```
**********************************************************
   50491 Fri Mar 25 18:28:30 2016
new/usr/src/man/man1m/zpool.1m
6781 zpool man page needs updated to remove duplicate entry of "cannot be" where
Reviewed by: Toomas Soome <tsoome@me.com>
Reviewed by: Robert Mustacchi <rm@joyent.com>
**********************************************************
```

```
.\"
.\" CDDL HEADER START
.\"
.\" The contents of this file are subject to the terms of the
.\" Common Development and Distribution License (the "License").
.\" You may not use this file except in compliance with the License.
.\"
.\" You can obtain a copy of the license at usr/src/OPENSOLARIS.LICENSE
.\" or http://www.opensolaris.org/os/licensing.
.\" See the License for the specific language governing permissions
.\" and limitations under the License.
.\"
.\" When distributing Covered Code, include this CDDL HEADER in each
.\" file and include the License file at usr/src/OPENSOLARIS.LICENSE.
.\" If applicable, add the following below this CDDL HEADER, with the
.\" fields enclosed by brackets "[]" replaced with your own identifying
.\" information: Portions Copyright [yyyy] [name of copyright owner]
.\"
.\" CDDL HEADER END
.\"
.\"
.\" Copyright (c) 2007, Sun Microsystems, Inc. All Rights Reserved.
.\" Copyright (c) 2013 by Delphix. All rights reserved.
.\" Copyright 2016 Nexenta Systems, Inc.
.\"
.Dd March 25, 2016
.Dd February 15, 2016
.Dt ZPOOL 1M
.Os
.Sh NAME
.Nm zpool
.Nd configure ZFS storage pools
.Sh SYNOPSIS
.Nm
.Fl \?
.Nm
.Cm add
.Op Fl fn
.Ar pool vdev Ns ...
.Nm
.Cm attach
.Op Fl f
.Ar pool device new_device
.Nm
.Cm clear
.Ar pool
.Op Ar device
.Nm
.Cm create
.Op Fl dfn
.Op Fl m Ar mountpoint
.Oo Fl o Ar property Ns = Ns Ar value Oc Ns ...
.Oo Fl O Ar file-system-property Ns = Ns Ar value Oc Ns ...
.Op Fl R Ar root
.Ar pool vdev Ns ...
.Nm
.Cm destroy
.Op Fl f
.Ar pool
```



```
.Nm
.Cm detach
.Ar pool device
.Nm
.Cm export
.Op Fl f
.Ar pool Ns ...
.Nm
.Cm get
.Op Fl Hp
.Op Fl o Ar field Ns Oo , Ns Ar field Oc Ns ...
.Sy all Ns | Ns Ar property Ns Oo , Ns Ar property Oc Ns ...
.Ar pool Ns ...
.Nm
.Cm history
.Op Fl il
.Oo Ar pool Oc Ns ...
.Nm
.Cm import
.Op Fl D
.Op Fl d Ar dir
.Nm
.Cm import
.Fl a
.Op Fl DfmN
.Op Fl F Op Fl n
.Op Fl c Ar cachefile Ns | Ns Fl d Ar dir
.Op Fl o Ar mntopts
.Oo Fl o Ar property Ns = Ns Ar value Oc Ns ...
.Op Fl R Ar root
.Nm
.Cm import
.Op Fl Dfm
.Op Fl F Op Fl n
.Op Fl c Ar cachefile Ns | Ns Fl d Ar dir
.Op Fl o Ar mntopts
.Oo Fl o Ar property Ns = Ns Ar value Oc Ns ...
.Op Fl R Ar root
.Ar pool Ns | Ns Ar id
.Op Ar newpool
.Nm
.Cm iostat
.Op Fl v
.Op Fl T Sy u Ns | Ns Sy d
.Oo Ar pool Oc Ns ...
.Op Ar interval Op Ar count
.Nm
.Cm list
.Op Fl Hpv
.Op Fl o Ar property Ns Oo , Ns Ar property Oc Ns ...
.Op Fl T Sy u Ns | Ns Sy d
.Oo Ar pool Oc Ns ...
.Op Ar interval Op Ar count
.Nm
.Cm offline
.Op Fl t
.Ar pool Ar device Ns ...
.Nm
.Cm online
.Op Fl e
.Ar pool Ar device Ns ...
.Nm
.Cm reguid
.Ar pool
.Nm
.Cm reopen
```

```
.Ar pool
.Nm
.Cm remove
.Ar pool Ar device Ns ...
.Nm
.Cm replace
.Op Fl f
.Ar pool Ar device Op Ar new_device
.Nm
.Cm scrub
.Op Fl s
.Ar pool Ns ...
.Nm
.Cm set
.Ar property Ns = Ns Ar value
.Ar pool
.Nm
.Cm split
.Op Fl n
.Oo Fl o Ar property Ns = Ns Ar value Oc Ns ...
.Op Fl R Ar root
.Ar pool newpool
.Nm
.Cm status
.Op Fl Dvx
.Op Fl T Sy u Ns | Ns Sy d
.Oo Ar pool Oc Ns ...
.Op Ar interval Op Ar count
.Nm
.Cm upgrade
.Nm
.Cm upgrade
.Fl v
.Nm
.Cm upgrade
.Op Fl V Ar version
.Fl a Ns | Ns Ar pool Ns ...
.Sh DESCRIPTION
The
.Nm
command configures ZFS storage pools. A storage pool is a collection of devices
that provides physical storage and data replication for ZFS datasets. All
datasets within a storage pool share the same space. See
.Xr zfs 1M
for information on managing datasets.
.Ss Virtual Devices (vdevs)
A "virtual device" describes a single device or a collection of devices
organized according to certain performance and fault characteristics. The
following virtual devices are supported:
.Bl -tag -width Ds
.It Sy disk
A block device, typically located under
.Pa /dev/dsk .
ZFS can use individual slices or partitions, though the recommended mode of
operation is to use whole disks. A disk can be specified by a full path, or it
can be a shorthand name
.Po the relative portion of the path under
.Pa /dev/dsk
.Pc .
A whole disk can be specified by omitting the slice or partition designation.
For example,
.Pa c0t0d0
is equivalent to
.Pa /dev/dsk/c0t0d0s2 .
When given a whole disk, ZFS automatically labels the disk, if necessary.
.It Sy file
```



```
A regular file. The use of files as a backing store is strongly discouraged. It
is designed primarily for experimental purposes, as the fault tolerance of a
file is only as good as the file system of which it is a part. A file must be
specified by a full path.
.It Sy mirror
A mirror of two or more devices. Data is replicated in an identical fashion
across all components of a mirror. A mirror with N disks of size X can hold X
bytes and can withstand (N-1) devices failing before data integrity is
compromised.
.It Sy raidz , raidz1 , raidz2 , raidz3
A variation on RAID-5 that allows for better distribution of parity and
eliminates the RAID-5
.Qq write hole
.Pq in which data and parity become inconsistent after a power loss .
Data and parity is striped across all disks within a raidz group.
.Pp
A raidz group can have single-, double-, or triple-parity, meaning that the
raidz group can sustain one, two, or three failures, respectively, without
losing any data. The
.Sy raidz1
vdev type specifies a single-parity raidz group; the
.Sy raidz2
vdev type specifies a double-parity raidz group; and the
.Sy raidz3
vdev type specifies a triple-parity raidz group. The
.Sy raidz
vdev type is an alias for
.Sy raidz1 .
.Pp
A raidz group with N disks of size X with P parity disks can hold approximately
(N-P)*X bytes and can withstand P device(s) failing before data integrity is
compromised. The minimum number of devices in a raidz group is one more than
the number of parity disks. The recommended number is between 3 and 9 to help
increase performance.
.It Sy spare
A special pseudo-vdev which keeps track of available hot spares for a pool. For
more information, see the
.Sx Hot Spares
section.
.It Sy log
A separate intent log device. If more than one log device is specified, then
writes are load-balanced between devices. Log devices can be mirrored. However,
raidz vdev types are not supported for the intent log. For more information,
see the
.Sx Intent Log
section.
.It Sy cache
A device used to cache storage pool data. A cache device cannot be configured
as a mirror or raidz group. For more information, see the
*A device used to cache storage pool data. A cache device cannot be cannot be*
*configured as a mirror or raidz group. For more information, see the*
.Sx Cache Devices
section.
.El
.Pp
Virtual devices cannot be nested, so a mirror or raidz virtual device can only
contain files or disks. Mirrors of mirrors
.Pq or other combinations
are not allowed.
.Pp
A pool can have any number of virtual devices at the top of the configuration
.Po known as
.Qq root vdevs
.Pc .
Data is dynamically distributed across all top-level devices to balance data
among devices. As new virtual devices are added, ZFS automatically places data
```

on the newly available devices.
.Pp
Virtual devices are specified one at a time on the command line, separated by whitespace. The keywords
.Sy mirror
and
.Sy raidz
are used to distinguish where a group ends and another begins. For example, the following creates two root vdevs, each a mirror of two disks:
.Bd -literal
zpool create mypool mirror c0t0d0 c0t1d0 mirror c1t0d0 c1t1d0
.Ed
.Ss Device Failure and Recovery
ZFS supports a rich set of mechanisms for handling device failure and data corruption. All metadata and data is checksummed, and ZFS automatically repairs bad data from a good copy when corruption is detected.
.Pp
In order to take advantage of these features, a pool must make use of some form of redundancy, using either mirrored or raidz groups. While ZFS supports running in a non-redundant configuration, where each root vdev is simply a disk or file, this is strongly discouraged. A single case of bit corruption can render some or all of your data unavailable.
.Pp
A pool's health status is described by one of three states: online, degraded, or faulted. An online pool has all devices operating normally. A degraded pool is one in which one or more devices have failed, but the data is still available due to a redundant configuration. A faulted pool has corrupted metadata, or one or more faulted devices, and insufficient replicas to continue functioning.
.Pp
The health of the top-level vdev, such as mirror or raidz device, is potentially impacted by the state of its associated vdevs, or component devices. A top-level vdev or component device is in one of the following states:
.Bl -tag -width "DEGRADED"
.It Sy DEGRADED
One or more top-level vdevs is in the degraded state because one or more component devices are offline. Sufficient replicas exist to continue functioning.
.Pp
One or more component devices is in the degraded or faulted state, but sufficient replicas exist to continue functioning. The underlying conditions are as follows:
.Bl -bullet
.It
The number of checksum errors exceeds acceptable levels and the device is degraded as an indication that something may be wrong. ZFS continues to use the device as necessary.
.It
The number of I/O errors exceeds acceptable levels. The device could not be marked as faulted because there are insufficient replicas to continue functioning.
.El
.It Sy FAULTED
One or more top-level vdevs is in the faulted state because one or more component devices are offline. Insufficient replicas exist to continue functioning.
.Pp
One or more component devices is in the faulted state, and insufficient replicas exist to continue functioning. The underlying conditions are as follows:
.Bl -bullet
.It
The device could be opened, but the contents did not match expected values.
.It
The number of I/O errors exceeds acceptable levels and the device is faulted to

prevent further use of the device.
.El
.It Sy OFFLINE
The device was explicitly taken offline by the
.Nm zpool Cm offline
command.
.It Sy ONLINE
The device is online and functioning.
.It Sy REMOVED
The device was physically removed while the system was running. Device removal detection is hardware-dependent and may not be supported on all platforms.
.It Sy UNAVAIL
The device could not be opened. If a pool is imported when a device was unavailable, then the device will be identified by a unique identifier instead of its path since the path was never correct in the first place.
.El
.Pp
If a device is removed and later re-attached to the system, ZFS attempts to put the device online automatically. Device attach detection is hardware-dependent and might not be supported on all platforms.
.Ss Hot Spares
ZFS allows devices to be associated with pools as
.Qq hot spares .
These devices are not actively used in the pool, but when an active device fails, it is automatically replaced by a hot spare. To create a pool with hot spares, specify a
.Sy spare
vdev with any number of devices. For example,
.Bd -literal
zpool create pool mirror c0d0 c1d0 spare c2d0 c3d0
.Ed
.Pp
Spares can be shared across multiple pools, and can be added with the
.Nm zpool Cm add
command and removed with the
.Nm zpool Cm remove
command. Once a spare replacement is initiated, a new
.Sy spare
vdev is created within the configuration that will remain there until the original device is replaced. At this point, the hot spare becomes available again if another device fails.
.Pp
If a pool has a shared spare that is currently being used, the pool can not be exported since other pools may use this shared spare, which may lead to potential data corruption.
.Pp
An in-progress spare replacement can be cancelled by detaching the hot spare. If the original faulted device is detached, then the hot spare assumes its place in the configuration, and is removed from the spare list of all active pools.
.Pp
Spares cannot replace log devices.
.Ss Intent Log
The ZFS Intent Log (ZIL) satisfies POSIX requirements for synchronous transactions. For instance, databases often require their transactions to be on stable storage devices when returning from a system call. NFS and other applications can also use
.Xr fsync 3C
to ensure data stability. By default, the intent log is allocated from blocks within the main pool. However, it might be possible to get better performance using separate intent log devices such as NVRAM or a dedicated disk. For example:
.Bd -literal
zpool create pool c0d0 c1d0 log c2d0
.Ed
.Pp

Multiple log devices can also be specified, and they can be mirrored. See the
.Sx EXAMPLES
section for an example of mirroring multiple log devices.
.Pp
Log devices can be added, replaced, attached, detached, and imported and
exported as part of the larger pool. Mirrored log devices can be removed by
specifying the top-level mirror for the log.
.Ss Cache Devices
Devices can be added to a storage pool as
.Qq cache devices .
These devices provide an additional layer of caching between main memory and
disk. For read-heavy workloads, where the working set size is much larger than
what can be cached in main memory, using cache devices allow much more of this
working set to be served from low latency media. Using cache devices provides
the greatest performance improvement for random read-workloads of mostly static
content.
.Pp
To create a pool with cache devices, specify a
.Sy cache
vdev with any number of devices. For example:
.Bd -literal
zpool create pool c0d0 c1d0 cache c2d0 c3d0
.Ed
.Pp
Cache devices cannot be mirrored or part of a raidz configuration. If a read
error is encountered on a cache device, that read I/O is reissued to the
original storage pool device, which might be part of a mirrored or raidz
configuration.
.Pp
The content of the cache devices is considered volatile, as is the case with
other system caches.
.Ss Properties
Each pool has several properties associated with it. Some properties are
read-only statistics while others are configurable and change the behavior of
the pool.
.Pp
The following are read-only properties:
.Bl -tag -width Ds
.It Sy available
Amount of storage available within the pool. This property can also be referred
to by its shortened column name,
.Sy avail .
.It Sy capacity
Percentage of pool space used. This property can also be referred to by its
shortened column name,
.Sy cap .
.It Sy expandsize
Amount of uninitialized space within the pool or device that can be used to
increase the total capacity of the pool. Uninitialized space consists of
any space on an EFI labeled vdev which has not been brought online
.Po e.g, using
.Nm zpool Cm online Fl e
.Pc .
This space occurs when a LUN is dynamically expanded.
.It Sy fragmentation
The amount of fragmentation in the pool.
.It Sy free
The amount of free space available in the pool.
.It Sy freeing
After a file system or snapshot is destroyed, the space it was using is
returned to the pool asynchronously.
.Sy freeing
is the amount of space remaining to be reclaimed. Over time
.Sy freeing
will decrease while
.Sy free

increases.
.It Sy health
The current health of the pool. Health can be one of
.Sy ONLINE , DEGRADED , FAULTED , OFFLINE, REMOVED , UNAVAIL .
.It Sy guid
A unique identifier for the pool.
.It Sy size
Total size of the storage pool.
.It Sy unsupported@ Ns Em feature_guid
Information about unsupported features that are enabled on the pool. See
.Xr zpool-features 5
for details.
.It Sy used
Amount of storage space used within the pool.
.El
.Pp
The space usage properties report actual physical space available to the
storage pool. The physical space can be different from the total amount of
space that any contained datasets can actually use. The amount of space used in
a raidz configuration depends on the characteristics of the data being
written. In addition, ZFS reserves some space for internal accounting
that the
.Xr zfs 1M
command takes into account, but the
.Nm
command does not. For non-full pools of a reasonable size, these effects should
be invisible. For small pools, or pools that are close to being completely
full, these discrepancies may become more noticeable.
.Pp
The following property can be set at creation time and import time:
.Bl -tag -width Ds
.It Sy altroot
Alternate root directory. If set, this directory is prepended to any mount
points within the pool. This can be used when examining an unknown pool where
the mount points cannot be trusted, or in an alternate boot environment, where
the typical paths are not valid.
.Sy altroot
is not a persistent property. It is valid only while the system is up. Setting
.Sy altroot
defaults to using
.Sy cachefile Ns = Ns Sy none ,
though this may be overridden using an explicit setting.
.El
.Pp
The following property can be set only at import time:
.Bl -tag -width Ds
.It Sy readonly Ns = Ns Sy on Ns | Ns Sy off
If set to
.Sy on ,
the pool will be imported in read-only mode. This property can also be referred
to by its shortened column name,
.Sy rdonly .
.El
.Pp
The following properties can be set at creation time and import time, and later
changed with the
.Nm zpool Cm set
command:
.Bl -tag -width Ds
.It Sy autoexpand Ns = Ns Sy on Ns | Ns Sy off
Controls automatic pool expansion when the underlying LUN is grown. If set to
.Sy on ,
the pool will be resized according to the size of the expanded device. If the
device is part of a mirror or raidz then all devices within that mirror/raidz
group must be expanded before the new space is made available to the pool. The
default behavior is

```
.Sy off .
This property can also be referred to by its shortened column name,
.Sy expand .
.It Sy autoreplace Ns = Ns Sy on Ns | Ns Sy off
Controls automatic device replacement. If set to
.Sy off ,
device replacement must be initiated by the administrator by using the
.Nm zpool Cm replace
command. If set to
.Sy on ,
any new device, found in the same physical location as a device that previously
belonged to the pool, is automatically formatted and replaced. The default
behavior is
.Sy off .
This property can also be referred to by its shortened column name,
.Sy replace .
.It Sy bootfs Ns = Ns Ar pool Ns / Ns Ar dataset
Identifies the default bootable dataset for the root pool. This property is
expected to be set mainly by the installation and upgrade programs.
.It Sy cachefile Ns = Ns Ar path Ns | Ns Sy none
Controls the location of where the pool configuration is cached. Discovering
all pools on system startup requires a cached copy of the configuration data
that is stored on the root file system. All pools in this cache are
automatically imported when the system boots. Some environments, such as
install and clustering, need to cache this information in a different location
so that pools are not automatically imported. Setting this property caches the
pool configuration in a different location that can later be imported with
.Nm zpool Cm import Fl c .
Setting it to the special value
.Sy none
creates a temporary pool that is never cached, and the special value
.Qq
.Pq empty string
uses the default location.
.Pp
Multiple pools can share the same cache file. Because the kernel destroys and
recreates this file when pools are added and removed, care should be taken when
attempting to access this file. When the last pool using a
.Sy cachefile
is exported or destroyed, the file is removed.
.It Sy comment Ns = Ns Ar text
A text string consisting of printable ASCII characters that will be stored
such that it is available even if the pool becomes faulted.  An administrator
can provide additional information about a pool using this property.
.It Sy dedupditto Ns = Ns Ar number
Threshold for the number of block ditto copies. If the reference count for a
deduplicated block increases above this number, a new ditto copy of this block
is automatically stored. The default setting is
.Sy 0
which causes no ditto copies to be created for deduplicated blocks. The miniumum
legal nonzero setting is
.Sy 100 .
.It Sy delegation Ns = Ns Sy on Ns | Ns Sy off
Controls whether a non-privileged user is granted access based on the dataset
permissions defined on the dataset. See
.Xr zfs 1M
for more information on ZFS delegated administration.
.It Sy failmode Ns = Ns Sy wait Ns | Ns Sy continue Ns | Ns Sy panic
Controls the system behavior in the event of catastrophic pool failure. This
condition is typically a result of a loss of connectivity to the underlying
storage device(s) or a failure of all devices within the pool. The behavior of
such an event is determined as follows:
.Bl -tag -width "continue"
.It Sy wait
Blocks all I/O access until the device connectivity is recovered and the errors
are cleared. This is the default behavior.
```



```
.It Sy continue
Returns
.Er EIO
to any new write I/O requests but allows reads to any of the remaining healthy
devices. Any write requests that have yet to be committed to disk would be
blocked.
.It Sy panic
Prints out a message to the console and generates a system crash dump.
.El
.It Sy feature@ Ns Ar feature_name Ns = Ns Sy enabled
The value of this property is the current state of
.Ar feature_name .
The only valid value when setting this property is
.Sy enabled
which moves
.Ar feature_name
to the enabled state. See
.Xr zpool-features 5
for details on feature states.
.It Sy listsnaps Ns = Ns Sy on Ns | Ns Sy off
Controls whether information about snapshots associated with this pool is
output when
.Nm zfs Cm list
is run without the
.Fl t
option. The default value is
.Sy off .
.It Sy version Ns = Ns Ar version
The current on-disk version of the pool. This can be increased, but never
decreased. The preferred method of updating pools is with the
.Nm zpool Cm upgrade
command, though this property can be used when a specific version is needed for
backwards compatibility. Once feature flags is enabled on a pool this property
will no longer have a value.
.El
.Ss Subcommands
All subcommands that modify state are logged persistently to the pool in their
original form.
.Pp
The
.Nm
command provides subcommands to create and destroy storage pools, add capacity
to storage pools, and provide information about the storage pools. The
following subcommands are supported:
.Bl -tag -width Ds
.It Xo
.Nm
.Fl \?
.Xc
Displays a help message.
.It Xo
.Nm
.Cm add
.Op Fl fn
.Ar pool vdev Ns ...
.Xc
Adds the specified virtual devices to the given pool. The
.Ar vdev
specification is described in the
.Sx Virtual Devices
section. The behavior of the
.Fl f
option, and the device checks performed are described in the
.Nm zpool Cm create
subcommand.
.Bl -tag -width Ds
```

```
.It Fl f
Forces use of
.Ar vdev Ns s ,
even if they appear in use or specify a conflicting replication level. Not all
devices can be overridden in this manner.
.It Fl n
Displays the configuration that would be used without actually adding the
.Ar vdev Ns s .
The actual pool creation can still fail due to insufficient privileges or
device sharing.
.El
.It Xo
.Nm
.Cm attach
.Op Fl f
.Ar pool device new_device
.Xc
Attaches
.Ar new_device
to the existing
.Ar device .
The existing device cannot be part of a raidz configuration. If
.Ar device
is not currently part of a mirrored configuration,
.Ar device
automatically transforms into a two-way mirror of
.Ar device
and
.Ar new_device .
If
.Ar device
is part of a two-way mirror, attaching
.Ar new_device
creates a three-way mirror, and so on. In either case,
.Ar new_device
begins to resilver immediately.
.Bl -tag -width Ds
.It Fl f
Forces use of
.Ar new_device ,
even if its appears to be in use. Not all devices can be overridden in this
manner.
.El
.It Xo
.Nm
.Cm clear
.Ar pool
.Op Ar device
.Xc
Clears device errors in a pool. If no arguments are specified, all device
errors within the pool are cleared. If one or more devices is specified, only
those errors associated with the specified device or devices are cleared.
.It Xo
.Nm
.Cm create
.Op Fl dfn
.Op Fl m Ar mountpoint
.Oo Fl o Ar property Ns = Ns Ar value Oc Ns ...
.Oo Fl O Ar file-system-property Ns = Ns Ar value Oc Ns ...
.Op Fl R Ar root
.Ar pool vdev Ns ...
.Xc
Creates a new storage pool containing the virtual devices specified on the
command line. The pool name must begin with a letter, and can only contain
alphanumeric characters as well as underscore
.Pq Qq Sy _ ,
```



```
dash
.Pq Qq Sy - ,
and period
.Pq Qq Sy \&. .
The pool names
.Sy mirror ,
.Sy raidz ,
.Sy spare
and
.Sy log
are reserved, as are names beginning with the pattern
.Sy c[0-9] .
The
.Ar vdev
specification is described in the
.Sx Virtual Devices
section.
.Pp
The command verifies that each device specified is accessible and not currently
in use by another subsystem. There are some uses, such as being currently
mounted, or specified as the dedicated dump device, that prevents a device from
ever being used by ZFS . Other uses, such as having a preexisting UFS file
system, can be overridden with the
.Fl f
option.
.Pp
The command also checks that the replication strategy for the pool is
consistent. An attempt to combine redundant and non-redundant storage in a
single pool, or to mix disks and files, results in an error unless
.Fl f
is specified. The use of differently sized devices within a single raidz or
mirror group is also flagged as an error unless
.Fl f
is specified.
.Pp
Unless the
.Fl R
option is specified, the default mount point is
.Pa / Ns Ar pool .
The mount point must not exist or must be empty, or else the root dataset
cannot be mounted. This can be overridden with the
.Fl m
option.
.Pp
By default all supported features are enabled on the new pool unless the
.Fl d
option is specified.
.Bl -tag -width Ds
.It Fl d
Do not enable any features on the new pool. Individual features can be enabled
by setting their corresponding properties to
.Sy enabled
with the
.Fl o
option. See
.Xr zpool-features 5
for details about feature properties.
.It Fl f
Forces use of
.Ar vdev Ns s ,
even if they appear in use or specify a conflicting replication level. Not all
devices can be overridden in this manner.
.It Fl m Ar mountpoint
Sets the mount point for the root dataset. The default mount point is
.Pa /pool
or
```

```
.Pa altroot/pool
if
.Ar altroot
is specified. The mount point must be an absolute path,
.Sy legacy ,
or
.Sy none .
For more information on dataset mount points, see
.Xr zfs 1M .
.It Fl n
Displays the configuration that would be used without actually creating the
pool. The actual pool creation can still fail due to insufficient privileges or
device sharing.
.It Fl o Ar property Ns = Ns Ar value
Sets the given pool properties. See the
.Sx Properties
section for a list of valid properties that can be set.
.It Fl O Ar file-system-property Ns = Ns Ar value
Sets the given file system properties in the root file system of the pool. See
the
.Sx Properties
section of
.Xr zfs 1M
for a list of valid properties that can be set.
.It Fl R Ar root
Equivalent to
.Fl o Sy cachefile Ns = Ns Sy none Fl o Sy altroot Ns = Ns Ar root
.El
.It Xo
.Nm
.Cm destroy
.Op Fl f
.Ar pool
.Xc
Destroys the given pool, freeing up any devices for other use. This command
tries to unmount any active datasets before destroying the pool.
.Bl -tag -width Ds
.It Fl f
Forces any active datasets contained within the pool to be unmounted.
.El
.It Xo
.Nm
.Cm detach
.Ar pool device
.Xc
Detaches
.Ar device
from a mirror. The operation is refused if there are no other valid replicas of
the data.
.It Xo
.Nm
.Cm export
.Op Fl f
.Ar pool Ns ...
.Xc
Exports the given pools from the system. All devices are marked as exported,
but are still considered in use by other subsystems. The devices can be moved
between systems
.Pq even those of different endianness
and imported as long as a sufficient number of devices are present.
.Pp
Before exporting the pool, all datasets within the pool are unmounted. A pool
can not be exported if it has a shared spare that is currently being used.
.Pp
For pools to be portable, you must give the
.Nm
```



```
command whole disks, not just slices, so that ZFS can label the disks with
portable EFI labels. Otherwise, disk drivers on platforms of different
endianness will not recognize the disks.
.Bl -tag -width Ds
.It Fl f
Forcefully unmount all datasets, using the
.Nm unmount Fl f
command.
.Pp
This command will forcefully export the pool even if it has a shared spare that
is currently being used. This may lead to potential data corruption.
.El
.It Xo
.Nm
.Cm get
.Op Fl Hp
.Op Fl o Ar field Ns Oo , Ns Ar field Oc Ns ...
.Sy all Ns | Ns Ar property Ns Oo , Ns Ar property Oc Ns ...
.Ar pool Ns ...
.Xc
Retrieves the given list of properties
.Po
or all properties if
.Sy all
is used
.Pc
for the specified storage pool(s). These properties are displayed with
the following fields:
.Bd -literal
        name          Name of storage pool
        property      Property name
        value         Property value
        source        Property source, either 'default' or 'local'.
.Ed
.Pp
See the
.Sx Properties
section for more information on the available pool properties.
.Bl -tag -width Ds
.It Fl H
Scripted mode. Do not display headers, and separate fields by a single tab
instead of arbitrary space.
.It Fl o Ar field
A comma-separated list of columns to display.
.Sy name Ns , Ns Sy property Ns , Ns Sy value Ns , Ns Sy source
is the default value.
.It Fl p
Display numbers in parsable (exact) values.
.El
.It Xo
.Nm
.Cm history
.Op Fl il
.Oo Ar pool Oc Ns ...
.Xc
Displays the command history of the specified pool(s) or all pools if no pool is
specified.
.Bl -tag -width Ds
.It Fl i
Displays internally logged ZFS events in addition to user initiated events.
.It Fl l
Displays log records in long format, which in addition to standard format
includes, the user name, the hostname, and the zone in which the operation was
performed.
.El
.It Xo
```

```
.Nm
.Cm import
.Op Fl D
.Op Fl d Ar dir
.Xc
Lists pools available to import. If the
.Fl d
option is not specified, this command searches for devices in
.Pa /dev/dsk .
The
.Fl d
option can be specified multiple times, and all directories are searched. If the
device appears to be part of an exported pool, this command displays a summary
of the pool with the name of the pool, a numeric identifier, as well as the vdev
layout and current health of the device for each device or file. Destroyed
pools, pools that were previously destroyed with the
.Nm zpool Cm destroy
command, are not listed unless the
.Fl D
option is specified.
.Pp
The numeric identifier is unique, and can be used instead of the pool name when
multiple exported pools of the same name are available.
.Bl -tag -width Ds
.It Fl c Ar cachefile
Reads configuration from the given
.Ar cachefile
that was created with the
.Sy cachefile
pool property. This
.Ar cachefile
is used instead of searching for devices.
.It Fl d Ar dir
Searches for devices or files in
.Ar dir .
The
.Fl d
option can be specified multiple times.
.It Fl D
Lists destroyed pools only.
.El
.It Xo
.Nm
.Cm import
.Fl a
.Op Fl DfmN
.Op Fl F Op Fl n
.Op Fl c Ar cachefile Ns | Ns Fl d Ar dir
.Op Fl o Ar mntopts
.Oo Fl o Ar property Ns = Ns Ar value Oc Ns ...
.Op Fl R Ar root
.Xc
Imports all pools found in the search directories. Identical to the previous
command, except that all pools with a sufficient number of devices available are
imported. Destroyed pools, pools that were previously destroyed with the
.Nm zpool Cm destroy
command, will not be imported unless the
.Fl D
option is specified.
.Bl -tag -width Ds
.It Fl a
Searches for and imports all pools found.
.It Fl c Ar cachefile
Reads configuration from the given
.Ar cachefile
that was created with the
```



```
.Sy cachefile
pool property. This
.Ar cachefile
is used instead of searching for devices.
.It Fl d Ar dir
Searches for devices or files in
.Ar dir .
The
.Fl d
option can be specified multiple times. This option is incompatible with the
.Fl c
option.
.It Fl D
Imports destroyed pools only. The
.Fl f
option is also required.
.It Fl f
Forces import, even if the pool appears to be potentially active.
.It Fl F
Recovery mode for a non-importable pool. Attempt to return the pool to an
importable state by discarding the last few transactions. Not all damaged pools
can be recovered by using this option. If successful, the data from the
discarded transactions is irretrievably lost. This option is ignored if the pool
is importable or already imported.
.It Fl m
Allows a pool to import when there is a missing log device. Recent transactions
can be lost because the log device will be discarded.
.It Fl n
Used with the
.Fl F
recovery option. Determines whether a non-importable pool can be made importable
again, but does not actually perform the pool recovery. For more details about
pool recovery mode, see the
.Fl F
option, above.
.It Fl N
Import the pool without mounting any file systems.
.It Fl o Ar mntopts
Comma-separated list of mount options to use when mounting datasets within the
pool. See
.Xr zfs 1M
for a description of dataset properties and mount options.
.It Fl o Ar property Ns = Ns Ar value
Sets the specified property on the imported pool. See the
.Sx Properties
section for more information on the available pool properties.
.It Fl R Ar root
Sets the
.Sy cachefile
property to
.Sy none
and the
.Sy altroot
property to
.Ar root .
.El
.It Xo
.Nm
.Cm import
.Op Fl Dfm
.Op Fl F Op Fl n
.Op Fl c Ar cachefile Ns | Ns Fl d Ar dir
.Op Fl o Ar mntopts
.Oo Fl o Ar property Ns = Ns Ar value Oc Ns ...
.Op Fl R Ar root
.Ar pool Ns | Ns Ar id
```

```
.Op Ar newpool
.Xc
Imports a specific pool. A pool can be identified by its name or the numeric
identifier. If
.Ar newpool
is specified, the pool is imported using the name
.Ar newpool .
Otherwise, it is imported with the same name as its exported name.
.Pp
If a device is removed from a system without running
.Nm zpool Cm export
first, the device appears as potentially active. It cannot be determined if
this was a failed export, or whether the device is really in use from another
host. To import a pool in this state, the
.Fl f
option is required.
.Bl -tag -width Ds
.It Fl c Ar cachefile
Reads configuration from the given
.Ar cachefile
that was created with the
.Sy cachefile
pool property. This
.Ar cachefile
is used instead of searching for devices.
.It Fl d Ar dir
Searches for devices or files in
.Ar dir .
The
.Fl d
option can be specified multiple times. This option is incompatible with the
.Fl c
option.
.It Fl D
Imports destroyed pool. The
.Fl f
option is also required.
.It Fl f
Forces import, even if the pool appears to be potentially active.
.It Fl F
Recovery mode for a non-importable pool. Attempt to return the pool to an
importable state by discarding the last few transactions. Not all damaged pools
can be recovered by using this option. If successful, the data from the
discarded transactions is irretrievably lost. This option is ignored if the pool
is importable or already imported.
.It Fl m
Allows a pool to import when there is a missing log device. Recent transactions
can be lost because the log device will be discarded.
.It Fl n
Used with the
.Fl F
recovery option. Determines whether a non-importable pool can be made importable
again, but does not actually perform the pool recovery. For more details about
pool recovery mode, see the
.Fl F
option, above.
.It Fl o Ar mntopts
Comma-separated list of mount options to use when mounting datasets within the
pool. See
.Xr zfs 1M
for a description of dataset properties and mount options.
.It Fl o Ar property Ns = Ns Ar value
Sets the specified property on the imported pool. See the
.Sx Properties
section for more information on the available pool properties.
.It Fl R Ar root
```



```
Sets the
.Sy cachefile
property to
.Sy none
and the
.Sy altroot
property to
.Ar root .
.El
.It Xo
.Nm
.Cm iostat
.Op Fl v
.Op Fl T Sy u Ns | Ns Sy d
.Oo Ar pool Oc Ns ...
.Op Ar interval Op Ar count
.Xc
Displays I/O statistics for the given pools. When given an
.Ar interval ,
the statistics are printed every
.Ar interval
seconds until ^C is pressed. If no
.Ar pool Ns s
are specified, statistics for every pool in the system is shown. If
.Ar count
is specified, the command exits after
.Ar count
reports are printed.
.Bl -tag -width Ds
.It Fl T Sy u Ns | Ns Sy d
Display a time stamp. Specify
.Sy u
for a printed representation of the internal representation of time. See
.Xr time 2 .
Specify
.Sy d
for standard date format. See
.Xr date 1 .
.It Fl v
Verbose statistics. Reports usage statistics for individual vdevs within the
pool, in addition to the pool-wide statistics.
.El
.It Xo
.Nm
.Cm list
.Op Fl Hpv
.Op Fl o Ar property Ns Oo , Ns Ar property Oc Ns ...
.Op Fl T Sy u Ns | Ns Sy d
.Oo Ar pool Oc Ns ...
.Op Ar interval Op Ar count
.Xc
Lists the given pools along with a health status and space usage. If no
.Ar pool Ns s
are specified, all pools in the system are listed. When given an
.Ar interval ,
the information is printed every
.Ar interval
seconds until ^C is pressed. If
.Ar count
is specified, the command exits after
.Ar count
reports are printed.
.Bl -tag -width Ds
.It Fl H
Scripted mode. Do not display headers, and separate fields by a single tab
instead of arbitrary space.
```

```
.It Fl o Ar property
Comma-separated list of properties to display. See the
.Sx Properties
section for a list of valid properties. The default list is
.Sy name , size , used , available , fragmentation , expandsize , capacity ,
.Sy dedupratio , health , altroot .
.It Fl p
Display numbers in parsable
.Pq exact
values.
.It Fl T Sy u Ns | Ns Sy d
Display a time stamp. Specify
.Fl u
for a printed representation of the internal representation of time. See
.Xr time 2 .
Specify
.Fl d
for standard date format. See
.Xr date 1 .
.It Fl v
Verbose statistics. Reports usage statistics for individual vdevs within the
pool, in addition to the pool-wise statistics.
.El
.It Xo
.Nm
.Cm offline
.Op Fl t
.Ar pool Ar device Ns ...
.Xc
Takes the specified physical device offline. While the
.Ar device
is offline, no attempt is made to read or write to the device. This command is
not applicable to spares.
.Bl -tag -width Ds
.It Fl t
Temporary. Upon reboot, the specified physical device reverts to its previous
state.
.El
.It Xo
.Nm
.Cm online
.Op Fl e
.Ar pool Ar device Ns ...
.Xc
Brings the specified physical device online. This command is not applicable to
spares.
.Bl -tag -width Ds
.It Fl e
Expand the device to use all available space. If the device is part of a mirror
or raidz then all devices must be expanded before the new space will become
available to the pool.
.El
.It Xo
.Nm
.Cm reguid
.Ar pool
.Xc
Generates a new unique identifier for the pool. You must ensure that all devices
in this pool are online and healthy before performing this action.
.It Xo
.Nm
.Cm reopen
.Ar pool
.Xc
Reopen all the vdevs associated with the pool.
.It Xo
```



```
.Nm
.Cm remove
.Ar pool Ar device Ns ...
.Xc
Removes the specified device from the pool. This command currently only supports
removing hot spares, cache, and log devices. A mirrored log device can be
removed by specifying the top-level mirror for the log. Non-log devices that are
part of a mirrored configuration can be removed using the
.Nm zpool Cm detach
command. Non-redundant and raidz devices cannot be removed from a pool.
.It Xo
.Nm
.Cm replace
.Op Fl f
.Ar pool Ar device Op Ar new_device
.Xc
Replaces
.Ar old_device
with
.Ar new_device .
This is equivalent to attaching
.Ar new_device ,
waiting for it to resilver, and then detaching
.Ar old_device .
.Pp
The size of
.Ar new_device
must be greater than or equal to the minimum size of all the devices in a mirror
or raidz configuration.
.Pp
.Ar new_device
is required if the pool is not redundant. If
.Ar new_device
is not specified, it defaults to
.Ar old_device .
This form of replacement is useful after an existing disk has failed and has
been physically replaced. In this case, the new disk may have the same
.Pa /dev/dsk
path as the old device, even though it is actually a different disk. ZFS
recognizes this.
.Bl -tag -width Ds
.It Fl f
Forces use of
.Ar new_device ,
even if its appears to be in use. Not all devices can be overridden in this
manner.
.El
.It Xo
.Nm
.Cm scrub
.Op Fl s
.Ar pool Ns ...
.Xc
Begins a scrub. The scrub examines all data in the specified pools to verify
that it checksums correctly. For replicated
.Pq mirror or raidz
devices, ZFS automatically repairs any damage discovered during the scrub. The
.Nm zpool Cm status
command reports the progress of the scrub and summarizes the results of the
scrub upon completion.
.Pp
Scrubbing and resilvering are very similar operations. The difference is that
resilvering only examines data that ZFS knows to be out of date
.Po
for example, when attaching a new device to a mirror or replacing an existing
device
```

```
.Pc ,
whereas scrubbing examines all data to discover silent errors due to hardware
faults or disk failure.
.Pp
Because scrubbing and resilvering are I/O-intensive operations, ZFS only allows
one at a time. If a scrub is already in progress, the
.Nm zpool Cm scrub
command terminates it and starts a new scrub. If a resilver is in progress, ZFS
does not allow a scrub to be started until the resilver completes.
.Bl -tag -width Ds
.It Fl s
Stop scrubbing.
.El
.It Xo
.Nm
.Cm set
.Ar property Ns = Ns Ar value
.Ar pool
.Xc
Sets the given property on the specified pool. See the
.Sx Properties
section for more information on what properties can be set and acceptable
values.
.It Xo
.Nm
.Cm split
.Op Fl n
.Oo Fl o Ar property Ns = Ns Ar value Oc Ns ...
.Op Fl R Ar root
.Ar pool newpool
.Xc
Splits devices off
.Ar pool
creating
.Ar newpool .
All vdevs in
.Ar pool
must be mirrors. At the time of the split,
.Ar newpool
will be a replica of
.Ar pool .
.Bl -tag -width Ds
.It Fl n
Do dry run, do not actually perform the split. Print out the expected
configuration of
.Ar newpool .
.It Fl o Ar property Ns = Ns Ar value
Sets the specified property for
.Ar newpool .
See the
.Sx Properties
section for more information on the available pool properties.
.It Fl R Ar root
Set
.Sy altroot
for
.Ar newpool
to
.Ar root
and automaticaly import it.
.El
.It Xo
.Nm
.Cm status
.Op Fl Dvx
.Op Fl T Sy u Ns | Ns Sy d
```



```
.Oo Ar pool Oc Ns ...
.Op Ar interval Op Ar count
.Xc
Displays the detailed health status for the given pools. If no
.Ar pool
is specified, then the status of each pool in the system is displayed. For more
information on pool and device health, see the
.Sx Device Failure and Recovery
section.
.Pp
If a scrub or resilver is in progress, this command reports the percentage done
and the estimated time to completion. Both of these are only approximate,
because the amount of data in the pool and the other workloads on the system can
change.
.Bl -tag -width Ds
.It Fl D
Display a histogram of deduplication statistics, showing the allocated
.Pq physically present on disk
and referenced
.Pq logically referenced in the pool
block counts and sizes by reference count.
.It Fl T Sy u Ns | Ns Sy d
Display a time stamp. Specify
.Fl u
for a printed representation of the internal representation of time. See
.Xr time 2 .
Specify
.Fl d
for standard date format. See
.Xr date 1 .
.It Fl v
Displays verbose data error information, printing out a complete list of all
data errors since the last complete pool scrub.
.It Fl x
Only display status for pools that are exhibiting errors or are otherwise
unavailable. Warnings about pools not using the latest on-disk format will not
be included.
.El
.It Xo
.Nm
.Cm upgrade
.Xc
Displays pools which do not have all supported features enabled and pools
formatted using a legacy ZFS version number. These pools can continue to be
used, but some features may not be available. Use
.Nm zpool Cm upgrade Fl a
to enable all features on all pools.
.It Xo
.Nm
.Cm upgrade
.Fl v
.Xc
Displays legacy ZFS versions supported by the current software. See
.Xr zpool-features 5
for a description of feature flags features supported by the current software.
.It Xo
.Nm
.Cm upgrade
.Op Fl V Ar version
.Fl a Ns | Ns Ar pool Ns ...
.Xc
Enables all supported features on the given pool. Once this is done, the pool
will no longer be accessible on systems that do not support feature flags. See
.Xr zpool-features 5
for details on compatibility with systems that support feature flags, but do not
support all features enabled on the pool.
```

```
.Bl -tag -width Ds
.It Fl a
Enables all supported features on all pools.
.It Fl V Ar version
Upgrade to the specified legacy version. If the
.Fl V
flag is specified, no features will be enabled on the pool. This option can only
be used to increase the version number up to the last supported legacy version
number.
.El
.El
.Sh EXIT STATUS
The following exit values are returned:
.Bl -tag -width Ds
.It Sy 0
Successful completion.
.It Sy 1
An error occurred.
.It Sy 2
Invalid command line options were specified.
.El
.Sh EXAMPLES
.Bl -tag -width Ds
.It Sy Example 1 No Creating a RAID-Z Storage Pool
The following command creates a pool with a single raidz root vdev that
consists of six disks.
.Bd -literal
# zpool create tank raidz c0t0d0 c0t1d0 c0t2d0 c0t3d0 c0t4d0 c0t5d0
.Ed
.It Sy Example 2 No Creating a Mirrored Storage Pool
The following command creates a pool with two mirrors, where each mirror
contains two disks.
.Bd -literal
# zpool create tank mirror c0t0d0 c0t1d0 mirror c0t2d0 c0t3d0
.Ed
.It Sy Example 3 No Creating a ZFS Storage Pool by Using Slices
The following command creates an unmirrored pool using two disk slices.
.Bd -literal
# zpool create tank /dev/dsk/c0t0d0s1 c0t1d0s4
.Ed
.It Sy Example 4 No Creating a ZFS Storage Pool by Using Files
The following command creates an unmirrored pool using files. While not
recommended, a pool based on files can be useful for experimental purposes.
.Bd -literal
# zpool create tank /path/to/file/a /path/to/file/b
.Ed
.It Sy Example 5 No Adding a Mirror to a ZFS Storage Pool
The following command adds two mirrored disks to the pool
.Em tank ,
assuming the pool is already made up of two-way mirrors. The additional space
is immediately available to any datasets within the pool.
.Bd -literal
# zpool add tank mirror c1t0d0 c1t1d0
.Ed
.It Sy Example 6 No Listing Available ZFS Storage Pools
The following command lists all available pools on the system. In this case,
the pool
.Em zion
is faulted due to a missing device. The results from this command are similar
to the following:
.Bd -literal
# zpool list
NAME    SIZE  ALLOC   FREE   FRAG  EXPANDSZ    CAP  DEDUP  HEALTH  ALTROOT
rpool  19.9G  8.43G  11.4G    33%         -    42%  1.00x  ONLINE  -
tank   61.5G  20.0G  41.5G    48%         -    32%  1.00x  ONLINE  -
zion       -      -      -      -         -      -      -  FAULTED -
```



```
.Ed
.It Sy Example 7 No Destroying a ZFS Storage Pool
The following command destroys the pool
.Em tank
and any datasets contained within.
.Bd -literal
# zpool destroy -f tank
.Ed
.It Sy Example 8 No Exporting a ZFS Storage Pool
The following command exports the devices in pool
.Em tank
so that they can be relocated or later imported.
.Bd -literal
# zpool export tank
.Ed
.It Sy Example 9 No Importing a ZFS Storage Pool
The following command displays available pools, and then imports the pool
.Em tank
for use on the system. The results from this command are similar to the
following:
.Bd -literal
# zpool import
  pool: tank
    id: 15451357997522795478
 state: ONLINE
action: The pool can be imported using its name or numeric identifier.
config:

        tank        ONLINE
          mirror    ONLINE
            c1t2d0  ONLINE
            c1t3d0  ONLINE

# zpool import tank
.Ed
.It Sy Example 10 No Upgrading All ZFS Storage Pools to the Current Version
The following command upgrades all ZFS Storage pools to the current version of
the software.
.Bd -literal
# zpool upgrade -a
This system is currently running ZFS version 2.
.Ed
.It Sy Example 11 No Managing Hot Spares
The following command creates a new pool with an available hot spare:
.Bd -literal
# zpool create tank mirror c0t0d0 c0t1d0 spare c0t2d0
.Ed
.Pp
If one of the disks were to fail, the pool would be reduced to the degraded
state. The failed device can be replaced using the following command:
.Bd -literal
# zpool replace tank c0t0d0 c0t3d0
.Ed
.Pp
Once the data has been resilvered, the spare is automatically removed and is
made available should another device fails. The hot spare can be permanently
removed from the pool using the following command:
.Bd -literal
# zpool remove tank c0t2d0
.Ed
.It Sy Example 12 No Creating a ZFS Pool with Mirrored Separate Intent Logs
The following command creates a ZFS storage pool consisting of two, two-way
mirrors and mirrored log devices:
.Bd -literal
# zpool create pool mirror c0d0 c1d0 mirror c2d0 c3d0 log mirror \e
  c4d0 c5d0
```

```
.Ed
.It Sy Example 13 No Adding Cache Devices to a ZFS Pool
The following command adds two disks for use as cache devices to a ZFS storage
pool:
.Bd -literal
# zpool add pool cache c2d0 c3d0
.Ed
.Pp
Once added, the cache devices gradually fill with content from main memory.
Depending on the size of your cache devices, it could take over an hour for
them to fill. Capacity and reads can be monitored using the
.Cm iostat
option as follows:
.Bd -literal
# zpool iostat -v pool 5
.Ed
.It Sy Example 14 No Removing a Mirrored Log Device
The following command removes the mirrored log device
.Sy mirror-2 .
Given this configuration:
.Bd -literal
  pool: tank
 state: ONLINE
 scrub: none requested
config:

         NAME        STATE     READ WRITE CKSUM
         tank        ONLINE       0     0     0
           mirror-0  ONLINE       0     0     0
             c6t0d0  ONLINE       0     0     0
             c6t1d0  ONLINE       0     0     0
           mirror-1  ONLINE       0     0     0
             c6t2d0  ONLINE       0     0     0
             c6t3d0  ONLINE       0     0     0
         logs
           mirror-2  ONLINE       0     0     0
             c4t0d0  ONLINE       0     0     0
             c4t1d0  ONLINE       0     0     0
.Ed
.Pp
The command to remove the mirrored log
.Sy mirror-2
is:
.Bd -literal
# zpool remove tank mirror-2
.Ed
.It Sy Example 15 No Displaying expanded space on a device
The following command dipslays the detailed information for the pool
.Em data .
This pool is comprised of a single raidz vdev where one of its devices
increased its capacity by 10GB. In this example, the pool will not be able to
utilize this extra capacity until all the devices under the raidz vdev have
been expanded.
.Bd -literal
# zpool list -v data
NAME         SIZE  ALLOC   FREE   FRAG  EXPANDSZ    CAP  DEDUP  HEALTH  ALTROOT
data        23.9G  14.6G  9.30G    48%         -    61%  1.00x  ONLINE  -
  raidz1    23.9G  14.6G  9.30G    48%         -
    c1t1d0      -      -      -      -         -
    c1t2d0      -      -      -      -       10G
    c1t3d0      -      -      -      -         -
.Ed
.El
.Sh INTERFACE STABILITY
.Sy Evolving
.Sh SEE ALSO
```



```
.Xr zfs 1M ,
.Xr attributes 5 ,
.Xr zpool-features 5
```