

# Oracle® SuperCluster Quorum Disk Manager

Administration Guide

**ORACLE®**

**Part No: E84120-02**  
September 2017



Oracle SuperCluster Quorum Disk Manager

**Part No: E84120-02**

Copyright © 2017, Oracle and/or its affiliates. All rights reserved.

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 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 installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. 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 and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon 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, Opteron, the AMD logo, and the AMD Opteron 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.

**Access to Oracle Support**

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

**Référence: E84120-02**

Copyright © 2017, Oracle et/ou ses affiliés. Tous droits réservés.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf stipulation expresse de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, accorder de licence, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est livré sous licence au Gouvernement des Etats-Unis, ou à quiconque qui aurait souscrit la licence de ce logiciel pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique :

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer un risque de dommages corporels. Si vous utilisez ce logiciel ou ce matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour des applications dangereuses.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée de The Open Group.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers, sauf mention contraire stipulée dans un contrat entre vous et Oracle. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation, sauf mention contraire stipulée dans un contrat entre vous et Oracle.

**Accès aux services de support Oracle**

Les clients Oracle qui ont souscrit un contrat de support ont accès au support électronique via My Oracle Support. Pour plus d'informations, visitez le site <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> ou le site <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> si vous êtes malentendant.

# Contents

---

|   |    |
|---|----|
| <b>Understanding Quorum Disk Manager</b> .....  | 7  |
| Overview of Quorum Disk Manager .....   | 7  |
| Software Requirements for Quorum Disk Manager .....   | 8  |
| ▼ Set Up Public Key Authentication for ZFS Storage Controllers .....                                    | 8  |
| <b>Adding Quorum Disks to Database Domains</b> .....  | 11 |
| New Deployments on Oracle SuperCluster 2.3 or Later .....   | 11 |
| Add Quorum Disks to Database Domains or Zones .....   | 11 |
| ▼ Add Quorum Disk to Database Domains (Global Zones) .....  | 12 |
| ▼ Add Quorum Disks to Database Zones (Non-Global Zones) .....   | 17 |
| ▼ Complete the Quorum Disk Configuration .....  | 21 |
| <b>Changing Existing Configurations</b> .....   | 27 |
| ▼ Removing the Quorum Disk Configuration .....  | 27 |
| Changing Elastic Configurations .....   | 28 |
| Adding a Database Domain .....  | 28 |
| Removing a Database Domain .....  | 28 |
| ▼ Adding an Oracle Exadata Storage Server and Expanding an Existing High<br>Redundancy Disk Group ..... | 29 |
| Removing an Oracle Exadata Storage Server .....   | 31 |
| <b>Quorum Disk Manager Reference</b> .....  | 33 |
| Syntax for the Quorum Disk Manager Utility .....  | 33 |
| quorumdiskmgr Objects .....   | 34 |
| Creating a Quorum Disk Configuration (--create --config) .....  | 34 |
| Syntax .....  | 34 |
| Parameters .....  | 34 |
| Example .....   | 35 |

|   |    |
|---|----|
| Creating a Target (--create --target) .....                   | 35 |
| Syntax .....  | 35 |
| Parameters .....  | 36 |
| Example .....   | 36 |
| Creating a Device (--create --device) .....                   | 36 |
| Syntax .....  | 36 |
| Parameters .....  | 37 |
| Example .....   | 37 |
| Listing Quorum Disk Configurations (--list --config) .....    | 37 |
| Syntax .....  | 37 |
| Sample Output .....   | 37 |
| Listing Targets (--list --target) .....                       | 37 |
| Syntax .....  | 38 |
| Parameters .....  | 38 |
| Sample Output .....   | 38 |
| Listing Devices (--list --device) .....                       | 38 |
| Syntax .....  | 39 |
| Parameters .....  | 39 |
| Sample Output .....   | 39 |
| Deleting Configurations (--delete --config) .....             | 39 |
| Syntax .....  | 40 |
| Deleting Targets (--delete --target) .....                    | 40 |
| Syntax .....  | 40 |
| Parameters .....  | 40 |
| Example .....   | 40 |
| Deleting Devices (--delete --device) .....                    | 40 |
| Syntax .....  | 41 |
| Parameters .....  | 41 |
| Example .....   | 41 |
| Changing Owner and Group Values (--alter --config) .....      | 41 |
| Syntax .....  | 42 |
| Parameters .....  | 42 |
| Example .....   | 42 |
| Changing the InfiniBand IP Addresses (--alter --target) ..... | 42 |
| Syntax .....  | 42 |
| Parameters .....  | 43 |
| Example .....   | 43 |

# Understanding Quorum Disk Manager

---

This chapter contains the following topics:

- [“Overview of Quorum Disk Manager” on page 7](#)
- [“Software Requirements for Quorum Disk Manager” on page 8](#)

## Overview of Quorum Disk Manager

The Quorum Disk Manager utility is a new feature in Oracle SuperCluster release 2.3.

Using this utility, you can create an iSCSI quorum disk on two of the database domains and store a voting file on those two quorum disks. These two additional voting files are used to meet the minimum requirement of five voting files for a high redundancy disk group. This feature is only applicable to Oracle SuperCluster systems that meet the following requirements:

- The Oracle SuperCluster rack has fewer than five storage servers
- The Oracle SuperCluster rack has at least two database domains
- You will have at least one high redundancy disk group

This feature allows for the voting files to be stored in a high redundancy disk group on Oracle SuperCluster racks with fewer than five storage servers by creating two extra failure groups.

Without this feature, on Oracle SuperCluster racks with fewer than five storage servers, voting files are stored in a normal redundancy disk group. This makes the Grid Infrastructure vulnerable to a double partner storage server failure that results in the loss of the voting file quorum, which then leads to a complete cluster and database outage. Refer to My Oracle Support note 1339373.1 to restart the clusterware and databases in this situation.

The iSCSI quorum disks are stored inside the ZFS storage appliance and are available for each domain, so they have the benefit of high availability. Systems with three Exadata storage servers require two quorum devices for high redundancy.

- In SuperCluster M8 and SuperCluster M7, the two ZFS storage controllers are set in an active-active configuration. On an active-active ZFS storage appliance system in

a correctly-configured SuperCluster M8 and SuperCluster M7 (one that has been set up using best practices), one iSCSI quorum disk device will be created on each ZFS storage controller. During controller failover, the quorum device associated with the failed controller will temporarily become unavailable.

- In SuperCluster T4-4, T5-8 and M6-32, the two ZFS storage controllers are set in an active-passive configuration. On an active-passive ZFS storage appliance system, both iSCSI quorum disk devices will be created in the same pool. During controller failover, both quorum devices will temporarily become unavailable.

The Quorum Disk Manager utility (`quorumdiskmgr`) is used to create and manage all the necessary components for implementing this feature, including the iSCSI configuration, the iSCSI targets, the iSCSI LUNs, and the iSCSI devices. For more information on the Quorum Disk Manager utility, see [“Quorum Disk Manager Reference” on page 33](#).

This feature is applicable to Oracle RAC clusters with Database I/O Domains, dedicated Database Domains, and zones on dedicated Database Domains or Database I/O Domains. For Oracle RAC clusters with zones on dedicated Database Domains or Database I/O Domains, the quorum disk devices are created and managed from its corresponding global zone.

## Software Requirements for Quorum Disk Manager

To use this feature, the following releases are required:

- Oracle SuperCluster ExaFamily software release 2.3 and above, which is available in the SuperCluster release 2.3 and in the January 2017 QFSDP
- Patch 23200778 for all Database homes
- Oracle Database 11.2.0.4 and above
- Oracle Clusterware/Grid version 12.1.0.2 BP161018 or higher
- OEDA January 2017 release and above

### ▼ Set Up Public Key Authentication for ZFS Storage Controllers

You must set up public key authentication on the ZFS storage appliance for the ZFS storage controllers for the first two global domains before using the `quorumdiskmgr` tool. You do not have to set up public key authentication from non-global zones, but it is required for its corresponding global zones.



### 1. Log in to the first ZFS storage controller using the CLI.

To log in remotely using the CLI, use an ssh client. If you have not configured other users to administer the appliance, you must log in as `root`. When you log in, the CLI prompt consists of the hostname, followed by a colon, followed by a greater-than sign:

```
% ssh root@zfs-controller1-hostname
Password:
Last login: Mon Oct 13 15:43:05 2009 from kiowa.sf.fishpo
zfs-controller1-hostname:>
```

where `zfs-controller1-hostname` is the host name for the first ZFS storage controller.

### 2. Set up the public key authentication on the first ZFS storage controller for the first global zone:

```
zfs-controller1-hostname:> configuration preferences keys
zfs-controller1-hostname:configuration preferences keys> create
zfs-controller1-hostname:configuration preferences key (uncommitted)> set type=DSA
zfs-controller1-hostname:configuration preferences key (uncommitted)> set key="DSA-key-text"
    key = DSA-key-text == (uncommitted)
zfs-controller1-hostname:configuration preferences key (uncommitted)> set comment="fw-log1"
    comment = fw-log1 (uncommitted)
zfs-controller1-hostname:configuration preferences key (uncommitted)> commit
zfs-controller1-hostname:configuration preferences keys> show
Keys:
```

| NAME    | MODIFIED            | TYPE | COMMENT |
|---------|---------------------|------|---------|
| key-000 | 10/12/2009 10:54:58 | DSA  | fw-log1 |

The `DSA-key-text` is the key text itself (usually hundreds of characters), without spaces.

3. Repeat [Step 2](#) for the second global zone.
4. Log out of the first ZFS storage controller:

```
zfs-controller1-hostname:> exit
```

---

**Note** - You do not have to repeat this procedure for the second ZFS storage controller. When you set up the public key authentication on the first ZFS storage controller, it appears when you log in to the second ZFS storage controller.

---

5. Log in to the first global zone in your system, if you are not logged in already.
6. Verify that you set up the public key authentication correctly:

```
% ssh root@zfs-controller1-hostname hostname
```

where *zfs-controller1-hostname* is the host name for the first ZFS storage controller.

You should be able to log in to the ZFS storage controller successfully and see the CLI prompt.

7. **Repeat [Step 5](#) and [Step 6](#) for the second global zone.**

# Adding Quorum Disks to Database Domains

---

These instructions describe how to add quorum disks to an existing SuperCluster system that has never had quorum disks set up previously:

- [“New Deployments on Oracle SuperCluster 2.3 or Later ” on page 11](#)
- [“Add Quorum Disks to Database Domains or Zones” on page 11](#)

If you have a SuperCluster system where you have quorum disks set up already but you want to make changes (such as removing quorum disk functionality or changing configurations where quorum disk functionality is in place), see [“Changing Existing Configurations” on page 27](#).

## New Deployments on Oracle SuperCluster 2.3 or Later

If you purchased a new SuperCluster system running release 2.3 or later, your Oracle installer will implement this feature by default when all of the following requirements are satisfied:

- The system has at least two database domains and fewer than five storage servers.
- You are running OEDA release January 2017 or later.
- You meet the software requirements listed in [“Software Requirements for Quorum Disk Manager” on page 8](#).
- You will have at least one high redundancy disk group.

If the system has three storage servers in place, then two quorum disks will be created on the first two database domains of the cluster selected by your Oracle installer.

If the system has four storage servers in place, then one quorum disk will be created on the first database domain selected by your Oracle installer.

## Add Quorum Disks to Database Domains or Zones

If you purchased a new SuperCluster system running release 2.3 or later, your Oracle installer should have already implemented the quorum disk feature as part of the initial installation. You

do not have to follow these instructions to manually add quorum disks to database domains in that case.

However, if you have an existing SuperCluster system that was set up with a software version earlier than version 2.3 and your SuperCluster system meets the following requirements:

- The SuperCluster system has fewer than five storage servers
- The SuperCluster system has at least two database domains
- You will have at least one high redundancy disk group

Then you can take advantage of the quorum disk feature by upgrading your software to release 2.3 or later and then implementing this feature manually using the instructions in this chapter. Your SuperCluster system must meet the requirements provided in [“Software Requirements for Quorum Disk Manager” on page 8](#) and must be running Oracle Database 11.2.0.4 or above to use the quorum disk feature.

The procedure to add quorum disks varies, depending on the type of database domain or zone that you are adding the quorum disks to:

- If you are adding quorum disks to dedicated Database Domains or Database I/O Domains (global zones), then follow the instructions in [“Add Quorum Disk to Database Domains \(Global Zones\)” on page 12](#).
- If you are adding quorum disks to zones in dedicated Database Domains or Database I/O Domains (non-global zones), then follow the instructions in [“Add Quorum Disks to Database Zones \(Non-Global Zones\)” on page 17](#).

Once you have completed the procedures in one of those two sections, you will complete the quorum disk configuration process using the instructions in [“Complete the Quorum Disk Configuration” on page 21](#).

## ▼ Add Quorum Disk to Database Domains (Global Zones)

The example in this section creates quorum disks for a small rack configuration with two database domains: db01 and db02. On both db01 and db02, there are two InfiniBand interfaces: bondib0 and bondib1.

You can find the network interfaces to be used for communication with the iSCSI devices using the following command:

```
$ oifcfg getif | grep cluster_interconnect | awk '{print $1}'
```

You can find the IP address of each interface using the following command:

```
ipadm show-addr interface_name
```

The InfiniBand IP addresses for this example are as follows:

- On the first database domain (*db01*):
  - Network interface: bondib0, IP address: 192.168.10.45
  - Network interface: bondib1, IP address: 192.168.10.46
- On the second database domain (*db02*):
  - Network interface: bondib0, IP address: 192.168.10.47
  - Network interface: bondib1, IP address: 192.168.10.48

The ASM disk group to which the quorum disks will be added is DATA1. The ASM owner is grid, and the user group is dba.

Initially, the voting files reside on a normal redundancy disk group RECOC1:

```
$ $GRID_HOME/bin/crsctl query css votedisk
## STATE      File Universal Id                File Name                     Disk group
--  -
1.  ONLINE     21f5507a28934f77bf3b7ecf88b26c47 (o/192.168.76.187;192.168.76.188/
RECOC1_CD_00_celadm12) [RECOC1]
2.  ONLINE     387f71ee81f14f38bfbd0693451e328 (o/192.168.76.189;192.168.76.190/
RECOC1_CD_00_celadm13) [RECOC1]
3.  ONLINE     6f7fab62e6054fb8bf167108cbbd2f64 (o/192.168.76.191;192.168.76.192/
RECOC1_CD_00_celadm14) [RECOC1]
Located 3 voting disk(s).
```

---

**Note** - For several of the steps in this procedure, you must perform a process first on the first database domain, wait for that process to finish on the first database domain, then perform the same process on the second database domain, then wait for that process to finish on the second database domain before moving on to the next step.

---

To create the quorum disks:

1. **Verify that your SuperCluster system fulfills the proper software requirements and that you have set up the public key authentication for the ZFS storage controllers, if necessary.**  
See [“Software Requirements for Quorum Disk Manager” on page 8](#) and [“Set Up Public Key Authentication for ZFS Storage Controllers” on page 8](#) for more information.
2. **Log into the first database domain (*db01*) and the second database domain (*db02*) as a root user.**

In this procedure, it is assumed that you have two separate terminal windows where you are able to run steps on the first database domain (*db01*) in one terminal window and steps on the second database domain (*db02*) in the second terminal window, and you are able to see when a process has completed in the *db01* terminal window before running that same step on the *db02* terminal window.

---

**Note** - The following commands use `--owner=oracle` and `--group=dba` as examples. Replace the owner and group information in the following commands with information that is valid for your setup.

---

**3. On the first database domain, create quorum disk configurations:**

```
db01# /opt/oracle.SupportTools/quorumdiskmgr --create --config --owner=oracle --group=dba --network-iface-list="bondib0, bondib1"
```

**4. Wait until the process has completed on the first database domain, then, on the second database domain, run the same command to create quorum disk configurations:**

```
db02# /opt/oracle.SupportTools/quorumdiskmgr --create --config --owner=oracle --group=dba --network-iface-list="bondib0, bondib1"
```

**5. Verify that the configurations have been successfully created on the first and second database domains.**

**a. Enter the following command on *db01*:**

```
db01# /opt/oracle.SupportTools/quorumdiskmgr --list --config
```

**b. Enter the same command on *db02*:**

```
db02# /opt/oracle.SupportTools/quorumdiskmgr --list --config
```

Output similar to the following appears for both *db01* and *db02*:

```
Owner: oracle
Group: dba
ifaces: bondib0,bondib1
```

```
[SUCCESS ] Successfully list quorumdisk config file
```

**6. On the first database domain, create a target for ASM disk group *DATA1* and make the target visible to both the first and second database domains.**

```
db01# /opt/oracle.SupportTools/quorumdiskmgr --
create --target --asm-disk-group=datac1 --visible-
to="db01_bondib0_IPaddress, db01_bondib1_IPaddress, db02_bondib0_IPaddress, db02_bondib1_IPaddress"
```

For example:

```
db01# /opt/oracle.SupportTools/quorumdiskmgr --create --target --asm-disk-group=datac1
--visible-to="192.168.10.45, 192.168.10.46, 192.168.10.47, 192.168.10.48"
```

7. **Wait until the process has completed on the first database domain, then, on the second database domain, run the same command to create a target for ASM disk group DATAC1 and make the target visible to both the first and second database domains.**

```
db02# /opt/oracle.SupportTools/quorumdiskmgr --
create --target --asm-disk-group=datac1 --visible-
to="db01_bondib0_IPaddress, db01_bondib1_IPaddress, db02_bondib0_IPaddress, db02_bondib1_IPaddress"
```

For example:

```
db02# /opt/oracle.SupportTools/quorumdiskmgr --create --target --asm-disk-group=datac1
--visible-to="192.168.10.45, 192.168.10.46, 192.168.10.47, 192.168.10.48"
```

8. **Verify the target has been successfully created on both database domains.**

- a. **Enter the following command on the first database domain:**

```
db01# /opt/oracle.SupportTools/quorumdiskmgr --list --target
```

On the first database domain, output similar to the following appears:

```
Name: iqn.1986-03.com.sun:02:080a9aa9-3c1a-cd96-b80e-a812ce744211
Size: 128 MB
Lun: QD_datac1_db01
Host name: DB01
ASM disk group name: DATAC1
Visible to: 192.168.10.45, 192.168.10.46, 192.168.10.47, 192.168.10.48
Discovered by:
```

```
[SUCCESS ] List lun(s)
```

- b. **Enter the same command on the second database domain:**

```
db02# /opt/oracle.SupportTools/quorumdiskmgr --list --target
```

On the second database domain, output similar to the following appears:

```
Name: iqn.1986-03.com.sun:02:080a9aa9-3c1a-cd96-b80e-a812ce744111
Size: 128 MB
Lun: QD_datac1_db02
Host name: DB02
ASM disk group name: DATAC1
Visible to: 192.168.10.45, 192.168.10.46, 192.168.10.47, 192.168.10.48
Discovered by:
```

```
[SUCCESS ] List lun(s)
```

- 9. On the first database domain, create devices from targets on both database domains.**

```
db01# /opt/oracle.SupportTools/quorumdiskmgr --create --device --target-ip-
list="db01_bondib0_IPAddress, db01_bondib1_IPAddress, db02_bondib0_IPAddress, db02_bondib1_IPAddress"
```

For example:

```
db01# /opt/oracle.SupportTools/quorumdiskmgr --create --device --target-ip-
list="192.168.10.45, 192.168.10.46, 192.168.10.47, 192.168.10.48"
```

- 10. Wait until the process has completed on the first database domain, then, on the second database domain, run the same command to create devices from targets on both database domains.**

```
db02# /opt/oracle.SupportTools/quorumdiskmgr --create --device --target-ip-
list="db01_bondib0_IPAddress, db01_bondib1_IPAddress, db02_bondib0_IPAddress, db02_bondib1_IPAddress"
```

For example:

```
db02# /opt/oracle.SupportTools/quorumdiskmgr --create --device --target-ip-
list="192.168.10.45, 192.168.10.46, 192.168.10.47, 192.168.10.48"
```

- 11. Verify the devices have been successfully created on both database domains.**

- a. Enter the following command on the first database domain:**

```
db01# /opt/oracle.SupportTools/quorumdiskmgr --list --device
```

- b. Enter the same command on the second database domain:**

```
db02# /opt/oracle.SupportTools/quorumdiskmgr --list --device
```

Output similar to the following appears on both db01 and db02:

```
Device path: /dev/exadata_quorum/QD_DATAC1_DB01
Size: 128 MB
```



```

Host name: DB01
ASM disk group name: DATA1

Device path: /dev/exadata_quorum/QD_DATA1_DB02
Size: 128 MB
Host name: DB02
ASM disk group name: DATA1

[SUCCESS ] List devices

```

**12. Complete the quorum disk configuration process.**

Go to [“Complete the Quorum Disk Configuration” on page 21](#) for those instructions.

## ▼ Add Quorum Disks to Database Zones (Non-Global Zones)

The example in this section creates quorum disks for a small rack configuration with two database zones: zone01 and zone02.

For non-global zones (zones within a dedicated Database Domain or Database I/O Domain), additional parameters are needed when you run the `quorumdiskmgr` tool, even though the quorum devices are managed from the global zone. Specify the non-global zone that will be affected to the `quorumdiskmgr` tool using the parameter `--zone-name`. Also, you will run the steps in the procedure in this section as the root user on the global zone (the Database Domain hosting the zones).

To create the quorum disks:

- 1. Verify that your SuperCluster system fulfills the proper software requirements and that you have set up the public key authentication for the ZFS storage controllers, if necessary.**

See [“Software Requirements for Quorum Disk Manager” on page 8](#) and [“Set Up Public Key Authentication for ZFS Storage Controllers” on page 8](#) for more information.

- 2. Get the necessary IP address information for each non-global zone.**

- On the first non-global zone, enter the following command:

```
zone01# ipadm show-addr bondib0
```

Output similar to the following appears - note the IP address provided in the ADDR column (in this case, 192.168.10.1):

| ADDROBJ    | TYPE   | STATE | ADDR            |
|------------|--------|-------|-----------------|
| bondib0/v4 | static | ok    | 192.168.10.1/22 |

- On the second non-global zone, enter the following command:

```
zone02# ipadm show-addr bondib0
```

Output similar to the following appears - note the IP address provided in the ADDR column (in this case, 192.168.10.2):

| ADDROBJ    | TYPE   | STATE | ADDR            |
|------------|--------|-------|-----------------|
| bondib0/v4 | static | ok    | 192.168.10.2/22 |

3. **Log into the global zone (the Database Domain hosting the zone) as a root user.**
4. **On the global zone, create the quorum disk configuration for the *first* non-global zone:**

```
/opt/oracle.SupportTools/quorumdiskmgr --create --config --owner=oracle --group=grid --network-iface-list="bondib0" --zone-name=zone01
```

Where *zone01* is the name of the first non-global zone.

5. **On the global zone, create the quorum disk configuration for the *second* non-global zone:**

```
/opt/oracle.SupportTools/quorumdiskmgr --create --config --owner=oracle --group=grid --network-iface-list="bondib0" --zone-name=zone02
```

Where *zone02* is the name of the second non-global zone.

6. **Verify that the configurations have been successfully created on the first and second non-global zones.**

- a. **Enter the following command on the non-global zone *zone01*:**

```
/opt/oracle.SupportTools/quorumdiskmgr --list --config --zone-name=zone01
```

- b. **Enter the same command on the non-global zone *zone02*:**

```
/opt/oracle.SupportTools/quorumdiskmgr --list --config --zone-name=zone02
```

Output similar to the following appears for both non-global zones *zone01* and *zone02*:

```
Owner: oracle
Group: dba
ifaces: bondib0,bondib1
```

[SUCCESS ] Successfully list quorumdisk config file

**7. On the global zone, create a target for the *first* non-global zone:**

```
/opt/oracle.SupportTools/quorumdiskmgr --create --target --asm-disk-group=datac1 --
visible-to="zone01_IP_address, zone02_IP_address" --zone-name=zone01
```

For example:

```
/opt/oracle.SupportTools/quorumdiskmgr --create --target --asm-disk-group=datac1 --
visible-to="192.168.10.1, 192.168.10.2" --zone-name=zone01
```

**8. On the global zone, create a target for the *second* non-global zone:**

```
/opt/oracle.SupportTools/quorumdiskmgr --create --target --asm-disk-group=datac1 --
visible-to="zone01_IP_address, zone02_IP_address" --zone-name=zone02
```

For example:

```
/opt/oracle.SupportTools/quorumdiskmgr --create --target --asm-disk-group=datac1 --
visible-to="192.168.10.1, 192.168.10.2" --zone-name=zone02
```

**9. Verify the target has been successfully created on the first and second non-global zones.**

**a. Enter the following command on the first non-global zone (*zone01*):**

```
/opt/oracle.SupportTools/quorumdiskmgr --list --target --zone-name=zone01
```

On the first non-global zone, output similar to the following appears:

```
Name: iqn.1986-03.com.sun:02:080a9aa9-3c1a-cd96-b80e-a812ce744211
Size: 128 MB
Lun: QD_datac1_zone01
Host name: ZONE01
ASM disk group name: DATA1
Visible to: 192.168.10.1, 192.168.10.2
Discovered by:
```

[SUCCESS ] List lun(s)

**b. Enter the same command on the second non-global zone (*zone02*):**

```
/opt/oracle.SupportTools/quorumdiskmgr --list --target --zone-name=zone02
```

On the second non-global zone, output similar to the following appears:

```
Name: iqn.1986-03.com.sun:02:080a9aa9-3c1a-cd96-b80e-a812ce744111
```

```
Size: 128 MB
Lun: QD_datac1_zone02
Host name: ZONE02
ASM disk group name: DATA1
Visible to: 192.168.10.1, 192.168.10.2
Discovered by:
```

```
[SUCCESS ] List lun(s)
```

**10. On the global zone, create devices for the *first* non-global zone:**

```
/opt/oracle.SupportTools/quorumdiskmgr --create --device --target-ip-
list="zone01_IP_address, zone02_IP_address" --zone-name=zone01
```

For example:

```
/opt/oracle.SupportTools/quorumdiskmgr --create --device --target-ip-list="192.168.10.1,
192.168.10.2" --zone-name=zone01
```

**11. On the global zone, create devices for the *second* non-global zone:**

```
/opt/oracle.SupportTools/quorumdiskmgr --create --device --target-ip-
list="zone01_IP_address, zone02_IP_address" --zone-name=zone02
```

For example:

```
/opt/oracle.SupportTools/quorumdiskmgr --create --device --target-ip-list="192.168.10.1,
192.168.10.2" --zone-name=zone02
```

**12. Verify the devices have been successfully created on the first and second non-global zones.**

**a. Enter the following command on the first non-global zone (*zone01*):**

```
/opt/oracle.SupportTools/quorumdiskmgr --list --device --zone-name=zone01
```

**b. Enter the same command on the second non-global zone (*zone02*):**

```
/opt/oracle.SupportTools/quorumdiskmgr --list --device --zone-name=zone02
```

Output similar to the following appears on both zone01 and zone02:

```
Device path: /dev/exadata_quorum/QD_DATA1_ZONE01
Size: 128 MB
Host name: ZONE01
ASM disk group name: DATA1

Device path: /dev/exadata_quorum/QD_DATA1_ZONE02
```

```
Size: 128 MB
Host name: ZONE02
ASM disk group name: DATA1
```

```
[SUCCESS ] List devices
```

**13. Complete the quorum disk configuration process.**

Go to [“Complete the Quorum Disk Configuration”](#) on page 21 for those instructions.

## ▼ Complete the Quorum Disk Configuration

Once you have added the quorum disks to the database zones or the database domains, follow these instructions to complete the quorum disk configuration.

**1. Determine if the database domain where the quorum disks are being created is a global zone or a non-global zone.**

- If the database domain is a global zone, then continue running the remaining steps in this procedure in this global zone.
- If the database domain is a non-global zone (for example, a zone within a dedicated Database Domain), then run the remaining steps in this procedure (steps [Step 2](#) through [Step 12](#)) in the *non-global* zone that is part of the database environment with high redundancy.

**2. Log in to the global zone or the non-global zone as an oracle user.**

**3. Set the ORACLE\_HOME and ORACLE\_SID environment variables:**

```
export ORACLE_HOME=$GRID_HOME
export ORACLE_SID=ASM_instance_id
```

**4. Alter the asm\_diskstring initialization parameter and add /dev/exadata\_quorum/\* to the existing string.**

For example:

```
SQL> alter system set asm_diskstring='o/*/DATA1_*, 'o/*/RECO1_*, '/dev/exadata_quorum/
*' scope=both sid='*';
```

**5. Verify the two quorum disk devices have been automatically discovered by ASM.**

```
SQL> set linesize 200
SQL> col path format a50
SQL> select inst_id, label, path, mode_status, header_status from gv$asm_disk where path
like '/dev/exadata_quorum/%';
```

Output similar to the following appears:

| INST_ID | LABEL           | PATH                                | MODE_STATUS | HEADER_STATUS |
|---------|-----------------|-------------------------------------|-------------|---------------|
| 1       | QD_DATAAC1_DB01 | /dev/exadata_quorum/QD_DATAAC1_DB01 | ONLINE      | MEMBER        |
| 1       | QD_DATAAC1_DB02 | /dev/exadata_quorum/QD_DATAAC1_DB02 | ONLINE      | MEMBER        |
| 2       | QD_DATAAC1_DB01 | /dev/exadata_quorum/QD_DATAAC1_DB01 | ONLINE      | MEMBER        |
| 2       | QD_DATAAC1_DB02 | /dev/exadata_quorum/QD_DATAAC1_DB02 | ONLINE      | MEMBER        |

**6. Add the two quorum disk devices to a high redundancy ASM disk group.**

Verify that there is at least one high redundancy disk group in place by listing the disk groups and their information:

```
$GI_HOME/bin/asmcmd lsdg
```

- If there is no high redundancy disk group, you have two options to create new diskgroup with high redundancy, both of which are described in My Oracle Support note 438580.1:
  - Drop the existing diskgroup after backing up data and create a new diskgroup with high redundancy.
  - Create a new high redundancy disk group and move the existing data to newly created diskgroup. For example:

```
SQL> create diskgroup DATAAC1 high redundancy quorum failgroup db01 disk '/dev/exadata_quorum/QD_DATAAC1_DB01' quorum failgroup db02 disk '/dev/exadata_quorum/QD_DATAAC1_DB02' ...
```

---

**Note** - The "..." at the end of the command above signifies the intentional omission of the ASM disks. Refer to [https://docs.oracle.com/cd/B28359\\_01/server.111/b31107/asmwithem.htm#OSTMG24000](https://docs.oracle.com/cd/B28359_01/server.111/b31107/asmwithem.htm#OSTMG24000) for more information on creating disk groups.

---

In both cases, add the two new quorum disks when you create the high redundancy ASM group.

- If a high redundancy disk group already exists, add the two new quorum disks. For example:

```
SQL> alter diskgroup dataac1 add quorum failgroup db01 disk '/dev/exadata_quorum/QD_DATAAC1_DB02' quorum failgroup db02 disk '/dev/exadata_quorum/QD_DATAAC1_DB01';
```

**7. Query the voting disks to ensure that they are in the desired disk group.**

```
$ $GRID_HOME/bin/crsctl query css votedisk
```

Output similar to the following appears:

```

## STATE File Universal Id File Name
Disk group
-- -----
-----
1. ONLINE ca2f1b57873f4ff4bf1dfb78824f2912 (o/192.168.10.42/DATAC1_CD_09_celadm12)
[DATAC1]
2. ONLINE a8c3609a3dd44f53bf17c89429c6ebe6 (o/192.168.10.43/DATAC1_CD_09_celadm13)
[DATAC1]
3. ONLINE cafb7e95a5be4f00bf10bc094469cad9 (o/192.168.10.44/DATAC1_CD_09_celadm14)
[DATAC1]
Located 3 voting disk(s).

```

8. Relocate the existing voting files from the normal redundancy disk group to the high redundancy disk group.

```
$ $GRID_HOME/bin/crsctl replace votedisk +DATAC1
```

9. Verify that the voting disks have been successfully relocated to the high redundancy disk group and that five voting files exist.

```
$ $GRID_HOME/bin/crsctl query css votedisk
```

The output should show three voting disks from the storage servers and two voting disks from the database domains:

```

## STATE File Universal Id File Name Disk
group
-- -----
-----
1. ONLINE ca2f1b57873f4ff4bf1dfb78824f2912 (o/192.168.10.42/DATAC1_CD_09_celadm12)
[DATAC1]
2. ONLINE a8c3609a3dd44f53bf17c89429c6ebe6 (o/192.168.10.43/DATAC1_CD_09_celadm13)
[DATAC1]
3. ONLINE cafb7e95a5be4f00bf10bc094469cad9 (o/192.168.10.44/DATAC1_CD_09_celadm14)
[DATAC1]
4. ONLINE 4dca8fb7bd594f6ebf8321ac23e53434 (/dev/exadata_quorum/QD_DATAC1_DB01)
[DATAC1]
5. ONLINE 4948b73db0514f47bf94ee53b98fdb51 (/dev/exadata_quorum/QD_DATAC1_DB02)
[DATAC1]
Located 5 voting disk(s).

```

10. Relocate the OCR files from the normal redundancy disk group to the high redundancy disk group.

```

$GI_HOME/bin/ocrcheck
$GI_HOME/bin/ocrconfig -add +DATAC1
$GI_HOME/bin/ocrconfig -delete +RECOC1

```

**11. Move the ASM password file and the ASM spfile to the high redundancy disk group.**

**a. Move the ASM password file:**

**i. Get the source ASM password file location.**

```
$ asmcmd pwget --asm
```

**ii. Move the ASM password file to the high redundancy disk group.**

```
$ asmcmd pwmove --asm full_path_of_source_file full_path_of_destination_file
```

For example:

```
asmcmd pwmove --asm +recoc1/ASM/PASSWORD/pwdasm.256.898960531 +datac1/  
asmpwdfile
```

**b. Move the ASM spfile.**

**i. Get the ASM spfile in use:**

```
$ asmcmd spget
```

**ii. Copy the ASM spfile to the high redundancy disk group.**

```
$ asmcmd spcopy full_path_of_source_file full_path_of_destination_file
```

**iii. Modify the Grid Infrastructure configuration to use the relocated spfile on the next restart.**

```
$ asmcmd spset full_path_of_destination_file
```

**iv. Determine if you can shut down the cluster at this time.**

- If you can shut down the cluster at this time, restart the Grid Infrastructure:

```
# $GI_HOME/bin/crsctl stop crs  
# $GI_HOME/bin/crsctl start crs
```

- If can not shut down the cluster at this time, repeat step [Step 11b](#) every time an initialization parameter modification to the ASM spfile is required, until the point where you can shut down the cluster and restart the Grid Infrastructure using the steps above.

**12. Relocate the MGMTDB to the high redundancy disk group.**



Move the mgmtdb (if running) to the high redundancy disk group using My Oracle Support note 1589394.1.

Configure the mgmtdb to not use hugepages:

```
export ORACLE_SID=-MGMTDB
export ORACLE_HOME=$GRID_HOME
sqlplus "sys as sysdba"
SQL> alter system set use_large_pages=false scope=spfile sid='*';
```



## Changing Existing Configurations

---

This chapter describes how to make changes to a SuperCluster system that has the quorum disk feature already set up:

- [“Removing the Quorum Disk Configuration” on page 27](#)
- [“Changing Elastic Configurations” on page 28](#)

### ▼ Removing the Quorum Disk Configuration

This section provides instructions for removing the quorum disk configuration, if you implemented the feature and you decide you would like to remove it for any reason.

To remove the quorum disk configuration:

1. **Ensure there is at least one normal redundancy disk group in place by listing the disk groups and their information:**

```
$GI_HOME/bin/asmcmd lsdg
```

If there is not at least one normal redundancy disk group, create one. Refer to My Oracle Support note 438580.1 for details on how to change disk group redundancy.

2. **Relocate the voting files to a normal redundancy disk group:**

```
$GI_HOME/bin/crsctl replace votedisk +normal_redundancy_diskgroup
```

3. **Drop the quorum disks from ASM.**

Run the following command for each quorum disk:

```
SQL> alter diskgroup diskgroup_name drop quorum disk quorum_disk_name force;
```

Wait for the rebalance operation to complete. You can tell it is complete when v \$asm\_operation returns no rows for the disk group.

4. **Delete the quorum devices.**

Run the following command from each database node that has quorum disks in place:

```
/opt/oracle.SupportTools/quorumdiskmgr --delete --device [--asm-disk-group  
asm_disk_group] [--host-name host_name]
```

**5. Delete the targets.**

Run the following command from each database domain that has quorum disks in place:

```
/opt/oracle.SupportTools/quorumdiskmgr --delete --target [--asm-disk-group  
asm_disk_group]
```

**6. Delete the configuration.**

Run the following command from each database domain that has quorum disks in place:

```
/opt/oracle.SupportTools/quorumdiskmgr --delete --config
```

## Changing Elastic Configurations

- [“Adding a Database Domain” on page 28](#)
- [“Removing a Database Domain ” on page 28](#)
- [“Adding an Oracle Exadata Storage Server and Expanding an Existing High Redundancy Disk Group ” on page 29](#)
- [“Removing an Oracle Exadata Storage Server ” on page 31](#)

## Adding a Database Domain

If the existing RAC cluster has fewer than two database domains and fewer than five storage servers, and the voting files are not stored in a high redundancy disk group, then you should add quorum disks to the database domains and relocate the voting files to a high redundancy disk group. See [“Adding Quorum Disks to Database Domains” on page 11](#) for those instructions. Note that the requirements listed in [“Software Requirements for Quorum Disk Manager” on page 8](#) must be met.

## Removing a Database Domain

If you are removing a database domain that is not hosting a quorum disk, then no action is required.

If you are removing a database domain that is hosting a quorum disk containing a voting file, and there are five or more storage servers in the RAC cluster, then no action is required. When you remove the database node hosting the voting file, the dropped voting file will be automatically relocated to a storage server that did not contain a voting file as part of the voting file refresh.

If you are removing a database domain that is hosting a quorum disk containing a voting file and there are fewer than five storage servers in the RAC cluster, then you must create a quorum disk on a different database domain before you can remove the database node.

## ▼ Adding an Oracle Exadata Storage Server and Expanding an Existing High Redundancy Disk Group

When you add a storage server that uses quorum disks, you should relocate a voting file from a database node to the newly added storage server.

### 1. Add the Exadata storage server.

Refer to the following document for those instructions:

[http://docs.oracle.com/cd/E50790\\_01/doc/doc.121/e51951/storage.htm#DBMMN22922](http://docs.oracle.com/cd/E50790_01/doc/doc.121/e51951/storage.htm#DBMMN22922)

In the example below, the new storage server that is being added is called "celadm04".

### 2. After the storage server is added, verify the new fail group from `v$asm_disk`.

```
SQL> select distinct failgroup from v$asm_disk;
FAILGROUP
-----
ADM01
ADM02
CELADM01
CELADM02
CELADM03
CELADM04
```

### 3. Verify at least one database node has a quorum disk containing a voting file.

```
$ crsctl query css votedisk
## STATE File Universal Id File Name Disk group
--  -----
1. ONLINE 834ee5a8f5054f12bf47210c51ecb8f4 o/192.168.12.125;192.168.12.126/
DATA5_CD_00_celadm01) [DATA5]
```

```

2. ONLINE   f4af2213d9964f0bbfa30b2ba711b475 (o/192.168.12.127;192.168.12.128/
   DATA5_CD_00_celadm02) [DATA5]
3. ONLINE   ed61778df2964f37bf1d53ea03cd7173 (o/192.168.12.129;192.168.12.130/
   DATA5_CD_00_celadm03) [DATA5]
4. ONLINE   bfe1c3aa91334f16bf78ee7d33ad77e0 (/dev/exadata_quorum/QD_DATA5_ADM01)
   [DATA5]
5. ONLINE   a3a56e7145694f75bf21751520b226ef (/dev/exadata_quorum/QD_DATA5_ADM02)
   [DATA5]
Located 5 voting disk(s).

```

The example above shows there are two quorum disks with voting files on two database nodes.

#### 4. Drop one of the quorum disks.

```
SQL> alter diskgroup datac5 drop quorum disk QD_DATA5_ADM01;
```

The voting file on the dropped quorum disk will be relocated automatically to the newly added storage server by the Grid Infrastructure as part of the voting file refresh. You can verify this as follows:

```

$ crsctl query css votedisk
## STATE      File Universal Id                File Name Disk group
--  -
1. ONLINE    834ee5a8f5054f12bf47210c51ecb8f4 (o/192.168.12.125;192.168.12.126/
   DATA5_CD_00_celadm01) [DATA5]
2. ONLINE    f4af2213d9964f0bbfa30b2ba711b475 (o/192.168.12.127;192.168.12.128/
   DATA5_CD_00_celadm02) [DATA5]
3. ONLINE    ed61778df2964f37bf1d53ea03cd7173 (o/192.168.12.129;192.168.12.130/
   DATA5_CD_00_celadm03) [DATA5]
4. ONLINE    a3a56e7145694f75bf21751520b226ef (/dev/exadata_quorum/QD_DATA5_ADM02)
   [DATA5]
5. ONLINE    ab5aefd60cf84fe9bfff6541b16e33787 (o/192.168.12.131;192.168.12.132/
   DATA5_CD_00_celadm04) [DATA5]

```

#### 5. Delete the quorum device.

Run the following command from the database node:

```
/opt/oracle.SupportTools/quorumdiskmgr --delete --device [--asm-disk-group asm_disk_group] [--host-name host_name]
```

#### 6. Delete the target.

Run the following command from the database domain:

```
/opt/oracle.SupportTools/quorumdiskmgr --delete --target [--asm-disk-group asm_disk_group]
```

#### 7. Delete the configuration.

Run the following command from the database domain:

```
/opt/oracle.SupportTools/quotumdiskmgr --delete -config
```

## Removing an Oracle Exadata Storage Server

If removing a storage server results in the number of storage servers being used by the Oracle RAC cluster dropping to fewer than five, and the voting files reside in a high redundancy disk group, then you should add quorum disks to the database nodes, if they are not in place already. See [“Adding Quorum Disks to Database Domains” on page 11](#).

Before removing the storage server, add the quorum disks so that five copies of the voting files are available immediately after removing the storage server.





## Quorum Disk Manager Reference

---

The quorum disk manager utility (`quorumdiskmgr`) runs on each database domain to enable you to create and manage iSCSI quorum disks on database servers. Use `quorumdiskmgr` to create, list, alter, and delete iSCSI quorum disks on database domains. The utility is installed on database domains when they are shipped.

This reference section contains the following topics:

- [“Syntax for the Quorum Disk Manager Utility” on page 33](#)
- [“quorumdiskmgr Objects” on page 34](#)
- [“Creating a Quorum Disk Configuration \(`--create --config`\)” on page 34](#)
- [“Creating a Target \(`--create --target`\)” on page 35](#)
- [“Creating a Device \(`--create --device`\)” on page 36](#)
- [“Listing Quorum Disk Configurations \(`--list --config`\)” on page 37](#)
- [“Listing Targets \(`--list --target`\)” on page 37](#)
- [“Listing Devices \(`--list --device`\)” on page 38](#)
- [“Deleting Configurations \(`--delete --config`\)” on page 39](#)
- [“Deleting Targets \(`--delete --target`\)” on page 40](#)
- [“Deleting Devices \(`--delete --device`\)” on page 40](#)
- [“Changing Owner and Group Values \(`--alter --config`\)” on page 41](#)
- [“Changing the InfiniBand IP Addresses \(`--alter --target`\)” on page 42](#)

### Syntax for the Quorum Disk Manager Utility

The quorum disk manager utility is a command-line tool. It has the following syntax:

```
quorumdiskmgr --verb --object [--options]
```

`verb` is an action performed on an object. It is one of: `alter`, `create`, `delete`, `list`.

`object` is an object on which the command performs an action.

options extend the use of a command combination to include additional parameters for the command.

When using the quorumdiskmgr utility, the following rules apply:

- Verbs, objects, and options are case-sensitive except where explicitly stated.
- Use the double quote character around the value of an option that includes spaces or punctuation.

## quorumdiskmgr Objects

| Object | Description  |
|--------|--|
| config | The quorum disk configurations include the owner and group of the ASM instance to which the iSCSI quorum disks will be added, and the list of network interfaces through which local and remote iSCSI quorum disks will be discovered. |
| target | A target is an endpoint on each database server that waits for an iSCSI initiator to establish a session and provides required IO data transfer.   |
| device | A device is an iSCSI device created by logging into a local or remote target.  |

## Creating a Quorum Disk Configuration (--create --config)

The `--create --config` action creates a quorum disk configuration. The configuration must be created before any targets or devices can be created.

### Syntax

```
quorumdiskmgr --create --config [--owner owner --group group] --network-iface-list
network-iface-list [--zone-name zone]
```

### Parameters

The following table lists the parameters for the `--create --config` action:

| Parameter          | Description  |
|--------------------|--|
| owner              | Specifies the owner of the ASM instance to which the iSCSI quorum disks will be added. This is an optional parameter. The default value is <code>grid</code> .     |
| group              | Specifies the group of the ASM instance to which the iSCSI quorum disks will be added. This is an optional parameter. The default value is <code>oinstall</code> . |
| network-iface-list | Specifies the list of network interface names through which the local and remote targets will be discovered.   |
| zone-name          | Specifies the zone's name if the quorum disk will be created inside a zone. This is an optional parameter  |

---

**Note** - For `--create --config` option, you can provide any value in the `--network-iface-list` parameter, since it will be ignored.

---

## Example

```
quorumdiskmgr --create --config --owner=oracle --group=oinstall --network-iface-list="
ib0, ib1" --zone-name dbadm010101
```

## Creating a Target (--create --target)

The `--create --target` action creates a target that can be accessed by database servers with an InfiniBand IP address in the specified InfiniBand IP address list and will be used to create devices that will be added to the specified ASM disk group.

After a target is created, its `asm-disk-group`, `host-name`, and `size` attributes cannot be changed.

## Syntax

```
quorumdiskmgr --create --target --asm-disk-group asm_disk_group --visible-to
infiniband_ip_list [--host-name host_name] [--size size] [--zone-name zone]
```

## Parameters

| Parameter      | Description  |
|----------------|--|
| asm-disk-group | Specifies the ASM disk group to which the device created from the target will be added. The value of asm-disk-group is not case-sensitive.   |
| visible-to     | Specifies a list of InfiniBand IP addresses. Database servers with an InfiniBand IP address in the list will have access to the target.  |
| host-name      | Specifies the host name of the database server on which quorumdiskmgr runs. The total length of asm-disk-group and host-name cannot exceed 26 characters. If the host name is too long, a shorter host name can be specified as long as a different host name is specified for each database server in the quarter rack. This is an optional parameter. The default value is the host name of the database server on which quorumdiskmgr runs. The value of host-name is not case-sensitive. |
| size           | Specifies the size of the target. This is an optional parameter. The default value is 128 MB.  |
| zone-name      | Specifies the zone's name if the quorum disk will be created inside a zone. This is an optional parameter  |

## Example

```
quorumdiskmgr --create --target --size=128MB --asm-disk-group=datacl --visible-to="
192.168.10.45, 192.168.10.46" --host-name=db01 --zone-name dbadm010101
```

## Creating a Device (--create --device)

The --create --device action creates devices by discovering and logging into targets on database servers with an InfiniBand IP address in the specified list of IP addresses.

The created devices will be automatically discovered by the ASM instance with the owner and group specified during configuration creation.

## Syntax

```
quorumdiskmgr --create --device --target-ip-list target_ip_list [--zone-name zone]
```

## Parameters

| Parameter      | Description  |
|----------------|--|
| target-ip-list | Specifies a list of InfiniBand IP addresses. quorumdiskmgr will discover targets on database servers with an InfiniBand IP address in the list and log into the targets to create devices. |
| zone-name      | Specifies the zone's name if the quorum disk will be created inside a zone. This is an optional parameter.   |

## Example

```
quorumdiskmgr --create --device --target-ip-list="192.168.10.45, 192.168.10.46" --zone-name dbadm010101
```

## Listing Quorum Disk Configurations (--list --config)

The `--list --config` action lists the quorum disk configurations.

## Syntax

```
quorumdiskmgr --list --config [--zone-name zone]
```

## Sample Output

```
Owner: grid
Group: oinstall
ifaces: ib1 ib0
```

## Listing Targets (--list --target)

The `--list --target` action lists the attributes of targets, including target name, size, host name, ASM disk group name, the list of IP addresses (a visible-to IP address list) indicating

which database servers have access to the target, and the list of IP addresses (a discovered-by IP address list) indicating which database servers have logged into the target.

If an ASM disk group name is specified, the action lists all local targets created for the specified ASM disk group. Otherwise, the action lists all local targets created for quorum disks.

## Syntax

```
quorumdiskmgr --list --target [--asm-disk-group asm_disk_group] [--zone-name zone]
```

## Parameters

| Parameter      | Description  |
|----------------|--|
| asm-disk-group | Specifies the ASM disk group. quorumdiskmgr displays all local targets for this ASM disk group. The value of asm-disk-group is not case-sensitive. |
| zone-name      | Specifies the zone's name if the quorum disk is inside a zone. This is an optional parameter.  |

## Sample Output

```
Name: iqn.2015-05.com.oracle:QD_DATA1_DB01
Size: 128 MB
Host name: DB01
ASM disk group name: DATA1
Visible to: 192.168.10.48, 192.168.10.49, 192.168.10.46, 192.168.10.47
Discovered by: 192.168.10.47, 192.168.10.46
```

## Listing Devices (--list --device)

The --list --device action lists the attributes of devices, including device path, size, host name and ASM disk group name.

If only an ASM disk group name is specified, the action lists all the devices that have been added to the ASM disk group.

If only a host name is specified, the action lists all the devices created from the targets on the host.

If both an ASM disk group name and a host name are specified, the action lists a single device created from the target on the host and that has been added to the ASM disk group.

If neither an ASM disk group name nor a host name is specified, the action lists all quorum disk devices.

## Syntax

```
quorumdiskmgr --list --device [--asm-disk-group asm_disk_group] [--host-name host_name]
[--zone-name zone]
```

## Parameters

| Parameter      | Description  |
|----------------|--|
| asm-disk-group | Specifies the ASM disk group to which devices have been added. The value of asm-disk-group is not case-sensitive.                    |
| host-name      | Specifies the host name of the database server from whose targets devices are created. The value of host-name is not case-sensitive. |
| zone-name      | Specifies the zone's name if the quorum disk is inside a zone. This is an optional parameter.  |

## Sample Output

```
Device path: /dev/exadata_quorum/QD_DATAAC1_DB01
Size: 128 MB
Host name: DB01
ASM disk group name: DATAAC1
```

```
Device path: /dev/exadata_quorum/QD_DATAAC1_DB02
Size: 128 MB
Host name: DB02
ASM disk group name: DATAAC1
```

## Deleting Configurations (--delete --config)

The `--delete --config` action deletes quorum disk configurations. The configurations can only be deleted when there are no targets or devices present.

## Syntax

```
quorumdiskmgr --delete --config [--zone-name zone]
```

## Deleting Targets (--delete --target)

The `--delete --target` action deletes the targets created for quorum disks on database servers.

If an ASM disk group name is specified, the action deletes all local targets created for the specified ASM disk group. Otherwise, the action deletes all local targets created for quorum disks.

## Syntax

```
quorumdiskmgr --delete --target [--asm-disk-group asm_disk_group] [--zone-name zone]
```

## Parameters

| Parameter      | Description   |
|----------------|---|
| asm-disk-group | Specifies the ASM disk group. Local targets created for this disk group will be deleted. The value of asm-disk-group is not case-sensitive. |
| zone-name      | Specifies the zone's name if the quorum disk is inside a zone. This is an optional parameter.   |

## Example

```
quorumdiskmgr --delete --target --asm-disk-group=datac1 --zone-name dbadm010101
```

## Deleting Devices (--delete --device)

The `--delete --device` action deletes quorum disk devices.



If only an ASM disk group name is specified, the action deletes all the devices that have been added to the ASM disk group.

If only a host name is specified, the action deletes all the devices created from the targets on the host.

If both an ASM disk group name and a host name are specified, the action deletes a single device created from the target on the host and that has been added to the ASM disk group.

If neither an ASM disk group name nor a host name is specified, the action deletes all quorum disk devices.

## Syntax

```
quorumdiskmgr --delete --device [--asm-disk-group asm_disk_group] [--host-name
host_name] [--zone-name zone]
```

## Parameters

| Parameter      | Description  |
|----------------|--|
| asm-disk-group | Specifies the ASM disk group whose device you want to delete. The value of asm-disk-group is not case-sensitive.   |
| host-name      | Specifies the host name of the database server. Devices created from targets on this host will be deleted. The value of host-name is not case-sensitive. |
| zone-name      | Specifies the zone's name if the quorum disk is inside a zone. This is an optional parameter.  |

## Example

```
quorumdiskmgr --delete --device --host-name=db01 --zone-name dbadm010101
```

## Changing Owner and Group Values (--alter --config)

The --alter --config action changes the owner and group configurations.

## Syntax

```
quorumdiskmgr --alter --config --owner owner --group group [--zone-name zone]
```

## Parameters

| Parameter | Description  |
|-----------|--|
| owner     | Specifies the owner of the ASM instance to which the iSCSI quorum disks will be added. This is an optional parameter. The default value is grid.     |
| Group     | Specifies the group of the ASM instance to which the iSCSI quorum disks will be added. This is an optional parameter. The default value is oinstall. |
| zone-name | Specifies the zone's name if the quorum disk is inside a zone. This is an optional parameter.  |

## Example

```
quorumdiskmgr --alter --config --owner=grid --group=oinstall --zone-name dbadm010101
```

## Changing the InfiniBand IP Addresses (--alter --target)

The `--alter --target` action changes the InfiniBand IP addresses of the database servers that have access to the local target created for the specified ASM disk group.

## Syntax

```
quorumdiskmgr --alter --target --asm-disk-group asm_disk_group --visible-to  
infiniband_ip_list [--zone-name zone]
```

## Parameters

| Parameter      | Description  |
|----------------|--|
| asm-disk-group | Specifies the ASM disk group to which the device created from the target will be added. The value of asm-disk-group is not case-sensitive. |
| visible-to     | Specifies a list of InfiniBand IP addresses. Database servers with an InfiniBand IP address in the list will have access to the target.    |
| zone-name      | Specifies the zone's name if the quorum disk is inside a zone. This is an optional parameter.  |

## Example

```
quorumdiskmgr --alter --target --asm-disk-group=datac1 --visible-to="192.168.10.45,  
192.168.10.47" --zone-name dbadm010101
```

