

===[ man1m/vxapslice.1m ]=====

Maintenance Commands

vxapslice(1M)

#### NAME

vxapslice - free a slice on a disk for an AP database

#### SYNOPSIS

```
/etc/vx/bin/vxapslice [-g diskgroup] [-o offset] [-p parti-  
tion] [-s size] accessname | medianame
```

#### DESCRIPTION

The vxapslice utility removes an area of a disk from VxVM control, allowing that area to be used as the location for an AP database. There must be sufficient free space (that is, not used by any subdisk) available. A partition is created that corresponds to the area and this partition is used to create and access the AP database (see apdb(1M)).

A subdisk is also created that corresponds to the area, and is tagged as being in use as an APDB. This subdisk is not usable by VxVM, and cannot be moved or associated.

accessname is a system-dependent disk device, while medianame is a disk media name.

#### OPTIONS

-g diskgroup

Use when accessname is a part of a disk group. diskgroup can be either a group name or group ID number.

-o offset Specifies an offset into the public region to use for the partition. This value may be rounded up for alignment purposes.

-p partition

Specifies a partition number to use. This may have a value from 0 to 7.

-s size Specifies a minimum size for the partition. Default is 300K. Size may be rounded up for alignment purposes.

#### EXIT CODES

On successful completion, an exit status of zero is returned. Additionally, the access path of the created partition is displayed. On an error, -1 is returned.

#### SEE ALSO

apdb(1M), vxintro(1M)

Alternate Pathing 2.0 User's Guide

===[ manlm/vxassist.1m ]=====

Maintenance Commands

vxassist(1M)

#### NAME

vxassist - create, layout, convert, mirror, backup, grow, shrink, delete, and move volumes

#### SYNOPSIS

vxassist [ options ] [ -b ] addlog volume [ attribute...]

vxassist [ options ] [ -b ] convert volume [ attribute...]

vxassist [ options ] [ -b ] growby volume lengthchange [ attribute...]

vxassist [ options ] [ -b ] growto volume newlength [ attribute...]

vxassist help [ alloc | attrs | layout | mirror | options | showattrs | space | usage]

vxassist [ options ] [ -b ] make volume length [ attribute...]

vxassist [ options ] [ -b ] maxgrow volume [ attribute...]

vxassist [ options ] [ -b ] maxsize [ attribute...]

vxassist [ options ] [ -b ] mirror volume [ attribute...]

vxassist [ options ] [ -b ] move volume storage-spec... [ attribute...]

vxassist [ options ] [ -b ] relay layout volume [ attribute...]

vxassist [ options ] remove {volume|mirror|log } volume [storage-spec...] [ attribute...]

vxassist [ options ] shrinkby volume lengthchange

vxassist [ options ] shrinkto volume newlength

```
vxassist [ options ] snapabort volume
vxassist [ options ] snapshot volume newvolume
vxassist [ options ] [ -b ] snapstart volume [ attribute... ]
vxassist [ options ] snapwait volume
```

#### DESCRIPTION

The vxassist utility is an interface to the VERITAS Volume Manager that finds space for and creates volumes, performs volume conversion, adds mirrors and logs to existing volumes, extends and shrinks existing volumes, provides for

VxVM 3.0

Last change: 15 Dec 1999

1

Maintenance Commands

vxassist(1M)

the migration of data from a specified set of disks, and provides facilities for the on-line backup of existing volumes.

The vxassist command is supplied a keyword that selects the action to be performed.

Each operation can be applied to only one disk group at a time, due to internal implementation constraints. For the make operation, attributes that name specific disks may be used to determine a default disk group, according to the standard disk group selection rules described in vxintro(1M). For other operations, the volume operand is used. A specific disk group can be forced with -g disk-group. When no disk attributes are specified, the make operation defaults to using the rootdg disk group. An alternate default disk group can be specified in a defaults file (/etc/default/vxassist).

Many vxassist operations can take a set of attributes that specify how volumes should be laid out, which sets of disks they should be built on, among other things. Attributes are of two basic types: storage specifications, and attribute settings. Storage specification attributes are either simple disk names (for example, disk01) or selected groups of disks (for example, ctrl:c1 specifies all disks on controller 1). Excluded (negated) storage specification can be specified with a prefix of !. For example, the arguments ctrl:c1 and !target:c1t5 can be combined to indicate that storage should be allocated from controller 1, but not from

SCSI target 5 on that controller.

Other attributes are of the form attrname=value; these additional attributes can specify the type of a volume (mirrored, RAID-5, striped, logged), layout policies (contiguous, spanning), mirroring requirements (mirror across controllers or SCSI targets), constraint parameters (constrain allocations to a single controller), and more. A complete list of attribute specifications is given later in this man page.

If no non-excluded storage specification attributes are given, then any non-reserved, non-volatile, non-spare, non-excluded disk may be used to allocate storage. Attributes may constrain the selection, particularly with respect to the selection of multiple disks. For example, the command:

```
vxassist make mirvol 500m layout=mirror,log mirror=ctrl
!ctrl:c2
```

requests that vxassist create a new mirrored volume on any disks that are not on controller 2. The selection of disks

#### Maintenance Commands

vxassist(1M)

is constrained by the mirror=ctrl attribute such that no disks within a mirror can be on the same controller as any disks on the other mirror.

#### KEYWORDS

addlog Add a log to a mirrored or RAID-5 volume. A log can be added to a mirrored volume only if the logging type of the volume is DRL (this is the default logging type). When adding the first log to a mirrored or RAID-5 volume, the loglen attribute can be used to specify the length of the log; otherwise, the loglen attribute will be ignored. Other attributes can be specified to indicate desired properties for the new allocations. Mirroring constraints do not apply to logs, though storage constraints can be specified that constrain the logs to disk sets already used by the volume.

Adding a log to a DRL mirrored volume involves creating a single log subdisk and a new plex to

contain that subdisk. The new plex will then be attached to the volume. Adding a log to a RAID-5 volume involves creating a new plex that will be attached to the volume as a log plex.

`convert` Converts a volume layout to a mirror-stripe from a stripe-mirror, or from a mirror-stripe to a stripe-mirror. Also converts to mirror-concat from a concat-mirror, or from a mirror-concat to a concat-mirror.

`growby` | `growto`  
Increase the length of the named volume to the length specified by `newlength` (`growto`), or by the length specified by `lengthchange` (`growby`). The new length or change in length is specified as a standard Volume Manager length (see `vxintro(1M)`). The `growto` operation fails if the new length is not greater than the current volume length. The length of the volume is increased by extending existing subdisks in the volume, or by adding and associating new subdisks. Plexes that are not fully populated already (that is, that are sparse) are left unchanged. Log-only plexes are also left unchanged.

When a volume is grown, the volume may be converted from one layout to another as a side effect. For example, a volume that has the mirror-stripe layout may be converted to stripe-mirror when it is grown. This kind of conversion

is done if `vxassist` determines that the new volume is too small or large for the original layout. If the new volume layout is inappropriate, the volume can be converted back using the `vxassist convert` command. See Layout Specifications for a description of the methods used to determine disk layouts.

Attributes can be specified to indicate various desired properties for the new allocations. Any mirroring constraints will still apply between all extensions to the existing mirrors and the other mirrors. Growing of a volume requires that the

volume be enabled.

No attempt will be made by vxassist to adapt the file system that resides on the volume (or other users of the volume). It is expected that any necessary application notifications will be made after the grow operation completes successfully.

- |         |   |
|---------|---|
| help    | Displays information on vxassist usage and VxVM attributes.   |
| make    | Create a volume with the specified name and the specified length. The length is specified as a standard Volume Manager length (see vxintro(1M)). Attributes can be specified to indicate various desired properties and storage locations for the created volume. See Layout Specifications for a description of the methods used to determine disk layouts.  |
| maxgrow | Reports the maximum size a volume could grow given its attributes and given the free storage available.   |
| maxsize | Reports the maximum size a volume could be created with given the attributes specified and given the free storage available.  |
| mirror  | Create a new mirror (or plex) and attach it to the volume. This operation is allowed only if the volume is enabled. Attributes can be specified to indicate various desired properties and storage locations for the created volume. Attributes that constrain mirroring (such as requiring that mirrors be on separate groups of disks) apply between the existing mirrors of the volume and the new mirror. |
| move    | Move subdisks within the named volume off the   |

excluded storage specified on the command line. Excluded storage is specified with a prefix of !. The move operation requires that at least one excluded storage specification be given.

If the volume is enabled, then subdisks within detached or disabled plexes (or detached log or RAID-5 subdisks) will be moved without recovery of data. If the volume is not enabled, then stale log or RAID-5 subdisks, or subdisks within STALE or OFFLINE plexes, will be moved without recovery; if there are other subdisks within a non-enabled volume that require moving, then the move operation will fail.

For enabled subdisks in enabled plexes within an enabled volume, the data within subdisks will be moved to the new location, without loss of availability (or redundancy) of the volume.

relayout Change a volume layout or properties. This operation changes the number of columns in a stripe or the stripe width of a volume. It also converts a volume to or from RAID-5, mirrored, concat-mirrored, striped, or any similar layout.

During the relayout process, a volume may also be converted into an intermediate layout. For example, to convert from a 4-column mirror-stripe to a 5-column mirror-stripe, the volume is first changed to a stripe-mirror before the new column is added.

To change volume voll from a 4-column mirror-stripe to a 5-column mirror-stripe, enter:

```
vxassist relayout voll ncol=5
vxassist convert voll layout=mirror-stripe
```

To change volume voll from a 4-column stripe-mirror to a 5-column stripe-mirror, enter:

```
vxassist relayout voll ncol=5
```

If you view the volume with vxprint during this layout conversion, the volume shows several subvolumes that are used during the change, and a more complex configuration than usual.

After the new column is added the volume remains

in the stripe-mirror layout. You must use vxassist convert to convert back to a mirror-stripe layout.

Note: If the system crashes during a relayout operation, the process continues. However, if the relayout was a two-stage operation, you must run a vxassist convert command to convert back to the original layout.

remove      Depending on the next keyword (volume, mirror, or log), vxassist deletes the entire volume, one or more mirrors, or one or more logs. When deleting a mirror or a log, the storage to be removed can be specified with the ! attribute. The alloc=, nmirror=, and nlog= also determine the number of mirrors or logs remaining after the remove operation completes. When the number of mirrors drops to 1, logs are also removed. By default, vxassist removes 1 mirror or 1 log.

shrinkby | shrinkto  
Decrease the length of the named volume to the length specified by newlength (shrinkto), or by the length specified by lengthchange (shrinkby). The new length or change in length is specified as a standard Volume Manager length (see vxintro(1M)). The shrinkto operation fails if the new length is not less than the current volume length.

The length of a volume is decreased by removing and shortening subdisks to leave each plex with the desired volume length. The freed space can then be allocated for use by other volumes. Log-only plexes are left unchanged.

When a volume is shrunk, the volume may be converted from one layout to another as a side effect. For example, a volume that has the mirror-stripe layout may be converted to stripe-mirror when it is shrunk. This kind of conversion is done if vxassist determines that the new volume is too small or large for the original layout. If the new volume layout is inappropriate, the volume can be converted back using the vxassist convert command. See Layout Specifications for a description of the methods used to determine disk layouts.

No attempt will be made by vxassist to adapt the file system that resides on the volume (or other



## Maintenance Commands

vxassist(1M)

users of the volume). It is expected that any necessary application notifications will be made before the shrink operation is initiated.

**snapabort** When a **snapstart** mirror has been started (and eventually completed), the **snapabort** can be used to remove the mirror snapshot.

**snapstart** and **snapshot**

Create a temporary mirror and attach it to the named volume. When the attach completes, the mirror will be considered a candidate for selection by the snapshot operation. The snapshot operation takes one of these attached temporary mirrors and creates a new volume with the temporary mirror as its one plex. Attributes can be specified to indicate desired properties of the snapshot mirror.

Some usage types will attempt to synchronize any in-memory data associated with the volume (such as unwritten file system modifications) when the snapshot operation is done. In particular, if the **fsgen** usage type is used with a volume containing a **VxFS** file system, then cooperating procedures ensure that all file system data is consistently flushed to the volume. For **ufs** and **s5**, the synchronization operation consists of a call to **sync**, which will make the snapshot a better image, but which may leave some inconsistencies between in-memory file system data and the data residing on the backup image.

**snapwait** If a **snapstart** mirror attach is done as a background task (such as using the **-b** option), it may be convenient to wait for an attached mirror to become available. The **snapwait** operation waits for such an attach to complete on the named volume. When a snapshot attach has completed, the operation exits.

## OPTIONS

**-b** Perform extended operations in the background. This applies to plex consistency recovery operations for **vxassist make**, **growto**, and **growby**. This flag also applies to plex attach operations started by **vxassist mirror**, **vxassist snapstart**,

and to Relayout operations started by vxassist  
Relayout.

-d defaults  
Specify a file containing defaults for various

VxVM 3.0

Last change: 15 Dec 1999

7

## Maintenance Commands

vxassist(1M)

attributes related to volume creation and space  
allocation. If not specified, this defaults to  
/etc/default/vxassist.

-f Forces operations not usually allowed by vxassist.

-g diskgroup  
Specify the disk group for the operation, either  
by disk group ID or by disk group name. By  
default, the disk group is chosen based on the  
medianame operands (if any) for the vxassist make  
operation, or based on the volume operands for all  
other operations.

-n Do not read the system defaults file.

-o useopt Pass in usage-type-specific options to the operation. A certain set of operations are expected to be implemented by all usage types:

slow[=iodelay]  
Reduce the system performance impact of  
copy operations. Copy and plex consistency  
recovery operations are usually a set of short  
operations on small regions of the volume  
(normally from 16 kilobytes to 128 kilobytes).  
This option inserts a delay between the  
recovery of each such region. A specific  
delay can be specified with iodelay as a  
number of milliseconds, or a default is  
chosen (normally 250 milliseconds).

iosize=size  
Perform copy and recovery operations in  
regions with the length specified by size,  
which is a standard Volume Manager length  
number (see vxintro(1M)). Speci-

fyng a larger number typically causes the operation to complete sooner, but with greater impact on other processes using the volume. The default I/O size is typically between 32 and 128 kilobytes.

- p For use with maxsize and maxgrow. Prints only the maximum size with no text wrappers.
- r Tells vxassist to include the spare disks (the disks reserved for relocation) in the calculation.

VxVM 3.0

Last change: 15 Dec 1999

8

Maintenance Commands

vxassist(1M)

- t tasktag  
If any tasks are registered to track the progress of an operation, mark them with the tag tasktag. This option is passed to utilities called by vxassist, so any child tasks are also tagged with tasktag.
- U usetype  
Limit the operation to apply to this usage type. Attempts to affect volumes with a different usage type will fail. For a vxassist make operation, this indicates the usage type to use for the created volume. Otherwise, the default is used and is determined by the existence of an entry in the /etc/default/vxassist file or else is set to the fsgen usage type.
- v  
Traces calls to other utilities to determine which VxVM commands vxassist is executing.

#### Hardware-specific Options

Some environments provide guidelines to optimize the Volume Manager's interaction with intelligent storage systems. If these guidelines are present, the Volume Manager follows the guidelines when creating volumes or allocating space for volumes. By default, vxassist only creates volumes that conform with these guidelines. The following options change the behavior of vxassist make:

- o override  
Creates the specified volume and ignores any

storage-specific guidelines. Overriding the guidelines is not recommended as it can result in incompatible objects, or objects that cannot be administered by some utilities.

-o verify Verifies that the specified volume can be created without violating any storage-specific guidelines, but does not create the volume. If any guidelines are violated, vxassist exits with an error message.

#### ATTRIBUTES

Attribute values for various purposes can be specified with arguments of the form attribute=value. Attributes can also be passed in through a defaults file. Default attribute values can be stored in the file /etc/default/vxassist.

Attributes are selected according to the order in which they are scanned. In general they are taken in decreasing priority of being specified on:

#### Maintenance Commands

vxassist(1M)

1. The command line.
2. The specified defaults file (as supplied with the -d command line argument).
3. The system defaults file (as specified in /etc/default/vxassist).

Attributes from all sources have the same form. However, in some cases, command-line attributes change default behaviors in ways that defaults-file supplied attributes do not. In particular, references to mirroring (such as specifying a mirror count) or logging (such as specifying a log count or length) on the command line will cause mirroring or logging to happen by default. If such attributes are specified in a defaults file, then they just indicate the attributes that would be used if mirroring or logging were enabled.

Attributes are either storage specifications (possibly negated), or are in the form attribute=value. In a defaults file, attributes are separated by a space or are on separate lines. Blank lines in a defaults file are ignored, and comments can be included with the standard # convention.

## Storage Specifications

Storage specification attributes have one of the following forms:

[!]diskname

Specify (or exclude) the given disk. diskname refers to a disk media record name in a Volume Manager disk group.

[!]diskclass:instance

Specify (or exclude) a set of disks. diskclass names a particular type of disk grouping (for example, ctrl to indicate groups of disks on a single controller), and instance specifies which grouping of that type (for example c1). Each type of disk class has a particular format for specifying instances of the class.

diskclass:same

Specify that the allocation must be constrained to include disks from the same instance of the given class. For example, a storage specification of diskparam:same indicates that all disks selected must have the same basic parameters (tracksize, cylinder size, size, and revolutions-per-minute), but does not specify which specific parameters to use.

## Maintenance Commands

vxassist(1M)

The defined disk classes (and alternate names, or aliases, for those classes) are as follows:

ctrl (alias c, ctrl, cntrl, controller)

Specify disks from a particular controller. The controller is identified in the form c#, where # is a decimal controller number starting at 0.

ctype (alias ctrlrtype, ctrlrtype, controllertype)

Specify disks from a particular type of controller. The controller type is usually a driver name. For example, on Sun machines, the regular SCSI controller type can be identified as

ctype:esp. As a special case, the SPARCstorage Array controller type is specified as ctype:ssa.

da Specify a specific Volume Manager disk by disk access record name. An example of a disk access record name is c2t1d0s2, which indicates a special Volume Manager disk defined on slice 2 of c2t1d0s2.

device (alias d) Specify all Volume Manager disks on a specific physical disk. The physical disk is specified in the form c#t#d#, which indicates the controller, target, and disk numbers. Typically, only one Volume Manager disk device is created for each physical disk. However, the vxdisk define operation (see vxdisk(1M)) can be used to create additional Volume Manager disk devices on selected partitions.

diskgeom (alias geom) Specify disks with a particular geometry. The geometry is specified in the form cyls.heads.sectors, to indicate the number of cylinders, heads, and sectors per track, of the disk. Many modern drives have variable geometries (or geometries that do not fit the standard conventions). For such drives, the geometry parameters used are whatever the drive reports.

diskparam Specify disks with particular parameters. The parameters are specified in the form cyls.heads.sectors.rpm, to indicate cylinders, heads, sectors per track, and the number of revolutions per minute of the drive. As with the diskgeom parameter, the instances of the diskparam disk class depend upon the values reported by the

drive, and may not be accurate.

diskrpm (alias rpm) Specify disks with a particular number of revolutions per minute. The revolutions per minute is given as a simple decimal number.

dm (alias disk)  
Specify a Volume Manager disk by disk media record name, in the disk group of the volume. Specifying a disk class type of dm or disk is equivalent to giving a storage specification with no disk class name.

probe\_granularity  
The granularity is a result of the maxsize request. The default value is 1M.

target (alias t)  
Specify disks on the same SCSI target address in the same controller. The target is specified in the form c#t#. For example, t:c2t4 selects disks on target 4 of controller 2.

tray  
Specify disks in the same removable tray (for SPARCstorage Array controllers). Trays in the SPARCstorage Array are grouped as two SCSI targets per tray (for example, targets 0 and 1 are in the same tray). Trays are specified in the form c#tray#, where c# identifies a particular SPARCstorage Array controller and tray# is the number of a tray (0 for SCSI targets 0 and 1, 1 for targets 2 and 3, or 2 for targets 4 and 5). For example, c2tray2 selects disks in tray 2 (targets 4 and 5) of controller 2.

The diskgeom and diskparam attributes are most often useful in the form diskgeom:same and diskparam:same, to indicate that all disks used for creating a volume should be on similar disk drives.

#### Other Attributes

Other attributes are of the form attribute=value. The attribute name in an attribute value pair will never contain a colon, so it is possible to specify a disk that has an equal sign in its name using the storage specification dm:disk01=a. Without the dm: prefix, disk01=a would yield an error indicating an unrecognized attribute.

Defined attributes (and common aliases) are:

alloc=storage-spec[,storage-spec,...]

This is provided as an alternate syntax for specifying storage as single attributes. It is useful in a defaults file, so that all attributes (including storage specifications) will be in the attribute=value format. Any number of storage specifications can be specified, separated by commas. More than one alloc attribute can be specified, in which case they are logically concatenated to form one longer list.

`comment=comment`

Specify a comment to give to a volume when it is created. This comment can be displayed with `vxprint -l`, and can be changed, at a later time, with `vxedit set`. This attribute is used only with the make operation.

`diskgroup=disk-group (alias: dg)`

Specify the disk group for an operation. If a disk group is specified in a defaults file, then it just specifies the default disk group to use for the make operation, if no other means of determining the disk group can be used. If specified as a command line attribute, it has the same effect as specifying a disk group with the `-g` option (the operation is forced to apply to the given disk group).

`excl=yes|y|on|true|no|n|off|false`

Sets or clears the EXCLUSIVE flag on the volume. A volume in exclusive open state can be opened by only one node in the cluster at a time. Multiple opens of an exclusive volume from the same node are permitted. Non-exclusive volumes can be simultaneously opened by more than one node. After a node opens an exclusive volume, every other node's open attempt fails until the last close of the volume by the first opener. Such an open failure returns a EBUSY error code. Available only if the Volume Manager cluster feature is enabled.

`fstype=file-system-type`

Specify the file system type for a new volume (this is used only with the make operation). The file system type for a volume is usually determined when needed by running the `fstyp` utility (see `fstyp(1M)`), but can be specified explicitly. The file system type parameter is used with the Volume Manager `vxresize` utility, as well as with volume snapshots, copies, and dissociates (to



select features that are used with the VERITAS VxFS file system).

The file system type can be changed at any later time with the vxedit set operation.

group=owning-group

Set the group ID for a new volume. The group ID can be specified numerically or with a system group name. This attribute is used only with the make operation. By default, volumes are created with group 0.

init=initialization-type

Specify the means for initializing a new volume. The default method (which can be selected explicitly with init=default) is to call vxvol start to do a usage-type-specific default initialization operation. A new volume can be left uninitialized with init=none. The most useful non-default initialization that can be specified is init=zero, to clear the volume before enabling it for general use.

layout=layout-spec[,layout-spec,...]

Specify a volume or plex layout type (RAID-5, mirrored, unmirrored, striped, cylinder alignment), and turn on or off some features (such as logging). The list of layout specifications is given in the next section. By default, unmirrored, non-striped volumes are created with no log.

loglen=length

Specify a log length to use for dirty-region log subdisks for mirrored volumes or for RAID-5 log plexes. If a log length is specified on the command line, then logging will be enabled by default.

logtype=drl|region|none

Specify the logging type for a mirrored volume. Dirty region logging (the default) can be selected with either logtype=drl or logtype=region. Logging can be disabled altogether with logtype=none.

max\_nraid5column=number

(alias: maxraid5columns, max\_nraidcolumn,  
max\_nraid5stripe, max\_nraidstripe, maxraidcolumns,

maxraid5stripes, maxraidstripes)

Specify the default maximum number of stripe

Maintenance Commands

vxassist(1M)

columns for a RAID-5 volume (default value is 8). The rules for selecting the default number of RAID-5 stripe columns are described in the nstripe attribute.

max\_ncolumn=number (alias: maxcolumns, max\_nstripe, maxstripes)

Specify the default maximum number of stripe columns, either for a RAID-5 volume (if max\_nraid5stripe is not also specified) or for a striped plex (default value is 8). The rules for selecting the default number of stripe columns is described in the nstripe attribute.

max\_regionloglen=length (alias: maxregionloglen)

Specify the maximum default dirty region logging log subdisk length. If the user does not specify the dirty region log length for a volume, when creating the first log for a mirrored volume, vxassist uses a simple formula based on the log length. The default length will not be bounded by max\_regionloglen (default value is 32K).

min\_nraid5column=number

(alias: minraid5columns min\_nraidcolumn, min\_nraid5stripe, min\_nraidstripe, minraidcolumns, minraid5stripes, minraidstripes)

Specify the default minimum number of stripe columns for a RAID-5 volume (default value is 3). The policy for selecting a default number of RAID-5 stripe columns will not be allowed to create a RAID-5 volume with fewer than min\_nraid5stripe stripe columns.

min\_nstripe=number (alias: minstripes)

Specify the default minimum number of stripe columns for either a RAID-5 volume (if min\_nraid5stripe is not also specified) or for a

striped plex (default value is 2). The policy for selecting a default number of stripe columns will not be allowed to select fewer than this number of columns.

mirror= +/- number|yes|no|diskclass[,...]  
Specify various mirroring parameters. Multiple mirroring parameters can be given, each separated by a comma. A decimal number indicates a specific number of mirrors to create, when creating a mirrored volume (equivalent to nmirror=number). yes

VxVM 3.0

Last change: 15 Dec 1999

15

Maintenance Commands

vxassist(1M)

indicates that volumes should be mirrored by default (equivalent to layout=mirror). no indicates that volumes should be unmirrored by default (equivalent to layout=nomirror).

Any other mirroring parameters specify that volumes should be mirrored across instances of a particular disk class. For example, mirror=target specifies that volumes should be mirrored between SCSI target addresses. Each mirror can contain disks from any number of instances of the disk class, but different mirrors cannot use disks from the same instance of the disk class.

mirrorconfine=diskclass[,diskclass,...]  
Specify a required list of disk classes for mirror confinement constraints. If this attribute is specified, a single mirror is confined to (that is, allocated from) a group of disks belonging to a specific disk class. Multiple mirrors can be allocated from a single disk class. Multiple disk classes cannot be used to allocate a mirror. The following example specifies that each mirror is allocated from devices belonging to a single SCSI target; and multiple mirrors can be allocated from devices belonging to a single SCSI target:

```
mirrorconfine=scsi_target
```

mode=permissions

Specify the permissions for the block and character device nodes created for a new volume. The

mode can be specified either as an octal number or symbolically. A symbolic mode is specified using the syntax given in `chmod(1)`. This attribute is used only with the `make` operation. The default mode for a volume gives read and write access only to the owner.

`nlog= +/- number, (nlogs, logs)`  
Specify the number of logs to create, by default, for a RAID-5 or mirrored volume (presuming that logs will be created). The number of logs to create can be specified independently for RAID-5 or mirrored volumes using the `nraid5log` and `nregionlog` attributes.

`nmirror= +/- number, (nmirrors, mirrors)`  
Specify the number of mirrors to create when mirroring a volume (default value is 2).

#### Maintenance Commands

`vxassist(1M)`

`nraid5log=number`  
(alias: `nraid5logs`, `raid5logs`, `nraidlog`, `nraidlogs`, `raidlogs`)

Specify the number of log plexes to create for a new RAID-5 volume (default value is 1). This attribute is used only with the `make` operation.

`nraid5stripe= +/- number`  
(alias: `raid5stripes`, `nraid5column`, `nraidstripe`, `raidstripes`, `raidcolumns`, `nraidcolumn`, `raid5columns`)

Specify the number of stripe columns to use when creating a new RAID-5 volume (the default is to adjust the number to available disks). This attribute is used only with `make` and `layout` operations. For the `layout` operation, the default is the same number of stripe columns the volume has. For the `layout`, this value can be preceded by a `+` or a `-` to add or subtract column(s).

`nregionlog=number` (alias: `nregionlogs`, `regionlogs`, `ndr1`)

Specify the number of log subdisks to create for a new mirrored volume (default value 1). This attribute is used only with the make operation, and only if logging is requested for the volume.

nstripe= +/- number  
(alias: stripes, ncolumn, ncolumns, ncol, ncols, columns, cols)

Specify the number of stripe columns to use when creating a new RAID-5 volume (with the make operation) or when creating a striped plex (with the make, relayout, mirror, and snapstart operations). The default is to adjust to the number of available disks. For the relayout operation, the default is the same number of stripe columns the volume has. For the relayout, this value can be preceded by a + or a - to add or subtract column(s).

raid5\_stripeunit=width  
(alias: raid5\_stwidth, raid5\_st\_width, raid5\_stwidth, raid\_st\_width, raid\_stripeunitsize, raid5\_stripeunitsize, raid5\_stripeunitwidth, raid\_stwid, raid\_stwidth, raid\_stripeunit, raid\_stripeunitwidth)

Specify the stripe unit size to use when creating a new RAID-5 volume (default value is 16K). This attribute is used only with the make operation.

raid5loglen=length, (alias: raidloglen)

Specify the log length to use when adding the first log to a RAID-5 volume. The default is four times the full stripe width (the stripe unit size times the number of stripe columns).

regionloglen=length, (alias: drlloglen, drllen)

Specify the log subdisk length to use when adding the first log subdisk to a mirrored volume. The default is chosen based on a formula involving the volume length.

stripe=diskclass[,...]

Specify that volumes should be striped across instances of a particular disk class. For example, stripe=target specifies that volumes should be striped between SCSI target addresses. Each column can contain disks from any number of instances of the disk class, but different columns cannot use disks from the same instance of the disk class.

stripe-mirror-col-trigger-pt=size

This size specification triggers the creation of a stripe-mirror volume. If the size of the volume is greater than the size specified in this attribute, then a stripe-mirror volume is created.

stripe-mirror-col-split-trigger-pt=size

This size specification triggers the creation of a stripe-mirror-sd volume. If the size of each column is greater than the size specified in this attribute, then a stripe-mirror-sd volume is created.

stripe\_stripeunit=width

(alias: stripe\_stwid, stripe\_stwidth, stripe\_st\_width, stripe\_stripeunitsize, stripe\_stripeunitwidth)

Specify the stripe unit size to use when creating striped plexes to attach to a volume. When attaching a new plex, the default is to use the same stripe unit size as any other striped plexes in the volume. If the volume does not yet contain striped plexes, the default value is 64K.

stripeunit=width

(alias: stwid, stripewidth, stwidth, st\_width, stripe\_width, stripeunitsize, stripeunitwidth)

Specify the stripe unit size to use for either a RAID-5 volume or for striped plexes. stripeunit can be used to set the values for both raid5\_stripeunit and stripe\_stripeunit.

tmpsize=size

Specify the space to be used as temporary storage in order to perform a relayout operation.

tmpalloc=storage-spec[,storage-spec,...]

Specify a set of storage specifications for the temporary storage required during a relayout operation. The format is the same as for the alloc attribute (a comma-separated list of storage specifications).

user=owning-user

Specify the user ID for a new volume (default value is root). The user ID can be specified numerically or it can be a system login name. This attribute is used only with the make operation.

usetype=volume-usage-type

Specify the usage type to use when creating a new volume (default value is raid5 for RAID-5 volumes; fsgen, otherwise). This attribute is used only with the make operation. The usage type can also be specified in the option list with -U.

wantalloc=storage-spec[,storage-spec,...]

Specify a set of desired storage specifications. This is useful in a defaults file to indicate desired storage specifications that should be discarded if they fail to yield a reasonable set of allocations. The format is the same as for the alloc attribute (a comma-separated list of storage specifications).

For example, a defaults file can name a specific controller type to use for allocations, if possible. As soon as all disks on that type of controller are full, other controllers will be used.

wantmirror=diskclass[,diskclass,...]

Specify a desired list of disk class mirroring

specifications. This is useful in a defaults file to indicate a set of desired mirroring constraints that can be dropped if they fail to yield a reasonable set of allocations.

For example, a defaults file can specify that volumes should be mirrored between disks on different controllers in the system. Then, if the set of volumes builds up in such a way that a new volume cannot be mirrored across SPARCstorage Array trays, this constraint will be dropped.

```
wantmirrorconfine=diskclass[,diskclass,...]
```

Specify a list of disk classes for mirror confinement constraints. A single mirror is confined to (that is, allocated entirely from) a group of disks belonging to a single disk class.

This is useful in a defaults file to indicate a set of mirror confinement constraints that can be dropped if they fail to yield a reasonable set of allocations. For example, a defaults file can specify that a single mirror be confined to disks of the same vendor:

```
wantmirrorconfine=vendor
```

If vxassist cannot allocate a mirror from disk devices of one vendor, this constraint is dropped and the mirror is allocated from disk devices of other vendors.

```
wantstripe=diskclass[,diskclass,...]
```

Specify a desired list of disk class stripping specifications. This is useful in a defaults file to indicate a set of desired stripping constraints that can be dropped if they fail to yield a reasonable set of allocations.

For example, a defaults file can specify that volumes should be striped between disks on different controllers in the system. Then, if the set of volumes builds up in such a way that a new volume cannot be striped across SPARCstorage Array trays, this constraint will be dropped.

#### Layout Specifications

The layout attribute specifies a comma-separated list of simple parameters (with no arguments) that apply to vxassist operations.



Two new layouts were added in the 3.0 release of Volume Manager: `stripe-mirror` and `concat-mirror`. In previous releases, whenever mirroring was used, the mirroring had to happen above striping or concatenation. Now there can be mirroring both above and below striping and concatenation.

Putting mirroring below striping mirrors each column of the stripe. If the stripe is large enough to have multiple subdisks per column, each subdisk can be individually mirrored. A similar concatenated volume would also mirror each subdisk individually. These new layouts enhance redundancy and reduce recovery time in case of an error. In a `mirror-stripe` layout, if a disk fails, the entire plex is detached, thereby losing redundancy on the entire volume. When the disk is replaced, the entire plex must be brought up to date. Recovering the entire plex can take a substantial amount of time. If a disk fails in a `stripe-mirror` layout, only the failing subdisk must be detached, and only that portion of the volume loses redundancy. When the disk is replaced, only a portion of the volume needs to be recovered.

The new volume layouts are more complex than the older volume layouts. Because the advantages of recovery time and improved redundancy are more important for larger volumes, it is best to continue using `mirror-stripe` and `mirror-concat` for most volumes and only use `stripe-mirror` and `concat-mirror` for very large volumes.

You can specify `layout=mirror-stripe` or `layout=stripe-mirror` to implement the desired layout. If you specify `layout=mirror-stripe`, `vxassist` automatically determines the best layout for the volume. Unless there is a reason to implement a particular layout, it is best to let `vxassist` create the layout for each volume. Because the advantages of the new layouts are related to the size of the volume, `vxassist` creates the simplest configuration for smaller volumes and the more complex `stripe-mirror` for larger volumes.

The attributes `stripe-mirror-col-trigger-pt` and `stripe-mirror-col-split-trigger-pt` control the selection. They can be set in `/etc/default/vxassist`. Volumes that are smaller than `stripe-mirror-col-trigger-pt` are created as `mirror-stripe`, and volumes that are larger are created as `stripe-mirror`. If `vxassist` creates the `stripe-mirror` and the columns are larger than `stripe-mirror-col-split-trigger-pt`, the individual subdisks are mirrored instead of mirroring the columns of the stripe. By default, both of these attributes are set to one gigabyte.

## Maintenance Commands

vxassist(1M)

The diskgroup must be created on a 3.0 or later release to use the new layouts, but older diskgroups can be updated. See the vxvg upgrade command for more information on upgrading diskgroups.

Defined layout specifications are:

mirror, nomirror, raid5

New volumes should be mirrored, unmirrored (default), or RAID-5, respectively. For mirror, the attributes stripe-mirror-col-trigger-pt and stripe-mirror-col-split-trigger-pt are applied.

mirror-stripe

New volumes should be mirrored and striped. The mirroring is done at the volume level. The attributes stripe-mirror-col-trigger-pt and stripe-mirror-col-split-trigger-pt are ignored with this layout.

mirror-concat

New volumes should be mirrored. The mirroring is done at the volume level. The attributes stripe-mirror-col-trigger-pt and stripe-mirror-col-split-trigger-pt are ignored with this layout.

stripe, nostripe

New plexes should be striped or unstriped, respectively. When creating a new volume, the default is nostripe. When adding a new plex to an existing volume, the default is stripe if the volume already has one or more striped plexes, and nostripe in other cases. For stripe (when combined with mirror, the attributes stripe-mirror-col-trigger-pt and stripe-mirror-col-split-trigger-pt are applied.

stripe-mirror

New volumes should be striped and mirrored. The mirroring is handled at each column level or sub-disk level depending on the attribute stripe-mirror-col-split-trigger-pt.

stripe-mirror-col

New volumes should be striped and mirrored. The mirroring is handled at each column level. The attributes stripe-mirror-col-trigger-pt and stripe-mirror-col-split-trigger-pt are ignored with this layout.

stripe-mirror-sd

New volumes should be striped and mirrored. The

Maintenance Commands

vxassist(1M)

mirroring is handled at each a subdisk level.

concat-mirror

New volumes should be concatenated and mirrored. The mirroring is handled at each a subdisk level. The attribute stripe-mirror-col-split-trigger-pt is applied.

span, nospan

Allow (default) or disallow plexes, regular stripe columns, or RAID-5 stripe columns from spanning multiple disks. If nospan is indicated, then plexes or columns can be formed from multiple regions of the same disk, but cannot be formed from more than one disk.

contig, nocontig

Disallow or allow (default) plexes, regular stripe columns, or RAID-5 stripe columns from using multiple regions of disk. If contig is specified, then plexes and columns must be allocated from a single contiguous region of disk. If this is not possible, the allocation fails. By default, vxassist tries to allocate space contiguously, but will use multiple regions or multiple disks if needed.

log, nolog

Create (or do not create) dirty region logs (for mirrored volumes) or log plexes (for RAID-5 volumes) when creating a new volume. This attribute can be specified independently for mirrored and RAID-5 volumes with the raid5log and regionlog layout specifications.

grow, nogrow  
 Allows a volume to grow during a relayout operation.

shrink, noshrink  
 Allows a volume to shrink during a relayout operation.

raid5log, noraid5log  
 Create (default) or do not create log plexes for RAID-5 volumes.

regionlog, noregionlog  
 Create or do not create (default) dirty-region log subdisks for mirrored volumes.

diskalign, nodiskalign

Maintenance Commands

vxassist(1M)

Align (default) or do not align subdisks on cylinder boundaries. When alignment is not disabled, subdisks are created beginning on cylinder boundaries and are extended to match the end of a cylinder. Dirty region log subdisks, however, are not cylinder aligned (they are usually only 2 or 3 blocks long). Instead, when creating log subdisks, spaces are located from the available disks that could not be turned into regular subdisks because the spaces aren't cylinder aligned. For example, once one cylinder is used for a log subdisk, that cylinder cannot be used to create an aligned data (or RAID-5 log) subdisk, so other log subdisks will be created there until that cylinder fills up.

FILES

/etc/default/vxassist            System default settings file  
    for vxassist attributes.

EXIT CODES

The vxassist utility exits with a nonzero status if the attempted operation fails. A nonzero exit code is not a complete indicator of the problems encountered, but rather denotes the first condition that prevented further execution of the utility.

See vxintro(1M) for a list of standard exit codes.

SEE ALSO

chmod(1), fstyp(1M), sync(1M), vxedit(1M), vxintro(1M),  
vxmake(1M), vxmend(1M), vxplex(1M), vxrelayout(1M),  
vxresize(1M), vxsd(1M), vxtask(1M), vxvol(1M)

VxVM 3.0

Last change: 15 Dec 1999

24

===[ man1m/vxbootsetup.1m ]=====

Maintenance Commands

vxbootsetup(1M)

NAME

vxbootsetup - set up system boot information on a Volume  
Manager disk

SYNOPSIS

/etc/vx/bin/vxbootsetup [medianame...]

DESCRIPTION

The vxbootsetup utility configures physical disks so that they can be used to boot the system. Before vxbootsetup is called to configure a disk, mirrors of the root, swap, /usr and /var volumes (if they exist) should be created on the disk. These mirrors should be restricted mirrors of the volume. The vxbootsetup utility configures a disk by writing a boot track at the beginning of the disk and by creating physical disk partitions in the UNIX VTOC that match the

mirrors of the root, swap, /usr and /var.

With no medianame arguments, all disks that contain usable mirrors of the root, swap, /usr and /var volumes are configured to be bootable. If medianame arguments are specified, only the named disks are configured.

vxbootsetup requires that the root volume is named rootvol and has a usage type of root. The swap volume is required to be named swapvol and to have a usage type of swap. The volumes containing /usr and /var (if any) are expected to be named usr and var, respectively.

Root, swap, /usr and /var volumes are created when the original system boot disk is encapsulated with the vxencap utility or when the Volume Manager is first set up on the system. (/usr and /var are only encapsulated if they existed on the system; see the VERITAS Volume Manager (VxVM) Installation Guide and the vxencap(1M) man page for details on encapsulating the boot disk.)

See the chapter on recovery in the VERITAS Volume Manager (VxVM) System Administrator's Reference Guide for detailed information on how the system boots and how the Volume Manager impacts the system boot process.

The following utilities call vxbootsetup automatically: vxmirror, vxrootmir, vxresize, and vxencap. If you use vxassist or vxmake and vxplex to create mirrors of the root volume on a disk, you must call vxbootsetup directly to make that disk bootable.

SEE ALSO

disksetup(1M), edvtoc(1M), vxassist(1M), vxevac(1M),  
vxintro(1M), vxmake(1M), vxmirror(1M), vxplex(1M),  
vxrootmir(1M), vxresize(1M)

VxVM 3.0

Last change: 26 Mar 1999

2

==[ man1m/vxconfigd.1m ]=====

Maintenance Commands

vxconfigd(1M)

## NAME

vxconfigd - Volume Manager configuration daemon

## SYNOPSIS

```
vxconfigd [-dfk ] [-D diag_portal] [-m mode] [-r reset] [-R
request_portal] [-x debug]
```

## DESCRIPTION

The Volume Manager configuration daemon, vxconfigd, maintains disk configurations and disk groups in the VERITAS Volume Manager. vxconfigd takes requests from other utilities for configuration changes, and communicates those changes to the kernel and modifies configuration information stored on disk. vxconfigd is also responsible for initializing the Volume Manager when the system is booted.

## OPTIONS

-d This is equivalent to -m disable, which starts vxconfigd in disabled mode.

-D diag\_portal  
Specify a rendezvous file pathname for diagnostic operation connections to vxconfigd. By default, /etc/vx/vold\_diag is used. The diagnostic portal exists in both the enabled and disabled operating modes.

-f Start vxconfigd in the foreground. This is often useful when debugging vxconfigd, or when tracing configuration changes. If this flag is not used, vxconfigd will fork a background daemon process. The foreground process will exit as soon as vxconfigd startup processing completes.

-k If a vxconfigd process is running already, then kill it before any other startup processing. This is useful for recovering from a hung vxconfigd process. Killing the old vxconfigd and starting a new one should not cause any problems for volume or plex devices that are being used by applications or that contain mounted file systems.

-m mode Set the initial operating mode for vxconfigd. Possible values for mode are:

boot Handle boot-time startup of the Volume Manager. This starts the rootdg disk group and the root and /usr file system volumes. This mode is capable of operating before the root file system is remounted to read-write. vxdctl enable should be called later in the boot



## Maintenance Commands

vxconfigd(1M)

sequence to trigger vxconfigd to rebuild the /dev/vx/dsk and /dev/vx/rdsk directories.

`disable` Start in disabled mode. This creates a rendezvous file for utilities that perform various diagnostic or initialization operations. This can be used with the `-r` reset option as part of a command sequence to completely reinitialize the Volume Manager configuration. Use the `vxctl enable` operation to enable vxconfigd.

`enable` Start fully enabled (default). This will use the `volboot` file to bootstrap and load in the `rootdg` disk group. It will then scan all known disks looking for disk groups to import, and will import those disk groups. This will also set up the /dev/vx/dsk and /dev/vx/rdsk directories to define all of the accessible Volume Manager devices. If the `volboot` file cannot be read or if the `rootdg` disk group cannot be imported, vxconfigd will be started in disabled mode.

`-r reset` Reset all Volume Manager configuration information stored in the kernel as part of startup processing. This will fail if any volume or plex devices are currently in use. This option is primarily useful for testing or debugging.

`-R request_portal` Specify a rendezvous file pathname for regular configuration and query requests. By default, this is /etc/vx/vold\_request. The regular request portal exists only when vxconfigd is operating in enabled mode.

`-x debug` Turn on various parameters used for debugging or other miscellaneous aspects of vxconfigd operation. The debug option argument is a decimal number, which will set a tracing output level, or one of the following strings:

boot=volboot\_path

Specify the pathname to use for the volboot file. This is primarily of use with the stub debug option. The volboot file contains an initial list of disks that are used to locate

VxVM 3.0

Last change: 26 Mar 1999

2

Maintenance Commands

vxconfigd(1M)

the root disk group. It also contains a host ID that is stored on disks in imported disk groups to define ownership of disks as a sanity check for disks that might be accessible from more than one host.

cleartempdir

This option removes and recreates the /var/vxvm/tempdb directory. This directory stores configuration information that is cleared on reboots (or cleared for specific disk groups on import and deport operations). If the contents of this directory become corrupt, such as due to a disk I/O failure, then vxconfigd will fail to start up if it is killed and restarted. Such a situation can be cleared by starting vxconfigd with -xcleartempdir. This option has no effect if vxconfigd is not started in enabled mode.

Note: It is advisable to kill any running operational utilities (vxvol, vxsd, or vxmend) before using the -x cleartempdir option. Failure to do so may cause those commands to fail, or may cause disastrous but unchecked interactions between those commands and the issuance of new commands. It is okay to use this option while running Volume Manager's Graphical User Interface, or while VxVM background daemons are running (vxsparecheck, vxnotify, or vxrelocd).

devprefix=prefixdir

Specify a directory pathname to prefix for any disk device accessed by vxconfigd. For example, with devprefix=/tmp, any access to a raw disk device named c2t1d0s2 would actually be directed to the file /tmp/dev/rdisk/c2t1d0s2. In stubbed-mode,

vxconfigd can operate with such files being regular files. vxconfigd only requires entries in the prefixdir /dev/rdisk directory in stubbed mode.

#### log or nolog

As an alternative to the use of syslog(), vxconfigd can directly log all of its console output to a file. This logging is reliable, in that any messages which are output just before a system crash will be available in the log file, presuming that the crash does not result in file system corruption.

VxVM 3.0

Last change: 26 Mar 1999

3

Maintenance Commands

vxconfigd(1M)

vxconfigd can be compiled such that logging defaults either to enabled or disabled. If it defaults to disabled, then it can be turned on with -x log; otherwise, it can be turned off with -x nolog.

Logging is disabled by default. If enabled, the default log file location is /var/vxvm/vxconfigd.log.

#### logfile=logfilename

Specify an alternate location for the vxconfigd logfile. This option implies -x log.

#### noautoconfig

Normally, vxconfigd will automatically configure disk devices that can be found by inspecting kernel disk drivers. These auto\_configured disk devices are not stored in persistent configurations, but are regenerated from kernel tables after every reboot. Invoking vxconfigd with -x noautoconfig prevents the automatic configuration of disk devices, forcing the Volume Manager to use only those disk devices configured explicitly using vxdisk define or vxdisk init.

#### nothreads

vxconfigd runs single-threaded.

stub This vxconfigd invocation will not communi-

cate configuration changes to the kernel. It is typically used as a demonstration mode of operation for vxconfigd. In most aspects, a stubbed vxconfigd will act like a regular vxconfigd, except that disk devices can be regular files and volume nodes are not created. A stubbed vxconfigd can run concurrently with a regular vxconfigd, or concurrently with any other stubbed vxconfigd processes, as long as different rendezvous, volboot, and disk files are used for each concurrent process.

Other Volume Manager utilities can detect when they are connected to a vxconfigd that is running in stubbed mode. When a utility detects a stubbed-mode vxconfigd, it will normally stub out any direct use of volume or plex devices, itself. This allows regular utilities to be used for making configuration changes in a testing environment that runs

## Maintenance Commands

## vxconfigd(1M)

without any communication with the kernel or creation of real volumes or plex devices.

### syslog or nosyslog

These options are provided only on systems where vxconfigd is compiled with support for the syslog() library calls (see syslog(3) and syslogd(1M)). vxconfigd can support use of syslog() to log all of its regular console messages. This support is either available by default (in which case -x nosyslog is required to disable its use), or it can support it as a run-time option (in which case -x syslog can be used to enable its use).

-x syslog must be supplied to vxconfigd to enable use of syslog(). The system can be changed to use this option at boot-time by editing the VxVM startup scripts.

Note: If the syslog option is enabled, then all console output will be directed through the syslog() interface. However, both syslog

and `log` (described below) can be used together to get reliable logging, into a private log file, along with distributed logging through `syslogd`.

`synctrace`

Flush tracefile data to disk, with `fsync`, to ensure that the last entry will be included in the file, even if the system crashes.

`timestamp` or `mstimestamp`

Attach a date and time-of-day timestamp to all messages written by `vxconfigd` onto the console. If `mstimestamp` is used, then a millisecond value is also displayed, allowing detailed timing of `vxconfigd`'s operation.

`tracefile=file`

Log all possible tracing information in the given file.

EXIT CODES

If errors are encountered, `vxconfigd` writes diagnostic messages to the standard error output. Some serious errors will cause `vxconfigd` to exit. If an error is encountered when importing the `rootdg` disk group during a normal startup, `vxconfigd` will enter disabled mode. See the VERITAS Volume Manager System Administrator's Reference Guide for a list of error messages and how to respond to them.

Defined exit codes for `vxconfigd` are:

- 0 The requested startup mode completed successfully. This is returned if `-f` is not used to startup `vxconfigd` as a foreground process. If `vxconfigd` is started as a foreground process, then it will exit with a zero status if `vxdctl stop` is used to cause `vxconfigd` to exit.
- 1 The command line usage is incorrect.
- 2 Enabled-mode operation was requested, but an error caused `vxconfigd` to enter disabled mode instead. This is also returned for boot-mode operation if startup failed. However, with boot-mode opera-

- tion, the background vxconfigd process exits as well.
- 3 The -k option was specified, but the existing vxconfigd could not be killed.
  - 4 A system error was encountered that vxconfigd cannot recover from. The specific operation that failed is printed on the standard error output.
  - 5 The background vxconfigd process was killed by a signal before startup completed. The specific signal is printed on the standard error output.
  - 6 A serious inconsistency was found in the kernel, preventing sane operation. This can also happen because of version mismatch between the kernel and vxconfigd.
  - 7 The -r reset option was specified, but the Volume Manager kernel cannot be reset. Usually this means that a volume is open or mounted.
  - 8 An interprocess communications failure (usually a STREAMS failure). Has made it impossible for vxconfigd to take requests from other utilities.
  - 9 Volumes that must be started early by vxconfigd could not be started. The reasons, and possible recovery solutions, are printed to the standard error output. The early-started volumes are the root and /usr file systems, if either or both of those file systems are defined on volumes.

#### FILES

/dev/vx/dsk                      Directory containing block device nodes for volumes.

VxVM 3.0

Last change: 26 Mar 1999

6

#### Maintenance Commands

vxconfigd(1M)

/dev/vx/rdisk                    Directory containing raw device nodes for volumes.

/etc/vx/volboot                 File containing miscellaneous boot information. See vxdctl(1M) for more information on this file.

/var/vxvm/tempdb                   Directory containing miscel-  
                                      laneous temporary files.  
                                      Files in this directory are  
                                      recreated after reboot.

/var/vxvm/vxconfigd.log.            Default log file.

SEE ALSO

syslogd(1M),     vxdctl(1M),     vxintro(1M),     vxnotify(1M),  
vxrelocd(1M),    vxsparecheck(1M),    fsync(2),     syslog(3),  
syslog.conf(4B)

## NAME

vxdctl - control the volume configuration daemon

## SYNOPSIS

vxdctl [ -c ] mode

vxdctl [ -k ] stop

vxdctl add disk accessname [ attr[=value] ]...

vxdctl disable

vxdctl enable

vxdctl hostid hostid

vxdctl init [ hostid ]

vxdctl initdmp

vxdctl license [init]

vxdctl list

vxdctl mode

vxdctl rm disk accessname...

vxdctl support

## DESCRIPTION

The vxdctl utility manages some aspects of the state of the volume configuration daemon, vxconfigd, and manages configuration aspects related to bootstrapping the rootdg disk group configuration.

A key part of the state of vxconfigd and of bootstrapping the rootdg disk group is the volboot file. This file contains a host ID that is used by the Volume Manager to establish ownership of physical disks. This host ID is used to ensure that two or more hosts that can access disks on a shared SCSI bus will not interfere with each other in their use of those disks. This host ID is also important in the generation of some unique ID strings that are used internally by the Volume Manager for stamping disks and disk groups.

Note: The purpose of the cluster environment is to allow hosts that can access disks on a shared bus to do so in a coordinated way. In this environment, the host ID is not used to prevent shared access; however, it is used for other internal functions.



## Maintenance Commands

vxdctl(1M)

The volboot file also contains a list of disks to scan in search of the rootdg disk group. At least one disk in this list must be both readable and a part of the rootdg disk group, or the Volume Manager will not be able to start up correctly.

vxconfigd operates in one of three modes: enabled, disabled, or booted. The enabled state is the normal operating state. Most configuration operations are allowed in the enabled state. Entering the enabled state imports all disk groups, and begins the management of device nodes stored in the /dev/vx/dsk and /dev/vx/rdsk directories.

In the disabled state, vxconfigd does not retain configuration information for the imported disk groups, and does not maintain the volume and plex device directories. Most operations are disallowed in the disabled state. Certain failures, most commonly the loss of all disks or configuration copies in the rootdg disk group, will cause vxconfigd to enter the disabled state automatically.

The booted state is entered as part of normal system startup, prior to checking the root file system (see fsck(1M)). Entering the booted mode imports the rootdg disk group, and then waits for a request to enter the enabled mode. The volume device node directories are not maintained in booted mode, because it may not be possible to write to the root file system.

The action performed by vxdctl depends upon the keyword specified as the first operand.

## KEYWORDS

vxdctl add disk

Add to the list of disks in the volboot file. Disks are specified based on their disk access name. This name identifies the physical address of the disk. For example, to add disk 0 at SCSI target 1 on SCSI controller 2, enter:

```
vxdctl add disk c2t1d0s2
```

This adds the disk c2t1d0s2. The controller s2 is assumed. If you add a disk with another slice,

type:

```
vxctl add disk c0t1d0s3
```

VxVM 3.0

Last change: 26 Mar 1999

2

Maintenance Commands

vxctl(1M)

If there is a disk access record in the rootdg configuration for the named disk, then configuration parameters are taken from that record. Otherwise, it may be necessary to specify some attributes to vxctl add disk.

vxctl disable

Request that vxconfigd enter disabled mode. This may be necessary to perform some maintenance operations. This does not disable any configuration state loaded into the kernel. It only prevents further configuration changes to loaded disk groups until vxconfigd is re-enabled.

vxctl enable

Request that vxconfigd enter enabled mode, import all disk groups, and rebuild the volume and plex device node directories. This operation can be used even if vxconfigd is already in enabled mode. The primary purpose for using this operation when in enabled mode is to rebuild the volume and plex device nodes. This operation will cause vxconfigd to scan for any disks that were newly added since vxconfigd was last started. In this manner, disks can be dynamically configured to the system and then recognized by the Volume Manager. If the Multipathing support is available, this option will also cause vxconfigd to rebuild the DMP internal database to reflect the new state of the system after addition of the disk devices. The new disk devices detected by vxconfigd are added in the DMP database with their associated subpaths and parent DMP device.

vxctl hostid

Change the host ID in the volboot file and on all disks in disk groups currently imported on this machine. It may be desirable to change the Volume

Manager host ID for your machine if you are also changing the network node name of your machine.

If some disks are inaccessible at the time of a `hostid` operation, it may be necessary to use the `vxdisk clearimport` operation to clear out the old host ID on those disks when they become reaccessible. Otherwise, you may not be able to re-add those disks to their disk groups.

Note: Be careful when using this command. If the system crashes before the `hostid` operation completes, some disk groups may not reimport automatically.

## Maintenance Commands

`vxctl(1M)`

### `vxctl init`

Reinitialize the `volboot` file with a new host ID, and with an empty list of disks. If a `hostid` operand is specified, then this string is used; otherwise, a default host ID is used. The default host ID is normally taken as the network node name for the host (see `uname(1M)`). On systems with a hardware-defined system ID, the default host ID might be derived from this hardware ID.

### `vxctl initdmp`

Create user level nodes for all the DMP devices that have been detected by the kernel. This will remove all the existing DMP nodes in `/dev/vx[r]dmp` directory, and create fresh nodes for the DMP devices that have been detected.

### `vxctl license [init]`

With an argument of `init`, request that `vxconfigd` re-read any persistently stored license information. If licenses have expired, then this may cause some features to become unavailable. If new licenses have been added, this will make the features defined in those licenses available.

With no arguments, `vxctl license` prints the list of features which are currently available based on known licensing information.

### `vxctl list`

List the contents of the volboot file. This list includes the host ID, some sequence numbers, and the list of disks and disk attributes stored in the volboot file.

#### vxctl mode

Print the current operating mode of vxconfigd. The output format is:

```
mode: operating_mode
```

where operating\_mode is either enabled, disabled, booted, or not-running.

If the -c flag is specified and vxconfigd is in the enabled mode, the clustered state of vxconfigd is displayed. The output format is:

```
mode: enabled: clustered_state
```

The clustered\_state is one of cluster inactive, cluster active - role not set, cluster active - MASTER, cluster active - SLAVE. Available only if the Volume Manager cluster feature is enabled.

#### vxctl rm disk

Remove one or more disks from the volboot file. Disks are specified based on the name used in the corresponding vxctl add disk operation.

#### vxctl support

Display information about what versions of various VxVM objects or components are supported by the configuration daemon currently running. Versions are printed as colon-separated name/value pairs. vold\_vrsn indicates the version of vxconfigd that is currently running; dg\_minimum and dg\_maximum indicate the lowest and highest diskgroup versions supported by vxconfigd; and the kernel: indicates the highest kernel version supported by vxconfigd.

#### vxctl stop

Request that vxconfigd exit. This may be necessary to reset the Volume Manager, such as using the -r reset option to vxconfigd. This does not disable any configuration state loaded into the kernel. It only affects the ability to make configuration changes until vxconfigd is restarted. If the -k option is used vxconfigd will be stopped by sending it a SIGKILL signal. The command will delay for up to one second to verify that vxconfigd has exited. After one second if vxconfigd has not exited an error will be returned.

#### FILES

/etc/vx/volboot/

#### SEE ALSO

fsck(1M), uname(1M), vxbootsetup(1M), vxconfigd(1M),  
vxdg(1M), vxdisk(1M), vxintro(1M), vxlicense(1M), signal(5),  
vxdmp(7)

===[ man1m/vxdg.1m ]=====

Maintenance Commands

vxdg(1M)

#### NAME

vxdg - manage Volume Manager disk groups

#### SYNOPSIS

vxdg [ -g diskgroup ] [ -kp ] [ -o verify | override ]  
addisk [ medianame= ] accessname...

vxdg [ -n newname ] [ -h newhostid ] deport diskgroup...

vxdg destroy diskgroup...

```

vxdg flush [ diskgroup...]

vxdg [ -g diskgroup ] [ -qa ] free [ medianame...]

vxdg [ -Cfst ] [ -n newname ] import diskgroup

vxdg [ -T version ] [ -s ] [ -o verify | override]
    init  groupname [ nconfig=config-copies ]
    [ nlog=log-copies ] [ minor=base-minor ]
    [ medianame =] accessname...

vxdg [ -q ] list [ diskgroup...]

vxdg [ -g diskgroup ] [ -f ] reminor
    [ diskgroup ] new-base-minor

vxdg [-g diskgroup ] [-k ] repldisk
    unassoc-medianame=spare-medianame...

vxdg [ -g diskgroup ] [ -k ] [ -o verify | override]
    rmdisk medianame...

vxdg [ -g diskgroup ] [ -q ] spare [ medianame...]

vxdg [ -T version ] upgrade diskgroup

```

#### DESCRIPTION

The vxdg utility performs basic administrative operations on disk groups. Operations include the creation of disk groups, the addition of disks to a disk group, and disk group imports and deports. The behavior of the vxdg utility depends upon the keyword specified as the first operand.

A diskgroup argument can be either a disk group name or a disk group ID. A groupname argument is a disk group name, not a disk group ID. An accessname argument refers to a system-dependent disk access name (also referred to as a disk device name), as stored in the root configuration by the vxdisk utility. If the slice number extension in the disk access record name is not included in the accessname, s2 is assumed by default. If any other slice is required

then it should be included in the accessname (as in c2t1d0s2). A medianame argument is an administrative name used to define a disk within a disk group.

## KEYWORDS

### vx dg adddisk

Add the specified disk(s) to a disk group (rootdg by default). The disk must not already be part of a disk group. The accessname component to a disk specification argument names a disk access record (essentially a device address specification) used to access the disk. If a medianame component is specified, then it names the disk media record used to define the disk within the disk group. If no medianame component is specified, then the disk media record will have the same name as the disk access record.

Adding a disk to a disk group causes the disk group's configuration to be copied onto the disk (if the disk has regions for configuration copies). Also, the disk is stamped with the system's host ID, as defined in the volboot file.

If the -k flag is specified, then the disk media name must represent a disk media record that was previously dissociated from its disk access record with -k rmdisk; otherwise, a new disk media record will be created to represent the disk. With the -k option, plexes requiring recovery will be flagged as stale.

Specifying the -p flag with -k packs contiguous subdisks into one subdisk and aligns them consecutively on their respective disks. Use the -p option when adding a root encapsulated disk or it may not boot correctly.

In a cluster environment, adding a disk to a cluster-shared disk group will fail if the disk is not physically accessible from all joined nodes in the cluster. If the addition is successful, the disk is stamped with the cluster ID and marked with the shared flag.

### vx dg deport

Disable access to the specified disk group. A disk group cannot be deported if any volumes in the disk group are currently open. When a disk group is deported, the host ID stored on all disks in the disk group will be cleared (unless a new host ID is specified with -h), so the disk group

will not be reimported automatically when the system is rebooted.

A disk group can be renamed on `deport` by specifying a new disk group name with `-n newname`. A lock can be assigned to an alternate host by specifying the host ID (see `vxdctl(1M)`) of the alternate host. This allows the disk group to be auto-imported when the alternate host reboots. For example, the `-n` and `-h` options can be combined to export a disk group to be used as the `rootdg` disk group for a new machine.

In a cluster environment, when a cluster-shared disk group is deported, the cluster ID and shared flag stored on all disks in the disk group are cleared, so the disk group is not imported automatically when the cluster is next started.

Trying to deport a shared disk group during a cluster reconfiguration will fail.

#### `vxdg destroy`

Removes a diskgroup from the system. Use this option when a disk group and the information on the disks is no longer needed. This frees up space for use by other disk groups. A disk group cannot be destroyed if any volumes in the disk group are open. `vxdg destroy` can be used only on imported disk groups.

#### `vxdg flush`

Rewrite all disk on-disk structures managed by the Volume Manager for the named disk groups. This rewrites all disk headers, configuration copies, and kernel log copies. Also, if any configuration copies were disabled (for example as a result of I/O failures), this will rewrite those configuration copies and attempt to enable them.

`vxdg free` List free space that can be used for allocating subdisks. If a disk group is specified, limit the output to the indicated disk group, otherwise list space from all disk groups. If disks are specified, by disk media name, then restrict the output to the indicated disks. A region of free space is identified by disk media name, a physical device tag, an offset relative to the beginning of the public region for the media, and a length.

The physical device tag is a reference that indicates which physical device the disk media is



## Maintenance Commands

vxdg(1M)

defined on. It appears as a truncated disk access name. If a particular physical device is split into several Volume Manager disk objects, then the device tag for each Volume Manager disk object will be the same. Device tags can be compared to identify space that is on the same or on different physical disks.

If the `-q` option is specified, then no header is printed describing output fields. If the `-a` option is specified, then space on spare disks (which is not really allocatable) is listed in addition to regular free space; otherwise, space on spare disks is not listed.

`vxdg import`

Import a disk group to make the specified disk group available on the local machine. This will make any configuration information stored with the disk group accessible, including any disk and volume configurations. The disk group to import is indicated by the `diskgroup` argument, which can be either an administrative disk group name or a disk group unique ID.

Typically, a disk group will not be imported if some disks in the disk group cannot be found by the local host. The `-f` option can be used to force an import if, for example, one of the disks is currently unusable or inaccessible.

Note: Be careful when using the `-f` flag because it can import the same disk group twice from disjointed sets of disks. This can make the disk group inconsistent.

When a disk group is imported, all disks in the disk group are stamped with the host's host ID. Typically, a disk group cannot be imported if any of its disks are stamped with a non-matching host ID. This provides a sanity check in cases where disks can be accessed from more than one host.

If it is certain that a disk is not in use by another host (such as because a disk group was not

cleanly deported), then the -C option can be used to clear the existing host ID on all disks in the disk group as part of the import. A host ID can also be cleared using vxdisk clearimport.

A new name can be given to the disk group on import using -n newname. If -n is used with the

## Maintenance Commands

vx dg(1M)

-t option, then the stored name of the disk group will remain unchanged, but the disk group will be known to the importing host under the new name; otherwise, the name change will be permanent.

Typically, an imported disk group will be reimported automatically when the system is rebooted, if at least some of the disks in the disk group remain accessible and usable. This can be disabled using the -t option, which causes the import to be persistent only until the system is rebooted.

As an example of the use of -n and -t, a rootdg disk group from one host can be imported on a second host, operations (such as making repairs to the root volume) can be performed on the second host and the disk group can be given back to the originating host, which can then be rebooted on the repaired disk group. To do this, identify the disk group ID for the rootdg disk group with vxdisk -s list, and use that disk group to import that rootdg using -C to clear import locks, -t for a temporary name, and -n to specify an alternate name (to avoid collision with the rootdg disk group on the second host). After repair, deport the disk group using -h (described below) to restore the import lock from the first host.

In a cluster environment, use the -s option to import a disk group as cluster-shareable. This is only valid if the cluster is active on the host where the import takes place. Ensure that all the disks in a shared disk group are physically accessible by all hosts. A host which cannot access all the disks in a shared disk group cannot join the cluster.

The disks in a shared disk group are stamped with the ID of the cluster to which the hosts belong and are marked with the shared flag. When a host joins a cluster, it automatically imports disk groups whose disks are stamped with the cluster ID.

Trying to import a shared disk group during a cluster reconfiguration will fail.

`vxdg init` Define a new disk group composed of the indicated disks, identified by disk access names. This involves assigning an internal unique ID to the group, storing a pointer to that group in the root

configuration, storing a reference to the group on all of the named disks that have a disk header, and storing a disk group record in the disk group's configuration database. At least one of the disks specified must have space allocated for a configuration copy.

The `init` cannot complete if a disk is being used by a disk group, deported or otherwise. If `vxdg` finds an unneeded disk group on the disk, it can be cleaned with the `vxdisk -f init` command. `vxdg init` can then be run again.

If a medianame is specified for use with a particular disk, then that medianame will name the disk media record used to reference the disk within the disk group (for operations such as `rmdisk` and sub-disk creations). If no medianame is specified, then the disk media name defaults to `accessname`. See `vxdisk(1M)` for a discussion of definition and initialization of disk access records.

The `init` operation can be used to initialize a root disk group configuration, which is identified by the special name `rootdg`. If any database locations are listed in the `volboot` file, then as a special case for initializing `rootdg`, no disk specifications are allowed. Disks should be initialized and added to the disk group as the first

operations after creating rootdg. Some or all disks added to the rootdg disk group should also be added to the volboot bootstrap file (see vxdctl(1M)).

The nconfig and nlog operands can be used to configure the number of configuration database copies and kernel log copies that are maintained for a disk group. The config-copies and log-copies values are either a decimal number (including 0 or -1) or set to all or default. A value of all or -1 signifies that all configuration or log copies on all disks in the disk group will be maintained. A value of default or 0 (this is also the default value) signifies that the Volume Manager will manage copies that are distributed in a reasonable pattern throughout the disks and controllers on the system. Any other number signifies that a particular number of copies should be maintained (or all copies, if that number is larger than the number of available configuration or log copies on all disks).

#### Maintenance Commands

vxdg(1M)

When a specific number (or default) is requested, configuration copies are scattered approximately evenly through the disk controllers on the system. If SCSI disks with multiple disks per target are found, then each such target is treated similarly to a controller (that is, configuration copies are evenly distributed between such targets). With the default policy, one configuration or log copy is maintained for each controller, and one configuration or log copy is also maintained for each SCSI target that has multiple disks; if this does not result in allocating at least 4 copies, then additional copies are spread through the controllers and targets.

Refer to vxdisk(1M) for more information on configuration and log copies, and for information on how to create them.

Note: If a policy other than all is used, then some disks will not have up-to-date, online confi-

guration and log copies. As a result, it is possible that some number of disk failures will leave a disk group unusable, even if some disks in the disk group remain usable. The default policy allocates a sufficient number of copies, in a sufficient spread of locations, that such a scenario is very unlikely to occur.

Since disk groups can be moved between systems, it is desirable that device numbers used for volumes be allocated in separate ranges for each disk group. That way, an administrator can choose ranges such that all disk groups in a group of machines can be moved around without causing device number collisions. Collisions may occur because the Volume Manager stores device numbers in disk group configurations, so that the same numbers can be used after a reboot (which is necessary for use with NFS, which requires persistency of device numbers). If two systems use the same device numbers for a set of volumes, and if a disk group from one machine is moved to the other, then the Volume Manager may be forced to temporarily remap some devices.

A base volume device minor number can be set for a disk group with the minor operand. Volume device numbers for a disk group will be chosen to have minor numbers starting at this base minor number. Minor numbers can range up to 131071, so if it is presumed that no more than 1000 volumes would ever

be created in any one disk group, 131 different ranges of minor numbers are available for different disk groups. A reasonably sized range should be left at the end for temporary device number remappings (in the event that two device numbers still conflict).

If no minor operand is specified on the init command line, then the Volume Manager chooses a random number of at least 1000 that is a multiple of 1000, and yields a usable range of 1000 device numbers. This default number is chosen such that it does not overlap within a range of 1000 of any

currently imported disk groups, and does not overlap any currently allocated volume device numbers.

Note: The default policy is likely to ensure that a small number of disk groups can be merged successfully between a set of machines. However, in cases where disk groups will be merged automatically using fail-over mechanisms, the administrator should select ranges that are known to avoid overlap.

In a cluster environment, the `-s` option defines a new disk group which is cluster-shareable while the cluster is active. It is the responsibility of the user to ensure that disks specified as members of a cluster-shareable disk group are physically accessible from the hosts that make up the cluster.

The disks in a shared disk group are stamped with the ID of the cluster to which the hosts belong and are marked with the shared flag. When a host joins a cluster, it automatically imports disk groups whose disks are stamped with the cluster ID.

Trying to create a shared disk group during a cluster reconfiguration will fail.

Note: Volumes in shared disk groups must have the same minor number on all nodes in the cluster. If there is a conflict when a node attempts to join the cluster, the join will fail. In that case, the administrator should use the `remminor` operation on the joined node(s) to resolve the conflict. In a cluster where more than one node is joined, the administrator should use a base minor number which does not conflict on any node.

If a version is specified with the `-T` option, the diskgroup is initialized with that diskgroup version. This limits the operations that can be performed and features that can be used to those supported by the specified diskgroup version. This

makes the diskgroup compatible with releases of VxVM that support that version. If no version is specified, the diskgroup is initialized with the highest versions supported by the release of VxVM currently running on the system. See the vxdbg upgrade operation for more information.

`vxdbg list` List the contents of disk groups. If no diskgroup arguments are specified, then all disk groups are listed in an abbreviated one-line format. If diskgroup arguments are specified, then a longer format is used to indicate the status of the disk group, and of the specified disk group configuration.

If the `-q` option is specified, then no header is printed describing output fields. This option has no effect with the long formats generated with diskgroup arguments.

In a cluster environment, if the `-s` option is specified, all cluster-shared disk groups are listed in a one-line format. If diskgroup arguments are specified, `-s` has no effect.

`vxdbg remminor`

Change the base minor number for a disk group, and renumber all devices in the disk group to a range starting at that number. If the device for a volume is open, then the old device number will remain in effect until the system is rebooted or until the disk group is deported and re-imported. Also, if you close an open volume, then the user can execute `vxdbg remminor` again to cause the renumbering to take effect without rebooting or reimporting.

A new device number may also overlap with a temporary renumbering for a volume device, which will also require a reboot or reimport for the new device numbering to take effect. A temporary renumbering can happen in the following situations: when two volumes (for example, volumes in two different disk groups) share the same permanently assigned device number, in which case one of the volumes is renumbered temporarily to use an alternate device number; or when the persistent

device number for a volume was changed, but the active device number could not be changed to match. The active number may be left unchanged after a persistent device number change either because the volume device was open, or because the new number was in use as the active device number for another volume.

vxdg will fail if you try to use a range of numbers that is currently in use as a persistent (not a temporary) device number. You can force use of the number range with use of the -f option. With -f, some device renumberings may not take effect until a reboot or a re-import (just as with open volumes). Also, if you force volumes in two disk groups to use the same device number, then one of the volumes will be temporarily renumbered on the next reboot. Which volume device will be renumbered should be considered random, except that device numberings in the rootdg disk group take precedence over all others.

The -f option should be used only when swapping the device number ranges used by two or more disk groups. To swap the number ranges for two disk groups, you would use -f when renumbering the first disk group to use the range of the second disk group. Renumbering the second disk group to the first range will not require use of -f.

#### vxdg repldisk

Dissociate the DA record from the DM record named by spare-medianame and reassociate it with the unassociated DM record named by unassoc-medianame. Both unassoc-medianame and spare-medianame must be members of the disk group named by the diskgroup argument (rootdg by default). However, if the -k flag is specified, then the disk media records for the spare-medianame will be kept, although in a removed state.

#### vxdg rmdisk

Remove the specified disk(s) from a disk group (rootdg by default). The last disk cannot be removed from its disk group. It is not possible to remove the last disk containing a valid disk group configuration or log copy from its disk group.

Typically, the rmdisk operation will fail if sub-disk records point to the named disk media records. However, if the -k flag is specified,



## Maintenance Commands

vxdg(1M)

then the disk media records will be kept, although in a removed state, and the subdisk records will still point to them. The subdisks, and any plexes that refer to them, will be unusable until the disk is re-added using the -k option to the adddisk operation. Any volumes that become unusable, because all plexes become unusable, will be disabled.

## vxdg spare

List spare space that can be used for relocating subdisks during recovery. If a disk group is specified, limit the output to the indicated disk group, otherwise list spare space from all disk groups. If disks are specified, by disk media name, then restrict the output to the indicated disks. A region of spare space is identified by disk media name, a physical device tag, an offset relative to the beginning of the public region for the media, and a length.

The physical device tag is a reference that indicates which physical device the disk media is defined on. It appears as a truncated disk access name.

If the -q option is specified, then no header is printed describing output fields.

## vxdg upgrade

Upgrades the disk group to the latest VxVM version. By default, the diskgroup version is updated to the running version of VxVM. The -T option upgrades the diskgroup to a specified version. The following section lists each diskgroup version, the features it supports, and the VxVM release that introduced it.

10 Supports only the most basic volume management features of mirroring and simple striping. This format was introduced in VxVM Release 1.2. Starting with VxVM Release 3.0, diskgroups of version 10 can be imported, but no operations can be performed on the objects it contains (for example, starting volumes or adding mirrors). The only operation supported is to upgrade the diskgroup to a later

release.

- 20 Introduced support for RAID-5 Volumes, new-style stripes, recovery checkpointing, disk-group configuration/klog copy limiting, and

VxVM 3.0

Last change: 15 Dec 1999

11

Maintenance Commands

vx dg(1M)

Dirty Region Logging. This version was introduced in VxVM Release 2.0 and is supported by all subsequent releases of VxVM.

- 30 Enabled support for the Oracle Resilvering Interface. This version was introduced in VxVM Release 2.2 and is supported by all subsequent releases of VxVM.
- 40 Support for Hot Relocation. Introduced in VxVM Release 2.3 and is supported by all subsequent releases of VxVM.
- 60 Support for Online Relayout, safe RAID-5 sub-disk moves, Striped Mirrors, and RAID-5 Snapshots. Introduced in Release 3.0.

You can determine a disk group version using the `vxprint -l` command, and the `vx disk list` operation prints a long listing of a disk group.

#### Hardware-specific Options

Some environments provide guidelines to optimize the Volume Manager's interaction with intelligent storage systems. If these guidelines are present, the Volume Manager follows the guidelines when creating disk groups and adding disks to disk groups. By default, `vx dg` only allows disk groups to contain disks that conform with these guidelines. The following options change the behavior of `vx dg`:

`-o override`

Performs the disk group task and ignores any storage-specific guidelines. Overriding the guidelines is not recommended as it can result in incompatible objects, or objects that cannot be administered by some utilities.

`-o verify` Verifies that the specified disk group task can be

performed without violating any storage-specific guidelines, but does not perform the task. If any guidelines are violated, vxdg exits with an error message.

SEE ALSO

vxconfigd(1M), vxdctl(1M), vxdisk(1M), vxintro(1M),  
vxplex(1M), vxprint(1M), vxvol(1M)

VxVM 3.0

Last change: 15 Dec 1999

12

===[ man1m/vxdisk.1m ]=====

Maintenance Commands

vxdisk(1M)

NAME

vxdisk - define and manage Volume Manager disks

SYNOPSIS

vxdisk [-g diskgroup] addregion region\_type disk offset  
length

vxdisk [-g diskgroup] check disk ...

vxdisk clearimport accessname ...

vxdisk [-f ] define accessname [attribute ...]

vxdisk [-f ] init accessname [attribute ...]

vxdisk [-g diskgroup] [-o alldgs] [-qs] list [disk ...]

vxdisk offline accessname ...

vxdisk -a online

vxdisk online accessname ...

vxdisk rm accessname ...

vxdisk [-g diskgroup] rmregion region\_type disk offset  
[length]

```
vxdisk [-g diskgroup] set disk [attribute ...]
```

#### DESCRIPTION

The vxdisk utility performs basic administrative operations on disks. Operations include initializing and replacing disks, as well as taking care of some book-keeping necessary for the disk model presented by the Volume Manager.

accessname refers to the disk access name, while disk represents the disk media name. vxdisk accesses disks based on disk access names, which are system-specific names that relate to disk addresses. Disk access names are in the form c#t#d#s#, which define a controller number(c#), a SCSI target ID (t#), a SCSI logical unit number (d#), and a partition number (s#). If the partition number is not specified, s2 is assumed (this represents the whole disk). If any other partition is required, it should be specified (as in c0t1d0s4). Special devices, such as internal RAM disks, may use different forms for disk access names. Disk access names relate directly to device node names in the /dev/dsk and /dev/rdisk directories.

Operations that take an accessname argument (see the SYNOPSIS section) accept only disk access names, as defined

VxVM 3.0

Last change: 26 Mar 1999

1

Maintenance Commands

vxdisk(1M)

in the previous paragraph. Operations that take a disk argument can take disk access names or disk media names (for example, disk01). For such operations, a disk group can be specified with -g to disambiguate disk media names that are used in more than one disk group.

Physical disks in the Volume Manager are presumed to be movable, and are usually identified by a unique disk ID stored on the physical disk, rather than by a disk device node. This allows disks to be moved to different SCSI target IDs or to different controllers without affecting correct operation.

The Volume Manager maintains known disk device address information in a set of disk access records, which are stored in the rootdg disk group configuration. These records are named based on the disk access name. These disk access records are normally used solely to identify which physical disks exist, based on disk IDs stored on the disks themselves. Operations for vxdisk other than init and

define require specification of defined disk access records.

Physical disks contain public regions, which are used for allocating subdisks. They can also contain private regions, which are used for storing private Volume Manager information. Private regions are structured regions, and are maintained entirely by the Volume Manager. Private regions contain the following structures:

#### Disk Header

Each private region contains exactly two copies of a disk header, which defines the unique disk ID, disk geometry information, and disk group association information. Two copies are created so that one copy can be lost (due to I/O failures) without causing use of the disk to be lost. The primary copy of the disk header is stored in block zero of the private region. The alternate copy is stored within the first 256 sectors. If the primary copy is unreadable or unusable, the Volume Manager will search the first 256 sectors of the private region for the alternate copy.

#### Table of Contents

A linked list of blocks, pointed to by the disk header, that define additional structures in the private and public regions. The table of contents blocks define disk group configuration copy locations, log copy locations, and reserved regions carved from the public region. Each link block in the table of contents is replicated at the beginning and end of the private region. If the

primary copy of any one link block is unreadable or unusable, the alternate copy of that link is used.

#### Configuration Copies

A disk normally contains one disk group configuration copy, according to the number specified when the disk was initialized using the `vxdisk init` operation (explained later). When a disk is added to a disk group, the disk group's persistent configuration records are written to each copy. For disks that are not associated with a disk group,

the space allocated for configuration copies is unused. Each disk group requires at least one usable configuration copy. Preferably there should be at least four copies, allocated between at least two disks. This allows one disk to be lost totally, while still preserving sufficient redundancy for recovering from simple read failures.

#### Disk Group Log Copies

A disk normally contains one disk group log copy. The number of log copies is set to the same as the number of configuration copies for the disk (as explained in the Configuration copies section above). These logs are written by the kernel when certain types of actions are performed: transaction commits, plex detaches resulting from I/O failures, total dirty region log (DRL) failures, the first write to a volume, and volume close. After a crash or a clean reboot, this log information is used to recover the state of a disk group just prior to the crash or reboot. Each disk group requires at least one usable disk group log copy. As with configuration copies, it is preferable to have at least four log copies, allocated between at least two disks.

For a single disk, the disk header and the table of contents blocks are critical data structures. At least one copy of the disk header, and at least one copy of each table of contents block, must be readable and usable, or else the disk itself is unusable and will have to be reinitialized.

Within disk groups, disk group configuration and log copies are critical data structures. At least one complete configuration copy and log copy must be readable and usable, or the disk group is unusable and will have to be reinitialized from scratch.

All disk group association information is stored in the disk header within private regions. This information consists of a disk group name, disk group unique ID, and a host ID. When the system boots, the Volume Manager scans for disks

that are stamped with the system's host ID. Each represented disk group is imported automatically. Disks with a non-matching host ID are not imported automatically, and cannot be used until the host ID is cleared with the clearimport operation.

The behavior of the vxdisk utility depends upon the keyword specified as the first operand.

#### KEYWORDS

##### vxdisk addregion

Add a new entry to the table of contents in a disk's private region. The new entry defines a region of disk that is relative to the public region, and that is reserved for a particular use. The offset and length operations indicate the location and extent of the region. Currently, the only region type that can be defined is:

reserve Mask out a region of disk that should be reserved for non Volume Manager purposes. This could be used, for example, to mask out a boot file system that cannot be used for subdisk allocation, or to mask out a region containing blocks that are used for bad-block or bad-track replacement.

Adding a region will fail if a subdisk or region is already allocated over the requested region.

The addregion functionality is currently unimplemented for any of the existing disk types.

##### vxdisk check

Validate the usability of the given disks. A disk is considered usable if the Volume Manager can write and read back at least one of the disk headers that are stored on the disk. If a disk in a disk group is found not to be usable, then it is detached from its disk group and all subdisks stored on the disk become invalid until the physical disk is replaced or the disk media record is reassigned to a different physical disk.

##### vxdisk clearimport

Clear the host-specific import information stored on the indicated disks, and in the configurations

stored on those disks. This command may be necessary in cases where import information stored for a disk group becomes unusable, due to host failures, or due to a disk group being moved from one machine to another.

This operation cannot be applied to disks that are in imported disk groups.

#### vxdisk define

Define a disk access record, but do not initialize it. In order for the Volume Manager to scan a disk, a disk access record must be defined for it. Thus, if you want to see what is on a new disk or you want to move a disk with a valid disk group from one system to another, you will need to use vxdisk define to make it accessible first. You can use vxdisk list to see what is on the disk, or vxdg import to import a disk group that is on the disk.

Attributes can be specified to define the access characteristics of the disk device. Some attributes that can be set are:

#### type=disk\_type

The disk device access type. See the init operation definition for more details.

The various disk types support additional attributes for the define operation. See the definition for each disk type, in the Disk Types section.

offline If specified, the disk will be created in the offline state.

Normally, a define operation will fail if the specified disk device is invalid, such as because no such disk currently exists. The -f option can be used to force definition of an unusable disk. This can be useful if, for example, the disk device could be used after a reboot. For example, if you intend to add a new controller and intend to move some existing disks to the new controller, you may need to define the new disk device addresses, even though they will not be usable until you shutdown and reconfigure your disks.

#### vxdisk init

Initialize regions of a disk used by the Volume



## Maintenance Commands

vxdisk(1M)

Manager. This involves installing a disk header and writing an empty configuration on the disk. The `accessname` operand identifies the disk. Normally, this command will fail if the disk already contains an apparently valid disk header. The `-f` option can be used to override this and to force initialization of the disk. A disk that is a member of an imported disk group cannot be initialized.

The `vxdisk init` operation creates a disk access record for a disk (if one does not already exist), and sets its state to online. Disks can be initialized when the root configuration is disabled, in which case the disk header will be initialized, but the disk will not be added to the permanent list of known disks until the root configuration is enabled.

Any attribute operands override default values assigned for various disk attributes. Some attributes that can be set are:

`type=disk_type`

The disk device access type, which is a system-specific name identifying a class of strategies for accessing disks and for managing private and public regions. For example, disk types could indicate network disks or a volatile RAM disk that may not require the storage of any private data. Disk types can also indicate a hard disk without separate partitions. If the disk access name is of the form `c#t#d#s2` or `c#t#d#`, the disk type defaults to `sliced`; otherwise, the disk type defaults to `simple`.

The various disk types support additional attributes for the `init` operation. See the definition for each disk in the Disk Types section.

`offline` The device will be left in the offline state, initially. This is used only if this operation is defining a new disk

access record.

#### vxdisk list

List path type and states along with the detailed information on the specified disks. The state will be listed as enabled or disabled. If no disk

VxVM 3.0

Last change: 26 Mar 1999

6

Maintenance Commands

vxdisk(1M)

arguments are specified, then print an one-line summary for all disk access records known to the system. If disk arguments are specified, then print a full description of the contents of the disk header and of the table of contents for each named disk. If no disk arguments are specified, but a disk group is specified with -g, then list only those disks added to the specified disk group.

If the -s option is specified, then list important information from the disk header. With -s, the output format is the same whether or not accessname arguments are specified. The information printed with -s includes the disk ID, the host ID (if the disk is or was imported), and the disk group ID and disk group name (if the disk is a member of a disk group).

If the -q option is specified, then no header is printed describing output fields. This option has no effect with the long formats generated with -s or with accessname arguments.

When -o alldgs is specified without -s and -g, a one line summary shows all disk to disk group associations. The disk group column shows imported disk groups as normal and shows all other disk groups in parentheses.

#### vxdisk offline

Declare the disk devices named by the accessname arguments to be in the offline state. This disables checking of the disk in searching for particular disk IDs, or for the set of disks in a particular disk group. This operation cannot be applied to disks that are members of an imported disk group.

A disk should be offlined if the disk is not currently accessible, and if accessing the disk has a negative impact on the system. For example, disk drivers on a few operating systems can cause system panics or hangs if an attempt is made to access disks that are not accessible. In other operating systems, attempts to access inaccessible drives may take several seconds or minutes before returning a failure.

#### `vxdisk online`

Clear the offline state for a disk device. This re-enables checking of the disk when searching for

#### Maintenance Commands

`vxdisk(1M)`

disk IDs, or for members of a disk group. This can be used for disks that are already in the online state, provided that they are not in imported disk groups. All internal information for an already online state disk is regenerated from the disk's private region.

If `-a` is specified, re-online all online disks that are not currently in an imported disk group. This can be used to force the Volume Manager to re-scan all disk headers, or to adapt to changes in the way a disk is partitioned.

`vxdisk rm` Remove the specified disk access records, by disk access name.

#### `vxdisk rmregion`

Free a region of space that is allocated in the private or public region for a particular use. Space that is freed from the public region becomes usable for subdisk creation. The arguments to `rmregion` must match the arguments used when adding the region with `vxdisk addregion` except for the optional length argument which can be excluded for the remove.

The `rmregion` functionality is currently unimplemented for any of the existing disk types.

#### `vxdisk set`

Change some set of attributes for a disk. The attributes are either simple names (used to turn on an on/off attribute), or can be of the form attrname=value, to indicate a value for a particular attribute.

#### Disk Types

Three disk types are provided with the base VERITAS Volume Manager. Additional types may be added for use with particular operating systems. The default is a sliced type for disk access names ending in s2. If the slice component is not specified, then the type defaults to sliced. If any other slice number is supplied, then the default type becomes simple.

#### nopriv Disks

The simplest disk type is nopriv, which defines a disk that has no private region, and that consists only of space for allocating subdisks. Configuration and log copies cannot be stored on such disks, and such disks do not support reserved regions defined with vxdisk addregion. Because nopriv disks are not self identifying, the Volume Manager cannot track

#### Maintenance Commands

vxdisk(1M)

the movement of such disks on a SCSI chain or between controllers.

nopriv devices are most useful for defining special devices (such as volatile RAM disks) that you wish to use with the Volume Manager, but that can't store private regions. A RAM disk cannot store a meaningful private region, because data written to a RAM disk may not survive a reboot.

Initializing a nopriv device with vxdisk init creates a disk access record in the rootdg configuration, but does not write to the disk. The disk ID for nopriv devices is stored in the disk access record in the rootdg configuration.

Attributes that can be used with the vxdisk init and define operations for the nopriv device type are:

publen=length or len=length

The usable length of the device. This is required if there is no system-defined procedure for determining the disk length; otherwise, a suitable default will be computed.

puboffset=offset or offset=offset

The offset within the device for the start of the usable region. This defaults to 1. This can be used if it is necessary to skip over some region reserved by the operating system. If an offset is specified, then the default disk length is adjusted accordingly.

volatile If this attribute is specified, the disk is considered to have volatile contents (that is, the disk contents are not expected to remain consistent across a system reboot). Subdisks and plexes defined on disks with the volatile attribute will inherit that attribute. The vxvol start operation interprets volatile plexes as requiring a complete revive from other plexes in the same volume.

The vxdisk define operation, with the nopriv device type, takes the same attributes as the init operation. In addition, define takes the following attribute:

diskid=newdiskid

This attribute will set the disk ID to the newdiskid value in the disk access record for the nopriv disk.

#### Sliced and Simple Disks

Two disk types are provided that support disk private

#### Maintenance Commands

vxdisk(1M)

regions: simple and sliced. The simple type presumes that the public and private regions are stored on the same disk partition, with the public region following the private region. The sliced type presumes that the public and private regions are stored on different disk partitions.

For the sliced type, if the disk access name ends in s2, then the the public and private regions default to disk partitions with the tags 14 and 15, respectively. If the disk access name ends in s followed by some octal digit other than 0 or 1, the public region defaults to the named partition, and the private region defaults to the immediately following partition. For example, with an access name of c0t1d0s4, the public region defaults to partition 4 (octal

4) and the private region defaults to partition 5. For an accessname of c0t1d0s1, the public region defaults to partition 1 and the private region defaults to partition 1.

Attributes that can be defined with vxdisk define for simple or sliced types are:

pubslice=number or slice=number

Define the partition number to use for the public partition. This can be used only with the sliced type.

privslice=number

Define the partition number to use for the private partition. This can be used only with the sliced type.

privoffset=offset

Specify the offset from the beginning of the partition containing the private region to the beginning of the private region. This defaults to zero for both simple and sliced types.

Note: It is strongly advised that this option never be used unless absolutely necessary. Most information for disks can be determined from the disk header stored at the beginning of the private region. The private region offset cannot be determined from the disk. As a result, specifying a private region offset adds an undesirable dependence between a disk access record and a specific physical disk.

In addition to the above attributes, vxdisk init takes the following attributes:

configlen=length

The size to reserve for each copy of the

configuration stored on the disk. The default size will be based on the size of the private area and the number of configuration copies requested, and will leave some space free for uses other than the configuration copies.

loglen=length

The size to reserve in the private region for each log region. This size limits the number of kernel-initiated detach operations that can be logged against the disk group. The default is about 15% of the size of the configuration copies. It is advised that the log sizes be kept as 15% of the configuration copy size.

nconfig=count

The number of configuration copies to store on the disk. This defaults to 1. This can be set to 0 to indicate that no configurations will be stored on the disk. The Volume Manager will automatically enable and disable the config copy. It will maintain a level of redundancy in configuration copies that will allow the configuration to be recovered from the loss of multiple disks. See the vxdg nconfig parameter for more information.

nlogs=count

The number of log regions to allocate on the disk. Logs regions are used for storing any plex detaches that happen within the disk group. This number defaults to 1. The Volume Manager will automatically enable and disable the config copy. It will maintain a level of redundancy in configuration copies that will allow the configuration to be recovered from the loss of multiple disks. Refer to the vxdg(1M) nlog parameter for more information.

publen=length or len=length

Specify the length of the public region. If not specified, the length of the public region is computed from available system-specific disk size information. If no such information is available, a public region length must be specified in this command. The default public region length is adjusted to account for the private region, or for any specified public or private region offsets.

puboffset=offset or offset=offset

Specify the offset from the beginning of the partition containing the public region to the beginning of the public region. For the sliced type,

this defaults to one. For the simple type, this defaults to the first block past the end of the private region.

`privlen=length`

Specify the length of the private region. If this is not specified, then a default is chosen. For the sliced type, the default is computed from available partition table information. For the simple type, the default size is 1024 blocks. With the sliced type, if no partition information is available, a private region length must be specified in this command.

#### Auto-Configured Disks

On some systems, the Volume Manager can ask the operating system for a list of known disk device addresses. On such systems, some device addresses will be auto-configured into the `rootdg` disk group when `vxconfigd` is started. Auto-configured disks will always be of type `sliced`, with default attributes.

Auto-configured devices can be removed, if necessary, using `vxdisk rm`. When removed, explicitly defined devices can be defined to override any auto-configured devices. When the system reboots, no auto-configured disk devices will be added to the `rootdg` disk group that would share a disk with an explicitly configured disk device.

Auto-configured devices can be disabled and reenabled using the `offline` and `online` operations. However, the `offline` state is not stored persistently. If you need to persistently `offline` a device at a particular address, you will need to convert the address to use an explicit device record. To do this, remove the auto-configured device, and use `vxdisk define` to create an explicitly configured device.

#### SEE ALSO

`vxconfigd(1M)`, `vxdg(1M)`, `vxintro(1M)`, `vxvol(1M)`



==[ man1m/vxdiskadd.1m ]=====

Maintenance Commands

vxdiskadd(1M)

NAME

vxdiskadd - add one or more disks for use with the Volume Manager

SYNOPSIS

vxdiskadd disk\_address\_pattern\_list

DESCRIPTION

The vxdiskadd utility configures disks for use by the Volume Manager. Disks must already have been formatted (see format(1M)) and have a valid VTOC (see prtvtoc(1M)). vxdiskadd prompts the user for a disk group name and disk name for the disks. If no disk group name specified, the disks will be left as unassigned replacement disks for future use. A new disk group may be created that will become the disks' disk group. If a disk group is specified for the disks, the user is prompted for whether the disks should be designated as spares for the disk group. If a disk is found to already contain non-Volume Manager partitioning, the user will be asked whether the disk should be encapsulated. Encapsulation will turn each partition of the disk into a volume. A disk should be encapsulated if it contains file systems or data that should not be overwritten. If encapsulation is not desired for a disk, the disk can be initialized as a new disk for use by the Volume Manager. For new disks, all space on the disk becomes free space in the disk's disk group. vxdiskadd interactively allows encapsulation or initialization to be done for all the disks specified or allows the user to ask to make the decision for each disk individually.

One or more disks may be specified with a disk\_address\_pattern\_list. The basic form of a disk address is c C t T d D. Any leading portion of this address may be specified to indicate that all disks that match that leading portion should be used. For example c2t0 may be used to specify all disks on controller 2, target 0. More than one disk address or address pattern may be specified on the command line. The word all may be used to specify all disks on the system. Disk address names relate directly to device node names in the /dev/dsk and /dev/rdsk directories. Here are more examples:

```
all:      all disks
c3 c4t2:  all disks on controller 3 and controller 4,
```

target 2

c3t4d2: a single disk

The files /etc/vx/cntrls.exclude and /etc/vx/disks.exclude may be used to exclude disks or controllers from use by vxdiskadd. Each line of cntrls.exclude specifies the address of a controller to exclude, for example, c2. Each

VxVM 3.0

Last change: 26 Mar 1999

1

Maintenance Commands

vxdiskadd(1M)

line of disks.exclude specifies a disk to exclude, for example, c0t3d0.

All courses of action that do not involve clear failure conditions require prompts from the user, with defaults supplied as appropriate. Help is made available at every prompt. Entering ? in response to a prompt displays a context-sensitive help message.

#### OPERATIONS

Reconnecting a drive that was temporarily inaccessible

This situation is detected automatically, by noting that the specified drive has a disk ID that matches a disk media record with no currently associated physical disk. After reconnection, any stale plexes referring the disk will be reattached, and any stopped volumes referring the disk will be restarted. This reattach action will be performed by calling the vxrecover script.

Encapsulation of disks that have existing contents

This is the default action for disks that do not have a valid, existing Volume Manager private region, but that do have a VTOC. Encapsulation is performed by calling vxencap.

Initialization of a disk with reserved regions and partitions

Initialization of the disk is performed by calling vxdisksetup.

Adding a disk to an existing disk group

This operation can be performed independently of the initialization of the disk drive, and will add the disk to the group so that its storage will be added to the free space pool in that disk group. The vxassist utility may subsequently allocate from that free space.

Creation of new disk groups in which to import a new disk  
If disks are added that are required to be put into a  
disk group that does not exist, then the option of  
creating the disk group will be offered.

FILES

/etc/vx/cntrls.exclude	Specifies the address of controllers to exclude from vxdiskadd operations.
/etc/vx/disks.exclude	Specifies the address of disks to exclude from vxdiskadd operations.

VxVM 3.0

Last change: 26 Mar 1999

2

Maintenance Commands

vxdiskadd(1M)

SEE ALSO

format(1M), prtvtoc(1M), vxassist(1M), vxdisk(1M),  
vxdiskadm(1M), vxdisksetup(1M), vxencap(1M), vxintro(1M),  
vxrecover(1M)

===[ man1m/vxdiskadm.1m ]=====

Maintenance Commands

vxdiskadm(1M)

NAME

vxdiskadm - menu driven Volume Manager disk administrator

SYNOPSIS

vxdiskadm

DESCRIPTION

vxdiskadm provides a menu-driven interface to perform common VxVM disk administration tasks.

The vxdiskadm script is interactive and prompts you for responses, supplying defaults where appropriate. There is a Help facility at every prompt. Enter a question mark ? at a prompt to display a context-sensitive help message.

To add or encapsulate disks, specify one or more disks with a disk-address-pattern-list. The basic format for disk

addresses is c#t#d#. You can specify just the controller and target to add or encapsulate all the disks at that SCSI address. For example, use c2t0 to specify all disks on controller two, target zero. You can specify more than one disk address or address pattern on the command line. The word all specifies all disks on the system. Disk address names relate directly to device node names in the /dev/dsk and /dev/rdisk directories.

You can use the files /etc/vx/cntrls.exclude and /etc/vx/disks.exclude to exclude disks or controllers from use by vxdiskadm. Each line of cntrls.exclude specifies the address of a controller to exclude, for example, c2. Each line of disks.exclude specifies a disk to exclude, for example, c0t3d0.

## OPERATIONS

Add or initialize one or more disks

Prompts for one or more disk device addresses.  
Prompts for a disk group (which defaults to rootdg). You are also given the alternative of initializing the disk but leaving it unallocated to a disk group. If a disk group is specified, you must name the disk. A default disk name in the form disk## is chosen for the root disk group. If you add a disk to a disk group other than rootdg, the name is in the form diskgroup##, such that the name are unique within all imported disk groups. If a disk group is specified for the disks, you are prompted to specify whether to designate the disks as spares for the disk group. The disks are then checked to ensure that there is no information already on them. If there is, you are given the option of encapsulating them.

Encapsulate one or more disks

Prompts for one or more disk addresses. It then calls vxencap to encapsulate partitions on the named disk.

Remove a disk

Prompts for a disk, by disk media name. The disk is checked to ensure that no subdisks reside on the disk. If the disk is in use, the operation

fails with a recommendation to first move all volumes off the disk. If this disk is the last disk in a group, you are also prompted whether to remove the disk group from the system.

The operation then calls `vxvg rmdisk` to remove the disk from its disk group. If this is the last disk in its disk group, `vxvg deport` is used instead.

#### Remove a disk for replacement

Prompts for a disk by disk media name. The disk is checked for volumes that would lose all mirrors as a result of the operation. If such volumes exist, those volumes are listed and you are prompted whether to continue the operation.

The operation calls `vxvg -k rmdisk` to dissociate the media record from the physical disk. If there are formatted disks available that have disk headers but no disk group, you are prompted whether to use one of these disks as a replacement.

#### Replace a failed or removed disk

Prompts for a disk media name. The named media record must be dissociated from a disk. If the media record is not in the removed state, unused disks are scanned for matching disk IDs. If a disk with a matching disk ID is found, you are prompted whether to reattach that disk.

If a matching disk is not used, you are prompted for a new disk, by device address. If the named replacement disk has a valid disk header, but is not allocated to a disk group, you are prompted whether to reinitialize the disk. If the named replacement disk is listed as allocated to a disk group or to another host, you are prompted whether to continue the operation.

If the device is initialized, `vxdisksetup` is called to set up public and private regions and to

create the disk header.

Given an initialized disk, the operation replaces the disk in a disk group with `vx dg -k adddisk`.

Mirror volumes on a disk

Prompts for a disk, by media name. It then prompts for a destination disk within the same disk group, also by media name. Specifying no destination disks indicates that any disk is can be the destination. The operation calls `vx mirror` to mirror the volumes.

Move volumes from a disk

Prompts for a disk, by media name. It then prompts for a possible list of destination disks, also by disk media name. Specifying no destination disks indicates that any disk is suitable. The operation calls `vx evac` to move subdisks off the disk.

Enable access to (import) a disk group

Prompts for a disk, by device address. The operation calls `vx dg import` to import the disk group stored on that disk.

Disable access to (deport) a disk group

Prompts for a disk group name. The prompt lists alternate disk groups and the disks (media name and access name) that they contain. The operation calls `vx dg deport`.

Enable (online) a disk device

Prompts for a disk device. The prompt lets you display the disks on the system. The operation functions only for disks currently offline. It then makes the disk accessible.

Disable (offline) a disk device

Prompts for a disk device. The prompt lets you display the disks on the system. The operation functions only for disks currently online, but not part of any disk group. It then marks the disk as offline such that Volume Manager no longer tries to access the disk.

Mark a disk as a spare for a disk group

Sets up a disk as a spare device for its disk group. A spare disk can be used to automatically replace a disk that has failed. No space can be used on a disk that is marked as a spare.

Turn off the spare flag for a disk

Removes a disk from spares list and returns its space to the general pool of available space.

FILES

/etc/vx/cntrls.exclude

Specifies the address of controllers to exclude from vxdiskadm operations.

/etc/vx/disks.exclude

Specifies the address of disks to exclude from vxdiskadm operations.

SEE ALSO

vx dg(1M), vx disk(1M), vx diskadd(1M), vx disksetup(1M), vx encap(1M), vx evac(1M), vx intro(1M) , vx rootmir(1M)



```
===[ man1m/vxdisksetup.1m ]=====
```

Maintenance Commands

vxdisksetup(1M)

#### NAME

vxdisksetup, vxdiskunsetup - configure a disk for use with the Volume Manager

#### SYNOPSIS

```
/etc/vx/bin/vxdisksetup [-i ] diskaddress [ attribute... ]
```

```
/etc/vx/bin/vxdiskunsetup [-C ] diskaddress..
```

#### DESCRIPTION

The vxdisksetup command configures the disk for use by the Volume Manager. vxdisksetup is called by the vxdiskadd command as part of initial disk configuration. It is not typically necessary to call this command directly. vxdisksetup creates two partitions on the specified disk, a Volume Manager private configuration region (called the private region partition), and a partition used for allocating sub-disks (called the public region partition). These two partitions have tags that identify them as appropriate to the Volume Manager.

The vxdiskunsetup command reverses the configuration done by vxdisksetup and makes the specified disks unusable by the Volume Manager. vxdiskunsetup can be applied only to disks that are not in use within an imported disk group.

The diskaddress argument passed to vxdisksetup or vxdiskunsetup directs the command to the disk being operated upon. The address should be provided in the form c # t # d #, which define a controller number (c#), a SCSI target ID (t#), and a SCSI logical unit number (d#). Each number is specified in decimal and can be multiple digits. The name must reference a valid disk with partition devices under the /dev/vx/rdmp directory.

The vxdiskunsetup command operates by removing the public and private regions that were created by the last invocation of vxdisksetup on the specified disks. After this operation, those disks are converted from the online to the error state (as displayed by vxdisk list). The command does not

operate on disks that are active members of an imported disk group.

#### OPTIONS

-C vxdiskunsetup usually does not operate on disks that appear to be imported by some other host (for example, a host that shares access to the disk). In this case, the -C option can be specified to force de-configuration of the disk, effectively removing the host locks that were detected by vxdiskunsetup.

VxVM 3.0

Last change: 26 Mar 1999

1

#### Maintenance Commands

vxdisksetup(1M)

-i Writes a disk header to the disk, making the disk directly usable (as a new disk in a disk group, for example).

#### ATTRIBUTES

You can specify attributes to affect the layout strategy used by vxdisksetup.

CAUTION: Be very careful when using these attributes as they can render a disk unusable by the Volume Manager.

config Sets up kernel logs or configuration databases on the disk. config is the converse of the noconfig attribute and is the default. config is ignored unless the -i option is specified.

All lengths and offsets are rounded up to cylinder boundaries due to restrictions on the layout of partitions.

noconfig Prevents setting up kernel logs or configuration databases on the disk. The size of the default private region partition is set to 80 blocks, which is the minimum allowed private region size. noconfig is ignored unless the -i option is specified.

privlen=length

Specifies the length of the private region partition of the disk. The default for this area is 1024 sectors.

privoffset=[-]offset

Indicates the sector offset of the private area on the disk. The default offset for the private area is at the beginning of the disk. A negative offset can be specified to relocate the private region at an offset relative to the end of the disk.

publen=length

Specifies the length of the public region partition of the disk. This defaults to the size of the disk minus the private area on the disk.

puboffset=offset

Sets the offset on the disk where the public region partition starts. This defaults to the end of the private region partition unless the private region partition is moved from the beginning of the disk, in which case the public region offset defaults to following the private region partition

SEE ALSO

vxdisk(1M), vxdiskadd(1M), vxintro(1M)

===[ man1m/vxdiskunsetup.1m ]=====

Maintenance Commands

vxdisksetup(1M)

NAME

vxdisksetup, vxdiskunsetup - configure a disk for use with the Volume Manager

SYNOPSIS

/etc/vx/bin/vxdisksetup [-i ] diskaddress [ attribute... ]

/etc/vx/bin/vxdiskunsetup [-C ] diskaddress..

DESCRIPTION

The vxdisksetup command configures the disk for use by the Volume Manager. vxdisksetup is called by the vxdiskadd command as part of initial disk configuration. It is not typi-

cally necessary to call this command directly. `vxdisksetup` creates two partitions on the specified disk, a Volume Manager private configuration region (called the private region partition), and a partition used for allocating sub-disks (called the public region partition). These two partitions have tags that identify them as appropriate to the Volume Manager.

The `vxdiskunsetup` command reverses the configuration done by `vxdisksetup` and makes the specified disks unusable by the Volume Manager. `vxdiskunsetup` can be applied only to disks that are not in use within an imported disk group.

The `diskaddress` argument passed to `vxdisksetup` or `vxdiskunsetup` directs the command to the disk being operated upon. The address should be provided in the form `c # t # d #`, which define a controller number (`c#`), a SCSI target ID (`t#`), and a SCSI logical unit number (`d#`). Each number is specified in decimal and can be multiple digits. The name must reference a valid disk with partition devices under the `/dev/vx/rdmp` directory.

The `vxdiskunsetup` command operates by removing the public and private regions that were created by the last invocation of `vxdisksetup` on the specified disks. After this operation, those disks are converted from the online to the error state (as displayed by `vxdisk list`). The command does not operate on disks that are active members of an imported disk group.

#### OPTIONS

`-C` `vxdiskunsetup` usually does not operate on disks that appear to be imported by some other host (for example, a host that shares access to the disk). In this case, the `-C` option can be specified to force de-configuration of the disk, effectively removing the host locks that were detected by `vxdiskunsetup`.

`-i` Writes a disk header to the disk, making the disk directly usable (as a new disk in a disk group, for example).

#### ATTRIBUTES

You can specify attributes to affect the layout strategy used by vxdisksetup.

**CAUTION:** Be very careful when using these attributes as they can render a disk unusable by the Volume Manager.

**config** Sets up kernel logs or configuration databases on the disk. config is the converse of the noconfig attribute and is the default. config is ignored unless the -i option is specified.

All lengths and offsets are rounded up to cylinder boundaries due to restrictions on the layout of partitions.

**noconfig** Prevents setting up kernel logs or configuration databases on the disk. The size of the default private region partition is set to 80 blocks, which is the minimum allowed private region size. noconfig is ignored unless the -i option is specified.

**privlen=length**  
Specifies the length of the private region partition of the disk. The default for this area is 1024 sectors.

**privoffset=[-]offset**  
Indicates the sector offset of the private area on the disk. The default offset for the private area is at the beginning of the disk. A negative offset can be specified to relocate the private region at an offset relative to the end of the disk.

**publen=length**  
Specifies the length of the public region partition of the disk. This defaults to the size of the disk minus the private area on the disk.

**puboffset=offset**  
Sets the offset on the disk where the public region partition starts. This defaults to the end of the private region partition unless the private region partition is moved from the beginning of the disk, in which case the public region offset defaults to following the private region partition

SEE ALSO

`vxdisk(1M)`, `vxdiskadd(1M)`, `vxintro(1M)`

```
===[ man1m/vxdmpadm.1m ]=====
```

Maintenance Commands

vxdmpadm(1M)

#### NAME

vxdmpadm - DMP subsystem administration

#### SYNOPSIS

```
/etc/vx/bin/vxdmpadm disable [ attribute...]  
  
/etc/vx/bin/vxdmpadm enable [ attribute...]  
  
/etc/vx/bin/vxdmpadm getdmpnode [ attribute...]  
  
/etc/vx/bin/vxdmpadm getsubpaths [ attribute...]  
  
/etc/vx/bin/vxdmpadm listctlr [ attribute...]  
  
/etc/vx/bin/vxdmpadm start restore [ attribute...]  
  
/etc/vx/bin/vxdmpadm stop restore
```

#### DESCRIPTION

The vxdmpadm utility is an administrative interface to the VERITAS Volume Manager Dynamic Multipathing (DMP) facility. It lists the paths under a DMP device, gets the DMP device corresponding to a path, lists all the disk controllers on the system, lists all the paths through a host disk controller, lists all the DMP nodes through a disk array, and enables or disables a host disk controller on the system.

vxdmpadm keywords select the action to perform, and attributes specify the objects on which the action is to be performed. For example, if the value of the boxno attribute is set to 3739a42622, information about the disk array with serial number "3739a42622" is displayed. Keyword type specifies the supported disk array types. Attributes are accepted in uppercase or lowercase.

#### KEYWORDS

**disable** Disables I/O to a host disk controller. This prevents DMP from issuing I/Os through the specified controller. Depending on the number of pending I/Os, this command can take some time to complete before returning control to the user. The command is blocked until all pending I/Os issued through the specified disk controller are completed.



Before detaching a system board all I/Os to the disk controller(s) connected to the board must be stopped. Execute the `vxdumpadm disable` command before executing the Dynamic Reconfiguration(DR) facility provided by Sun. The command must be executed for every controller connected to the system board being detached, followed by the DR

## Maintenance Commands

`vxdumpadm(1M)`

detach facility.

The disable operation fails if it is issued to a controller connected to the root disk through a single path.

This keyword supports the attributes: `ctlr`, `boxno`, and `type`.

`enable` Enables I/O to a host disk controller that was previously disabled. `enable` allows operations on the paths through this controller. This option succeeds only if the controller is available to the system and I/O can be done on this controller, or if one or more paths through this controller can be enabled. The serial number of the disk array, and the type of the disk array, can be specified with this keyword to operate on a disk array within the DMP database. Note: The disk controller name, disk array serial number, and type of disk array, if available, can be obtained by the `listctlr` keyword option.

### `getdmpnode`

Gets the DMP node that controls a particular path by setting the `nodename` attribute to a valid path under the `/dev/rdisk` directory. `getdmpnode` can also obtain all DMP nodes under a particular disk array by specifying the serial number of the disk array in the `boxno` attribute.

### `getsubpaths`

Gets the paths controlled by a DMP node. The `dmpnodename` specified must be a valid node in the `/dev/vx/rdmp` directory. `getsubpaths` can also obtain all the paths through a particular host

disk controller if the `ctlr` attribute specifies a controller name.

`listctlr` Lists disk controllers on the system. The `all` attribute specifies the controllers to list. Other attributes such as `boxno`, `type`, or `both` can be specified to list controllers on a particular disk array, or on a particular disk array type.

`start restore`

Starts the DMP restore daemon. The restore daemon analyzes the health of paths every interval seconds. The type of analysis depends on the policy specified by the user.

If the policy is `check_disabled`, it will check the

health of paths that were previously disabled due to hardware failures and revive them if they are back online. If the policy is `check_all`, it will analyze all paths in the system and revive the paths that are back online as well as disable the paths that are inaccessible. The default policy is `check_disabled`.

The policy can be set as shown in the example:

```
vxdmpadm start restore policy=check_all
```

If controller is disabled using the command:

```
vxdmpadm disable ctlr=ctlr
```

the restore daemon does not change the disabled state of the path through this controller.

Set the interval as shown in the example:

```
vxdmpadm start restore interval=400
```

The default interval is 300 seconds. Decreasing this interval could adversely affect system per-

formance.

In order to change the interval or policy the restore daemon must be stopped and restarted with new attributes specified.

stop restore

Stops the DMP restore daemon. Automatic path fail-back halts after the restore daemon is stopped.

#### ATTRIBUTES

Attribute values can be specified with keywords in the form attribute=value.

all Can be specified with listctlr to extract information about all host I/O controllers on the system. No other values can be specified with it. Specifying all results in all the disk controllers on the system being listed along with controller such as which disk array type the controller belongs to or the state of the controller.

boxno Specifies a disk array serial number. To obtain information about a disk array, set boxno to the

VxVM 3.0

Last change: 26 Mar 1999

3

Maintenance Commands

vxdmppadm(1M)

disk array cabinet serial number. This can be obtained by the listctlr all command.

boxno can be used with the getdmpnode, listctlr, enable, and disable keywords.

ctlr Specifies the host I/O controller about which information is to be printed. For example, set ctlr to /pci@1f,4000/scsi@3 to list the information for the controller /pci@1f,4000/scsi@3

You can combine the ctlr attribute with the get-subpaths, listctlr, enable, and disable keywords.

dmpnodename

Specifies the DMP node name. This node must be a valid node under the /dev/vx/rdmp directory.

interval Specifies the DMP restore daemon polling interval in seconds. The daemon checks the health of the

paths after every specified number of seconds.

**nodename** Specifies the name of the path. This node must be a valid node under the /dev/rdisk directory.

This attribute can be used with the getdmpnode keyword.

**policy** Specifies the policy to be used by the restore daemon. Value can be set to check\_all or check\_disabled.

**type** Specifies the disk array type. The disk array type must be one of the disk arrays supported by DMP. See the VERITAS Volume Manager Hardware Application Note for information on DMP supported disk arrays.

type is used with the keywords: getsubpaths, listctlr, enable and disable keywords.

#### EXIT CODES

vxddmpadm exists with a zero status if the attempted operation succeeds. A non-zero exit code is not a complete indicator of the problems encountered, but rather denotes the first condition that prevented further execution of the utility.

See vxintro(1M) for a list of standard exit codes.

#### SEE ALSO

vxintro(1M), vxdmp(7)

VxVM 3.0

Last change: 26 Mar 1999

5

===[ man1m/vxedit.1m ]=====

Maintenance Commands

vxedit(1M)

NAME

vxedit - create, remove, and modify Volume Manager records

#### SYNOPSIS

```
vxedit [-dfGprsv ] [-e pattern] [-g diskgroup] cc  
/search/replace/ [gp] [name...]
```

```
vxedit [-dpsv ] [-g diskgroup] rename oldname newname
```

```
vxedit [-dfGprsv ] [-g diskgroup] rm name...
```

```
vxedit [-dfGprsv ] [-e pattern] [-g diskgroup] set  
attribute=value... [name...]
```

#### DESCRIPTION

The vxedit utility sets and changes attributes for Volume Manager configuration records that do not depend upon volume usage types. See vxvol(1M) for operations that can set attributes that are dependent upon usage types. In particular, setting the length and logging type for a volume requires use of the vxvol set operation.

Each invocation can be applied to only one disk group at a time, due to internal implementation constraints. Any name or oldname operands will be used as record names to determine a default disk group, according to the standard disk group selection rules described in vxintro(1M). If no name or oldname operands are given, then the disk group defaults to rootdg. A specific disk group can be forced with -g diskgroup.

#### KEYWORDS

cc Change a comment using a search-replacement specification similar to that used by sed in volume, plex, subdisk, disk media, or disk group records within the selected disk group. The records to be changed are those that match the pattern specified with -e pattern option and those specified by the name operands. See vxintro(1M) for a description of Volume Manager search patterns. If no search pattern is specified with -e, and no name operands are given, then the change is made to all records whose comment field matches the search regular expression.

The search string is a regular expression, in the form accepted by the function regcmp(3G). This regular expression is used to determine which substring of the comment field is to be changed. The replace string represents the new string to use as a replacement for the matched part of the comment.

An ampersand (&) in the replace string is replaced by the substring of the comment matched by the regular expression. An occurrence of \n in the replace string, where n is a single digit between 1 and 9, will be replaced by the substring matched by a parenthetical section of the regular expression; the regular expression is followed by \$n.

The / character following the replace string is optional. If the / is given, then it can be followed by the letters g or p, or both. If a g is given, then all matches in a comment are replaced, rather than just the first match. If the letter p is given, then the resulting comment strings are written to the standard output, immediately preceded (on the same line) by the name of the record.

If the -r option is given, the operation is applied recursively to records associated with the selected records (to plexes and subdisks for selected volume records, and to subdisks for selected plex records). Recursion (when selected) applies regardless of the -p, and -s options.

Each record to be changed is changed only once, even if the record could be matched several times through combinations of name arguments, search patterns, and the -r option.

For example, the following command changes all subdisk comments that begin with ``Henry'' and a second word beginning with an uppercase letter to begin with ``Frank'' and the same second word:

```
vxedit -s cc '/^Henry ([A-Z])$1/Frank \1/p'
```

This command also lists the resulting comment fields.

- |        |   |
|--------|---|
| rename | Change the name of a volume, plex, subdisk, or disk media record from oldname to newname. A record cannot be renamed if the tutil0 field is set, which indicates that an operation is in progress that involves the record. |
| rm     | Remove volume, plex, or subdisk records from the selected disk group. Disk media records can be removed with vxdg rmdisk. Disk access records can   |

be removed with `vxdisk rm`.

Maintenance Commands

`vxedit(1M)`

Removing a subdisk requires that the subdisk be dissociated. Removing a plex requires that the plex be dissociated and that it have no associated subdisks. Removing a volume requires that it have no associated plexes. The `-r` option can be specified to recursively remove a volume and all plex and subdisk records associated with it, or to remove a plex and all subdisk records associated with it. If the `-r` options is provided, subvolumes are also removed. Even when removing with `-r`, a named plex or subdisk cannot be associated with a volume or plex, respectively.

The `-f` option is required to remove an enabled volume. A volume cannot be removed, even with `-f`, if the corresponding volume block or raw device, or if any of the volume's plex devices, is open or mounted. Likewise, a plex cannot be removed if the corresponding plex device is open.

`set` Set a field within a volume, plex, subdisk, disk media, or disk group record in the selected disk group. The records to be changed are those that match the pattern specified with the `-e` pattern option and those specified by the name operands.

The attribute names specify the field to set within the selected records. More than one attribute can be specified in a single invocation. The operands that indicate attribute settings end at the first operand that does not contain an equal sign. An operand of `--` can be used to separate the attribute list from record names, even if the first record name contains an equal sign.

If the `-r` option is given, the operation is applied recursively to records associated with the selected records (to plexes, subdisks, and subvolumes for selected volume records, and to subdisks and subvolumes for selected plex records). Recursion applies regardless of whether the `-p` and `-s` options are specified.



The following values can be set for all record types except for disk access records:

comment or c

Set the comment string for the selected records to the given value. The comment string cannot be longer than 40 characters and cannot contain a newline character.

VxVM 3.0

Last change: 26 Mar 1999

3

Maintenance Commands

vxedit(1M)

putil0, putil1, putil2

Set one of the non-persistent (temporary) utility fields in the record. The six utility fields cannot be longer than 14 characters and cannot contain a newline character. The tutil0 and putil0 fields are reserved for use by the the base Volume Manager utility set, and by usage types. The tutil1 and putil1 fields are reserved for use by higher-level utilities. The tutil2 and putil2 fields are reserved for any desired use by system administrators.

The putil0 field for a plex and subdisk record can be set to prevent utilities from associating the record to a volume or plex. This is a convenient means for reserving a plex, or for reserving a region of disk space (a subdisk).

Changing a non-empty putil0 or tutil0 field requires use of -f. Setting these fields for dissociated plex and subdisk records is generally not a problem. However, for an associated plex or subdisk that is associated (through a plex) with a volume, setting either of these fields can be dangerous, as it can affect the operation of usage types that expect to manage these fields themselves during an operation.

tutil0, tutil1, tutil2

Set one of the non-persistent (temporary) utility fields in the record.

The following attributes can only be set for volume records:

`fstype`

If a volume contains a file system, `fstype` can be used to indicate what type of file system it is. Under most circumstances, if a file system type is not specified for a volume, the Volume Manager will determine the usage type by running the `fstyp` program (see `fstyp(1M)`). However, it is preferable to set `fstype`, to avoid problems when the `fstyp` program returns ambiguous results.

`group`

Set the group that owns a volume record to the group specified as the attribute value.

Maintenance Commands

`vxedit(1M)`

The attribute value can be either a group name from the group database, or a numeric group ID.

`mode` Set the access permissions for the volume to the permission mode specified in the attribute value. The attribute value can be a symbolic permission mode or an octal mode. The format is compatible with permission modes as used by the `chmod` utility (see `chmod(1)`).

`user` Set the user that owns a volume record to the user specified as the attribute value. The attribute value can be either a login name from the `/etc/passwd` database, or a numeric user ID.

`writeback`

Set (on) or clear (off) a volume policy that affects recovery after read failures on a mirrored volume. If the `writeback` flag is set (which is normally the default), then a read failure for a plex will cause data to be read from an alternate plex and then written back to the plex that got the read failure. This will usually fix the error. Only if the

writeback fails will the plex be detached for having an unrecoverable I/O failure.

If this flag is clear, then data from an alternate plex will be read to satisfy the volume read operation, but the failing plex will be detached with no action taken to try to fix the problem.

There is seldom (if ever) a reason to turn off this feature.

#### writecopy

Set (on) or clear (off) a volume policy that affects consistency of data written to a volume when dirty region logging is in effect on the volume. When the operating system hands off a write request to the volume driver, the operating system may continue to change the memory that is being written to disk. The Volume Manager cannot detect that the memory is changing, so it can inadvertently leave plexes with inconsistent contents. This is not normally a problem, because the operating system ensures that any

such modified memory is rewritten to the volume before the volume is closed, such as by a clean system shutdown. However, if the system crashes, plexes may be inconsistent. Since the dirty region logging feature prevents recovery of the entire volume, it may not ensure that plexes are entirely consistent.

Setting the writecopy flag often causes the Volume Manager to copy the data for a write request to a new section of memory before writing it to disk. Because the write is done from the copied memory, it can't change, so the data written to each plex is guaranteed to be the same. Default is writecopy off.

The following special attributes can be set for

## subdisk records:

**len** Set the length of the subdisk to the given length. The attribute value is a standard Volume Manager length number (see vxintro(1M)). The length of a subdisk can be changed only if the subdisk is dissociated. The length of a subdisk cannot be increased to the point where it would extend past the end of the disk, or to where it would overlap a reserved disk region or another subdisk.

The following special attribute can be set for disk media records:

### failing

Set (on) or clear (off) the disk failing flag. If the failing flag is set for a disk, then the disk space is not used as free space or used by the hot relocation facility.

### reserve

Set (on) or clear (off) the disk reservation flag. If the reserve flag is set for a disk, then vxassist will not allocate a subdisk on that disk unless the disk is specified on the vxassist command line.

### spare

Set (on) or clear (off) the disk spare flag. If the spare flag is set for a disk, then that disk is designated for use by the hot-relocation facility. A DM record with the

spare flag set will be used only for hot-relocation. vxassist will not allocate a subdisk on that disk unless forced to by command line arguments.

## OPTIONS

**-d, -G, -p, -s, -v**

Select only disk media, disk group, plex, subdisk, or volume records, respectively. If more than one of these options are specified, records of any of the indicated types may be selected.

- e pattern  
Use a Volume Manager configuration search expression to select records from the selected disk group configuration. Search patterns are limited to a selection of volume, plex, and subdisk records.
  
- f  
Force an operation that the Volume Manager considers potentially dangerous or is not a normal operation for the command. This enables a limited set of operations that would otherwise be disallowed. Some operations may be disallowed even with this flag. The vxedit operations that are allowed with this flag are changing a non-empty tutil0 or putil0 field, and removing enabled volumes.
  
- g diskgroup  
Specify the disk group for the operation, either by disk group ID or by disk group name. By default, the disk group is chosen based on the name and oldname operands.
  
- r  
Operate recursively on records associated with the selected records. For selected volume records, this affects associated plex, subdisk, and subvolumes records. For selected plex records, this affects associated subdisk and subvolume records.

EXIT CODES

vxedit exits with a nonzero status if the operation fails. A nonzero exit code is not a complete indicator of the problems encountered, but rather denotes the first condition that prevented further execution of the utility.

See vxintro(1M) for a list of standard exit codes.

SEE ALSO

chmod(1), fstyp(1), sed(1), vxedit(1M), vxintro(1M), vxmake(1M), vxmend(1M), vxvol(1M), regcmp(3G)

=== [ man1m/vxencap.1m ] =====

## NAME

vxencap - encapsulate partitions on a new disk

## SYNOPSIS

```
/etc/vx/bin/vxencap [-c ] [-g diskgroup] [-N ]  
[medianame=]accessname
```

```
/etc/vx/bin/vxencap -k -a | accessname
```

## DESCRIPTION

The vxencap script identifies any partitions on the specified disk that could be used for file systems or special areas such as swap devices, and then generates volumes to cover those areas on the disk. This process is known as encapsulation. In the case of partitions that contain file systems, the entries in /etc/vfstab will be modified to use the new volume devices.

vxencap is generally called from the vxdiskadm menus or else is called at installation time from the vxinstall script. A reboot may be necessary to complete the process of encapsulation.

## OPTIONS

- a Specifies that all encapsulation requests are to be handled by this command. This option is only useful with the -k flag specified.
- c Creates a disk group if the specified disk does not exist. Only useful if a disk group is also supplied using the -g option.
- g diskgroup Places the encapsulated disk into the disk group specified by disk group ID or disk group name.
- k Aborts encapsulation requests. If used with the -a flag, all encapsulation requests are aborted. Otherwise, specify a list to abort.
- N Specifies that the changes logged to be made during the next reboot are not committed. Usually any system failure before the reboot causes the operations to be completed on the next reboot. With the -N flag specified, the changes are left uncompleted and the operation aborted should the system fail at any point before the successful configuration change. This is accomplished by not creating the usual /etc/vx/reconfig.d/state.d/reconfig commit file.

Maintenance Commands

vxencap(1M)

SEE ALSO

vxdiskadm(1M), vxintro(1M)

VERITAS Volume Manager (VxVM) Installation Guide

===[ man1m/vxevac.1m ]=====

Maintenance Commands

vxevac(1M)

#### NAME

vxevac - evacuate all volumes from a disk

#### SYNOPSIS

```
/etc/vx/bin/vxevac [-g diskgroup] [-t tasktag] medianame  
[new_medianame...]
```

#### DESCRIPTION

The vxevac utility moves subdisks off the specified Volume Manager disk (medianame) to the specified destination disks (new\_medianame...). If no new\_medianame operands are specified, then any non-volatile, non-reserved disks can be used as destination disks. Subdisks that are part of unmirrored striped plexes are moved by moving the entire plex to a new location.

This operation assumes there is sufficient space in the disk group for it to complete the operation. If the process runs out of space, then some of the volumes on the disk may not be evacuated.

This utility is usually called from the vxdiskadm menus.

A medianame argument is an administrative name used to define a disk within a disk group.

#### OPTIONS

-g diskgroup

Specifies the disk group for the operation, either by disk group ID or by disk group name.

-t tasktag

Specifies an administrative operation to perform by the numeric identifier tasktag. See the vxtask(1M) man page for information on VxVM tasks.

#### SEE ALSO

vxdiskadm(1M), vxintro(1M), vxtask(1M)



===[ man1m/vxinfo.1m ]=====

Maintenance Commands

vxinfo(1M)

#### NAME

vxinfo - print accessibility and usability of volumes

#### SYNOPSIS

vxinfo [-pV] [-g diskgroup] [-o useopt] [-U usetype]  
[volume...]

#### DESCRIPTION

The vxinfo utility reports a usage-type-dependent condition on one or more volumes in a disk group. A report for each volume specified by the volume operand is written to the standard output. If no volume operands are given, then a volume condition report is provided for each volume in the selected disk group.

Each invocation can be applied to only one disk group at a time, due to internal implementation constraints. Any volume operands will be used to determine a default disk group, according to the standard disk group selection rules described in vxintro(1M). A specific disk group can be forced with -g diskgroup.

#### OPTIONS

-g diskgroup

Specify the disk group for the operation, either by disk group ID or by disk group name. By default, the disk group is chosen based on the volume operands. If no volume operands are speci-

fied, the disk group defaults to rootdg.

- o useopt Pass in usage-type-specific options to the operation.
- p Report the name and condition of each plex in each reported volume.
- U usetype Specify the usage type for the operation. If no volume operands are specified, then the output is restricted to volumes with this usage type. If volume operands are specified, then this will result in a failure message for all named volumes that do not have the indicated usage type.
- V Write a list of utilities that would be called from vxinfo, along with the arguments that would be passed. The -V performs a ``mock run'' so the utilities are not actually called.

#### Volume Conditions

The volume condition is a usage-type-dependent summary of the state of a volume. This condition is derived from the

#### Maintenance Commands

vxinfo(1M)

volume's kernel-enabled state and the usage-type-dependent states of the volume's plexes. The vxinfo utility reports the following conditions for volumes:

**Startable** A vxvol startall operation would likely succeed in starting the volume.

**Started** The volume has been started and can be used.

#### Started Unusable

The volume has been started but is not operationally accessible. This condition may result from errors that have occurred since the volume was started, or may be a result of administrative actions, such as vxdg -k rmdisk.

#### Unstartable

The volume is not started and either is not correctly configured or doesn't meet the prerequisites for automatic startup (with volume

startup) because of errors or other conditions.

#### OUTPUT FORMAT

Summary reports for each volume are printed in one-line output records. Each volume output line consists of blank-separated fields for the volume name, volume usage type, and volume condition. The following example shows the volume summary:

```
bigvol      fsgen      Startable
vol2        fsgen      Startable
brokenvol   gen        Unstartable
```

Each plex output line consists of blank-separated fields for the plex name and the plex condition. The plex records are accompanied by their volume records, as the following example shows:

```
vol bigvol      fsgen      Startable
plex bigvol-01  ACTIVE
vol vol2        fsgen      Startable
plex vol2-01    ACTIVE
vol brokenvol   gen        Unstartable
```

#### FSGEN and GEN Usage Types

The fsgen and gen usage types provide identical semantics for the vxinfo utility. The fsgen and gen usage types do not support any options passed in with -o. Plex conditions (reported with -p) can be one of the following:

#### Maintenance Commands

vxinfo(1M)

ACTIVE	Either the volume is started and the plex is enabled, or the volume was not stopped cleanly and the plex was valid when the volume was stopped.
CLEAN	The plex contains valid data and the volume was stopped cleanly.
EMPTY	The plex is part of a volume that has not yet been initialized.
IOFAIL	The plex was detached because of an uncorrectable I/O failure on one of the subdisks in the plex.

NODAREC No physical disk was found for one of the subdisks in the plex. This implies either that the physical disk failed, making it unrecognizable, or that the physical disk is no longer attached through a known access path.

OFFLINE The plex was disabled using the vxmend off operation.

REMOVED A physical disk used by one of the subdisks in the plex was removed through administrative action with vxdg -k rmdisk.

SNAPATT The plex is being attached as part of a backup operation by the vxassist snapstart operation. When the attach is complete, the condition will change to SNAPDONE. A system reboot or manual starting of the volume will remove the plex and all of its subdisks.

SNAPDIS A vxassist snapstart operation completed the process of attaching the plex. It is a candidate for selection by the vxplex snapshot operation. A system reboot or manual starting of the volume will dissociate the plex.

SNAPDONE A vxassist snapstart operation completed the process of attaching the plex. It is a candidate for selection by the vxassist snapshot operation. A system reboot or manual starting of the volume will remove the plex and all of its subdisks.

SNAPTMP The plex is being attached as part of a backup operation by the vxplex snapstart operation. When the attach is complete, the condition will change to SNAPDIS. A system reboot or manual starting of the volume will dissociate the plex.

STALE The plex does not contain valid data, either as a

result of a disk replacement affecting one of the subdisks in the plex, or as a result of an administrative action on the plex such as vxplex det.

TEMP The plex is associated temporarily as part of a current operation, such as vxplex cp or vxplex att. A system reboot or manual starting of a volume will dissociate the plex.

TEMPRM The plex was created for temporary use by a current operation. A system reboot or manual starting of a volume will remove the plex.

TEMPRMSD The plex and its subdisks were created for temporary use by a current operation. A system reboot or manual starting of the volume will remove the plex and all of its subdisks.

Volume conditions for these usage types are reported as follows:

Startable This condition is reported if the volume is not enabled and if any of the plexes have a reported condition of ACTIVE or CLEAN.

Started This condition is reported if the volume is enabled and at least one of the associated plexes is enabled in read-write mode (which is normal for enabled plexes in the ACTIVE and EMPTY conditions).

Started Unusable This condition is reported if the volume is enabled, but the volume does not meet the criteria for being Started.

Unstartable This condition is reported if the volume is not enabled, but the volume does not meet the criteria for being Startable.

#### RAID-5 Usage Type

ACTIVE Either the volume is started and the plex is enabled, or the volume was not stopped cleanly and the plex was valid when the volume was stopped.

CLEAN The plex contains valid data and the volume was stopped cleanly. The raid5 usage type does not support any options passed in with -o.

Plexes of RAID-5 volumes can be either data plexes (that is, RAID-5 plexes) or log plexes. RAID-5 plex conditions are as

follows:

- IOFAIL The plex was detached from use as a result of an uncorrectable I/O failure on one of the subdisks in the plex.
- NODAREC No physical disk was found for one of the subdisks in the plex. This implies either that the physical disk failed, making it unrecognizable, or that the physical disk is no longer attached through a known access path.
- REMOVED A physical disk used by one of the subdisks in the plex was removed through administrative action with `vx dg -k rmdisk`.
- OFFLINE The plex was disabled using the `vxmend off` operation.

RAID-5 plexes can also have these additional conditions:

- DEGRADED Due to subdisk failures, the plex is in degraded mode. This indicates a loss of data redundancy in the RAID-5 volume and any further failures could cause data loss.
- STALEPRTY The parity is not in sync with the data in the plex. This indicates a loss of data redundancy in the RAID-5 volume and any further failures could cause data loss.
- UNUSABLE This indicates a double-failure occurred within the plex. The plex is unusable due to subdisk failures and/or stale parity.

Log plexes of RAID-5 volumes can have the following conditions:

- ACTIVE Either the volume is started and the plex is enabled, or the volume was not stopped cleanly and the plex was valid when the volume was stopped.
- CLEAN The plex contains valid data and the volume was stopped cleanly.
- IOFAIL The plex was detached from use as a result of an uncorrectable I/O failure on one of the subdisks in the plex.
- NODAREC No physical disk was found for one of the subdisks in the plex. This implies either that the physical disk failed, making it unrecognizable, or that

## Maintenance Commands

vxinfo(1M)

the physical disk is no longer attached through a known access path.

OFFLINE The plex was disabled using the vxmend off operation.

REMOVED A physical disk used by one of the subdisks in the plex was removed through administrative action with vxdg -k rmdisk.

Log plexes of RAID-5 volumes can also have the following additional condition:

BADLOG The contents of the plex are not usable as logging data.

Volume conditions for volumes of the raid5 usage type are the same as for the fsgen and gen usage types (Startable, Unstartable, Started and Started Unusable). In addition, the following conditions can modify the conditions:

Degraded This condition indicates that the RAID-5 plex of the volume is in degraded mode due to the unavailability of a subdisk in that plex.

Staleprty This condition indicates that some of the parity in the RAID-5 plex is stale and requires recovery.

## FILES

/etc/vx/type/usetype/vxinfo The utility that performs vxinfo operations for a particular volume usage type.

## EXIT CODES

The vxinfo utility exits with a nonzero status if the attempted operation fails. A nonzero exit code is not a complete indicator of the problems encountered, but rather denotes the first condition that prevented further execution of the utility. See vxintro(1M) for a list of standard exit codes.

## SEE ALSO

vxassist(1M), vxintro(1M), vxmend(1M), vxplex(1M), vxsd(1M), vxvol(1M)

===[ manlm/vxintro.1m ]=====

Maintenance Commands

vxintro(1M)

#### NAME

vxintro - introduction to Volume Manager utilities

#### SYNOPSIS

vxapslice, vxassist, vxbootsetup, vxconfigd, vxdctl, vxdg, vxdisk, vxdiskadd, vxdiskadm, vxedit, vxencap, vxevac, vxinfo, vxiod, vxmake, vxmend, vxmirror, vxnotify, vxplex, vxprint, vxr5check, vxreattch, vxrecover, vxrelayout, vxrelocd, vxresize, vxrootmir, vxsd, vxsparecheck, vxstat, vxtask, vxtrace, vxunroot, vxvol

#### DESCRIPTION

The Volume Manager utilities provide a shell-level interface for use by system administrators and high-level applications and scripts to query and manipulate objects managed through the VERITAS Volume Manager (VxVM).

#### GLOSSARY

concatenated plex

A plex whose subdisks are associated at specific offsets within the address range of the plex, and extend into the plex address range for the length of the subdisk. This layout allows regions of one or more disks to create a plex, rather than a single big region.

disk

Disks exist as two entities. One is the physical disk on which all data is ultimately stored and which exhibits all the behaviors of the underlying technology. The other is the Volume Manager presentation of disks which, while mapping one-to-one with the physical disks, are just presentations of units from which allocations of storage are made. As an example, a physical disk presents the image of a device with a definable geometry



with a definable number of cylinders, heads etc. whereas a Volume Manager disk is simply a unit of allocation with a name and a size.

#### disk access record

A configuration record that defines a pathway to a disk. The disk access records most commonly name a particular controller number, target ID, and logical unit number. The list of all disk access records stored in a system is used to find all disks attached to the system. Disk access records do not identify particular physical disks.

Disk access records are identified by their disk access names (also known as DA names).

Through the use of disk IDs, the Volume Manager

allows disks to be moved between controllers, or to different locations on a controller. When a disk is moved, a different disk access record will be used when accessing the disk, although the disk media record will continue to track the actual physical disk.

On some systems, the Volume Manager will build a list of disk access records automatically, based on the list of all devices attached to the system. On these systems, it is not necessary to define disk access records explicitly. On other systems, disk access records must be defined explicitly with the `vxdisk define` operation. Specialty disks (such as RAM disks or floppy disks) are likely to require explicit `vxdisk define` operations on all systems.

#### disk group

A group of disks that share a common configuration. A configuration consists of a set of records describing objects (including disks, volumes, plexes, and subdisks) that are associated with one particular disk group. Each disk group has an administrator-assigned name that can be used by the administrator to reference that disk group. Each disk group has an internally defined

unique disk group ID, which is used to differentiate two disk groups with the same administrator-assigned name.

Disk groups provide a method to partition the configuration database, so that the database size is not too large and so that database modifications do not affect too many drives. They also allow the Volume Manager to operate with groups of physical disk media that can be moved between systems.

Disks and disk groups have a circular relationship: disk groups are formed from disks, and disk group configurations are stored on disks. All disks in a disk group are stamped with a disk group ID, which is a unique identifier for naming disk groups. Some or all disks in a disk group also store copies of the configuration database of the disk group.

#### disk group configuration

A disk group configuration is a small database that contains all volume, plex, subdisk, and disk media records. These configurations are

replicated onto some or all disks in the disk group, usually with one copy on each disk. Because these databases are stored within disk groups, record associations cannot span disk groups. Thus, a subdisk defined on a disk in one disk group cannot be associated with a volume in another disk group.

#### disk header

A block stored in a private region of a disk and that defines several properties of the disk. The disk header defines the size of the private region, the location and size of the public region, the unique disk ID for the disk, the disk group ID and disk group name (if the disk is currently associated with a disk group), and the host ID for a host that has exclusive use of the disk.

disk ID A 64-byte universally unique identifier that is assigned to a physical disk when its private region is initialized with the `vxdisk init` operation. The disk ID is stored in the disk media record so that the physical disk can be related to the disk media record at system startup.

#### disk media record

A reference to a physical disk or a disk partition. This record can be thought of as a physical disk identifier for the disk or partition. Disk media records are configuration records that provide a name (known as the disk media name or DM name) that you can use to reference a particular disk, independent of its location on the system's various disk controllers. Disk media records reference particular physical disks through a disk ID, which is a unique identifier that is assigned to a disk when it is initialized for use with the Volume Manager.

Operations are provided to set or remove the disk ID stored in a disk media record. Such operations have the effect of removing or replacing disks, with any associated subdisks being removed or replaced along with the disk.

host ID A name, usually assigned by the administrator, that identifies a particular host. Host IDs are used to assign ownership to particular physical disks. When a disk is part of a disk group that is in active use by a particular host, the disk is stamped with that host's host ID. If another

system attempts to access the disk, it will detect that the disk has a non-matching host ID and will disallow access until the first system discontinues use of the disk. To allow for system failures that do not clear the host ID, the `vxdisk clearimport` operation can be used to clear the host ID stored on a disk.

If a disk is a member of a disk group and has a host ID that matches a particular host, then that host will import the disk group as part of system

startup.

kernel log

A log kept in the private region on the disk and that is written by the Volume Manager kernel. The log contains records describing the state of volumes in the disk group. This log provides a mechanism for the kernel to persistently register state changes so that vxconfigd can be guaranteed to detect the state changes even in the event of a system failure.

layered volume

A virtual volume used and managed directly by the Volume Manager, not by a user. A layered volume provides storage for an upper-level volume, which can itself be a layered volume. Up to four levels of volumes can be layered. At the top level is a user-controlled volume. In VxVM 3.0, a layered volume provides storage for only one upper-level volume.

Layered volumes allow the Volume Manager to implement the following features:

- o Striped mirrors and concatenated mirrors are new types of volume layouts that are less likely to fail and that provide faster recovery times because there is less data to recover (one column or part of a column versus the whole volume).
- o Online relayout lets you change the volume layout while the volume is online. See the vxassist(1M) and vxrelayout(1M) man pages for more information.
- o RAID-5 subdisk moves and RAID-5 snapshots (see vxassist(1M)).

plex

A copy of a volume's logical data address space,

also called a mirror. A volume can have up to 32 plexes associated with it. Each plex is, at least conceptually, a copy of the volume that is main-

tained consistently in the presence of volume I/O and reconfigurations. Plexes represent the primary means of configuring storage for a volume. Plexes can have a striped, concatenated, or RAID-5 organization (layout).

#### plex consistency

If the plexes of a volume contain different data, then the plexes are said to be inconsistent. This is only a problem if the Volume Manager is unaware of the inconsistencies, as the volume can return differing results for consecutive reads. Plex inconsistency is a serious compromise of data integrity. This inconsistency can be caused by write operations that start around the time of a system failure, if parts of the write complete on one plex but not the other. Plexes can also be inconsistent after creation of a mirrored volume, if the plexes are not first synchronized to contain the same data. An important part of Volume Manager operation is ensuring that consistent data is returned to any application that reads a volume. This may require that plex consistency of a volume be ``recovered'' by copying data between plexes so that they have the same contents. Alternatively, the volume can be put into a state such that reads from one plex are automatically written back to the other plexes, thus making the data consistent for that volume offset.

#### private region

Disks used by the Volume Manager contain two special regions: a private region and a public region. Usually, each region is formed from a complete partition of the disk; however, the private and public regions can be allocated from the same partition.

The private region of a disk contains various on-disk structures that are used by the Volume Manager for various internal purposes. Each private region begins with a disk header which identifies the disk and its disk group. Private regions can also contain copies of a disk group's configuration, and copies of the disk group's kernel log.

#### public region

The public region of a disk is the space reserved

for allocating subdisks. Subdisks are defined with offsets that are relative to the beginning of the public region of a particular disk. Only one contiguous region of disk can form the public region for a particular disk.

root disk group

Each system requires one special disk group, named rootdg, which is generally the default for most utilities. In addition to defining the regular disk group information, the configuration for the root disk group (rootdg) contains local information that is specific to a disk group and that is not intended to be movable between systems.

striped plex

A plex that scatters data evenly across each of its associated subdisks. A plex has a characteristic number of stripe columns (consisting of associated subdisks) and a characteristic stripe unit size. The stripe unit size defines how data with a particular address is allocated to one of the associated subdisks. Given a stripe unit size of 128 blocks and two stripe columns, the first group of 128 blocks is allocated to the first subdisk, the second group of 128 blocks is allocated to the second subdisk, the third group to the first subdisk, and so on.

subdisk

A region of storage allocated on a disk for use with a volume. Subdisks are associated to volumes through plexes. One or more subdisks are laid out to form plexes based on the plex layout (striped, concatenated, or RAID-5). Subdisks are defined relative to disk media records.

volboot file

The volboot file is a special file (usually stored in `/etc/vx/volboot`) that is used to bootstrap the root disk group and to define a system's host ID. In addition to a host ID, the volboot file contains a list of disk access records. On system startup, this list of disks is scanned to find a disk that is a member of the rootdg disk group and that is stamped with this system's host ID. When such a disk is found, its configuration is read and is used to get a more complete list of disk access records that are used as a second-stage bootstrap of the root disk group, and to locate all other disk groups.

volume

A virtual disk device that looks to applications

## Maintenance Commands

vxintro(1M)

and file systems like a regular disk partition device. Volumes present block and raw device interfaces that are compatible in their use with disk partition devices. However, a volume is a virtual device that can be mirrored, spanned across disk drives, moved to use different storage, and striped using administrative commands. The configuration of a volume can be changed, using the Volume Manager utilities, without causing disruption to applications or file systems that are using the volume.

## CONVENTIONS

The VERITAS Volume Manager employs certain conventions to provide a degree of similarity between various operations. The conventions are used in the following areas:

- o Disk Group Selection
  
- o Standard Length Numbers
  
- o Command Syntax

## Disk Group Selection

Most commands operate upon only one disk group per invocation. Each disk group has a separate configuration from every other disk group and it is possible for two disk groups to contain two objects that have the same name. This can happen, in particular, if a disk group is moved from one system to another. However, most utilities make no attempt to ensure that names between disk groups are unique, so name collisions can occur.

System administrators who try to avoid name collisions should be able to use most of the utilities without having to specify disk groups except when creating objects. Administrators cannot use single-command invocations that reference objects in more than one disk group, but disk groups will be selected automatically, based on objects specified in the command.

The standard rules that most commands use for selecting the disk group for a command are as follows:

- o Given a particular set of object names specified on the command, look for the disk group of each object. If all objects are in the same disk group, use that disk group. If any named object is not unique between all disk

## Maintenance Commands

vxintro(1M)

groups, and if one of those object names is not in the rootdg disk group, then fail.

- o To force use of a particular disk group, use `-g diskgroup` to indicate the group. Non-unique names do not cause errors when a disk group is specified explicitly. The `diskgroup` specification is either a disk group ID or a disk group name.
- o A special case is provided for the `rootdg` disk group. Any set of objects in the `rootdg` disk group can be specified without specifying `-g rootdg`, even if the given object name is used in another disk group.

If a set of object names is given on the command line, and if some are unique but some are not unique, then the command will still fail according to the rules listed above.

## Standard Length Numbers

Many basic properties of objects that are managed by the Volume Manager require specification of lengths, either as a pure object length or as an offset relative to some other object. The Volume Manager supports volume lengths up to 2,147,483,647 disk sectors (one terabyte on most systems). Typing such large numbers, or even much smaller numbers, can be annoying. The Volume Manager provides a uniform syntax for representing such numbers, and uses suffixes to provide convenient multipliers. Numbers can be specified in decimal, octal, or hexadecimal. Also, numbers can be specified as a sum of several numbers, as a convenience to avoid using a calculator.



A hexadecimal (base 16) number is introduced using a prefix of 0x. For example, 0xffff is the same as decimal 4095. An octal (base 8) number is introduced using a prefix of 0. For example, 0177777 is the same as decimal 65535.

A number can be followed by a single-character suffix to indicate a multiplier for the number. A length number with no suffix character represents a count of standard disk sectors. The length of a standard disk sector can vary between systems; it is typically 512 bytes or 1024 bytes. On systems where disks can have different sector sizes, one of the sector sizes is chosen as the ``standard'' size. Supported suffix characters are:

b      Multiply the length by 512 bytes (blocks)

#### Maintenance Commands

vxintro(1M)

s      Multiply the length by the standard sector size (default)

k      Multiply the length by 1024 bytes (Kilobytes)

m      Multiply the length by 1,048,576 (1024K) bytes (Megabytes)

g      Multiply the length by 1,073,741,824 (1024M) bytes (Gigabytes)

Numbers are represented internally as an integer number of sectors. As a result, if the standard disk sector size is larger than 512 bytes, numbers can be specified that will need to be rounded to a sector. Rounding is always done to the next lowest, not the nearest, multiple of the sector size.

Because the letter b is a valid hexadecimal character, there is a special case for the b suffix where a single blank character can separate a number from the b suffix character. Use of a blank within a number, when invoking commands from the shell, usually requires quoting the number. For example:

```
vxassist make vol01 "0x1000 b"
```

Numbers can be added or subtracted by separating two or more

numbers by a plus or minus sign, respectively. A plus sign is optional. As an example, the largest allowed number that can be represented on a system with a 512 byte sector size can be entered as:

```
1023g+1023m+1023k+1
```

Note that 1024g-1 cannot be used, because the implementation cannot handle the intermediate representation of 1024g (which is greater than the largest number that can be represented) internally.

In output, the Volume Manager reports length numbers as a simple count of sectors, with no suffix character.

Case is ignored in length specification. Hexadecimal numbers and suffix characters can be specified using any reasonable combination of uppercase and lowercase letters.

#### Utility Syntax

Most utilities in the Volume Manager provide more than one operation, with operations grouped into utilities primarily

#### Maintenance Commands

vxintro(1M)

by object type. Utilities that provide multiple operations are typically invoked with the following form:

```
utility_name [ options ] keyword [ operands ]
```

utility\_name is the name of the Volume Manager command and keyword specifies the specific operation to perform. Any options that are introduced in the standard -letter form precede the operation keyword. This is in keeping with standard System V UNIX utility syntax, which provides for a set of options as the first arguments to a command, followed by any non-option operands. The keyword is considered an operand under standard System V syntax.

All of the utilities provide an extended usage message that lists all the options and operation keywords supported by the utility. For utilities that are keyword-based, this extended usage message can be displayed using the keyword help. Utilities that use operands for purposes other than operation selection provide a reserved option of -H to

display the extended usage information.

#### RECORD TYPES

Disk group configurations contain six types of records:

- o Disk Access Records
  
- o Disk Group Records
  
- o Disk Media Records
  
- o Plex Records
  
- o Subdisk Records
  
- o Volume Records

Disk access records are specific to the root disk group and are stored in configurations only because there is no other convenient place to store them; otherwise, they are logically separate from all disk groups. Because they are

specific and meaningful to the local system only, the logical place for their storage is the rootdg since that is the only disk group guaranteed to exist on the system.

#### Disk Access Records

Disk access records define an address, or access path, that can be used to access a disk. The list of all disk access records defines the list of all disk addresses that the Volume Manager can use to locate physical disks. Disk access records do not define specific physical disks, since physical disks can be moved on a system. When a physical disk is moved, a different disk access record may be necessary to locate it.

Disk access records are stored in the rootdg disk group configuration. Unlike all other record types, the names of disk access records can conflict with the names of other records. For example, a specialty disk (such as a RAM disk) can use the same name for both the disk access record and the disk media record that points to it. It is typically advisable to use different names for the access and media records, to avoid additional confusion if disks are moved.

Disk access records can be defined explicitly. Some (sometimes all) disk access records may be configured automatically by the Volume Manager, based on available information in the operating system. Such automatically-configured disks are not stored persistently in the on-disk root disk group configuration, but are instead regenerated every time the Volume Manager starts up.

Disk access records have the following fundamental attributes:

disk access name

The name of the disk access record is typically a disk address. Disk access names are usually of the form `c#t#d#s2`, indicating use of the entire disk on controller `c#`, with SCSI target ID `t#`, and logical unit `d#`. The `s2` suffix is used as a convention indicating the use of standard partitioning. Other systems are likely to have different conventions for the disk access name.

type

Each disk access record has a type, which identifies key characteristics of the Volume Manager's interaction with the disk. Available types are: `sliced`, `simple`, and `nopriv`. See `vxdisk(1M)` for more information on disk types. Typically, most or all of the disks will be of type `sliced`. It may be desirable to create specialty disks (such as RAM disks) with type `nopriv`.

If the physical disk represented by the disk access record is associated with a disk media record, then the following fields are defined:

disk group name

The name of the disk group containing the disk media record.

disk media name

The name of the disk media record that points to the physical disk.

Additional attributes can be added, arbitrarily, by disk types. See `vxdisk(1M)` for a list of additional attributes defined by the standard disk types.

#### Disk Group Records

Disk group records define several different types of names for a disk group. The different types of names are:

alias name

This is the standard name that the system uses when referencing the disk group. References to the disk group name usually mean the alias name. Volume directories are structured into subdirectories based on the disk group alias name. Typically, the disk group's alias name and real name are identical. A local alias can be useful for gaining access to a disk group with a name that conflicts with other disk groups in the system, or that conflicts with records in the `rootdg` disk group.

disk group ID

A 64-byte identifier that represents the unique ID of the disk group. All disk groups on all systems should have a different disk group ID, even if they have the same real name. This identifier is stored in the disk headers of all disks in the disk group. It is used to ensure that the Volume Manager does not confuse two disk groups which were created with the same name.

real name This is the name of the disk group, as defined on disk. This name is stored in the disk group configuration, and is also stored in the disk headers of all disks in the disk group.

#### Disk Media Records

Disk media records define a specific disk within a disk group. The name of a disk media record (the disk media name) is assigned when a disk is first added to a disk group

(using the vxdbg adddisk operation). Disk media records can be assigned to specific physical disks by associating the disk media record with the current disk access record for the physical disk.

Disk media records have the following fundamental attributes:

disk access name

The disk access name that is currently used to access the physical disk referenced by the disk ID. If the disk ID is defined, but no physical disk with that ID could be found, the disk access name will be clear. A disk where the physical disk could not be found is considered to be in the NODAREC, or inaccessible, state. A disk can become inaccessible either because the indicated disk is not currently attached to the system, or because I/O failures on the physical disk prevented the Volume Manager from identifying or using the physical disk.

disk ID A 64-byte unique identifier representing the physical disk to which the media record is associated. This can be cleared to indicate that the disk is considered in the removed state. A removed disk has no current association with any physical disk.

A disk media record that has an active association with a physical disk (both the disk ID and the disk access name attributes are defined), inherits several properties from the underlying physical disk. These attributes are taken from the disk header, which is stored in the private region of the the disk. These inherited attributes are:

atomic I/O size

This is the fundamental I/O size for the disk, in bytes, also known as the sector size. All I/O operations to this disk must be multiples of this size. Volume Manager requires that all disks have the same sector size. On most systems, this size is 512 bytes.

private length

The length of the region of the physical disk that is reserved for storing private Volume Manager information.

public length

The length of the region of the physical disk that is available for subdisk allocations.

## Maintenance Commands

vxintro(1M)

## Plex Records

Plex records define the characteristics of a particular plex of a volume. A plex can be in either an associated state or a dissociated state. In the dissociated state, the plex is not a part of a volume. A dissociated plex cannot be accessed in any way. An associated plex can be accessed through the volume.

Plexes have the following fundamental attributes:

**comment** An administrator-assigned string of up to 40 characters that can be set and changed using the `vxedit` utility. The Volume Manager does not interpret the comment field. The comment cannot contain newline characters.

**condition flags**

Various condition flags are defined for the plex that define state which is recognized automatically, rather than managed by the volume usage type. Defined flags are:

**IOFAIL** The plex was detached as a result of an I/O failure detected during normal volume I/O. The plex is out-of-date with respect to the volume, and in need of complete recovery. However, this condition also indicates a likelihood that one of the disks in the system should be replaced.

**NODAREC** No physical disk could be found corresponding to the disk ID in the disk media record for one of the subdisks associated with the plex. The plex cannot be used until the condition is fixed or the affected subdisk is dissociated.

**RECOVER** A disk for one of the disk media records was replaced or was reattached too late to prevent the plex from becoming out-of-date with respect to the volume. The plex requires complete recovery from another plex in the volume to synchronize the plex with the correct contents of the volume.

REMOVED One of the disk media records was put into the removed state through explicit administrative action. The plex cannot be used until the disk is replaced or the affected subdisk is dissociated.

Maintenance Commands

vxintro(1M)

contiguous length

The offset of the first block in the plex address space that is not backed by a subdisk. If the plex has no holes, the contiguous length matches the plex length. If the contiguous length is equal to or greater than the length of the associated volume, the plex is considered complete, otherwise it is sparse.

I/O mode

Each plex is in read-write, read-only, or write-only mode. This mode affects read and write operations directed to the volume, if the plex is enabled. For read-write and read-only modes, volume read operations can be directed to the plex. For read-write and write-only modes, volume write operations are directed to the plex.

Plexes are normally in read-write mode. Write-only mode is used to recover a plex that failed, and whose contents have thus become out-of-date with respect to the volume. It is also used when attaching a new plex to a volume. In write-only mode, writes to the volume will update the plex, causing written regions to be up-to-date. Typically, a set of special copy operations will be used to update the remainder of the plex.

layout

The organization of associated subdisks with respect to the plex address space. The layout is striped, concatenated, or RAID-5.

length

The length of a plex is the offset of the last subdisk in the plex plus the length of that subdisk. In other words, the length of the plex is defined by the last block in the plex address space that is backed by a subdisk. This value may or may not relate to the length of the volume, depending on whether the plex is completely contiguously allocated.



#### log subdisk

Each plex can have at most one associated log subdisk. A log subdisk is used with the dirty region logging feature to improve the time required to recover consistency of a volume after a system failure.

#### plex state

Each plex is either enabled, disabled, or detached. When enabled, normal read and write operations from the volume can be directed to the plex. When disabled, no I/O operations will be

#### Maintenance Commands

vxintro(1M)

applied to the plex. When detached, normal volume I/O will not be directed to the plex.

I/O failures encountered during normal volume I/O may move the enabled state for a plex directly from enabled to detached. See the description of volume exception policies (earlier in this manual page) for more information.

#### subdisks

Each plex has zero or more associated subdisks. Subdisks are associated at offsets relative to the beginning of the plex address space. Subdisks for concatenated plexes may not cover the entire length of the plex, in which case they leave holes in the plex. A plex that is not as long as the volume to which it is associated is considered to have a hole extending from the end of the plex to the end of the volume. A plex with a hole is considered incomplete, and is sometimes called sparse.

#### usage-type state

Volume usage types maintain a private state field related to the the operations that have been performed on the plex, or to failure conditions that have been encountered. This state field contains a string of up to 14 characters.

#### volatile state

A plex is considered to have ``volatile'' contents if the disk for any of the plex's subdisks is con-

sidered to be volatile. The contents of a volatile disk are not presumed to survive a system reboot. The contents of a volatile plex are always considered out-of-date after a recovery and in need of complete recovery from another plex.

#### Subdisk Records

Subdisk records define a region of disk, allocated from a disk's public region. Subdisks have very little state associated with them, other than the configuration state that defines which region of disk the subdisk occupies. Subdisks cannot overlap each other, either in their associations with plexes, or in their arrangement on disk public regions.

Subdisks have the following fundamental attributes:

comment    An administrator-assigned string of up to 40 characters that can be set and changed using the vxedit utility. The Volume Manager does not interpret the comment field. The comment cannot contain newline characters.

#### Maintenance Commands

vxintro(1M)

disk media name

The name of the disk media record that the subdisk is defined on.

disk offset

The offset, from the beginning of the disk's public region, to the start of the subdisk.

length      The length of the subdisk.

plex offset

For associated subdisks, this is the offset (from the beginning of the plex) of the subdisk association. For subdisks associated with striped plexes, the plex offset defines relative ordering of subdisks in the plex, rather than actual offsets within the plex address space.

#### Volume Records

Volume records define the characteristics of particular volume devices. The name of a volume record defines the node name used for files in the /dev/vx/dsk and /dev/vx/rdsk directories. The block device for a particular volume

(which can be used as an argument to the mount command (see mount(1M)) has the path:

```
/dev/vx/dsk/groupname/volume
```

where groupname is the name assigned by the administrator to the disk group containing the volume. The raw device for a volume, typically used for application I/O and for issuing I/O control operations (see ioctl(2)), has the path:

```
/dev/vx/rdsk/groupname/volume
```

For convenience, volumes assigned to the root disk group are accessible under the rootdg subdirectories of /dev/vx/dsk and /dev/vx/rdsk, but are also under /dev/vx/dsk/volume and /dev/vx/rdsk/volume.

Reads to a volume device are directed to one of the read-write or read-only plexes associated with the volume. Writes to the volume are directed to all of the enabled read-write and write-only plexes associated with the volume.

During a write operation, two plexes of a volume may become out of sync with each other, due to the fact that writes directed to two disks can complete at different times. This is not normally a problem. However, if the system were to crash or lose power during a write operation, the two plexes

could have different contents.

Most applications and file systems are not written with the presumption that two separate reads of a device can return different contents without an intervening write operation. Since plexes with different contents could cause such a situation where two read operations of a block return different contents, the Volume Manager expends considerable effort to ensure that this is avoided.

Volumes have the following fundamental attributes:

comment An administrator-assigned string of up to 40 characters that can be set and changed using the vxedit utility. The Volume Manager does not

interpret the comment field. The comment cannot contain newline characters.

exception policy

There are several modes that can be set on the volume, by utilities according to the usage type of the volume. These modes affect operation of a volume in the presence of I/O failures. Currently only one of these policies, called GEN\_DET\_SPARSE is ever used. This policy tracks complete and incomplete plexes in a volume (an incomplete plex does not have a backing subdisk for all blocks in the volume). If an unrecoverable error occurs on an incomplete plex, the plex is detached (disabled from receiving regular volume I/O requests). If an unrecoverable error occurs on a complete plex, the plex is detached unless it is the last complete plex. If the plex is the last complete read-write plex, any incomplete plexes that overlap with the error will be detached but the plex with the error will remain attached.

This default policy is chosen to ensure that an I/O that fails on one plex will not, in the future, be directed to that plex again unless that plex is the last complete plex remaining attached to the volume. In that case, the policy ensures that the volume will return the error consistently, even in the presence of incomplete plexes.

length

Each volume has a length, which defines the limiting offset of read and write operations. The length is assigned by the administrator, and may or may not match the lengths of the associated plexes.

log type A policy to use for logging changes to the volume, which can be assigned by the administrator. Policies that can be specified are:

none Do not perform any special actions when writing to the volume. Just write the requested data to all read-write or

write-only plexes.

dirty-region-log

A volume is divided into regions. A bit-map where each bit corresponds to a region is maintained. When a write to a particular region occurs, the respective bit is set to on. When the system is restarted after a crash, this region bitmap is used to limit the amount of data copying that is required to recover plex consistency for the volume. The region changes are logged to special log subdisks associated with each of the plexes associated with the volume. Use of dirty region logging can greatly speed recovery of a volume, but it also degrades performance of the volume under normal operation.

plexes      Each volume has between zero and 32 associated plexes.

read policy

A configurable policy for switching between plexes for volume reads. When a volume has more than one enabled associated plex, the Volume Manager can distribute reads between the plexes to distribute the I/O load and thus increase total possible bandwidth of reads through the volume. The read policy can be set by the administrator. Possible policies are:

preferred plex

This read policy specifies a particular named plex that is used to satisfy read requests. In the event that a read request cannot be satisfied by the preferred plex, this policy changes to round-robin.

round-robin

For every other read operation, distribute the operation across all of the available plexes. Given three plexes,

this will switch between each of the three plexes, so each plex gets one third of the read requests.

select This read policy is the default policy, and adjusts to use an appropriate read policy based on the set of plexes associated with the volume. If exactly one enabled read-write striped plex is associated with the volume, then that plex is chosen automatically as the preferred plex; otherwise, the round-robin policy is used. If a volume has one striped plex and one non-striped plex, preferring the striped plex often yields better throughput.

#### read/write-back recover mode

This is a mode that applies to the volume, which is managed by utilities as part of plex consistency recovery. When this mode is enabled, each read operation will recover plex consistency for the region covered by the read. Plex consistency is recovered by reading data from blocks of one plex and writing that data to all other writable plexes. This ensures that a future read operation covering the same range of blocks will read the same data.

#### start options

This is a string that is organized as a set of usage-type options to apply when starting (enabling) a volume. See vxvol(1M) for details.

#### usage type

Each volume has a usage type, which defines a particular class of rules for operating on the volume, typically based on the expected content of the volume. Several utilities can apply extensions or limitations that apply to volumes with a particular usage type. Several usage types are included with the base release of the Volume Manager: fsgen, for use with volumes that contain file systems; gen, for use with volumes that are used as swap devices or for other applications that do not use file systems; and special root and swap usage types which are specifically for use with the root file system volume and the primary swap device.

#### usage-type state

Usage types maintain a private state field related

to the volume that relate to operations that have been performed on the volume, or to failure conditions that have been encountered. This state field contains a string of up to 14 characters.

user|group|mode

These attributes are the user, group, and file permission modes used for the volume device nodes. The user and group are normally root. The mode usually allows read and write permission to the owner, and no access by other users.

volume state

Each volume is either enabled, disabled, or detached. When enabled, normal read and write operations are allowed on the volume, and any file system residing on the volume can be mounted, or used in the usual way. When disabled, no access to the volume or any of its associated plexes is allowed.

write-back-on-read-failure mode

This is a mode that applies to the volume, which can be enabled or disabled by the administrator using vxedit. If this mode is enabled, then a read failure for a plex will cause data to be read from an alternate plex and then written back to the plex that got the read failure. This will usually fix the error. Only if the writeback fails will the plex be detached for having an unrecoverable I/O failure.

writecopy mode

This is a mode that applies to the volume, which can be enabled or disabled by the administrator using vxedit. This mode takes affect only if dirty region logging is in effect. When the operating system hands off a write request to the volume driver, the operating system may continue to change the memory that is being written to disk. The Volume Manager cannot detect that the memory is changing, so it can inadvertently leave plexes with inconsistent contents. This is not normally a problem, because the operating system ensures that any such modified memory is rewritten to the volume before the volume is closed (such as by a clean system shutdown). However, if the system crashes, plexes may be inconsistent. Since the dirty region logging feature prevents recovery

of the entire volume, it may not ensure that plexes are entirely consistent.

Maintenance Commands

vxintro(1M)

Turning on the writecopy mode (which is normally set by default) often causes the Volume Manager to copy the data for a write request to a new section of memory before writing it to disk. Because the write is done from the copied memory, it cannot change and so the data written to each plex is guaranteed to be the same if the write completes.

VOLUME USAGE TYPES

The usage type of a volume represents a class of rules for operating on a volume. Each usage type is defined by a set of executables under the directory /etc/vx/type/usage\_type, where usage\_type is the name given to the usage type. The required executables are: vxinfo, vxmake, vxmend, vxplex, vxsd, and vxvol. These executables are invoked by the Volume Manager administrative utilities with the same names. The executables under /etc/vx/type should not, normally, be executed directly.

Five usage types are provided with the Volume Manager: gen, fsgen, root, swap, and raid5. It is possible for third-party products to install additional usage types.

The usage types provided with the Volume Manager store state information in the volume and plex usage-type state fields. The state fields defined for volumes are:

- ACTIVE      The volume has been started and is running normally, or was running normally when the system was stopped. If the system crashes in this state, then the volume may require plex consistency recovery.
- CLEAN      The volume has been stopped and the contents for all plexes are consistent.
- EMPTY      The volume is not yet initialized. This is the initial state for volumes created by vxmake.
- NEEDSYNC   The volume requires recovery. This is typically set after a system failure to indicate that the



plexes in the volume may be inconsistent, so that they require recovery (see the resync operation in vxvol(1M)).

SYNC Plex consistency recovery is currently being done on the volume. vxvol resync sets this state when it starts to recovery plex consistency on a volume that was in the NEEDSYNC state.

The state fields defined for plexes are:

VxVM 3.0

Last change: 15 Dec 1999

22

Maintenance Commands

vxintro(1M)

ACTIVE The plex is running normally on a started volume. The plex condition flags (NODAREC, REMOVED, RECOVER, and IOFAIL) may apply if the system is rebooted and the volume restarted.

CLEAN The plex was running normally when the volume was stopped. The plex will be enabled without requiring recovery when the volume is started.

EMPTY The plex is not yet initialized. This state is set when the volume state is also EMPTY.

OFFLINE The plex was disabled by the vxmend off operation. See vxmend(1M) for more information.

SNAPATT This is a snapshot plex that is being attached by the vxassist snapstart operation. When the attach is complete, the state for the plex will be changed to SNAPDONE. If the system fails before the attach completes, the plex and all of its sub-disks will be removed.

SNAPDIS This is a snapshot plex created by vxplex snapstart that is fully attached. A plex in this state can be turned into a snapshot volume with vxplex snapshot. See vxplex(1M) for more information. If the system fails before the attach completes, the plex will be dissociated from the volume.

SNAPDONE This is a snapshot plex created by vxassist snapstart that is fully attached. A plex in this state can be turned into a snapshot volume with

vxassist snapshot. See vxassist(1M) for more information. If the system fails before the attach completes, the plex and all of its subdisks will be removed.

SNAPTMP This is a snapshot plex being attached by the vxplex snapstart operation. When the attach is complete, the state for the plex will be changed to SNAPDIS. If the system fails before the attach completes, the plex will be dissociated from the volume.

STALE The plex was detached, either by vxplex det or by an I/O failure. vxvol start will change the state for a plex to STALE if any of the plex condition flags are set. STALE plexes will be reattached automatically, when starting a volume, by calling vxplex att.

#### Maintenance Commands

vxintro(1M)

TEMP This is a plex that is being associated and attached to a volume with vxplex att. If the system fails before the attach completes the plex will be dissociated from the volume.

TEMPRM This is a plex that is being associated and attached to a volume with vxplex att. If the system fails before the attach completes the plex will be dissociated from the volume and removed. Any subdisks in the plex will be kept.

TEMPRMSD This is a plex that is being associated and attached to a volume with vxplex att. If the system fails before the attach completes, the plex and its subdisks will be dissociated from the volume and removed.

#### EXIT CODES

The majority of the Volume Manager utilities use a common set of exit codes, which can be used by shell scripts or other types of programs to react to specific problems detected by the utilities. For C programmers, these exit status codes are defined in the include file volclient.h. The number and macro name for each distinct exit code is described below. Shell script writers must directly compare

against the numbers specified.

- (0) VEX\_OK  
The utility is not reporting any error through the exit code.
- (1) VEX\_USAGE  
Some command line arguments to the utility were invalid.
- (2) VEX\_SYNTAX  
A syntax error occurred in a command or description, or a specified record name is too long or contains invalid characters. This code is returned only by utilities that implement a command or description language. This code may also be returned for errors in search patterns.
- (3) VEX\_NOVOLD  
The volume daemon does not appear to be running.
- (4) VEX\_IPC  
An unexpected error was encountered while communicating with the volume daemon.
- (5) VEX\_OSERR  
An unexpected error was returned by a system call

or by the C library. This can also indicate that the utility ran out of memory.

- (6) VEX\_LOST  
The status for a commit was lost because the volume daemon was killed and restarted during the commit of a transaction, but after restart the volume daemon did not know whether the commit succeeded or failed.
- (7) VEX\_UTILERR  
The utility encountered an error that it should not have encountered. This generally implies a condition that the utility should have tested for but did not, or a condition that results from the volume daemon returning a value that did not make sense.

(8) VEX\_TIMEOUT

The time required to complete a transaction exceeded 60 seconds, causing the transaction locks to be lost. As most utilities will reattempt the transaction at least once if a timeout occurs, this usually implies that a transaction timed out two or more times.

(9) VEX\_NODG

No disk group could be identified for an operation. This results either from naming a disk group that does not exist, or from supplying names on a command line that are in different disk groups or in multiple disk groups.

(10) VEX\_CHANGED

A change made to the database by another process caused the utility to stop. This code is also returned by a usage-type-dependent utility if it is given a record that is associated with a different usage type. If this situation occurs when the usage-type-dependent utility is called from a switchout utility, then the database was changed after the switchout utility determined the proper usage type to invoke.

(11) VEX\_NOENT

A requested subdisk, plex, or volume record was not found in the configuration database. This may also mean that a record was an inappropriate type.

(12) VEX\_EXIST

A name used to create a new configuration record matches the name of an existing record.

(13) VEX\_BUSY

A subdisk, plex, or volume is locked against concurrent access. This code is used for inter-transaction locks associated with usage type utilities. The code is also used for the dissociated plex or subdisk lock convention, which writes a non-blank string to the tutil[0] field in a plex or subdisk structure to indicate that the record is being used.

- (14) VEX\_NOUSETYPE  
No usage type could be determined for a utility that requires a usage type.
- (15) VEX\_BADUSETYPE  
An unknown or invalid usage type was specified.
- (16) VEX\_ASSOC  
A plex or subdisk is associated, but the operation requires a dissociated record.
- (17) VEX\_DISASSOC  
A plex or subdisk is dissociated, but the operation requires an associated record. This code can also be used to indicate that a subdisk or plex is not associated with a specific plex or volume.
- (18) VEX\_\_LAST  
A plex or subdisk was not dissociated because it was the last record associated with a volume or plex.
- (19) VEX\_TOOMANY  
Association of a plex or subdisk would surpass the maximum number that can be associated to a volume or plex.
- (20) VEX\_INVALID  
A specified operation is invalid within the parameters specified. For example, this code is returned when an attempt is made to split a subdisk on a striped plex, or to use a split size that is greater than the size of the plex.
- (21) VEX\_IOERR  
An I/O error was encountered that caused the utility to abort an operation.
- (22) VEX\_NOPLEX  
A volume involved in an operation did not have any associated plexes, although at least one was required.

- (23) VEX\_NOSUBDISK

A plex involved in an operation did not have any associated subdisks, although at least one was required.

- (24) VEX\_UNSTARTABLE  
A volume could not be started by the vxvol start operation, because the configuration of the volume and its plexes prevented the operation.
- (25) VEX\_STARTED  
A specified volume was already started.
- (26) VEX\_UNSTARTED  
A specified volume was not started. For example, this code is returned by the vxvol stop operation if the operation is given a volume that is not started.
- (27) VEX\_DETACHED  
A volume or plex involved in an operation is in the detached state, thus preventing a successful operation.
- (28) VEX\_DISABLED  
A volume or plex involved in an operation is in the disabled state, thus preventing a successful operation.
- (29) VEX\_ENABLED  
A volume or plex involved in an operation is in the enabled state, thus preventing a successful operation.
- (30) VEX\_UNKNOWN  
An unknown error was encountered. This code may be used, for example, when the volume daemon returns an unrecognized error number.
- (31) VEX\_OPEN  
An operation failed because a volume device was open or mounted, or because a subdisk was associated with an open or mounted volume or plex.
- (65) VEX\_ARRAY  
An array-related API failed.
- (66) VEX\_ARRAY\_INCOMPAT  
An array-specific guideline was violated.

Exit codes greater than 32 are reserved for use by usage types. Codes greater than 64 can be reserved for use by

specific utilities.

**SEE ALSO**

mount(1M), vxapslice(1M), vxassist(1M), vxbootsetup(1M),  
vxconfigd(1M), vxdctl(1M), vxdg(1M), vxdisk(1M),  
vxdiskadd(1M), vxdiskadm(1M), vxdisksetup(1M), vxdmpadm(1M),  
vxedit(1M), vxencap(1M), vxevac(1M), vxinfo(1M), vxiod(1M),  
vxlicense(1M), vxmake(1M), vxmend(1M), vxmirror(1M),  
vxnotify(1M), vxplex(1M), vxprint(1M), vxr5check(1M),  
vxreattach(1M), vxrecover(1M), vxrelayout(1M), vxrelocd(1M),  
vxresize(1M), vxrootmir(1M), vxsd(1M), vxsparecheck(1M),  
vxstat(1M), vxtask(1M), vxtrace(1M), vxunroot(1M),  
vxvol(1M), ioctl(2), plexrec(4), sdrec(4), vol\_pattern(4),  
volrec(4), vxmake(4)

For information about other aspects of the Volume Manager,  
see:

vxintro(3X), vxio(7), vxinfo(7)

VERITAS Volume Manager Getting Started Guide

===[ manlm/vxiod.lm ]=====

## Maintenance Commands

vxiod(1M)

### NAME

vxiod - start, stop, and report on Volume Manager kernel daemons

### SYNOPSIS

vxiod

vxiod [-f ] [set count]

### DESCRIPTION

The vxiod utility starts, stops, or reports on Volume Manager I/O daemons. An I/O daemon provides a process context for performing Volume Manager I/O. Volume Manager I/O daemons are not required for correct operation, but not having I/O daemons can adversely affect system performance.

When invoked with no arguments, vxiod prints the current number of volume I/O daemons on the standard output.

The number of daemons to create for general I/O handling is dependent on system load and usage. It is generally not necessary to start more than one daemon for each CPU on the system. If volume recovery seems to proceed slower at times, it may be worthwhile to create more daemons.

Each I/O daemon starts in the background and creates an asynchronously-running process, which detaches itself from the controlling terminal and becomes a volume I/O daemon. The vxiod utility does not wait for these processes to complete.

### KEYWORDS

set

When invoked with the set keyword, vxiod creates the number of daemons specified by count. If more volume I/O daemons exist than are specified by count, the excess processes terminate. If more than the maximum number are created, the requested number is silently truncated to that maximum.

### OPTIONS



-f Forces the kill of the last I/O daemon. Without this option, the I/O daemons can only be reduced to one.

#### EXIT CODES

The vxiod utility prints a diagnostic on the standard error, and exits if an error is encountered. If an I/O occurs within a forked I/O daemon process, then the I/O is not reflected in the exit status for vxiod. Otherwise, vxiod returns a nonzero exit status on errors.

VxVM 3.0

Last change: 26 Mar 1999

1

Maintenance Commands

vxiod(1M)

Usage errors result in an exit status of 1 and a usage message. If the requested number of daemons cannot be created, then the exit status is 2, and the number of daemons that were successfully started is reported. If any other error occurs, the exit status is 3.

#### FILES

/dev/vx/vxiod The device used to report on and start volume I/O daemon processes.

#### NOTES

Volume Manager I/O daemons cannot be killed directly through the use of signals.

The number of Volume Manager I/O daemons currently running can only be determined by running vxiod; I/O daemons do not appear in the list of processes produced by the ps command.

#### SEE ALSO

ps(1), vxconfigd(1M), vxdctl(1M), vxintro(1M), fork(2), vxio(7), vxiod(7)

===[ man1m/vxlicense.1m ]=====

Maintenance Commands

vxlicense(1M)

NAME

vxlicense - VERITAS licensing key utility

SYNOPSIS

vxlicense [-c] [-p] [-t feature]

DESCRIPTION

vxlicense, the VERITAS licensing key utility, maintains the VERITAS license key file.

OPTIONS

- c           Creates a license key file. vxlicense prompts for a license key.
  
- p           Prints available VERITAS licenses and features in a system.
  
- t feature   Tests a VERITAS license in a system.

EXAMPLES

The following example shows how to use vxlicense to initialize the VERITAS licensing key file. After entering the command, vxlicense prompts for a license key. It creates the license key file after you enter the key.

```
vxlicense -c
```

To list available features, enter:

```
vxlicense -p
```

```
Feature name: CURRSET [95]
Number of licenses: 1 (non-floating)
Expiration date: 12/31/99
Release Level: 20
Machine Class: 934986342
```

```
Feature name: RAID [96]
Number of licenses: 1 (non-floating)
Expiration date: 12/31/99
Release Level: 20
Machine Class: 934986342
```

VERITAS

Last change: 26 Mar 1999

1

===[ man1m/vxmake.1m ]=====

Maintenance Commands

vxmake(1M)

NAME

vxmake - create Volume Manager configuration records

SYNOPSIS

```
vxmake [-pV ] [-g diskgroup] [-o useopt] [-U usetype] [-T
a|o] [-d descfile]
```

```
vxmake [-pV ] [-g diskgroup] [-o useopt] plex name
[attribute=value...]
```

```
vxmake [-pV ] [-g diskgroup] [-o useopt] sd name  
[attribute=value...]
```

```
vxmake [-pV ] [-g diskgroup] [-o useopt] [-U usetype] vol  
name [attribute=value...]
```

#### DESCRIPTION

The vxmake utility creates subdisk, plex, and volume records for the Volume Manager. Records can be created entirely from parameters specified on the command line, or they can be created using a description file.

Note: The vxprint or vxprivutil utilities can be used with vxmake to perform the functions of a rudimentary configuration backup and restore. This should be the primary use of vxmake. Use of vxmake other than through a description file containing content mostly gathered from either the vxprint or the vxprivutil utilities is discouraged in Volume Manager releases starting with 3.0. The vxassist utility is the recommended means of creating new or modified volume hierarchies.

If no operands are specified, then a description file is used to specify what records to create. By default, this description file is read from the standard input. If operands are specified, then the first operand is a keyword that determines the object to build, the second operand is the name of the object to be created, and additional operands specify attributes for the object to be built. An attribute given on the command line is specified in the same manner used with the description file (with multiple attributes given, one attribute per operand), except that the double-quote convention is not needed. See the EXAMPLES section for sample uses.

For convenience, a special attribute form can be used for subdisks specified directly from the command line. This special attribute form is:

```
devicepath,offset[,len]
```

medianame,offset[,len]

The first form specifies a subdisk location by device path, offset within that device, and length. The second form specifies the subdisk location by a disk media record name, offset within the disk's public region, and length. Optionally, the length can be specified with a len=length attribute instead. If a device path is specified, then it must match the device path for the public region of a known disk in the indicated disk group.

More than one record, including more than one volume or plex hierarchy, can be specified in a single description file. vxmake attempts to create all records within a single transaction, so that either all records are created or no records are created. This is the default behavior.

A new option, -T allows users to select transaction control that differs from the default behavior when restoring records within a volume hierarchy from a description file. This may be particularly helpful with volume hierarchies that include layered volumes.

See vxmake(4) for a specification of vxmake descriptions.

Descriptions of records within a volume hierarchy are filtered through a usage-type-specific utility, as appropriate for the usage type, before the records are actually created. Subdisks and plexes that are not specified within a volume hierarchy are filtered by the gen usage type. The usage-type utilities may change the descriptions so that certain fields are ignored or set up with initial values. Plex or subdisk records to be associated with a volume or plex may also be changed if the association is indicated in the description, even if the plex or subdisk record is not explicitly specified.

A usage type must be specified for every volume record, either through setting the usetype field in the volume record description, or through the -U option to vxmake.

#### OPTIONS

-d descfile

Specify the name of a description file to use for building subdisks, plexes, and volumes. If the descfile option argument is -d, then the description file is read from the standard input. If no operands are specified, and no descfile is specified, then a description file is read from the

standard input. This option is ignored for records that are specified directly on the command line.

- g diskgroup  
Specify the disk group for the operation. The disk group can be specified either by name or by disk group ID. See vxdg(1M) for more information on disk groups.
- o useopt Give usage-type-specific options to the usage type utility.
- p Print a list of all changes that would be made to the standard output, but do not make those changes. The output is in the vxmake description file format. All computable and potentially useful attributes are printed for new records. For records that would be modified, only those attributes that would change are printed, along with an extra attribute declaration, CHANGED=yes. In this way the -p option performs a ``mock run'' of the utilities.

This option is intended for use with higher-level tools that preview record creations, particularly with respect to usage-type-dependent attribute modifications.

- T This option controls the unit of work that vxmake uses to make records within a single transaction when restoring records from a description file. If the option argument is a or the -T option is absent altogether, then vxmake attempts to create all records within a single transaction, so that either all records are created or no records are created. This behavior is considered the default behavior and it is backward-compatible with previous releases of the Volume Manager.

If the option argument is o, then one complete top-level object (Volume hierarchy, dissociated plex hierarchy, or dissociated subdisk hierarchy) is created at a time, and if multiple top-level objects are to be created from the same description file, then vxmake attempts to create each one a top-level object at a time in its own transaction. Note that should some objects be creatable and some not, the configuration results in those objects that vxmake was actually able to create. This option can be used for any record hierarchy. Because record hierarchies with layered volumes

## Maintenance Commands

vxmake(1M)

may contain many subordinate records, the `-T o` option is particularly useful for safely creating these kinds of objects.

`-U usetype`

Use `usetype` as the default usage type for volumes. This option has no effect if only subdisks and plexes are created, or if the usage type for a volume is specified directly with the description file variable `usetype`.

`-V`

Display a list of utilities that would be called from `vxmake`, along with the arguments that would be passed. The `-V` option performs a ``mock run'' of the utilities; the utilities are not actually called, and no changes are made to the volume configuration database.

## Attribute Specification Rules

The only attributes that must be specified are the `path` and `len` attributes for subdisk records, and a usage type for volume records. In addition, if the layout attribute for a plex record is set to `STRIPE`, then a positive value must be specified for the `stwidth` attribute. Also, if no subdisks are specified for a plex of the `STRIPE` layout, the `ncolumn` attribute must be specified.

Attempts to specify certain attributes are ignored by `vxmake`. Some attributes are ignored by the `switchout` (usage-type independent) `vxmake` itself. The usage-type dependent `vxmake` may cause additional attributes to be ignored.

The sections that follow provide information on the attributes that can be specified for subdisk, plex, and volume creation.

## Subdisks

Fields that can be specified for subdisks are: `tutil0`, `tutil1`, `tutil2`, `putil0`, `putil1`, `putil2`, `daname`, `dmname` (or `disk`), `path`, `comment`, `devoffset`, `dmoffset`, `len` and `dmrid`.

The disk specification for subdisks uses one of following fields, in order of precedence: `dmrid`, `dmname`, `daname`, or

path. The dmrid field is used, if non-zero. One of the other fields is used if defined to a non-empty string. If the path field is the only field used, it must define the block device path for a partition containing the public region of a disk in the requested disk group.

With path, the devoffset field is used as an offset into the partition; otherwise doffset is used as an offset into the

## Maintenance Commands

vxmake(1M)

private region. This distinction is important if the public region does not begin at the beginning of its partition.

Specification of a disk in some form (including possibly the special forms devicepath,offset,length or medianame,offset,length) is required. Specification of the len field is also required. All other fields default to zero (for numbers) or empty strings (for strings).

## Plexes

Fields that can be specified for plexes are: tutil0, tutil1, tutil2, putil0, putil1, putil2, comment, layout, ncolumn, stwidth, sd, and logsd.

There are no required fields for plex records. All fields default to zero (for numbers) or empty strings (for strings). The layout field defaults to concat.

If the layout field is set to stripe or raid5, then the stwidth attribute must be specified. If the layout field is set to concat, the ncolumn field cannot be specified.

An sd specification names the subdisks to associate with the plex. The specification optionally names the offsets of the subdisks within the plex. If no offsets are specified, then the subdisks are directly concatenated in the order indicated. If the sd attribute is specified for a striped plex and ncolumn is not specified, ncolumn will default to the column of the subdisk with the highest column specification. As a special case, if ncolumn is not specified and the sd does not specify columns or offsets for the subdisks, then each subdisk will be placed in its own column at an offset of zero in the order they are specified and the plex's ncolumn attribute is set to the number of subdisks specified.



A logsd specification names a subdisk to associate as a special log subdisk for recording volume activity as part of the dirty region logging feature.

#### Volumes

Fields that can be specified for volumes are: tutil0, tutil1, tutil2, putil0, putil1, putil2, comment, usetype, startopts, readpol, prefname, minor, user, group, mode, len, writeback, writecopy, exclusive, logtype, and plex.

Specification of a usage type is required, either using the -U option on the command line, or using the usetype attribute. Specification of a usetype attribute overrides a value specified with -U. All other fields have default values: readpol defaults to round; user and group both default to 0 (root); mode defaults to u=rw (mode 600);

VxVM 3.0

Last change: 26 Mar 1999

5

#### Maintenance Commands

vxmake(1M)

writeback and writecopy default to on; exclusive defaults to off; len defaults to the length of the shortest associated plex (or zero, if no plex associations are requested); all other numeric fields default to zero, and all other string fields default to empty strings.

A plex specification names the plexes to associate with the volume. The order of plexes in this list is not important.

The value of the prefname field is relevant only if the readpol attribute is set to prefer.

#### FSGEN and GEN Usage Types

The fsgen and gen usage-type-specific utilities that support vxmake are used by the switchout vxmake utility as filters that set defaults for plex and volume states.

Subdisk fields are set, checked, or modified as follows:

len            Subdisk lengths are required to be set to a nonzero value.

tutil0        This field is cleared.

Plex fields are set, checked, or modified as follows:

tutil0        This field is cleared.

state If the plex is to be associated with a volume, the state field is set to EMPTY; otherwise it is cleared.

Volume fields are set, checked, or modified as follows:

tutil0 This field is cleared.

state This field is set to EMPTY.

Attempts to associate an existing subdisk or plex with a new plex or volume will be refused, unless the putil0 field for the existing record is empty.

The putil0 fields are not set, and their contents are preserved.

#### RAID-5 Description

The raid5 usage-type-specific utilities that support vxmake are used by the switchout vxmake utility as filters that set defaults for plex and volume states.

Subdisk fields are set, checked or modified as follows:

#### Maintenance Commands

vxmake(1M)

len Subdisk lengths are required to be set to a nonzero value.

tutil0 This field is cleared.

Plex fields are set, checked, or modified as follows:

tutil0 This field is cleared.

state If the plex is to be associated with a volume, the state field is set to EMPTY; otherwise it is cleared.

Volume fields are set, checked, or modified as follows:

tutil0 This field is cleared.

state This field is set to EMPTY.

Attempts to associate an existing subdisk or plex with a new

plex or volume will be refused, unless the putil0 field for the existing record is empty.

The putil0 fields are not set, and their contents are preserved.

The raid5 specific vxmake utility will only allow one plex with a layout of raid5 to be associated with a volume. Plexes not of the raid5 layout are associated as RAID-5 log plexes.

#### EXAMPLES

The following is an example of a vxmake description file:

```
#rectyp  #name      #options
sd        disk3-01  disk=disk3 offset=0 len=10000
sd        disk3-02  disk=disk3 offset=25000 len=10480
sd        disk4-01  disk=disk4 offset=0 len=8000
sd        disk4-02  disk=disk4 offset=15000 len=8000
sd        disk4-03  disk=disk4 offset=30000 len=4480
plex      db-01      layout=STRIPE ncolumn=2 stwidth=16k

sd=disk3-01:0/0,disk3-02:0/10000,disk4-01:1/0,
      disk4-02:1/8000,disk4-03:1/16000
sd        ramd1-01  disk=ramd1 len=640
      comment="Hot spot for dbvol
plex      db-02      sd=ramd1-01:40320

vol       db        usetype=gen plex=db-01,db-02
      readpol=prefer prefname=db-02
      comment="Uses mem1 for hot spot in last 5m
```

This description specifies a gen type volume that contains two plexes: a volatile memory disk plex (db-02), which is preferred, and a physical disk plex (db-01). The memory disk plex is sparse and covers only the last 640 sectors of the 40960-sector length of the volume. The physical disk plex is striped across two 20480-sector columns. Column zero contains two subdisks and column one contains three subdisks. The plex offset specified for each subdisk concatenates the subdisks contiguously in their respective columns.

For striped or RAID-5 plex subdisk associations, if one

number is specified for the the column and column offset field, the number is interpreted as a column number and the subdisk is associated at the end of the column in the order it appears in the subdisk list. The same subdisk association list for plex db-01 could be re-written as follows:

```
sd=disk3-01:0,disk3-02:0,disk4-01:1,disk4-02:1,disk4-03:1
```

The following is a simple sequence of commands to create a hierarchy with one subdisk, one plex, and one volume:

```
vxmake sd disk07-01 disk=disk07 len=10000 comment="disk07
subdisk 1"
vxmake plex demo-1 sd=disk07-01 comment="Demo volume, plex
1"
vxmake -U gen vol demo plex=demo-1 comment="Demo gen
volume"
```

#### EXIT CODES

The vxmake utility exits with a nonzero status if the attempted operation fails. A nonzero exit code is not a complete indicator of the problems encountered, but rather denotes the first condition that prevented further execution of the utility.

See vxintro(1M) for a list of standard exit codes.

#### FILES

```
/etc/vx/type/usetype/vxmake Usage-type-specific utility
for filtering volume hierarchy
descriptions.
```

#### SEE ALSO

```
vxassist(1M), vxdisk(1M), vxdg(1M), vxintro(1M), vxmake(4)
```

## NAME

vxmend - mend simple problems in configuration records

## SYNOPSIS

```
vxmend [-fprsvV] [-g diskgroup] [-o useopt] [-U usetype]
clear field name ...
```

```
vxmend [-fprsvV] [-g diskgroup] [-o useopt] [-U usetype] fix
how name [arg...]
```

```
vxmend [-fprsvV] [-g diskgroup] [-o useopt] [-U usetype] off
name...
```

```
vxmend [-fprsvV] [-g diskgroup] [-o useopt] [-U usetype] on
name...
```

## DESCRIPTION

The vxmend utility performs various Volume Manager usage-type-specific operations on subdisk, plex, and volume records. The first operand is a keyword that determines the specific operation to perform. The remaining operands specify the configuration objects to which the operation is applied.

Each invocation can be applied to only one disk group at a time, due to internal implementation constraints. Any name operands will be used as record names to determine a default disk group, according to the standard disk group selection rules described in vxintro(1M). A specific disk group can be forced with -g diskgroup.

## KEYWORDS

**clear** Clears specified utility fields for each named record in the disk group. An option of -v, -p, or -s specifies that the utility operates only on volumes, plexes, or subdisks, respectively. If a record is a volume, or is associated directly or indirectly with a volume, then the clear operation is performed according to rules used by the usage type appropriate for that volume.

The field operand is a comma-separated list of keywords specifying fields to be cleared. Each keyword in the field operand is one of the following:

**all** Clears all clearable utility fields. Normally, this clears all of the persistent and non-persistent utility fields. A usage-type utility may choose a different set of fields, as appropriate.

## Maintenance Commands

vxmend(1M)

putil Clears all persistent utility fields.

putil0, putil1, or putil2  
Clears a specific persistent utility field.

tutil Clears all non-persistent (temporary) utility fields.

tutil0, tutil1, or tutil2  
Clears a specific non-persistent utility field.

Note: This particular functionality in conjunction with -f and -r flag can be used to recursively clear tutil0 fields in objects undergoing relayout, provided the top most object undergoing relayout is used to initiate it.

Usage types may implement additional field keywords. A usage type may also limit the set of clear operations that can be performed.

fix Changes the state of a volume or plex, named by the name operand, in a manner specified by the how string. The meaning of this operation is entirely usage-type specific. This operation applies only to volumes, or to plexes associated with a volume. Usage type rules appropriate for the volume are used to interpret the command. Additional arguments, after name, are interpreted according to rules defined by the usage type.

off Puts the named plexes or volumes into an offline state. This operation can be applied only to volumes, or to plexes associated with a volume. Usage type rules appropriate for the volume are used to perform the operation.

on Takes the named plexes or volumes out of the offline state. This operation can be applied only to volumes, or to plexes associated with a volume. Usage type rules appropriate for the volume are used to perform the operation.

## OPTIONS

- f Forces an operation that the Volume Manager considers potentially dangerous or unnecessary. This enables a limited set of additional operations that would normally not be allowed. Some operations may be disallowed even with this flag.

## Maintenance Commands

vxmend(1M)

- g diskgroup  
Specifies the disk group for the operation, either by disk group ID or by disk group name. By default, the disk group is chosen based on the name operands.
- o useopt Passes in usage-type-specific options to the operation.
- p Requires that name operands name plex records.
- r Operates recursively on records associated with the named volume or plex record. Operations applied to a volume will apply to the associated plexes and subdisks. Likewise, operations applied to a plex may be applied to the associated subdisks.
- s Requires that name operands name subdisk records.
- U usetype  
Limits the operation to apply to the specified usage type. Attempts to affect volumes with a different usage type will fail.
- v Requires that name operands name volume records.
- V Displays a list of utilities that would be called from vxmend, along with the arguments that would be passed. The -V performs a ``mock run'' so the utilities are not actually called.

## FSGEN and GEN Usage Types

The fsgen and gen usage types provide identical semantics for all operations of the vxmend utility. These usage types provide the following options as arguments to -o:

- force Forces an operation that internal consistency checks consider to be questionable. This applies

to attempts to use vxmend fix empty to uninitialize a volume that has plexes in the ACTIVE state, and also to attempts to disable the last plex, or the last complete (non-sparse) plex, in a volume. This flag is the same as -f.

plex=plexname

Requires that any named subdisk record be associated with a plex named plexname. Several plex options can be specified to indicate a list of allowed plex names.

vol=volume

VxVM 3.0

Last change: 26 Mar 1999

3

Maintenance Commands

vxmend(1M)

Requires that any named plex or subdisk record be associated, directly or indirectly, with a volume named volume. Several vol options can be specified to indicate a list of allowed volume names.

Limitations and extensions for the fsgen and gen usage types consist of the following:

clear      Keywords supported in the field operands include all standard keywords. In addition, a keyword of the form attnumber can be used to remove locks on a volume for a particular number of concurrent plex attach operations. For example, if the command:

```
vxplex att v1 v1-01
```

is aborted by pressing the INTERRUPT key (or equivalent) several times (which prevents a clean abort of the operation), you may have to clear the operation with:

```
vxmend clear tutil0 v1-01
vxmend clear att1 v1
vxplex dis v1-01
```

Interruption of the vxplex att command will display a list of commands to run, which will include the previous list.



The number of plexes currently being attached, either directly or as part of a compound operation, is stored in the tutil0 field of the volume record as a string in the form ATTnumber. If the number drops to zero, the tutil0 field is cleared. Some operations on a volume require that the tutil0 field be cleared.

Similar to the above example, if an online relay-out operation, either started using vxassist or the vxrelayout command is aborted by pressing the the INTERRUPT key (or equivalent) several times (which prevents a clean abort of the operation), you may have to clear the operation with:

```
topmost_plex          vxmend -rf clear tutil0 topmost_volume |
```

Interruption of an online relayout operation will

#### Maintenance Commands

vxmend(1M)

display a command similar to above with the appropriate object name to clear the tutil0 field. Online relayout stores a string RELAYOUT in the the tutil0 of all objects subject to relayout. If an interruption caused a cleanup to happen correctly, the tutil0 of all objects used by online relayout will be cleared, else they may or may not be cleared. The vxmend utility of the form mentioned above can be used to clear the tutil0 field. It is important that the object name used in the operation is the topmost object name or any object above the topmost object undergoing relayout.

fix The fsgen and gen usage types support the following vxmend fix operations:

```
vxmend fix active plex
```

Sets the state for the named plex to ACTIVE. The state for the volume is set to SYNC. The associated volume must be disabled, and the named plex must be in

the STALE state.

When starting a volume in the SYNC state, all ACTIVE plexes are enabled and are synchronized to have the same contents using a special read/write-back recovery mode. Any STALE plexes are then recovered by copying data from the ACTIVE plexes.

`vxmend fix clean plex`

Sets the state for the named plex to CLEAN. The associated volume must be disabled, the named plex must be in the STALE state, and the volume must have no additional plexes in the CLEAN state.

A volume is not startable if one plex is in the CLEAN state and some plexes are in the ACTIVE state. Thus, several `vxmend fix` operations are normally used in conjunction to set all plexes in a volume to STALE and then to set one plex to CLEAN. A volume start operation will then enable the CLEAN plex and recover the STALE plexes by copying data from the one CLEAN plex.

`vxmend fix empty volume`

Sets the named volume and all of its

associated plexes to the EMPTY state. The volume can then be re-initialized using `vxvol start` or any of the `vxvol init` operations. This operation requires that the volume be disabled.

`vxmend fix stale plex`

Sets the state for the named plex to STALE. The associated volume must be disabled, and the named plex must be ACTIVE or CLEAN. This operation names plexes that will be recovered by copying data from other plexes by a `vxvol start` operation.

off           The fsgen and gen usage types allow volumes and plexes to be specified as operands to vxmend off. A volume can be named only if -r is used to specify recursion, and is applied to all plexes in the volume. Taking a plex offline disables the plex and sets its state to OFFLINE.

              Taking the last enabled read-write plex in a volume offline, or the last complete plex (if there are additional sparse plexes), normally fails unless -f is specified. Applying this operation to a volume also disables the volume, and does not require use of -f.

              If a volume is disabled, then the checks for the last plex or the last complete plex are performed on the set of ACTIVE and CLEAN plexes, rather than on enabled read-write plexes.

              An offline state for a plex can be cleared with vxmend on. Alternatively, an offline plex can be reattached with vxplex att.

on            The fsgen and gen usage types allow volumes and plexes to be specified as operands to vxmend on. The vxmend on operation applied to a plex will change the OFFLINE state for a plex to STALE, allowing the plex to be recovered by the next vxvol start or vxvol startall. Applying vxmend on to a volume will change the state for all associated OFFLINE plexes to STALE.

#### RAID-5 Usage Type

The raid5 usage type provides the following options as arguments to -o:

force        Forces an operation that internal consistency

checks consider to be questionable.

Operations not documented in this section are not supported for the raid5 usage type. In particular, off and on are not supported, but other fsgen and gen operations are supported. In addition, the raid5 usage type also supports the follow-

ing extension:

fix            In addition to those supported for the fsgen and gen usage types, the raid5 usage type supports the following vxmend fix operation:

    vxmend fix unstale subdisk

              Clears any flags indicating that a subdisk of a RAID-5 plex is invalid. This prevents the data on the subdisk from being recovered when the volume is started.

#### EXIT CODES

The vxmend utility exits with a nonzero status if the attempted operation fails. A nonzero exit code is not a complete indicator of the problems encountered but rather denotes the first condition that prevented further execution of the utility. See vxintro(1M) for a list of standard exit codes.

#### FILES

    /etc/vx/type/usetype/vxmend    The utility that performs vxmend operations for a particular volume usage type.

#### SEE ALSO

    vxintro(1M), vxplex(1M), vxvol(1M)

## NAME

vxmirror - mirror volumes on a disk or control default mirroring

## SYNOPSIS

```
/etc/vx/bin/vxmirror [-g diskgroup ] [-d yes|no ] [-t task-tag ] medianame [new_medianame...]
```

```
/etc/vx/bin/vxmirror [-g diskgroup ] [-d yes|no ] [-t task-tag ] -a [new_medianame...]
```

```
/etc/vx/bin/vxmirror [-g diskgroup ] [-d yes|no]
```

```
/etc/vx/bin/vxmirror [-g diskgroup] -D
```

## DESCRIPTION

The vxmirror command provides a mechanism to mirror the contents of a specified disk, to mirror all currently unmirrored volumes in the specified disk group, or to change or display the current defaults for mirroring. All volumes that have only a single plex (mirror copy), are mirrored by adding an additional plex.

Volumes containing subdisks that reside on more than one disk are not mirrored by vxmirror.

vxmirror is generally called from the vxdiskadm menus. It is not an interactive command, and after it is called, continues until completion of the operation or until a failure is detected.

Note: Generating mirror copies of volumes can take a considerable time to complete.

Mirroring volumes from the boot disk produces a disk that can be used as an alternate boot disk. This is done by calling the vxrootmir command. This mirrors root and swap volumes. If usr and var are also volumes, they too are mirrored even though they may not be on the same disk.

In the first listed form of this command, the disk media name is supplied on the command line to vxmirror. That name is assumed to be the only disk from which volumes are mirrored. In the case of mirroring volumes from a specified disk, only simple single-subdisk volumes are mirrored.

In the first and second listed forms of the command, new\_medianame... identifies a new disk media name (or set of names). The mirroring operation uses these names as targets on which to allocate the mirrors. An error results if the same disk is specified for both the source and target disk and if no other viable targets are supplied.

## Maintenance Commands

vxmirror(1M)

## OPTIONS

- a Mirrors all existing volumes for the specified disk group.
- d yes | no  
Changes the default for subsequent volume creation, depending on the option argument. If yes, then all subsequent volumes created automatically become mirrored volumes. If no, then mirroring is turned off for future volumes created.
- D Displays current default status for mirroring.
- g diskgroup  
Limits operation of the command to the given disk group, as specified by disk group ID or disk group name. The medianame operands are evaluated relative to the given disk group. If no disk group is supplied, then rootdg is assumed.
- t tasktag  
Specifies using a tasktag as the tag for any tasks created to perform the mirror operations.

## EXAMPLES

The following command mirrors the disk disk01 to any available space on any available disk. Subsequent calls to vxassist mirror created volumes by default.

```
vxmirror -d yes disk01
```

This command displays the current status of default mirroring. It outputs the string yes if mirroring is currently enabled, or no if mirroring is not enabled.

```
vxmirror -D
```

This command mirrors any volumes on disk02 to disk03.

```
vxmirror disk02 disk03
```

## FILES

/etc/default/vxassist

The defaults file for vxassist parameters.

SEE ALSO

vxassist(1M), vxdiskadm(1M), vxintro(1M), vxrootmir(1M), vxtask(1M)

VxVM 3.0

Last change: 26 Mar 1999

2

===[ man1m/vxnotify.1m ]=====

Maintenance Commands

vxnotify(1M)

NAME

vxnotify - display Volume Manager configuration events

SYNOPSIS

vxnotify [-cdDfims] [-g diskgroup] [-n number ] [-t timeout] [-w wait-time]

DESCRIPTION

The vxnotify utility displays events related to disk and configuration changes, as managed by the Volume Manager configuration daemon, vxconfigd. If vxnotify is running on a system where the Volume Manager cluster feature is active, it displays events related to changes in the cluster state of the system on which it is running. vxnotify displays requested event types until killed by a signal, until a given number of events have been received, or until a given number of seconds have passed.

CONFIGURATION EVENTS

Each event is displayed as a single-line output record on the standard output.

connected A connection was established with vxconfigd. This event type is displayed immediately after successful startup and initialization of vxnotify. A connected event is also displayed if the connection to vxconfigd is lost, and then regained. A connected event displayed after a reconnection indicates that some events may have been lost.

disconnected

The connection to vxconfigd was lost. This normally results from vxconfigd being stopped (such as by vxdctl stop) or killed by a signal. In

response to a disconnection, vxnotify displays a disconnected event and then waits until a reconnection succeeds. A connected event is then displayed.

A disconnected event is also printed if vxconfigd is not accessible at the time vxnotify is started. In this case, the disconnected event precedes the first connected event.

#### more events

Due to internal buffer overruns, or other possible problems, some events may have been lost.

#### vxconfigd disabled

vxconfigd was changed to disabled mode. Most configuration information will be unavailable until vxconfigd is changed back to enabled mode.

#### vxconfigd enabled

vxconfigd was changed to enabled mode. All configuration information should now be retrievable. vxconfigd disabled and vxconfigd enabled events can be retrieved only when using diagnostic-mode connections to the vxconfigd diagnostic portal. Use -D to obtain a regular diagnostic mode connection.

#### waiting ...

If the -w option is specified, a waiting event is displayed after a defined period with no other events. Shell scripts can use waiting messages to collect groups of related, or at least nearly simultaneous, events. This can make shell scripts more efficient. This can also provide some scripts with better input since sets of detach events, in particular, often occur in groups that scripts can relate together. This is particularly important given that a typical shell script will block until vxnotify produces output, thus requiring output to indicate the end of a possible sequence of related events.

import dg groupname dgid groupid



The disk group named groupname was imported. The disk group ID of the imported disk group is groupid.

deport dg groupname dgid groupid

The named disk group was deported.

disable dg groupname dgid groupid

The named disk group was disabled. A disabled disk group cannot be changed, and its records cannot be printed with vxprint. However, some volumes in a disabled disk group may still be usable, although it is unlikely that the volumes will be usable after a system reboot. A disk group will be disabled as a result of excessive failures. A disk group will be disabled if the last disk in the disk group fails, or if errors occur when writing to all configuration and log copies in the disk group.

change dg groupname dgid groupid

A change was made to the configuration for the named disk group. The transaction ID for the

Maintenance Commands

vxnotify(1M)

update was groupid.

groupid

detach subdisk subdisk plex plex volume volume dg groupname  
dgid

The named subdisk, in the named disk group, was detached as a result of an I/O failure detected during normal volume I/O, or disabled as a result of a detected disk failure. Failures of a subdisk in a RAID-5 volume or a log subdisk within a mirrored volume result in a subdisk detach; other subdisk failures generally result in the subdisk's plex being detached.

detach plex plex volume volume dg groupname dgid groupid

The named plex, in the named disk group, was

detached as a result of an I/O failure detected during normal volume I/O, or disabled as a result of a detected total disk failure.

detach volume volume dg groupname dgid groupid

The named volume, in the named disk group, was detached as a result of an I/O failure detected during normal volume I/O, or as a result of a detected total disk failure. Usually, only plexes or subdisks are detached as a result of volume I/O failure. However, if a volume would become entirely unusable by detaching a plex or subdisk, then the volume may be detached.

detach disk accessname dm medianame dg groupname dgid groupid

The named disk, with device access name accessname and disk media name medianame was disconnected from the named disk group as a result of an apparent total disk failure. Total disk failures are checked for automatically when plexes or subdisks are detached by kernel failures, or explicitly by the vxdisk check operation (see vxdisk(1M)).

log-detach volume volume dg groupname dgid groupid

All log copies for the volume (either log plexes for a RAID-5 volume or log subdisks for a regular mirrored volume) have become unusable, either as a result of I/O failures or as a result of a detected total disk failure.

change disk accessname dm medianame dg groupname dgid groupid

The disk header changed for the disk with a device access name of accessname. The disk group name and ID of the disk are groupname and groupid, respectively. The displayed groupname and groupid strings will be ``-' or blank if the disk is not currently in an imported disk group.

degraded volume volume dg groupname dgid groupid

The RAID-5 volume has become degraded due to the loss of one subdisk in the raid5 plex of the volume. Accesses to some parts of the volume may be slower than to other parts depending on the location of the failed subdisk and the subsequent I/O patterns.

enabled path pathname to dmpnode dmpnodename

The path named pathname has been enabled. It is a path to the dmpnode named dmpnodename.

disabled path pathname to dmpnode dmpnodename

The path named pathname has been disabled. It is a path to the dmpnode named dmpnodename.

enabled dmpnode dmpnodename

The dmpnode named dmpnodename has been enabled. At least one of the paths under this node is now available for I/O.

disabled dmpnode dmpnodename

The dmpnode named dmpnodename has been disabled. It is no longer going to respond to I/O. All paths under this dmpnode are disabled.

enabled controller controllername disk array da\_serial\_no

The controller named controllername belonging to disk array with disk array serial number da\_serial\_no has been enabled. This means that at least one path through this controller is available for I/O.

disabled controller controllername disk array da\_serial\_no

The controller named controllername belonging to disk array with disk array serial number da\_serial\_no has been disabled. This means that I/Os will not be allowed through any of the paths that go through this controller.

removed disk array da\_serial\_no

The disk array with serial number da\_serial\_no has been removed from the system.

added disk array da\_serial\_no

The disk array with serial number da\_serial\_no has been added to the system.

joined cluster clustername as slave node nodeid

This system has joined the cluster named clustername as a slave node. Its node ID is nodeid. Available only if the Volume Manager cluster feature is enabled.

joined cluster clustername as master node nodeid

This system has joined the cluster named clustername as a master node. Its node ID is nodeid. If the system was already in the cluster as a slave, it has now become the master node. Available only if the Volume Manager cluster feature is enabled.

left cluster

This system has left the cluster of which it was previously a member. Available only if the Volume Manager cluster feature is enabled.

#### OPTIONS

- c Display disk group change events.
- d Display disk change events.
- D Use a diagnostic-mode connection to vxconfigd. This allows the receipt of events when vxconfigd is running in disabled mode. Access to configuration information is limited when vxconfigd is running in disabled mode. For most applications, it is better to let vxnotify print events only when vxconfigd is running in enabled mode.
- g diskgroup Restrict displayed events to those in the indicated disk group. The disk group can be specified either as a disk group name or a disk group ID.
- f Display plex, volume, and disk detach events.
- i Display disk group import, deport, and disable events.

- m            Display multipath database change related events.
- n number    Display the indicated number of vxconfigd events, then exit. Events that are not generated by vxconfigd (that is, connect, disconnect, and waiting events) do not count towards the number of counted events and will not cause an exit to occur.
- s            Display cluster change events. If the -i option is also specified, the imports and deports of shared disk groups will be displayed when a cluster change takes place. Available only if the Volume Manager cluster feature is enabled.
- t timeout    Display events for up to timeout seconds, then exit. The -n and -t options can be combined to specify a maximum number of events and a maximum timeout to wait before exiting.
- w wait\_time    Display waiting events after wait\_time seconds with no other events.

If none of the -c, -d, -f, -i, or -s options are specified, then default to printing all event types. If a disk group is specified with -g, display only disk group-related events.

#### EXAMPLES

The following example shell script sends mail to root for all detected plex, volume, and disk detaches:

```

checkdetach() {
    d=`vxprint -AQdF '%name %nodarec' | awk '$2=="on" {print
" " $1}`
    p=`vxprint -AQpe 'pl_kdetach || pl_nodarec' -F ' %name'`
    v=`vxprint -AQvF ' %name' -e \
        "((any aslist.pl_kdetach==true) ||
        (any aslist.pl_nodarec)) &&
        !(any aslist.pl_stale==false)"`

    if [ ! -z "$d" ] || [ ! -z "$p" ] || [ ! -z "$v" ]
    then
        (
            cat <<EOF
Failures have been detected by the VERITAS Volume Manager:
EOF
            [ -z "$d" ] || echo "\\nfailed disks:\\n$d"
            [ -z "$p" ] || echo "\\nfailed plexes:\\n$p"
        )
    fi
}

```

```
    [ -z "$v" ] || echo "\\nfailed volumes:\\n$v"  
  ) | mailx -s "Volume Manager failures" root
```

VxVM 3.0

Last change: 26 Mar 1999

6

Maintenance Commands

vxnotify(1M)

```
    fi  
  }  
  
  vxnotify -f -w 30 | while read code more  
do  
    case $code in  
      waiting) checkdetach;;  
    esac  
done
```

#### EXIT CODES

The vxnotify utility exits with a nonzero status if an error is encountered while communicating with vxconfigd. See vxintro(1M) for a list of standard exit codes.

#### SEE ALSO

vxconfigd(1M), vxdisk(1M), vxintro(1M), vxtrace(1M),  
vx\_notify(3X)

===[ man1m/vxplex.1m ]=====

## Maintenance Commands

vxplex(1M)

## NAME

vxplex - perform Volume Manager operations on plexes

## SYNOPSIS

vxplex [-fV] [-g diskgroup] [-o useopt ] [-t tasktag] [-U  
usetype ] att volume plex...

vxplex [-fV] [-g diskgroup] [-o useopt ] [-t tasktag] [-U  
usetype ] cp volume plex...

vxplex [-fV] [-g diskgroup] [-o useopt ] [-t tasktag] [-U  
usetype ] [-v volume ] det plex...

vxplex [-fV] [-g diskgroup] [-o useopt ] [-t tasktag] [-U  
usetype ] [-v volume ] dis plex...

vxplex [-fV] [-g diskgroup] [-o useopt ] [-t tasktag] [-U  
usetype ] snapabort plex...

vxplex [-fV] [-g diskgroup] [-o useopt ] [-t tasktag] [-U  
usetype ] snapshot plex newvolume...

vxplex [-fV] [-g diskgroup] [-o useopt ] [-t tasktag] [-U  
usetype ] snapstart volumeplex...

vxplex [-fV] [-g diskgroup] [-o useopt ] [-t tasktag] [-U  
usetype ] [-v volume ] mv oldplex newplex

## DESCRIPTION

The vxplex utility performs Volume Manager operations on plexes and on volume-and-plex combinations. The first

operand is a keyword that determines the specific operation to perform. The remaining operands specify the configuration objects to which the operation is to be applied.

Each operation can be applied to only one disk group at a time, due to internal implementation constraints. Any volume or plex operands will be used to determine a default disk group, according to the standard disk group selection rules described in vxintro(1M). A specific disk group can be forced with -g diskgroup.

#### KEYWORDS

att Attaches each named plex to the named volume. This can be applied to dissociated plexes, or to non-enabled plexes already associated with the named volume. If the volume is enabled, then the result of the successful operation will be to associate the plex (if needed) and to recover the plex to have the same contents as all other attached plexes in the volume. The rules for performing the attach depend upon the usage type of

VxVM 3.0

Last change: 26 Mar 1999

1

Maintenance Commands

vxplex(1M)

the named volume.

Attaching a plex is the normal means of recovering a plex after a disk replacement, or after a plex offline.

cp Copies the named volume to the named plexes. The volume cannot be enabled, and the named plexes must not be associated. The results of the operation will be a set of dissociated plexes that are an exact copy of the volume at the time of completion of the operation. The rules for performing the attach depend upon the usage type of the named volume. To improve the quality of the copies, some usage types attempt to make the detached plex consistent with respect to in-memory data.

This operation can be used to make a copy of a volume, for backup purposes, without mirroring the volume in advance.

det Detaches each of the named plexes. Detaching a plex leaves the plex associated with its volume,



but prevents normal volume I/O from being directed to the plex. This operation can be applied to plexes that are enabled or disabled. The rules for performing the detach depend upon the usage types of the volumes involved. The operation does not apply to dissociated plexes.

`dis` Dissociates each of the named plexes. Dissociating a plex breaks the link between the plex and its volume. A dissociated plex is inaccessible until it is reassociated, which can be done either with `vxplex att` or with `vxmake`. Any checks and synchronizations that apply to the `det` operation also apply to the `dis` operation.

Plex dissociation is the normal means of unmirroring a volume, or reducing the mirror count for a volume. To support this use, `-o rm` can be used to dissociate and remove the plex (and its associated subdisks) in the same operation. This makes the space used by those subdisks usable for new allocations (such as with `vxassist` or with `vxmake`).

Plex dissociation can also be used for file system backups of volumes that are normally mirrored. Plex devices are not directly mountable, so the backup method described for the `det` operation will not work if the backup program requires a mounted file system. To support such backup programs, a

Maintenance Commands

`vxplex(1M)`

plex can be dissociated and can then be allocated to a new volume as in the following example:

```
vxmake -U gen vol volume01 plex=plex01
```

The created volume can then be started and mounted for use by the backup program.

`mv` Attach the plex `newplex` to the volume that `oldplex` is associated with and dissociate `oldplex`. The volume cannot be disabled, and `newplex` must name a dissociated plex. The operation ensures seamless replacement of the dissociated plex without loss of data in the volume and without significant

delays in volume accessibility.

A primary purpose for the plex move operation is to move a plex that is using a disk to another location. In support of this purpose for the operation, `-o rm` can be specified to remove the original plex after completion of the operation.

For concatenated or striped plexes, the `vxsd mv` operation can be used to move individual subdisks off a disk. The rules for performing the move depend upon the usage types of the volume to which `oldplex` is associated.

`snapabort` This operation can be used in order to cancel the effects of a `snapstart`.

`snapstart` and `snapshot`

These two operations form the two parts of a preferred means of copying a volume to a plex for backup purposes. The `snapstart` operation attaches a plex to a volume and, when the operation is complete, leaves the plex associated as a temporary plex. After the operation completes, the administrator can convert the plex attached by `snapstart` into a new volume using `vxplex snapshot`. To improve the quality of the copies, some usage types attempt to make the detached plex consistent with respect to in-memory data.

This method of backup is preferable to using `vxplex cp` because it allows the administrator to coordinate breaking off the plex from the original volume at a well-defined point in time. This is important, since attaching a plex to a volume can take a considerable amount of time, and it is difficult to know when it will complete. Also,

directly converting the plex into a new volume is more convenient than requiring additional steps.

#### OPTIONS

`-f` Forces an operation that the Volume Manager considers potentially dangerous or of questionable use. This permits a limited set of operations

that would otherwise be disallowed. Some operations may be disallowed even with this flag.

-g diskgroup

Specifies the disk group for the operation, either by disk group ID or by disk group name. By default, the disk group is chosen based on the name operands.

-o useopt Passes in usage-type-specific options to the operation. A certain set of operations are expected to be implemented by all usage types:

iosize=size

Performs copy operations in regions with the length specified by size, which is a standard Volume Manager length number (see vxintro(1M)). Specifying a larger number typically causes the operation to complete sooner, but with greater impact on other processes using the volume. The default I/O size is typically 32 kilobytes.

rm Removes the plexes after successful completion of a vxplex dis operation. Remove the source plex after successful completion of vxplex mv.

slow[=iodelay]

Reduces the system performance impact of copy operations. Copies operations are usually a set of short copy operations on small regions of the volume (normally from 16 kilobytes to 128 kilobytes). This option inserts a delay between the recovery of each such region. A specific delay can be specified with iodelay as a number of milliseconds; otherwise, a default is chosen (normally 250 milliseconds).

-t tasktag

If any tasks are registered to track the progress of the operation, mark them with the tag tasktag.

-U usetype

Limits the operation to apply to this usage type. Attempts to affect volumes with a different usage type will fail.

`-v volume` Requires that the plex named by a plex or oldplex operand be associated with the named volume. This option can be used as a sanity check, to ensure that the specified plex is actually the plex desired for the operation.

`-V` Displays a list of utilities that would be called from vxplex, along with the arguments that would be passed. The `-V` option performs a ``mock run'' so the utilities are not actually called.

#### FSGEN and GEN Usage Types

The fsgen and gen usage types provide similar, though not identical, semantics for all operations of the vxplex utility. In particular, the fsgen usage type will attempt to flush in-memory data cached for the file system residing on the volume. For most file systems, this consists of calling `sync(1M)` to attempt to flush all in-memory data to disk. For the vxfs file system type, this will use special ioctls to ensure a reliable flush of the involved volume.

If a vxplex operation is interrupted by a signal, then an attempt is made to restore the disk group configuration to a state that is roughly equivalent to its original state. If this attempt is interrupted (such as through another signal) then the user may need to perform some cleanup. The specific cleanup actions that are needed are written to the standard error before vxplex exits.

The fsgen and gen usage types provide the following options as arguments to `-o` in addition to the required options:

`force` Forces an operation that the Volume Manager considers potentially dangerous or of questionable use. This applies to attempts to detach or dissociate the last (complete) plex in a volume, or to attempts to move a plex to a plex that has a different size. This flag is the same as `-f`.

`mapzero` If a plex is moved to a new plex that has regions that are mapped to a subdisk in the destination, but are not mapped to a subdisk for any enabled, readable plex in the volume, then zero out that mapped region in the destination plex. Without this flag, the mapped region may be left unchanged from its original contents.

`rerr` Ignore volume or plex read errors when copying

data onto a plex. A warning message is written to standard error if a read error occurs, but the error does not affect success of the operation. This operation can be used only with the cp operation; the operation is ignored if used with other operations.

werr Ignore plex write errors when copying data onto a plex. A warning message is written to standard error if a write error occurs, but the error does not affect success of the operation. This operation can be used only with the cp operation; the operation is ignored if used with other operations.

Limitations and extensions for the fsgen and gen usage types consist of the following:

att If the volume is enabled and one of the named plexes is associated with the volume, then the plex must be STALE, EMPTY, ACTIVE, or OFFLINE. If the operation succeeds in attaching a plex, then any I/O fail condition for the plex is cleared. Also, attaching to an enabled volume requires that the volume have at least one enabled, read-write plex.

If the volume is not enabled, then the named plexes are associated with the volume (if not already associated) and are set to the STALE state, so that the plex will be fully attached by the next vxvol start or vxvol startall operation that applies to the volume.

If the log type of the volume is UNDEF and an unassociated plex with a log subdisk is attached, the volume is automatically converted to have a log type of DRL. Logging of volume changes is enabled when the volume has at least one enabled, associated plex with an enabled log subdisk and at least two read-write mode plexes.

An attempt to attach an unassociated plex fails if the putil0 field is not empty. This makes it possible to prevent use of a plex by using vxedit set to set the putil0 field to a non-empty string. The putil0 field can then be cleared with either vxedit set or with vxmend clear putil0.

cp           The fsgen and gen usage types do not add any specific restrictions to the cp operation.

## Maintenance Commands

vxplex(1M)

### dis and det

A detach or dissociate of a plex in an enabled volume fails if applied to a plex that is the last complete, enabled, read-write plex in the volume and the volume contains two or more non-complete, enabled, read-write plexes. In other words, a volume cannot be left with two enabled, non-complete plexes. A complete plex is one that is at least as long as the volume, and has subdisks mapped to the plex for all blocks up to the length of the volume. The -f option is required to reduce a volume to containing one enabled, read-write, non-complete plex, or to having no enabled, read-write plexes at all.

The det operation changes the state for an ACTIVE or CLEAN plex to STALE. The next time the volume is started, the plex will be re-attached automatically.

mv           If the destination plex has unmapped regions (a range of blocks in the plex with no backing sub-disk) that are not mapped in the source plex, or if the destination plex is shorter than the source plex, then the -f option is required. Even with -f, the operation will prevent the plex from being sparsed such that the volume would be left with two or more sparse, enabled, read-write plexes, but no complete plexes.

### RAID5 Usage-Type

The raid5 usage type provides the following options as arguments to -o in addition to the required options:

force       Forces an operation that the Volume Manager considers potentially dangerous or of questionable use. This applies to attempts to dissociate the RAID-5 plex of a non-EMPTY volume or to remove the last RAID-5 log plex of a non-EMPTY volume.

As with other usage types, if a vxplex operation is inter-

rupted by a signal, then an attempt is made to restore the disk group configuration to a state that is roughly equivalent to its original state. If this attempt is interrupted (such as through another signal) then the user may need to perform some cleanup. The specific cleanup actions that are needed are written to the standard error before vxplex exits.

The raid5 usage type supports only the following keywords:

att           Attaches the named plexes to the named volume. If

## Maintenance Commands

vxplex(1M)

a plex has a layout of RAID, the plex will be attached as the RAID-5 plex of the RAID-5 volume. To attach a RAID-5 plex to the volume, the volume must be disabled and be in the EMPTY state, and the RAID-5 plex will be given a state of EMPTY.

If a plex has a layout other than RAID, the plex will be attached as a RAID-5 log plex for the RAID-5 volume. If the volume has no RAID-5 log plexes, the log length for the volume will be set to the length of the smallest log plex being attached. If the volume already has at least one log plex, a plex can only be attached as a log plex if its contiguous length is at minimum the volume's log length. RAID-5 log plexes cannot be sparse in respect to the volume's log length; attempts to attach a sparse log plex will fail.

If the RAID-5 volume is not enabled, log plexes are attached and marked as STALE. If the RAID-5 volume is enabled and has no log plexes, attaching a log plex will cause plexes being attached as log plexes to be zeroed before they are enabled. Otherwise, the new log plexes are attached write-only and the contents of the existing log plexes are copied to the new log plexes using ATOMIC\_COPY ioctls, after which the logs are enabled.

dis           Dissociates the named plex from the RAID-5 volume to which it is attached. If the plex is the RAID-5 plex of the volume and the volume is not EMPTY, this requires the -o force option, as any data on the volume would be lost. If the plex is a log

plex for the volume and will leave the RAID-5 volume with no usable log plexes, the -o force option is required.

Note that the RAID-5 usage type does not support the det, cp, snapstart, snapshot, snapabort, or copy keywords; these operations are either inappropriate or impossible to perform within the operational concepts of RAID-5.

#### FILES

/etc/vx/type/usetype/vxplex The utility that performs vxplex operations for a particular volume usage type.

/etc/vx/type/fsgen/fs/fstype/vxsync Path to a program used with the fsgen usage type for synchronizing in-memory file system data with a volume, for

the file system type fstype. The program is given arguments of a volume name and one or more plex names. For the ufs and s5 file system types, this is a link to sync. For the vxfs file system type, this program uses the VxFS file system freeze feature to ensure a perfect synchronized detach.

#### EXIT CODES

The vxplex utility exits with a nonzero status if the attempted operation fails. A nonzero exit code is not a complete indicator of the problems encountered but rather denotes the first condition that prevented further execution of the utility. See vxintro(1M) for a list of standard exit codes.

#### SEE ALSO

sync(1M), vxassist(1M), vxedit(1M), vxintro(1M), vxmend(1M), vxtask(1M), vxvol(1M)



===[ man1m/vxprint.1m ]=====

Maintenance Commands

vxprint(1M)

#### NAME

vxprint - display records from the Volume Manager configuration

#### SYNOPSIS

vxprint [-aAdfGhHlLmnpqQrstv ] [-D database ] [-e pattern ]  
[-F [type:] format-spec ] [-g diskgroup ] [name...]

vxprint [-AqS] [-g diskgroup ]

#### DESCRIPTION

The vxprint utility displays complete or partial information from records in Volume Manager disk group configurations. Records can be selected by name or with special search

expressions. Additionally, record association hierarchies can be displayed in an orderly fashion so that the structure of records is more apparent.

Dashes (-) are displayed in the output wherever there is no applicable record value.

If no options are specified, the default output uses -f, -h, -r, and -A. Specifying other options can override these defaults, making it necessary to explicitly specify the normal default options.

The default output format consists of single-line records, each of which includes a record type, name, usage type or object association, enabled state, length, and other fields. A one-line header is written before the record information.

When no disk group is specified, selected records are retrieved from rootdg. Subdisks/Subvolumes are sorted primarily by their device/volume, subdisks may also include a device offset. Plex and volume records are sorted by name.

Note 1: The vxprint utility can display disk group, disk media, volume, plex, and subdisk/subvolume records. It cannot display disk access records. Use the vxdisk list operation to display disk access records, or physical disk information.

Note 2: The subvolume "record type" appears as sv as a convenience in displaying views of the configuration database. For manipulation purposes, these records are accessible as subdisk record types, i.e., records tagged as type sd .

#### OPTIONS

The -v, -p, -s, -d, and -G options may be combined to specify that more than one record type is allowed. Specifying all these options restores the default behavior of

retrieving all record types.

-a            Display all information about each selected record, one record per line. The contents are similar to the -m option, with the following exceptions: the -a option format appears on a sin-

gle line with one space character between each field, the list of associated records is not displayed, and the `-m` option retains the `sd` record type rather than `sv`. This format is useful for processing output through filters such as `sed` and `grep` that operate exclusively on one-line records, although the fields are not readily distinguishable. It isn't a practical format from the viewpoint of human readability.

`-A` Print records from all active (imported) disk groups. Each disk group represented in the output is separated from other disk groups by blank lines. A short header line introduces each disk group.

`-d` Display only disk media records.

`-D database`

Get a configuration from the specified location. The database option argument can be one of:

`vxconfigd` Get a configuration from the volume configuration daemon (default).

`-` Read a configuration from the standard input. The standard input is expected to be in standard `vxmake` input format.

`-e pattern`

Use a volume configuration search expression to select records to be displayed. See `vol_pattern(4)` for a description of search patterns.

`-f` Display information about each record as one-line output records containing the following fields, from left to right. A one-line header is written before any record information.

1. Record type

2. Record name

3. Usage-type, volume association, or plex association (or `-` for unassociated plexes and

subdisks)

4. Enabled state (or - for subdisks, disks, or disk groups)
5. Length, in units of system sectors
6. Plex association offset (or - for volumes, plexes, disks, or disk groups). This field will appear as LOG for log subdisks.
7. Usage-dependent state (or - for subdisks). If an exception condition is recognized (a plex I/O failure, removed or inaccessible disk, or an unrecovered stale data condition), then that condition is listed instead of any usage-type-dependent state.
8. The tutil[0] field. This field is set by usage-types as a lockout mechanism.
9. The putil[0] field. This field can be set to prevent associations of plex or subdisk records.

`-F [type :]format_spec`

Set a literal format string to use for displaying record information. If the option argument begins with a comma-separated list of zero or more record types (sd, plex, or vol) followed by a colon, then the format\_spec after the colon is used when printing the indicated record types. If no record types are specified, then all record types are assumed.

The order of `-F` options is significant, with specifications later in the option list overriding earlier specifications. Any use of `-F` overrides any other option letter specifying a type of format for the indicated record types. Thus, `-F vol:format_spec` can be used with the `-t` option to change the format used for volumes, while still using the `-t` format for plex and subdisk records.

The format-spec string consists of literal text with embedded configuration record variables. Configuration record variables are introduced with a percent sign (%). The percent sign is followed by a variable name or by a variable name and optional field width in braces. The following

formats are allowed for a variable specification:

```
%field_name  
%{field_name}  
%{field_name : [[-] width ][ * ] }  
%{field_spec | field_spec [ |... ] }
```

The first format specifies the exact field name. The second format allows a field to be specified with immediately surrounding text that would otherwise be taken as part of the field name. The third format allows the specification of a justification and a field width. The fourth format allows alternate specifications to be used, either with or without justification and width specifications. For the fourth, the first specification is used if the specified field name is applicable to the record and is non-empty; otherwise, the next available specification is used. Any number of alternate specifications can be used.

If no field width is specified, then the number of output column positions used for the field is the smallest possible to contain the value; otherwise spaces are added in the output to make it width columns in length. A field is not truncated if the minimum number of column positions necessary for a value is greater than width.

If a field width is specified with a leading dash (-) character, then an output field is lengthened by adding spaces after the field value, yielding a left-justified field. Otherwise, spaces are added before the value, yielding a right-justified field.

If a field width is followed or replaced by an asterisk (\*) character, then an unrecognized or inappropriate field yields either no output for the field or a field containing all blanks. Without the asterisk, the printed field contains the character -.

A percent sign (%) can be displayed by including two percent characters (%%) in format-spec.

See the Record Fields section for a description of the field names that can be specified. An invalid format string may yield unexpected output, but does not generate an error.

Maintenance Commands

vxprint(1M)

- g diskgroup  
Display records from the specified disk group. The diskgroup option argument can be either a disk group name or disk group ID.
- G  
Display only disk group records.
- h  
List hierarchies below selected records. For volumes, this list includes all associated plexes and subdisks. For plexes, this list includes all associated subdisks/subvolumes. Hierarchies are separated in the output by a blank line. Each object listed occupies its own line. The order of output is a volume name, followed by one associated plex, followed by all of the subdisks/subvolumes for that plex, followed by another associated plex, followed by all of the subdisks/subvolumes for the second plex, and so on.  
  
The -v, -p, and -s options limit the selection only of the head of a hierarchy. They do not prevent the display of associated records through the -h option.  
  
Unless objects are named explicitly with name operands, a record is never displayed in two separate hierarchies. Thus, a selected plex is not displayed as a separate hierarchy if the volume that is associated with the plex is also selected.
- H  
Print help information on usage.
- l  
Display all information from each selected record. This information is in a free format that is not intended for use by scripts. This format is more convenient than the -m format for looking at

records directly, because the density of information is more appropriate for human viewing.

- m Display all information about each selected record in a format that is useful as input to both the vxmake utility and to awk(1) scripts. The format used is the vxmake description format (see vxmake(4)). In addition to record information, the list of plex or subdisk/subvolume records associated with selected volume or plex records is displayed. Each field is output on a separate line, indented by a single tab. Values for fields

#### Maintenance Commands

vxprint(1M)

that contain comment-style strings are always preceded by one double-quote character.

- n Display only the names of selected records.
- p Display only plexes with associated subdisks. If a name operand names a volume or subdisk, then a diagnostic is written to the standard error.
- q Suppress headers that would otherwise be printed for the default and the -t and -f output formats.
- Q Suppress the disk group header that separates each disk group. A single blank line still separates each disk group.
- R Rendezvous point. This is a reserved option to indicate the communications point used with vxconfigd.
- r Display related records of a volume containing subvolumes. Grouping is done under the highest level volume.
- L Useful when used with the -r parameter. Display related records of a volume containing subvolumes, but grouping is done under any volume.
- s Display only subdisks/subvolumes. If a name operand names a volume or plex, then a diagnostic is written to the standard error.

-S Display configuration summary information. The output consists of a header line followed by a line containing the total number of subdisks/subvolumes, plexes and volumes; the number of unassociated subdisks/subvolumes; and the number of unassociated plexes.

-t Print single-line output records that depend upon the configuration record type. For disk groups, the output consists of the record type (dg) the disk group name, and the disk group ID.

For disk media records, the output consists of the following fields, in order from left to right:

1. Record type (dm)
2. Record name
3. Underlying disk access record

Maintenance Commands

vxprint(1M)

4. Disk access record type (sliced, simple, or nopriv)
5. Length of the disk's private region
6. Length of the disk's public region
7. Path to use for accessing the underlying raw disk device for the disk's public region.

For subdisks, the output consists of the following fields, from left to right.

1. Record type (sd)
2. Record name
3. Associated plex, or dash (-) if the subdisk is dissociated
4. Name of the disk media record used by the subdisk
5. Device offset in sectors



6. Subdisk length in sectors
7. Plex association offset, optionally preceded by subdisk column number for subdisks associated to striped plexes, LOG for log subdisks, or the putil[0] field if the subdisk is dissociated. The putil[0] field can be non-empty to reserve the subdisk's space for non-volume uses. If the putil[0] field is empty, - is displayed for dissociated subdisks.
8. A string representing the state of the subdisk (ENA if the subdisk is usable; DIS if the subdisk is disabled; RCOV if the subdisk is part of a RAID-5 plex and has stale content; DET if the subdisk has been detached; KDET if the subdisk has been detached in the kernel due to an error; RMOV if the media record on which the subdisk is defined has been removed from its disk access record by a utility; RLOC if a subdisk has failed and is waiting to be relocated, or NDEV if the media record on which the subdisk is defined has no access record associated).

For subvolumes, the output consists of the

following fields, from left to right.

1. Record type (sv)
2. Record name
3. Associated plex, or dash (-) if the subvolume is dissociated.
4. Name of the underlying (layered) volume record used by the subvolume.
5. Number of layers used in the subvolume.
6. Subvolume length in sectors

7. Plex association offset, optionally preceded by subvolume column number for subvolumes associated to striped plexes.
8. Number of active plexes, followed by the number of plexes in the underlying (layered) volume.
9. A string representing the state of the subvolume (ENA if the subvolume is usable; DIS if the subvolume is disabled; KDET if the subvolume has been detached in the kernel due to an error)

For plexes, the output consists of the following fields, from left to right:

1. Record type (pl)
2. Record name
3. Associated volume, or - if the plex is dissociated
4. Plex kernel state
5. Plex utility state. If an exception condition is recognized on the plex (an I/O failure, a removed or inaccessible disk, or an unrecovered stale data condition), then that condition is listed instead of the value of the plex record's state field.
6. Plex length in sectors
7. Plex layout type

8. Number of columns and plex stripe width, or - if the plex is not striped
9. Plex I/O mode, either RW (read-write), WO (write-only), or RO (read-only)

For volumes, the output consists of the following fields, from left to right:

1. Record type (v)
2. Record name
3. Associated usage type
4. Volume kernel state
5. Volume utility state
6. Volume length in sectors
7. Volume read policy
8. The preferred plex, if the read-policy uses a preferred plex

A header line is printed before any record information, for each type of record that could be selected based on the -v, -p, -s, and -h options. These header lines are followed by a single blank line.

-v Display only volumes with associated plexes and subdisks/subvolumes. This restricts the records matched with search patterns. Also, the default selection of all records in the database is then restricted to all volumes. If a name operand names a plex or subdisk/subvolume, then a diagnostic is written to the standard error.

#### Record Fields

The field names that can be used with the format\_spec string of the -F option and that are produced for the -m or -a options are the same as those that can be provided as input to the vxmake utility. For a list of these field names, see vxmake(4). Some additional pseudo fields are also supported. These are:

##### admin\_state

The persistent state for a plex or volume record, accounting for any exceptional conditions. For volume records, this displays the state field.

For plex records, this displays one of the following in the given precedence order: NODEVICE if an expected underlying disk could not be found; REMOVED if an underlying disk is in the removed state; IOFAIL if an unrecovered I/O failure caused the plex to be detached; RECOVER if a disk replacement left the plex in need of recovery, either from another plex or from a backup.

- aslist      A comma-separated list of subdisks or plexes that are associated with a plex or volume record.
- assoc      The name of the volume or plex to which a plex or subdisk record is associated. If the record is not associated, this field is empty.
- column\_pl\_offset      For a subdisk associated with a striped plex, the column number and column offset of the subdisk separated by a / or the plex offset (if the subdisk is associated in a non-striped plex) or - (if the plex is not associated).
- dgname or dg\_name      This is the name of the disk group containing the record.
- name      The name of the record being displayed. Because the record name is specified positionally within vxmake description formats, the vxmake utility and the -m and -a options to vxprint do not explicitly provide this field name.
- ncolumn\_st\_width      For a striped plex, the number of columns and stripe unit size for a plex, separated by a /, or - if the plex is not striped.
- plname or pl\_name      The name of an associated plex record. For a plex record, this is the plex's name; for a subdisk record, this is the associated plex's name (if any).
- rec\_type or rtype      This is either dg (disk group), dm (disk), vol (volume), plex (plex), or sd (subdisk), depending on the record being displayed.
- sd\_flag      For a subdisk associated with a RAID-5 plex, this will display flags relating to the status of the subdisk. An S indicates that the subdisk is

considered to contain stale data. A `d` indicates that the subdisk has been detached from the RAID-5 plex.

`sdaslist` A comma-separated list of subdisks associated with a plex. Each subdisk name is followed by a colon and the subdisk's plex association offset, in sectors. For volume records, this field is equivalent to `aslist`.

`short_type` or `type`

This is either `dg`, `dm`, `v`, `pl`, or `sd`, depending upon the record type. This pseudo variable can be used in a 2-character field, if a full 4 character field (required by `rec_type`) is too large.

`use_assoc` The usage type for volume records and the association name for associated plexes and subdisks. For dissociated plexes and subdisks, this is an empty string.

`vname` or `v_name`

The name of an associated volume record. For a volume record, this is the volume's name; for a plex record, this is the associated volume's name (if any); for a subdisk record, this is the associated volume of the associated plex (if any).

Displaying a boolean value always yields `on` or `off`. If a field containing a length or offset is specified in a `format_spec` string, then the result is the length or offset in sectors. When the field is displayed with `-m` or `-a`, the length or offset is displayed in sectors with a suffix of `s`.

#### EXIT CODES

The `vxprint` utility exits with a nonzero status if the attempted operation fails. A nonzero exit code is not a complete indicator of the problems encountered but rather denotes the first condition that prevented further execution of the utility.

See `vxintro(1M)` for a list of standard exit codes.

#### EXAMPLES

To display all records in all disk groups, with clearly displayed associations and with output lines tailored to each record type, enter:

```
vxprint -Ath
```

## Maintenance Commands

vxprint(1M)

To avoid looking at the 5-line header and the extra disk group headers generated by this command, you can remove all the headers by adding a `-q`.

To display all subdisks and all disk groups, in sorted order by disk, enter:

```
vxprint -AGts
```

If all plexes are named based on volumes, this can be a convenient means of viewing large configurations. The association field for each of the subdisks names the plex, and the plex name will normally imply a volume association by the form of the plex name.

To display the names of all unassociated plexes, use the command

```
vxprint -n -p -e !assoc
```

To print all subdisks, including the subdisk name and either the subdisk plex association offset or the `putil0` field for dissociated subdisks, enter:

```
vxprint -s -F "%{name:-14} %{pl_offset|putil0}"
```

The `vxprint` command can be used to back up objects from one disk group configuration with the use of the combined options `-mvpsh` or `-mvpshr`. The output should be saved into a file that is maintained outside of the system being backed up. The command:

```
vxmake -d file
```

can be used with the saved file to restore the configuration.

SEE ALSO

awk(1), grep(1), sed(1), vxinfo(1M), vxintro(1M),  
vxmake(1M), vxmake(4)

VxVM 3.0

Last change: 26 Mar 1999

12

===[ man1m/vxr5check.1m ]=====

Maintenance Commands

vxr5check(1M)

#### NAME

vxr5check - verify RAID-5 volume parity

#### SYNOPSIS

/etc/vx/bin/vxr5check [-iv ] [-g diskgroup] volume

#### DESCRIPTION

The vxr5check utility compares the parity of each stripe of a RAID-5 volume specified by the volume operand. vxr5check reads the data for each stripe, generates the parity for this stripe, and compares this parity with the existing parity.

vxr5check can be executed against the entire RAID-5 volume, or incrementally on RAID-5 stripe boundaries, by using the -i option.

#### OPTIONS

- g diskgroup  
Specifies the VM disk group name for the RAID-5 volume name for verification. If a disk group name is not specified, the rootdg disk group is assumed as the default.
- i  
Verify the RAID-5 volume incrementally per stripes. If any parity mismatch is found, that stripe location is displayed.
- v  
Verbose output for the incremental vxr5check verification. The verbose option outputs each stripe number that is being verified.

## OUTPUT FORMAT

In verbose mode and incremental mode, summary reports for each stripe of the RAID-5 volume are printed in output records. If an error is returned for a stripe, then an error message and stripe number are displayed. In non-verbose mode, if an error is returned, then an error message is displayed.

If an parity mismatch error is determined on a stripe, then `vxr5check` exits on that stripe, and does not continue for the remaining stripes in the RAID-5 volume.

## FILES

`/usr/lib/vxvm/bin/vxr5vrfy` The utility that `vxr5check` calls to perform RAID-5 parity verification operations for the specified RAID-5 volume.

## EXIT CODES

The `vxr5check` utility exits with a nonzero status if the

VxVM 3.0

Last change: 26 Mar 1999

1

Maintenance Commands

`vxr5check(1M)`

attempted operation fails. A nonzero exit code is not a complete indicator of the problems encountered, but rather denotes the first condition that prevented further execution of the utility. See `vxintro(1M)` for a list of standard exit codes.

## SEE ALSO

`vxevac(1M)`, `vxintro(1M)`, `vxmend(1M)`, `vxvol(1M)`



===[ man1m/vxreattach.1m ]=====

Maintenance Commands

vxreattach(1M)

NAME

vxreattach - reattach disk drives that have once again become accessible

SYNOPSIS

/etc/vx/bin/vxreattach [-br ] [accessname...]

/etc/vx/bin/vxreattach -c accessname

DESCRIPTION

The vxreattach utility reattaches disks to the disk group they were in and retains the same media name.

This operation may be necessary if a disk has a transient failure, or if the Volume Manager starts with some disk drivers unloaded and unloadable. Disks then enter the failed state. If the problem is fixed, vxreattach may be able to reattach the disks without plexes being flagged as stale, as long as the reattach happens before any volumes on the disk are started.

vxreattach tries to find a disk in the same disk group with the same disk ID for the disk(s) to be reattached. The reattach operation may fail even after finding the disk with the matching disk ID if the original cause (or some other cause) for the disk failure still exists.

vxreattach is called as part of disk recovery from the vxdiskadm menus.

#### OPTIONS

- b Performs the reattach operation in the background.
- c Checks if a reattach is possible. No operation is performed, but the name of the disk group and disk media name at which the disk can be reattached is displayed.
- r Tries to recover stale plexes of any volumes on the failed disk. It does this by calling vxrecover.

#### EXIT CODES

A zero exit status is returned if the reattach is performed; non-zero is returned otherwise. See vxintro(1M) for a list of standard exit codes.

#### SEE ALSO

vxdiskadm(1M), vxintro(1M), vxrecover(1M)

===[ man1m/vxrecover.1m ]=====

vxrecover - perform volume recovery operations

#### SYNOPSIS

```
vxrecover [-bcEnpsvV] [-g diskgroup ] [-o options ] [-t
tasktag ] [ volume | medianame...]
```

#### DESCRIPTION

The vxrecover program performs plex attach, RAID-5 subdisk recovery, and resynchronize operations for the named volumes, or for volumes residing on the named disks (medianame). If no medianame or volume operands are specified, then the operation applies to all volumes (or to all volumes in the specified disk group). If -s is specified, then disabled volumes will be started. With -s and -n, volumes are started, but no other recovery takes place.

Recovery operations will be started in an order that prevents two concurrent operations from involving the same disk. Operations that involve unrelated disks will run in parallel.

#### OPTIONS

- b Perform recovery operations in the background. With this option, vxrecover will put itself in the background to attach stale plexes and subdisks, and to resynchronize mirrored volumes and RAID-5 parity. If this is used with -s, then volumes will be started before recovery begins in the background.
- c Operate on cluster-shareable disk groups. This option is used by the cluster reconfiguration utility to effect recovery on volumes in shared disk groups after the failure of a node in the cluster. Do not use -c as a command-line option. Available only if the Volume Manager cluster feature is enabled.
- E Start disabled volumes or plexes even when they are in the EMPTY state. This is useful for starting up volumes restored by the vxmake utility when specified along with the -s option.
- g diskgroup  
Limit operation of the command to the given disk group, as specified by disk group ID or disk group name. If no volume or medianame operands are given, then all disks in this disk group will be recovered; otherwise, the volume and medianame operands will be evaluated relative to the given disk group.

Without the `-g` option, if no operands are given, then all volumes in all imported disk groups will be recovered; otherwise, the disk group for each medianame operand will be determined based on name uniqueness within all disk groups.

`-n` Do not perform any recovery operations. If used with `-s`, then volumes will be started, but no other actions will be taken. If used with `-p`, then the only action of `vxrecover` will be to print a list of startable volumes.

`-o options`

Pass the given option arguments to the `-o options` for the `vxplex att` and `vxvol start` operations generated by `vxrecover`. An option argument of the form `prefix:options` can be specified to restrict the set of commands that the `-o options` should be applied to. Defined prefixes are: `vol` applies to all invocations of the `vxvol` utility (volume starts, mirror resynchronizations, RAID-5 parity rebuilds, and RAID-5 subdisk recoveries); `plex`, applies to all invocations of the `vxplex` utility (currently used only for attaching plexes); `attach` applies specifically to `plex attach` operations; `start` applies specifically to volume start operations; `recover` applies to subdisk recoveries; and `resync` applies to mirror resynchronization and RAID-5 parity recovery.

`-p` Print the list of selected volumes that are startable. For each startable volume, a line is printed containing the following information: the volume name, the disk group ID of the volume, the volume's usage type, and a list of state flags pertaining to mirrors of the volume. State flags include the following: `kdetach` indicates that one of the mirrors was detached by an I/O failure; `stale` means that one of the mirrors needs recovery, but that the recovery is related to an administrative operation, not an I/O failure; `stopped` is printed if neither `kdetach` nor `stale` is appropriate for the volume.

`-s` Start disabled volumes that are selected by the operation. Volumes are started before any other recovery actions are taken. Volumes are started with the `-o delayrecover start` option. This requests that any operations that can be delayed in starting a volume will be delayed. In other words, only those operations necessary to make a

volume available for use will occur. Other

VxVM 3.0

Last change: 26 Mar 1999

2

Maintenance Commands

vxrecover(1M)

operations, such as mirror resynchronization, attaching of stale plexes and subdisks, and recovery of stale RAID-5 parity will normally be delayed.

-t tasktag

If any tasks are registered to track the progress of the operation, mark them with the tag tasktag. This option is also passed to any utilities called to perform recoveries, so all tasks registered by any utility to perform the recoveries will also be tagged with tasktag.

-v

Display information about each task started by vxrecover. For recovery operations (as opposed to start operations), print a completion status when each task completes.

-V

Display each command executed by vxrecover. The output generated is more verbose than the job descriptions printed with the -v option.

#### EXAMPLES

To recover, in the background, any detached subdisks or plexes that resulted from replacement of a specified disk, use the command:

```
vxrecover -b medianame
```

If you want to monitor the operations, use the command:

```
vxrecover -v medianame
```

#### SEE ALSO

vxintro(1M), vxplex(1M), vxvol(1M)

===[ man1m/vxrelayout.1m ]=====

Maintenance Commands

vxrelayout(1M)

#### NAME

vxrelayout - convert online storage from one layout to another

#### SYNOPSIS

vxrelayout [-f] [-g diskgroup] [-o useopt]  
[-U usetype] reverse volume

vxrelayout [-f] [-g diskgroup] [-o useopt]  
[-U usetype] start volume

vxrelayout status volume

#### DESCRIPTION

Online relayout is a process where storage in a volume or a plex is converted from one layout to another maintaining data availability at all times. Thus a user can convert any supported layout in the Volume Manager to another and/or change the stripe width or the number columns. The user can also do just stripe width and/or column changes if so desired.

Such a transformation of a volume or a plex has to be initiated using the vxassist(1M) utility. vxassist creates necessary infrastructure and storage needed to perform the layout transformation and invokes vxrelayout internally to actually perform the transformation. Since vxrelayout depends on vxassist to do all the setup before a layout transformation can be commenced, the vxassist utility has to be used to initiate any layout transformation. The vxtask(1M) utility can be used monitor the state of an ongoing transformation.

Since these transformations are time consuming (mainly depends on the volume size and chunk size), they can be interrupted because of a system crash, I/O failure, or the user voluntarily stopping the transformation process either by aborting the task associated with the transformation using `vxtask` or by killing the transformation process. Using `kill -9` is not advisable to stop such transformations. The `vxrelayout` utility can be used to continue or reverse a transformation process because of such interruptions.

The reaction of the Volume Manager to such interruptions are:

I/O Failure                    `vxrelayout` utility exits with an error on `stderr`. The `vxrelocd(1M)` daemon tries to relocate the failed subdisks. If it succeeds, it invokes `vxrecover(1M)`, which in turn invokes `vxrelayout` to continue the transformation. An I/O failure

VxVM 3.0

Last change: 26 Mar 1999

1

Maintenance Commands

`vxrelayout(1M)`

is transparent to the user if the relocation succeeds. If relocation fails, the transformation process is not restarted again.

System Crash

On a reboot after a system crash, `vxrecover` starts all startable volumes, invoking `vxrelayout` on appropriate volumes and plexes.

If the user stops `vxrelayout`, then the transformation process can be reversed or continued by invoking `vxrelayout` manually.

The first operand is a keyword that determines whether a layout conversion is to be continued, reversed, or to display the status of any existing transformation in the volume. For example, if a RAID-5 volume is being converted to STRIPED volume, reversal of such an operation will revert all data and storage in STRIPED format to the RAID-5 format. A reversal of a reversal is equivalent to a continue.

Even though such a capability of direction change can be

used ad infinitum, it has been provided with an objective to react to unforeseen situations. It is advisable that such direction changes are done with care.

The direction of a transformation can be determined using the status operation provided in the vxrelayout utility. This command will determine if any plex in the volume is being transformed and display its characteristics.

The second operand is the name of the topmost volume to which the plex undergoing relayout is attached. If there are multiple plexes in the volume and only one of them is being transformed, the vxrelayout utility will select the appropriate plex using the supplied volume name. It should be noted that there can only be one relayout operation happening in a volume at any instant of time. A specific disk group can be forced with -g diskgroup.

This operation can be applied only to disk groups with version 4.4 or above (see vxdg(1M)) and only one disk group at a time due to internal implementation constraints.

#### KEYWORDS

reverse Reverses a discontinued layout conversion. This operation determines the current state and direction of a layout transformation and reverse where it was left off since it was last stopped. It is imperative that any existing layout transformation

#### Maintenance Commands

#### vxrelayout(1M)

process is stopped before this operation is applied.

start Continues a discontinued layout conversion. This operation determines the current state and direction of a layout transformation and continue where it was left off since it was last stopped. It is imperative that any existing layout transformation process is stopped before this operation is applied.

status Displays the status of a discontinued or an ongoing layout conversion. This operation displays the characteristics of the source and the destination layout, such as the layout, number of columns, and stripe width. It also displays whether the



transformation is ongoing or stopped and the percentage completed.

#### OPTIONS

- f Forces a continue or reverse of a layout transformation. This is considered potentially dangerous or questionable to use as using this flag may result in loss of data. This permits a limited set of operations that would otherwise be disallowed:
- o It is possible that an I/O failure during a layout transformation caused some plexes to detach, If the user still wants to continue or reverse such a transformation, this flag will have to be used to tell vxrelayout to ignore the fact some plexes have failed. There is no guarantee that this operation will succeed.
  - o vxrelayout shrinks any existing filesystem on a volume if the layout transformation causes the volume size to decrease. If the user does not want vxrelayout to shrink the filesystem but wants to still perform the layout conversion, this flag must be used.
- g diskgroup Specifies the disk group for the operation, either by disk group ID or by disk group name. By default, the disk group is chosen based on the name operands.
- o useopt Passes in usage-type-specific options to the operation. A certain set of operations are expected to be implemented by all usage types:

#### slow[=iodelay]

Reduces the system performance impact of copy operations. Copy operations are usually a set of short copy operations on small regions of the volume (normally from 16 kilobytes to 128 kilobytes). This option inserts a delay between the recovery of each such region. A specific delay can be specified with iodelay

as a number of milliseconds; otherwise, a default is chosen (normally 250 milliseconds).

`iosize=size`

Performs copy operations in regions with the length specified by `size`, which is a standard Volume Manager length number (see `vxintro(1M)`). Specifying a larger number typically causes the operation to complete sooner, but with greater impact on other processes using the volume. The default I/O size is typically 32 kilobytes.

Note: These parameters can also be changed using the `vxtask` utility during an online relayout.

`BG` Run the `vxrelayout` operation in the background.

`-U usetype`

Limits the operation to apply to this usage type. Attempts to affect volumes with a different usage type will fail.

#### EXIT CODES

`vxrelayout` exits with a nonzero status if the operation fails. A nonzero exit code is not a complete description of the problem; it only indicates the first condition that prevented further execution.

See `vxintro(1M)` for a list of standard exit codes.

#### SEE ALSO

`kill(1)`, `vxassist(1M)`, `vxdg(1M)`, `vxintro(1M)`, `vxrecover(1M)`, `vxrelocd(1M)`, `vxtask(1M)`

## NAME

vxrelocd - monitor the Volume Manager for failure events and relocate failed subdisks

## SYNOPSIS

```
/etc/vx/bin/vxrelocd [-o vxrecover_argument]
[mail_address...]
```

## DESCRIPTION

The vxrelocd command monitors the Volume Manager by analyzing the output of the vxnotify command, waiting for failures to occur. When a failure occurs, vxrelocd sends mail via mailx to root (by default) or to other specified users and relocates failed subdisks. After completing the relocation, vxrelocd sends more mail indicating the status of each subdisk replacement. The vxrecover utility is then run on volumes with relocated subdisks to restore data.

## OPTIONS

-o The -o option and its argument are passed directly to vxrecover if vxrecover is called. This allows specifying -o slow[=iodelay] to keep vxrecover from overloading a busy system during recovery. The default value for the delay is 250 milliseconds.

## Mail Notification

By default, vxrelocd sends mail to root with information about a detected failure and the status of any relocation and recovery attempts. To send this mail to other users, add the user logins to the vxrelocd startup line in the startup script /etc/rc2.d/s95vxvm-recover and reboot the system. Alternatively, you can kill the vxrelocd process and restart it as vxrelocd root mail\_address, where mail\_address is a user's login. Do not kill the vxrelocd process while a relocation attempt is in progress.

The mail notification that is sent when a failure is detected follows this format:

```
Failures have been detected by the VERITAS Volume Manager:
```

```
failed disks:
medianame
...
failed plexes:
plexname
...
failed log plexes:
plexname
...
failing disks:
```

## Maintenance Commands

vxrelocd(1M)

```
medianame
...
failed subdisks:
subdiskname
...
```

The Volume Manager will attempt to find spare disks, relocate failed subdisks and then recover the data in the failed plexes.

The medianame list under failed disks: specifies disks that appear to have completely failed; the medianame list under failing disks: indicates a partial disk failure or a disk that is in the process of failing. When a disk has failed completely, the same medianame list appears under both failed disks: and failing disks:. The plexname list under failed plexes: shows plexes that have been detached due to I/O failures experienced while attempting to do I/O to subdisks they contain. The plexname list under failed log plexes: indicates RAID-5 or DRL log plexes that have experienced failures. The subdiskname list specifies subdisks in RAID-5 volumes that have been detached due to I/O errors.

## Spare Space

A disk can be marked as ``spare.'' This makes the disk available as a site for relocating failed subdisks. Disks that are marked as spares are not used for normal allocations unless you explicitly specify them. This ensures that there is a pool of spare space available for relocating failed subdisks and that this space won't get consumed by normal operations. Spare space is the first space used to relocate failed subdisks. However, if no spare space is available, or the available spare space is not suitable or sufficient, free space is also used. See the vxedit(1M) and vxdiskadm(1M) manual pages for more information on marking a disk as a spare.

## Replacement Procedure

After mail is sent, vxrelocd relocates failed subdisks (those listed in the subdisks list). This involves finding appropriate spare and/or free space in the same disk group as the failed subdisk. A disk is eligible as replacement space if it is a valid Volume Manager disk and contains enough space to hold the data contained in the failed subdisk. If no space is available on spare disks, the relocation uses free space instead.

If no spare or free space is found, mail is sent explaining the disposition of volumes that had storage on the failed disk:

VxVM 3.0

Last change: 26 Mar 1999

2

Maintenance Commands

vxrelocd(1M)

Hot-relocation was not successful for subdisks on disk dm\_name in volume v\_name in disk group dg\_name. No replacement was made and the disk is still unusable.

The following volumes have storage on medianame:

volumename

These volumes are still usable, but the redundancy of those volumes is reduced. Any RAID-5 volumes with storage on the failed disk may become unusable in the face of further failures.

If any non-RAID-5 volumes were made unusable due to the disk failure, the following message is included:

The following volumes:

volumename

have data on medianame but have no other usable mirrors on other disks. These volumes are now unusable and the data on them is unavailable. These volumes must have their data restored.

If any RAID-5 volumes were made unavailable due to the disk failure, the following message is included:

The following RAID-5 volumes:

volumename

had storage on medianame and have experienced other failures. These RAID-5 volumes are now unusable and data on them is unavailable. These RAID-5 volumes

must

have their data restored.

If there is spare space available, subdisk relocations are attempted. This involves setting up a subdisk on the spare or free space and using it to replace the failed subdisk. If this is successful, the vxrecover command is run in the background to recover the data in volumes that had storage on the disk.

If the relocation fails, the following message is sent:

```
Hot-relocation was not successful for subdisks
on disk dm_name in volume v_name in disk
group dg_name. No replacement was made
```

VxVM 3.0

Last change: 26 Mar 1999

3

Maintenance Commands

vxrelocd(1M)

and the disk is still unusable.

error message

If any volumes (RAID-5 or otherwise) are rendered unusable due to the failure, the following message is included:

The following volumes:

volumename

other  
is  
have data on dm\_name but have no other usable mirrors on  
disks. These volumes are now unusable and the data on them  
unavailable. These volumes must have their data restored.

If the relocation procedure was successful and recovery is under way, the following mail message is sent:

```
Volume v_name Subdisk sd_name relocated to
newsd_name, but not yet recovered.
```

After recovery completes, a mail message is sent relaying the outcome of the recovery procedure. If the recovery is successful, the following message is included in the mail:

Recovery complete for volume v\_name in disk  
group dg\_name.

If the recovery was not successful, the following message is  
included in the mail:

Failure recovering v\_name in disk group dg\_name.

#### Disabling vxrelocd

If you do not want automatic subdisk relocation, you can  
disable the hot-relocation feature by killing the relocation  
daemon, vxrelocd, and preventing it from restarting. How-  
ever, do not kill the daemon while it is doing the reloca-  
tion. To kill the daemon, run the command:

```
ps -ef
```

from the command line and find the two entries for vxrelocd.  
Execute the command:

VxVM 3.0

Last change: 26 Mar 1999

4

Maintenance Commands

vxrelocd(1M)

```
kill -9 PID1 PID2
```

(substituting PID1 and PID2 with the process IDs for the two  
vxrelocd processes). To prevent vxrelocd from being started  
again, you must comment out the line that starts up vxrelocd  
in the startup script /etc/rc2.d/s95vxvm-recover.

#### FILES

/etc/rc2.d/s95vxvm-recover     The startup file for vxrelocd.

#### SEE ALSO

kill(1), mailx(1), ps(1), vxdiskadm(1M), vxedit(1M),  
vxintro(1M), vxnotify(1M), vxrecover(1M)

===[ man1m/vxresize.1m ]=====

Maintenance Commands

vxresize(1M)

NAME

vxresize - change the length of a volume containing a file system

SYNOPSIS

/etc/vx/bin/vxresize [-bsx] [-F fstype] [-g diskgroup] [-t tasktag] volume new\_length [medianame...]

DESCRIPTION



The `vxresize` command will either grow or shrink both the file system and its underlying volume to match the specified new volume length. The ability to grow or shrink is file system dependent. Some file system types may require that the file system be unmounted for the operation to succeed.

Note: `vxresize` works with VxFS and UFS file systems only.

In some situations, when resizing large volumes, `vxresize` may take a long time to complete.

The `new_length` operand can begin with a plus (+) or minus (-) to indicate that the new length is added to or subtracted from the current volume length.

The `medianame` operands name disks to use for allocating new space for a volume. These arguments can be a simple name for a disk media record, or they can be of the form `medianame, offset` to specify an offset within the named disk. If an offset is specified, then regions from that offset to the end of the disk are considered candidates for allocation.

#### OPTIONS

- b Performs the resize operation in the background. The command returns quickly, but the resize will be in progress. Use the `vxprint` command to determine when the operation completes.
- F `fstype` Supplies the type of the file system to be resized.
- g `diskgroup` Limits operation of the command to the given disk group, as specified by disk group ID or disk group name. The volume operand is evaluated relative to the given disk group.
- s Requires that the operation represent a decrease in the volume length.
- t `tasktag` If any tasks are registered to track the progress

-x            Requires that the operation represent an increase  
              in the volume length. Fail the operation other-  
              wise.

SEE ALSO

vxintro(1M), vxprint(1M), vxtask(1M)

===[ manlm/vxrootmir.1m ]=====

Maintenance Commands

vxrootmir(1M)

NAME

vxrootmir - mirror areas necessary for booting to a new disk

SYNOPSIS

/etc/vx/bin/vxrootmir [-vV] [-t tasktag] medianame

DESCRIPTION

The vxrootmir utility creates mirrors of volumes required in booting. It creates a mirror for rootvol. The mirror is created on the specified disk media device. The subdisks for this volume are placed on cylinder boundaries and a partition is created for them.

The specified disk media device must have enough space to contain the mirror for rootvol or else the create fails. Also, slice zero must be free because it is used to create the partition for root.

All disk regions required for booting are set up by invoking the installboot command and partitions for the new volume mirrors are created.

This script is called by the vxmirror command if the root disk is required to be mirrored. It is also called from the vxdiskadm menus through the choice of the mirror volumes on a disk operation.

SEE ALSO

installboot(1M), vxdiskadm(1M), vxintro(1M), vxmirror(1M), vxtask(1M)

===[ man1m/vxsd.1m ]=====

Maintenance Commands

vxsd(1M)

#### NAME

vxsd - perform Volume Manager operations on subdisks

#### SYNOPSIS

vxsd [-Vf] [-g diskgroup] [-o useopt ] [-t tasktag] [-U use-  
type ] [-v volume ] aslog plex subdisk

vxsd [-Vf] [-g diskgroup] [-l offset ] [-o useopt ] [-t  
tasktag] [-U usetype ] [-v volume ] assoc plex subdisk...

vxsd [-Vf] [-g diskgroup] [-l column[/offset ] [-o useopt ]  
[-t tasktag] [-U usetype ] [-v volume ] assoc plex subdisk  
[:column[/offset ]]....

vxsd [-Vf] [-g diskgroup] [-o useopt ] [-p plex ] [-t task-  
tag] [-U usetype ] [-v volume ] dis subdisk...

vxsd [-Vf] [-g diskgroup] [-o useopt ] [-p plex ] [-t task-  
tag] [-U usetype ] [-v volume ] join sd1 sd2...newsd

vxsd [-Vf] [-g diskgroup] [-o useopt ] [-p plex ] [-t task-  
tag] [-U usetype ] [-v volume ] mv oldsd newsd [newsd...]

vxsd [-Vf] [-g diskgroup] [-o useopt ] [-p plex ] [-s size]  
[-t tasktag] [-U usetype ] [-v volume ] split subdisk newsd  
[newsd2]...

#### DESCRIPTION

The vxsd utility performs Volume Manager operations on subdisks and on plex-and-subdisk combinations. The first operand is a keyword that determines the specific operation to perform. The remaining operands specify the configuration objects to which the operation is to be applied.

Each operation can be applied to only one disk group at a time, due to internal implementation constraints. Any plex or subdisk name operands will be used to determine a default

disk group, according to the standard disk group selection rules described in vxintro(1M). A specific disk group can be selected with -g diskgroup.

If a vxsd operation is interrupted by a signal, then an attempt is made to restore the disk group configuration to a state that is roughly equivalent to its original state. If this attempt is interrupted, such as through another signal, then the user may need to perform some cleanup. The specific cleanup actions that are needed are written to the standard error before vxsd exits.

#### KEYWORDS

aslog Associates the named subdisk with the named plex as a log area for the plex. At most, one log

VxVM 3.0

Last change: 26 Mar 1999

1

Maintenance Commands

vxsd(1M)

subdisk can be associated with a plex at any given time. Currently, log subdisks can be used only with the dirty region logging feature, as defined by the DRL volume logging type. If the named plex is associated with a volume, then the rules for performing the operation depend upon the usage type of the volume. A subdisk cannot be associated if the putil0 field is set on the subdisk, just as with vxsd assoc.

assoc Associates each named subdisk operand with the specified plex. The first form applies to concatenated plexes (that is, with a layout of concat). The offset within the plex for the association can be specified with -l, which takes a standard Volume Manager length number (see vxintro(1M)). If no offset is specified, then the default is to associate the subdisk at the end of the plex, thus extending the length of the plex by the length of the new subdisk.

For striped plexes, a column number for the subdisk association may be specified. The offset is interpreted as the column offset for the subdisk. If only one number is specified with -l for striped plexes, the number is interpreted as a column number and the subdisk is associated at the end of the column.

The column or column/offset at which a subdisk is to be associated can also be specified as part of the subdisk name in the same manner as subdisk associations are specified for plex creations in vxmake (see vxmake(4)). When specifying multiple subdisks, if no column or column/offset is specified for a subdisk, it is associated after the previous subdisk.

A subdisk cannot be associated to overlap with another associated subdisk in the same plex.

If the named plex is associated with a volume, then the rules for performing the operation depend upon the usage type of the volume. A subdisk cannot be associated to a plex if the putil0 field for the subdisk is not empty. Creating a subdisk with the putil0 field set to a non-empty value is a sufficient means of ensuring that no Volume Manager operation will write to the region of disk blocks allocated to the subdisk because the subdisk cannot be associated through any means to a plex, and because subdisks cannot be used directly

## Maintenance Commands

vxsd(1M)

to read from or write to a disk.

**dis** Dissociates each specified subdisk from the plex that it is associated with. If a subdisk is associated (through its plex) with a volume, then the rules for performing the operation depend upon the usage type of the volume.

Subdisk dissociation can be used as part of tearing down a plex, or as part of reorganization of disk space usage. Typically, the subdisk is no longer needed after dissociation. To support this type of use, -o rm can be specified to remove the named subdisks after successful dissociation.

**join** Joins the subdisks named by the sd operands to form a new subdisk named newsd. The sd operands must specify subdisks that represent contiguous sections of the same device, and of the same plex (if they are associated). For a striped plex, the sd operands must be in the same column. At least

two sd operands are required. At the end of the operation, the sd configuration objects are removed. The newsd operand can have the same name as one of the sd operands, or it can have a different name.

If the sd operands are associated with an associated plex, then the rules for performing the operation depend upon the usage type of the volume.

mv Moves the contents of oldsd onto the new subdisks and replaces oldsd with the new subdisks for any associations. If multiple new subdisks are specified, they are associated starting where the old subdisk began and placed consecutively with no space between them. The operation requires that oldsd be associated with an associated plex and that all new subdisks be dissociated. The operation can be used on a subdisk that is used by an active volume, and will ensure that data is copied and associations are changed without loss or corruption of data. The rules for performing the operation depend upon the usage type of the volume.

Moving a subdisk is the normal means of reorganizing disk space. For example, move regions of disk used by one volume to another disk to reduce contention on the original disk. Typically, once the operation completes, the original subdisk is no

longer needed and can be removed. To support this use of the operation, -o rm can be specified to remove oldsd after successful completion of the operation.

split Split the subdisk subdisk into two subdisks that reside on the same section of the same device, and that have contiguous plex associations (if the named subdisk is associated). The first of the two resultant subdisks will have a length of size, and the second will take up the remainder of the space used by the original subdisk. If both newsd and newsd2 are specified, then the resulting sub-

disks are newsd and newsd2. If no newsd2 operand was specified, then the resultant subdisks are named subdisk and newsd.

If the named subdisk is associated with an associated plex, then the rules for performing the operation depend upon the usage type of the volume. Log subdisks cannot be split.

#### OPTIONS

-f Force an operation that the Volume Manager considers potentially dangerous or of questionable use. This permits a limited set of operations that would otherwise be disallowed. Some operations may be disallowed even with this flag.

-g diskgroup Specify the disk group for the operation, either by disk group ID or by disk group name. By default, the disk group is chosen based on the name operands.

-l [column]offset Specify the offset of a subdisk within a plex address space for the vxsd assoc operation. For striped plexes, a column number may be optionally specified. If one number is given for striped plexes the number is interpreted as a column number and the subdisk is associated at the end of the column. The offset is a standard Volume Manager length number (see vxintro(1M)).

-o useopt Pass in usage-type-specific options to the operation. A certain set of operations are expected to be implemented by all usage types:

iosize=size Perform copy operations in regions with the

length specified by size, which is a standard Volume Manager length number (see vxintro(1M)). Specifying a larger number typically causes the operation to complete sooner, but with greater impact on other



processes using the volume. The default I/O size is typically 32 kilobytes.

rm Remove the subdisks after successful completion of a vxsd dis operation. Remove the source subdisk after successful completion of vxsd mv.

slow[=iodelay]

Reduce the system performance impact of copy operations. Copy operations are usually a set of short copy operations on small regions of the volume (normally from 16 kilobytes to 128 kilobytes). This option inserts a delay between the recovery of each such region. A specific delay can be specified with iodelay as a number of milliseconds, or a default is chosen (normally 250 milliseconds).

-p plex Require that a named subdisk (source subdisk for vxsd split, join, and mv) be associated with this plex.

-s size Specify the size for the subdisk split operation. This option takes a standard Volume Manager length number (see vxintro(1M)).

-t tasktag  
If any tasks are registered to track the progress of the operation, mark them with the tag tasktag.

-U usetype  
Limit the operation to apply to this usage type. Attempts to affect volumes with a different usage type will fail.

-v volume  
Require that a named plex be associated with this volume, or that a named subdisk (source subdisk for vxsd split, join, and mv) be associated with a plex that is associated with this volume.

-V  
Write a list of utilities that would be called from vxsd, along with the arguments that would be passed. The -V performs a ``mock run'' so the utilities are not actually called.

## FSGEN and GEN Usage Types

The fsgen and gen usage types provide identical semantics for all operations of the vxsd utility.

In addition to the standard -o options required for all usage types, the fsgen and gen usage types provide the following additional option:

**force** Forces an operation that the Volume Manager considers potentially dangerous or of questionable use. This applies to attempts to dissociate subdisks (making a plex sparse) and to attempts to move subdisks onto subdisks that have a different size. This flag is the same as -f.

Limitations and extensions for the fsgen and gen usage types consist of the following:

**aslog** If a log subdisk is associated with a plex that is associated with a volume that has a logging type of UNDEF, then the logging type of the volume is converted to DRL. Logging of volume changes is not enabled until there are at least two read-write mode plexes attached to the volume.

**assoc** If the named plex is enabled, and is associated with an enabled plex, then the named plex must be ACTIVE or EMPTY. Subdisks can be associated with a non-enabled plex only if the utility state of the plex is EMPTY, STALE, or OFFLINE, or if the plex is CLEAN and no other plexes associated with the volume are CLEAN or ACTIVE.

If the subdisk is associated with a non-enabled plex, or if it is associated with the only enabled, read-write plex in a volume, then the operation completes without copying any data onto the subdisk. If the subdisk is associated with an enabled plex in a mirrored volume, then the operation may have to copy data from the volume onto the new subdisk before the operation can complete.

**dis** Dissociating a subdisk requires use of -f if it would cause an enabled plex in an enabled volume to become sparse relative to the volume. Even with -f, it is not possible to make two plexes sparse if no complete, enabled, read-write plexes would remain associated. For disabled volumes, a similar check is made with respect to ACTIVE and CLEAN plexes.

**mv** If the total size of the destination subdisks

## Maintenance Commands

vxsd(1M)

differs from that of the source subdisk, then the -f option must be specified. The operation still fails if the total size of the destination subdisks is larger than the source subdisk and if the address range of any destination subdisk would conflict with another subdisk that is associated with the plex. The total size of the destination subdisks cannot be larger than the source subdisk if the kernel state of the volume or plex is detached.

The operation fails if the total size of the destination subdisks is smaller than the source subdisk and the operation would cause the total number of complete, enabled, read-mode plexes in the volume to drop to zero, while leaving more than one sparse, enabled, read-write plex.

## split and join

The fsgen and gen usage types apply no additional restrictions and add no extensions to the split and join operations.

## RAID-5 Usage Type

In addition to the standard -o options required for all usage types, the raid5 usage type provides the following additional option:

**force** Forces an operation that the Volume Manager considers potentially dangerous or of questionable use. This applies to attempts to move a subdisk in a RAID-5 plex if the volume the plex is associated with does not have a log plex. This flag is the same as -f.

The raid5 usage type supports the following keywords:

**assoc** Associate the named subdisks with the named RAID-5 plex. If plex is enabled and is associated with an enabled volume, then any data that maps onto the subdisk will be regenerated from the other columns of the RAID-5 plex. This is done by marking the subdisk as stale and writeonly, regenerating the data via VOL\_R5\_RECOVER ioctls, and then turning off the stale and writeonly flags.

If the RAID-5 plex is not associated or the RAID-5

volume is not ENABLED, the subdisk is associated and marked as stale. The subdisk's contents will be recovered when the volume is started.

The assoc operation may not be used on a log plex.

Maintenance Commands

vxsd(1M)

**dis** Dissociate the named subdisks from the named RAID-5 plex. If removing the subdisk would make the volume unusable (because other subdisks in other columns at the same altitude are unusable or missing) and the volume is not disabled and EMPTY, the operation is not allowed. If the volume is disabled and non-EMPTY the operation requires use of -f.

The dis operation may not be used on a log plex.

**mv** If the old subdisk is associated with a RAID-5 plex that is associated to a RAID-5 volume, the volume must be enabled for the move operation to complete. The mv operation creates a redundant subvolume structure and synchronizes the new subdisk(s) using VOL\_R5\_RECOVER ioctls. This makes the operation safe in case of failures in the new subdisks, and does not put the integrity of the data at risk.

The mv operation may be used on a log plex. Similar rules as those defined for the fsgen and gen usage type plexes apply to log plexes. The force or -f must be used to make a log plex sparse.

Note that there is no aslog operation for the raid5 usage type. Logging is done on a plex level and therefore vxsd aslog is not needed. Log plexes can be associated with RAID-5 volumes using the vxplex att command.

**split and join**

The raid5 usage type applies no additional restrictions and adds no extensions to the split and join operations. These operations may not be used on a log plex.

/etc/vx/type/usetype/vxsd      The utility that performs vxsd operations for a particular volume usage type.

#### EXIT CODES

The vxsd utility exits with a nonzero status if the attempted operation fails. A nonzero exit code is not a complete indicator of the problems encountered, but rather denotes the first condition that prevented further execution of the utility. See vxintro(1M) for a list of standard exit codes.

VxVM 3.0

Last change: 26 Mar 1999

8

Maintenance Commands

vxsd(1M)

#### SEE ALSO

vxintro(1M), vxmake(1M), vxplex(1M), vxtask(1M), vxvol(1M)

=== [ man1m/vxserial.1m ] =====

Maintenance Commands

vxlicense(1M)

NAME

vxlicense - VERITAS licensing key utility

SYNOPSIS

vxlicense [-c] [-p] [-t feature]

DESCRIPTION

vxlicense, the VERITAS licensing key utility, maintains the VERITAS license key file.

OPTIONS

- c           Creates a license key file. vxlicense prompts for a license key.
  
- p           Prints available VERITAS licenses and features in a system.
  
- t feature   Tests a VERITAS license in a system.

## EXAMPLES

The following example shows how to use vxlicense to initialize the VERITAS licensing key file. After entering the command, vxlicense prompts for a license key. It creates the license key file after you enter the key.

```
vxlicense -c
```

To list available features, enter:

```
vxlicense -p
```

```
Feature name: CURRSET [95]
Number of licenses: 1 (non-floating)
Expiration date: 12/31/99
Release Level: 20
Machine Class: 934986342
```

```
Feature name: RAID [96]
Number of licenses: 1 (non-floating)
Expiration date: 12/31/99
Release Level: 20
Machine Class: 934986342
```

VERITAS

Last change: 26 Mar 1999

1

```
===[ man1m/vxsparecheck.1m ]=====
```

Maintenance Commands

vxsparecheck(1M)

### NAME

vxsparecheck - monitor Volume Manager for failure events and replace failed disks

### SYNOPSIS

```
/etc/vx/bin/vxsparecheck [mail-address...]
```

### DESCRIPTION

The vxsparecheck command monitors the Volume Manager by

analyzing the output of the vxnotify command, waiting for failures to occur. It then sends mail via mailx to the login(s) specified on the command line, or (by default) to root. It then replaces any failed disks. After an attempt at replacement is complete, mail will be sent indicating the status of each disk replacement.

The mail notification that is sent when a failure is detected follows this format:

```
Failures have been detected by the VERITAS Volume Manager:
```

```
failed disks:  
medianame  
...  
failed plexes:  
plexname  
...  
failed subdisks:  
subdiskname  
...  
failed volumes:  
volumename  
...
```

The Volume Manager will attempt to find hot-spare disks to replace any failed disks and attempt to reconstruct any data in volumes that have storage on the failed disk.

The medianame list specifies disks that appear to have completely failed. The plexname list show plexes of mirrored volumes that have been detached due to I/O failures experienced while attempting to do I/O to subdisks they contain. The subdiskname list specifies subdisks in RAID-5 volumes that have been detached due to I/O errors. The volumename list shows non-RAID-5 volumes that have become unusable because disks in all of their plexes have failed (and are listed in the ``failed disks'' list) and shows those RAID-5 volumes that have become unusable because of multiple failures.



If any volumes appear to have failed, the following paragraph will be included in the mail:

The data in the failed volumes listed above is no longer available. It will need to be restored from backup.

#### Replacement Procedure

After mail has been sent, vxsparecheck finds a hot spare replacement for any disks that appear to have failed (that is, those listed in the medianame list). This involves finding an appropriate replacement for those eligible hot spares in the same disk group as the failed disk. A disk is eligible as a replacement if it is a valid Volume Manager disk, has been marked as a hot-spare disk and contains enough space to hold the data contained in all the subdisks on the failed disk.

To determine which disk from among the eligible hot spares to use, vxsparecheck first checks the file /etc/vx/sparelist (see Sparelist File below). If this file does not exist or lists no eligible hot spares for the failed disk, the disk that is ``closest'' to the failed disk is chosen. The value of ``closeness'' depends on the controller, target and disk number of the failed disk. A disk on the same controller as the failed disk is closer than a disk on a different controller; and a disk under the same target as the failed disk is closer than one under a different target.

If no hot spare disk can be found, the following mail is sent:

still           No hot spare could be found for disk medianame in  
                  diskgroup. No replacement has been made and the disk is  
                  unusable.

The mail then explains the disposition of volumes that had storage on the failed disk. The following message lists disks that had storage on the failed disk, but are still usable:

The following volumes have storage on medianame:

volumename

further        These volumes are still usable, but the redundancy of  
                  those volumes is reduced. Any RAID-5 volumes with storage  
                  on the failed disk may become unusable in the face of  
                  failures.

## Maintenance Commands

vxsparecheck(1M)

If any non-RAID-5 volumes were made unusable due to the failure of the disk, the following message is included:

The following volumes:

volumename

have data on medianame but have no other usable mirrors on other disks. These volumes are now unusable and the data on them is unavailable.

If any RAID-5 volumes were made unavailable due to the disk failure, the following message is included

The following RAID-5 volumes:

volumename

had storage on medianame and have experienced other failures. These RAID-5 volumes are now unusable and data on them is unavailable.

If a hot-spare disk was found, a hot-spare replacement is attempted. This involves associating the device marked as a hot spare with the media record that was associated with the failed disk. If this is successful, the vxrecover(1M) command is used in the background to recover the contents of any data in volumes that had storage on the disk.

If the hot-spare replacement fails, the following message is sent:

Replacement of disk medianame in group diskgroup failed. The error is:

error message

If any volumes (RAID-5 or otherwise) are rendered unusable due to the failure, the following message is included:

The following volumes:

volumename

available occupy space on the failed disk and have no other  
mirrors or have experienced other failures. These volumes  
are unusable, and the data they contain is unavailable.

VxVM 3.0

Last change: 26 Mar 1999

3

Maintenance Commands

vxsparecheck(1M)

If the hot-spare replacement procedure completed successfully and recovery is under way, a final mail message is sent:

Replacement of disk medianame in group diskgroup  
with disk device sparedevice has successfully completed  
and recovery is under way.

If any non-RAID-5 volumes were rendered unusable by the failure despite the successful hot-spare procedure, the following message is included in the mail:

The following volumes:

volumename

enabled occupy spare on the replaced disk, but have no other  
mirrors on other disks from which to perform recovery.  
These volumes must have their data restored.

If any RAID-5 volumes were rendered unusable by the failure despite the successful hot-spare procedure, the following message is included in the mail:

The following RAID-5 volumes:

volumename

have subdisks on the replaced disk and have experienced other failures that prevent recovery. These RAID-5 volumes must have their data restored.

If any volumes (RAID-5 or otherwise) were rendered unusable, the following message is also included:

To restore the contents of any volumes listed above, the volume should be started with the command:

```
vxvol -f start volumename
```

and the data restored from backup.

#### Sparelist File

The sparelist file is a text file that specifies an ordered list of disks to be used as hot spares when a specific disk fails. The system-wide sparelist file is located in /etc/vx/sparelist. Each line in the sparelist file specifies a list of spares for one disk. Lines beginning with

VxVM 3.0

Last change: 26 Mar 1999

4

#### Maintenance Commands

vxsparecheck(1M)

the pound (#) character and empty lines are ignored. The format for a line in the sparelist file is:

```
[ diskgroupB:] diskname : spare1 [ spare2 ... ]
```

The diskgroup field, if present, specifies the disk group within which the disk and designated spares reside. If not present, rootdg is presumed. The diskname specifies the disk for which spares are being designated. The spare list after the colon lists the disks to be used as hot spares. The list is order dependent; in case of failure of diskname, the spares are tried in order. A spare will be used only if it is a valid hot spare (see above). If the list is exhausted without finding any spares, the default policy of using the closest disk is used.

#### FILES

/etc/vx/sparelist	Specifies a list of disks to serve as hot spares for a disk.
-------------------	--

#### NOTES

The sparelist file is not checked in any way for correctness until a disk failure occurs. It is possible to inadvertently specify a non-existent disk or inappropriate disk or disk group. Malformed lines are also ignored.

SEE ALSO

mailx(1), vxintro(1M), vxnotify(1M), vxrecover(1M)

VxVM 3.0

Last change: 26 Mar 1999

5

===[ man1m/vxstat.1m ]=====

Maintenance Commands

vxstat(1M)

NAME

vxstat - Volume Manager statistics management utility

SYNOPSIS

vxstat [-dpsv ] [-f fields ] [-g diskgroup] [-i interval [ -c count ]]

[-n node [,node...]] [-r ] [object...]

DESCRIPTION

The vxstat utility prints and resets statistics information on one or more volumes, plexes, subdisks, or disks.

The vxstat utility reads statistics from the volume device files in the directory /dev/vx/rdisk and prints them to standard output. These statistics represent volume, plex, sub-

disk, and disk activity since boot time. If no object operands are given, then statistics from all volumes in the configuration database are reported. object can be the name of a volume, plex, subdisk, or disk.

In a Volume Manager cluster environment, vxstat gathers the requested statistics from all nodes in the cluster, and reports the aggregate totals. You can specify a node list to indicate a subset of nodes from which to gather statistics.

#### OPTIONS

-c count Stop after printing interval statistics count times.

-d Display statistics for disks on which the object specified on the command line is fully or partially located.

-f fields Select the display of statistics collected. The following options are available:

0 | 0 Statistics for the VOL\_R5\_ZERO operation. Each operation represents one call to the VOL\_R5\_ZERO ioctl. The number of blocks is based on the number of zeroed blocks written to the array. The average time is the time taken to complete the entire ioctl operation.

a Statistics on atomic copies performed (has meaning only for mirrored volumes). Displays the number of operations, number of blocks, and the average time spent per operation.

b Displays the statistics on read-

writeback mirror consistency recovery operations (has meaning only for mirrored volumes). Displays three fields: the number of read-writeback operations, the number of blocks involved in read-writeback operations, and the average time for completing a read-writeback operation. While in recovery mode, most

read operations to a mirrored volume invoke read-writeback consistency recovery.

c Statistics on corrected (fixed) read operations (has meaning only for mirrored or RAID-5 volumes). Displays the number of fixed read and write operations.

Note: Currently, only read operations are ever corrected, so the number of fixed writes will always be zero.

C Statistics for the VOL\_R5\_RECOVER operation. Each operation count represents one call to the VOL\_R5\_RECOVER ioctl. The number of blocks represents the resulting number of blocks that were written to the missing column region as part of the data recovery. The read operations are not counted towards the total.

f Displays the number of failed read and write operations.

F Statistics for full-stripe writes on a RAID-5 volume. The number of operations represents the number of write operations within a stripe that were conducted as a full-stripe write optimization. Full-stripe writes represent considerably less overhead than read-modify-writes in terms of overall I/O time, latency and CPU overhead. The total number of blocks represents the total size of the written data and the average time is the time taken for a full-stripe write operation. Since the I/O may be larger than a single stripe, more than one stripe operation may be seen for a single logical I/O request.

- M Read-modify-write statistics. Each operation represents a read-modify write operation performed within a stripe. I/O crossing a stripe boundary will be represented by more than a single read-modify write operation. The number of blocks counted represents only the size of the requested write. The read portion of the I/O can be derived.
- R Reconstruct read operations. Each operation is a separate reconstruct read operation. A single stripe read or write operation can lead to numerous reconstruct read operations since each reconstruction takes place at the sub-disk level. A detached column can consist of several subdisks each of which will lead to a reconstruct read operation.
- s Statistics on read and write operations. Displays six fields: the number of read operations, the number of write operations, the number of blocks read, the number of blocks written, the average time spent on read operations in the interval, and the average time spent on write operations in the interval. These statistics are displayed as the default output format.
- S Statistics for the VOL\_R5\_RESYNC operation. Each operation count represents one call to the VOL\_R5\_RESYNC ioctl. The number of blocks represents the resulting number of blocks that were written to the parity regions as part of the resynchronization of parity. The read operations are not counted towards the total.
- v Statistics on verified reads and writes (has meaning only for mirrored volumes). Displays six fields: the number of verified read operations, the number of verified write operations, the number of blocks read, the number of blocks written, the average time spent on verified read operations in the interval, and the average time spent on verified write operations in the interval.



- V            Statistics for the VOL\_R5\_VERIFY operation. Each operation count represents one call to the VOL\_R5\_VERIFY ioctl. The number of blocks represents the resulting number of blocks that were read as part of the RAID-5 stripes consistency verification.
- W            Reconstruct write statistics. Each operation counted is for a reconstruct write operation performed as an optimization of a write operation within a stripe. The number of blocks counted represents the count of data blocks written not including parity or read operations.
- g diskgroup  
Select records from the specified disk group. The diskgroup option argument can be either a disk group name or disk group ID.
- i interval  
Print the change in volume statistics that occurs after every interval seconds. The first interval is assumed to encompass the entire previous history of objects. Subsequent displays will show statistics with a zero value if there has been no change since the previous interval.
- n node[,node...]  
Gathers and displays statistics for the nodes specified in node, where node is a comma separated list of integer node numbers.
- p            Display statistics for plexes on the object specified on the command line. For subdisk objects, displays information about a plex with which it is associated.
- r            Reset statistics instead of printing them. This option will follow the same selection rules as printing for any type selection arguments or for any named objects. If an interval was specified on the command line, then the first set of statistics will not be printed since they will have been reset to zero. Subsequent activity will cause printing of statistics as normal.
- s            Display statistics for subdisks on the objects

specified on the command line.

Maintenance Commands

vxstat(1M)

-v            Display statistics for volumes on the objects specified on the command line. For an object that is a plex or a subdisk, displays information about the volume with which the object is associated. If an object supplied is a disk, then any volumes that occupy any part of the disk will be selected.

OUTPUT FORMAT

Summary statistics for each object are printed in one-line output records, preceded by two header lines. The output line consists of blank-separated fields for the object type, object name (standard), and the fields requested by the -f switch in the order they are specified on the command line.

If the -i interval option was supplied, then statistics will be prefaced with a time-stamp showing the current local time on the system.

EXIT CODES

The vxstat utility exits with a nonzero status if the attempted operation fails. A nonzero exit code is not a complete indicator of the problems encountered, but rather denotes the first condition that prevented further execution of the utility. See vxintro(1M) for a list of standard exit codes.

EXAMPLES

To display statistics for all subdisks associated with all volumes, use the command:

```
vxstat -s
```

To display statistics for the plexes and subdisks of a volume named blob, use the following:

```
vxstat -ps blob
```

To reset all statistics for a disk group named foodg, type the following command:

```
vxstat -g foodg -r
```

To display 5 sets of disk statistics at 10 second intervals,  
use the following:

```
vxstat -i 10 -c 5 -d
```

VxVM 3.0

Last change: 26 Mar 1999

5

Maintenance Commands

vxstat(1M)

SEE ALSO

vxintro(1M), vxtrace(1M), vxtrace(7)

=== [ man1m/vxtask.1m ] =====

Maintenance Commands

vxtask(1M)

#### NAME

vxtask - list and administer Volume Manager tasks

#### SYNOPSIS

vxtask abort taskid

vxtask [-ahlpr ] [-g dg\_name [-g dg\_name...]] [[-G dg\_name]  
-v volume [-v volume...]] [-i taskid] list [taskid...]

vxtask [-c count] [-ln ] [-t time] [-w wait\_interval] moni-  
tor [taskid[taskid...]]

vxtask pause taskid

vxtask resume taskid

vxtask [-i taskid] set name=value taskid

#### DESCRIPTION

The vxtask utility performs basic administrative operations on Volume Manager tasks that are running on the system. Operations include listing tasks (subject to filtering), modifying the state of a task (pausing, resuming, aborting)

and modifying the rate of progress of a task.

Volume Manager tasks represent long-term operations in progress on the system. Every task gives information on the time the operation started, the size and progress of the operation, and the operation's state and rate of progress (throttle). The administrator can change the state of a task, giving coarse-grained control over an operation's progress. For those operations that support throttling, the task's rate of progress can be changed, giving more fine-grained control over the task.

The majority of tasks represent I/O being performed to objects. Operations such as read-writeback recovery for mirrored volumes, parity recalculation for RAID-5 volumes, volume relayouts, etc. all involve moderate to large amounts of I/O. Tasks that represent operations that perform I/O directly are referred to as basic tasks. The task progress information for basic tasks consists of the starting and ending block of the I/O to be performed and the offset to which the I/O has currently completed.

Some operations, such as vxrecover and starting RAID-5 volumes, may require multiple tasks to complete the operation. vxrecover may determine that many recoveries are required, some that must be performed serially. Starting a RAID-5 volume requires that its log (if any) be replayed and cleared, its parity be recalculated (if necessary), any stale subdisks be resynchronized (if needed), and that these

## Maintenance Commands

vxtask(1M)

tasks be performed serially. In these and similar cases, an overall task is created to keep track of the underlying basic tasks being performed, and these are known as parent tasks. The progress information for parent tasks are the total number of subtasks required to complete the overall operation; the number of subtasks completed; and the number of subtasks currently running.

Every task is given a unique task identifier. This is a numeric identifier for the task that can be specified to the vxtask utility to specifically identify a single task.

Tasks also contain the following information:

task tag A task's tag is a string that the administrator

can specify to make administration easier. It is usually set by the command that initiates the task. For most utilities, the tag is specified with the -t tag option.

**task type** The task's type describes the specific operation that is being performed. For example, attaching a plex to a volume results in an atomic copy loop, which is represented by an ATOMIC\_COPY task.

**description**  
Describes the utility operation on whose behalf the I/O is being performed. For example, both a volume start and a plex attach can result in an atomic copy loop and thus an ATOMIC\_COPY task. The description of the operation attempts to disambiguate between tasks that are performing the same basic work.

**object** Most tasks are related to a specific object, usually the volume within which the operation is taking place.

**parent** Some utilities require multiple operations to fulfill a request. For example, the vxrecover utility may need to start many different volumes. In these cases, the subtasks list their parent.

**progress information**  
Each task has an indication of its starting point and ending point, and an indication as to what point it has progressed so far. Combined with the work time, this allows the amount of time remaining until completion to be estimated for some tasks.

**state** Each task has one of three states: RUNNING, PAUSED

or ABORTING.

**time** Each task notes its starting time and the total amount of time spent doing work.

#### KEYWORDS

The vxtask utility supports the following operations:

abort|pause|resume

These three operations request that the specified task change its state. pause will put a running task in the paused state, causing it to suspend operation. resume will cause a paused task to continue operation. abort will cause the specified task to cease operation. In most cases, the operations will "back out" in a manner similar to as if an I/O error occurred, reversing to the extent possible what they had done so far.

list

List tasks running on the system in one-line summaries. By default, all tasks running on the system are printed. If a taskid argument is supplied, the output is limited to those tasks whose taskid or task tag match taskid. The remaining options filter and limit the listed tasks as follows:

- a Limits the output to tasks in the aborting state.
- g dg\_name Limits the listed tasks to those tasks running on objects in the diskgroup dg\_name.
- G Distinguishes between volumes with the same name in different diskgroups; see vxintro(1M).
- h Prints tasks hierarchically with a task's child tasks following the parent task.
- l Prints tasks in long format.
- p Limits the output to tasks in the paused state.
- r Limits the output to tasks in the running state.
- v volume Limits the output to tasks whose object

is volume.

Any task meeting any of the criteria specified are printed. That is, if `-v foovol` and `-p` are specified, then any task operating on foovol or that is paused are printed. If `-h` were also specified, then those tasks and all their children would be listed.

`monitor` The monitor operation causes information about a task or group of tasks to be printed continuously as task information changes. This allows the administrator to track progress on an ongoing basis. Specifying `-l` causes a long listing to be printed; by default, short one-line listing are printed.

The monitor command accepts the following options:

`-c count` Causes the program to print exactly count sets of task information and then exit.

`-n` Causes the program to not only monitor the tasks specified on the command line, but to also monitor any newly registered tasks while the program is running.

`-t time` Causes the program to exit after time seconds.

`-w interval` Causes the string "waiting..." to be printed when interval seconds have passed with no output activity.

In addition to printing task information when a task's state changes, output is also generated when the task completes. When this occurs, the task's state is printed as EXITED (see the OUTPUT section).

`set` The set operation is used to change modifiable parameters of a task. Currently, there is only one modifiable parameter for tasks: the slow attribute, which represents a throttle on the task progress. The larger the slow value, the slower the progress of the task and the fewer system resources it consumes in a given time. (The slow attribute is exactly the same attribute that many commands, such as `vxplex`, `vxvol` and `vxrecover` accept on their command lines.)



## OUTPUT

There are two output formats printed by vxtask: a short, one-line summary format per task, and a long task listing. The short listing attempts to provide the most used task information for quick perusal, displaying the following fields from left to right:

1. The task ID for the task.
2. The task ID of the task's parent, if any. If the task has no parent, this field is left blank.
3. The task type and state, separated by a slash (/). The type field is a description of the work being performed, such as ATOMIC COPY for atomic copy I/O and RELAYOUT for an online relay layout operation. The state is a single letter representing the task state: R for running, P for paused, and K for aborting. In the special case of the completion of a task that is being monitored, the state is printed as EXITING.
4. The percentage of the operation that has been completed to this point.
5. The task's progress information. For basic tasks, this is the starting offset, ending offset and current offset for the operation, separated by slashes (/). For parent tasks, this is the starting number of child tasks, total number of child tasks required for the operation to complete, and the number of child tasks already completed to this point, separated by slashes, followed by the number of children currently running in parentheses.
6. A description of the command for which the task is performing work. This is normally a single word, such as START for a volume start operation or RELAYOUT for an online relay layout, followed by object names, indicating what the arguments were for the command. See EXAMPLES below.

The long output format prints all available information for a task, spanning multiple lines. If more than one task is printed, the output for different tasks is separated by a single blank line. Each line in the long output format contains a title for the line, followed by a colon (:), followed by the information. The possible information printed for a task includes:

Operation:

A description of the operation on whose behalf the work is being performed. This is essentially a

VxVM 3.0

Last change: 26 Mar 1999

5

Maintenance Commands

vxtask(1M)

description of the command executed along with the VxVM objects specified as arguments.

Progress: A description of the progress of the task. The progress is stated as a percentage of completion, along with the actual progress values for the task.

Started: The time that the task was started.

Task: General information for the task. The task's ID, followed by its parent's ID (if any) in parentheses, followed by the task's state, either RUNNING, PAUSED or ABORTING.

Throttle The throttle value for the task, if applicable.

Type: The task type, describing the work being performed.

Work time:

States how much time has been spent performing the work, and a rough estimate of the time remaining for the task to complete.

As a special case, when a task being monitored completes only the Task: line is printed for the task, with the task state represented by EXITED.

EXAMPLES

To list all tasks currently running on the system:

```
vxtask list
```

To trace all tasks in the diskgroup foodg that are currently paused, as well as any tasks with the tag sysstart:

```
vxtask -g foodg -p -i sysstart list
```

To list all tasks on the system that are currently paused:

```
vxtask -p list
```

To monitor all tasks with the tag myoperation:

```
vxtask monitor myoperation
```

VxVM 3.0

Last change: 26 Mar 1999

6

Maintenance Commands

vxtask(1M)

To cause all tasks tagged with recovall to exit:

```
vxtask abort recovall
```

#### FILES

/dev/vx/info

Used to get task information from the kernel.

/dev/vx/taskmon

Task monitoring device.

#### SEE ALSO

vxintro(1M), vxplex(1M), vxrecover(1M), vxsd(1M), vxvol(1M)

===[ man1m/vxtrace.1m ]=====

Maintenance Commands

vxtrace(1M)

#### NAME

vxtrace - trace operations on volumes

#### SYNOPSIS

```
vxtrace [-aeE1 ] [-b buffersize] [-c eventcount]
        [-d outputfile] [-f inputfile] [-g diskgroup]
        [-o objtype [,objtype]...] [-t timeout]
        [-w waitinterval] [name | device]...
```

#### DESCRIPTION

The vxtrace utility prints kernel error or I/O trace event records on the standard output or writes them to a file in binary format. Binary trace records written to a file can be read back and formatted by vxtrace as well.

If no operands are given, then either all error trace data or all I/O trace data on all virtual disk devices are reported. With error trace data, it is possible to select all accumulated error trace data, to wait for new error trace data, or both (the default). Selection can be limited to a specific disk group, to specific Volume Manager kernel I/O object types, or to particular named objects or devices.

## OPTIONS

- a Append to the outputfile rather than truncating it. By default, the output file is truncated.
  
- b buffersize  
Set the kernel I/O trace buffer size or set the read buffer size when used with the -f option. The Volume Manager kernel allocates a private kernel space to buffer the I/O trace records for each vxtrace command. The default buffer size is 8K bytes. Some trace records may be discarded if the trace buffer is too small. This option can be used to set a larger or a smaller kernel trace buffer size. The buffer size is specified as a standard Volume Manager length (see vxintro(1M)). Depending on the Volume Manager kernel configuration, usually only a maximum of 64K bytes buffer size will be granted.
  
- c eventcount  
Accumulate at most eventcount events and then exit. The timeout and eventcount options can be used together.
  
- d outputfile  
Write (dump) binary trace data to the specified output file.
  
- e Select new error trace data. The default is to

## Maintenance Commands

vxtrace(1M)

- select I/O trace data.
  
- E Select pre-existing error trace data. This can be combined with -e to get both pre-existing trace data and new trace data.
  
- f inputfile  
Read binary trace data from the specified input file, rather than from the Volume Manager kernel.
  
- g diskgroup  
Select objects from the specified disk group. The disk group can be specified either by disk group ID or by disk group name. With no name or device

operands, all appropriate objects in the disk group are selected. With name operands, diskgroup names the disk group that is expected to contain the named configuration record.

-l Long format. Print all available fields for all tracing records, rather than a subset of the available fields. The default is to use the short format.

name | device  
If name or device operands are provided, Volume Manager kernel objects of the requested types are selected if they are associated with the configuration records or virtual disk devices indicated by those operands.

-o objtype [,objtype]...  
Select object based on the objtype option arguments. Multiple types of objects can be specified with one or several -o options. The possible object selection types are:

all | ALL Select all possible virtual disk devices, kernel objects, and physical disks.

dev | logical  
Select virtual disk devices.

disk | physical  
Select VxVM physical disks.

log Select all log objects.

logplex Select RAID-5 log plexes.

logsd Select DRL or RAID-5 log subdisks.

logvol Select DRL or RAID-5 log volumes.

m | mv | mirror  
Select mirrored volume kernel objects.

p | pl | plex

Select striped or concatenated plex kernel objects.

s | sd | subdisk  
Select subdisk kernel objects.

v | vol | volume  
Select mirrored or RAID-5 volume kernel objects.

-t timeout  
Accumulate trace data for at most timeout seconds and then exit.

-w waitinterval  
If vxtrace waits for waitinterval seconds without receiving any new events, print waiting... to allow scripts to wake up and process previously accumulated events. This is mostly of use for processing errors. The waiting... message does not count as an event for the purposes of the -c option.

#### OPERANDS

Operands specify configuration record names, or physical or virtual disk device nodes (by device path). If no object types were selected with the -o option, then only trace records corresponding to the indicated configuration records or devices are selected; otherwise, objects of the requested types are selected if they are associated in any way with the named configuration record or device.

If a name argument does not match a regular configuration record but does match a disk access record, then the indicated physical disk is selected. Physical disks can also be selected by the device path of the public region partition, or by the disk media record name.

By default, name arguments are searched for in all disk groups or in the disk group specified using the -g option. Without the -g option, a record that is found in more than one disk group will cause an error unless the record is in the rootdg disk group (in which case, the record in the rootdg disk group is selected). The disk group for any individual name operand can be overridden using the form:

diskgroup/recordname

Note: When reading trace data from a file with the `-f` option, association information is not available.

#### EXAMPLES

To trace all physical disk I/Os, enter:

```
vxtrace -o disk
```

To trace virtual disk device I/Os to the device associated with volume testvol, use either of the commands:

```
vxtrace -o dev testvol  
vxtrace /dev/vx/dsk/testvol
```

To trace all log subdisks associated with volume testvol, enter:

```
vxtrace -o logsd testvol
```

To trace all log objects, enter:

```
vxtrace -o log
```

To accumulate ten seconds worth of trace data for disk04 and then format that data, use:

```
vxtrace -t 10 -d /tmp/tracedata disk04  
vxtrace -l -f /tmp/tracedata
```

To read error trace data into a script for processing, using ten second pauses to generate mail messages, use the command:

```
vxtrace -leE -w 10 | while read ...
```

#### FILES

/dev/vx/trace

#### SEE ALSO

vxintro(1M), vxstat(1M), vxtrace(7)



===[ man1m/vxunroot.1m ]=====

Maintenance Commands

vxunroot(1M)

NAME

vxunroot - remove Volume Manager hooks for rootable volumes

SYNOPSIS

/etc/vx/bin/vxunroot

DESCRIPTION

The vxunroot utility converts the root, swap, usr, and var file systems back to being accessible directly through disk partitions instead of through volume devices. Other changes made to ensure the booting of the system from the root volume are also removed so that the system boots with no dependency on the Volume Manager.

For vxunroot to work properly, all but one plex of rootvol, swapvol, usr, and var must be removed. The plexes left behind for the above volumes must be the ones created by vxrootmir or the original ones created when the root disk was encapsulated. This ensures that the underlying subdisks have equivalent partitions defined for them on the disk. If none of these conditions is met, the vxunroot operation fails and none of the volumes is converted to disk partitions.

The unwanted plexes can be removed using either vxedit or vxplex.

SEE ALSO

vxdiskadm(1M), vxedit(1M), vxintro(1M), vxplex(1M), vxrootmir(1M), vxunroot(1M)

===[ man1m/vxvol.1m ]=====

## Maintenance Commands

vxvol(1M)

## NAME

vxvol - perform Volume Manager operations on volumes

## SYNOPSIS

vxvol [-fqV] [-g diskgroup ] [-U usetype] [-o useopt] [-t tasktag] init init\_type volume [arg ...]

vxvol [-fqV] [-g diskgroup ] [-U usetype] [-o useopt] [-t tasktag] maint volume ...

vxvol [-fqV] [-g diskgroup ] [-U usetype] [-o useopt] [-t tasktag] noderecover volume ...

vxvol [-fqV] [-g diskgroup ] [-U usetype] [-o useopt] [-t tasktag] rdpol policy volume [plex]

vxvol [-fqV] [-g diskgroup ] [-U usetype] [-o useopt] recover volume [subdisk] ...

vxvol [-fqV] [-g diskgroup ] [-U usetype] [-o useopt] [-t tasktag] resync volume ...

vxvol [-fqV] [-g diskgroup] [-U usetype] [-o useopt] [-t tasktag] set attribute=value ... [ -- ] volume ...

vxvol [-fqV] [-g diskgroup ] [-U usetype] [-o useopt] [-t tasktag] start volume ...

vxvol [-fqV] [-g diskgroup ] [-U usetype] [-o useopt] [-t tasktag] startall

vxvol [-fqV] [-g diskgroup ] [-U usetype] [-o useopt] [-t tasktag] stop volume ...

```
vxvol [-fqV] [-g diskgroup ] [-U usetype] [-o useopt] [-t
tasktag] stopall
```

#### DESCRIPTION

The vxvol utility performs Volume Manager operations on volumes. The first operand is a keyword that determines the specific operation to perform. The remaining operands specify configuration records to which the operation is to be applied.

Each operation can be applied to only one disk group at a time, due to internal implementation constraints. Any volume operands will be used to determine a default disk group, according to the standard disk group selection rules described in vxintro(1M). A specific disk group can be selected with -g diskgroup.

VxVM 3.0

Last change: 26 Mar 1999

1

Maintenance Commands

vxvol(1M)

#### KEYWORDS

init Perform an initialization action on a volume. This can be applied to volumes that were created by vxmake and that have not yet been initialized, or volumes that have been set to the uninitialized state with vxmend fix empty. The action to perform is specified by the init\_type operand, which is usage-type-dependent. The volume operand determines which usage type to use for performing the operation.

maint Detach each volume named by the volume operands. When a volume is detached, normal read and write operations to the volume fail, although most volume ioctl operations can still be used.

noderecover When a node in a cluster aborts or crashes, it is not allowed to rejoin the cluster until the active logs of the volumes affected by the crash are recovered. The noderecover operation examines all volumes named by the volume operands and recovers the active log maps corresponding to the crashed nodes, but does not initiate mirror resynchronization. Mirrors are subsequently resynchronized by a vxvol -o force resync volume... command.

This operation is currently applicable only to the volumes with DRL logs and is relevant only in clustered systems.

- rdpol Set the read policy for a volume based on the policy operand. These are the recognized read policies:
- prefer Read preferentially from the plex named by the plex operand. If the plex is enabled, readable, and associated with the volume, then any read operation on the volume results in a read from that plex if all blocks requested in the read are contained in the plex. The plex operand is required for the prefer read-policy type.
  - round Use a round-robin read order among the enabled, readable plexes associated with the volume. No plex operand should be specified for the round read-policy type.
  - select Select a default policy based on plex

## Maintenance Commands

vxvol(1M)

- associations to the volume. For a volume that contains one enabled, striped plex, the default is to prefer that plex. For any other set of plex associations, the default is to use a round-robin policy. No plex operand should be specified for the select read-policy type.
- recover Some usage types support the concept of recovery of data for objects within a volume different from the mirror resynchronization model. The recover operation is used to specify this type of recovery. The exact procedure for this operation is usage type specific.
- resync Examine all volumes named by the volume operands and perform any synchronization operations that

are required. The exact procedure for this operation is usage-type specific.

**set** Change specific volume characteristics. The changes to be made are given by arguments immediately after the **set** keyword of the form **attribute=value**. The set of volumes affected by the operation are given after these operands; thus the attribute list ends with an operand that does not contain an equal sign. To allow for volume names that contain an equal sign, an operand of **--** can be used to terminate the attribute list. Each usage type represented by the list of volume operands is called once, with the set of all volumes with that usage type.

The set of **attribute=value** attribute arguments that are recognized depends upon the volume usage type (see the sections on usage types later in this manual page for information on available attributes). However, an attribute argument of the form **len=number** is expected to be interpreted (if at all) as requesting a change in the length of a volume, regardless of the volume's usage type. The number value is interpreted as a standard length number (see **vxintro(1M)**).

**start** Enable disabled or detached volumes named by the volume operands. The process of enabling a volume is a highly usage-type-dependent operation and may result in transfers of data between plexes associated with the volume.

If the **start** operation is applied to an

uninitialized volume (for example, a volume just created by **vxmake**), a default initialization will be used to initialize and enable the volume.

If the volume is not normally started because failures and disk removals have left all associated plexes with invalid data, the **-f** option can be used to try to start the volume, anyway. This can be used after replacing disks to enable the volume so that its contents can be restored from

backup or reinitialized.

`startall` Attempt to start all volumes that are disabled. If a `-U usetype` option is specified, then attempt to start all disabled volumes with the indicated usage type. This operation will not start uninitialized volumes. By default, start all volumes in the `rootdg` disk group. A different disk group can be specified with the `-g` option.

`stop` Disable the enabled or detached volumes named by the volume operands.

The `stop` operation provides an interface to the usage type of a volume for shutting down operations on a volume in a clean manner. The specific method for cleanly stopping a volume, and the precise meaning of `clean` are both highly usage-type-dependent. By convention, `-f` can be used to force stopping of a volume that is in use, forcing I/O failures to be returned for any further volume device operations.

`stopall` Attempt to stop all volumes that are enabled. If a `-o usetype` option is specified, then attempt to stop all disabled volumes with the indicated usage type. By default, stop all volumes in the `rootdg` disk group. A different disk group can be specified with the `-g` option.

#### OPTIONS

`-f` Force an operation in some situations where the operation has questionable semantics. For example, `-f` may be used to reduce the length of a volume with `vxvol set`, to stop a volume that is currently open or mounted as a file system, or to attempt to start a volume that has no plexes with valid data.

`-g diskgroup` Specify the disk group for the operation. The disk group can be specified either by name or by

disk group ID. See `vx dg(1M)` for more information on disk groups.

-o useopt Pass in usage-type-specific options to the operation. By convention, the following usage-type-specific options should be implemented by all usage types:

bg Perform any extended revive operations in background processes after the volume and one or more plexes have been enabled. A volume that is started or whose length is changed successfully with this option will be usable immediately after the operation completes, although recovery operations may affect performance of the volume for an extended period of time.

delayrecover

Do not perform any plex revive operations when starting a volume. Simply enable the volume and any plexes. This may leave some stale plexes, and may leave a mirrored volume in a special read-writeback (NEEDSYNC) recover state that performs limited plex recovery for each read to the volume.

iosize=size

Perform recovery operations in regions with the length specified by size, which is a standard Volume Manager length number (see vxintro(1M)). Specifying a larger number typically causes the operation to complete sooner, but with greater impact on other processes using the volume. The default I/O size is typically 32 kilobytes.

plexfork[=count]

Perform up to the specified number of plex revive operations simultaneously. If no count is specified, then a suitable small number is used (normally 10).

slow[=iodelay]

Reduce the system performance impact of plex recovery operations and volume length changes. Startup recovery and length change consistency operations are usually a set of short operations on

small regions of the volume (normally from 16K bytes to 128K bytes). This option inserts a delay between the recovery of each such region. A specific delay can be specified with `iodelay` as a number of milliseconds, or else a default is chosen (normally 250 milliseconds).

- `verbose` Print a message for each volume that is successfully started. Without this option, messages appear only for volumes that fail to start.
- `-q` Wait for the volume daemon if it is not running or available for transactions. If `-q` is not specified, `vxvol` aborts if the volume daemon is not available.
- `-t tasktag` If any tasks are registered to track the progress of the operation, mark them with the tag `tasktag`.
- `-U usetype` Force the operation to be performed by the usage-type utility for this usage type.
- `-V` Display a list of utilities that would be called from `vxvol`, along with the arguments that would be passed. The `-V` option performs a ``mock run'' so the utilities are not actually called, and no changes are made to the volume configuration database.

#### FSGEN and GEN Usage Types

The `fsgen` and `gen` usage types provide the same semantics for all operations of the `vxvol` utility. However, some options are provided only by the `gen` usage type.

In addition to the standard `-o` options required for all usage types, the `fsgen` and `gen` usage types provide the following additional options:

- `force` Force an operation that is not normally performed as part of the operational model of the Volume Manager and may have adverse effects on data. This is the same as `-f`.
- `norecov` This can only be used with the `gen` usage type. Prevent the start operation from recovering plexes through the `vxplex` utility. Instead, all STALE and ACTIVE plexes are simply treated as equivalent



## Maintenance Commands

vxvol(1M)

to CLEAN plexes, and are thus enabled without being made consistent. This can be used for volumes whose contents are recreated for each use.

An example of a possible use for this attribute is a swap area and the /tmp file system. In the case of /tmp, the model assumes that mkfs is used to create an empty file system after the volume has been started.

Limitations and extensions for the fsgen and gen usage types consist of the following:

init        These are the recognized uses of the vxvol init operation:

vxvol init active volume

Set the state for all plexes associated with volume to ACTIVE and enable the volume and its plexes. This is used to initialize a single or multiple-plex volume where all plexes are known to have identical contents.

vxvol init clean volume [plex]

Set the state for the specified plex to CLEAN, and set all other plexes to STALE. The vxvol start operation can then be used to recover the volume from the CLEAN plex. This operation requires that the volume not be enabled.

If the specified volume has only one plex, then the plex argument is not required as it defaults to that plex. If specified, then the plex argument must represent a plex that is associated with the volume.

vxvol init enable volume

Enable the volume and its plexes but leave the volume uninitialized. This operation can be used only for non-enabled volumes. It is used to tem-

porarily enable a volume so that data can be loaded onto it to make it consistent. Once the data has been loaded, `init active` should be used to fully enable the volume. `init active` could be used, for example, if a complete image of the volume is to be loaded from a tape.

## Maintenance Commands

`vxvol(1M)`

### `vxvol init zero volume`

Write zero blocks to all plexes in the volume, up to the length of the volume. After the writes complete, the state of each plex is set to `ACTIVE` and the volume and its plexes are enabled. `init zero volume` could be used, for example, before running `mkfs` to put a file system on the volume.

If this operation is interrupted by a signal, then an attempt is made to restore all affected records to their original state, or to a state that is roughly equivalent to their original state. If this attempt is interrupted, such as through another signal, then the user may need to perform some cleanup. A set of commands to perform this cleanup are written to the standard error before the volume utility exits.

`maint` The `-f` option is required to detach an enabled volume. Also, a warning is written to the standard error for volumes that are open or mounted.

`resync` Volumes that have possibly differing plex contents will be re-synchronized to contain consistent data. Any such volumes that are in the `NEEDSYNC` state will be recovered using a read/write-back recovery mode and then put into the `ACTIVE` state.

Plexes in the `SYNC` state may already be under recovery and the volume command will take no action to recover them unless the command was invoked with the `-o force` option.

set           The attributes that can be changed are:

exclusive=yes|y|on|true|no|n|off|false

Sets or clears the EXCLUSIVE flag on the volume. A volume in exclusive open state can be opened by only one node in the cluster at a time. Multiple opens of an exclusive volume from the same node are permitted. Non-exclusive volumes can be simultaneously opened by more than one node. After a node opens an exclusive volume, every other node's open attempt fails until the last close of the volume by the first opener. Such an open failure returns a EBUSY error

VxVM 3.0

Last change: 26 Mar 1999

8

Maintenance Commands

vxvol(1M)

code. Available only if the Volume Manager cluster feature is enabled.

len=number

Change the length of each volume specified by the volume operands to number sectors. number is a standard Volume Manager length number (see vxintro(1M)). Decreasing the length of a volume requires use of -f.

If the volume is enabled, then count the number of enabled, read-write plexes that would remain complete after the length change. The operation fails if this number would become zero, but the number of sparse plexes would become greater than 1. Changing the length of a volume with one enabled plex beyond the length of the plex requires use of the -f option.

If the volume is not enabled, count the number of CLEAN and ACTIVE plexes that would remain complete after the length change, then use the algorithm mentioned previously for determining whether the operation is allowed or requires use of

-f.

To ensure that the new region of the volume is consistent across all plexes of the volume, the volume is put into a SYNC state and read/write-back mode, and a read loop is then performed against the volume. Once this loop has completed, the volume is put back into the ACTIVE state.

loglen=size

Set the size for logs used with the volume. The size value is a standard Volume Manager length numbers (see vxintro(1M)).

logtype=type

Set the type of logging to be used on the volume. This change can be applied only to volumes that are stopped and that have no ACTIVE plexes. Allowed log types are drl (logs the regions involved in all volume writes), none (never does

logging), and undef (never does logging). If the logging type is set to undef, then a future vxsd aslog or vxplex att operation will change it to drl. See the fsgen and gen sections of vxsd(1M) and vxplex(1M) for more information.

startopts=volume\_options

Set options that are applied to the volume every time the volume is started, independently of options specified with the volume start command. This is a set of comma-separated options of the same form used with the -o option letter. At the present time, only the noattach and verbose options can be applied to volumes in this manner. Unrecognized or inappropriate options are ignored.

start Starting an uninitialized gen or fsgen volume enables the volume and its plexes, sets the plexes to the ACTIVE state, and recovers the plexes to ensure that each plex has the same contents. If the volume has only one plex, then the volume is immediately set to the ACTIVE state; otherwise, the volume is set to the SYNC state and a special read/write-back mode is set to recover regions of the volume on every read operation. The volume is then read from beginning to end to make all plexes consistent, then the volume is set to the ACTIVE state.

Starting a volume with no active dirty region logging involves enabling all CLEAN and ACTIVE plexes and putting them in the ACTIVE state. If an I/O failure was logged against the plex, or if a disk replacement caused a plex to become stale, then the plex is considered STALE. If any of the subdisks for the plex reside on a removed or inaccessible disk, then the plex is ignored for the purposes of starting the volume.

If two or more plexes were enabled, and if the volume was active at the time the system went down, then the state for the volume is set to SYNC and a special read/write-back recovery mode is used to recover consistency of the volume, segment-by-segment, on every read. A process is then started (in the background with the -o bg option) to recover consistency for the entire length of the volume.

If any plexes were considered STALE, then those plexes are attached by calling vxplex att. The number of concurrent plex attach operations are limited based on the rules for -o plexfork.

Recovery of plexes with a dirty region log uses the same rules as for volumes without a valid dirty region log, except that recovery of non-stale plexes is done by scanning the contents of the dirty region log and recovering consistency for those regions listed in the log as requiring recovery.

In addition to enabling the volume and managing the recovery of plex consistency, starting a volume clears any transient operations that were being applied to a volume before the system was rebooted. Starting a volume dissociates and removes temporary plexes or subdisks, and dissociates plexes that were being attached if the attach operation did not complete. Snapshot plexes created by vxassist are also removed.

If the volume is unstartable because there are no valid, non-stale plexes and the -f flag is then specified, all STALE plexes that do not contain unusable subdisks (subdisks on failed or removed disks) will be changed to ACTIVE. The volume will then be started and synchronized from those plexes.

stop            Stopping an fsgen or gen volume disables the volume and its associated plexes. In addition, the utility state for each ACTIVE plex is changed as follows:

- o If the plex is detached or disabled, set the state for the plex to STALE. If all plexes are set to STALE, then the volume cannot be started until vxmend is used to change the state of one or more plexes to CLEAN or ACTIVE. A plex normally becomes detached as a result of an I/O error on the plex, or a disk failure or replacement. I/O failures will not normally detach the last remaining enabled plex in a volume, so disk removal operations are the only normal operational method of making a volume unstartable.
- o If the plex is volatile, that is, one of the subdisks in the plex is defined on a disk with the volatile attribute (see vxdisk(1M)), then

set the plex state to STALE.

- o If the volume is enabled and the plex is also enabled, then set the plex state to CLEAN.

- o If the volume is detached and the plex is enabled, then the plex state is left as ACTIVE. A volume can be left detached, with remaining valid plexes, only as a result of calling vxvol maint to detach an enabled volume.

Normally, the stop operation fails if any extended operations are using the volume or any of its associated plexes. Such operations are detected as a nonempty value for the tutil0 field in a volume or plex record. If the -f option is specified, then the stop operation ignores volume and plex tutil0 fields.

The -f option must also be given to force the stopping of a volume that is open or mounted as a file system. In this case, a warning message is still written to the standard error, but the stop operation is not otherwise affected. Stopping an open or mounted volume is not normally advisable.

#### RAID-5 Usage Type

In addition to the standard -o options required for all usage types, the raid5 usage type provides the following options:

##### checkpt=size

Set the checkpoint size for a volume. A complete resynchronization of a volume via VOL\_R5\_RESYNC ioctls can take an extended amount of time. It is conceivable in some circumstances that the operation could be stopped before it complete (such as by a system crash). To avoid having to restart the synchronization from the beginning of the volume (after a certain amount of the volume has been synchronized), a transaction is issued to record the offset to which the resynchronization has completed. This size is called the checkpoint length and can be set using the checkpt option. The default checkpoint length is 64 megabytes.

##### delayrecover

This prevents the start operation from undergoing some recovery operations. RAID-5 logs, if any are valid, will still be replayed; however, no parity resynchronizations or subdisk recoveries will be performed.

`force` Force an operation that is not normally performed as part of the operational model of the Volume Manager. It may have adverse effects on the data contained by the volume. This is the same as `-f`.

`syncstartok` This allows the `delayrecover` option to be ignored if the volume must undergo parity resynchronizations or subdisk recoveries before the volume can be enabled.

`unsafe` This allows access to certain volumes earlier in the starting process than is normally allowed by the operating process of RAID-5 volumes. This can have adverse effects on the data, and can also result in the RAID-5 volume becoming unusable after a system crash or a power failure.

Limitations and extensions for the `raid5` usage type consist of the following:

`init` The `raid5` usage type recognizes the following uses of the `init` operation:

`vxvol init active volume`  
Zero the RAID-5 log plexes, if any, and make the volume available for use. The parity in the volume is marked as stale, though no parity resynchronization is performed; the volume is left with a state of `NEEDSYNC`.

`vxvol init zero volume`  
Write zeros to the RAID-5 log plexes, if any, and write zeros to the entire length of the volume. This is achieved by issuing the `VOL_R5_ZERO` ioctl for the entire altitude of the volume. The volume is left in the `ACTIVE` state.

`recover` The `raid5` usage type supports the following invocations of the `recover` operation:

`recover [subdisk ...]`  
The `recover` operation can be used to initiate a recovery of subdisks containing invalid data. If subdisks are specified and are stale, they are recovered in the order specified. This is done by setting the stale and write-only flags on the subdisks and issuing `VOL_R5_RECOVER` ioctls to regenerate the



## Maintenance Commands

vxvol(1M)

data. After a successful recovery, the subdisk is marked as non-stale and read-write.

If no subdisk arguments are specified, the subdisks of the RAID-5 plex of the volume are checked to see if they are stale or have invalid contents. If any are found, they are recovered as described above.

**resync** The resync operation examines the named volumes to see if they are enabled and if the parity in any part of a volume is stale; this is normally indicated by a volume state of NEEDSYNC. If so, the volume is placed in the SYNC state and VOL\_R5\_RESYNC ioctls are issued to resynchronize the parity in those regions. Upon completion, the volume is placed in the ACTIVE state.

**set** The attributes that can be set for raid5 volumes are:

**len=number**

Change the length of the volumes specified to be number sectors. number is a standard Volume Manager length specification (see vxintro(1M)). Decreasing the length of a volume requires -f.

The volume length cannot be increased such that the RAID-5 plex is sparse in respect to the new volume length; this would make the volume unusable.

In order to assure that the new region of the volume is consistent, the new region of the volume (from the old length to the new length) is filled with zeros by issuing VOL\_R5\_ZERO ioctls before the length is reset.

**loglen=size**

Set the size of the RAID-5 log for the volume. This cannot be set if the volume has no logs. If the length is being

increased, the operation will not be allowed if it would cause any of the RAID-5 log plexes to become sparse in respect to the new length.

startopts=volume\_options

VxVM 3.0

Last change: 26 Mar 1999

14

Maintenance Commands

vxvol(1M)

Set options that are applied to the volume every time the volume is started, independently of options specified with the volume start command. This is a set of comma-separated options of the same form used with the -o option letter. Unrecognized or inappropriate options are ignored.

start Starting an uninitialized volume (one with a state of EMPTY) zeros any RAID-5 log plexes and then resynchronizes the parity of the volume by issuing VOL\_R5\_RESYNC ioctls. All subdisks are marked as non-stale and read-write. The volume and RAID-5 plex are then enabled and marked as ACTIVE, and all valid RAID-5 log plexes are marked as LOG. If any RAID-5 log plex proves to be invalid (such as having its NODAREC flag set) its state is set to BADLOG.

Starting a volume that has been shut down cleanly or is not marked as dirty enables the RAID-5 plex and RAID-5 log plexes, and sets the volume kernel state to detached to zero the RAID-5 log plexes for the volume, if any. Once this is completed, all valid RAID-5 log plexes are set to LOG and the volume is enabled and put in the ACTIVE state.

Starting a volume that was not shut down cleanly requires that the parity be resynchronized. If the volume has valid RAID-5 log plexes, the volume is first detached and has its state set to REPLAY, and all log plexes and the RAID-5 plex are enabled. If there are any valid RAID-5 log plexes, their contents are read and their data is written to the appropriate regions of the RAID-5 plex. If reading the RAID-5 logs fails, the logs are marked as invalid and the parity is resynchronized as if

there were no logs. Once the replay is complete, the RAID-5 logs are enabled and the volume is enabled and its state is set to ACTIVE.

If the volume needs resynchronization and no valid log plexes exist, the parity must be fully resynchronized. The volume is enabled and its state is set to RESYNC, and the RAID-5 plex is enabled. If usable RAID-5 plexes are available, but contain invalid data, they are zeroed. The parity is then resynchronized by issuing VOL\_R5\_RESYNC ioctls for the entire length of the volume. Once this is completed, the volume's state is set to ACTIVE. Any usable RAID-5 logs are enabled and set to the LOG

#### Maintenance Commands

vxvol(1M)

state.

If a volume requires full resynchronization (that is, has no usable logs) and the RAID-5 plex has stale or unusable subdisks, the volume is unusable and the start operation will fail. This can be overridden by using the -f flag or the -o force option. In this case, any stale subdisks are marked as non-stale and a full resynchronization is performed; however, this may result in some invalid data being introduced into the volume. If multiple subdisks at the same altitude in the RAID-5 plex are unusable (such as because they have their NODEVICE flag set), the volume is unusable and cannot be overridden.

Once any parity resynchronization has been completed, any subdisks still marked as stale are recovered. This is done by marking the subdisk as stale and write-only and issuing VOL\_R5\_RECOVER ioctls to regenerate the data on the stale subdisks. The subdisk is then marked as non-stale and read-write.

If the -o delayrecover option is specified, the only recoveries that will be performed are log replays. If the volume requires a parity resynchronization, it is enabled and left in the NEEDSYNC operation, and its parity is marked as stale. Any subdisk recoveries that are needed are

not performed, and the stale subdisks are marked as stale.

Normally, if a volume has no RAID-5 logs, it will not be enabled with a stale subdisk or an unusable subdisk because were the system to crash or a power failure occurred while the volume was in use, the parity could become stale and the volume would be unusable. This behavior can be overridden by specifying the `-o unsafe` option, which will cause the volume to be enabled during the above situations. As the name suggests, this is considered unsafe because doing so could cause data loss.

If only the `-o delayrecover` option is specified to start a volume with a stale subdisk or an unusable subdisk, the start operation will fail. In this case, the `delayrecover` option can be ignored by also specifying the `-o syncstartok` option.

`stop` Stopping a raid5 volume disables the volume and

## Maintenance Commands

`vxvol(1M)`

its associated plexes. If the volume is in the SYNC state, it is changed to the NEEDSYNC state so that recovery will be performed at the next start. Any invalid or detached RAID-5 logs are set to the BADLOG state so that they will not be used during the next start.

Normally, the `stop` operation will fail if any extended operations are using the volume or any of its plexes. Such operations are detected as a non-empty value for the `tutil0` field in a volume or plex record. If the `-f` option is specified, then the `stop` operation ignores volume and plex `tutil0` fields.

## FILES

`/etc/vx/type/usetype/vxvol` The utility that performs volume operations for a particular volume usage type.

`/dev/vx/dsk/group/volume` The device node that can be used for mounting a file sys-

tem created on the volume named volume in the disk group named group. Volumes in group rootdg are also directly under the /dev/vx/dsk directory.

/dev/vx/rdsk/group/volume

The device node that can be used for issuing raw I/O requests and also for issuing ioctl requests to the volume named volume in disk group named group. Volumes in group rootdg are also directly under the /dev/vx/rdsk directory.

#### EXIT CODES

The vxvol utility exits with a nonzero status if the attempted operation fails. A nonzero exit code is not a complete indicator of the problems encountered, but rather denotes the first condition that prevented further execution of the utility.

See vxintro(1M) for a list of standard exit codes.

#### SEE ALSO

vxassist(1M), vxdg(1M), vxinfo(1M), vxintro(1M), vxmend(1M), vxplex(1M), vxrecover(1M)

VxVM 3.0

Last change: 26 Mar 1999

17

===[ man3x/libvxvm3.3x ]=====

Miscellaneous Library Functions

libvxvm3(3X)

#### NAME

libvxvm3 - Volume Manager Library for SunCluster

#### SYNOPSIS

```
cc [ options ] files ... -lvxvm3 [ library ...]
```

```
#include <libvxvm3.h>
```

```
typedef vx_u32_t      vm_major_t;
typedef vx_u32_t      vm_minor_t;
typedef char          vm_name_t[];
```

```

typedef char          vm_path_t[];

int libvxvm_get_version(int version );

int libvxvm_get_conf(int parameter );

int libvxvm_get_dgs(int len, vm_name_t namep []);

int libvxvm_get_dginfo(vm_name_t dgname,vm_name_t dgid,
    vm_major_t *majorp, vm_minor_t *minorp,
    vx_u32_t *maxvolp, vx_u32_t *nvolp);

int libvxvm_get_disks(vm_name_t dgname, intlen, vm_path_t
pathp[]);

int libvxvm_get_volumes(vm_name_t dgname, int len, vm_name_t
namep[]);

int libvxvm_get_volinfo(vm_name_t dgname, vm_name_t volname,
    struct stat *statp);

int libvxvm_import_dg(vm_name_t dgname, vm_name_t newname,
    vx_u32_t flags, vm_minor_t newminor);

int libvxvm_deport_dg(vm_name_t dgname, vm_name_t newname,
    vm_name_t newhostid);

```

AVAILABILITY  
 VRTSvmdev

#### DESCRIPTION

The VERITAS Volume Manager libvxvm library provides a private interface for Sun Microsystem's SunCluster product. Using libvxvm functions, you can obtain generic VM information, such as the names of imported diskgroups, the disks belonging to a specific diskgroup, and diskgroup volumes and volume attributes (for example, owner, group, and mode). While most of this information is available using libvxvm, using libvxvm requires that you have some understanding of Volume Manager data structures. And because libvxvm is a static library, it requires recompiling every time the library changes. libvxvm uses some of the modules of libvxvm, but provides an interface that hides the VM-

VxVM 3.0

Last change: 05 Dec 1999

1

Miscellaneous Library Functions

libvxvm(3X)

specific data structures. And libvxvm is a shared

library, so applications remain binary compatible across releases.

## FUNCTIONS

### libvxvm\_get\_version

Call `libvxvm_get_version` before any other function in the `libvxvm` library. Arguments to `libvxvm_get_version` are the version number the application is compiled with. The return value from `libvxvm_get_version` is the version number the library supports, or -1 if the library and application versions are incompatible. The version number is made up of a major and minor component. For an application to be compatible with the `libvxvm` library, its major version (value of `LIBVXVM_MAJOR_VERSION`) must be same as that of the library. The macros `LIBVXVM_GET_MAJOR_VERSION` and `LIBVXVM_GET_MINOR_VERSION` return the major and minor component of the version. The macro `LIBVXVM_MAKE_VERSION` combines the major and minor version in a format that can be passed as argument to `libvxvm_get_version()`.

### libvxvm\_get\_conf

`libvxvm_get_conf` returns the value of various configuration parameters, or -1 (with `errno` set to `EINVAL`) if an invalid parameter was passed. The supported values for parameter are:

`LIBVXVM_NAME_SIZE` Returns the size of `vm_name_t`

`LIBVXVM_PATH_SIZE` Returns the size of `vm_path_t`.

### libvxvm\_get\_dgs

Returns the total number of imported disk groups (including `rootdg` and shared disk groups, but excluding disabled disk groups). If the parameter `namep` is not `NULL`, `libvxvm_get_dgs` also stores the name of imported disk groups (at the most, this is `len` entries) in the buffer pointed to by `namep` (which must be at least `len * libvxvm_get_conf(LIBVXVM_NAME_SIZE)` bytes).

### libvxvm\_get\_dginfo

Stores information (such as the base minor number and total number of volumes) about a specified disk group if the argument passed is a not `NULL` address. Due to resource constraints, the actual number of volumes created in a disk group can be substantially lower than the maximum volumes returned by this function. `libvxvm_get_dginfo`

returns zero upon successful completion, or else -1 and sets the appropriate errno (EINVAL if dgroup is NULL, ENOENT if the specified disk group is not imported).

#### libvxvm\_get\_disks

Returns the total number of disks belonging to the specified disk group. If the parameter pathp is not NULL, it also stores the name of disk access path for the disks (at most len entries) in the buffer pointed to by pathp (which must be at least len \* libvxvm\_get\_conf(LIBVXVM\_PATH\_SIZE) bytes). The path returned corresponds to the raw device in which the public and private regions are located (typically /dev/rdisk/cctxdxs2 or /dev/vx/rdmp/cctxdxs2).

#### libvxvm\_get\_volumes

Returns the total number of volumes in the specified disk group. If the parameter namep is not NULL, it also stores the name of volumes (at most len entries) in the buffer pointed to by namep (which must be at least len \* libvxvm\_get\_conf(LIBVXVM\_NAME\_SIZE) bytes).

#### libvxvm\_get\_volinfo

Returns zero upon successful completion, or else -1 and sets the appropriate errno (EINVAL if dgroup volname is NULL, ENOENT if the specified disk group is not imported, ENODEV if the specified volume does not exist. If the statp argument is not NULL, libvxvm\_get\_volinfo stores information pertaining to the volume in the st\_uid, st\_gid, st\_mode, st\_rdev, st\_size fields of the stat structure.

#### libvxvm\_import\_dg

Imports the disk group dgroup using the vxvg command (see vxvg(1M)). If newname is not NULL, the disk group is renamed to newname during import. If newminor is a valid minor number and the IMPORT\_OPT\_REMINOR flag is set, libvxvm\_import\_dg performs a vxvg remminor. The supported flags are:

IMPORT\_OPT\_TEMPNAME Setting this flag disables the auto-import flag.

IMPORT\_OPT\_CLEARIMPORT Setting this flag clears all import locks.



IMPORT\_OPT\_FORCE      Setting this flag performs the

VxVM 3.0

Last change: 05 Dec 1999

3

Miscellaneous Library Functions

libvxvm(3X)

import even if some disks are not accessible.

Caution: Incorrect use of IMPORT\_OPT\_CLEARIMPORT or IMPORT\_OPT\_FORCE flag. can cause data corruption.

libvxvm\_deport\_dg

Deports the disk group specified by dgname using the vxdg command. If newname is not NULL, the disk group is renamed to newname during deport. If newhostid is not NULL, libvxvm\_deport\_dg sets the hostid field of the disks belonging to this disk group during deport.

#### RETURN VALUES

Upon successful completion, the libvxvm library functions return a non-negative value. If there is an error, libvxvm returns a -1 and sets the appropriate error number. Possible error codes include:

EBUSY      The specified disk group is busy (cannot be deported).

EEXIST     The specified disk group exists (cannot be imported).

EINVAL     An invalid or inappropriate parameter was passed.

ENODEV     The specified volume does not exist.

ENOENT     The specified disk group does not exist.

ESRCH      The configuration daemon is not accessible.

ETIMEDOUT The configuration daemon is not responding or is busy.

#### SEE ALSO

vxdg(1M), vxintro(1M)

```
===[ man4/vol_pattern.4 ]=====
```

File Formats

vol\_pattern(4)

#### NAME

vol\_pattern - Volume Manager disk group search specifications

#### DESCRIPTION

The vol\_pattern search expression language can be used in vxprint and vxedit, and with some Volume Manager support library functions, to locate Volume Manager configuration records with particular characteristics. Searches are limited to operation on one disk group at a time. In the current release, search expressions are limited to volume, plex, and subdisk records. Search expressions cannot be applied to disk group, disk media, or disk access records.

Search patterns provide an expression-based language for specifying record attributes. In simple cases, patterns can be specified that match records whose fields have specific values. In more complex cases, records can be matched based on attributes of associated records.

Search patterns are expressions that do not have side effects (that is, variables cannot be set or changed). Expressions evaluate either to true or false. If evaluation of an expression within the context of a configuration record yields a value of true, then the expression matches the record; otherwise, the expression does not match the record.

The search pattern language does not have operators that cause changes to be made, and does not have a concept of local variables that can be set and used within or between expressions.

## Types of Expressions

Expressions in the pattern language are similar to expressions in awk(1). Expressions can have the following types for constants, variables, and subexpressions:

- boolean** A boolean variable or expression is either true or false. Numbers and strings can be used as truth values if the context requires a boolean value. For a number, zero implies false and nonzero implies true. For a string, an empty string implies false and a nonempty string implies true.
- string** A string is a literal string in quotes or a character array field from a configuration record. Strings can be implicitly promoted to boolean, but are never promoted to other types. For example, a string-valued field in a record that contains only digits cannot be compared with a number. String literals are described later.

## File Formats

vol\_pattern(4)

- number** A number is an integer quantity. The maximum size for a number is the largest size that may be used for values in configuration records. In the reference porting base for the Volume Manager, these numbers are 32-bit signed numbers.

Numbers can be specified in decimal, octal, or hexadecimal with a suffix to indicate the unit of the number when it indicates an object length or offset. Hexadecimal numbers begin with 0x, octal numbers begin with 0; otherwise numbers are in decimal. When used as a length, numbers with no suffix are taken as sectors.

Defined suffix characters are: b for 512-byte blocks, s for sectors (typically 512 bytes), k for kilobytes, m for megabytes, and g for gigabytes. Suffix characters also can be specified in upper case.

Length numbers are actually defined in terms of sectors. As a result, on a system with 512 byte sectors, 1k is the same as 2. Thus, the expression:

nassoc=1k

tests for a record with 2 associated records, rather than 1024 associated records.

#### enumeration

An enumeration field is a number that represents one of the values for a field that has symbolic names. The `vol_pattern` expression grammar handles enumeration values exactly like numbers. Enumerations are usually only useful for comparing compatible enumeration field and enumeration constants. For example, the enumeration field `pl_layout` enumeration field is a number whose possible values are represented by the symbols `CONCAT` and `STRIPE`.

#### regular expression

A regular expression is a value, much like a regular expression in `awk` or `ed`, that can be used to specify a pattern for matching strings. Regular expressions can only be used in the context of the boolean `~` and `!~` operators.

#### sequence numbers and record IDs

Each time a transaction is issued against a disk

#### File Formats

`vol_pattern(4)`

group configuration, a 64-bit configuration sequence number is increased. When a configuration records is created or changed in a transaction, the record is stamped with the configuration sequence number associated with that transaction.

When a record is created, it is given a 64-bit record ID. These record IDs are never reused within a disk group configuration.

Configuration sequence numbers and record IDs can be specified as constants of the form `number`, where both numbers are specified in decimal.

Two higher-level types are also included, which can represent multiple values of the same base type. These

higher-level types are:

**list**        A list is a set of expressions and sets within parenthesis that are separated by commas. List expressions can only be used for the right-hand expression of the in operator.

**set**         A set is a multi-valued variable (an array). Sets are normally used within the right-hand expression for the any and in operators. Reference to a set expression yields all of the currently valid members of the set, of which there may be none.

#### EXPRESSION GRAMMAR

A pattern is an expression with the following components and operators, in increasing order of precedence. Groups of operators described together have the same precedence.

**expr1 ? expr2 : expr3**

This is a conditional expression. Expression **expr1** is evaluated as a boolean value. If it has a truth value of true, then the value for the conditional expression is the value of **expr2**, otherwise the value for the conditional expression is the value of **expr3**.

**expr1 || expr2**

The logical or operator. **expr1** and **expr2** are evaluated as boolean expressions. If either **expr1** or **expr2** have a truth value of true, the expression is true; otherwise, the value of the expression is false.

**expr1 && expr2**

The logical and operator. **expr1** and **expr2** are evaluated as boolean expressions. If both **expr1**

and **expr2** have a truth value of true, the expression is true; otherwise, the value of the expression is false.

**any expr**

The expression **expr** is evaluated as a boolean expression for all combinations of elements for all sets within **expr** that represent sets. If any of these combinations yields a truth value of

true, then the value of the any expression is true; otherwise, the value is false.

For example, pl\_sd is a set defining the names of subdisk records that are associated with a plex. The expression:

```
any pl_sd.sd_disk="disk01"
```

will search for plex records that have an associated subdisk on disk disk01. Each name in the pl\_sd\_name set will be referenced as a subdisk and checked to see if that subdisk is on disk disk01. The expression is true if it is true for any subdisk.

expr in list

Expression expr is evaluated, and the resulting value is compared to all of the values in list. If any of the values in list match expr, then the value of the in expression is true; otherwise, the value is false. The list can either be a list of expressions within parentheses that are separated by commas, or it can be a single expression. If any of the expressions in list represent sets, then the comparison is made against all combinations of values in that set.

Type promotions are not performed on the expressions in list, or on expr. Types for all these expressions must match exactly.

For example, the expression:

```
"disk01" in  
(pl_sd.sd_disk,pl_log_sd.sd_disk)
```

will search for plexes that have either a regular subdisk or a subdisk on disk disk01.

expr ~/regexp/

expr!~/regexp/

The binary ~ and !~ operators are used to match

string values against a regular expression. The value of `expr` must be a string. For the `~` operator, if the value of `expr` matches the regular expression `regexp`, then the value for the match expression is true; otherwise the value is false. For the `!~` operator, if `expr` matches `regexp` the result is false; otherwise the result is true.

Regular expressions are in the format described in `regcmp(3G)`. The slash (`/`) character in the regular expression must be escaped with a backslash (`\`) if it does not occur within a character range.

For example, to search for subdisks that are defined on partition 14 (slice `e`) of any device, use:

```
sd_path=/^\dev/vx\dmp\[^\]*se/
```

```
expr1 = expr2  
expr1 != expr2
```

These expressions compare string, number, sequence number, record ID, or boolean values. If either `expr1` or `expr2` is a boolean value, then the other is promoted to a boolean as well. For the `=` operator, if the two values are equal then the result is true; otherwise, the result is false. For the `!=` operator, if the two values are equal, then the result is false; otherwise, the result is true.

For consistency with `awk(1)` and many other languages, `==` is allowed as an alias for the `=` operator.

```
expr1 < expr2  
expr1 > expr2  
expr1 <= expr2  
expr1 >= expr2
```

These expressions compare the magnitudes of numbers or sequence numbers. They cannot be applied to any other types of expressions. The result is a boolean.

```
expr1 + expr2  
expr1 - expr2
```

These expressions add or subtract two numbers, yielding another number.

expr1 \* expr2  
expr1 / expr2  
expr1 % expr2

These expressions multiply, divide, or take the modulo of two numbers, yielding another number.

+ expr  
- expr

The unary + operator applied to any expression yields the value of that expression. The unary - operator negates the value of a number-valued expression expr.

! expr

expr is evaluated as a boolean expression. If its value is true, then the value for the ! expression is false; otherwise, the value for the ! expression is true.

expr . field-name

expr is evaluated as a string expression. Its value is used as the name of a record in the disk group configuration. If the record exists, and if that record contains a field named field-name, then the value of this expression is the value in that field within the referenced record. If the record does not exist, or if the record does not contain a field named field-name, then the behavior is context-dependent.

If this occurs within the context of an any or in operator, then the next combination of values for the right-hand side of the any or in operator is evaluated. If this occurs within the context of an && or || operator, then the left or right-hand expression of that operator yields false, whichever side the . operator occurs on. The innermost any, in, &&, or || operator determines the context. If no such context exists, then the result of the entire search pattern is that the record being tested does not match the pattern.

The effect of this short-circuiting behavior is that an expression which is found not to make sense because a record does not exist, or does not contain a specific field, will cause that expression to be false. The effect is similar to implicit asserts, which are described later in this section under field name.

For compatibility with earlier releases of VxVM,



-> can be used instead of . as the operator.

( expr ) Parentheses can be used to override precedence

VxVM 3.0

Last change: 26 Mar 1999

6

File Formats

vol\_pattern(4)

rules.

literal A string, number, sequence number, or record ID literal, as defined in the section Types of Expressions.

assert An assert is a special type of variable name that is true within the context of a specific type of record. Outside of that context, an assert causes the expression that it is in to be false. The scope of an assert is the right or left-hand side of the smallest expression within a || or && operator. If an assert does not occur within a || or && operator, then the search pattern does not match the record being evaluated. The names that specify asserts are vol, plex and sd which specify volume, plex and subdisk records, respectively.

constant There is a set of symbolic constants that can be specified in search patterns. For example, true and false are the two truth values. Also, if a field within a record has a specific set of symbolic values, then these symbolic values are available as numeric constants. See the section Constants.

variable There is a small set of global variables that can be referenced in search patterns. Such variables can be used within any context. See the section Global Variables.

field name

Referencing the name of a field in a configuration record yields the value of that field within the configuration record being evaluated.

Not all field names are defined for all record types. If a field name is referenced that is not valid in the context of the record type, then the smallest expression on either side of a && or || operator becomes false. This is an implicit

assert. For example, the expression:

```
v_kstate!=V_ENABLED
```

will match only volume records because the v\_kstate field is only defined for volume records.

Some fields in a record only make sense if another field contains a specific value. For example, the pl\_st\_width field in plex records is not defined

VxVM 3.0

Last change: 26 Mar 1999

7

File Formats

vol\_pattern(4)

unless the value of the pl\_layout field is PL\_STRIPE. These cases are treated as implicit asserts.

#### CONSTANTS

The names of constants that are recognized within search patterns are:

true, yes, on

Constants representing boolean true.

false, no, off

Constants representing boolean false.

ROUND, PREFER, SELECT

The enumeration constants that represent values for the v\_read\_pol field in volume records. The v\_current\_read\_pol field can be set to ROUND or PREFER. These constants can also be specified with a prefix of V\_ to match the enumeration values in the include file sys/vol.h.

DISABLED, DETACHED, ENABLED

The enumeration constants that represent values for the v\_kstate field in volume and plex records. These constants can also be specified with a prefix of V\_ or PL\_.

NO\_OP, FAIL\_OP, DET\_PL, FAIL\_DET\_PL,

DET\_VOL, FAIL\_DET\_VOL, GEN\_DET,

GEN\_DET\_SPARSE, GEN\_FAIL, GEN\_DET2

The enumeration constants that represent values for the v\_r\_all, v\_r\_some, v\_w\_all, and v\_w\_some

fields in volume records. These constants can also be specified with a prefix of V\_.

LOGUNDEF, LOGNONE, LOGBLKNO

The enumeration constants that represent values for the v\_log\_type field in volume records. These enumeration constants can also be specified with a prefix of V\_PFLAG\_ to match the enumeration values in the include file /usr/include/sys/vol.h.

STRIPE, CONCAT

The enumeration constants that represent values for the pl\_layout field in plex records. These constants can also be specified with a prefix of PL\_.

RO, WO, RW

The enumeration constants that represent values for the pl\_iomode field in plex records. These

VxVM 3.0

Last change: 26 Mar 1999

8

File Formats

vol\_pattern(4)

enumeration constants can also be specified with a prefix of PL\_TFLAG\_ to match the values in the include file /usr/include/sys/vol.h.

NODEV The value for a device number field indicating that the device number is undefined.

#### GLOBAL VARIABLES

The global variables that can be used within search patterns are:

config\_tid or tid

The configuration sequence number of the disk group configuration.

util\_num or nutil

The number of fields in a tutil or putil arrays in a volume, plex, or subdisk structure. This is always 3, but it is treated as a global variable anyway.

#### GENERIC FIELDS

Some fields can be found in all types of configuration records. These fields can be referenced without a v\_, pl\_, or sd\_ suffix, and do not generate implicit asserts. These

fields can be used either to name fields in the record being evaluated, or as a field name reference for the . operator.

A generic field name can be specified with a prefix of v\_, pl\_, or sd\_ to indicate that implicit asserts should be used and that the field requires a context of a volume, plex, or subdisk record, respectively. Currently, use of a prefix for a generic field name is ignored for the right-hand side of a . operator.

The recognized generic field names are:

tutil        A set of strings representing all of the temporary utility fields.

tutil0 or t0        A string representing the first temporary utility field.

tutil1 or t1        A string representing the second temporary utility field.

tutil2 or t2        A string representing the third temporary utility field.

## File Formats

vol\_pattern(4)

putil        A set of strings representing all of the permanent utility fields.

putil0 or p0        A string containing the first permanent utility field.

putil1 or p1        A string containing the second permanent utility field.

putil2 or p2        A string containing the third permanent utility field.

name or n        A string containing the record name.

comment or c           A string containing the comment field.

len or l           A number representing the length of the volume, plex or subdisk.

open           A boolean value indicating whether the volume, plex, or subdisk is open. A subdisk is considered open if the volume or plex device for the subdisk is open. A plex is considered open only if the plex device is open. A volume is considered open if the volume or one of its plex devices is open.

rid           The record ID of the record being evaluated.

update\_tid or update           The disk group configuration sequence number corresponding to the transaction that created or that last modified the record.

lock           A boolean value indicating whether the record is locked in the current transaction of the caller. This variable is not useful when used with vxprint or vxedit. It is only useful for C programs.

Two special fields exist that do not obey the prefix convention, but which apply to more than one type of record. These are:

assoc           This is the name of a field to which a record is associated. In the context of a plex or subdisk record, assoc yields the name of the volume or plex to which the record is associated. If the record is not associated, assoc yields an empty string. Use of assoc is not appropriate in the

context of a volume record. This causes an implicit assert when applied to a volume.

aslist           This is the set of associated record names. In the context of a volume or plex record, this yields the names of the associated plexes and subdisks, of which there may be none. This is not appropriate in the context of a subdisk record. This causes an implicit assert when applied to a

subdisk.

naslist This is the number of records in the aslist set. This causes an implicit assert when applied to a subdisk.

#### VOLUME RECORD FIELDS

The fields that can be specified for volume records are:

v\_rec\_lock

A boolean value indicating whether the record is locked within the current transaction of the caller.

v\_data\_lock

A boolean value indicating whether the volume is locked against I/O activity within the current transaction of the caller.

v\_kstate

An enumeration field for the kernel-enabled state for the volume. This can have one of the values DISABLED, DETACHED, or ENABLED.

v\_r\_all, v\_r\_some, v\_w\_all, v\_w\_some

These are numbers that represent the exception policies for the volume. Each of these fields can have one of the exception policy values:

NO\_OP  
FAIL\_OP  
DET\_PL  
FAIL\_DET\_PL  
DET\_VOL  
FAIL\_DET\_VOL  
GEN\_DET  
GEN\_DET\_SPARSE  
GEN\_FAIL  
GEN\_DET2

v\_lasterr

A number representing a sequence number for the last error that occurred on the volume.

detach\_tid

The sequence number of the last transaction that

detached the volume due to a failure of some kind.

This is also set when all copies of the block-change-log for the volume are detached by I/O failures.

- `v_use_type` A string containing the usage type of the volume.
- `v_fstype` A string containing the file system type for the volume, as set by the administrator.
- `v_state` A string representing the usage-type-dependent state for the volume.
- `v_writeback` A boolean value indicating whether the writeback-on-read-error policy is turned on for the volume. This policy is set by default and can be turned off or back on using `vxedit`. See `vxedit(1M)` for more information.
- `v_writecopy` A boolean value indicating whether volume writes are copied to a private buffer in the kernel before being written to disk. See `vxedit(1M)` for more information.
- `v_specify_writecopy` A boolean value indicating whether the value of `v_writecopy` was specified explicitly by the administrator using `vxmake` or `vxedit`. If `v_specify_writecopy` is false then the value for `v_writecopy` will be derived from the configuration of the volume: `v_writecopy` will be turned on if block-change-logging is in operation on the volume, and will be turned off otherwise.
- `v_pl_num` or `v_nplex` A number representing the number of plexes associated with the volume.
- `v_pl_name` or `v_plex` or `v_plexes` The set of plex names that are associated with the volume.
- `v_read_pol` An enumeration representing the read policy for the volume. This can be one of the values `ROUND`, `PREFER`, or `SELECT`. The value `SELECT` indicates that read policy will be adjusted to match the configuration of the volume: if there is one associated striped plex, it will be preferred;

otherwise the policy will be round-robin.

`v_current_read_pol` or `v_creadpol`

An enumeration representing the current read policy that is in effect for the volume. This is the same as `v_read_pol` unless `v_read_pol` is `SELECT`, in which case `v_current_read_pol` matches the chosen policy.

`v_pref_name`

A string containing the preferred plex for the volume. This is only defined if the value for `v_current_read_pol` is `PREFER`. Otherwise, use of this variable causes an implicit assert.

`pref_plex_rid`

The record ID of the preferred plex.

`v_minor`

A number representing the minor number of the volume block and character devices. This is the persistent device number for the volume. If two volumes on the system have the same value for `v_minor`, then one of the volumes will be given a different device number, which will be reflected in `v_bdev` and `v_cdev`.

`v_bdev` and `v_cdev`

The block and character device numbers for the volume.

`force_minor`

A boolean value indicating whether the volume can use one of the reserved minor numbers (minor numbers zero through four). The root and swap usage types set this flag to allow them to grab minor numbers 0 and 1 for the root and swap volumes.

`v_uid`

A number representing the user ID for the volume. There is currently no way to represent a user name in a search pattern.

`v_gid`

A number representing the group ID for the volume. There is currently no way to represent a group name in a search pattern.

`v_mode`

A number representing the file permission modes for the volume. There is currently no way to represent a symbolic mode in a search pattern.

`v_log_len` A number representing the length of logs associ-



ated with the volume, if any.

File Formats

vol\_pattern(4)

v\_log\_type

An enumeration representing the logging type for the volume. This can have one of the values LOGUNDEF (no specific type specified), LOGNONE (no logging), or LOGBLKNO (block-change logging).

v\_logging A boolean value that is true if logging is enabled for the volume. Logging is enabled if the volume is enabled, and it has at least one associated, enabled plex with an enabled, associated log sub-disk.

v\_log\_serial\_hi, v\_log\_serial\_lo

Two numbers that together represent the most significant and least significant long values for the current log serial number for a volume. The log serial number, represented as a combined binary integer, is incremented every time a new log is written for a volume that has logging enabled.

badlog A boolean value indicating that all block-change-log copies for the value were detached due to I/O failures.

v\_start\_opts

A string value that is used by usage types to represent options to the volume start utility. See the usage-type-specific options in the vxintro(1M), vxmake(1M), vxmend(1M), vxplex(1M), vxprint(1M), vxrelayout(1M), vxsd(1M), and vxvol(1M) manual pages for more information.

v\_active A boolean value indicating whether the volume was in use at the time of a system failure. This will not be set if the volume was not written prior to the crash. This flag indicates that plex synchronization recovery is required for the volume. v\_active is normally cleared during system startup when the volume is started.

v\_rwback A boolean value indicating that read-writeback plex synchronization recovery has been started for the volume. This is normally set when a mirrored

volume is started when the v\_active flag was set.

v\_krwbck A boolean value indicating whether the read-writeback recovery mode is still enabled in the kernel. This flag is turned on when v\_rwbck is turned on, but is cleared by the kernel when recovery for the volume is complete.

v\_rwbck\_offset

VxVM 3.0

Last change: 26 Mar 1999

14

File Formats

vol\_pattern(4)

A number indicating the offset within the volume to where recovery has completed. A read operation on a volume performs a plex synchronization recovery if the read is after this point in the volume. A read that occurs at the current read-writeback offset advances the offset to the end of the read. When the offset reaches the end of the volume, recovery is complete and v\_krwbck is turned off.

v\_iosize A number indicating the largest sector size, in bytes, of any disk associated with the volume.

#### PLEX RECORD FIELDS

The fields that can be specified for plex records are:

pl\_compact

A boolean value that is true if the plex is compact. A plex is compact if its entire length is mapped by subdisks, leaving no holes. Note that the length of a plex is computed as the length to the end of the last block in the plex address space that is mapped by a subdisk.

pl\_contig\_len

The offset to the first unmapped block in the plex. If the plex is compact, this will match the plex length. If the plex is not mapped, this will be the offset of the first hole in the plex.

pl\_iomode An enumeration representing the I/O mode of the plex. This can have one of the values RO (read-only), WO (write-only), or RW (read-write).

pl\_log A boolean value that is true if block-change-

logging is enabled for the plex.

pl\_volatile

A boolean value that is true if one of the subdisks in the plex is defined on a volatile disk. Volatile disks are disks whose contents are not guaranteed to be preserved across a system reboot. Volatile disks can be created using the nopriv disk type (see vxdisk(1M)).

pl\_kstate An enumeration for the kernel enabled state for the plex. This can have one of the values DISABLED, DETACHED or ENABLED.

pl\_lasterr

A number representing a sequence number for the last error that occurred on the plex.

VxVM 3.0

Last change: 26 Mar 1999

15

File Formats

vol\_pattern(4)

pl\_minor A number representing the minor number of the plex device. This will be -1 if the plex is dissociated.

pl\_dev A number representing the device number of the plex device (a character device). This will be NODEV if the plex is dissociated.

pl\_noerror

A boolean value that is true if the plex is accounted for in volume error handling. If this is set, then I/O errors on the plex are ignored if they result from normal volume I/O.

pl\_v\_name or pl\_volume

The name of the associated volume. If the plex is not associated, this is an empty string.

pl\_sd\_num or pl\_nsd

A number representing the count of subdisks that are associated with the plex.

pl\_sd\_name, pl\_sd, pl\_subdisk, or pl\_subdisks

A set containing the names of subdisks that are associated with the plex.

pl\_state A string representing the usage-type-dependent

state for the plex.

`pl_layout` An enumeration representing the subdisk layout policy for the plex. This can be one of the values `CONCAT` or `STRIPE` to represent concatenation or striping, respectively.

`pl_st_width`  
A number representing the stripe width for the plex. This is only defined if the value for the `pl_layout` field in the plex is `STRIPE`. Otherwise, use of this variable causes an implicit assert.

`log_sd` A string representing the name of a subdisk associated with the plex as a log subdisk. This string is empty if no log subdisk is associated with the plex.

`pl_removed`  
A boolean indicating whether any subdisk in the plex is defined on a removed disk. See `vxdg(1M)` for information on removed disks.

`pl_nodarec` or `pl_failed`  
A boolean indicating whether any subdisk in the

plex is defined on a failed disk. A failed disk is one where the underlying physical disk could not be found, or was detached by an I/O failure.

`pl_stale` or `pl_recover`  
A boolean indicating whether a disk replacement or a temporary failure caused the plex to become disabled and out-of-date with respect to the volume contents. Such plexes require recovery either from other plexes in the volume or from a backup of the volume.

`pl_kdetach` or `pl_iofail`  
A boolean indicating that the plex was detached by the kernel due to an I/O failure. The stale and `kdetach` flags are normally both cleared during a system reboot as part of starting the volume.

`pl_vol_rid` or `pl_v_rid`

The record ID of the volume to which the plex is associated. This will be 0.0 if the plex is not associated.

pl\_log\_sd\_rid

The record ID of the log subdisk for the plex. This will be 0.0 if the plex has no log subdisk.

detach\_tid

The configuration sequence number of the disk group that was in effect when the plex was last detached or disabled due to an I/O failure or a disk failure.

#### SUBDISK RECORD FIELDS

The fields that can be specified for subdisk records are:

sd\_dm\_name or sd\_disk

The name of the disk media record on which the subdisk is defined.

sd\_da\_name or sd\_device

The disk access record underlying the disk media record of the subdisk. This will be empty if the disk is currently failed or removed.

sd\_dev

The device number of the device that the subdisk resides on. This number is a combination of the major and minor number for the device. This will be NODEV if the underlying disk is removed or failed.

sd\_path

A string containing the path to the device that

the subdisk resides on.

sd\_pl\_name or sd\_plex

The name of the associated plex. If the subdisk is not associated, this is an empty string.

sd\_dm\_offset or sd\_offset

A number representing the offset of the subdisk within the public region of the disk that the subdisk resides on.

sd\_dev\_offset

A number representing the offset of the subdisk within the disk partition containing the public region for the disk. This can differ from sd\_dm\_offset if the public region of the disk does not begin at the beginning of its partition.

sd\_pl\_offset

A number representing the offset of the subdisk within the address space of the plex to which it is associated.

sd\_volatile

A boolean value that is true if the subdisk is defined on a volatile disk. The contents of a volatile disk are not guaranteed to be preserved across a reboot. Volatile disks can be created using the nopriv disk type (see vxdisk(1M)).

sd\_is\_log or sd\_log

A boolean value indicating that this is a log subdisk for a plex.

sd\_plex\_rid or sd\_pl\_rid

The record ID of the plex to which the subdisk is associated.

sd\_dm\_rid or sd\_disk\_rid

The record ID of the disk media record on which the subdisk is defined.

#### EXAMPLES

The following example matches any plex or volume that does not have any associated subdisks or plexes. Because pl\_nsd only applies to plexes, while v\_nplex only applies to volumes, there is no ambiguity concerning which types of records can be matched by the expression.

```
pl_nsd = 0 || v_nplex = 0
```

The following example matches all plexes and subdisks that have a comment field beginning with the string Henry. For example, this could be a convention describing ownership.

```
(plex || sd) && comment ~ /^Henry/
```

The following example matches all volumes for which the preferred plex is striped:

```
v_pref_name.pl_layout = STRIPE
```

Because `v_pref_name` is only defined for volumes that have `v_read_pol` set to `PREFER`, the preceding expression is equivalent to:

```
vol && v_read_pol = PREFER && v_pref_name.pl_layout =  
STRIPE
```

The following expression matches any volume that contains a plex with a subdisk longer than 100 megabytes:

```
any v_plex.pl_subdisk.sd_len > 100M
```

If an expression contains a set but does not use the `any` or `in` operator, then an `any` operator is assumed to surround the expression. Thus, the preceding example can also be specified as:

```
v_plex.pl_subdisk.sd_len > 100M
```

The following example matches any record that was changed in the most recent set of changes to the disk group:

```
update_tid = config_tid
```

Alternatively, to get all records that have been changed since a past disk group configuration (with an configuration transaction ID of 0.23482), use:

```
update_tid > 0.23482
```

#### SEE ALSO

```
awk(1), vxdg(1M), vxdisk(1M), vxedit(1M), vxintro(1M),  
vxmake(1M), vxmend(1M), vxplex(1M), vxprint(1M),  
vxrelayout(1M), vxsd(1M), vxvol(1M), regcmp(3G)
```

===[ man4/vxmake.4 ]=====

File Formats

vxmake(4)

#### NAME

vxmake - format of a vxmake description file

#### DESCRIPTION

Descriptions of configuration records can be given to the vxmake utility either on the command line or in a file. Descriptions given on the command line can specify only one new record. Descriptions given in a description file can specify more than one record. The description file format is a convenient way of indicating record contents for other uses as well. For example, the vxprint utility can print records in this format to provide input to a shell or awk script that can be more easily parsed than other verbose formats.

NOTE: The vxprint or vxprivutil utilities can be used with vxmake to perform the functions of a rudimentary configuration backup and restore. This should be the primary use of vxmake. Use of vxmake other than through a description file containing content mostly gathered from either the vxprint or the vxprivutil utilities is discouraged in Volume Manager releases starting with 3.0. The vxassist utility is the recommended means of creating new or modified volume hierarchies.

In many places within the utility set where a user can specify a field name and a value, a format is used that specifies only one configuration field, rather than complete configuration records.

#### Record Description Format

vxmake description files specify configuration records, one after another. Each record description specifies a record type, name, and zero or more fields.

The record type is a single keyword from the set vol, plex, sd, dm, or dg, to indicate a volume, plex, subdisk, disk media, or disk group record, respectively. The record name must be specified on the same line as the record type keyword.

The vxmake utility does not take disk media or disk group record descriptions on input. The vxprint utility can print descriptions of disk media and disk group records in this format though.

Fields are specified as a field name and a field value separated by a single = character. A record description can span more than one line if its continuation lines begin with



a character that is a space or tab. Empty lines are ignored and do not end a record description.

VxVM 3.0

Last change: 26 Mar 1999

1

File Formats

vxmake(4)

Comments can be included in a record description. A comment is introduced with the character #, either at the beginning of a line, or after a blank that is not used as part of a field value. Comments continue until the next new line character or until the end-of-file.

A field specification normally ends with the next white-space character. A field can be extended until the end of a line by preceding the field value with one double-quote character ("). Such quoted strings extend until the next newline, rather than until a double quote. This makes vxmake descriptions easier to generate within shell scripts.

An example of a valid configuration file is

```
vol v1 usetype=gen      # simple gen volume with no plexes
plex p1                 # plex with no subdisk
plex p2 sd=s1:0,s2:1   # striped plex with two columns
                        # and a comment

    comment="This is a comment
    layout=stripe ncolumn=2 stwidth=64k

plex p2 sd=s3:16384 # free plex with hole at beginning

# subdisks
sd s1 disk=disk01 len=1g+512m      offset=0
sd s2 disk=disk02 len=512m-1      offset=0
sd s3 path=disk03 len=32768        offset=512m-1
```

#### Field Description Format

Sometimes a utility gets field descriptions from somewhere other than a vxmake description file and uses that field description to set a field in a volume, plex, or subdisk record. In these cases, the utility itself has a quoting convention for separating one field from another. An example of a utility that does this is vxedit.

The vxedit utility takes field descriptions from the command line. For example, the following vxedit command sets a com-

ment for the record named v1:

```
vxedit set comment="Henry"
```

Shell quoting conventions are sufficient to allow for double quotes and spaces, so no quoting conventions are defined for these simple field descriptions.

#### FIELDS

Within vxmake description files any volume, plex, or subdisk

VxVM 3.0

Last change: 26 Mar 1999

2

#### File Formats

vxmake(4)

field can be specified. However, some field settings are ignored. This allows a user to use a command such as:

```
"vxprint -vpshm v1"
```

to get a set of record descriptions that can be fed back into vxmake, possibly with some changes. If this is done, then any field that cannot be set is ignored, but the resulting record hierarchy, plex layout, and subdisk offsets are the same, along with most or all fields that are normally administrator-definable.

Different fields have different properties with respect to the description language. The classes of fields are

**boolean** Boolean fields contain a truth value. The value for a boolean field can be either off or on. Alternately, yes or no, or true or false can be used. vxprint prints truth values as on or off.

#### device-number

Device numbers define the major and minor number of a device. These numbers are in the form major-number.minor-number where major-number and minor-number are both simple numbers. In some cases, the device number will be NODEV, indicating that the device number is undefined.

#### enumeration

An enumeration field can be set to one of a set of symbolic values. The set of symbolic values for

enumeration fields are specified in the field descriptions below. The enumeration fields are listed in this manual page in upper case. These enumeration constants can be given as input in either all upper-case letters, or all lower-case letters.

#### length-number

Fields which define object sizes of some type, such as volume lengths, plex offsets, or stripe widths, have the syntax of standard Volume Manager length numbers, as defined in vxintro(1M). This syntax consists of one or more numbers separated by + or - to indicate addition or subtraction.

Each separate number in the string can have a prefix of 0x to indicate that the number is in hexadecimal, or a prefix of 0 to indicate octal. Suffix characters can be used to specify the unit for

the number. Defined suffixes are: b for 512-byte blocks, s for standard system sectors, k for kilobytes, m for megabytes, and g for gigabytes. Suffix letters can be specified either upper case or lower case.

The size of a sector is system dependent. On most systems it is 512 bytes, the same as blocks. On systems that support disks with differing sector sizes, one standard sector size is chosen as the standard.

The unit for a number is in sectors by default. vxprint prints length numbers as a decimal number of sectors, with no unit suffix.

Since b and B are valid digits in a hexadecimal number, a single space must be used to separate the hexadecimal number from the suffix. Double quote are not necessary in Record Description Format, unlike the rules for string-valued fields that contain spaces.

Unless otherwise indicated, all length numbers must be greater than or equal to zero.

name           A record name field can be an arbitrary string of up to fourteen characters. However, a name cannot contain a null byte, white-space characters, or the character /.

record ID or sequence number

These types are 64-bit numbers encoded as two 32-bit unsigned simple numbers. These are values assigned to records or to configuration changes, respectively. The record ID assigned to a record is guaranteed to be unique within that disk group. Also, that record ID will never be reused if the record is deleted.

Every time the configuration of a disk group is updated, the configuration update sequence number is increased. The configuration update sequence number only increases; it never decreases. This can be used to determine the order in which records were changed, or as a quick test for what has changed since a known time in the past.

Record IDs and sequence numbers have the form high-part.low-part, where both high-part and low-part are simple unsigned numbers.

## File Formats

vxmake(4)

simple-number

Simple numbers are used for counts or for other purposes besides representation of a length of offset. Simple numbers can be specified in decimal (default), hexadecimal, or octal using the 0x or 0 prefixing convention. No suffixes or arithmetic are defined for simple numbers.

string

Other types of string fields can contain any characters other than a null byte or a newline. Different string fields have different length restrictions.

Some additional one-of-a-kind field types also exist. These are described in the description for the corresponding field.

Many fields have aliases that can be used in vxmake descriptions. In all cases, the vxprint -m format uses the first name given for each field. However, alias field names can be used in format strings specified to vxprint with the -F option. All field names containing underscores have an alias that is identical but lacks the underscores. For example, start\_opts has an alias of startopts.

#### Volume Record Fields

These are the fields that can be used for specifying volume records:

active      A boolean. This is set to on if the volume was open and had at least one pending write at the time of the last system failure. Utilities can use this flag to indicate that plex consistency recovery should be performed. If this flag is off, the volume should not require any recovery.

badlog      A boolean. This flag is set to on if write failures occur on all logging subdisks for a volume that has dirty region logging enabled. Dirty region logging cannot be reenabled until this flag is turned off. This flag can be turned off by changing the logging type of the volume to NONE and setting it back to DRL.

bdev (alias: dev)      The block device number for the volume. This is a read-only field and attempts to set it are ignored.

cdev      The character (raw) device number for the volume. This is a read-only field and attempts to set it are ignored.

#### File Formats

vxmake(4)

comment      A comment-type field of up to 40 characters. This is reserved for use by local installations.

config\_id, config\_epoch      Simple unsigned number fields. These are the lower and upper 32 bits, respectively, of the update\_tid field. These fields are provided for backward compatibility with earlier releases of the Volume Manager. This field is not printed by

the -m flag to vxprint.

These are read-only fields and attempts to set them will be ignored.

detach\_tid

The configuration update sequence number in effect the last time the plex was detached by an I/O failure. This is a read-only field and attempts to set it are ignored.

forceminor

A boolean. If this is on, then the volume can use one of the reserved minor numbers (zero through four). The root and swap usage types set this automatically to grab volume minor numbers 0 and 1 for the root and swap volumes, respectively.

fstype

A name field. This is the file system type associated with the volume contents. A usage type is free to use or ignore this value.

group

This is either a group name or a decimal group ID. The value specifies the user that owns the volume block and character devices.

iosize

The smallest I/O size allowed on the volume. This corresponds to the largest sector size for any disk included in the volume. The current release of the Volume Manager does not support disks with differing sector sizes. As a result, in the current release this value will always match the system default sector size (normally 512 bytes).

kstate

An enumeration field. This is the kernel-enabled state of the volume. Possible values are ENABLED, DISABLED, and DETACHED.

lasterr

A simple number field. This is the sequence number of the last error on the volume. It is a read-only field and attempts to set it are ignored. This number does not have a default multiplier, so lack of a suffix specifies a simple

integer.

len           A length number field. This is the volume length. If a volume is created in vxmake with a plex attribute that associates plexes, the volume length will default to the length of the shortest associated plex. If no plexes are associated at creation time, the volume length defaults to zero, though a warning will be printed.

log\_len       The length for the volume's logging area. This is primarily intended for use with future logging types. When used with dirty region logging, the only currently available method of logging, the value of log\_len must be an even number of blocks between two and ten blocks long.

log\_serial\_lo and log\_serial\_hi  
               Simple unsigned number fields. These values, taken together, yield a value that is increased for every log write that occurs to a volume with logging enabled.

              These are read-only fields and attempts to set them are ignored.

log\_type     An enumeration field. This is the logging type for the volume. Possible values are: UNDEF, NONE, and DRL. The first two possible values indicate that logging is disabled, and the third value indicates use of dirty region logging, where block numbers for write operations are logged before the corresponding blocks are actually written to disk.

              The UNDEF value indicates that no logging policy has yet been chosen. If a log subdisk is associated to a volume with the UNDEF policy, the policy will be changed to DRL automatically.

logging      A boolean field. This is on if logging is enabled on the volume. Logging is enabled for a volume if the log type for the volume supports logging, the volume is enabled, the badlog flag is not set, and the volume has at least one enabled, associated plex with an enabled log subdisk. Normally, usage types do not enable logging if only one plex in the volume has a log subdisk.

              This is a read-only field and attempts to set it are ignored.

minor        A simple number field. This is the minor number

to use for the volume block and character devices. If volumes in two disk groups attempt to use the same minor number, one of the volumes will be given a different one. To determine the actual device number used for a particular volume, get the cdev and bdev fields.

If a minor number is specified as input to vxmake, the number will be used if it is not already in use, or is reserved; otherwise, a minor number will be chosen that is not currently in use or reserved. Minor numbers from zero to four are reserved for internal purposes.

**mode** This is a symbolic or numeric file permission mode. This can be any string that is acceptable to `chmod(1)` for setting the mode of a file.

**open** A boolean field. This is on if the volume is open or mounted. This is a read-only field and attempts to set it are ignored.

**pl\_num** (alias: `nplex`)  
A number field. This is the number of plexes associated with the volume. This is a read-only field and attempts to set it are ignored.

**plex** A list of names for plexes to associate with the volume. The names are separated in the list by a comma or by blanks. If a list is specified as input to vxmake, then the indicated plexes will be associated with the volume.

**putil0, putil1, putil2**  
Comment-type fields of up to fourteen characters. These are the permanent utility fields, which are preserved across a reboot. The first field is reserved for usage-type utilities. The second field is reserved for utilities, such as the Storage Administrator, that use the usage-type utilities and that need to store their own state into records. The third field is reserved for local installations.

**read\_pol, pref\_name, pref\_plex\_rid,**  
**current\_read\_pol** (alias: `creadpol`)  
`read_pol` is an enumeration field. This is the read policy for the volume. Possible values are `ROUND`, `PREFER`, `SELECT`. A value of `ROUND` indicates round-robin read scheduling, `PREFER` indicates a preferred plex. Round-robin scheduling scatters



reads evenly between all plexes. Preferred-plex

File Formats

vxmake(4)

read scheduling attempts to use a single plex for all reads.

Typically, if a volume consists of one striped plex and one non-striped plex, it is better to read only from the striped plex. This is because the striped plex can handle more write traffic and is thus more likely to be free to accept read requests with a minimum of delay.

If the SELECT read policy is used, the actual read policy (round-robin or preferred plex), is chosen automatically. The policy chosen with SELECT is to prefer a striped plex if there is exactly one striped plex in the volume, and to use round-robin otherwise. The read policy that is actually in effect is stored in `current_read_pol`.

If a preferred plex policy is in effect, either because PREFER is set or because the SELECT policy chose the preferred-plex policy, the value of `pref_name` indicates the name of the preferred plex. The SELECT policy sets the value of `pref_name` automatically. The ROUND policy ignores any value stored in `pref_name`.

The record ID for the preferred plex can be set or displayed with `pref_plex_rid`.

`r_all, r_some, w_all, w_some`

Enumeration fields. These are the exception policies for the volume, for read errors on all plexes, read errors on some plexes, write errors on all plexes, and write errors on some plexes. Possible values for these fields are NO\_OP, FAIL\_OP, DET\_PL, FAIL\_DET\_PL, DET\_VOL, FAIL\_DET\_VOL, GEN\_DET, GEN\_DET\_SPARSE, GEN\_FAIL, and GEN\_DET2.

Usage types normally ignore any attempts to set these fields in `vxmake`.

`rid`

The record ID for the volume record.

rwback, krwback, rwback\_offset

rwback and krwback are booleans. rwback\_offset is a length number. rwback is set to on to enable read-writeback mode on a volume. If this is set, then reads from the volume past rwback\_offset perform a recovery procedure that ensures that all plexes are consistent before the read operation completes. Reads that occur at the current read-

## File Formats

vxmake(4)

writeback offset will advance the offset. If the read-writeback offset reaches the end of the volume, read-writeback will be turned off. The value of krwback indicates whether read-writeback is actually in effect.

These fields are set by usage types. Attempts to set these fields in vxmake are ignored.

start\_opts

A comment-type field of up to 32 characters. This field is used in a usage-type-dependent manner. The existing usage types use this field as a set of options to apply to the volume for the vxvol start operation.

state

A comment field of up to fourteen characters. This field represents a usage-type-dependent volume state.

Usage types normally ignore attempts to set this field in vxmake.

tutil0, tutil1, tutil2

Comment-type fields of up to fourteen characters. These are the temporary utility fields, which are cleared by a reboot. The first field is reserved for usage-type utilities, and is typically used to define operation locks. The second field is reserved for utilities, such as the VERITAS Storage Administrator, that use the usage-type utilities and that need to store their own state into records. The third field is reserved for use by local installations.

update\_tid The configuration update sequence number for the last transaction to update this record. This is a read-only field and attempts to set it are ignored.

use\_type A name field. This is the name of the usage type for the volume.

user This is either a user name or a decimal user ID. The value specifies the user that owns the volume block and character devices.

writeback A boolean field. If on, then if a read failure on one plex of a multiplex volume the block will be read from another plex and written back to the plex with the failure. This often corrects the

## File Formats

vxmake(4)

I/O failure. If the writeback succeeds, the failure is ignored and no action is taken to detach the plex.

writecopy, specify\_writecopy  
Boolean fields. If writecopy is on, then writes to the volume will be copied prior to being written to disk. This prevents pages of memory that are undergoing changes from causing plexes in a volume to become inconsistent.

Normally, writecopy is not needed because the system pager ensures that a consistent version of each page is written to the volume prior to a clean system shutdown. For most volumes, recovery after an unclean system shutdown (that is, after a panic, unintended reset, or power failure) recovers all blocks in a volume to have consistent data between each plex.

However, volumes which use the dirty region logging feature recover only those regions which were being written at the time of an unclean shutdown. If some regions were written inconsistently at some point prior to a shutdown, then recovery may not ensure consistency of these regions. The dirty region logging feature is enabled when

log\_type is set to DRL and at least one plex has a log subdisk and the volume has at least two read-write mode plexes.

If specify\_writecopy is off then the value of writecopy is set automatically to be on if block-change-logging is enabled on the volume, and off otherwise. specify\_writecopy is set to on by utilities whenever the administrator explicitly sets the value of writecopy.

This combination of writecopy and specify\_writecopy eliminates the need for administrators to set this field explicitly, in most cases while allowing the administrator to override the default value, if desired. It is often desirable to turn off writecopy for block-change-logging volumes, to reduce the overhead of write operations. Whether plexes can become inconsistent, and whether inconsistencies pose problems, are matters that depend upon how the volume is used.

#### Plex Record Fields

These are the fields that can be used for specifying plex

VxVM 3.0

Last change: 26 Mar 1999

11

#### File Formats

vxmake(4)

##### records:

comment A comment-type field of up to 40 characters. This is reserved for use by local installations.

compact A boolean field. This is on if the plex is compact. A compact plex has no holes; that is, has no regions within the length of the plex that aren't backed by a subdisk. This is a read-only field and attempts to set it are ignored.

complete A boolean field. This is set to on to indicate to the GEN\_DET, GEN\_DET2, and GEN\_DET\_SPARSE exception policies (see volrec(4)) that the plex should be considered to be complete, relative to any volume to which the plex is associated.

##### config\_len

A length number field. This is the offset of the

first block in the plex that doesn't map to a sub-disk. If the plex is compact, `contig_len` will equal `len`; otherwise, `contig_len` will identify where the first hole in the plex is.

This is a read-only field and attempts to set it are ignored.

`dev` The device number of the plex device. This is a read-only field and attempts to set it are ignored.

`iomode` A pseudo enumeration field. This field does not represent an explicit enumeration field in the plex record. Rather it represents the possible values in the plex `pl_tflag` structure element (which is not accessible through the `vxmake` description format) that form the I/O mode of the plex. Possible values are `NONE`, `RO`, `WO`, and `RW`, which mean no I/O is possible, read-only, write-only and read-write, respectively. In general, only `RW` and `WO` are useful.

`kdetach` (alias: `iofail`)  
A boolean field. This is set to on if the plex becomes detached as a result of an I/O failure. If this is set, the plex contents may be out-of-date and will require recovery from another plex in the volume.

`kstate` An enumeration field. This is the kernel-enabled state of the plex. Possible values are `ENABLED`, `DISABLED`, and `DETACHED`.

## File Formats

`vxmake(4)`

`lasterr` A number field. This is the sequence number of the last error on the volume. It is a read-only field and attempts to set it are ignored.

`layout` An enumeration field. This specifies how blocks in the plex address space map onto blocks. Possible values are `CONCAT` and `STRIPE`. A value of `CONCAT` indicates that subdisks are simply concatenated together to form the blocks in the plex address space. A value of `STRIPE` indicates that columns made up of subdisks are mapped onto the

plex address space.

len           A length number field. This is the length of the plex. The length of a plex is computed from the end of the associated subdisk that has the largest plex offset.

This is a read-only field and attempts to set it are ignored.

log           A boolean field. This is set to on to enable use of any log subdisk associated with the plex. If off, then any associated log subdisk is ignored.

log\_sd        A name field. This is the name of an associated log subdisk, if such a subdisk has been associated to the plex with the vxsd aslog operation. This can be specified as input to vxmake to specify that the subdisk should be associated to the plex by vxmake.

log\_sd\_riid   The record ID of the log subdisk. Attempts to set this field in vxmake are ignored.

minor         A simple number. This is the minor number of the plex device.

ncolumn       A number of stripe columns. If the value of layout is STRIPE, then this specifies the number of columns in the striped plex. This number must be nonnegative.

nodarec (alias: nodevice)  
A boolean field. This is set to on if the disk for one of the subdisks in the plex has the nodarec flag set. This implies that the physical disk backing a subdisk either could not be found after a reboot, or failed entirely during operation.

This is a read-only field and attempts to set it are ignored.

noerror A boolean field. This field is on if I/O errors on the plex should not result in actions based on the volume exception policies. This maps onto setting the PL\_PFLAG\_NOERROR flag in the plex pl\_pflag structure element.

open A boolean field. This is on if the plex device is open. This is a read-only field and attempts to set it are ignored.

putil0, putil1, putil2  
 Comment-type fields of up to fourteen characters. These are the permanent utility fields. The first field is reserved for use by usage-type utilities. The second field is reserved for use by utilities, such as the Storage Administrator, that use the usage-type utilities and that need to store their own state into records. The third field is reserved for use by local installations.

removed A boolean field. This is set to on if the disk for one of the subdisks in the plex has the removed flag set. This implies that the physical disk backing a subdisk was removed by the vxvg rmdisk operation.

This is a read-only field and attempts to set it are ignored.

rid The record ID for the plex record.

sd\_num (alias: nsd)  
 A simple number field. This is the number of subdisks associated with the plex. It is a read-only field and attempts to set it are ignored.

sd A list of subdisks to associate with the plex. Each subdisk is separated by a comma. A specific plex offset can be given for a subdisk by following the subdisk name with a colon and a number. If no plex offset is specified, then the subdisk is associated at the end of the previous subdisk in the list. The default plex offset for the first subdisk in the list is zero.

st\_width A length number field. If the value of layout is STRIPE, then this specifies the length of each stripe. This length must be non-negative and must be an even multiple of a block (512 bytes on most

machines).

stale (alias: recover)

A boolean field. This is set to on if a disk is removed that containing one of the subdisks in the plex, or if the disk for one of the subdisks fails while the plex is in use. If this is set, the plex contents may be out-of-date and will require recovery from another plex in the volume.

state A comment field of up to fourteen characters. This is a usage-type-dependent volume state. Typically, usage types will ignore any attempts to set this field in vxmake.

tutil0, tutil1, tutil2

Comment-type fields of up to fourteen characters. These are the temporary utility fields. The first field is reserved for usage-type utilities. The second field is reserved for utilities, such as the Storage Administrator, that use the usage-type utilities and that need to store their own state into records. The third field is reserved for local installations.

update\_tid

The configuration update sequence number for the last transaction to update this record. This is a read-only field and attempts to set it are ignored.

v\_name (alias: volume)

A name field. This is the name of the volume to which the plex is associated. This is a read-only field when used with the vxmake utility and attempts to set it are ignored.

vol\_rid The record ID of the associated volume. If the plex isn't associated, the value of this field is 0.0.

volatile A boolean field. This is on if any subdisk associated with the plex has its volatile flag set to on; otherwise, this is off. This is a read-only field and attempts to set it are ignored.

#### Subdisk Record Fields

These are the fields that can be used for specifying subdisk records:

comment A comment-type field of up to 40 characters. This is reserved for use by local installations.



## File Formats

vxmake(4)

`config_id, config_epoch`

Simple unsigned number fields. These are the lower and upper 32 bits, respectively, of the `update_tid` field. These fields are provided for backward compatibility with earlier releases of the Volume Manager. This field is not printed by the `-m` flag to `vxprint`.

These are read-only fields and attempts to set them will be ignored.

`da_name` (alias: device or accessname)

The name of the disk access record used to access the physical disk that backs the subdisk. If the disk media record is in the `nodarec` or `removed` state, then this will be empty. This can be used as an alternative to setting `dm_name` as a means of specifying the disk to `vxmake`.

`dev`

This is the device number of the public region for the disk that the subdisk resides on. If the disk is removed or failed, the value for this field is `NODEV`. This is a read-only field and attempts to set it are ignored.

`dev_offset`

A length number. This is the offset of the subdisk from the beginning of the device containing the private region of the disk that backs it. This is not defined (prints as an empty string) if the disk media record for the subdisk is in the `removed` or `nodarec` state. This is a read-only field and attempts to set it are ignored.

`dm_name` (alias: disk or medianame)

The name of the disk media record that the subdisk is defined on. This field is the primary means of defining the disk as input to `vxmake`

`dm_offset` (alias: offset)

A length number. This is the offset of the subdisk from the beginning of the private region of the disk that backs it. This is a read-only field and attempts to set it are ignored.

dm\_rid      The record ID of the disk media record that the subdisk is defined on. This can be used as an alternative to setting dm\_name as a means of specifying the disk to vxmake.

is\_log (alias: log)  
A boolean field. This is on if the subdisk is

VxVM 3.0

Last change: 26 Mar 1999

16

File Formats

vxmake(4)

associated with a subdisk as a log subdisk.

len      A length number. This is the subdisk length. The length must be a nonnegative number, and must be a multiple of a standard disk block (512 bytes on most machines).

open      A boolean field. This is on if the associated plex device is open, or if the associated volume is open or mounted. This is a read-only field and attempts to set it are ignored.

path      A comment-type string field of up to 32 characters. This is the path to the block device for the public region of the disk that the subdisk resides on. If the disk is currently failed or removed, this field is empty. This field can be specified as an alternative to setting dm\_name when creating a subdisk record, in which case it will be used to determine the disk.

pl\_offset A length number. This is the offset of the subdisk in address space of the plex address. If the subdisk is not associated, the field is not defined (prints as an empty string). This is a read-only field and attempts to set it are ignored.

pl\_name (alias: plex)  
A name field. This is the name of the plex to which the subdisk is associated. This is a read-only field and attempts to set it are ignored.

plex\_rid      The record ID of the plex to which the subdisk is associated. If the subdisk is not associated, this is 0.0. This is a read-only field and attempts to set it are ignored.

putil0, putil1, putil2

Comment-type fields of up to fourteen characters. These are the permanent utility fields. The first field is reserved for usage-type utilities. The second field is reserved for utilities, such as the Storage Administrator, that use the usage-type utilities and that need to store their own state into records. The third field is reserved for local installations.

rid The record ID for the subdisk record.

tutil0, tutil1, tutil2"

Comment-type fields of up to fourteen characters.

VxVM 3.0

Last change: 26 Mar 1999

17

File Formats

vxmake(4)

These are the temporary utility fields. The first field is reserved for usage-type utilities. The second field is reserved for utilities, such as the Storage Administrator, that use the usage-type utilities and that need to store their own state into records. The third field is reserved for local installations.

update\_tid

The configuration update sequence number for the last transaction to update this record. This is a read-only field and attempts to set it are ignored.

volatile

A boolean field. This is set to on if the disk that backs the subdisk is defined to be volatile. See vxdisk(1M) for information on how to create a volatile disk with the nopriv disk type. Typically, such disks represent RAM disks defined in memory that is reset by a reboot or a power failure.

Disk Media Record Fields

These are the fields that can be used for specifying disk media records (Note: disk media records cannot be created using vxmake):

atomic\_iosize (alias: iosize)

The sector size of the device that backs the disk

media record. This field is undefined (prints as an empty string) if the disk media record is in the removed or nodarec state.

autoconfig

A boolean. This is on if the disk access record that underlies the disk media record was auto-configured during system startup. This is off if the disk media record is in the nodarec or removed state, or if the disk access record was created explicitly with vxdisk define or vxdisk init.

See vxdisk(1M) for more information.

comment A comment-type field of up to 40 characters. This is reserved for use by local installations.

component A boolean. This is on if the public and private regions are on separate partition devices. This is off if the disk media record is in the removed or nodarec state.

da\_name (alias: device or accessname)

File Formats

vxmake(4)

The name of the disk access record that can be used to access the disk that backs the disk media record. This is empty if the disk has the removed or nodarec flags set.

da\_type (alias: type)

The disk type for the disk access record that backs the disk media record. See vxdisk(1M) for a description of disk types. This field is empty if the disk media record is in the removed or nodarec state.

device\_tag

A name identifying the device that the disk access record references. If two or more disk access records identify the same actual device (for example, if they reference different sets of partitions of the same physical disk), the device\_tag will be the same for both disk access records. The device\_tag field is inherited by the disk media record from the disk access record.

This field is empty if the disk media record is in the removed or nodarec state.

`diskid` The disk ID for the physical disk that backs the disk media record. This is a 64-byte string that matches the disk ID stored in the disk header of the disk on which this disk media record is defined. This field is empty if the disk is in the removed state. `diskid` will be set to a new value if the disk media record is reassociated using `vxchg -k adddisk`.

`last_da_name` (alias: `last_device`)  
The name of the last disk access record used by the disk media record. This is set whenever the disk access record changes. The disk access record used for a disk media record can change if disks are moved during a reboot, or if a disk is replaced. The `last_da_name` field is preserved if a disk fails or is removed.

`nodarec` (alias: `nodevice`)  
A boolean. This is on if the physical disk underlying the disk media record could not be found during system startup, or if it failed completely during system operation.

`priv_bpath`  
The path to the block device containing the private region for the disk. The private region

is the region of the disk used to allocate configuration information, disk headers, and other on-disk structures that are used by the Volume Manager. None of the private region can be allocated to volumes or used for uses outside of the Volume Manager.

This field is empty if the disk media record is in the removed or nodarec state, or if the disk does not have a private region (for example, on `nopriv`-type disks).

`priv_cpath`

The path to the character (raw) device containing the private region for the disk. This field is empty if the disk media record is in the removed or nodarec state, or if the disk does not have a private region (for example, on nopriv-type disks).

priv\_offset, priv\_len

Length number fields. These define the offset and length of the private region of the disk within the c2t1d0s2 device containing the private region. These fields are not defined (print as empty strings) if the disk media record is in the removed or nodarec state, or if the disk does not have a private region (for example, on nopriv-type disks).

private A boolean. This is on if the disk underlying the disk media record has a private region (for example, if the disk is not a nopriv-type disk). This is off if the disk media record is in the removed or nodarec state.

pub\_bdev (alias: pub\_dev, bdev, or dev)

priv\_bdev (alias: priv\_dev) The block device numbers of c2t1d0s2 containing the public and private regions of the disk, respectively. These fields are both NODEV if the disk media record is in the removed or nodarec state. The priv\_bdev field is NODEV if the disk does not have a private region (for example, on nopriv-type disks).

pub\_bpath (alias: bpath or path)

The path to the block device containing the public region for the disk. The public region is the region of the disk from which subdisks are allocated. The public region and the private region of a disk are often defined on different partitions, so different device nodes must be used to access

them. This field is empty if the disk media record is in the removed or nodarec state.

pub\_cpath (alias: cpath)

The path to the character (raw) device containing

the public region for the disk. This field is empty if the disk media record is in the removed or nodarec state.

pub\_offset, pub\_len (alias: len)

Length number fields. These define the offset and length of the public region of the disk within the device c2tld0s2 containing the public region. These fields are not defined (print as empty strings) if the disk media record is in the removed or nodarec state.

putil0, putil1, putil2

Comment-type fields of up to fourteen characters. These are the permanent utility fields. The first field is reserved for internal Volume Manager purposes. The second field is reserved for utilities, such as the Storage Administrator, that use the usage-type utilities and that need to store their own state into records. The third field is reserved for local installations.

removed A boolean. This is set to on if the disk is detached from its backing store using the vx dg -k rmdisk operation. Detaching the disk also clears the diskid f

reserved A boolean. This can be set or cleared (see vxedit(1M)) to reserve or unreserve the disk. Reserving a disk prevents vxassist from allocating subdisks from the disk, unless the disk is specified explicitly on the vxassist command line.

rid The record ID for the disk media record.

sd\_num (alias: nsd)

A simple number field. This is the number of subdisks that are defined on the disk media record.

tutil0, tutil1, tutil2

Comment-type fields of up to fourteen characters. These are the temporary utility fields. The first field is reserved for internal Volume Manager purposes. The second field is reserved for utilities, such as the Storage Administrator, that use the usage-type utilities and that need to store their own state into records. The third field is

reserved for local installations.

update\_tid

The configuration update sequence number for the last transaction to update this record.

volatile

A boolean. This is on if the disk access record underlying the disk media record was created with the volatile flag set. See voldisk(1M) for information on how to create a nopriv-type disk with the volatile flag set.

#### Disk Group Record Fields

These are the fields that can be used for specifying disk group records (Note: disk group records cannot be created using vxmake):

comment

A comment-type field of up to 40 characters. This is reserved for use by local installations.

dgid

The disk group ID of the disk group. This is a 64-byte string that is assigned to the disk group when it is created. The form of the string is chosen to substantially eliminate the chance that any two disk groups ever created will have the same disk group ID.

import\_id

A sequence number associated with the current import of the disk group. Each time a disk group is imported, it is assigned a sequence number from an increasing sequence stored in the rootdg disk group. This number can be saved and then used to find out whether a disk group has been deported and then re-imported again.

putil0, putil1, putil2

Comment-type fields of up to fourteen characters. These are the permanent utility fields. The first field is reserved for internal Volume Manager purposes. The second field is reserved for utilities, such as the Storage Administrator, that use the usage-type utilities and that need to store their own state into records. The third field is reserved for local installations.

real\_name

This is the name of the disk group as it is defined within the disk group. Future releases of the Volume Manager may support a temporary aliasing feature, where a disk group is imported using a disk group name that differs from the name as it appears in the disk group configuration itself. Normally, this will be the same as the disk



## File Formats

vxmake(4)

imported disk name.

real\_name will always match the import name of the disk group.

rid           The record ID for the disk group record.

tutil0, tutil1, tutil2

Comment-type fields of up to fourteen characters. These are the temporary utility fields. The first field is reserved for internal Volume Manager purposes. The second field is reserved for utilities, such as the Storage Administrator, that use the usage-type utilities and that need to store their own state into records. The third field is reserved for local installations.

update\_tid

The configuration update sequence number for the last transaction to update this record.

## SEE ALSO

awk(1), chmod(1), sh(1), vxassist(1M), vxdg(1M), vxdisk(1M), vxedit(1M), vxintro(1M), vxprint(1M), passwd(4), group(4), volrec(4)

===[ man7/vxconfig.7 ]=====

Device and Network Interfaces

vxconfig(7)

NAME

vxconfig - volume management configuration device

SYNOPSIS

/dev/vx/config

DESCRIPTION

The vxconfig device is the interface through which all changes to the volume driver state are performed. This device can only be opened by one process at a time. The initial volume configuration is downloaded into the kernel through this device. The vxconfigd process updates the configuration through this device. The vxconfigd process was designed to be the sole and exclusive owner of this device. As such, the ioctl interface between vxconfigd and the device is considered an internal interface and is not documented here.

FILES

/dev/vx/config

SEE ALSO

vxconfigd(1M), vxintro(1M), vxintro(3X), vol(7), vxiod(7), vxtrace(7)

=== [ man7/vxdmp.7 ] =====

Device and Network Interfaces

vxdmp(7)

NAME

vxdmp - Volume Manager multipathing disk devices

DESCRIPTION

Multipathed disk devices in the Volume Manager represent virtual devices with one or more physical access paths to a physical disk. The Multipathing feature provides disk access reliability by dynamically switching to another physical path if there is a path failure.

Every disk that is visible to Volume Manager is represented by a multipath device node in the /dev/vx/dmp and /dev/vx/rdmp directories for block type and character type interfaces. A simple disk is represented by a multipath device having one physical path, whereas a disk that is part of a multipathed disk array is represented by a multipath device with many device access paths depending upon the type of the disk array.

Volume Manager uses standard disk device naming conventions to name multipath devices. For example, if there is a disk device with two physical access paths such as, c1t0d0 and c2t0d0, the multipathing device node for this particular disk is created under /dev/vx/dmp and /dev/vx/rdmp as c1t0d0 or c2t0d0.

To provide multipathing, Volume Manager maintains an internal database of all the multipathed devices detected by the system and their respective paths. The database is created at system boot time. When disks are added or removed from the system, Volume Manager provides support to reconfigure the multipath device database to reflect new states of the hardware, while the system is online. To reconfigure the DMP database, run `ioscan` followed by `insf` and invoke `vxctl enable` (see `vxctl(1M)`).

The various access paths to a multipathed device can be seen by using a `vxdisk` command (see `vxdisk(1M)`). These commands open the multipathing configuration device `/dev/vx/dmpconfig` to obtain the required information about the various paths to a multipathed device and their states.

Volume Manager also provides the facility to enable or disable a disk controller. This feature can be used to stop all I/Os through a particular controller to perform maintenance operations on it. Use the `vxmpadm` command to do this and also perform other operations on the multipath device database.

Volume Manager can provide notification about events that occur in the DMP database due to a change in hardware. See the `vxnotify(1M)` man page for more information.

## Device and Network Interfaces

`vxvmp(7)`

Volume Manager is capable of multipathing both Active/Active and Active/Passive disk arrays.

### IOCTLS

The `ioctl` commands supported by the multipathing device interface are:

#### `DMP_GET_SUBPATHS`

Find the various subpaths of a multipathing device.

#### `DMP_GET_MULTIPATH`

Find the multipathing device for a given subpath.

#### `DMP_CHANGE_CTLR_STATE`

Change the state of a controller.

DMP\_GET\_CTLRS

List all disk controllers on the system.

DMP\_GET\_PATHS\_THRU\_CTLR

Find all paths going through a controller.

FILES

/dev/vx/dmp                      Multipathing block devices.  
  
/dev/vx/rdmp                    Multipathing character dev-  
ices.

NOTES

Do not invoke the ioctls directly. An API for all the functionality is available in the Volume Manager developer's package.

SEE ALSO

vxdctl(1M), vxdisk(1M), vxdmpadm(1M), vxnotify(1M), ioctl(2)

===[ man7/vxinfo.7 ]=====

Device and Network Interfaces

vxinfo(7)

NAME

vxinfo - Volume Manager general information device

SYNOPSIS

/dev/vx/info

## DESCRIPTION

The vxinfo device gathers performance statistics from the kernel for a supplied object. Statistics can be retrieved for volume, plex, subdisk, and disk media objects. Statistics can also be cleared through this interface.

## IOCTLS

The format for calling each ioctl command is:

```
#include <sys/types.h>
#include <vvxvm/volstats.h>
struct tag arg;
int ioctl (int fd, int cmd, struct tag *arg);
```

The value of cmd is the ioctl command code, and arg is usually a pointer to a structure containing the arguments that need to be passed to the kernel.

The return value for all the ioctls, with some exceptions, is 0 if the command was successful, and -1 if it was rejected. If the return value is -1, errno is set to indicate the cause of the error.

The following ioctl commands are supported:

### GET\_DAEMON

This ioctl returns the process ID of the process with the dev/vx/config device open, or 0 if the dev/vx/config device is closed. The value of arg is undefined and should be NULL.

### GET\_VOLINFO

This command accepts a pointer to a volinfo structure as an argument. It fills in the volinfo structure with the corresponding values from the kernel. The members of a volinfo structure are:

number*/	long	version;	/*kernel version
volprivmem area*/	long	max_volprivmem;	/*max size of
number*/	major_t	volbmajor;	/*volume blk dev major
number*/	major_t	volcmajor;	/*volume char dev major
number*/	major_t	plexmajor;	/*plex device major
supported*/	long	maxvol;	/*max # of volumes
plexes*/	long	maxplex;	/*max # of associated
volume*/	long	plexnum;	/*max plexes per
plex*/	long	sdnum;	/*max subdisks per

## Device and Network Interfaces

vxinfo(7)

```

data*/          long      max_ioctl;          /*max size of ioctl
op*/           long      max_specio;         /*max size of ioctl I/O
operation*/    long      max_io;           /*max size of I/O
allowed*/      long      vol_maxkiocount;    /*max # top level I/Os
utils*/       long      dflt_iodelay;       /*default I/O delay for
allowed*/      long      max_parallelio;     /*max # voldios
regions"/     long      voldrl_min_regionsz; /*min DRL region size*/
              long      voldrl_max_drtregs; /*max # of DRL dirty
volume*/      long      vol_is_root;        /*if set, root is
region size*/ long      mvrmaxround;        /*max round-robin
system*/      long      prom_version;      /*PROM version of the
buffer*/      long      vol_maxstablebufsize; /*max size of copy
buf spc*/     size_t    voliot_iobuf_limit; /*max total I/O trace
buffer*/      size_t    voliot_iobuf_max;   /*max size of I/O trace
buf size*/    size_t    voliot_iobuf_default; /*default I/O trace
buf size*/    size_t    voliot_errbuf_default; /*default error trace
channels*/    long      voliot_max_open;    /*max # of trace
size*/       size_t    vol_checkpoint_default; /*default checkpoint
RSRs*/       long      volraid_rsrtransmax; /*max # of transient

```

## VOL\_CLR\_IO\_STATS

This command accepts a pointer to a `vol_io_stats_list` structure as an argument and will clear the kernel statistics for the identified objects.

## VOL\_GET\_IO\_STATS

This command accepts a pointer to a `vol_io_stats_list` structure as an argument so that a single `ioctl` can operate on any reasonable number of objects. The `ioctl`

then fills in the performance statistics for the objects whose IDs appear in the structure.

The members of the `vol_io_stats_list` structure are:

```
    size_t    vsl_nobject;          /*# of objects*/
    struct    vol_io_stats          /*vsl_stats;*/
```

The members of the `vol_io_stats` structure are:

```
    volrid_t  vs_dg_iid;           /*disk group rid*/
    volrid_t  vs_rid;             /*vs_rid; *rid of
object*/
    int       vs_object_type;     /*object type*/
    size_t    vs_stats_size;     /*size of stats
buffer*/
    structure vol_stats          /*vs_stats; *stats
buffer*/
```

The `vs_dg_iid` field identifies the disk group in which the `vs_rid` object exists. The `vs_rid` value can be obtained from either a `volrec`, `plexrec`, or `sdrec` structure. See the `volrec(4)`, `plexrec(4)`, and `sdrec(4)` man pages.

Most statistics are formed as a concatenation of the

## Device and Network Interfaces

`vxinfo(7)`

`vol_stats` structure. The members of the `vol_stats` structure are:

```
    size_t    vs_op;              /*count of I/O
operations*/
    size_t    vs_fail_op;        /*count of failed
I/Os*/
    size_t    vs_blocks;         /*total blocks of I/O*/
    time_t    vs_time;          /*accumulated I/O
times*/
```

The `vol_io_stats_list`, `vol_io_stats`, and `vol_stats` structures are defined in the `</usr/include/vxvm/volstats.h>` header file.

## DIAGNOSTICS

The following errors are returned by the `VOL_GET_IO_STATS` and `VOL_CLR_IO_STATS` ioctls:



EFAULT            A pointer passed to the kernel was invalid,  
                  causing a bad memory reference.

ENOENT            The supplied vs\_dg\_iid does not match any  
                  existing imported disk groups, or the sup-  
                  plied vs\_rid object could not be found.

FILES

  /dev/vx/info                            Volume Manager info device.

SEE ALSO

  vxintro(1M), ioctl(2), plexrec(4), sdrec(4), volrec(4),  
  vxconfig(7), vxio(7)

===[ man7/vxio.7 ]=====

Device and Network Interfaces

vxio(7)

NAME

  vxio - Volume Manager virtual disk devices

DESCRIPTION

Volume devices are the virtual disk devices for Volume Manager. The volume devices support a virtual disk access method with disk mirroring and disk striping. A volume is a logical entity composed of one or more plexes. A read can be satisfied from any plex, while a write is directed to all plexes. The virtual disk devices have a wide variety of behaviors, which are programmable through the /dev/vx/config device. For volume devices, both block- and character-special devices are implemented.

Each plex in the volume is a copy of the volume address space. The plex has subdisks associated with it. These subdisks provide backup storage for the volume address space.

It is possible to create a sparse plex, which is a plex without backup storage for some of the volume address space. The areas of a sparse plex that do not have a backing subdisk are called holes. An attempt to read a hole in a sparse plex fails. If there are other plexes that have backup storage for the read, then one of those plexes is read. Otherwise, the read fails. A write to a hole in a sparse plex is considered a success even though the data can't be read back.

In addition, a plex may be designated as a logging plex. This means that a log of blocks that are in transition will be kept, which enables fast recovery after a system failure. This feature is known as DRL, or dirty region logging. The log for each plex consists of a specially designated subdisk that is not part of the normal plex address space.

#### IOCTLS

The ioctl commands supported by the volume virtual disk device interface are discussed later in this section. The only ioctls supported for plex devices are GET\_DAEMON and GET\_VOLINFO. The format for calling each ioctl command is:

```
#include <sys/types.h>
#include <sys/volclient.h>
struct tag arg;
int ioctl (int fd, int cmd, struct tag *arg);
```

The value of cmd is the ioctl command code and arg is usually a pointer to a structure containing the arguments that need to be passed to the kernel.

The return value for all these ioctls, with some exceptions, is zero if the command was successful, and -1 if it was rejected. If the return value is -1, then errno is set to indicate the cause of the error.

The following ioctl commands are supported:

#### GET\_DAEMON

This ioctl returns the pid of the process with the /dev/vx/config device open, or 0 if it is closed. The value of arg is undefined and should be NULL.

#### GET\_VOLINFO

This command accepts a pointer to a volinfo structure as an argument. It fills in the volinfo structure with the corresponding values from the kernel. The members of a volinfo structure are:

number*/	long	version;	/*kernel version
volprivmem area*/	long	max_volprivmem;	/*max size of
number*/	major_t	volbmajor;	/*volume blk dev major
number*/	major_t	volcmajor;	/*volume char dev major
number*/	major_t	plexmajor;	/*plex device major
supported*/	long	maxvol;	/*max # of volumes
plexes*/	long	maxplex;	/*max # of associated
volume*/	long	plexnum;	/*max plexes per
plex*/	long	sdnum;	/*max subdisks per
data*/	long	max_ioctl;	/*max size of ioctl
op*/	long	max_specio;	/*max size of ioctl I/O
operation*/	long	max_io;	/*max size of I/O
allowed*/	long	vol_maxkiocount;	/*max # top level I/Os
utils*/	long	dflt_iodelay;	/*default I/O delay for
allowed*/	long	max_parallel_io;	/*max # voldios
regions*/	long	voldrl_min_regionsz;	/*min DRL region size*/
	long	voldrl_max_drtregs;	/*max # of DRL dirty
volume*/	long	vol_is_root;	/*if set, root is
	long	mvrmaxround;	/*max round-robin

```

region size*/
system*/      long      prom_version;          /*PROM version of the
buffer*/      long      vol_maxstablebufsize; /*max size of copy
buf spc*/     size_t    voliot_iobuf_limit; /*max total I/O trace
buffer*/     size_t    voliot_iobuf_max;   /*max size of I/O trace
buf size*/   size_t    voliot_iobuf_default; /*default I/O trace
buf size*/   size_t    voliot_errbuf_default; /*default error trace
channels*/   long      voliot_max_open;     /*max # of trace
size*/       size_t    vol_checkpt_default; /*default checkpoint
RSRs*/       long      volraid_rsrtransmax; /*max # of transient

```

#### PLEX\_DETACH

This command is used to force a plex to be detached from a volume. The name of the plex is passed in as an argument. The volume is the volume device against which the ioctl is being performed.

VxVM 3.0

Last change: 26 Mar 1999

2

Device and Network Interfaces

vxio(7)

#### VOL\_LOG\_WRITE

This command forces a dirty region log to be flushed to disk. This is used by the vxconfigd process to flush an initial log to disk before starting the volume.

#### VOL\_READ, VOL\_WRITE

These commands provide a mechanism by which I/O can be issued to volumes larger than 2 gigabytes in length. Current UNIX read and write system calls on 32-bit processors limit sizes to 2 gigabytes because of the signed byte-offset value used to perform the I/O. These ioctl commands provide a method of providing sector offsets to an I/O and raise the limit to one sector less than 1024 gigabytes.

The required I/O is identified to the command by the use of a vol\_rdwr structure containing the following:

```

ulong_t      vrw_flags;          /*flags*/

```

```

                voff_t    vrw_off;           /*offset in volume
(sectors)*/
                size_t    vrw_size;         /*number of sectors to
Xfer*/
                caddr_t   vrw_addr;        /*user address for
Xfer*/

```

The vrw\_flags field is currently unused; other fields are explained in the comments.

#### VOLUME-SPECIAL IOCTLs

The ATOMIC\_COPY, VERIFY\_READ, and VERIFY\_WRITE ioctls perform special I/O operations against the volume. They use the vol\_io structure to initiate I/O requests and receive the status information back. The members of the vol\_io structure are:

```

                voff_t    vi_offset;        /*0x00 offset on plex*/
                size_t    vi_len;          /*0x04 amount of data to
read/write*/
                caddr_t   vi_buf;          /*0x08 ptr to buffer*/
                size_t    vi_nsrcplex;     /*0x0c number of source
plexes*/
                size_t    vi_ndestplex;    /*0x10 number of destination
plexes*/
                struct plx_ent              *vi_plexptr; /*0x14 ptr to
array of plex entries*/
                ulong_t   vi_flag;        /*0x18 flags associated with
op*/

```

The members of the plx\_ent structure are:

```

                char      pe_name[NAME_SZ]; /*name of plex*/
                int       pe_errno;        /*error number against
plex*/

```

The vi\_offset value specifies the sector offset of the I/O within the volume. It must be within the address range of the volume. Also, the entire range of the I/O from vi\_offset to vi\_offset + vi\_len must be within the address range of the volume.

The vi\_len field specifies the length of the I/O in sectors. It must be a between 0 and 120 sectors (VOL\_MAXSPECIALIO).

The `vi_buf` field is a pointer to a buffer of `vi_len` sectors. The `VERIFY_WRITE` ioctl writes the data stored in this buffer.

The `vi_nsrcplex` field is the number of source plexes available for the operation and the `vi_ndestplex` field is the number of destination plexes available for the operation. The `vi_nsrcplex` and `vi_ndestplex` values must be between 0 and 8 (`PLEX_NUM`).

The `vi_plexptr` is a pointer to an array of `plx_ent` structures. The first `vi_nsrcplex` entries in the array are source plexes. The `pe_name` contains the name of the plex. If the name of the first source entry is the null string, then the kernel selects all plexes available for reading as part of the volume and fills in the `pe_name` fields.

After the source plexes, the next `vi_ndestplex` entries are the destination plexes. If the name of the first destination entry is the null string, then the kernel selects all plexes available for writing as part of the volume and fills in the `pe_name` fields.

After the I/O operations are performed, the `plx_ent` structures are copied back to the user. If the kernel selected the plexes, the names of the selected plexes are in the `pe_name` fields. The status of the operations on each plex are stored in the `pe_errno` field of the `plx_ent` structure for that plex.

The `pe_errno` field is 0 if the operation succeeded against the plex. If `pe_errno` isn't 0, then the error code indicates what happened to the plex. The possible values for `pe_errno` are:

<code>EACCES</code>	The specified plex is in the disabled state, so no I/O can be performed against it.
<code>EFAULT</code>	A source plex is sparse and doesn't have blocks that map the entire I/O request.
<code>EIO</code>	A read I/O error was returned against a source plex.
<code>ENOENT</code>	The specified plex isn't associated with the volume the ioctl was issued against.
<code>ENXIO</code>	An error was detected in the operation, so no I/O operation was attempted to the plex. If

one `plx_ent` structure in a list contains a bad name, then no I/O is done. All `plx_ent` structures with a valid name have their `pe_errno` set to `ENXIO` to indicate no I/O was attempted.

- `EROFS` A write I/O error was returned against a destination plex.
- `ESRCH` The `VERIFY_READ` and `VERIFY_WRITE` operations compare the data from different plexes against each other to verify the consistency. If the comparison detects an error, the plex that was read first is considered correct. An `ESRCH` error is returned against the plex that was read second to indicate it contains bad data.

For the `ATOMIC_COPY`, `VERIFY_READ`, and `VERIFY_WRITE` ioctls, if the entire operation is a success, then a 0 is returned. If there is a fatal error, a -1 is returned and the external variable `errno` indicates the reason for failure. If a 1 is returned, there was some sort of failure and the net results must be determined by examining the `pe_errno` fields of all the plexes.

The ioctls that do volume-special I/O are:

#### `ATOMIC_COPY`

This ioctl takes a pointer to a `vol_io` structure as an argument. It reads `vi_len` sectors of data, at offset `vi_offset`, from one of the plexes specified by the first `vi_nsrcplex` plex entries into a buffer. Then it writes the contents of the buffer onto all the plexes specified by the next `vi_ndestplex` plex entries at the same offset. This entire operation is atomic with respect to the I/O stream of the volume. The `vi_buf` field is unused and should be `NULL`.

If the first source plex entry has a null string for the name, the kernel selects from any plex of the volume that is enabled for read access. The names of any selected plexes are copied into the appropriate plex entries.

If the first destination plex entry has a null string for the name, the kernel will write to all plexes of the volume that are enabled for write access. The names of selected plexes are copied into the appropriate plex entries.

When the list of source plexes has been compiled, the

## Device and Network Interfaces

vxio(7)

kernel tries to read each plex in order. A plex can't be read if it doesn't have backup storage covering the entire operation. Once a plex has been successfully read, all the destination plexes are written. The writes can succeed even if the destination plex is sparse and doesn't have backup storage to cover the entire write.

If the ioctl returns a value of -1, then some error has occurred which prevented the ATOMIC\_COPY from working. If the return value is 0, then everything worked fine. If the return value is 1, then the pe\_errno field must be examined to determine the errors on each individual plex. The status of the overall operation depends on these individual errors.

## VERIFY\_READ

This command accepts a pointer to a vol\_io structure as an argument. It reads vi\_len sectors, from offset vi\_offset, on the first vi\_nsrcplex plexes specified by the vi\_plexptr array of plex entries. It compares the data from each plex. If any plex is different from the previous plexes read, the pe\_errno value for that plex is set to ESRCH. This entire operation is atomic with respect to the I/O stream of the volume.

The vi\_ndestplex field must be 0. If the vi\_buf field is not NULL, then the data read from the plexes not marked with ESRCH is copied into the buffer specified by vi\_buf.

As each plex is read, any data that has already been read is compared against the previous reads. If all the data for the plex passes, then any data that was read for the first time is copied into the comparison buffer. This allows VERIFY\_READ operations to work if the volume contains sparse plexes. The data not represented by backup storage is not compared against anything.

## VERIFY\_WRITE

This command accepts a pointer to a vol\_io structure as an argument. It writes vi\_len sectors to offset vi\_offset on the first vi\_ndestplex plexes specified by



the `vi_plexptr` array of plex entries. The data to be written is stored in the buffer pointed to by `vi_buf`. Then, the data is read back from each plex that was successfully written and is compared against the data written. If the data from any plex doesn't match the data written, the `pe_errno` value for the plex is set to `ESRCH`. This entire operation is atomic with respect to the I/O stream of the volume.

## Device and Network Interfaces

`vxio(7)`

The `vi_nsrcplex` field must be 0.

As each plex is read, any data that doesn't have backup storage on that plex is filled in from the write buffer. This allows `VERIFY_WRITE` operations to work if the volume contains sparse plexes. The data not represented by backup storage will always succeed.

### `VOL_LOG_WRITE`

This command takes no argument, and causes the log for a volume to be written to disk immediately. This command is useful for making sure that the on-disk images of the log have been written.

This command returns -1 with `errno` set to `EINVAL` if the specified volume does not have logging enabled.

### `PLEX_DETACH`

This command allows an enabled plex to be detached. The argument is the name of the plex to detach.

## DIAGNOSTICS

The following errors are returned by the volume and plex virtual disk device interfaces:

<code>EAGAIN</code>	A needed kernel resource couldn't be obtained.
<code>EBADF</code>	An attempt was made to write a volume that wasn't opened for writing or read a volume that wasn't open for reading.
<code>EFAULT</code>	A pointer passed to the kernel was invalid, causing a bad memory reference.
<code>EINVAL</code>	Invalid data was passed to the kernel. Some

field in a vol\_io structure failed a sanity check.

EIO A physical I/O error occurred during an operation. If this error is returned, the driver tried an I/O and it failed.

EMFILE The kernel was asked to supply a list of destination plexes for an ATOMIC\_COPY or VERIFY\_WRITE ioctl. If more than vi\_ndestplex enabled plexes available in write mode are found, an EMFILE error is returned.

ENFILE The kernel was asked to supply a list of source plexes for an ATOMIC\_COPY or

VxVM 3.0

Last change: 26 Mar 1999

7

Device and Network Interfaces

vxio(7)

VERIFY\_READ ioctl. If more than vi\_nsrcplex enabled plexes available in read mode are found, an ENFILE error is returned.

ENOENT An object named in GET\_VOL\_STATS ioctl was not associated with the volume.

ENOENT The kernel was asked to supply a list of destination plexes for an ATOMIC\_COPY or VERIFY\_WRITE ioctl. If no enabled plexes available in write mode are found, an ENOENT error is returned.

ENXIO A validation error occurred during an I/O operation. If this error is returned, the driver attempted no I/O.

ESRCH The kernel was asked to supply a list of source plexes for an ATOMIC\_COPY or VERIFY\_READ ioctl. If no enabled plexes available in read mode are found, an ESRCH error is returned.

FILES

/dev/vx/dsk/\*... Volume block device files.

/dev/vx/rdisk/\*... Volume character (raw) device files.

SEE ALSO

ioctl(2), vxintro(3X), vxconfig(7), vxiod(7), vxtrace(7)

VxVM 3.0

Last change: 26 Mar 1999

8

===[ man7/vxiod.7 ]=====

Device and Network Interfaces

vxiod(7)

NAME

vxiod - control Volume Manager I/O daemon process

DESCRIPTION

The vxiod device is used to control the number of volume I/O daemons active on the system. A process context is necessary to implement the plex consistency recovery and writeback error handling policies for multi-plex volumes, and for continuing normal I/O after a log write if the volume has logging enabled. It is also required for the plex recovery performed with a mirrored volume in the read/writeback mode.

There are three aspects of I/O daemon operations:

- o General I/O
- o Error handling

- o Log handling

I/O handling is achieved by an ioctl command that does not return, but instead calls the vxiod routine to wait for errors or I/O requests and process them. When an error occurs, if there are no I/O daemons active, the I/O simply turns into a failure on that plex. If a general I/O request is queued up when no daemons exist, then the I/O will hang forever until a daemon process is created. If I/O daemons are active, then the I/O is put on a work queue and the daemons are awakened. A daemon takes an error request and tries to read other plexes until a read succeeds or all plexes have been tried. Then, if the writeback facility is enabled, the daemon tries to write the good data to each plex that failed on the read. If the write is successful, the read error is nullified. An I/O request is handled in a similar manner.

Logging is handled in a similar manner. An ioctl command, which does not return, is issued to create a daemon for each volume which has logging enabled. This daemon monitors two queues: one queue of I/O which was started while the log was busy (the ``log'' queue), and another queue of requests which have been logged and now need to be started (the ``ready'' queue). I/O requests are taken from the log queue when the log is no longer busy, and another log write is started. Completion of a log write results in all I/O requests which have just been logged being placed on the I/O daemon's ready queue where they are immediately started.

One mechanism finds out how many I/O error daemons are running, and another mechanism allows a process to become an I/O daemon. Before a process becomes an I/O daemon, it should close all open files and detach from the controlling

tty. An I/O or logging daemon cannot be killed except through an explicit ioctl.

#### IOCTLS

The format for calling each ioctl command is:

```
#include <sys/types.h>
#include <sys/vol.h>
struct tag arg;
```

```
int ioctl (int fd, int cmd, struct tag arg);
```

The value of `cmd` is the `ioctl` command code, and `arg` is usually a pointer to a structure containing the arguments that need to be passed to the kernel.

The return value for all these `ioctls` is 0 if the command was successful, and -1 if it was rejected. If the return value is -1, then `errno` is set to indicate the cause of the error.

The following `ioctl` commands are supported:

#### GET\_DAEMON

This `ioctl` returns the `pid` of the process with the `volconfig` device open, or 0 if the `volconfig` device is closed. The value of `arg` is undefined and should be `NULL`.

#### GET\_VOLINFO

This command accepts a pointer to a `volinfo` structure as an argument. It fills in the `volinfo` structure with the corresponding values from the kernel. The members of the `volinfo` structure are:

<code>number*/</code>	<code>long</code>	<code>version;</code>	<code>/*kernel version</code>
<code>area*/</code>	<code>long</code>	<code>voldinfo_sz;</code>	<code>/*size of voldinfo</code>
<code>number*/</code>	<code>major_t</code>	<code>volbmajor;</code>	<code>/*volume blk dev major</code>
<code>number*/</code>	<code>major_t</code>	<code>volcmajor;</code>	<code>/*volume char dev major</code>
<code>number*/</code>	<code>major_t</code>	<code>plexmajor;</code>	<code>/*plex device major</code>
<code>supported*/</code>	<code>long</code>	<code>maxvol;</code>	<code>/*max # of volumes</code>
<code>plexes*/</code>	<code>long</code>	<code>maxplex;</code>	<code>/*max # of associated</code>
<code>volume*/</code>	<code>long</code>	<code>plexnum;</code>	<code>/*max plexes per</code>
<code>plex*/</code>	<code>long</code>	<code>sdnum;</code>	<code>/*max subdisks per</code>
<code>data*/</code>	<code>long</code>	<code>max_ioctl;</code>	<code>/*max size of ioctl</code>
<code>op*/</code>	<code>long</code>	<code>max_specio;</code>	<code>/*max size of ioctl I"O</code>
<code>operation*/</code>	<code>long</code>	<code>max_io;</code>	<code>/*max size of I"O</code>
<code>utils*/</code>	<code>long</code>	<code>max_ilocksleep;</code>	<code>/*max sleeps on ilock*/</code>
<code>to keep*/</code>	<code>long</code>	<code>dflt_iodelay;</code>	<code>/*default I"O delay for</code>
<code>to keep*/</code>	<code>long</code>	<code>max_freebuf;</code>	<code>/*max number of volbufs</code>
<code>to keep*/</code>	<code>long</code>	<code>max_freevcb;</code>	<code>/*max number of volcbs</code>
<code>to keep*/</code>	<code>long</code>	<code>max_parallel_io;</code>	<code>/*max # voldios</code>

allowed\*/

VxVM 3.0

Last change: 26 Mar 1999

2

Device and Network Interfaces

vxiod(7)

#### VOL\_IO\_DAEMON

This command turns the process into an I/O daemon, for the purpose of error handling. This process is killed on receipt of a non-ignored signal.

#### VOL\_IO\_DAEMON\_CNT

This ioctl returns the number of I/O daemons currently running for the purpose of error handling. The value of arg is undefined and should be NULL.

#### VOL\_IO\_DAEMON\_KILL

This command causes the passed argument to be the maximum number of active I/O daemon processes to be left alive. All others will be killed.

#### VOL\_LOG\_IO\_DAEMON

This command creates an I/O daemon for the purpose of log handling for the indicated volume. The argument is the device number (as returned by stat) of the volume for which logging is handled. Only one such process can exist for each volume. A logging daemon can be created for a volume which does not yet exist, or for which logging is not used. In this case, the logging daemon is unused until the volume is created or logging is enabled.

This command is killed when the logging daemon is stopped explicitly by the VOL\_LOGIOD\_KILL ioctl.

#### VOL\_LOGIOD\_CHECK

Returns a positive number if a volume has a logging daemon running against it; otherwise, returns 0. The argument is the device number of the volume.

#### VOL\_LOGIOD\_KILL

Causes the logging daemon for a volume to exit or return (see VOL\_LOG\_IO\_DAEMON). The argument is the device number of the volume. This ioctl does not return until the logging daemon has exited, or until it is interrupted by a signal.

#### VOL\_LOGIOD\_WAIT

Sleeps until a logging daemon exists for a volume, then

returns 0. The argument is the device number of the volume. The volume does not need to exist yet. This ioctl allows the parent of a caller to VOL\_LOG\_IO\_DAEMON to wait until the logging daemon is known to exist. This ioctl is interruptible.

FILES

/dev/vxiod

VxVM 3.0

Last change: 26 Mar 1999

3

Device and Network Interfaces

vxiod(7)

SEE ALSO

vxiod(1M), volinfo(1M), ioctl(2), volinfo(3X), volconfig(7), volevent(7), vol(7)

=== [ man7/vxtrace.7 ] =====

Device and Network Interfaces

vxtrace(7)

#### NAME

vxtrace - Volume Manager I/O Tracing Device

#### SYNOPSIS

/dev/vx/trace

#### DESCRIPTION

The vxtrace device implements the Volume Manager I/O tracing and the error tracing. An I/O tracing interface is available that users or processes can use to get a trace of I/Os for specified sets of kernel objects. Each separate user of the I/O tracing interface can specify the set of desired trace data independent of all other users. I/O events include regular read and write operations, special I/O operations (ioctl's), as well as special recovery operations (for example, recovery reads). A special tracing mechanism exists for getting error trace data. The error tracing mechanism is independent of any I/O tracing and is always enabled for all pertinent kernel I/O objects. It is possible for a process to get both a set of saved errors and to wait for new errors.

#### IOCTLS

The format for calling each ioctl command is:



```
#include <sys/types.h>
#include <vxvm/voltrace.h>
struct tag arg;
int ioctl (int fd, int cmd, struct tag arg);
```

The first argument `fd` is a file descriptor which is returned from opening the `/dev/vx/trace` device. Each tracing device opened is a cloned device which can be used as a private kernel trace channel. The value of `cmd` is the `ioctl` command code, and `arg` is usually a pointer to a structure containing the arguments that need to be passed to the kernel.

The return value for all these `ioctl`s is 0 if the command was successful, and -1 if it was rejected. If the return value is -1, `errno` is set to indicate the cause of the error.

The following `ioctl` commands are supported:

`VOLIOT_ERROR_TRACE_INIT`

This command accepts no argument. The `VOLIOT_ERROR_TRACE` `ioctl` initializes a kernel trace channel to return error trace data. The trace channel will be initialized to return any previously accumulated error trace data that has not yet been discarded. The accumulated trace data can be skipped by issuing `VOLIOT_DISCARD` on the channel. This call can be issued

on a trace channel that was previously initialized either for error tracing or for regular I/O tracing. In this case, the channel is effectively closed down and then reinitialized as described above. To get the error trace data, issue the `read(2)` system call. The error trace data consists of a set of variable length trace event records. The first byte of each record indicates the length, in bytes, of the entire record (including the length byte), the second byte indicates the type of the entry (which can be used to determine the format of the entry). Each call to `read()` returns an integral number of trace event records, not to exceed the number of bytes requested in the `read()` call; the return value from `read()` will be adjusted to the number of bytes of trace data actually returned. If the `O_NONBLOCK` flag is set on the trace channel, and no

trace data is available, EAGAIN will be returned; otherwise, the read will block interruptibly until at least one trace record is available. When some trace data is available, the available unread trace records will be returned, up to the limit specified in the call to read(). If more trace records are available, subsequent reads will return those records.

#### VOLIOT\_IO\_TRACE\_INIT

The VOLIOT\_IO\_TRACE\_INIT ioctl initializes a kernel trace channel to return I/O trace data. This command accepts bufsize as the argument. Initially, no objects are selected for I/O tracing. To select objects to trace, issue the VOLIOT\_IO\_TRACE ioctl. The bufsize argument specifies the kernel buffer size to use for gathering events. A larger size reduces the chance that events are lost due to scheduling delays in the event reading process. A bufsize value of 0 requests a default size which is considered reasonable for the system. The value of bufsize will be silently truncated to a maximum value to avoid extreme use of system resources. A bufsize value of (size\_t)-1 will yield the maximum buffer size.

#### VOLIOT\_IO\_TRACE, VOLIOT\_IO\_UNTRACE

The VOLIOT\_IO\_TRACE and VOLIOT\_IO\_UNTRACE ioctls enable and disable, respectively, I/O tracing for particular sets of objects on an I/O tracing channel. They both accept a voliot\_want\_list structure tracelist as the argument. The tracelist argument specifies object sets. The voliot\_want\_list structure specifies an array of desired object sets. Each object set is identified by a union of structures (the voliot\_want\_set union), each representing different types of object sets. See the declaration of these structures in voltrace.h for more detail.

#### FILES

/dev/vx/trace

#### SEE ALSO

vxintro(1M), vxtrace(1M), vxvol(1M) ioctl(2), read(2),  
vxintro(3X), vxconfig(7), vxiod(7),

