

VERITAS Volume Manager™ 4.0 Cross-platform Data Sharing

Administrator's Guide

Disclaimer

The information contained in this publication is subject to change without notice. VERITAS Software Corporation makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. VERITAS Software Corporation shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this manual.

VERITAS Software Corporation
350 Ellis Street
Mountain View, CA 94043
USA
Phone 650-527-8000
Fax 650-527-2908
<http://www.veritas.com>

Copyright

Copyright © 2004 VERITAS Software Corporation. All rights reserved. VERITAS, the VERITAS Logo, and all other VERITAS product names and slogans are trademarks or registered trademarks of VERITAS Software Corporation. VERITAS and the VERITAS Logo Reg. U.S. Pat. & Tm. Off. Other product names and/or slogans mentioned herein may be trademarks or registered trademarks of their respective companies.

Third-Party Copyrights

Data Encryption Standard (DES) Copyright

Copyright © 1990 Dennis Ferguson. All rights reserved.

Commercial use is permitted only if products that are derived from or include this software are made available for purchase and/or use in Canada. Otherwise, redistribution and use in source and binary forms are permitted.

Copyright 1985, 1986, 1987, 1988, 1990 by the Massachusetts Institute of Technology. All rights reserved.

Export of this software from the United States of America may require a specific license from the United States Government. It is the responsibility of any person or organization contemplating export to obtain such a license before exporting.

WITHIN THAT CONSTRAINT, permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of M.I.T. not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. M.I.T. makes no representations about the suitability of this software for any purpose. It is provided as is without express or implied warranty.

Contents

Preface	vii
Audience	vii
Organization	vii
Related Documents	viii
Conventions	ix
Getting Help	x
Downloading and Running VRTSexplorer	x
Chapter 1. Overview	1
General Concepts	1
Sector Size	1
Block Size	2
Operating System Data	2
CDS Disk Access and Format	2
CDS Disk Types	3
Private and Public Regions	3
Minor Device Numbers	3
Disk Access Type <i>auto</i>	4
Platform Block	4
AIX Coexistence Label	4
HP-UX Coexistence Label	4
VxVM ID Block	4
CDS Disk Groups	4
Device Quotas	5
Non-CDS Disk Groups	6
Disk Group Alignment	6



Alignment Values	7
Alignment Value 1	7
Alignment Value 8K	7
CDS Maps and Logs	7
DCO Map Layout	8
Region Sizes and Disk Group Alignment	8
Traditional DRL and logmap_len	8
Chapter 2. Setting up Your System	9
Creating CDS Disks from Uninitialized Disks	9
Using vxdisksetup	9
Using vxdiskadm	10
Creating CDS Disks from Initialized VxVM Disks	10
Disk Not in a Disk Group	10
Disk Already in a Disk Group	10
Creating CDS Disk Groups	11
Using vxdg init to Create CDS Disk Groups	11
Using vxdiskadm to Create CDS Disk Groups	11
Converting a Non-CDS Disk to a CDS Disk	12
Converting a Non-CDS Disk Group to a CDS Disk Group	13
Notes	14
Examples	14
Licensing	15
Defaults File	16
Chapter 3. Maintaining Your System	17
Disk Tasks	18
Changing Default Disk Format Setting for Disk Creation	18
Converting an LVM Disk	18
Encapsulating Existing Partitions/Slices	18
Restoring CDS Disk Labels	18



Disk Group Tasks	20
Setting Alignment for a Non-CDS Disk Group	20
Changing CDS Disk Group Alignment	21
Reducing Alignment of a CDS Disk Group	21
Increasing Alignment of a non-CDS Version 110 Disk Group	22
Determining if the CDS Attribute Is Set	22
Joining CDS Disk Groups	22
Moving Objects Between CDS and Non-CDS Disk Groups	23
Moving Objects Between CDS Disk Groups	23
Changing Default CDS Setting for Disk Group Creation	23
Creating Non-CDS Disk Groups	23
Upgrading to a CDS Disk Group	23
Before Upgrading (Pre-Version 110 Disk Group)	24
After Upgrading (Version 110 Disk Group)	24
Replacing a Disk in a CDS Disk Group	24
Setting Device Quotas for CDS Disk Groups	24
Setting DRL Map Size and Log Size	25
Notes	25
Creating a DRL Log	26
Using vxassist to Create a DRL Log	26
Using vxvol to Set a DRL Map Length	26
Displaying Information	27
Displaying Traditional DRL Map Size and Log Size	27
Displaying Disk Group Alignment	27
Displaying Volume Log Map Values	28
Listing Offset and Length Information	28
Listing CDS Disk Groups	29
Listing Disks	29
Displaying Device Quotas for CDS Disk Groups	29
Default Activation of Shared Disk Groups	30



Creating CDS Disk Groups Using vxdg split	30
Appendix A. Cross-platform Transfer — HP-UX	31
Default Activation of Shared Disk Groups	31
Block Size	31
Alignment Value	31
Appendix B. Cross-platform Transfer — Linux	33
Importing Disk Groups Between Linux and Non-Linux Machines	33
Alignment Value	33
Appendix C. Cross-platform Transfer — Solaris	35
Disk Group alignment and Encapsulated Disks	35
Alignment Value	35
Appendix D. Error Codes and Error Recovery	37
Error Codes	37
Glossary	39
Index	43



Preface

The *VERITAS Volume Manager™ Cross-platform Data Sharing Administrator's* guide describes how to use the VERITAS Cross-platform Data Sharing (CDS) product.

Audience

This guide is intended for system administrators responsible for configuring and maintaining systems using the VERITAS CDS product. This guide assumes the following:

- ◆ A basic understanding of system administration
- ◆ A working knowledge of the operating systems on all supported platforms.
- ◆ Familiarity with VxVM 4.0 and storage management concepts

Organization

This guide is organized as follows:

- ◆ Chapter 1, “[Overview](#),” gives an introduction to CDS concepts.
- ◆ Chapter 2, “[Setting up Your System](#),” presents considerations when transferring data between platforms.
- ◆ Chapter 3, “[Maintaining Your System](#),” describes how to use CDS.
- ◆ Appendix A, “[Cross-platform Transfer — Linux](#),” contains cross-platform considerations for the Linux platform.
- ◆ Appendix B, “[Cross-platform Transfer — HP-UX](#),” contains cross-platform considerations for the HP-UX platform.
- ◆ Appendix C, “[Cross-platform Transfer — Solaris](#),” contains cross-platform considerations for the Solaris platform.
- ◆ Appendix D, “[Error Codes and Error Recovery](#),” describes the error messages output by CDS.
- ◆ The [Glossary](#) defines terms that relate to CDS.



Related Documents

The following documents provide related information:

- ◆ *VERITAS Volume Manager Administrator's Guide*
- ◆ *VERITAS Volume Manager Troubleshooting Guide*
- ◆ *VERITAS Volume Manager Hardware Notes*
- ◆ *VERITAS Volume Manager Installation Guide*
- ◆ *VERITAS Volume Manager Release Notes*
- ◆ *VERITAS Volume Manager User's Guide*
- ◆ VERITAS Volume Manager manual pages



Conventions

The following table describes the typographic conventions used in this guide.

Typeface	Usage	Examples
monospace	Computer output, file contents, files, directories, software elements such as command options, function names, and parameters	Read tunables from the <code>/etc/vx/tunefstab</code> file. See the <code>ls(1)</code> manual page for more information.
<i>italic</i>	New terms, book titles, emphasis, variables to be replaced by a name or value	The variable <code>ncsize</code> determines the value of...
monospace (bold)	User input; the “#” symbol indicates a command prompt ¹	#mount -F vxfs /h/filesys
<i>monospace (bold and italic)</i>	Variables to be replaced by a name or value in user input ¹	#mount -F <i>fstype</i> <i>mount_point</i>

Symbol	Usage	Examples
%	C shell prompt	
\$	Bourne/Korn/Bash shell prompt	
#	Superuser prompt (all shells)	
\	Continued input on the following line ¹	#mount -F vxfs \ /h/filesys
[]	In a command synopsis, brackets indicates an optional argument	<code>ls [-a]</code>
	In a command synopsis, a vertical bar separates mutually exclusive arguments	<code>mount [suid nosuid]</code>



Getting Help

If you have any comments or problems with the VERITAS products, contact VERITAS Technical Support:

- ◆ U.S. and Canadian Customers: 1-800-342-0652
- ◆ International Customers: +1 (650) 527-8555
- ◆ E-mail: support@veritas.com

For license information (U.S. and Canadian Customers):

- ◆ Phone: 1-650-527-0300
- ◆ E-mail: license@veritas.com
- ◆ Fax: 1-650-527-0952

For software updates:

- ◆ E-mail: swupdate@veritas.com

For additional technical support information, such as TechNotes, product alerts, and hardware compatibility lists, visit the VERITAS Technical Support Web site at:

- ◆ <http://support.veritas.com> (U.S. and Canada Customers)

For information on purchasing VERITAS products:

- ◆ Phone: 1-800-327-2232
- ◆ Email: sales.mail@veritas.com

For additional information about VERITAS and VERITAS products, visit the Web site at:

- ◆ <http://www.veritas.com>

Downloading and Running VRTSexplorer

If you have access to the Internet, you can use the VRTSexplorer program to assist Technical Support in diagnosing the cause of your problem. Refer to the *VERITAS Volume Manager Release Notes* for information on downloading and using the VRTSexplorer program.

This chapter presents an overview of Cross-platform Data Sharing (CDS):

- ◆ [General Concepts](#)
- ◆ [CDS Disk Access and Format](#)
- ◆ [CDS Disk Groups](#)

General Concepts

CDS provides you with a foundation for sharing data concurrently across heterogeneous systems (Solaris, HP-UX, AIX (RS/6K), and Linux (x86)). With CDS, all systems have direct access to physical devices holding data. It allows VERITAS products and applications to access data storage independently of Operating System (OS) or platform, enabling them to work transparently in heterogeneous environments.

For data sharing to occur, the following levels in the device hierarchy, from disk through file system, must provide support for data sharing:

- ◆ Operating system - device level
- ◆ VERITAS Volume Manager (VxVM) — volume level
- ◆ VERITAS File System™ (VxFS) — file system level
- ◆ End-user applications — application level

CDS utilizes a new disk type (*auto:cdsdisk*). To effect data sharing, VxVM supports a new disk group attribute (*cds*) and also supports different OS block sizes.

CDS is a license-enabled feature applied at the disk group level.

Sector Size

Sector size is an attribute of a disk drive (or SCSI LUN for an array-type device), which is set when the drive is formatted. Sectors are the smallest addressable unit of storage on the drive, and are the units in which the device performs I/O. The sector size is significant



because it defines the atomic I/O size at the device level. Any multi-sector writes which VxVM submits to the device driver are not guaranteed to be atomic (by the SCSI subsystem) in the case of system failure.

Block Size

The block size is a platform-dependent value that is greater than or equal to the sector size. Each platform accesses the disk on block boundaries and in quantities that are multiples of the block size. Data that is created on one platform, and then accessed by a platform of a different block size, can suffer from the following problems:

- ◆ Addressing issues
 - The data may not have been created on a block boundary compatible with that used by the accessing platform.
 - The accessing platform cannot address the start of the data.

- ◆ Bleed-over issues

The size of the data written may not be an exact multiple of the block size used by the accessing platform. Therefore the accessing platform cannot constrain its I/O within the boundaries of the data on disk.

Operating System Data

Some Operating Systems (OS) require OS-specific data on disks in order to recognize and control access to the disk.

CDS Disk Access and Format

For a disk to be accessible by multiple platforms, the disk must be consistently recognized by the platforms, and all platforms must be capable of performing I/O on the disk. CDS disks contain specific content at specific locations to identify or control access to the disk on different platforms. The same content and location are used on all CDS disks, independent of the platform on which the disks are initialized.

In order for a disk to be converted to a CDS disk, it must satisfy the following requirements:

- ◆ Must be a SCSI disk that supports Mode Sense
- ◆ Cannot be an EFI disk
- ◆ Must be the entire physical disk (LUN)
- ◆ Only one Volume Manager can manage a physical disk (LUN)

- ◆ There can be no disk partition (slice) which is not contained within the CDS disk
- ◆ Must exist in either a CDS disk group or a non-CDS disk group whose version is at least 110 (but not a disk group whose version is less than 110)
- ◆ Cannot contain a volume whose use-type is either `root` or `swap` (for example, it cannot be a boot disk)

CDS Disk Types

When CDS disks are created, they are given the default format `cdsdisk`. You can also use `vxcdsconvert` to convert other VxVM disk formats and types to CDS format. CDS disk format, `cdsdisk`, is recognized by all VxVM platforms (including Windows).

You can change the default format used when creating disks by editing the `/etc/default/vxdisk` file.

If there is a conflict between the default disk format and default disk group (DG) type created by the system (for example, CDS DG, and non-CDS disk), then you can either:

- ◆ Create a non-CDS DG, by setting `cds=off`
- or*
- ◆ Create a CDS disk, by explicitly re-initializing the disk as a CDS disk

Note Disks with format `cdsdisk` can only be added to disk groups Version 110 or later

Private and Public Regions

All regions within the private area are aligned and sized in multiples of 8K.

The public region is aligned and sized in multiples of 8K.

The alignment of objects within the public region is directed by the disk group alignment attribute.

Note With other (non-CDS) VxVM disk formats, the private and public regions are aligned to the platform-specific OS block size.

Minor Device Numbers

Importing a disk group will fail if it will exceed the maximum devices for that platform.

Note There is a large disparity between the maximum number of devices allowed for devices on the Linux platform, and that for other supported platforms.



Disk Access Type *auto*

The disk access (DA) disk type *auto* supports multiple disk formats, including `cdsdisk`, which is supported across all platforms. It is associated with the DA records created by the VxVM auto-configuration mode. Disk type *auto* automatically determines which format is on the disk.

Platform Block

The Platform Block resides on disk sector 0, and contains data specific to the Operating System for the platforms. It is necessary for proper interaction with each of those platforms. The platform block allows a disk to perform as if it was initialized by each of the specific platforms.

AIX Coexistence Label

The AIX Coexistence label resides on the disk, and identifies the disk to the AIX volume manager (LVM) as being controlled by VxVM.

HP-UX Coexistence Label

The HP-UX Coexistence label resides on the disk, and identifies the disk to the HP volume manager (LVM) as being controlled by VxVM.

VxVM ID Block

The VxVM ID block resides on the disk, and indicates the disk is under VxVM control. It provides dynamic VxVM private region location and other information.

CDS Disk Groups

CDS disk groups are available only with disk group versions 110 and greater, and with 8K alignment. Pre-version 110 disk groups do not support CDS disks.

A CDS disk group has the following qualities:

- ◆ Composed only of CDS disks (that is, disks with the VxVM disk format `cdsdisk`).
- ◆ Contains the CDS attribute `cds` set to `on`. The CDS attribute indicates the disk group's ability to be shared across platforms. See [“Maintaining Your System”](#) on page 17. By default, newly initialized disk groups created using VxVM 4.0 are created with `cds` set to `on`.

A CDS disk group (DG) can be:

- ◆ Initialized on one system and then used “as-is” by VxVM on a system employing a different type of platform.
- ◆ Imported (in a serial fashion) by systems of the platforms (Linux, Solaris, and HP-UX).
- ◆ Private, shared, or distributed.

Note CDS disk groups can be shared by CVM.

You cannot include the following disks or volumes in a CDS disk group:

- ◆ Volumes of usage type `root` and `swap`. (This means that you cannot use CDS to share boot devices.)
- ◆ Encapsulated disks.

Objects in a CDS disk group have the following properties:

- ◆ Allow data written on any of the supported platforms to be accessed by any of the other supported platforms. Since each platform accesses an object using its concept of a system I/O block, all objects in the CDS disk group must be aligned and sized such that any system can access the object using the system’s concept of an I/O block. Alignment of objects is set using the platform-independent *DG alignment* attribute.
- ◆ Support OS block size up to 8K.

Device Quotas

Device quotas limit the number of objects in the disk group which create associated device nodes in the file system. (This is useful for disk groups which share serially between Linux and other supported platforms.)

You can limit the number of devices that can be created in a given CDS disk group by setting the device quota (refer to [“Setting Device Quotas for CDS Disk Groups”](#) on page 24).

When you create a device, an error is returned if the number of devices would exceed the device quota. You then either need to increase the quota, or remove some objects using device numbers, before the device can be created.

See also [“Displaying Device Quotas for CDS Disk Groups”](#) on page 29 for instructions on displaying the device quota value.



Non-CDS Disk Groups

Any version 110 (or greater) disk group (DG) can contain both CDS and non-CDS disks. However, only version 110 (or greater) disk groups composed entirely of CDS disks have the ability to be shared across platforms. Whether or not that ability has been enabled is a license-controlled attribute of the disk group (the *cds* attribute). Enabling that attribute causes a non-CDS disk group to become a CDS disk group.

Although a non-CDS disk group can contain a mixture of CDS and non-CDS disks having dissimilar private region alignment characteristics, its disk group alignment will still direct how all subdisks are created.

Note The term *disk encapsulation* implies encapsulation using *vxencap* on Solaris/Linux and LVM conversion on AIX and HP-UX.

Disk Group Alignment

One of the attributes of the disk group is the block *alignment*, which represents the largest block size supported by the disk group. The alignment constrains attributes of the objects within the disk group, such as:

- ◆ Subdisk; how subdisks are positioned on the drive (*offset*), and their size granularity (*length*)
- ◆ Plex offset
- ◆ Volume length
- ◆ Log length
- ◆ Stripe width

The disk group alignment is assigned at disk group creation time (see “[Maintaining Your System](#)” on page 17).

Alignment Values

The disk group block alignment has two values, one and 8K.

Alignment Value 1

All disk group versions before version 110 require alignment value 1.

Disk groups have alignment value 1 after upgrading to VM 4.0.

Disk groups have alignment value 1 after upgrading from pre-version 110 to version 110 or later.

Encapsulated disks and other features which require disk cylinder alignment have alignment value 1.

Non-CDS disk groups version 110 and later can have either a value of 1 block or 8k (8KB).

Alignment Value 8K

Non-CDS disk groups version 110 and later can have either a value of 1 block or 8k (8KB).

All CDS disk groups have an alignment value of 8k.

The alignment for all newly initialized disk groups in VxVM 4.0 is 8k. This value (which is used when creating the disk group) cannot be changed. However, the disk group alignment can be subsequently changed (refer to [“Changing CDS Disk Group Alignment”](#) on page 21).

Note The default usage of `vxassist` is to set the `layout=diskalign` attribute on all platforms. This is ignored on 8K-aligned disk groups, which means that scripts relying on the default may fail.

CDS Maps and Logs

With VxVM 4.0, *Dirty Region Logging* (DRL) is available as part of a DCO Version 20 volume. (Please refer to the *VERITAS Volume Manager Administrator's Guide* for information on DCO and DRL.)

Note Traditional DRL is also still available.



DCO Map Layout

The DCO map layout changed with VxVM 4.0, though the original format is still available. The old layout is available in DCO Version 0 objects, and the new layout is available in DCO Version 20 objects.

Region Sizes and Disk Group Alignment

The location and size of each DRL map within a DRL log must not violate the disk group alignment for the disk group (containing the volume to which the DRL log is associated). This means that the region size and alignment of each DRL map must be a multiple of the disk group alignment. (Features utilizing the region size can impose additional minimums and size granule over and above this restriction, but cannot violate it.)

Version 110 disk groups with traditional DRL have:

- ◆ Minimum region size for DRL volumes is 512KB
- ◆ Increment region size for DRL volumes is 64KB

Version 110 disk groups with a Version 0 DCO have:

- ◆ Minimum region size for DCO volumes is 16KB.
- ◆ Incremental region size for DCO volumes is 8KB.

Traditional DRL and logmap_len

The `logmap_len` attribute applies only to traditional DRL.

Note For CVM, each active map can be concurrently updated by I/O threads on each node of the cluster.

Each map must be a multiple of the disk group alignment (that is, a minimum of 8K map size on 8K-aligned or CDS disk groups).

The `logmap_len` attribute can be used to adjust how much of the on-disk bitmap is used, changing its effective size.

A value of "0" means use the entire on-disk map.

On newly created volumes, `vxassist` sets `logmap_len` to a reasonable value, but it can be altered using `vxvol set`.

This chapter describes how to set up your initial CDS system. It includes information on:

- ◆ [Creating CDS Disks from Uninitialized Disks](#)
- ◆ [Creating CDS Disks from Initialized VxVM Disks](#)
- ◆ [Creating CDS Disk Groups](#)
- ◆ [Converting a Non-CDS Disk to a CDS Disk](#)
- ◆ [Converting a Non-CDS Disk Group to a CDS Disk Group](#)
- ◆ [Licensing](#)
- ◆ [Defaults File](#)

Creating CDS Disks from Uninitialized Disks

To create a CDS disk from an uninitialized disk, use one of the following commands:

- ◆ `vxdisksetup`
- ◆ `vxdiskadm`

Using `vxdisksetup`

You can use the `vxdisksetup` command to create a CDS disk.

```
# vxdisksetup -i disk [format=format]
```

The `format` attribute is needed only if you want to override the default format.

Note The `-i` option is required if you want to write the format. Without this option, the only activity is to write the VTOC, and to carry out some error checking.

Using the `format=cdsdisk` attribute overrides the default setting in the `/etc/default/vxdisk` file.



Using vxdiskadm

Use the `disk init` option. You are prompted to specify the format.

Creating CDS Disks from Initialized VxVM Disks

If the disk is already initialized, there are two cases to consider:

- ◆ [Disk Not in a Disk Group](#)
- ◆ [Disk Already in a Disk Group](#)

Disk Not in a Disk Group

If the disk is not in a disk group:

1. Run the command:

```
# vxdiskunsetup disk
```

This sets `auto:none` format for the disk. This is necessary as non-auto types cannot be reinitialized by `vxdisksetup`.

2. If the DA record is persistent (that is, if the disk is listed in `/etc/vx/darecs`), remove the disk using this command:

```
# vxdisk rm disk
```

3. Rescan for the disk using this command:

```
# vxdisk scandisks
```

4. Finally, run this command to set up the disk:

```
# vxdisksetup -i disk
```

Disk Already in a Disk Group

If the disk is already in a disk group, run `vxcdsconvert`, as described in [“Converting a Non-CDS Disk to a CDS Disk”](#) on page 12.

Creating CDS Disk Groups

You can create a CDS disk group (DG) in two ways:

- ◆ [Using vx dg init to Create CDS Disk Groups](#)
- ◆ [Using vx diskadm to Create CDS Disk Groups](#)

Using vx dg init to Create CDS Disk Groups

You can use the `vx dg init` command to create a CDS disk group:

```
# vx dg init diskgroup disklist [cds=on|off]
```

Use the `cds` attribute only if you do not want the default `cds` value.

The `cds` attribute has the values `on` or `off`. The default is `on` (create a CDS disk group), unless specified otherwise in the `/etc/default/vx dg` file.

Note You cannot specify `cds=on` and `-T version` for a non-CDS disk group with a version that is less than 110.

You can have a non-CDS disk group created by default when the `cds` attribute is not specified on the `vx dg init` command by using a line of the form `cds=off` in the `/etc/default/vx dg` defaults file.

A CDS disk group has an 8KB alignment by default. There is no mechanism for specifying an alternate disk group alignment when the disk group is created.

See the `vx dg(1M)` manual page for more information on the `vx dg` command.

Using vx diskadm to Create CDS Disk Groups

You can create a CDS disk group when using `vx diskadm` to initialize a disk. Specify that the disk group should be a CDS disk group when prompted.

Note You cannot create a CDS disk group using the `vx diskadm` command as part of encapsulating, or LVM conversion.

When the `vx diskadm` command initializes a disk into an existing CDS disk group, the disk must have format `cdsdisk`.

The CDS attribute for the disk group remains unchanged (that is, you are not given the option to change it).



Note When initializing a disk, if the target disk group is an existing CDS disk group, `vxdiskadm` will only allow the disk to be initialized as a CDS disk. If the target disk group is a non-CDS disk group, the disk can be initialized as either a CDS disk or a non-CDS disk.

Converting a Non-CDS Disk to a CDS Disk

Use the CDS conversion utility (`vxcdsconvert`) to convert non-CDS disks to CDS disks, to make them portable between different operating systems that are running VxVM with the CDS feature.

```
# vxcdsconvert -g diskgroup [-A] [-d defaults_file] \  
[-o novolstop] alldisks [attribute=value] ...  
  
# vxcdsconvert -g diskgroup [-A] [-d defaults_file] \  
[-o novolstop] disk name [attribute=value] ...
```

Note The existing DA record needs to have a type of `auto` in order to be re-initialized as a CDS disk.

You can use `vxcdsconvert` options to convert CDS disks either singly, or by disk group:

- ◆ Specify a single disk. You might consider this option, rather than converting the entire disk group, if you do not want a CDS disk group. For example:
 - Where a disk in the non-CDS disk group has cross-platform exposure, and you want other VxVM nodes to recognize the disk and not assume it is available for initialization.
 - LVM on HP-UX and AIX, and Windows VxVM nodes need to recognize CDS disks, but not attempt to initialize or manage them.
 - The intention is to move the disk into an existing CDS disk group.
- ◆ All the non-CDS disks in a disk group are converted.

You can use `vxcdsconvert` options to convert CDS disks either on-line, or off-line:

- ◆ Perform the conversion on-line (that is, while use of the disk continues). However, this may greatly increase the amount of time (and work) required for disk conversion as it may involve evacuating objects from the disk, reinitializing the disk, and relocating objects back to disk.
- ◆ Perform the conversion off-line. This requires minimal offline time.

Note Disk conversion may take significantly longer if the `-o novolstop` option is specified, depending on the subdisk layout of the disk.

Refer to the `vxcdsconvert(1M)` manual page for further information on options, attributes, and keywords.

Converting a Non-CDS Disk Group to a CDS Disk Group

Use the CDS conversion utility (`vxcdsconvert`) to make a VxVM non-CDS disk group (DG) portable between different operating systems that are running versions of VxVM with the CDS feature. This allows existing data to be made shareable in a CDS environment.

```
# vxcdsconvert -g diskgroup [-A] [-d defaults_file] \
  [-o novolstop] alignment [attribute=value] ...

# vxcdsconvert -g diskgroup [-A] [-d defaults_file] \
  [-o novolstop] group [attribute=value] ...
```

CDS provides the following options for disk group conversion:

- ◆ Specify alignment conversion. In this case, disks are not converted, and an object relayout is performed on the disk group. A successful completion results in an 8K-aligned disk group. You might consider this option, rather than converting the entire disk group, if you want to reduce the amount of work to be done for a later full conversion to CDS disk group.
- ◆ Specify group conversion. This implies `alldisk`, and will perform that function prior to object relayout. All the non-CDS disks in the disk group are converted.
- ◆ Perform the conversion on-line (that is, while use of the disk group continues). However, for a group conversion, this may greatly increase the amount of time (and work) required for conversion.
- ◆ Perform the conversion off-line. This requires minimal offline time.

Note the following:

- ◆ The disk group must be in pristine condition. That is:
 - It has no dissociated or disabled objects.
 - No sparse plexes are present.
 - There are no volumes requiring recovery or having pending snapshot operations.
 - There are no objects in an error state.
- ◆ Stopped (but startable) volumes will be started, for the duration of the conversion only.



- ◆ Conversion has the following side effects:
 - Any objects created with `layout=diskalign` can no longer be disk aligned.
 - Encapsulated disks may lose the ability to be encapsulated.
 - Fine performance tuning may be lost as data may have migrated (and even migrated to different disks).

Please refer to the `vxcdsconvert(1M)` manual page for information on options, attributes, and keywords.

Notes

- ◆ A non-CDS disk group will be upgraded (using the `vx dg upgrade` command). Also, if the disk group was originally created by the conversion of an LVM Volume Group (VG), then rolling back to the original LVM VG is not an option. If you decide to go through with the conversion, the rollback records for the disk group will be removed, so that an accidental rollback to an LVM VG cannot be done.
- ◆ If the non-CDS disk group has one or more disks that are not CDS disks, these disks are converted to CDS disks. If you do not want one or more of these disks to be converted, you need to move them out of the disk group (for example, using `vx dg move` or `split`) prior to invoking the `vxcdsconvert group` command.
- ◆ If the non-CDS disk group does not have a CDS-compatible disk group alignment, the objects need to go through relayout, so they have a CDS-compatible alignment.
- ◆ Unless the `novolstop` option is specified, applications using disks which require format conversion must be terminated for the duration of the disk conversion process.
- ◆ Use of the `novolstop` option for the disk conversion piece may add a large amount of work, as objects may need to be evacuated and then unrelocated back.

Examples

- ◆ Check if the disk group `mydg` and all its disks can be made portable, using the command:

```
# vxcdsconvert -g mydg -A group
```
- ◆ Make the disk group `mydg` and all its disks portable while its volumes are still online. You can avoid using the `novolstop` option during disk conversion by using the following sequence of commands:
 - Stop all applications that are using volumes on the disks requiring format conversion.

- Enter the command:


```
# vxcdsconvert -g mydg alldisks
```
- Restart applications.
- Enter the command:


```
# vxcdsconvert -g mydg group
```
- ◆ Make the disk group, anodg, and all its disks, portable while its volumes are still online, and allow subdisks to be evacuated to disks anodg11 through anodg14 if required:


```
# vxcdsconvert -g anodg -o novolstop group \  
  move_subdisks_ok=yes evac_subdisks_ok=yes \  
  evac_disk_list=anodg11,anodg12,anodg13,anodg14
```

Licensing

The ability to create or import a CDS disk group is controlled by a CDS license. CDS licenses are included as part of the VERITAS Storage Foundation license.

An example of a CDS enabling license (output by the `vxlicrep` command) is:

```
License Key                               =
P2CE-I8ZR-USLU-RR8D-YN8P-NN8H
  Product Name                             = VERITAS Storage
Foundation Standard
  License Type                             = DEMO
  OEM ID                                    = 58
  Demo End Date                             = Tue Sep 30 01:00:00 2003
                                           (21.5 days from now).
  Editions Product                         = YES
Features :=
  VxVM#VERITAS Volume Manager              = Enabled
  CPU Count#VERITAS Volume Manager         = Not Restricted
  VXFS#VERITAS File System                 = Enabled
  CPU Count#VERITAS File System            = Not Restricted
  QLOG#VERITAS File System                 = Enabled
  VERITAS Foundation Suite                 = Enabled
SPC Lite#VERITAS Storage Manager for Storage Networks = Enabled
  Storage Expert#VERITAS Volume Manager    = Enabled
  Platform                                  = un-used
  Version                                    = un-used
```



Defaults File

Apart from the `vxdisk` system default file, there are two defaults files associated with `vxcdsconvert` that are used to reset `vxcdsconvert` options to values other than the defaults:

- ◆ The `vxcdsconvert` system default file, in `/etc/default/vxcdsconvert`.
- ◆ A user supplied default file, which can be specified by using the `-d` option.

The syntax of each file is as follows:

- ◆ A line that is empty, or begins with a `#` in the first column, is treated as a comment, and is ignored.
- ◆ Apart from comment lines, all other lines must be of the form `attribute=value`, where `attribute` is the name of an option. Each line starts in the first column, and is terminated by the value. No white space is allowed around the `=` sign; otherwise, the `vxcdsconvert` command will exit.

The following attributes are defined:

```
evac_subdisks_ok
min_split_size
move_subdisks_ok
privlen
split_subdisks_ok
evac_disk_list
```

The section describes the following CDS tasks performed by the systems administrator:

- ◆ Disk Tasks
 - Changing Default Disk Format Setting for Disk Creation
 - Restoring CDS Disk Labels
- ◆ Disk Group Tasks
 - Setting Alignment for a Non-CDS Disk Group
 - Changing CDS Disk Group Alignment
 - Determining if the CDS Attribute Is Set
 - Joining CDS Disk Groups
 - Moving Objects Between CDS and Non-CDS Disk Groups
 - Changing Default CDS Setting for Disk Group Creation
 - Creating Non-CDS Disk Groups
 - Upgrading to a CDS Disk Group
 - Replacing a Disk in a CDS Disk Group
 - Setting Device Quotas for CDS Disk Groups
- ◆ Setting DRL Map Size and Log Size
- ◆ Displaying Information
 - Displaying Disk Group Alignment
 - Displaying Volume Log Map Values
 - Listing Offset and Length Information
 - Listing CDS Disk Groups
 - Listing Disks
 - Displaying Device Quotas for CDS Disk Groups



- ◆ [Creating CDS Disk Groups Using vxdg split](#)
- ◆ [Default Activation of Shared Disk Groups](#)

Disk Tasks

The following disk tasks are supported:

- ◆ [Changing Default Disk Format Setting for Disk Creation](#)
- ◆ [Restoring CDS Disk Labels](#)

Changing Default Disk Format Setting for Disk Creation

You can change the default disk format setting from `cdsdisk`; for example you might want to do this for disk initialization, by modifying the `/etc/default/vxdisk` file.

Converting an LVM Disk

When converting an LVM disk (AIX/HP-UX) to a VxVM disk, a CDS disk is created whenever possible. This default behavior can be overridden using the `format` attribute in the `/etc/default/vxencap` file.

Encapsulating Existing Partitions/Slices

Encapsulating existing partitions/slices on a (Linux/Solaris) disk as a VxVM disk will create a CDS disk whenever possible. You can override this default behavior using the `format` attribute in the `/etc/default/vxencap` file.

Restoring CDS Disk Labels

CDS disks have three labels:

- ◆ Platform block
- ◆ AIX coexistence label
- ◆ HP coexistence/VxVM ID block

There are also backup copies of each. If any of the primary labels become corrupted, VxVM will not bring the disk online and user intervention is required.

If two of the three labels are intact, the disk is still be recognized as a `cdsdisk` (though in the 'error' state) and `vxdisk flush` can be used to restore the CDS disk labels from their backup copies.

Primary labels are at sectors 0, 7, and 16; and a normal flush will not flush sectors 7 and 16. Also, the private area is not updated as the disk is not in a disk group. There is no means of finding a “good” private region to flush from. In this case, it is possible to restore the CDS disk labels from the existing backups on disk using the flush operation.

If a corruption happened after the labels were read and the disk is still online and part of a disk group, then a flush operation will also flush the private region.

Caution Caution and knowledge must be employed because the damage could involve more than the CDS disk labels. If the damage is constrained to the first 128K (as would be the case if some entity on the fabric - such as a Windows box - wrote a disk label to a disk which was actually a `cdsdisk` being used in some disk group), then the disk flush would fix it.

Use the following command to rewrite the CDS ID information to a specific disk (this rewrites all labels except sectors 7 and 16):

```
# vxdisk flush disk_accessname
```

Use the following command to rewrite all the disks in a CDS disk group (this rewrites all labels except sectors 7 and 16):

```
# vxdg flush dg_name
```

Use the `-f` option with the `vxdisk` command to forcibly rewrite the AIX coexistence label (sector 7), VxVM ID block (sector 16), HP-UX Coexistence Labels:

```
# vxdisk -f flush disk_accessname
```

This command rewrites all labels if there exists a valid VxVM ID block that points to a valid private region.

Note Sectors 7 and 16 are only rewritten if the `-f` flag is given. The reason for this is that these sectors lay within the first track of the disk, and Windows systems use the first track for their own purposes. In the example above where the disk was taken offline due to label corruption, if it involved sectors 7 or 16, then `-f` is required to fix the problem.



Disk Group Tasks

The following disk group tasks are supported:

- ◆ [Setting Alignment for a Non-CDS Disk Group](#)
- ◆ [Changing CDS Disk Group Alignment](#)
- ◆ [Determining if the CDS Attribute Is Set](#)
- ◆ [Joining CDS Disk Groups](#)
- ◆ [Moving Objects Between CDS and Non-CDS Disk Groups](#)
- ◆ [Changing Default CDS Setting for Disk Group Creation](#)
- ◆ [Creating Non-CDS Disk Groups](#)
- ◆ [Upgrading to a CDS Disk Group](#)
- ◆ [Replacing a Disk in a CDS Disk Group](#)
- ◆ [Setting Device Quotas for CDS Disk Groups](#)

Setting Alignment for a Non-CDS Disk Group

When you encapsulate a disk into a disk group using the `vxdiskadm` command, if the disk group alignment is greater than 1, then alignment is reduced to 1. As part of the encapsulation process, you are asked to confirm that any reduction of disk group alignment is acceptable.

Note If you are using the `vxencap` command to perform the encapsulation, the alignment is carried out automatically, with no confirmation prompt.

The `vxassist` attribute, `dgalign_checking` can have a value of `strict` or `round`, and the attribute may be specified either on the command-line or in a defaults file. If the attribute is not specified, the default value of `round` will be used.

Note If the value is `strict`, specifying an unaligned attribute value will result in an error message. Use of the `round` value will result in all unaligned attribute values being silently rounded.

Use of the `diskalign` layout attribute of `vxassist` is honored only with disk groups whose alignment is 1.

Changing CDS Disk Group Alignment

Use the `alignment` attribute (`align`) with the `vxchg set` command to set the alignment value of a disk group (DG). For example, to set the CDS disk group `cdsdg1` to alignment 8KB you would use a command of the form:

```
# vxchg -g cdsdg1 set alignment=8k
```

Note For a CDS disk group, `alignment` can only take a value of 8k (8KB).

Attempts to set or change a CDS disk group to alignment 1 will fail.

Note Increasing the alignment could involve object relayout, and would then require `vxcdsconvert` to be run.

The sequence: 8KB -> 1 -> 8KB is possible only using `vxchg set` as long as the configuration does not change after 8KB -> 1 transition.

When changing alignment, there are several things that could prevent the change from occurring:

- ◆ Once the CDS attribute is set, the DG alignment cannot be changed to a value less than 8KB.
- ◆ Increasing the alignment may fail if any object in the DG does not already conform to the higher alignment.

Reducing Alignment of a CDS Disk Group

Use the `alignment` attribute with the `vxchg set` command to change the disk group alignment value for a current-rev DG. The only allowed values are 1 and 8KB; therefore, attempting to change the DG alignment of a CDS disk group will result in an error.

Note Attempting to change the disk group alignment value of a pre-version 110 disk group will result in an error.

For example to set the non-CDS disk group `dg2` to an alignment of 1 block, use a command of the form:

```
# vxchg -g dg2 set alignment=1
```

Then, to set the alignment back to the 8KB value used for CDS disk group compatibility, use a command of the form:

```
# vxchg -g dg2 set alignment=8k
```



This operation fails if you attempt to increase the alignment value, but objects exist which do not already adhere to the new alignment restrictions. Attempts to set the CDS disk group attribute with a disk group alignment less than 8KB also fail.

Increasing Alignment of a non-CDS Version 110 Disk Group

You can increase the alignment of a non-CDS version 110 disk group if all objects meet the new alignment restrictions. If not, you need to run `vxcdsconvert`.

Determining if the CDS Attribute Is Set

You can use the `vxprint` or `vx dg list` commands to determine the state of the CDS attribute, as follows:

```
# vx dg list
NAME                STATE                ID
dgTestSol2         en abled,cds        1063238039.206.vmescl
# vx dg list dgTestSol2
Group:              dgTestSol2
dgid:               1063238039.206.vmescl
import-id:          1024.205
flags:              cds
version:            12 0
alignment:          8192 (bytes)
...

# vxprint -F %cds dgTestSol2
on
```

Joining CDS Disk Groups

Use the `vx dg join` command to join two CDS disk groups. Joining disk groups is allowed only if both disk groups have the same `cds` attribute value. That is, you cannot join a CDS disk group and a non-CDS disk group.

The command results in a disk group alignment for the new (target) disk group that is the minimum of the original (source) and new disk group alignments. The command outputs an informational message if the alignment of the new disk group is reduced (that is, if the original DG has a smaller DG alignment than the new disk group).

Please refer to the *VERITAS Volume Manager Administrator's Guide* for examples of using this command.

Moving Objects Between CDS and Non-CDS Disk Groups

Use the `vxldg move` command to move objects between CDS and non-CDS disk groups. If the new (target) disk group is CDS and original (source) non-CDS, then the original must be 8KB aligned. The disk group alignment of the target disk group is unchanged.

Moving Objects Between CDS Disk Groups

Use the `vxldg move` command to move objects between CDS disk groups. The command results in a disk group alignment for the target disk group is the minimum of the original and target disk group alignments.

The command will output an informational message if the alignment of the target disk group is reduced (that is, if the original disk group has a smaller disk group alignment than the target disk group).

Changing Default CDS Setting for Disk Group Creation

You can change the default CDS attribute setting used in creating disk groups by modifying the `/etc/default/vxldg` file.

Creating Non-CDS Disk Groups

Use the `vxldg` command to create a non-CDS pre-version 110 disk group :

```
# vxldg -T vers init dg disk_name=device_name
```

A pre-version 110 DG is given an alignment value equal to 1.

Note For a pre-version 110 DG, since the `dg_align` value is not stored in the configuration database, it is set to 1 when the disk group is imported.

Upgrading to a CDS Disk Group

You might want to upgrade a non-CDS pre-version 110 disk group to a version 110 disk group in order to use new 4.0 features (except CDS).

Note You must always perform a disk group conversion (using the `vxcdsconvert` utility) to use the CDS feature.



Before Upgrading (Pre-Version 110 Disk Group)

Alignment is always 1 (that alignment is written to the configuration database as part of the upgrade transaction).

After Upgrading (Version 110 Disk Group)

Alignment is still 1 after upgrade but alignment is now stored in configuration database.

Use the `vxdg` command to upgrade the non-CDS pre-version 110 disk group:

```
# vxdg upgrade dg
```

Note The `cds` attribute is set to `off`; and the disk group keeps its 1 alignment.

Replacing a Disk in a CDS Disk Group

To replace a disk in a CDS disk group, use the command sequence:

```
# vxdg -k rmdisk disk_name
# vxdg -k adddisk
```

(Please refer to the *VERITAS Volume Manager Administrator's Guide* for an explanation of the `-k` option.)

For example:

```
# vxdg -k rmdisk disk1
# vxdg -k adddisk disk1=c1t0d1
```

Where `c1t0d1` is the device name of a Solaris disk that is currently being reassigned to disk `disk1`.

Note When replacing a disk in a CDS disk group, you cannot use a non-CDS disk as the replacement.

Setting Device Quotas for CDS Disk Groups

Use the `vxdg set` command to limit the number of devices that can be created in a given CDS disk group:

```
# vxdg set maxdev=max-devices
```

`maxdev` can have any positive integer value greater than the number of devices currently in the disk group. For example:

```
# vxdg -g dg set maxdev=1000
```

Setting DRL Map Size and Log Size

If newly created using `vxassist` without specifying log or map size, the map size is set to a reasonable value, such as for pre-version 110 disk groups.

Use the command line attributes `logmap_len` and `loglen` in conjunction with the `vxassist`, `vxvol`, and `vxmake` commands to set the DRL map and DRL log sizes. The attributes can be used independently, or together.

You can change the DRL map size and DRL log size only when the volume is disabled and the DRL maps are not in use. Changes can be made to the DRL map size only for volumes in a CDS disk group.

The `logmap_len` attribute specifies the required size, in bytes, for the DRL log. It cannot be greater than the number of bytes available in the map on the disk.

If you need to change the size after creating, you need to remove and rebuild the logs, using the following commands:

```
# vxassist -g diskgroup remove log volume nlog=0
# vxassist -g diskgroup addlog volume nlog=nlogs logtype=drl \
  logmap_len=len-bytes [drlloglen=len]
```

Notes

- ◆ If only `logmap_len` is specified, the DRL log size is set to the default value (33*DG alignment).
- ◆ If `logmap_len` is greater than $(DRL\ log\ size)/2$, the command fails, and you need to either provide a sufficiently large `loglen` value or reduce `logmap_len`.
- ◆ For CDS disk groups, the DRL map and log sizes are set to a minimum of $2*(DG\ alignment)$.



Creating a DRL Log

You can create a DRL log using one of the following commands:

- ◆ `vxassist`
- ◆ `vxvol`

Using `vxassist` to Create a DRL Log

When using the `vxassist` command to create a DRL log, it creates logging subdisks equal to the size of the DRL log.

```
# vxassist -g dgTestSol2 make drlvol 100m mirror=2 logtype=drl \  
loglen=264k logmap_len=2048
```

If neither `logmap_len` nor `loglen` is specified, then:

- ◆ `loglen` is set to a default value based on disk group alignment.
- ◆ `maplen` is set to a reasonable value.

If only `loglen` is specified, then:

- ◆ For pre-version 110 disk groups, `maplen` is set to zero.
- ◆ For version 110 and greater disk groups, `maplen` set to use all the bytes available in on-disk map.

If only `logmap_len` is specified (this applies only to disk groups with a version of 110 or greater), `maplen` must be less than number of available bytes in on-disk map for the default log length.

Using `vxvol` to Set a DRL Map Length

You can use `vxvol` only if the volume is stopped (that is, the DRL is inactive).

```
# vxvol -g dgTestSol2 set logmap_len=512 drlvol
```

The `vxvol set` command does not change the existing DRL map size.

Note When specifying a `loglen` attribute, specifying a value less than the minimum required (twice the disk group alignment value) results in an error message.

The value of `loglen` is constrained by size of the logging subdisk.

The value of `logmap_len` is constrained by the log length.

If both `logmap_len` and `loglen` are specified, then, if `logmap_len` is greater than `loglen/2`, `vxvol` fails with an error message, and you then have to either increase `loglen` to a sufficiently large value or decrease `logmap_len` to a sufficiently small value.

The value of `logmap_len` is in units of bytes, and the value of `loglen` is in units of blocks. `logmap_len` cannot exceed the number of bytes in the on-disk map. The number of bytes in the on-disk map can be calculated as follows:

$$[\text{ROUND_DOWN}(\text{loglen}/n\text{maps}) - \text{VOLDRL_HEADER_SIZE}]$$

where `nmaps` is 2 for a private disk group or 33 for a shared disk group, `ROUND_DOWN()` represents rounding down to a log map alignment boundary, and the value of `VOLDRL_HEADER_SIZE` is 24.

Displaying Information

This section describes:

- ◆ [Displaying Traditional DRL Map Size and Log Size](#)
- ◆ [Displaying Disk Group Alignment](#)
- ◆ [Displaying Volume Log Map Values](#)
- ◆ [Listing Offset and Length Information](#)
- ◆ [Listing CDS Disk Groups](#)
- ◆ [Listing Disks](#)
- ◆ [Displaying Device Quotas for CDS Disk Groups](#)

Displaying Traditional DRL Map Size and Log Size

Use the `vxprint` command to display the map length and map alignment of traditional DRL logs:

```
# vxprint -g dg1 -vl drlvols
# vxprint -g dg1 -vF '%name %logmap_len %logmap_align' drlvols
```

Displaying Disk Group Alignment

Utilities that print information regarding a disk group record (such as `vxprint` and `vxvg list`) can also output the disk group alignment. For example, to print information for disk group `dg1`, you can use the command:

```
# vxprint list dg1
```



Displaying Volume Log Map Values

Use the `vxprint` command to determine the log map alignment and log map length values. For example, to print information for disk group `dg1`, you can use the command:

```
# vxprint [ -g dg1 ] -lv volname
```

Output is of the form:

```
logging: type=REGION loglen=0 serial=0/0 mapalign=0 maplen=0
(disabled)
```

This indicates a log map alignment (`logmap_align`) value of 0, and a log map length (`logmap_len`) value of 0.

If the log map is set and enabled, the command and results may be of the form:

```
# vxprint -lv drlvol
Disk group: dgTestSol

Volume:  drlvol
info:    l en=20480
type:    u setype=fsgen
state:   state=ACTIVE kernel=ENABLED cdsrecovery=0/0 (clean)
assoc:   plexes=drlvol-01,drlvol-02,drlvol-03
policies: read=SELECT (round-robin) exceptions=GEN_DET_SPARSE
flags:   closed writecopy writeback
logging: type=REGION loglen=528 serial=0/0 mapalign=16 maplen=512
(enabled)
apprecov: seqno=0/0
recovery: mode=default
recov_id=0
device:  minor=46000 bdev=212/46000 cdev=212/46000
path=/dev/vx/dsk/dgTestSol/drlvol
perms:   user=root group=root mode=0600
guid: {d968de3e-1dd1-11b2-8fc1-080020d223e5}
```

Listing Offset and Length Information

Use the `vxprint` command with the `-b` option to output all offset and length information in units of 512 bytes. This enables consistent output from different platforms. Without the `-b` option, this information is output with a value of sectors that may differ between platforms.

When `vxprint -bm` is used, the output also contains the `b` suffix, so that the output can be fed back to `vxmake`.

Listing CDS Disk Groups

You can use the following commands to list all disk groups; then look for flag `cds` in the listing to locate CDS disk groups.

```
# vxdbg list
# vxprint -G1
```

Listing Disks

Use the `vxdisk list` command with the `-b` option to output all offset and length information in units of 512 bytes. This enables consistent output from different platforms. Without the `-b` option, this information is output with a value of sectors that may differ between platforms.

When `vxdisk list -bm` is used, the output also contains the `b` suffix, so that the output can be fed back to `vxmake`.

The disk format is put in the `TYPE` column along with the `auto` type.

Displaying Device Quotas for CDS Disk Groups

Use the `vxprint -lG` command to display the maximum number of devices in a CDS disk group. For example, to display the maximum number of devices in the CDS disk group `cdsdg1`, you can use a command of the form:

```
# vxprint -lG cdsdg1
```

The output is similar to the following:

```
group:dg1
info:dgid=1027365608.1064.vmesec6
version: 110
alignment: 8192 (bytes)
activation: read-write
detach-policy: global
copies: nconfig=default nlog=default
devices: max=32000 cur=0
minors: >= 16000
cds=on
```

The value of `max` on the `devices` line indicates the maximum number of devices. In this example, the maximum number of devices is 32,000.

You can also use the following command to obtain the maximum number of devices:

```
# vxprint -g dg1 -GF %maxdev
32000
```



Default Activation of Shared Disk Groups

Default activation of shared disk groups involves a local in-kernel policy that differs between platforms. This means that, regardless of the platform the disk group was created on, the importing platform will have platform-specific behavior with respect to activation of shared disk groups. Specifically, with the exception of HP-UX, importation of a shared disk group results in the volumes being active and enabled for shared-write. In the case of HP-UX, the shared volumes will be inactive and require other actions to activate them for shared-write operations.

Creating CDS Disk Groups Using `vxdg split`

You can use the `vxdg split` command to create a CDS disk group from an existing CDS disk group. The new (target) and original (source) disk group retain the same CDS attribute; that is, if the source is CDS, then the new disk group is also CDS. The new disk group also has the same alignment as the original disk group.

Default Activation of Shared Disk Groups

This is a local in-kernel policy that differs between platforms. This means that, regardless of the platform the disk group was created on, the importing platform will have platform-specific behavior with respect to activation of shared disk groups. Specifically, with the exception of HP-UX, importation of a shared disk group will result in the volumes being active and enabled for shared-write. In the case of HP-UX, the shared volumes will be inactive and require other actions to activate them for shared-write operations.

Block Size

The block size on HP is different from that on other supported platforms, so that output without the `-b` option from commands such as `vxdisk` and `vxprint` looks different on HP for the same disk group.

Alignment Value

The alignment value 1 refers to an HP-UX block size of 1024 bytes.





Importing Disk Groups Between Linux and Non-Linux Machines

A disk group created on non-Linux platforms typically has device numbers above 1000. When that disk group is imported on a Linux machine, the devices are reassigned minor numbers below 256.

If this disk group is now imported to a non-Linux machine, all device numbers will be less than 256. If those devices are available (that is, they do not conflict with devices in an imported boot disk group) they will be used. Otherwise new device numbers will be reassigned.

A single disk group could contain a number of devices exceeding the maximum number of devices for a given platform. In this case, the disk group cannot be imported on that platform because importation would exhaust available minor devices for the VxVM driver. Although the case of minor number exhaustion is possible in a homogeneous environment, it will be more pronounced between platforms with different values for the maximum number of devices supported, such as Linux. This difference will render platforms with low maximum devices supported values less useful as heterogeneous disk group failover or recovery candidates.

Note Using the disk group `maxdev` attribute may reduce the likelihood that a CDS disk group import on Linux will exceed the maximum number of devices.

Alignment Value

The alignment value 1 refers to a block size of 512 bytes on Linux.





Disk Group alignment and Encapsulated Disks

On Solaris, all native file systems are cylinder aligned. Encapsulating such a disk results in subdisks that are also cylinder aligned. Such alignment will normally not be 8K aligned, but it will be 1K aligned. For the encapsulation process, there is no flexibility as to where on the disk the subdisks must be since the data location is pre-defined. If an alignment conflict occurs, user intervention is required. If the disk group alignment is 8K this operation will probably fail because this would require the cylinder to be an even number of 8K blocks in size.

Alignment Value

The alignment value 1 refers to a block size of 512 bytes on Solaris.





Error Codes and Error Recovery



Error Codes

Error Number	Message	Action
329	Cannot join a non-CDS disk group and a CDS disk group	Change the non-CDS disk group into a CDS disk group (or vice versa), then retry the join operation.
330	Disk group is for a different platform	Import the disk group on the correct platform. It cannot be imported on this platform.
331	Volume has a log which is not CDS compatible	To get a log which is CDS compatible; you need to stop the volume, if currently active, then start the volume. After the volume has been successfully started, retry setting the CDS attribute for the disk group.
332	License has expired, or is not available for CDS	Obtain a license from VERITAS that enables the usage of CDS disk groups.
333	Non-CDS disk cannot be placed in a CDS disk group	Do one of the following: <ol style="list-style-type: none">1. Add the disk to another disk group that is a non-CDS disk group.2. Re-initialize the disk as a CDS disk so that it can be added to the CDS disk group.3. Change the CDS disk group into a non-CDS disk group and then add the disk.
334	Disk group alignment not CDS compatible	Change the alignment of the disk group to 8K and then retry setting the CDS attribute for the disk group.
335	Subdisk length violates disk group alignment	Ensure that sub-disk length value is a multiple of 8K.
336	Subdisk offset violates disk group alignment	Ensure that sub-disk offset value is a multiple of 8K.



Error Number	Message	Action
337	Subdisk plex offset violates disk group alignment	Ensure that sub-disk plex offset value is a multiple of 8K.
338	Plex stripe width violates disk group alignment	Ensure that plex stripe width value is a multiple of 8K.
339	Volume or log length violates disk group alignment	Ensure that volume length value is a multiple of 8K. For volume log length; if you are using <code>vxassist</code> to create the volume, then you should set <code>dalign_checking</code> to <code>round</code> . This will ensure that the log length is silently rounded to a valid value.
340	Last disk media offset violates disk group alignment	Reassociate the DM record prior to upgrading.
341	Too many device nodes in disk group	Increase the number of device nodes allowed in the disk group, if not already at the maximum. Otherwise, you need to remove volumes from the disk group, possibly by splitting the disk group.
342	Map length too large for current log length	Use a smaller map length for the DRL/DCM log, or increase the log length and retry.
343	Volume log map alignment violates disk group alignment	Remove the DRL/DCM log, then add it back after changing the alignment of the disk group.
345	Disk group contains an old-style RVG which cannot be imported on this platform	Import the disk group on the platform that created the RVG. To import the disk group on this platform, first remove the RVG on the creating platform.
346	Cache object autogrow by <code>max_autogrow</code> violates disk group alignment	Ensure that cache attribute value is a multiple of 8K.
347	User transactions are disabled for the disk group	Retry the command as it was temporarily disallowed by the <code>vxcdsconvert</code> command executing at the same time.
348	Disk is in use	Contact VERITAS Technical Support.



Glossary

AIX coexistence label

Data on disk which identifies the disk to the AIX volume manager (LVM) as being controlled by VxVM. The contents has no relation to VxVM ID Blocks.

back-rev disk group

A disk group created using a version of VxVM released prior to the release of CDS. Adding CDS functionality rolls over to the latest disk group version number; see also current-rev disk group.

CDS (Cross-platform Data Sharing)

Sharing data between heterogeneous systems (such as SUN and HP), where each system has direct access to the physical devices used to hold the data, and understands the data on the physical device. Sharing in this sense should not be confused with the sharing provided with CVM by means of a shared disk group.

CDS disk

A disk whose contents and attributes are such that the disk can be used for CDS as part of a CDS disk group. In contrast, a non-CDS disk cannot be used for CDS, nor can it be part of a CDS disk group. CDS disk also contains a set of AIX Coexistence Labels, HP-UX Coexistence Labels/VxVM ID Blocks, and Platform Blocks.

CDS disk group

A VxVM disk group whose contents and attributes are such that the disk group can be used to provide CDS. In contrast, a non-CDS disk group (that is, a back-rev disk group or a current-rev disk group) cannot be used for CDS. A CDS disk group is a current-rev disk group with the CDS attribute set for the disk group. A CDS disk group can only contain CDS disks.

CFS

Cluster file system. A VxFS file system mounted on a selected volume in cluster (*shared*) mode.



children

Objects that belong to an object group.

cluster

A set of host machines (*nodes*) that shares a set of disks.

cluster file system

See CFS.

current-rev disk group

A disk group created using a version of VxVM providing CDS functionality; however, the CDS attribute is not set. If the CDS attribute is set for the disk group, the disk group is called a CDS disk group.

data change object

See DCO.

DCO (Data Change Object)

A VxVM object that is used to manage information about the FastResync maps in the DCO volume. Both a DCO object and a DCO volume must be associated with a volume to implement Persistent FastResync on that volume.

device name

The physical disk device name (or *disk access name*).

dirty region logging

See DRL.

disk group

A set of disks that are under VxVM control and share a common configuration. A disk group configuration is a set of records containing detailed information on existing VERITAS Volume Manager objects (such as disk and volume attributes) and their relationships. Each disk group has an administrator-assigned name. Volumes can only be created on disks that belong to disk groups.

disk name

The VERITAS Volume Manager disk name (or *disk media name*). This is the logical name given to the disk by VxVM or the administrator.



DRL (Dirty Region Logging)

The method by which the VxVM monitors and logs modifications to a plex as a bitmap of changed regions. For a volumes with a new-style DCO volume, the dirty region log is maintained in the DCO volume. Otherwise, the dirty region log is allocated to an associated subdisk called a *log subdisk*.

enclosure

A disk array.

gap

A disk region that does not contain VERITAS Volume Manager objects (subdisks).

HP-UX coexistence label

Data on disk which identifies the disk to the HP volume manager (LVM) as being controlled by VxVM. The contents of this label are identical to the contents of the VxVM ID block.

mirror

A copy of a volume and its data. There can be several mirrors per volume. The terms *mirror* and *plex* are used synonymously.

node

In the VxVM tree, a node is an element attached to the tree.

In a cluster environment, a node is a host machine in a cluster.

object group

A group of objects of the same type. Each object group has a group icon and a group name. In VxVM, object groups include disk groups, disks, volumes, controllers, free disk pool disks, uninitialized disks, and file systems.

object tree

A dynamic hierarchical display of VERITAS Volume Manager objects and other objects on the system. Each node in the tree represents a group of objects of the same type.

platform block

Data placed in sector 0, which contains OS-specific data for a variety of platforms that require its presence for proper interaction with each of those platforms. The platform block allows a disk to masquerade as if it was initialized by each of the specific platforms.



plex

A copy of a volume and its data. There can be several plexes per volume. The terms *mirror* and *plex* are used synonymously.

sector size

Sector size is an attribute of a disk drive (or SCSI LUN for an array-type device), which is set when the drive is formatted. Sectors are the smallest addressable unit of storage on the drive, and are the units in which the device performs I/O.

subdisk

A set of contiguous disk blocks that form a logical disk segment. Subdisks are associated with plexes (mirrors) to form volumes.

uninitialized disks

Disks that are not under VxVM control.

volume

A virtual disk or entity that is made up of portions of one or more physical disks.

VxFS

VERITAS File System.

VxVM

VERITAS Volume Manager.

VxVM ID block

Data on disk that indicates the disk is under VxVM control. The VxVM ID Block provides dynamic VxVM private region location, GUID, and other information.



Index

A

- access type 4
- activation
 - default 29
- AIX coexistence label 4
- alignment 6
 - changing 21
 - reducing 21
 - setting 20
- alignment value 7
- attribute
 - CDS 22
 - logmap_len 8
- auto disk type 4

B

- block size 2

C

- CDS
 - attribute 22
 - changing setting 23
 - creating DGs 11
 - creating disks 10
 - disk group alignment 2
 - disk group device quotas 5
 - disks 2
 - log 7
 - map 7
- CDS disk groups
 - alignment 22, 27
 - increasing alignment 22
 - joining 22
 - moving 23
 - reducing alignment 21
 - setting alignment 21
- CDS disks
 - creating 9
- changing CDS setting 23

- changing default CDS setting 23
- changing default disk format setting 18
- changing disk format 18
- co-existence label 4
- concepts 1
- converting non-CDS disks to CDS 11
- converting non-CDS disks to CDS disks 12
- creating a DRL log 26
- creating CDS disk groups 11
- creating CDS disks 9, 10
- creating DRL logs 26
- creating non-CDS disk groups 23
- creating pre-version 110 disk groups 23
- current-rev disk groups 6

D

- DCO
 - map layout 8
- default activation 29
- default CDS setting
 - changing 23
- defaults file 12, 16
- device quota 5
- device quotas 5, 29
 - displaying 29
 - setting 24
- disk
 - access type 4
 - change format 18
 - labels 18
 - LVM 18
 - replacing 24
- disk access 2
- disk format 2
- disk group 4
 - alignment 6
 - non-CDS 6
- disk group alignment 21
- displaying 27



- disk groups
 - creating 23
 - joining 22
 - upgrading 23
- disk quotas
 - setting 24
- disk types 3
- displaying device quotas 29
- displaying disk group alignment 27
- displaying DRL log size 27
- displaying DRL map size 27
- displaying log map values 27
- displaying log size 27
- displaying v_logmap values 27, 28
- displaying volume log map values 27
- DRL log
 - creating 26
- DRL log size
 - displaying 27
 - setting 25
- DRL logs
 - creating 26
- DRL map length 26
- DRL map size
 - displaying 27
 - setting 25

E

- encapsulation 18, 20

I

- I/O block size 2
- ID block 4
- increasing alignment 22

J

- joining CDS disk groups 22
- joining disk groups 22

L

- length listing 28
- licensing 15
- listing disk groups 28
- listing disks 29
- listing offset and length information 23
- log size
 - displaying 27
 - setting 25
- logmap_len attribute 8
- LVM disks 18

M

- minor device numbers 3
- moving CDS disk groups 23
- moving disk group objects 23

N

- non-CDS disk groups
 - alignment 20
 - setting alignment 20

O

- objects
 - moving 23
- offset
 - listing 28
- offset information 28
- operating system data 2

P

- platform block 4
- private region 3
- public region 3

R

- reducing alignment 21
- reducing alignment of CDS disk groups 21, 22
- replacing disks 24
- restoring CDS disk labels 18
- restoring disk labels 18

S

- setting CDS disk group alignment 21
- setting device quotas 24
- setting disk quotas 24
- setting DRL log size 25
- setting DRL map length 26
- setting DRL map size 25
- setting log size 25
- system
 - setting up 9

U

- upgrading disk groups 23
- upgrading pre-version 110 disk groups 23

V

- v_logmap
 - displaying 27, 28
- VRTExplorer x
- vxcdsconvert 12
- vx dg init 11
- vx dg split 30



vxdiskadm 10, 11
vxdisksetup 9
VxVM

devices 1
obtaining system information x
vxvol 26

