is Boolean, and the default is `true` (enabled).
- `--flash-cache-min`, `--xrmem-cache-min`: Specifies the guaranteed minimum fraction of flash cache or XRMEM cache space available to each client associated with the resource profile.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Starting with Oracle Exadata System Software release 26.1.0, each value is nominally a percentage with up to two decimal places. The range of valid values is `0-100`, and the default value is `0` (no set minimum).

Before Oracle Exadata System Software release 26.1.0, each value represents a fraction out of 10000. The range of valid values is `0-10000`, and the default value is `0` (no set minimum).

For each cache type, if the sum of all values across all resource profiles exceeds the maximum, then all the values are proportionally scaled down. Each value applies only when the corresponding cache is enabled in the resource profile (for example, `enableFlashCache=true`).

- `--flash-cache-size`, `--xrmem-cache-size`: Specifies the maximum fraction of flash cache or XRMEM cache space available to each client associated with the resource profile.

  Starting with Oracle Exadata System Software release 26.1.0, each value is nominally a percentage with up to two decimal places. The range of valid values is `0-100`, and the default value is `100`.

  Before Oracle Exadata System Software release 26.1.0, each value represents a fraction out of 10000. The range of valid values is `0-10000`, and the default value is `10000`.

  Each value applies only when the corresponding cache is enabled in the resource profile (for example, `enableFlashCache=true`).

- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

- `-j`, `--json`: Displays the output in JSON format.
- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

Examples

Example 7-31 Create a Resource Profile with Default Attribute Values

This example shows how to create a resource profile with default attribute values, which is named `silver` and is associated with the vault named `vault2`.

```
$ xsh mkresourceprofile --vault vault2 --name silver
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Example 7-32 Create a Resource Profile with Specific Attribute Values

This example shows how to create a resource profile that contains specific attribute settings.

```
$ xsh mkresourceprofile --vault vault2 --name gold --iops-share-hc 20 --
enable-flash-log true
```

Related Topics

- About Exascale Resource Profiles

7.2.20 mktemplate

Create an Exascale template.

Syntax

```
mktemplate { --file-type file-type | --name template-name }
           [ --vault vault | --cluster ]
           { --media-type media-type | --storage-pool storage-pool-name ]
           --redundancy redundancy
           --content-type content-type
           [{ -w | --wallet } wallet-location ]
           [{ -T | --trace } trace-level ]
           [{ -j | --json } [ --compact ]]
```

Command Options

The options for the `mktemplate` command are:

- `--file-type`: Creates a template associated with the specified file type.
- `--name`: Creates a user-defined template having the specified name.
- `--vault`: Identifies the vault associated with the template being created.
- `--cluster`: Specifies the creation of a cluster template.
- `--media-type`: Specifies the physical media type used to store files created with the template. Possible values are:
  - `HC`: Identifies high capacity storage, using hard disk drives (HDDs) on high-performance Exadata storage servers.
  - `EF`: Identifies extreme flash storage, using low-latency, high-throughput flash devices.
  - `XT`: Identifies extended storage, using HDDs on Exadata storage servers optimized for cost-effective online storage of less frequently accessed data.
- `--storage-pool`: Specifies the name of a storage pool used to store files created with the template.

  This option is available starting with Oracle Exadata System Software release 26.1.0, and is only permitted in conjunction with a vault template.
- `--redundancy`: Specifies the number of data copies that are maintained for files created with the template. Currently, the only permitted value is:
  - `high`: Indicates three mirrored copies of the file data.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `--content-type`: Specifies the type of content in files created with the template. Exascale internally uses this attribute to place file extents on physically separate devices in a manner that maximizes availability in the event of a failure. Possible values are:
  - `DATA`: Principally associated with user data.
  - `RECO`: Primarily for data used in backup and recovery operations.
- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`
- `-j`, `--json`: Displays the output in JSON format.
- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

Usage Notes

Note the following information when using this command:

- If neither `--vault` nor `--cluster` is specified, the command creates a cluster template.

Examples

Example 7-33 Create an Exascale template

The following example creates a user-defined template named `my-template`, which is associated with the `my-data` vault and contains the specified media type, redundancy, and content type settings.

```
$ xsh mktemplate --name my-template --vault my-data --media-type EF --
redundancy high --content-type DATA
```

Related Topics

- About Exascale Templates

7.2.21 mv

Move Exascale files.

Syntax

```
mv source1 target1 [ sourceN targetN ]...
   [{ -w | --wallet } wallet-location ]
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
[{ -T | --trace } trace-level ]
[{ -j | --json } [ --compact ]]
```

Command Options

The options for the `mv` command are:

- *source1-N*: Specifies the source file being moved.
- *target1-N*: Specifies the target location of the file. The value cannot be a directory.
- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

- `-j`, `--json`: Displays the output in JSON format.
- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

Usage Notes

Note the following information when using this command:

- The source and target locations must reside in the same Exascale vault.
- You can use a wildcard (`%`) in the source location to specify multiple source files, in which case the corresponding target location must also contain a matching wildcard. For example, see Example 7-34.

Examples

Example 7-34 Move Multiple Files on Exascale Storage

The following example uses the wildcard character (`%`) to move multiple files from `@my-data/location1/prefix1%` to `@my-data/location2/another-location/prefix2%`. Using the example command, a file named `@my-data/location1/prefix1mydata` would be moved to `@my-data/location2/another-location/prefix2mydata`.

```
$ xsh mv @my-data/location1/prefix1% @my-data/location2/another-location/
prefix2%
```

Related Topics

- About Exascale Files

Copyright © 2022, 2026, Oracle and/or its affiliates.

7.2.22 rm

Remove Exascale files.

Syntax

```
rm [ -f ] [ -i ] [ -v ] file-name
   [{ -w | --wallet } wallet-location ]
   [{ -T | --trace } trace-level ]
   [{ -j | --json } [ --compact ]]
```

Command Options

The options for the `rm` command are:

- *file-name*: Specifies the files being removed.

  You can use a wildcard (`%`) in the *file-name* to specify multiple files.
- `-f`: Forces removal of open files.
- `-i`: Interactively prompt for confirmation before deleting a file.
- `-v`: Print the names of the files being deleted.
- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`
- `-j`, `--json`: Displays the output in JSON format.
- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

Examples

Example 7-35 Delete an Exascale file

The following example deletes `@my-data/my-file`.

```
$ xsh rm @my-data/my-file
```

Related Topics

- About Exascale Files

Copyright © 2022, 2026, Oracle and/or its affiliates.

7.2.23 rmresourceprofile

Delete a resource profile.

Syntax

```
rmresourceprofile --vault vault-name --name resource-profile-name
                  [{ -w | --wallet } wallet-location ]
                  [{ -T | --trace } trace-level ]
                  [{ -j | --json } [ --compact ]]
```

Command Options

The options for the `rmresourceprofile` command are:

- `--vault`: Identifies a specific vault.
- `--name`: Identifies the name of the resource profile being deleted.
- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

- `-j`, `--json`: Displays the output in JSON format.
- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

Examples

Example 7-36 Delete a Resource Profile

This example shows how to delete a resource profile named `silver` that is associated with the vault named `vault2`.

```
$ xsh rmresourceprofile --vault vault2 --name silver
```

Related Topics

- About Exascale Resource Profiles

Copyright © 2022, 2026, Oracle and/or its affiliates.

7.2.24 rmtemplate

Delete an Exascale template.

Syntax

```
rmtemplate { --vault vault | --cluster }
           { --file-type file-type | --name template-name }
           [{ -w | --wallet } wallet-location ]
           [{ -T | --trace } trace-level ]
           [{ -j | --json } [ --compact ]]
```

Command Options

The options for the `rmtemplate` command are:

- `--vault`: Identifies the vault associated with the template being deleted.
- `--cluster`: Specifies the deletion of a cluster template.
- `--file-type`: Deletes the template associated with the specified file type.
- `--name`: Deletes the user-defined template having the specified name.
- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

- `-j`, `--json`: Displays the output in JSON format.
- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

Usage Notes

Note the following information when using this command:

- Deletion of a cluster template that is associated with a file type it is not permitted. Consequently, you cannot specify the `--cluster` and `--file-type` options in the same command.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Examples

Example 7-37 Delete an Exascale template

The following example deletes the user-defined template named `my-template`, which is associated with the `my-data` vault.

```
$ xsh rmtemplate --vault my-data --name my-template
```

Related Topics

- About Exascale Templates

7.2.25 rmxattr

Delete an extended attribute for a file or vault.

Syntax

```
rmxattr { file-name | vault-name }
        --name xname
        [{ -w | --wallet } wallet-location ]
        [{ -T | --trace } trace-level ]
        [{ -j | --json } [ --compact ]]
```

Command Options

The options for the `rmxattr` command are:

- `{ file-name | vault-name }`: Specifies the file or vault that is the subject of the command.
- `--name`: Specifies the name of the extended attribute that is being deleted.
- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

- `-j`, `--json`: Displays the output in JSON format.
- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Related Topics

- About Exascale Extended Attributes

7.2.26 scrub

On-demand Exascale file scrubbing for logical consistency.

Syntax

```
scrub [ -a ] [ -i interval] [ -n requests] [ -r count]
      [ -f offset] [ -t offset] [ -s ] [ -v[v[v]] ]
      file-name
      [{ -w | --wallet } wallet-location ]
      [{ -T | --trace } trace-level ]
```

Command Options

The options for the `scrub` command are:

- *file-name*: Specifies the files being scrubbed.

  You can use a wildcard (%) in the *file-name* to specify multiple files.
- `-a`: Specifies that the scrub operation performs HARD checks on all regions, regardless of the file type or mirror status.
- `-i`: Specifies the interval (in seconds) to wait before re-scrubbing a region.
- `-n`: Specifies the queue depth for the scrubbing operation. This option controls the number of concurrent read requests performed by the scrubbing operation. Larger values enable faster scrubbing but consume more resources, which may affect other workloads.
- `-r`: Specifies the number of re-scrubbing attempts permitted while awaiting concurrent writes on the region. After the specified number of re-scrubbing attempts is exhausted, the corresponding region is reported in the command output and the scrubbing operation continues.
- `-f`: Enables scrubbing on part of a file by specifying the file offset that scrubbing starts from. The offset value is specified in megabytes from the start of the file.
- `-t`: Enables scrubbing on part of a file by specifying the file offset that scrubbing operates to. The offset value is specified in megabytes from the start of the file.
- `-s`: Specifies that when considering multiple sets of data as the source for a repair, the scrub operation should use the highest system change number (SCN) as the repair source.
- `-v[v[v]`: Increases the verbosity of the command output:
  - `-v`: Generates detailed output.
  - `-vv`: Generates more detailed output.
  - `-vvv`: Generates the most detailed output.
- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

Copyright © 2022, 2026, Oracle and/or its affiliates.

1. If the `$ADR_BASE` environment variable is set:

   `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

Usage Notes

Logical file scrubbing is a valuable tool for addressing certain issues that may affect the integrity or consistency of Exascale files. However, regular or periodic scrubbing is not required or recommended.

Examples

Example 7-38 Scrub an Exascale file

The following example performs logical scrubbing on `@my-data/my-file`.

```
$ xsh scrub @my-data/my-file
```

Related Topics

- About Exascale Files

7.2.27 snap

Snapshot Exascale files.

Syntax

```
snap source1 target1 [ sourceN targetN ]...
     [{ -w | --wallet } wallet-location ]
     [{ -T | --trace } trace-level ]
     [{ -j | --json } [ --compact ]]
```

Command Options

The options for the `snap` command are:

- *source1-N*: Specifies the snapshot source file(s).
- *target1-N*: Specifies the location of the snapshot(s).
- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

Copyright © 2022, 2026, Oracle and/or its affiliates.

1. If the `$ADR_BASE` environment variable is set:

   `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

- `-j`, `--json`: Displays the output in JSON format.
- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

Usage Notes

Note the following information when using this command:

- You can use a wildcard (%) in the *source* to specify multiple source files, in which case the corresponding *target* must also contain a matching wildcard.
- All files in a snapshot operation must be in the same vault.
- Multiple *source* and *target* pairs are permitted. In this case, the source file specifications are considered in order, and only the first match is used.
- All snapshots created in the same operation are point-in-time consistent.

Examples

Example 7-39 Snapshot a File

This example shows creating a snapshot. The source file is `@my-data/file1`, and the snapshot is `@my-data/snap-file1`.

```
$ xsh snap @my-data/file1 @my-data/snap-file1
```

Example 7-40 Snapshot a Group of Files

This example shows creating snapshots for a group of files using a wildcard. Assuming the existence of files named `@my-data/file1`, `@my-data/file2`, and so on, the example creates a snapshot of `@my-data/file1` named `@my-data/snap-dir/file1`, `@my-data/file2` named `@my-data/snap-dir/file2`, and so on.

```
$ xsh snap @my-data/file* @my-data/snap-dir/file*
```

Example 7-41 Snapshot Multiple File Groups

This example shows creating snapshots for multiple groups of files using wildcards. The source file groups are `@my-data/a*` and `@my-data/b*`, with the resulting snapshots located under `@my-data/snap/`.

```
$ xsh snap @my-data/a* @my-data/snap/a* @my-data/b* @my-data/snap/b*
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Example 7-42 Ordering Significance

These examples show the significance of the order in which the *source* and *target* pairs are specified.

In the first command, the snapshots for files matching `@my-data/a*` are created in `@my-data/snap/a`, and the snapshots for the other files are created in `@my-data/snap/other`.

```
$ xsh snap @my-data/a* @my-data/snap/a/a* @my-data/* @my-data/snap/other/*
```

In the following command, all of the snapshots are created in `@my-data/snap/other` because all of the files match `@my-data/*`. In this case, the second *source* and *target* pair is never used.

```
$ xsh snap @my-data/* @my-data/snap/other/* @my-data/a* @my-data/snap/a/a*
```

Example 7-43 Invalid Commands

The following examples show invalid commands. The first and second commands are invalid because multiple vaults are referenced. The final command is invalid because the snapshot destination does not contain a wildcard to match the source specification.

```
$ xsh snap @my-data/a* @my-data-backup/a*
```

```
$ xsh snap @my-data/a* @my-data/snap/a* @vault2/a* @vault2/snap/a*
```

```
$ xsh snap @my-data/withwc* @my-data/snap/withoutwc
```

Related Topics

- About Exascale File Snapshots

7.2.28 strings

Display strings of printable characters in an Exascale file.

Syntax

```
strings [ -n num ] [ --aio=naio ]
        file-name
        [{ -w | --wallet } wallet-location ]
        [{ -T | --trace } trace-level ]
```

Command Options

The options for the `strings` command are:

- *file-name*: Specifies the file being processed.
- `-n`: Optionally specifies the minimum output string length. The default value is 4.
- `--aio`: Specifies the number of asynchronous I/O channels used to process the command. Asynchronous I/O enables the command to initiate I/O operations and continue executing other tasks without waiting (blocking) for the I/Os to complete. The default value is 4, which generally provides an optimal balance of performance and resource utilization.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

Examples

Example 7-44 Display printable strings in an Exascale file

The following example displays printable strings contained in `@my-data/my-file` to standard output.

```
$ xsh strings @my-data/my-file
```

Related Topics

- About Exascale Files

7.2.29 touch

Create a file or modify attributes of an existing file.

Syntax

```
touch [ -s num | --size=num ]
      [ -t file-type | --type=file-type ]
      [ --bs=block-size ]
      [ -r { true | false } | --read-only={ true | false }]
      [ -e expiry-time | --lockExpiryTime=expiry-time
      [ -i { true | false } | --lockImmutable={ true | false } ]]
      file-name
      [{ -w | --wallet } wallet-location ]
      [{ -T | --trace } trace-level ]
      [{ -j | --json } [ --compact ]]
```

Command Options

The options for the `touch` command are:

- *file-name*: Specifies the file being created or modified.
- `-s`, `--size`: Optionally specifies the size of the file in bytes.

  The file size can be also specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.

Copyright © 2022, 2026, Oracle and/or its affiliates.

When modifying an existing file, you can grow or shrink the file by specifying a new size.

> **⚠ Caution**
>
> If you specify a new size that is smaller than the current file size, the file is truncated to the new size and any data beyond that point is discarded.

- `-t`, `--ftype`: Optionally specifies the Oracle Database file type for a new file. The valid values are:
  - `ctrl` - control file
  - `data` - data file
  - `olog` - on-line log file
  - `alog` - archived log file
  - `temp` - temporary sort file
  - `init` - initialization parameter file
  - `pswd` - password file
  - `flog` - flashback log file
  - `ctrk` - change tracking file
- `--bs`: Optionally specifies the block size for a new file.
- `-r`, `--read-only`: Optionally specifies that the file is read-only. When set to `true`, the file contents and metadata cannot be modified.
- `-e`, `--lockExpiryTime`: Optionally specifies the snapshot lock expiry date and time.

  This option applies only to snapshot files.

  The value can be an ISO 8601-formatted future date and time. Or, to unlock a locked snapshot, you can set the expiry time to 'null' (`lockExpiryTime='null'`). You can modify this attribute only if the snapshot lock is not immutable (`lockImmutable=FALSE`).
- `-i`, `--lockImmutable`: Indicates whether the lock is immutable.

  When a lock is immutable, the snapshot cannot be unlocked before the lock expires, and the lock cannot be modified to expire sooner.

  This option applies only to snapshot files and only in conjunction with setting the lock expiry time. If not specified, the default value is `false`.
- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

Copyright © 2022, 2026, Oracle and/or its affiliates.

3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

- `-j`, `--json`: Displays the output in JSON format.
- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

Examples

Example 7-45 Create or resize an Exascale file

Assuming that `@my-data/my-file` does not already exist, the following example creates a 1 GB file at `@my-data/my-file`.

```
$ xsh touch --size=1G @my-data/my-file
```

If `@my-data/my-file` already exists, the command changes the size of the existing file to 1 GB.

Related Topics

- About Exascale Files

7.2.30 version

Print XSH version information to standard output.

Syntax

```
version
```

7.2.31 volumeiops

Create and manage volume groups and their allocated I/O bandwidth (IOPS).

Syntax

```
volumeiops create { -i | --id } group
                  { -o | --iops } total-IOPS
                  { -a | --add } volume [{ -a | --add } volume ]...
                  [{ -w | --wallet } wallet-location ]
                  [{ -T | --trace } trace-level ]
                  [{ -j | --json } [ --compact ]]


volumeiops update { -i | --id } group
                  [{ -o | --iops } total-IOPS ]
                  [{ -a | --add } volume ]...
                  [{ -r | --remove } volume ]...
                  [{ -w | --wallet } wallet-location ]
                  [{ -T | --trace } trace-level ]
                  [{ -j | --json } [ --compact ]]


volumeiops list [{ -i | --id } group ]
                [{ -w | --wallet } wallet-location ]
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

```
                [{ -T | --trace } trace-level ]
                [{ -j | --json } [ --compact ]]


volumeiops delete { -i | --id } group
                  [{ -w | --wallet } wallet-location ]
                  [{ -T | --trace } trace-level ]
                  [{ -j | --json } [ --compact ]]
```

Command Options

The `volumeiops` command supports the following actions:

- `create`: Creates a volume group containing one or more volumes and configures the group with an I/O bandwidth allocation expressed in I/O operations per second (IOPS).
- `update`: Modifies the specified volume group by changing volume membership, adjusting the I/O bandwidth allocation, or both.
- `list`: Displays details for all volume groups or for the specified volume group.
- `delete`: Deletes the specified volume group and its associated I/O bandwidth allocation.

The options for the `volumeiops` command are:

- `-i`, `--id`: Identifies the name of the volume group that is the subject of the command.
- `-o`, `--iops`: Specifies the I/O bandwidth allocation expressed in I/O operations per second (IOPS).
- `-a`, `--add`: Adds a volume to the specified group. The volume is identified using a device path, such as: `/dev/exc/exa01adm01vm01_cfg_vol0128_8b144f7d0`
- `-r`, `--remove`: Removes a volume from the specified group. The volume is identified using a device path, such as: `/dev/exc/exa01adm01vm01_cfg_vol0128_8b144f7d0`
- `-w`, `--wallet`: Optionally specifies the path to the Exascale wallet directory.
- `-T`, `--trace`: Optionally enables tracing, with the trace level (*trace-level*) set to `1` (minimum tracing), `2` (medium tracing), or `3` (maximum tracing). If the trace level is not specified, then minimum tracing is enabled by default.

  The trace file is written to the first accessible location in the following list:

  1. If the `$ADR_BASE` environment variable is set:

     `$ADR_BASE/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  2. `/var/log/oracle/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

  3. `/tmp/diag/EXC/xsh_<user-name>/<host-name>/trace/xsh_<date>.trc`

- `-j`, `--json`: Displays the output in JSON format.
- `--compact`: Displays JSON formatted output in a compact format, without white space and line breaks.

Usage Notes

Note the following information when using this command:

Copyright © 2022, 2026, Oracle and/or its affiliates.

- For `update`, at least one of `{ -o | --iops }`, `{ -a | --add }`, or `{ -r | --remove }` must be specified.

Related Topics

- About Exascale Volume Groups

Copyright © 2022, 2026, Oracle and/or its affiliates.

8 Using Exascale Utility Programs

This chapter describes how to use various utility programs included with Exascale:

- edvmkvol
  Runs ESCLI commands to make an Exascale volume and associated EDV attachment.
- edvrmvol
  Runs ESCLI commands to remove an EDV attachment and associated Exascale volume previously created using the `edvmkvol` command.

8.1 edvmkvol

Runs ESCLI commands to make an Exascale volume and associated EDV attachment.

Purpose

The `edvmkvol` command provides a simple interface to create an Exascale volume and an associated Exascale Direct Volume (EDV) attachment. The command performs the ESCLI `mkvolume` and `mkvolumeattachment` commands.

Syntax

```
$ edvmkvol volume-name vault-name volume-size
          [ --local | --cluster ] [ --verbose ]
```

Command Options

The options for the `edvmkvol` command are:

- *volume-name*: Specifies the name of the volume. The value is also used as the device name for the volume attachment.
- *vault-name*: Specifies the vault that the volume is created in.
- *volume-size*: Specifies the size of the volume. The *volume-size* can be specified using suffixes `K`, `KB`, `M`, `MB`, `G`, `GB`, `T`, `TB`. The suffix is not case-sensitive.
- `--local`: Specifies the creation of a node-specific EDV attachment residing on the server running the `edvmkvol` command.
- `--cluster`: Specifies the creation of a cluster-wide EDV attachment on the Oracle Grid Infrastructure (GI) cluster that contains the server running the `edvmkvol` command.
- `--verbose`: Instructs the command to display verbose output, including the underlying ESCLI commands and their output.

Usage Notes

- The command must run on an EDV client node where the attachment will reside.
- The command uses the first Exascale wallet found in the following search path:
  1. `$OSSCONF/eswallet`

Copyright © 2022, 2026, Oracle and/or its affiliates.

2. `$ORACLE_BASE/admin/eswallet`
3. `/etc/oracle/cell/network-config/eswallet`

   The wallet provides the Exascale user credentials and endpoint address that the command uses to connect to Exascale control services, also known as the ERS endpoint. The command cannot function if the credentials are invalid or the client cannot connect to the ERS endpoint.

- If neither `--local` nor `--cluster` is specified, then `--cluster` is the default option when the host running the command contains an active Oracle Grid Infrastructure (GI) cluster. Otherwise, `--local` is the default option.

Example 8-1 Create a Volume and Cluster-Wide EDV Attachment

The example shows a simple command to create a volume and cluster-wide EDV attachment. In the example, the volume is named `my-vol1` and is created in the vault named `testVault`. The volume size is 1 TB.

```
$ edvmkvol my-vol1 testVault 1T --cluster

Created volume with id vol0002_1312220a4a9841a1be28a6c60762702a
Created edv attachment with id att0001_57172474ab234589a411178f92e4acc8
Created device: /dev/exc/my-vol1
```

Example 8-2 Create a Volume and Cluster-Wide EDV Attachment with Verbose Output

The example shows a command to create a volume and cluster-wide EDV attachment with verbose command output. The example is the same as the previous example, except that the command is also directed to display verbose output.

```
$ edvmkvol my-vol1 testVault 1T --cluster --verbose

edvutil lsinitiator
id:                                 b68b97b1-ab59-6a6e-b68b-97b1ab596a6e
hostName:                           db01vm01
state:                              ONLINE
giClusterID:                        deadbeef-badc-0fee-dead-beefbadc0fee
giClusterName:                      edvTestCluster
EDV Driver Base Version Info:
  EDV Driver Version                25.1.0.0.0.241108
EDV Driver Online Patch Version Info:
  EDV Online Patch Driver Version: None

ESCLI> mkvolume 1T --vault testVault --attributes name=my-vol1
Created volume with id vol0002_1312220a4a9841a1be28a6c60762702a

ESCLI> mkvolumeattachment vol0002_1312220a4a9841a1be28a6c60762702a my-vol1 --
giClusterId deadbeef-badc-0fee-dead-beefbadc0fee
Created edv attachment with id att0001_57172474ab234589a411178f92e4acc8
Created device: /dev/exc/my-vol1
```

Related Topics

- mkvolume
  Create a volume or volume clone.

Copyright © 2022, 2026, Oracle and/or its affiliates.

- mkvolumeattachment
  Create an attachment for an Exascale volume.
- edvrmvol
  Runs ESCLI commands to remove an EDV attachment and associated Exascale volume previously created using the `edvmkvol` command.

8.2 edvrmvol

Runs ESCLI commands to remove an EDV attachment and associated Exascale volume previously created using the `edvmkvol` command.

Purpose

The `edvrmvol` command provides a simple interface to remove an Exascale Direct Volume (EDV) attachment and associated Exascale volume that were created using the `edvmkvol` command. The command performs the ESCLI `rmvolumeattachment` and `rmvolume` commands.

Syntax

```
$ edvrmvol volume-name [ --verbose ]
```

Command Options

The options for the `edvrmvol` command are:

- *volume-name*: Specifies the name of the volume being removed. This can be either the name specified in the `edvmkvol` command (for example, `my-vol1`), or the device path (for example, `/dev/exc/my-vol1`).
- `--verbose`: Instructs the command to display verbose output, including the underlying ESCLI commands and their output.

Usage Notes

- The command uses the first Exascale wallet found in the following search path:
  1. `$OSSCONF/eswallet`
  2. `$ORACLE_BASE/admin/eswallet`
  3. `/etc/oracle/cell/network-config/eswallet`

  The wallet provides the Exascale user credentials and endpoint address that the command uses to connect to Exascale control services, also known as the ERS endpoint. The command cannot function if the credentials are invalid or the client cannot connect to the ERS endpoint.

Example 8-3 Remove an EDV Attachment and Associated Volume

The example shows a simple command to remove an EDV attachment and associated volume. In the example, the volume is named `my-vol1`.

```
$ edvrmvol my-vol1

Removed edv attachment with id att0001_57172474ab234589a411178f92e4acc8
Removed volume with id vol0002_1312220a4a9841a1be28a6c60762702a
Removed device: /dev/exc/my-vol1
```


Copyright © 2022, 2026, Oracle and/or its affiliates.

Example 8-4 Remove an EDV Attachment and Associated Volume with Verbose Output

The example shows a command to remove an EDV attachment and associated volume with verbose output. In the example, the volume is identified using the EDV device name `/dev/exc/my-vol1`.

```
$ edvrmvol /dev/exc/my-vol1 --verbose

ESCLI>lsvolumeattachment --filter deviceName=my-vol1
id
volume                                    deviceName attachTime
att0001_70b18560f4714f178d80940dc643f453
vol0002_87e95b8e9e4a4989a543c713769a867b my-vol1    2024-11-19T14:35:30+00:00

ESCLI>rmvolumeattachment att0001_70b18560f4714f178d80940dc643f453
Removed edv attachment with id att0001_70b18560f4714f178d80940dc643f453

ESCLI>rmvolume vol0002_87e95b8e9e4a4989a543c713769a867b
Removed volume with id vol0002_87e95b8e9e4a4989a543c713769a867b
```

Related Topics

- edvmkvol
  Runs ESCLI commands to make an Exascale volume and associated EDV attachment.
- rmvolumeattachment
  Delete a volume attachment.
- rmvolume
  Delete volumes.

Copyright © 2022, 2026, Oracle and/or its affiliates.

A

Exascale-Specific Information in the Oracle Data Dictionary

- Oracle Database Dictionary Views

A.1 Oracle Database Dictionary Views

The Oracle Database dictionary includes views that contain information about specific objects in Oracle Exadata Exascale.

- V$EXA_FILE
- V$EXA_TEMPLATE
- V$EXA_VAULT

A.1.1 V$EXA_FILE

The `V$EXA_FILE` view contains information about Oracle Database files stored in Exascale.

Table A-1 V$EXA_FILE Columns and Descriptions

Column	Datatype	Description
FULL_PATH	VARCHAR2(1024)	Full file path, including the vault name and file path inside the vault.
VAULT_NAME	VARCHAR2(256)	Name of the Exascale vault containing the file.
FILE_NAME	VARCHAR2(1024)	File path inside the vault.
REDUNDANCY	VARCHAR2(20)	File redundancy setting. Possible values are: • `high`: Indicates three mirrored copies of the file data.
CONTENT_TYPE	VARCHAR2(20)	File content type. Possible values are: • `DATA`: Principally associated with user data. • `RECO`: Primarily for data used in backup and recovery operations.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table A-1 (Cont.) V$EXA_FILE Columns and Descriptions

Column	Datatype	Description
MEDIA_TYPE	VARCHAR2(20)	Storage media type used by the file. The supported media types are: • HC: Identifies high capacity storage, using hard disk drives (HDDs) on high-performance Exadata storage servers. • EF: Identifies extreme flash storage, using low-latency, high-throughput flash devices. • XT: Identifies extended storage, using HDDs on Exadata storage servers optimized for cost-effective online storage of less frequently accessed data.
CREATE_TIME	DATE	File creation time.
FILE_TYPE	VARCHAR2(30)	Oracle Database file type.
SIZE_IN_BYTES	NUMBER	File size in bytes.
SPACE_USED	NUMBER	Storage space used by the file.
USED_BY	VARCHAR2(64)	Identifies the Oracle Grid Infrastructure cluster or Oracle Database associated with the file.
CON_ID	NUMBER	Oracle Database container ID.

A.1.2 V$EXA_TEMPLATE

The V$EXA_TEMPLATE view contains information about Exascale templates.

Table A-2 V$EXA_TEMPLATE Columns and Descriptions

Column	Datatype	Description
TEMPLATE_NAME	VARCHAR2(32)	Exascale template name.
VAULT_NAME	VARCHAR2(256)	Name of the Exascale vault containing the template.
REDUNDANCY	VARCHAR2(20)	File redundancy setting. Possible values are: • high: Indicates three mirrored copies of the file data.
CONTENT_TYPE	VARCHAR2(20)	File content type. Possible values are: • DATA: Principally associated with user data. • RECO: Primarily for data used in backup and recovery operations.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table A-2 (Cont.) V$EXA_TEMPLATE Columns and Descriptions

Column	Datatype	Description
MEDIA_TYPE	VARCHAR2(20)	Storage media type used by the file. The supported media types are: • HC: Identifies high capacity storage, using hard disk drives (HDDs) on high-performance Exadata storage servers. • EF: Identifies extreme flash storage, using low-latency, high-throughput flash devices. • XT: Identifies extended storage, using HDDs on Exadata storage servers optimized for cost-effective online storage of less frequently accessed data.
CON_ID	NUMBER	Oracle Database container ID.

A.1.3 V$EXA_VAULT

The V$EXA_VAULT view contains information about Exascale vaults.

Table A-3 V$EXA_VAULT Columns and Descriptions

Column	Datatype	Description
VAULT_NAME	VARCHAR2(256)	Exascale vault name.
CREATE_TIME	DATE	Vault creation time.
EF_SPACE_USED	NUMBER	Vault space used on extreme flash (EF) storage.
EF_SPACE_PROV	NUMBER	Vault space provisioned on EF storage.
EF_IOPS_PROV	NUMBER	Vault I/Os per second (IOPS) provisioned on EF storage.
HC_SPACE_USED	NUMBER	Vault space used on high capacity (HC) storage.
HC_SPACE_PROV	NUMBER	Vault space provisioned on HC storage.
HC_IOPS_PROV	NUMBER	Vault I/Os per second (IOPS) provisioned on HC storage.
XT_SPACE_USED	NUMBER	Vault space used on extended (XT) storage.
XT_SPACE_PROV	NUMBER	Vault space provisioned on XT storage.
XT_IOPS_PROV	NUMBER	Vault I/Os per second (IOPS) provisioned on XT storage.
FLASH_CACHE_PROV	NUMBER	Flash Cache space provisioned for the vault.

Copyright © 2022, 2026, Oracle and/or its affiliates.

Table A-3 (Cont.) V$EXA_VAULT Columns and Descriptions

Column	Datatype	Description
XRMEM_CACHE_PROV	NUMBER	Exadata RDMA Memory (XRMEM) Cache space provisioned for the vault.
CON_ID	NUMBER	Oracle Database container ID.

Copyright © 2022, 2026, Oracle and/or its affiliates.

B

Exascale Resources and Attributes

The ESCLI `describe` command displays information about Exascale resources and their attributes.

The following tables show information from the ESCLI `describe --all` command. The information is current for Exascale resources and their attributes in Oracle Exadata System Software release 26.1.0 (May 2026).

acfsFileSystems

Attribute	Type	Description	Settable on creation	Modifiable after creation
acfsCompat	String	ACFS compatibility (default is the ACFS software version under which the file system was created). String size: 1 - 31 characters.	Optional	false
acfsVersion	String	ACFS version under which file system was created	No	false
auSize	Integer (ub4)	Allocation unit size	No	false
autoShrinkEnabled	Boolean	State of auto-shrink on this file system	No	false
compressionEnabled	Boolean	State of compression on this file system	No	false
creationTime	Date	Time file system created	No	false
encryptionAlgorithm	String	Encryption algorithm for this file system	No	false
encryptionEnabled	Boolean	State of ACFS-based encryption on this file system	Optional	false
encryptionKeyLength	Integer (ub4)	Encryption key length (in bytes) for this file system	No	false
fileEntryTableAllocation	SpaceSize	Size of the File Entry Table	No	false
flags	String	File system state flags	No	false
giClusterId	String	Global unique identifier (GUID) for the Oracle Grid Infrastructure (GI) cluster where the file system is attached.	No	false
label	String	Volume label (default is no label). String size: 0 - 63 characters.	Optional	false
logSectorSize	Integer (ub4)	Logical sector size for this volume	No	false
metadataSize	Integer (ub4)	Metadata block size (in bytes). Valid values: '512', or '4096'. Default: 4096.	Optional	false
mountLeafGroup	Integer (ub4)	Group ownership of the mount point leaf directory (specified as a numeric ID; default is zero (the root group)).	Optional	false
mountLeafMode	String	Mode of the mount point leaf directory (specified as an octal number). String size: 3 - 4 characters. Default: 755.	Optional	false

Copyright © 2022, 2026, Oracle and/or its affiliates.

Attribute	Type	Description	Settable on creation	Modifiable after creation
mountLeafOwner	Integer (ub4)	User ownership of the mount point leaf directory (specified as a numeric ID; default is zero (the root user)).	Optional	false
mountNoRootSUID	Boolean	Disallow set user ID binary execution by non-root user (default is false; i.e., allow set user ID binary execution).	Optional	false
mountPath	String	Absolute directory path on which volume is mounted. String size: 2 - 1023 characters.	Mandatory	false
mountReadOnly	Boolean	Mount the file system read-only (default is false; i.e., mount with read-write access).	Optional	false
mountTime	Date	Time of last mount	No	false
name	String	File system name (case-sensitive). If not specified during file system creation, the name is automatically generated. String size: 1 - 63 characters.	Optional	false
nodeCount	Integer (ub8)	Number of nodes having this file system mounted	No	false
nodeId	String	Identifier for the EDV node where the file system is attached.	No	false
onDiskVersion	String	ACFS on-disk version	No	false
replicationEnabled	Boolean	State of replication on this file system	No	false
size	SpaceSize	Size of the file system in bytes (default is the size of the containing block volume). If specified, must be less than or equal to the size of the containing block volume.	Optional	false
snapShots	JSON	JSON array of ACFS snapshots of this file system	No	false
snapshotCount	Integer (ub4)	Number of ACFS snapshots associated with this file system	No	false
snapshotSpaceUsage	SpaceSize	Number of bytes consumed by ACFS snapshots of this file system	No	false
state	String	State of the volume	No	false
storageInterestTrackingEnabled	Boolean	State of storage interest tracking on this file system	No	false
totalFree	SpaceSize	Number of free bytes on the file system	No	false
user	String	Exascale user ID associated with the file system. Default: 'null'.	Optional	false
volume	String	Volume ID. String size: 1 - 63 characters.	Mandatory	false
volumeType	String	Type of volume hosting the file system	No	false

cellDisks

Attribute	Type	Description	Settable on creation	Modifiable after creation
comment	String	User-specified comment	No	false
creationTime	String	Creation timestamp	No	false

Copyright © 2022, 2026, Oracle and/or its affiliates.

Attribute	Type	Description	Settable on creation	Modifiable after creation
deviceName	String	Operating system device name for the cell disk LUN	No	false
devicePartition	String	Operating system device name for the cell disk partition	No	false
diskType	String	Disk type	No	false
errorCount	String	Total error count	No	false
freeSpace	String	Total unused space (in bytes)	No	false
name	String	Cell disk name	No	false
physicalDisk	String	Physical disk that contains the cell disk	No	false
size	String	Total cell disk size (in bytes)	No	false
status	String	Cell disk status. For example: 'normal' or 'importRequired'	No	false
uuid	String	Cell disk global unique identifier (GUID)	No	false

cells

Attribute	Type	Description	Settable on creation	Modifiable after creation
accessLevelPerm	String	Cell access level: 'remoteLoginEnabled' or 'remoteLoginDisabled'	No	false
allCellPartners	String	All partner cells	No	false
bbuStatus	String	Status of the hard disk controller battery-backed unit (BBU)	No	false
bootDeviceStatus	String	Status of the boot device	No	false
bsmStatus	String	Status of the Exascale block store manager service	No	false
bswStatus	String	Status of the Exascale block store worker service	No	false
cellVersion	String	Cell software label	No	false
cellsrvStatus	String	Status of cell service	No	false
cpuCount	String	Number of CPUs on the cell	No	false
diagHistoryDays	String	Automatic diagnostic repository (ADR) retention period (in days)	No	false
doNotServiceLEDStatus	String	Status of the cell DoNotService LED: 'on' or 'off'	No	false
efDataCellPartners	String	Extreme flash (EF) data ring partner cells	No	false
efRecoCellPartners	String	Extreme flash (EF) recovery ring partner cells	No	false
egsStatus	String	Status of Exascale cluster services	No	false
ersStatus	String	Status of Exascale REST services	No	false
events	String	Trace events set	No	false
fanCount	String	Number of fans (working/total)	No	false
fanStatus	String	Fan status: 'normal', 'warning', or 'critical'	No	false
flashCacheMode	String	Flash cache mode: 'writethrough' or 'writeback'. The default is 'writethrough'.	No	false

Copyright © 2022, 2026, Oracle and/or its affiliates.

Attribute	Type	Description	Settable on creation	Modifiable after creation
hcDataCellPartners	String	High capacity (HC) data ring partner cells	No	false
hcRecoCellPartners	String	High capacity (HC) recovery ring partner cells	No	false
httpsAccess	String	IP addresses or subnets with access to Exadata REST services using HTTPs. The value can be 'ALL', 'NONE', or a comma-separated list of IP addresses or subnet masks.	No	false
ifdStatus	String	Exascale Instant Failure Detection (IFD) status	No	false
ilomIpAddress	String	Integrated Lights Out Management (ILOM) IP address	No	false
ilomSyslogClients	String	The remote destination to forward Syslog messages from the Integrated Lights Out Manager (ILOM) service processor (SP). The value is the hostname or IP address of a remote log server or a comma-separated list of up to two log servers.	No	false
interconnect1	String	Network interface for interconnect1	No	false
interconnect2	String	Network interface for interconnect2	No	false
interconnect3	String	Network interface for interconnect3	No	false
interconnect4	String	Network interface for interconnect4	No	false
interconnectCount	String	Number of active RDMA Network Fabric interfaces	No	false
iormBoost	String	Ratio of the cumulative number of positions in the I/O queue skipped because of IORM scheduling, relative to the total number of scheduled I/Os. The value is recalculated every minute.	No	false
ipaddress1	String	ipaddress1 for the cell	No	false
ipaddress2	String	ipaddress2 for the cell	No	false
ipaddress3	String	ipaddress3 for the cell	No	false
ipaddress4	String	ipaddress4 for the cell	No	false
kernelVersion	String	Operating system kernel software version	No	false
locatorLEDStatus	String	Status of the cell Locate LED: 'on' or 'off'	No	false
makeModel	String	Cell hardware make and model	No	false
managementIpAddress	String	Management IP address	No	false
memoryGB	String	Cell memory (GB)	No	false
metricCollection	String	Specifies if Management Server (MS) performs metric collection. If 'FALSE', then no metric collection or alert mining occurs in the cell. The default setting is 'TRUE'.	No	false

Copyright © 2022, 2026, Oracle and/or its affiliates.

Attribute	Type	Description	Settable on creation	Modifiable after creation
metricFGCollIntvlInSec	String	Controls fine-grained metric collection. If the value is between 1 and 60, fine-grained metric collection is enabled, and the value represents the collection interval (in seconds). If the value is zero, then fine-grained metric collection is disabled.	No	false
metricHistoryDays	String	Retention period (in days) for metric history files	No	false
metricStreamEndPoint	String	Specifies one or more collection endpoints that automatically receive the metric stream. The value uses the following syntax: ((host='<endpoint-URL>'[,type='<stream-format>'][,token='<authentication-token>'][,{httpProxy|httpsProxy}='<proxy-server>'])[,(host='<endpoint-URL>'[,type='<stream-format>'][,token='<authentication-token>'][,{httpProxy|httpsProxy}='<proxy-server>'])]...)	No	false
metricStreamIntvlInSec	String	Controls automatic metric upload, otherwise known as streaming. If the value is zero, then streaming is disabled. Otherwise, the value represents the upload interval (in seconds) for automatic uploads to the metric streaming endpoints specified by the metricStreamEndPoint attribute.	No	false
msStatus	String	Status of Management Server	No	false
name	String	Unique name for the cell	No	false
offloadGroupEvents	String	Used only under the guidance of Oracle Support	No	false
powerCount	String	Number of power supplies (working/total)	No	false
powerStatus	String	Power status: 'normal', 'warning', or 'critical'	No	false
ramCacheMaxSize	String	Maximum possible size for the cell RAM cache	No	false
ramCacheMode	String	Cell RAM cache mode: 'on' or 'off' (disabled)	No	false
ramCacheSize	String	Size of the cell RAM cache	No	false
releaseImageStatus	String	Cell imaging success indicator	No	false
releaseTrackingBug	String	Cell software patch number. For example: 11223344	No	false
releaseVersion	String	Cell software release number. For example: 21.2.0.0.0	No	false
restSubscriber	String	REST subscriber	No	false
rollbackVersion	String	Inactive image version maintained on the cell, which is reinstated if a patchmgr rollback operation is invoked	No	false
rpmVersion	String	Cell RPM version	No	false

Copyright © 2022, 2026, Oracle and/or its affiliates.

Attribute	Type	Description	Settable on creation	Modifiable after creation
rsStatus	String	Status of Restart Server	No	false
snmpSubscriber	String	List of hosts that subscribe to the SNMP alert notifications	No	false
status	String	Status of the cell	No	false
sysEdsStatus	String	Status of the Exascale system vault manager service	No	false
syslogConf	String	Specification for forwarding Syslog messages to remote management servers. In the value, each <selector> defines a message type, and each <node> is a target server. The <selector> and <node> specifications follow the standard syntax rules for syslog.conf. The value uses the following syntax: ('<selector> @<node>' [, '<selector> @<node>']...)	No	false
syslogInput	String	Syslog input	No	false
temperatureReading	String	Current cell temperature (Celsius) from the baseboard management controller (BMC)	No	false
temperatureStatus	String	Temperature status: 'normal', 'warning', or 'critical'	No	false
traceLevel	String	Level of detail for trace messages	No	false
upTime	String	Time since system startup (days, hours:minutes)	No	false
usbStatus	String	Status of the USB device	No	false
usrEdsStatus	String	Status of the Exascale user vault manager service	No	false
uuid	String	Global unique identifier (GUID) for the cell hardware	No	false
xrmemCacheMode	String	Exadata RDMA memory (XRMEM) cache mode	No	false
xtDataCellPartners	String	Extended storage (XT) data ring partner cells	No	false
xtRecoCellPartners	String	Extended storage (XT) recovery ring partner cells	No	false

clusterAlerts

Attribute	Type	Description	Settable on creation	Modifiable after creation
alertAction	String	Alert action	No	false
alertDescription	String	Short description	No	false
alertMessage	String	Alert message	No	false
alertShortName	String	Short name	No	false
alertType	String	Alert type (case-sensitive). Valid values: 'Stateful', or 'Stateless'.	No	false
beginTime	Date	Alert begin time	No	false
clusterAlertName	String	Name	No	false

Copyright © 2022, 2026, Oracle and/or its affiliates.

Attribute	Type	Description	Settable on creation	Modifiable after creation
examinedBy	String	Name of the user that examined the alert	No	true
metricObjectName	String	Metric object name	No	false
notificationState	String	Alert notification state	No	false
releaseVersion	String	Cell image release version	No	false
rpmVersion	String	Cell image RPM version	No	false
sequenceBeginTime	Date	Alert sequence begin time	No	false
severity	String	Alert severity	No	false

clusters

Attribute	Type	Description	Settable on creation	Modifiable after creation
autoFileEncryption	Boolean	Automatic file encryption	No	true
autoGenVaultNames	Boolean	Automatic vault name generation	No	true
controlServicesFrontendPriorities	String	The priority storage servers for running front-end control service instances (expressed as <cellname1>:<priority>[;<cellnameN>:<priority>]...). Valid priority values: 1 - 255. String size: up to 660 characters.	No	true
controlServicesState	String	The availability state of control services. The possible values are: 'Available' - Indicates that the expected number of service instances are running; 'Available - reduced redundancy' - Indicates that some instances of the service are running, but the number is fewer than expected; 'Disabled' - Indicates that no instances are configured to run for this service; 'Unavailable' - Indicates that no instances of the service are running.	No	false
exaRootLeaderId	Integer (ub4)	Exascale cluster services leader ID	No	false
exaRootLeaderInc	Integer (ub4)	Exascale cluster services leader incarnation number	No	false
exaRootLeaderUrl	String	Exascale cluster services leader uniform resource locator (URL)	No	false
exaRootUrl	String	Exascale cluster services URL	No	false
guid	String	Exascale cluster globally unique identifier (GUID)	No	false

Copyright © 2022, 2026, Oracle and/or its affiliates.

Attribute	Type	Description	Settable on creation	Modifiable after creation
inclusionFilterControlServicesFrontends	String	The filter used by automatic service placement to govern placement of front-end control service instances. Any storage server hostname that matches the filter is a candidate for inclusion. The filter governs how automatic service placement starts Exascale processes on a storage server. The filter does not impact stopping an Exascale process on a storage server. A filter is a comma-separated list of sub-filters. A sub-filter string may contain any alphanumeric character. The following characters have special meanings: '*' matches any zero or more alphanumeric characters. '?' matches any one alphanumeric character. A character class is specified using a sequence of characters or character ranges within square brackets '[]'. For example, '[4-6a-c]' matches any of the following characters: '4', '5', '6', 'a', 'b', or 'c'. Preceding a character class with '!' or '^' applies a logical 'not' to the character class. For example, '[!4-6a-c]' matches any character except: '4', '5', '6', 'a', 'b', or 'c'. Examples of valid patterns include:'' (an empty string) - Matches any string. '*' - Matches any string. '?' - Matches any single character. 'pandora*' - Matches any string prefixed with 'pandora'. 'pandora[0-1][0-9]' - Matches any string starting with 'pandora' followed by numbers ranging from 00 to 19. String size: up to 255 characters.	No	true
minSupportedReleaseVersion	String	Minimum supported release version. String size: up to 255 characters.	No	true
name	String	Exascale cluster name	No	false
numControlServicesBackends	Integer (ub1)	The number of back-end control service instances to run. Valid values: 0 - 2. Default: 0.	Optional	true
numControlServicesFrontends	Integer (ub1)	The number of front-end control service instances to run. Valid values: 3 - 3. Default: 3.	Optional	true
numNodes	Integer (ub4)	Number of storage servers in the Exascale Cluster	No	false
numStorageVolumeWorkers	Integer (ub1)	The number of volume worker service instances to run on storage nodes. Valid values: 0 - 3. Default: 0.	Optional	true
numSystemVaultManagers	Integer (ub1)	The number of system vault manager service instances to run. Valid values: 3 - 5. Default: 3.	Optional	true
numUserVaultManagers	Integer (ub1)	The number of user vault manager service instances to run. Valid values: 3 - 5. Default: 3.	Optional	true

Copyright © 2022, 2026, Oracle and/or its affiliates.

Attribute	Type	Description	Settable on creation	Modifiable after creation
numVolumeManagers	Integer (ub1)	The number of volume manager service instances to run. Valid values: 0 - 3. Default: 3.	Optional	true
restEndPoint	String	Exascale REST services endpoint address. String size: up to 40 characters.	No	true
restNetworkID	Integer (ub1)	A unique id to help distinguish exascale clusters deployed on the same subnet. Valid values: 1 - 255.	No	true
storageVolumeWorkersState	String	The availability state of volume workers on storage nodes. The possible values are: 'Available' - Indicates that the expected number of service instances are running; 'Available - reduced redundancy' - Indicates that some instances of the service are running, but the number is fewer than expected; 'Disabled' - Indicates that no instances are configured to run for this service; 'Unavailable' - Indicates that no instances of the service are running.	No	false
systemVaultManagersState	String	The availability state of system vault managers. The possible values are: 'Available' - Indicates that the expected number of service instances are running; 'Available - reduced redundancy' - Indicates that some instances of the service are running, but the number is fewer than expected; 'Disabled' - Indicates that no instances are configured to run for this service; 'Unavailable' - Indicates that no instances of the service are running.	No	false
telemetryGenKeyPair	Boolean	Indicates whether to generate an authentication key pair	No	true
telemetryKeyId	String	Authentication key ID (case-sensitive). String size: up to 127 characters.	No	true
telemetryLogFileList	String	Log file list (case-sensitive). String size: up to 2047 characters.	No	true
telemetryProxyUrl	String	URL that receives Exascale alerts (case-sensitive). String size: up to 255 characters.	No	true
telemetryPubKey	String	Authentication public key (case-sensitive)	No	false
telemetryRemoveKeyPair	Boolean	Indicates whether to remove the existing authentication key pair	No	true
telemetryRootCertBundle	String	Root certificate bundle (case-sensitive). String size: up to 12287 characters.	No	true
telemetryRsyslogRelay	String	Rsyslog server for log messages (expressed as <host>:<port>). String size: up to 255 characters.	No	true
telemetryRsyslogRelayIlom	String	Rsyslog server for Integrated Lights Out Management (ILOM) log messages (expressed as <host>:<port>). String size: up to 255 characters.	No	true

Copyright © 2022, 2026, Oracle and/or its affiliates.

Attribute	Type	Description	Settable on creation	Modifiable after creation
trustStore	String	Trust store content in PEM format	No	false
trustStoreDigest	String	Trust store digest in HEX format	No	false
userVaultManagersState	String	The availability state of user vault managers. The possible values are: 'Available' - Indicates that the expected number of service instances are running; 'Available - reduced redundancy' - Indicates that some instances of the service are running, but the number is fewer than expected; 'Disabled' - Indicates that no instances are configured to run for this service; 'Unavailable' - Indicates that no instances of the service are running.	No	false
volumeManagersState	String	The availability state of volume managers. The possible values are: 'Available' - Indicates that the expected number of service instances are running; 'Available - reduced redundancy' - Indicates that some instances of the service are running, but the number is fewer than expected; 'Disabled' - Indicates that no instances are configured to run for this service; 'Unavailable' - Indicates that no instances of the service are running.	No	false

computes

Attribute	Type	Description	Settable on creation	Modifiable after creation
accessLevelPerm	String	Compute node access level: 'remoteLoginEnabled' or 'remoteLoginDisabled'	No	false
bswStatus	String	Status of the Exascale block store worker service	No	false
cpuCount	String	Number of CPUs on the compute node	No	false
diagHistoryDays	String	Automatic diagnostic repository (ADR) retention period (in days)	No	false
doNotServiceLEDStatus	String	Status of the server DoNotService LED: 'on' or 'off'	No	false
edvStatus	String	Status of the Exascale direct volume service	No	false
egsStatus	String	Status of Exascale cluster services	No	false
esnpStatus	String	Status of the Exascale node proxy service	No	false
events	String	Trace events set	No	false
fanCount	String	Number of fans (working/total)	No	false
fanStatus	String	Fan status: 'normal', 'warning', or 'critical'	No	false
flashCacheMode	String	Flash cache mode: 'writethrough' or 'writeback'. The default is 'writethrough'.	No	false

Copyright © 2022, 2026, Oracle and/or its affiliates.

Attribute	Type	Description	Settable on creation	Modifiable after creation
httpsAccess	String	IP addresses or subnets with access to Exadata REST services using HTTPs. The value can be 'ALL', 'NONE', or a comma-separated list of IP addresses or subnet masks.	No	false
ilomIpAddress	String	Integrated Lights Out Management (ILOM) IP address	No	false
interconnect1	String	Network interface for interconnect1	No	false
interconnect2	String	Network interface for interconnect2	No	false
interconnect3	String	Network interface for interconnect3	No	false
interconnect4	String	Network interface for interconnect4	No	false
interconnectCount	String	Number of active RDMA Network Fabric interfaces	No	false
iormBoost	String	Ratio of the cumulative number of positions in the I/O queue skipped because of IORM scheduling, relative to the total number of scheduled I/Os. The value is recalculated every minute.	No	false
ipaddress1	String	ipaddress1 for the compute node	No	false
ipaddress2	String	ipaddress2 for the compute node	No	false
ipaddress3	String	ipaddress3 for the compute node	No	false
ipaddress4	String	ipaddress4 for the compute node	No	false
kernelVersion	String	Operating system kernel software version	No	false
locatorLEDStatus	String	Status of the server Locate LED: 'on' or 'off'	No	false
makeModel	String	Compute hardware make and model	No	false
managementIpAddress	String	Management network IP address	No	false
memoryGB	String	Compute node memory (GB)	No	false
metricHistoryDays	String	Retention period (in days) for metric history files	No	false
msStatus	String	Status of Management Server	No	false
msVersion	String	Exadata software release number. For example: 21.2.0.0.0	No	false
name	String	Unique compute node name	No	false
offloadGroupEvents	String	Used only under the guidance of Oracle Support	No	false
powerCount	String	Number of power supplies (working/total)	No	false
powerStatus	String	Power status: 'normal', 'warning', or 'critical'	No	false
ramCacheMaxSize	String	Maximum possible RAM cache size	No	false
ramCacheMode	String	RAM cache mode: 'on' or 'off' (disabled)	No	false
ramCacheSize	String	RAM cache size	No	false
releaseImageStatus	String	Imaging success indicator	No	false
releaseTrackingBug	String	Exadata software patch number. For example: 11223344	No	false

Copyright © 2022, 2026, Oracle and/or its affiliates.

Attribute	Type	Description	Settable on creation	Modifiable after creation
releaseVersion	String	Exadata software release number. For example: 21.2.0.0.0	No	false
rpmVersion	String	Compute node RPM version	No	false
rsStatus	String	Status of Restart Server	No	false
snmpSubscriber	String	List of hosts that subscribe to the SNMP alert notifications	No	false
status	String	Compute node status	No	false
temperatureReading	String	Current temperature (Celsius) from the baseboard management controller (BMC)	No	false
temperatureStatus	String	Temperature status: 'normal', 'warning', or 'critical'	No	false
traceLevel	String	Level of detail for trace messages	No	false
upTime	String	Time since system startup (days, hours:minutes)	No	false
usbStatus	String	Status of the USB device	No	false
uuid	String	Global unique identifier (GUID) for the compute hardware	No	false
xrmemCacheMode	String	Exadata RDMA memory (XRMEM) cache mode	No	false

containerResources

Attribute	Type	Description	Settable on creation	Modifiable after creation
flashCacheProv	SpaceSize	Provisioned flash cache space (in bytes). Valid values: up to 4503599626321920.	No	true
flashLogProv	Boolean	Specifies whether flash log is provisioned	No	true
guid	String	Storage pool global unique identifier (GUID)	No	true
iopsProv	IopsCount	IOPS Provisioned. Valid values: up to 4294967295.	No	true
mediaType	String	Storage pool media type. Valid values: 'HC', 'EF', or 'XT'.	No	true
spaceProv	SpaceSize	Provisioned storage space (in bytes). Valid values: up to 4503599626321920.	No	true
spaceUsed	SpaceSize	Used space (in bytes). Valid values: up to 4503599626321920.	No	true
storagePool	String	Storage pool name. String size: 1 - 64 characters.	No	true
vault	String	The name of the vault containing the provisioned resources	No	true
xrmemCacheProv	SpaceSize	Provisioned XRMEM cache space (in bytes). Valid values: up to 4503599626321920.	No	true
xrmemLogProv	Boolean	Specifies whether XRMEM log is provisioned	No	true

Copyright © 2022, 2026, Oracle and/or its affiliates.

datasets

Attribute	Type	Description	Settable on creation	Modifiable after creation
datasetSpaceUsedEF	SpaceSize	Total space in EF media used by files in this and descendant datasets.	No	false
datasetSpaceUsedHC	SpaceSize	Total space in HC media used by files in this and descendant datasets.	No	false
datasetSpaceUsedXT	SpaceSize	Total space in XT media used by files in this and descendant datasets.	No	false
fileSpaceUsedEF	SpaceSize	Total space in EF media used by files in this dataset.	No	false
fileSpaceUsedHC	SpaceSize	Total space in HC media used by files in this dataset.	No	false
fileSpaceUsedXT	SpaceSize	Total space in XT media used by files in this dataset.	No	false
name	String	Dataset name. String size: up to 255 characters.	Optional	false
storagePoolSpaceUsed	JSON	JSON array of storage pool space usage statistics.	No	false

edvNodes

Attribute	Type	Description	Settable on creation	Modifiable after creation
edvBaseDriverVersion	String	Base version of the EDV driver resulting from the last system update or patch installation with reboot	No	false
edvEffectiveDriverVersion	String	Currently effective version of the EDV driver. Contains the base version plus any Ksplice updates that are in effect	No	false
giClusterId	String	Global unique identifier (GUID) for the Oracle Grid Infrastructure (GI) cluster that contains the EDV node	No	false
giClusterName	String	Name of the GI cluster that contains the EDV node	No	false
hostName	String	EDV node name	No	false
lastHeartbeat	Date	EDV node last heartbeat time	No	false
registerTime	Date	EDV node registration time	No	false
state	String	EDV node state. Valid values: 'INITIAL', 'REGISTERING', 'OFFLINING', or 'ONLINE'.	No	false
user	String	EDV node storage owner (Exascale user ID)	No	false
version	String	Software version of the EDV node	No	false

Copyright © 2022, 2026, Oracle and/or its affiliates.

edvVolumeAttachments

Attribute	Type	Description	Settable on creation	Modifiable after creation
atomicToken	String	Identifier for an atomic operation that involves multiple volume attachments. Default: 'FALSE'.	Optional	true
attachTime	Date	Creation timestamp	No	false
deviceName	String	Name of the storage device associated with the attachment. String size: 1 - 127 characters.	Mandatory	false
devicePath	String	Absolute path of the storage device on the EDV node	No	false
enabledOnRestart	Boolean	Indicates whether the volume attachment will be enabled if the initiator restarts	No	false
giClusterId	String	Global unique identifier (GUID) for the Oracle Grid Infrastructure (GI) cluster that contains the EDV attachment. When creating an EDV attachment, either 'giClusterId' or 'initiator' must be specified. String size: 32 - 36 characters.	Optional	false
giClusterName	String	Name of the GI cluster that contains the EDV attachment	No	false
hostName	String	Host name associated with the attachment	No	false
initiator	String	Identifier for the EDV node that contains the attachment. When creating an EDV attachment, either 'giClusterId' or 'initiator' must be specified. String size: 32 - 36 characters.	Optional	false
kernelDeviceName	String	Kernel device name of the storage device associated with the attachment	No	false
logicalSectorSize	Integer (ub4)	Volume logical sector size	No	false
user	String	Exascale user ID associated with the attachment. Default: 'null'.	Optional	false
volume	String	Attached volume ID. String size: 1 - 63 characters.	Mandatory	false
volumeSnapshot	String	Snapshot identifier. Used only for snapshot attachments.	No	false

files

Attribute	Type	Description	Settable on creation	Modifiable after creation
acl	String	Access control list. String size: up to 16384 characters.	No	true
contentType	String	Content type. If not specified during file creation, the value is derived from the file type and associated template. Valid values: 'data', or 'reco'.	Optional	false
createTime	Date	File creation time	No	false

Copyright © 2022, 2026, Oracle and/or its affiliates.

Attribute	Type	Description	Settable on creation	Modifiable after creation
createdBy	String	File owner	No	false
fileType	String	Oracle file type: 'ARCHIVELOG', 'CONTROLFILE', 'DATAFILE', 'ONLINELOG', and so on. Default: 'null'.	Optional	false
mediaType	String	Media type. If not specified during file creation, the value is derived from the file type and associated template. Valid values: 'HC', 'EF', or 'XT'.	Optional	false
name	String	File name	Optional	false
redundancy	String	Redundancy setting. If not specified during file creation, the value is derived from the file type and associated template.	Optional	false
size	SpaceSize	File size (in bytes). Default: 0.	Optional	true
snapLockExpiryTime	Date	Snapshot lock expiry date and time (applies only to snapshot files). To unlock a locked snapshot, set the expiry time to 'null' (snapLockExpiryTime='null'). You can modify this attribute only if the snapshot lock is not immutable (snapLockImmutable=FALSE).	No	true
snapLockImmutable	Boolean	Indicates whether the lock is immutable. When a lock is immutable, the file cannot be unlocked before the lock expires, and the lock cannot be modified to expire sooner (applies only to snapshot files).	No	true
snapLockStatus	String	Indicates the snapshot lock status: 'locked' if the lock is active, or 'expired' if the lock expiry time has passed (applies only to snapshot files with a lock expiry time).	No	false
spaceUsed	SpaceSize	Vault space used by the file (in bytes).	No	false
storagePool	String	Name of the storage pool that contains the file. String size: 1 - 64 characters.	Optional	false
templateName	String	Name of the template used only during file creation (case-sensitive). The attribute value is not retained after file creation. String size: 1 - 31 characters.	Optional	false
usedBy	String	Oracle database that uses the file	No	false
vaultName	String	Name of the vault containing the file. The value may also be derived from the URL for some APIs. For example, myVault is the vault name in the following API call: PUT /files/@myVault/myFile	Optional	false
xAttrs	String	Comma-separated extended file attribute names	No	false

Copyright © 2022, 2026, Oracle and/or its affiliates.

gridDisks

Attribute	Type	Description	Settable on creation	Modifiable after creation
asmDiskGroupName	String	Oracle ASM disk group containing the grid disk	No	false
asmDiskName	String	Oracle ASM disk name	No	false
asmFailGroupName	String	Oracle ASM failure group	No	false
availableTo	String	Clients with access to the grid disk	No	false
cachedBy	String	Flash disks currently used by write-back flash cache to cache grid disk data	No	false
cachingPolicy	String	Flash cache policy: 'DEFAULT' indicates that the grid disk uses flash cache. 'NONE' indicates that the grid disk does not use flash cache.	No	false
cell	String	Name of the cell containing the grid disk	No	false
cellDisk	String	Name of the cell disk containing the grid disk	No	false
comment	String	User-specified comment	No	false
creationTime	String	Creation timestamp	No	false
diskType	String	Disk type	No	false
errorCount	String	Total error count	No	false
name	String	Grid disk name	No	false
poolDiskStatus	String	Status of the associated pool disk	No	false
size	String	Total grid disk size (in bytes)	No	false
sizeAllocated	String	Total used space (in bytes) (applies only to sparse grid disks)	No	false
spaceUsed	String	Total used space (in bytes) (applies only to Exascale grid disks)	No	false
status	String	Grid disk status. For example: 'ACTIVE', 'INACTIVE', 'NOT PRESENT', or 'IMPORTREQUIRED'.	No	false
storagePool	String	Storage pool that contains the grid disk	No	false
uuid	String	Grid disk global unique identifier (GUID)	No	false

poolDisks

Attribute	Type	Description	Settable on creation	Modifiable after creation
cellId	Integer (ub4)	Cell ID for the cell containing the pool disk	No	false
gdName	String	Grid disk name associated with the pool disk	No	false
healthy	Boolean	Health indicator	No	false
incarnation	Integer (ub4)	Incarnation number	No	false
name	String	Pool disk name	No	false
size	SpaceSize	Total pool disk size (in bytes)	No	false

Copyright © 2022, 2026, Oracle and/or its affiliates.

Attribute	Type	Description	Settable on creation	Modifiable after creation
state	String	Life-cycle state	No	false
status	String	Pool disk status. You can only modify the pool disk status to 'ONLINE' or 'OFFLINE'. Valid values: 'BEING_ADDED', 'BEING_ADDED_FORCE', 'BEING_ONLINED', 'ONLINE', 'BEING_OFFLINED', 'OFFLINE', 'BEING_OFFLINED_MAINTENANCE', 'OFFLINE_MAINTENANCE', 'BEING_DROPPED', 'DROPPED', 'BEING_DROPPED_FORCE', 'DROPPED_FORCE', 'BEING_DROPPED_PERMANENT', or 'BEING_DROPPED_FORCE_PERMANENT'.	No	true
storagePool	String	Storage pool that contains the pool disk	No	false

resourceProfiles

Attribute	Type	Description	Settable on creation	Modifiable after creation
enableFlashCache	Boolean	Enable or disable use of flash cache for clients associated with the resource profile. A 'true' value means enabled. Default: 'true'.	Optional	true
enableFlashLog	Boolean	Enable or disable use of flash log for clients associated with the resource profile. A 'true' value means enabled. Default: 'true'.	Optional	true
enableXrmemCache	Boolean	Enable or disable use of XRMEM cache for clients associated with the resource profile. A 'true' value means enabled. Default: 'true'.	Optional	true
flashCacheMin	String	Guaranteed minimum share of flash cache available to each client. The value is a percentage with up to two decimal places.Valid values: 0 - 100. Default: 0. If the sum of all minimum values exceeds 100% across all clients and resource profiles, then all values are scaled down proportionally to ensure that share guarantees can be honored. Valid only if 'enableFlashCache=true'. A zero (0) value means no set minimum.	Optional	true
flashCacheSize	String	Upper limit of flash cache available to each client. The value is a percentage with up to two decimal places.Valid values: 0 - 100. Default: 100. Valid only if 'enableFlashCache=true'.	Optional	true

Copyright © 2022, 2026, Oracle and/or its affiliates.

Attribute	Type	Description	Settable on creation	Modifiable after creation
iopsLimitEF	String	Upper limit of extreme flash (EF) media IOPS available to each client. The value is a percentage with up to two decimal places.Valid values: 0 - 100. Default: 100.	Optional	true
iopsLimitHC	String	Upper limit of high capacity (HC) media IOPS available to each client. The value is a percentage with up to two decimal places.Valid values: 0 - 100. Default: 100.	Optional	true
iopsLimitXT	String	Upper limit of extended storage (XT) media IOPS available to each client. The value is a percentage with up to two decimal places.Valid values: 0 - 100. Default: 100.	Optional	true
iopsShareEF	Integer (ub4)	Share of IOPS from extreme flash (EF) media available to each client. Each client's share is relative to all other client shares. A higher share value implies higher priority. For example, a client with a share value of 2 gets twice as many IOPS as a client with a share value of 1, but both clients get far less than another client with a share value of 50. Valid values: 0 - 100. Default: 1.	Optional	true
iopsShareHC	Integer (ub4)	Share of IOPS from high capacity (HC) media available to each client. Each client's share is relative to all other client shares. A higher share value implies higher priority. For example, a client with a share value of 2 gets twice as many IOPS as a client with a share value of 1, but both clients get far less than another client with a share value of 50. Valid values: 0 - 100. Default: 1.	Optional	true
iopsShareXT	Integer (ub4)	Share of IOPS from extended storage (XT) media available to each client. Each client's share is relative to all other client shares. A higher share value implies higher priority. For example, a client with a share value of 2 gets twice as many IOPS as a client with a share value of 1, but both clients get far less than another client with a share value of 50. Valid values: 0 - 100. Default: 1.	Optional	true
name	String	Resource profile name. String size: up to 32 characters.	Mandatory	false
vault	String	Vault name. String size: up to 30 characters.	Mandatory	false

Copyright © 2022, 2026, Oracle and/or its affiliates.

Attribute	Type	Description	Settable on creation	Modifiable after creation
xrmemCacheMin	String	Guaranteed minimum share of XRMEM cache available to each client. The value is a percentage with up to two decimal places. Valid values: 0 - 100. Default: 0. If the sum of all minimum values exceeds 100% across all clients and resource profiles, then all values are scaled down proportionally to ensure that share guarantees can be honored. Valid only if 'enableXrmemCache=true'. A zero (0) value means no set minimum.	Optional	true
xrmemCacheSize	String	Upper limit of XRMEM cache available to each client. The value is a percentage with up to two decimal places.Valid values: 0 - 100. Default: 100. Valid only if 'enableXrmemCache=true'.	Optional	true

services

Attribute	Type	Description	Settable on creation	Modifiable after creation
cellName	String	Cell name on the server running the service	No	false
cssBirthIncarnation	Integer (ub4)	CSS cluster birth incarnation. Only for Exascale Node Proxy (ESNP).	No	false
cssClusterName	String	CSS cluster name. Only for Exascale Node Proxy (ESNP).	No	false
cssGUID	String	CSS cluster GUID. Only for Exascale Node Proxy (ESNP).	No	false
cssNodeID	Integer (ub4)	CSS node ID. Only for Exascale Node Proxy (ESNP).	No	false
cssNodeIncarnation	Integer (ub4)	CSS node incarnation. Only for Exascale Node Proxy (ESNP).	No	false
egsName	String	Service name chosen by Exascale cluster services	No	false
frontend	Boolean	Frontend deployed. Only for Exascale REST services (controlServices).	No	false
frontendPriority	Integer (ub4)	Frontend priority. Only for Exascale REST services (controlServices).	No	false
hostName	String	Host name of the server running the service	No	false
incarnation	Integer (ub4)	Incarnation number for the service	No	false
leader	Boolean	Service has leader role. Only for Exascale cluster services (rootServices) and volume manager services.	No	false
name	String	Service name	No	false
rackId	String	Rack identifier for the service	No	false

Copyright © 2022, 2026, Oracle and/or its affiliates.

Attribute	Type	Description	Settable on creation	Modifiable after creation
serviceType	String	Service type. Valid values: 'cellServices', 'systemVaultManagers', 'userVaultManagers', 'controlServices', 'rootServices', 'volumeManagers', or 'volumeWorkers'.	No	false
status	String	Service status	No	false
url	String	Service uniform resource locator (URL)	No	false
virtualRouterId	Integer (ub4)	Keepalived virtual router ID. Only for Exascale REST services (controlServices).	No	false

storagePoolOperations

Attribute	Type	Description	Settable on creation	Modifiable after creation
bytesCompleted	SpaceSize	Bytes completed	No	false
bytesRemaining	SpaceSize	Bytes remaining	No	false
contentType	String	Content Type	No	false
estimatedCompletionTime	String	Estimated completion time of the operation	No	false
name	String	Storage pool name	No	false
opType	String	Operation Type	No	false
startTime	String	Start time of the operation	No	false
state	String	State of the operation	No	false

storagePools

Attribute	Type	Description	Settable on creation	Modifiable after creation
cellOfflineTimerInMins	Integer (ub4)	Cell offline timer (in minutes). The timer starts when a cell goes offline. If the cell comes back online before the timer expires, the affected disks are automatically updated and returned to the storage pool. Otherwise, when the timer expires, the disks are dropped and the storage pool is reconfigured.	No	true
cellsCount	Integer (ub4)	Number of storage servers that host the storage pool	No	false
diskOfflineTimerInMins	Integer (ub4)	Disk offline timer (in minutes). The timer starts when a disk goes offline. If the disk comes back online before the timer expires, the disk is automatically updated and returned to the storage pool. Otherwise, when the timer expires, the disk is dropped and the storage pool is reconfigured.	No	true

Copyright © 2022, 2026, Oracle and/or its affiliates.

Attribute	Type	Description	Settable on creation	Modifiable after creation
disksCount	Integer (ub4)	Number of disks used by the storage pool	No	false
flashCacheProvisionableSize	SpaceSize	Total flash cache space (in bytes) available for provisioning	No	false
flashCacheProvisionedSize	SpaceSize	Provisioned flash cache space (in bytes)	No	false
flashCacheSize	SpaceSize	Total flash cache space (in bytes) used by the storage pool	No	false
guid	String	Storage pool global unique identifier (GUID)	No	false
highDataStatus	String	Status of the high-redundancy data ring	No	false
highRecoStatus	String	Status of the high-redundancy recovery ring	No	false
iopsProvisionable	Integer (ub8)	Total IOPS available for provisioning	No	false
iopsProvisioned	Integer (ub8)	Provisioned IOPS	No	false
iopsRaw	Integer (ub8)	Total IOPS (I/Os per second) provisioned for the storage pool	No	false
mediaType	String	Media type. Valid values: 'HC', 'EF', or 'XT'.	No	false
name	String	Storage pool name. String size: 1 - 64 characters.	Mandatory	false
normalDataStatus	String	Status of the normal-redundancy data ring	No	false
normalRecoStatus	String	Status of the normal-redundancy recovery ring	No	false
spaceProvisionable	SpaceSize	Total storage space (in bytes) available for provisioning	No	false
spaceProvisioned	SpaceSize	Provisioned storage space (in bytes)	No	false
spaceRaw	SpaceSize	Total storage space (in bytes) used by the storage pool	No	false
spaceUsed	SpaceSize	Used space (in bytes)	No	false
sysDataStorePresent	Boolean	Indicator that the storage pool contains the system metadata store. Default: 'false'.	Optional	false
sysDataStoreShardsCount	Integer (ub4)	Number of system metadata store shards (if applicable)	No	false
xrmemCacheProvisionableSize	SpaceSize	Total XRMEM cache space (in bytes) available for provisioning	No	false
xrmemCacheProvisionedSize	SpaceSize	Provisioned XRMEM cache space (in bytes)	No	false
xrmemCacheSize	SpaceSize	Total Exadata RDMA Memory (XRMEM) cache space (in bytes) used by the storage pool	No	false

Copyright © 2022, 2026, Oracle and/or its affiliates.

users

Attribute	Type	Description	Settable on creation	Modifiable after creation
category	String	User category. Valid values: 'compute_server', 'storage_server', 'internal', or 'external'.	Optional	false
name	String	User name. The value must start with an alphanumeric character and contain only alphanumeric, hyphen ('-'), and period ('.') characters. If not specified during user creation, the user name matches the user ID. When creating a user, you must specify either 'id' or 'name'. String size: 1 - 63 characters. Default: 'null'.	Optional	true
privilege	String	List of user privileges. Valid combinations of: 'no_privilege', 'cellsrv', 'egs', 'ers', 'syseds', 'usreds', 'bsm', 'bsw', 'ms', 'vlt_manage', 'vlt_use', 'vlt_read', 'vlt_inspect', 'cl_admin', 'cl_operator', 'cl_monitor', 'on_behalf_of', 'user_create', and 'system_restore'. Each privilege is separated by '|'. For example: 'vlt_read | cl_operator'. Default: 'vlt_inspect'.	Optional	true
pubKeyFingerprint1	String	User public key fingerprint (ID 1)	No	false
pubKeyFingerprint2	String	User public key fingerprint (ID 2)	No	false
pubKeyFingerprint3	String	User public key fingerprint (ID 3)	No	false
publicKey1	String	User public key (ID 1). String size: up to 511 characters. Default: 'null'.	Optional	true
publicKey2	String	User public key (ID 2). String size: up to 511 characters. Default: 'null'.	Optional	true
publicKey3	String	User public key (ID 3). String size: up to 511 characters. Default: 'null'.	Optional	true

vaults

Attribute	Type	Description	Settable on creation	Modifiable after creation
acl	String	Access control list. String size: up to 16384 characters.	No	true
createTime	Date	Vault creation time	No	false
createdBy	String	Vault owner	No	false
flags	String	Vault flags: 'being_deleted'	No	false

Copyright © 2022, 2026, Oracle and/or its affiliates.

Attribute	Type	Description	Settable on creation	Modifiable after creation
flashCacheProv	SpaceSize	In sharedCloud deployment mode, this attribute represents the additional flash cache provisioned (in bytes) on top of the auto provisioned value. A zero value in this case indicates no additional flash cache is provisioned. In other deployment modes, this attribute represents the total space provisioned in flash cache (in bytes). A zero value disables flash cache use for the vault in such cases. Valid values: up to 4503599626321920.	Optional	true
flashLogProv	Boolean	Enable/disable flash log usage. Default: 'true'.	Optional	true
iopsProvEF	IopsCount	IOPS provisioned in EF media. Valid values: up to 4294967295, or 'unlimited'. Default: 'unlimited'.	Optional	true
iopsProvHC	IopsCount	IOPS provisioned in HC media. Valid values: up to 4294967295, or 'unlimited'. Default: 'unlimited'.	Optional	true
iopsProvXT	IopsCount	IOPS provisioned in XT media. Valid values: up to 4294967295, or 'unlimited'. Default: 'unlimited'.	Optional	true
name	String	Vault name. If not specified during vault creation, a unique name is automatically generated. String size: up to 30 characters.	Optional	false
spaceProvEF	SpaceSize	Vault space (in bytes) provisioned in EF media. Valid values: up to 4503599626321920, or 'unlimited'. Default: 'unlimited'.	Optional	true
spaceProvHC	SpaceSize	Vault space (in bytes) provisioned in HC media. Valid values: up to 4503599626321920, or 'unlimited'. Default: 'unlimited'.	Optional	true
spaceProvXT	SpaceSize	Vault space (in bytes) provisioned in XT media. Valid values: up to 4503599626321920, or 'unlimited'. Default: 'unlimited'.	Optional	true
spaceUsedEF	SpaceSize	Extreme flash (EF) media space used (in bytes)	No	false
spaceUsedHC	SpaceSize	High capacity (HC) media space used (in bytes)	No	false
spaceUsedXT	SpaceSize	Extended storage (XT) media space used (in bytes)	No	false
xAttrKeys	String	Comma-separated extended vault attribute names	No	false
xAttrs	JSON	JSON array of vault extended attributes, which is only displayed when the extended attribute query string ('?withXAttrs') is specified	No	false

Copyright © 2022, 2026, Oracle and/or its affiliates.

Attribute	Type	Description	Settable on creation	Modifiable after creation
xrmemCacheProv	SpaceSize	Space provisioned in Exadata RDMA Memory (XRMEM) cache (in bytes). A zero value disables XRMEM cache use for the vault. Use only on systems with HC or EF media, but not both. For systems with both HC and EF media, use xrmemCacheProvHC and xrmemCacheProvEF. This attribute cannot be specified in sharedCloud deployment mode. Valid values: up to 4503599626321920, or 'unlimited'. Default: 'unlimited'.	Optional	true
xrmemCacheProvEF	SpaceSize	Space provisioned in XRMEM cache for data stored on EF media (in bytes). A zero value disables the vault from using XRMEM cache for data stored on EF media. This attribute cannot be specified in sharedCloud deployment mode. Valid values: up to 4503599626321920, or 'unlimited'. Default: 'unlimited'.	Optional	true
xrmemCacheProvHC	SpaceSize	Space provisioned in XRMEM cache for data stored on HC media (in bytes). A zero value disables the vault from using XRMEM cache for data stored on HC media. This attribute cannot be specified in sharedCloud deployment mode. Valid values: up to 4503599626321920, or 'unlimited'. Default: 'unlimited'.	Optional	true

volumeBackups

Attribute	Type	Description	Settable on creation	Modifiable after creation
backupType	String	Backup type. Valid values: 'incremental', or 'full'.	Optional	false
contentType	String	Content type (only applies when an Exascale vault is the backup destination). Valid values: 'DATA', or 'RECO'.	Optional	false
creationTime	Date	Creation timestamp	No	false
destination	String	Backup destination	Optional	false
name	String	Volume backup name (case-sensitive). String size: 1 - 255 characters.	Optional	true
progress	String	Completion percentage	No	false
size	SpaceSize	Backup size (in bytes)	No	false
state	String	Backup state. Valid values: 'CREATING', 'AVAILABLE', 'DELETING', 'DELETED', 'RESTORING', 'FAILED', or 'UNKNOWN'.	No	false
volume	String	Volume ID	No	false
volumeSnapshot	String	Volume snapshot ID	Mandatory	false

Copyright © 2022, 2026, Oracle and/or its affiliates.

volumeGroups

Attribute	Type	Description	Settable on creation	Modifiable after creation
creationTime	Date	Volume group creation timestamp	No	false
iopsProvisioned	IopsCount	Provisioned throughput available to the volume group (in IOPS). Valid values: 0 - 4294967295. Default: 0.	Optional	true
name	String	Volume group name (case-sensitive). If not specified during creation, the name is automatically generated. String size: 1 - 255 characters.	Optional	true
owners	String	List of volume group owners. Each Volume has at least one owner and up to two owners. Owners are separated by ','. For example: 'user1,user2'.	No	true
resourceProfile	String	Resource profile associated with the volume group. String size: 1 - 32 characters.	No	true
resourceSharing	String	Specifies the resource sharing policy applicable to the group. Valid values: 'aggregate', 'specified', or 'none'. Default: 'none'.	Optional	true
state	String	Volume group state. Valid values: 'UNAVAILABLE', 'AVAILABLE', 'DELETING', 'DELETED', 'RENAMING', or 'CHANGING IOPS'.	No	false
volumes	String	List of volumes associated with the volume group.	No	true

volumeSnapshots

Attribute	Type	Description	Settable on creation	Modifiable after creation
contentType	String	Content type (derived from the parent volume). For a volume snapshot, 'data' is the only valid value.	No	false
creationTime	Date	Creation timestamp	No	false
fileName	String	Name of the Exascale file that contains the volume snapshot data	No	false
filePath	String	File path of the Exascale file that contains the volume snapshot data	No	false
iopsProvisioned	Integer (ub4)	Provisioned throughput (in IOPS) available to the snapshot	No	false
lockExpiryTime	Date	Lock expiry date and time. To unlock a locked volume snapshot, set the expiry time to 'null' (lockExpiryTime='null'). You can modify this attribute only if the lock is not immutable (lockImmutable=FALSE).	No	true

Copyright © 2022, 2026, Oracle and/or its affiliates.

Attribute	Type	Description	Settable on creation	Modifiable after creation
lockImmutable	Boolean	Indicates whether the lock is immutable. When a lock is immutable, the volume snapshot cannot be unlocked before the lock expires, and the lock cannot be modified to expire sooner.	No	true
lockStatus	String	Indicates the volume snapshot lock status: 'locked' if the lock is active, or 'expired' if the lock expiry time has passed (applies only to volume snapshots with a lock expiry time).	No	false
mediaType	String	Media type (derived from the parent volume). Valid values: 'HC', 'EF', or 'XT'.	No	false
name	String	Volume snapshot name (case-sensitive). String size: 1 - 255 characters.	Optional	true
numAttachments	Integer (ub4)	Number of attachments associated with the snapshot	No	false
redundancy	String	Redundancy setting (derived from the parent volume)	No	false
size	SpaceSize	Snapshot size (derived from the parent volume)	No	false
spaceUsed	SpaceSize	Space (in bytes) used by the snapshot (including all mirrors)	No	false
state	String	Snapshot state. Valid values: 'UNAVAILABLE', 'AVAILABLE', 'DELETING', 'DELETED', 'ATTACHING', or 'DETACHING'.	No	false
user	String	Volume snapshot owners. Each snapshot has at least one owner and up to two owners. Owners are separated by ','. For example: 'user1,user2'.	No	false
vault	String	Name of the vault containing the snapshot (derived from the parent volume)	No	false
volume	String	Parent volume ID. String size: 32 - 64 characters.	Mandatory	false
volumeBackup	String	Backup ID associated with the snapshot	No	false
volumeGroup	Boolean	'true', create snaps for a volume group. 'false', otherwise.'. Default: 'FALSE'.	Optional	false
volumeGroupSnapshot	String	ID of the volume group snapshot operation associated with the volume snapshot	No	false

volumes

Attribute	Type	Description	Settable on creation	Modifiable after creation
acfsFileSystem	String	ACFS File System (if any) hosted by the volume	No	false
contentType	String	Content type. Valid values: 'DATA', or 'RECO'. Default: 'DATA'.	Optional	false
createdBy	String	The user that created the volume	No	false

Copyright © 2022, 2026, Oracle and/or its affiliates.

Attribute	Type	Description	Settable on creation	Modifiable after creation
creationTime	Date	Creation timestamp	No	false
fileName	String	Name of the Exascale file that contains the volume data (case-sensitive). String size: 1 or more characters.	Optional	false
filePath	String	File path of the Exascale file that contains the volume data	No	false
iopsInherited	Boolean	'true' share the IOPS with the parent volume. 'false' the clone volume IOPS are independent from its parent. Default: 'TRUE'.	Optional	true
iopsProvisioned	IopsCount	Provisioned throughput available to the volume (in IOPS). Valid values: 1 - 4294967295, or 'unlimited'. Default: 'unlimited'.	Optional	true
log	Boolean	Indicates whether the volume contains Oracle Database redo log files.	Optional	false
mediaType	String	Media type. Valid values: 'HC', 'EF', or 'XT'. Default: 'HC'.	Optional	false
name	String	Volume name (case-sensitive). If not specified during volume creation, the name is automatically generated. String size: 1 - 255 characters.	Optional	true
numAttachments	Integer (ub4)	Number of attachments associated with the volume	No	false
numBackups	Integer (ub4)	Total number of backups of the volume	No	false
numSnapshots	Integer (ub4)	Total number of snapshots of the volume	No	false
owners	String	List of volume owners. Each Volume has at least one owner and up to two owners. Owners are separated by ','. For example: 'user1,user2'.	No	true
redundancy	String	Redundancy setting. Default: 'high'.	Optional	false
resourceProfile	String	Resource profile associated with this volume. String size: 1 - 32 characters. Default: 'none'.	Optional	true
restoreProgress	String	Progress percentage during volume restoration	No	false
size	SpaceSize	Volume size (in bytes) Cannot be specified for volume clones or when restoring a backup to an existing volume.	Optional	true
spaceUsed	SpaceSize	Space (in bytes) used by the volume (including all mirrors)	No	false
state	String	Volume state. Valid values: 'UNAVAILABLE', 'AVAILABLE', 'DELETING', 'DELETED', 'ATTACHING', or 'DETACHING'.	No	false
vault	String	Name of the vault containing the volume. Must not be specified when creating a volume clone.	Optional	false

Copyright © 2022, 2026, Oracle and/or its affiliates.

Attribute	Type	Description	Settable on creation	Modifiable after creation
volume	String	ID of the existing volume where volume backup data will be restored. Used only for volumes restored from a backup in MD recovery mode. String size: 32 - 64 characters. Default: 'null'.	Optional	false
volumeBackup	String	Source volume backup ID. Used only for volumes restored from a backup. Default: 'null'.	Optional	false
volumeGroup	String	ID of the volume group associated with the volume	No	true
volumeSnapshot	String	Source volume snapshot ID. Used only for volumes cloned from a snapshot. Default: 'null'.	Optional	false
volumeSource	String	Source volume ID. Used only for volumes cloned directly from an existing volume. Default: 'null'.	Optional	false

Copyright © 2022, 2026, Oracle and/or its affiliates.