

# ZFS

# Implementation and Operations Guide

December 2016 (Edition 1.0)  
Fujitsu Limited

## ■ Purpose

- This document presents methods of building and operating ZFS (Zettabyte File System), which is the standard file system of Oracle Solaris 11.

## ■ Audience

- People who have a basic knowledge of Oracle Solaris and RAID
- People who are referring to the *ZFS Overview and Design Guide*

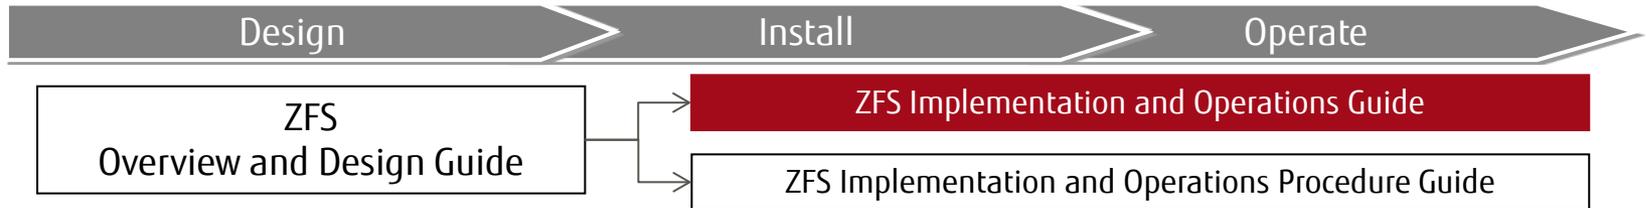
## ■ Notes

- This document describes the procedure for building a Solaris environment on Oracle VM Server for SPARC. Note that the contents depend somewhat on Oracle VM Server for SPARC.  
Example: Device names such as for disks and network interfaces are specific to Oracle VM Server for SPARC, differing from a physical server environment.
- This document is based on Oracle Solaris 11.3. For the latest information on Oracle Solaris 11, see the manuals from Oracle.
  - Oracle Solaris 11 Documentation  
<http://www.oracle.com/technetwork/documentation/solaris-11-192991.html>
- Fujitsu M10 is sold as SPARC M10 Systems by Fujitsu in Japan. Fujitsu M10 and SPARC M10 Systems are identical products.

## ■ Positioning of documents

- ZFS

<http://www.fujitsu.com/global/products/computing/servers/unix/sparc/downloads/documents/>



\* Read this document together with the *ZFS Implementation and Operations Procedure Guide*.

## ■ Descriptions in this document

- The section numbers of commands are omitted.

Example:

- ls(1) => ls command
- shutdown(1M) => shutdown command

- The following table lists terms that may be abbreviated.

Abbreviation	Formal Name
Solaris	Oracle Solaris

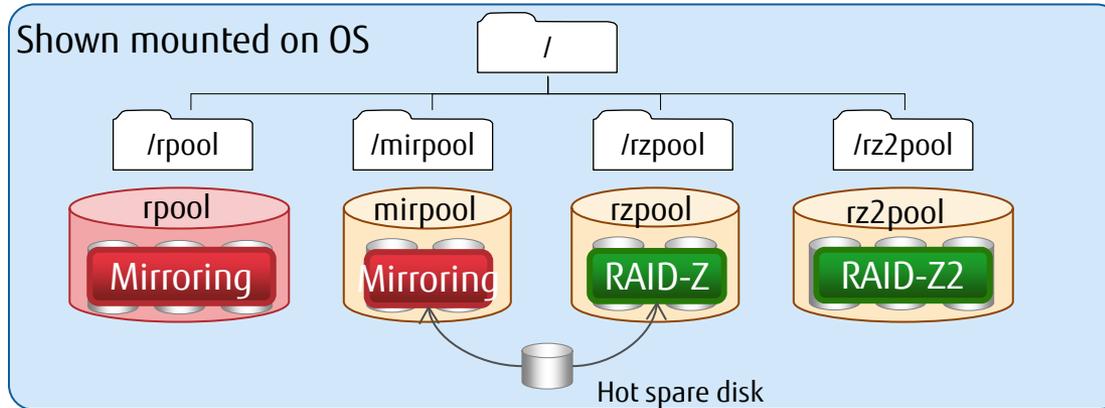
1. Building a Storage Pool
  2. Building a File System
  3. Creating Snapshots and Clones
  4. Backup/Restore
  5. Releasing a Mirror Disk
- Reference Information
  - Appendix

# 1. Building a Storage Pool

This chapter describes how to change the storage pool configuration and how to create a storage pool.

## ■ Storage pool configuration to be created

This chapter describes how to change the storage pool configuration and how to create a storage pool. The following figure shows the configuration of the storage pool to be built.



## ■ Storage pool operations

- Changing the root pool (system area) configuration
  - Add a disk to the root pool, `rpool`, to change it to a mirror configuration.
- Creating a storage pool (user area)
  - Create storage pools in a mirror (RAID 1) configuration (`mirpool`), RAID-Z configuration (`rzpool`), and RAID-Z2 configuration (`rz2pool`).
- Registering a hot spare disk
  - Share a hot disk between `mirpool` and `rzpool`.
- Checking a storage pool
  - Check the storage pool status, using `rpool` as an example.

# Changing the Root Pool (System Area) Configuration

## ■ Add a disk to change from the single configuration to a multi-way mirror configuration.

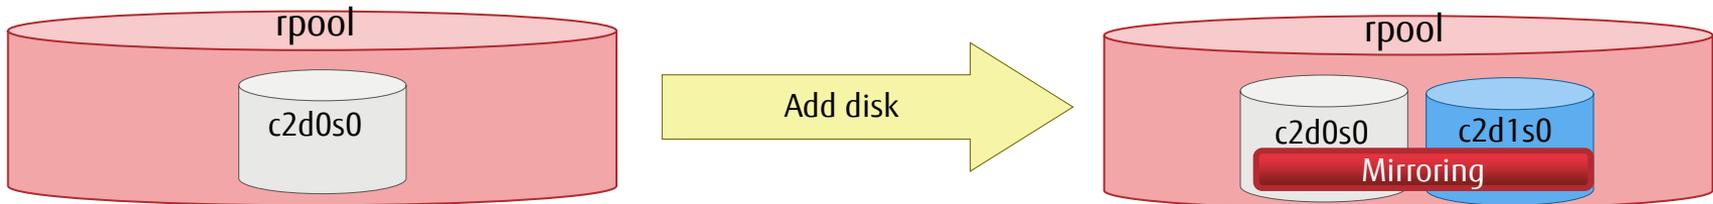
– The default (at OS installation) configuration of the root pool has only one disk. After OS installation, you can change the single configuration to a mirror configuration.

### ✓ Add a mirror disk (zpool attach command).

Syntax: `zpool attach pool_name mirror_source_disk mirror_disk`

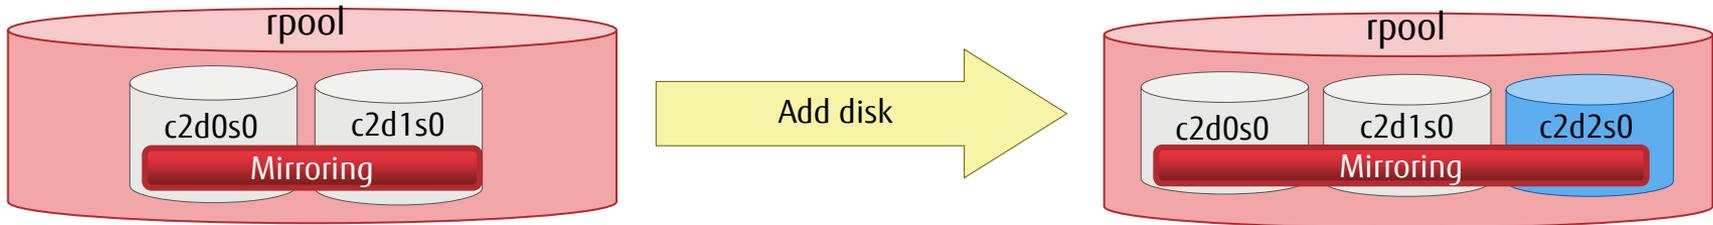
– Single-disk configuration to two-way mirror configuration

```
# zpool attach rpool c2d0s0 c2d1s0
```



– Two-way mirror configuration to three-way mirror configuration

```
# zpool attach rpool c2d0s0 c2d2s0
```

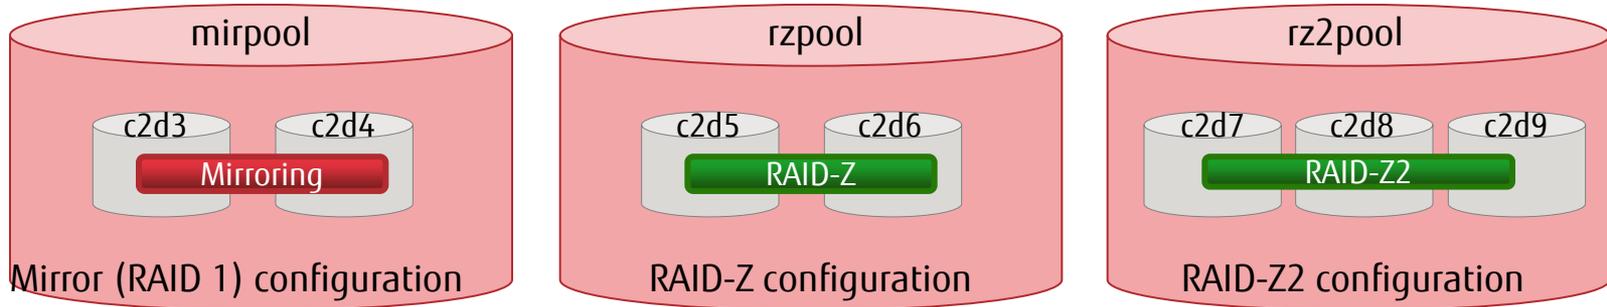


– If the firmware on the SPARC M10 is XCP 2230 or later, you can use a whole disk (EFI label) for the root pool. If it is earlier than XCP 2230, you need to use a disk slice (SMI label).

# Creating a Storage Pool (User Area)

## ■ Create a storage pool (user area).

- Specify a RAID level (RAID 0, RAID 1, RAID-Z, RAID-Z2, etc.) for the storage pool when creating it.



## ✓ Create a storage pool (zpool create command).

Syntax: `zpool create pool_name [RAID] disk_name ...`

\* If [RAID] is omitted, the stripe configuration is specified.

- Mirror (RAID 1) configuration

```
# zpool create mirpool mirror c2d3 c2d4
```

- RAID-Z configuration

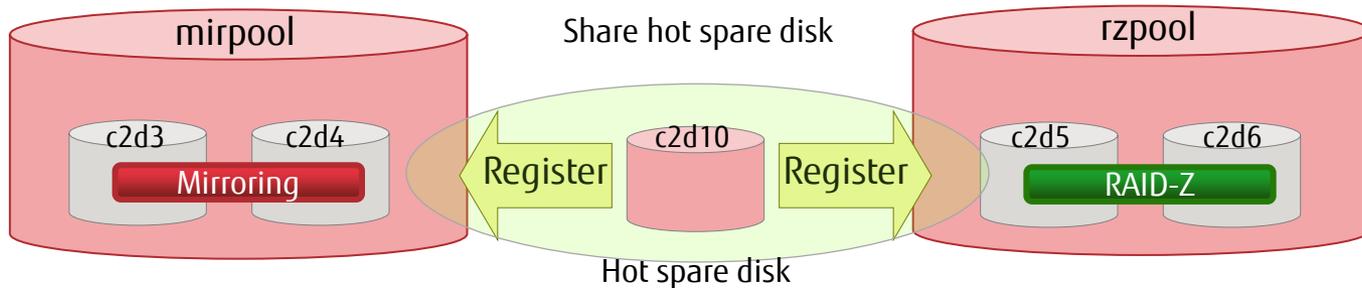
```
# zpool create rzpool raidz c2d5 c2d6
```

- RAID-Z2 configuration

```
# zpool create rz2pool raidz2 c2d7 c2d8 c2d9
```

# Registering a Hot Spare Disk

- You can set a hot spare disk to prepare for a disk failure in the storage pool.
  - Multiple storage pools can share a hot spare disk.



- ✓ Register a hot spare disk (zpool add command).

Syntax: `zpool add pool_name spare hot_spare_disk_name_to_register [added_hot_spare_disk]`

```
# zpool add mirpool spare c2d10
# zpool add rzpool spare c2d10
```

\* You can register multiple hot spare disks at the same time.

- ✓ Unregister a hot spare disk (zpool remove command).

Syntax: `zpool remove pool_name hot_spare_disk_name_to_unregister`

```
# zpool remove mirpool c2d10
```



– You can share a hot spare disk by registering the same disk as a hot spare disk with multiple storage pools.

# Checking a Storage Pool 1/3

## ■ Use mainly the following commands to check the storage pool status.

– zpool list command

Checks basic information such as the storage pool name and capacity used.

– zpool status command

Checks detailed information, including error data and whether a failure occurred, in addition to the above information.

## ✓ Check the basic storage pool information (zpool list command).

```
# zpool list
```

NAME	SIZE	ALLOC	FREE	CAP	DEDUP	HEALTH	ALTROOT
rpool	11.9G	6.09G	5.78G	51%	1.00x	ONLINE	/mnt

(1) (2) (3) (4) (5) (6) (7) (8)

### Output format

(1): Storage pool name

(2): Storage pool size

(3): Allocated physical capacity

(4): Unallocated capacity

(5): Used disk space

(6): Amount of deduplication

(7): State

ONLINE

Normal state

OFFLINE

Offline state set manually by administrator

FAULTED

Virtual device inaccessible

DEGRADED

Virtual device experienced failure but is available

UNAVAILABLE

Device or virtual device inaccessible

REMOVE

Device physically removed while system is running

(8): Mount point of ZFS alternative root pool or alternative root pool

\* An alternative root pool is a boot image, which is used for startup from the alternative root pool when startup from the root pool has failed.

- ✓ Check detailed storage pool information (zpool status command).

```
# zpool status
pool: rpool (1)
state: ONLINE (2)
scan: resilvered 70.0G in 10m15s with 0 errors on Wed Dec 31 19:00:00 1969 (3)

config:
(4)
NAME                STATE      READ    WRITE   CKSUM
rpool                ONLINE    0       0       0
  mirror-0
    c0t50000394083213E0d0 ONLINE    0       0       0
    c0t500003942823F558d0 ONLINE    0       0       0

errors: No known data errors (9)
```

## Output format

- (1): Storage pool name
- (2): Storage pool state
- (3): Scrubbing and resynchronization state
- (4): Storage pool name, RAID, disk name
- (5): State
- (6): Number of read errors
- (7): Number of write errors
- (8): Number of checksum errors
- (9): Error information  
"No known data errors" is output for a normal situation.

- You can check the property information for a storage pool.

- ✓ Check the storage pool property information (zpool get command).

```
# zpool get all rpool
NAME      PROPERTY      VALUE      SOURCE
rpool    size          11.9G     -
rpool    used          5.95G     -
rpool    available     5.92G     -
rpool    capacity      50%       -
```

- ✓ Check individually specified property information.

Specify properties delimited by a comma (,).

```
# zpool get bootfs,listsnapshots rpool
NAME      PROPERTY      VALUE      SOURCE
rpool    bootfs        rpool/ROOT/SRU1111_BE  local
rpool    listsnapshots on            default
```

Displays only the specified properties.

# <<Reference>> Storage Pool Properties

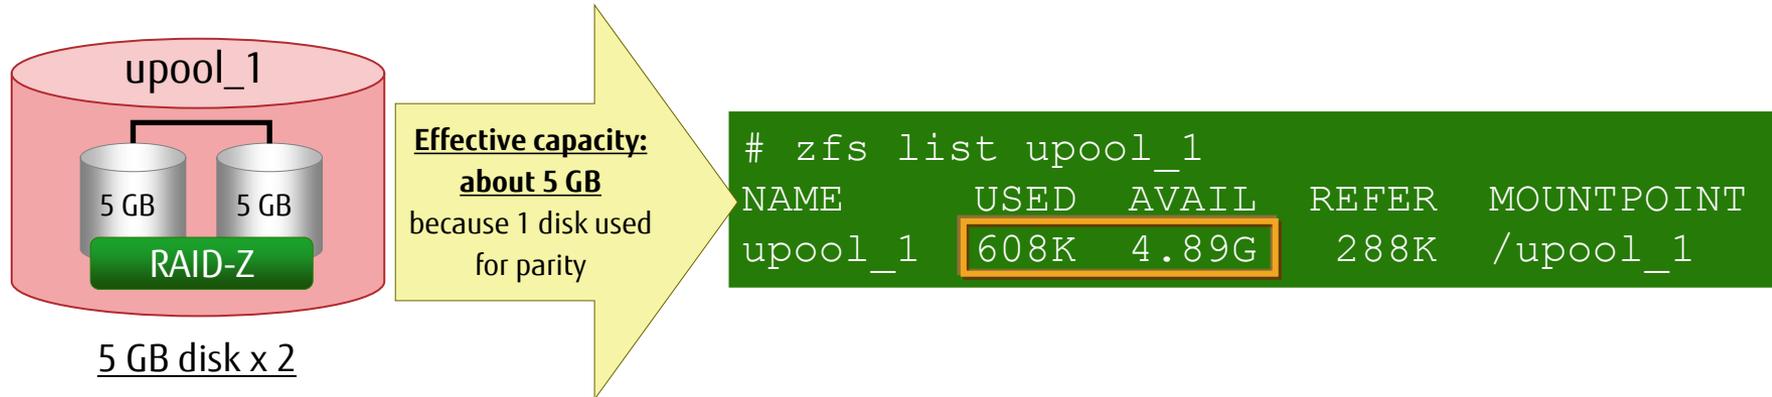
Property Name	Description
allocated	Read-only value that identifies the amount of the physically allocated storage area in the pool
altroot	Identifies the alternative root directory.
autoexpand	Controls automatic pool expansion.
autoreplace	Controls automatic device replacement.
bootfs	Identifies the default bootable file system of the root pool.
cachefile	Controls the cache location of the pool configuration information.
capacity	Read-only value that identifies the percentage of the pool area for use (%)
dedupditto	Sets the threshold of the reference count of deduplicated blocks.
dedupratio	Read-only value that is the deduplication ratio achieved for the pool
delegation	Controls whether to grant the access right defined for the file system to a non-privileged user.
failmode	Controls system behavior in case a catastrophic pool failure occurs.
free	Read-only value that identifies the number of unallocated blocks in the pool
guid	Read-only property that identifies a unique identifier
health	Read-only property that identifies the current soundness of the pool
listshares	Controls whether the zfs list command displays the shared information in the pool.
listsnapshots	Controls whether the zfs list command displays the snapshot information associated with the pool.
readonly	Specifies whether to allow changes to the pool.
size	Read-only property that identifies the total size of the storage pool
version	Identifies the current version on the pool disk.

## ■ Use the zfs list command to check the effective capacity of a storage pool.

- ✓ Check the effective capacity of a storage pool (zpool list command).

Syntax: `zfs list storage_pool_name`

### Execution example: Storage pool in RAID-Z configuration



- USED and AVAIL display the currently used capacity and available capacity, respectively.
- The effective capacity is the sum of the USED and AVAIL values.

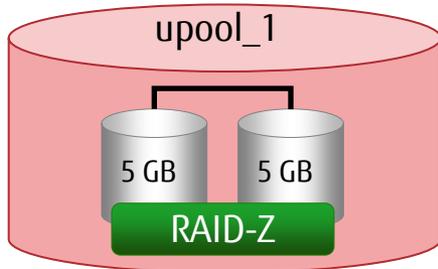


- The zpool list command may not be able to check the effective capacity, depending on the RAID configuration.  
-> For details, see [<<Reference>> How to Check the Effective Capacity of a Storage Pool 2/2](#).
- For details, see "Resolving ZFS Space Issues" at the following:  
[https://docs.oracle.com/cd/E53394\\_01/html/E54801/gbby.html](https://docs.oracle.com/cd/E53394_01/html/E54801/gbby.html)

# <<Reference>> How to Check the Effective Capacity of a Storage Pool 2/2

Note that the capacity displayed by the zpool list command has different meanings depending on the RAID configuration.

## RAID-Z configuration

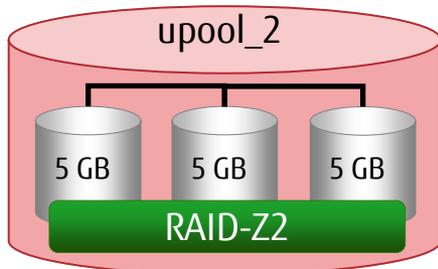


**Effective capacity:**  
**5 GB**  
because 1 disk  
used for parity

```
# zpool list upool_1
NAME      SIZE  ALLOC   FREE  CAP  DEDUP  HEALTH  ALTROOT
upool_1   9.94G  2.16M  9.94G   0%  1.00x  ONLINE  -
```

Displays **total capacity of disks**

## RAID-Z2 configuration

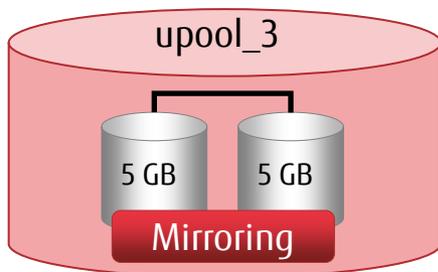


**Effective capacity:**  
**5 GB**  
because 2 disks  
used for parity

```
# zpool list upool_2
NAME      SIZE  ALLOC   FREE  CAP  DEDUP  HEALTH  ALTROOT
upool_2   14.9G  3.23M  14.9G   0%  1.00x  ONLINE  -
```

Displays **total capacity of disks**

## Mirror configuration



**Effective capacity:**  
**5 GB**

```
# zpool list upool_3
NAME      SIZE  ALLOC   FREE  CAP  DEDUP  HEALTH  ALTROOT
upool_3   4.97G  1.38M  4.97G   0%  1.00x  ONLINE  -
```

Displays **effective capacity**

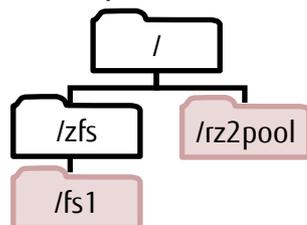
## 2. Building a File System

This chapter describes how to create a file system and how to make various settings.

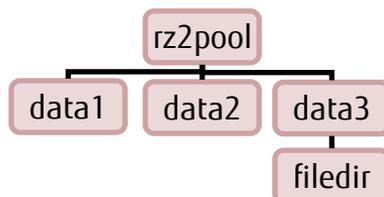
## ■ File system configuration to be created

This chapter describes how to create a file system and how to set file system properties. The following figure shows the configuration of the file system to be built.

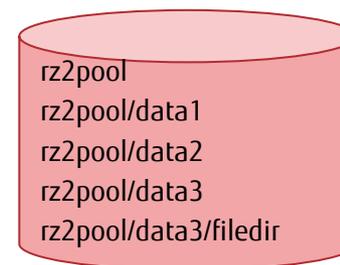
Directory tree



File system tree



Storage pool



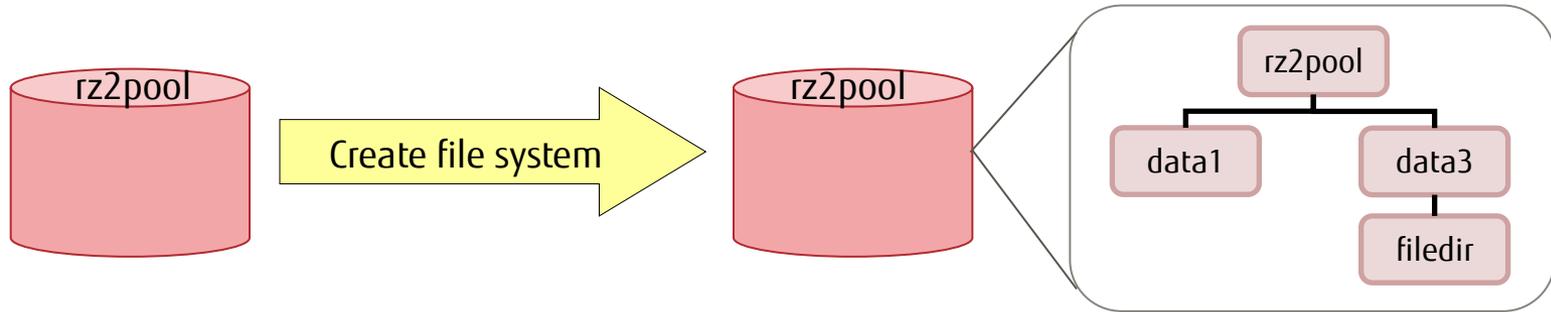
## ■ File system operations

- Creating a file system  
Create a file system in rz2pool, an existing storage pool.
- Changing the file system name, mounting/unmounting, and checking the file system usage
- Setting a property  
Set file system properties to implement various functions. (The functions are mount point change, NFS share setting, file system assignment limit and reservation, user/group assignment limit, detection and elimination data duplication, and data encryption.)
- Checking a file system  
Check the file system usage and properties.
- Deleting a file system

# Creating a File System

## ■ Specify a file system name and create a file system.

- Creating a file system also automatically creates a mount point and mounts the file system.



\* The mount point is created directly under / (root), with the same name as the file system, which is then mounted there.

## ✓ Create a file system (zfs create command).

```
# zfs create rz2pool/data1
```

## ✓ Create a file system (and create an intermediate layer at the same time).

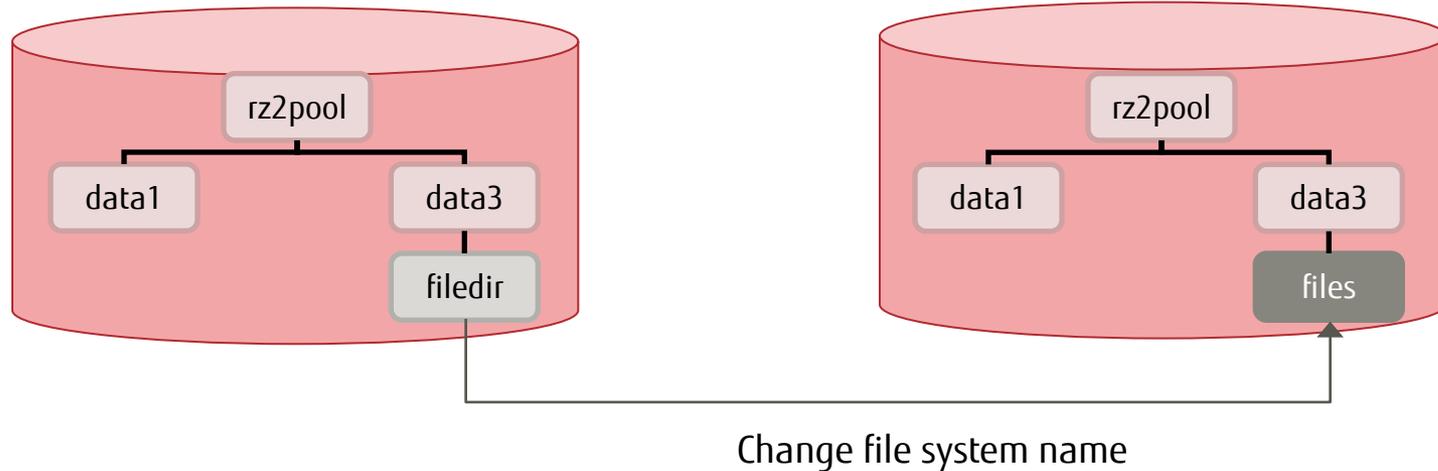
```
# zfs create -p rz2pool/data3/filedir
```



- You can also specify a mount point when creating a file system (mountpoint property).  
# zfs create -o mountpoint=/zfs/fs4 rz2pool/data4

# Changing the File System Name

- You can change the name of a created file system.



- ✓ Change the file system name (zfs rename command).

```
# zfs rename rz2pool/data3/filedir rz2pool/data3/files
```



- Changing the file system name also changes the mount point.

- ZFS is automatically mounted when the OS starts up or a file system is created.
- You can also mount/unmount ZFS manually.
  - ✓ Mount the file system (zfs mount command).

```
# zfs mount rz2pool/data1
```

- ✓ Unmount the file system (zfs unmount command).

```
# zfs unmount rz2pool/data1
```



- For unmounting, you can specify the file system or the mount point. For mounting, you can specify only the file system.
- Mounting can also be done in the same way as in UFS.
  - > For details, see "[Legacy Mount Setting](#)."

This section describes how to set representative file system properties.

- ✓ Set a property (zfs set command and others).

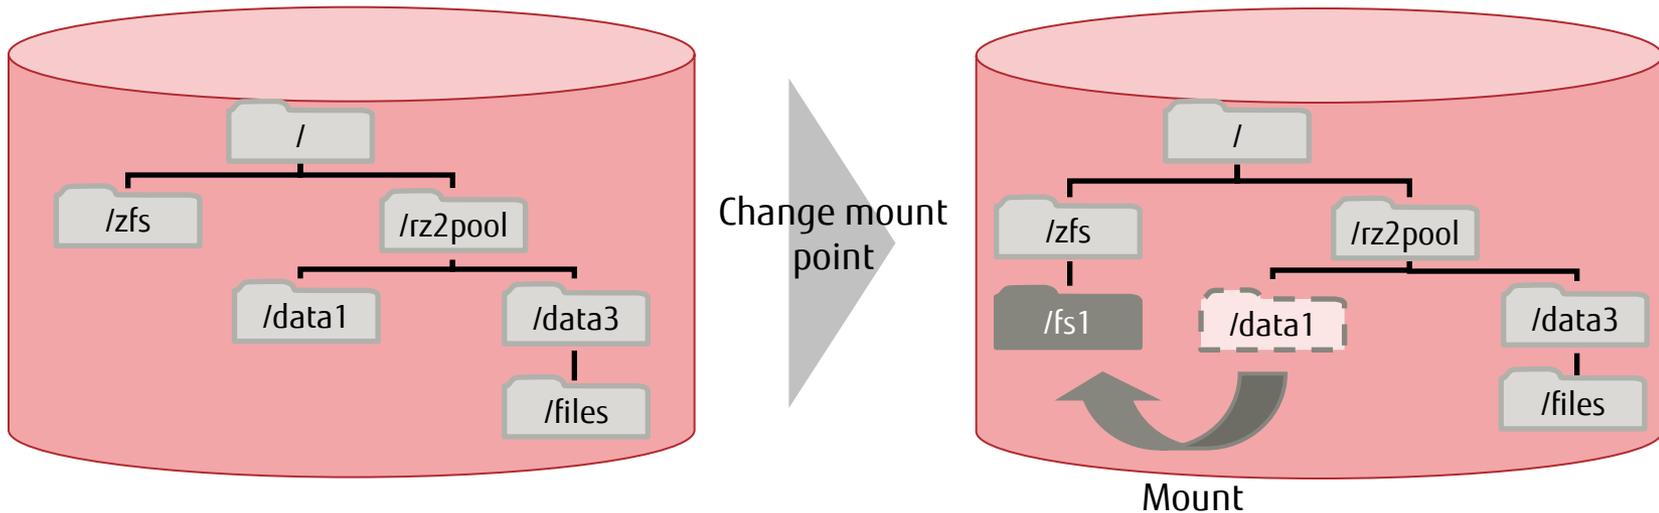
```
# zfs set <property>=value ...
```

- \* Specify a property name and a value. The method of specifying the value depends on the property.
- \* Excluding the zfs set command, commands such as zfs create (creating a file system) may also be able to set properties.

## ■ Representative properties

- Changing a mount point  
mountpoint property
- Configuring NFS sharing  
share.nfs property
- File system assignment limit and reservation  
quota and reservation properties
- User/Group assignment limit  
defaultuserquota and defaultgroupquota properties  
userquota@user and groupquota@group properties
- Detection and elimination of data duplication  
dedup property
- Data encryption  
encryption property

## ■ Changing a mount point



- ✓ Change a mount point (mountpoint property).

```
# zfs set mountpoint=/zfs/fs1 rz2pool/data1
```

\* If no value is specified for the mountpoint property, a mount point with the same name as the file system is created by default directly under / (root).



- Changing the mount point does not change the file system name.

## ■ Configuring NFS sharing

### Solaris 11.1 or later

\* The sharenfs property was changed to the share.nfs property.

- ✓ Configure sharing (share.nfs property).

Syntax: share.nfs=[on|off] *data\_set*

```
# zfs set share.nfs=on rz2pool/data5
```

### Solaris 11 11/11

- Configure sharing in two steps. The first step creates a share (zfs set share), and the second step publishes the NFS share (sharenfs=on).

\* You can use this setting method in Solaris 11.1 too.

- (1) Create sharing (zfs set share command).

Syntax: name=*share\_name*, path=*NFS\_share\_path*, prot=*protocol\_data\_set* (e.g., NFS, SMB)

```
# zfs set share=name=d5,path=/rz2pool/data5,prot=nfs rz2pool/data5
name=d5,path=/rz2pool/data5,prot=nfs
```

- (2) Publish sharing (sharenfs property).

Syntax: sharenfs=[on|off] *data\_set*

```
# zfs set sharenfs=on rz2pool/data5
```



– Data set:  
It is a file system, volume, snapshot, or clone created in a storage pool.

## ■ File system assignment limit and reservation

- ✓ Reserve the area available in the file system (reservation property).

You can reserve the area in advance.

```
# zfs set reservation=500M rz2pool/data3
```

- ✓ Set the upper limit on the available area (quota property).

You can set quota (upper limit setting).

```
# zfs set quota=600M rz2pool/data3
```

## ■ User/Group assignment limit

- ✓ Set the default assignment limit of users (defaultuserquota property).

You can set the upper limit on usage capacity for one user.

```
# zfs set defaultuserquota=25gb rz2pool/data3
```

- ✓ Set the default assignment limit of groups (defaultgroupquota property).

You can set the upper limit on usage capacity for one group.

```
# zfs set defaultgroupquota=50gb rz2pool/data3
```

- ✓ Assignment limit on a specific user (userquota@user property)

You can set the upper limit on usage capacity for a specific user.

```
# zfs set userquota@user1=300M rz2pool/data3
```

If the user belongs to multiple groups, the respective assignment limits of the user, primary group, and secondary group apply. In this case, the lowest value among these property values takes precedence.

Example: Suppose that the assignment limits have the following values. Then, the lowest value is the 300 M specified for the user, and it takes precedence.

User	300 M	<- The lowest value, 300 M, takes precedence.
Primary group	500 M	
Secondary group	700M	

- ✓ Assignment limit on a specific group (groupquota@group property)

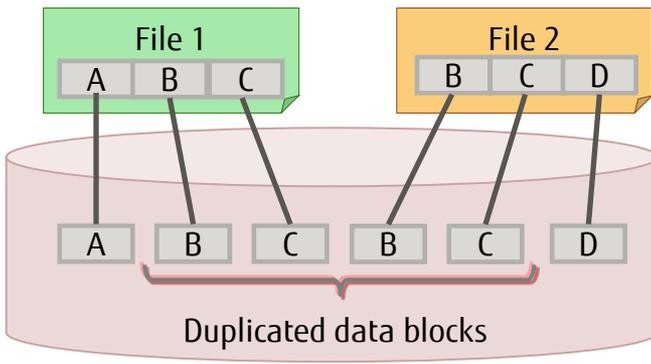
You can set the upper limit on usage capacity for a specific group.

```
# zfs set groupquota@group1=500M rz2pool/data3
```

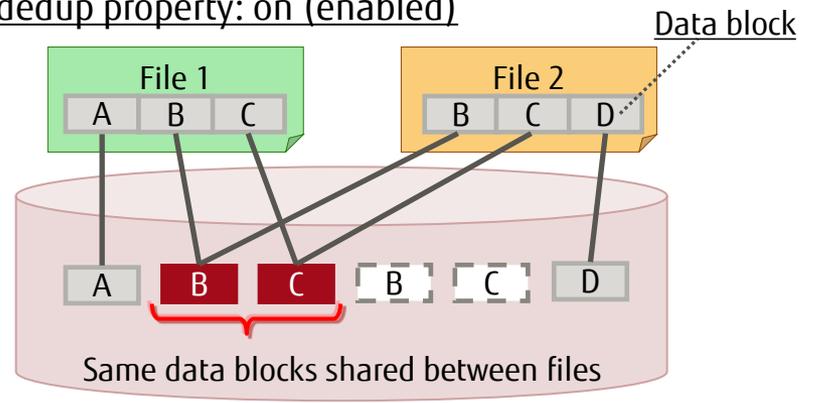
## ■ Detection and elimination of data duplication

- You can set the deduplication function for each data set.
- The function conserves storage pool capacity by detecting and eliminating duplicated data areas (data block) among multiple files.

dedup property: off (disabled) \* Default



dedup property: on (enabled)



- ✓ Set the deduplication function (dedup property).

Syntax: `zfs set dedup=[on | off] [target_file_system_name]`

```
# zfs set dedup=on rpool/data1
```



- This property enables effective utilization of the disk capacity at virtualization and backup.

## ■ Notes on enabling ("on" setting) the dedup property

– Use the zdb command to calculate the deduplication ratio (dedup) and required memory quantity. If the deduplication ratio is 2 or larger and the mounted memory includes an allowance for the memory required, a significant saving of capacity can be achieved with minimum impact on performance.

### ✓ How to calculate the memory required

Memory required = Number of allocated blocks x 320 (\*1)

\*1 320 is a coefficient (dedup table size (bytes)).

```
# zdb -S rpool (Storage pool targeted for deduplication)
Simulated DDT histogram:
bucket          allocated          referenced
-----
refcnt  blocks  LSIZE  PSIZE  DSIZE  blocks  LSIZE  PSIZE  DSIZE
-----
      1   165K  6.93G  6.93G  7.50G   165K  6.93G  6.93G  7.50G
      2    9.66K  337M  337M  379M   20.9K  763M  763M  853M
:
Total   177K  7.29G  7.29G  7.92G   267K  16.3G  16.3G  17.0G
dedup = 2.15, compress = 1.00, copies = 1.05, dedup * compress / copies = 2.05
```

### Left example

Deduplication ratio (dedup): 2.15  
Memory required: 56.64 M  
(177 K x 320 = 56.64 M)

It has a deduplication ratio of at least 2 and an allowance for the memory required.

Enabling the dedup property will have an effect.

- \* Deduplication is done within or between data sets that have the dedup=on setting in the same storage pool. Even if the file systems are different, data sets in the same storage pool are deduplicated.
- \* Deduplication is in effect for the data written after dedup is set to on. Data written before this setting is not deduplicated.
- \* The capacity displayed by the zfs list command does not take deduplicated parts into account. Since reference capacities are displayed, they are larger than the actual amount used.



– For details on how to calculate the required memory capacity, see the following:

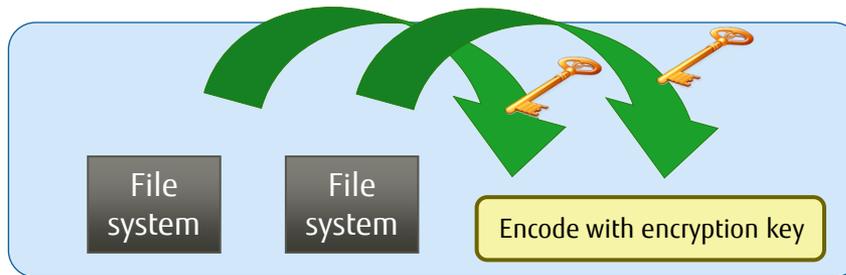
*Managing File Systems in Oracle® Solaris 11.3* (Oracle)

[https://docs.oracle.com/cd/E53394\\_01/pdf/E54801.pdf](https://docs.oracle.com/cd/E53394_01/pdf/E54801.pdf)

## ■ Data encryption

- You can set an encryption policy (\*1) for each data set when creating a file system with the zfs create command.
- The encryption policy (\*1) cannot be changed.

\*1 The encryption policy has an on/off (use/do not use) setting for encryption.



- ✓ Encrypt a file system (encryption property).

```
# zfs create -o encryption=on rz2pool/data4
Enter passphrase for 'rpool/home':*****
Enter again:*****
```

- ✓ Confirm that encryption is enabled for a file system.

```
# zfs get encryption rz2pool/data4
NAME                PROPERTY  VALUE      SOURCE
rz2pool/data4/     encryption on         local
```

## ■ Read-only setting

- ✓ Set read-only for a file system (readonly property).

You can set a file system to read-only.

```
# zfs set readonly=on rz2pool/data1
```

## ■ Data compression

- ✓ Compress file system data (compression property setting).

File system data is automatically compressed.

You can specify the following compression formats:

[on | lzjb | gzip | gzip-[0-9]]

```
# zfs set compression=on rz2pool/data1
```

- \* The values of on and lzib are the same. The values of gzip and gzip-6 are the same.

# Checking a File System 1/2

## ■ You can check the usage and mount point of a file system.

- ✓ Check the information for all file systems (zfs list command).

```
# zfs list
NAME                                USED    AVAIL    REFER    MOUNTPOINT
rpool                                12.2G   12.3G    384K     /rpool
:
rz2pool/data3                        590K    4.86G    303K     /rz2pool/data3
rz2pool/data3/filedir                287K    4.86G    287K
/rz2pool/data3/filedir
```

(1) (2) (3) (4) (5)

### Output format

(1): Data set name                   (2): Size of area used               (3): Size of available area  
(4): Size of area used for data set   (5): Mount point

- ✓ Check the information for the specified file system.

```
# zfs list rz2pool
NAME                                USED    AVAIL    REFER    MOUNTPOINT
rz2pool                             1.65M   4.86G    319K     /rz2pool
```

- ✓ Check the information for the specified file system and everything under it.

```
# zfs list -r rz2pool
NAME                                USED    AVAIL    REFER    MOUNTPOINT
rz2pool                             1.65M   4.86G    319K     /rz2pool
rz2pool/data1                       287K    4.86G    287K     /rz2pool/data1
rz2pool/data3                       590K    4.86G    303K     /rz2pool/data3
rz2pool/data3/filedir                287K    4.86G    287K
/rz2pool/data3/filedir
```

## ■ You can check the property information for a file system by data set.

- ✓ Check the file system property information (zfs get command).

```
# zfs get all rz2pool
NAME      PROPERTY      VALUE          SOURCE
rz2pool   aclinherit    restricted     default
rz2pool   aclmode       discard        default
:
-- <<Omitted>> --
```

(1)	(2)	(3)	(4)
NAME	PROPERTY	VALUE	SOURCE
rz2pool	aclinherit	restricted	default
rz2pool	aclmode	discard	default
:			
-- <<Omitted>> --			

- ✓ Check individually specified property information.

```
# zfs get type rz2pool
NAME      PROPERTY      VALUE          SOURCE
rz2pool   type          filesystem     -
```

(1)	(2)	(3)	(4)
NAME	PROPERTY	VALUE	SOURCE
rz2pool	type	filesystem	-

Displays only the specified property.

### Output format:

(1): Data set name  
(4): Property state

(2): Property name

(3): Property value

default  
local  
-

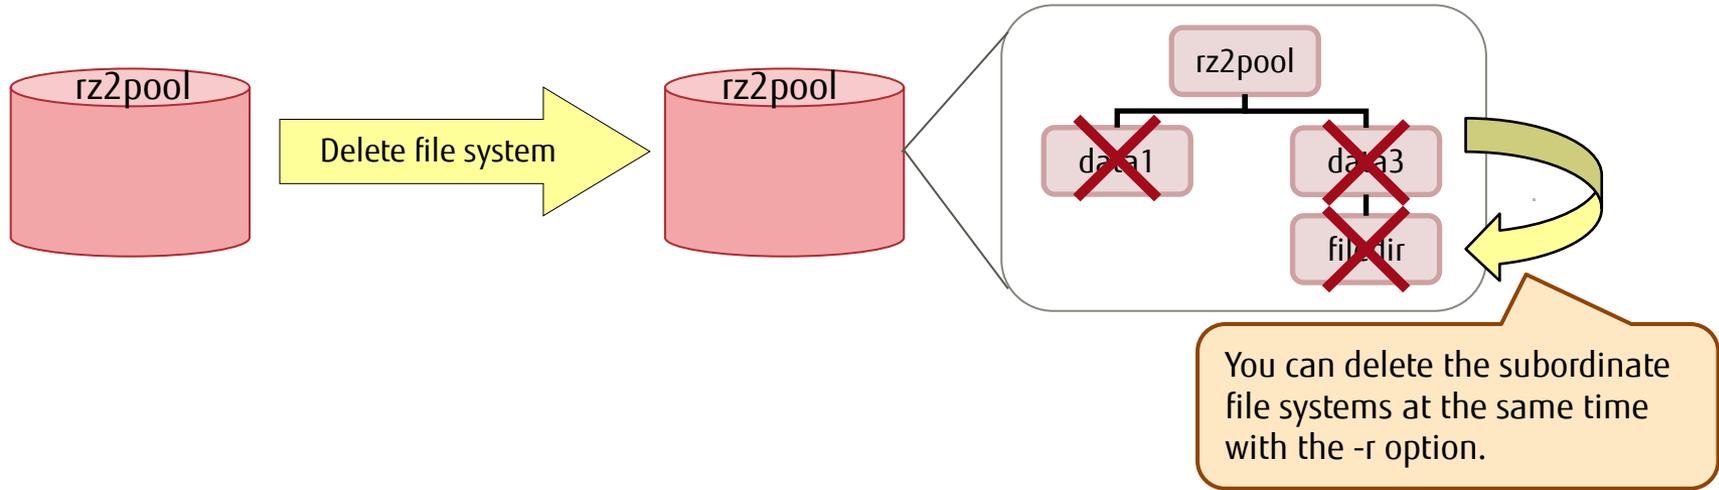
inherited form *data\_set\_name*

Not set explicitly (Default value)  
Explicitly set value  
Read-only

Inherited from displayed data set name

# Deleting a File System

- Use the `zfs destroy` command to delete a file system.



\* The specified file system is unmounted and deleted.

- ✓ Delete a file system (`zfs destroy` command).

```
# zfs destroy rz2pool/data1
```

- ✓ Delete subordinate file systems recursively.

```
# zfs destroy -r rz2pool/data3
```

\* If `data3` is specified and deleted, `filedir` under it is also deleted.

# <<Reference>> File System Properties 1/4

Property Name	Description	Remarks
available	Capacity available for the data set and all its subordinate entities	
compressratio	Compression ratio achieved for the data set	
creation	Time when this data set was created	
defer_destroy	If a snapshot is marked as deferred destruction by the <code>zfs destroy -d</code> command, this property is on. Otherwise, the property is off.	
keychangedate	Date of the last wrap key change made by the <code>zfs key -c</code> command operation on the specified data set. If no key change operation was performed, <code>keychangedate</code> indicates the creation date.	
keystatus	Identifies the status of the encryption key of the data set.	Solaris 11 11/11 and later
mounted	Indicates whether the file system is currently mounted.	
origin	Snapshot from which a clone of a file system or volume was created. This source of the clone cannot be destroyed (even if the <code>-r</code> or <code>-f</code> option is used) as long as the clone exists.	
referenced	Amount of accessible data in this data set	
rekeydate	Date of the last encryption key change made by the <code>zfs key -K</code> or <code>zfs clone -K</code> command operation on this data set	Solaris 11 11/11 and later
type	Type of data set. It is filesystem, volume, or snapshot.	
used	Checks the capacity consumed by this data set and all subordinate entities.	
usedbychildren	Capacity used by the subordinates of this dataset	Solaris 10 10/09 and later
usedbydataset	Capacity used by this data set itself	Solaris 10 10/09 and later
usedbyreservation	Capacity used by the reservation set of this data set	Solaris 10 10/09 and later
usedbysnapshots	Capacity consumed by snapshots of this data set	Solaris 10 10/09 and later

# <<Reference>> File System Properties 2/4

Property Name	Description	Remarks
userused@user	Capacity consumed in this data set by the specified user	Solaris 10 10/09 and later
userrefs	Number of holds on this snapshot by users. A user hold is set by the zfs hold command.	
groupused@group	Capacity consumed in this data set by the specified group	Solaris 10 10/09 and later
volblocksize	Specifies the volume block size, for a volume.	
aclmode	Controls how ACL is changed during chmod(2) execution.	
aclinherit	Controls how ACL entries are inherited when files and directories are created.	
atime	Controls whether to update the file access time when reading a file.	
canmount	Controls whether the specified file system can be mounted by the zfs mount command.	
checksum	Controls the checksum used to verify the integrity of data.	
compression	Controls the compression algorithm used for this data set.	
copies	Controls the number of copies of data saved for this data set.	
dedup	Controls whether to apply deduplication to the data set.	Solaris 11 11/11 and later
devices	Controls whether to allow this file system to open device nodes.	
exec	Controls whether to allow execution of processes from inside this file system.	
logbias	Controls how ZFS optimizes synchronization requests to this data set.	
mlslabel	Secret label that specifies whether the data set can be mounted on a zone	When Trusted Extensions is enabled
mountpoint	Controls the mount points used by the file system.	
nbmand	Controls whether the file system should be mounted by nbmand (non-blocking mandatory lock).	
primarycache	Controls the cache content of the primary cache (ARC).	Solaris 10 10/09 and later

# <<Reference>> File System Properties 3/4

Property Name	Description	Remarks
quota	Limits the capacity consumed by this data set and subordinate entities.	
sync	Specifies the degree of synchronization of file system transactions.	
userquota@user	Limits the capacity consumed by the specified user.	Solaris 10 10/09 and later
groupquota@group	Limits the capacity consumed by the specified group.	Solaris 10 10/09 and later
readonly	Controls whether to allow changes to this data.	Solaris 11 11/11 and later
recordsize	Specifies the recommended block size of files stored in the file system.	
refquota	Limits the capacity that can be consumed by 1 data set.	
reservation	Minimum capacity guaranteed for the data set	
reservation	Minimum capacity guaranteed for the data set and subordinate entities	
rstchown	Specifies whether the file system restricts users from granting ownership of a file with the chown(1) or chown(2) system call.	Solaris 11 11/11 and later
secondarycache	Controls the cached contents in the secondary cache (L2ARC).	Solaris 10 10/09 and later
setuid	Controls whether the set UID bit is respected on this system.	
shadow	Identifies the ZFS file system as the shadow of the file system specified by a URI.	
sharenfs	Controls whether to create and publish the ZFS data set as an NFS share.	
share.nfs	Same as above	Solaris 11.1 and later
sharesmb	Controls whether to create and publish the ZFS data set as an SMB share.	
share.smb	Same as above	Solaris 11.1 and later
snapdir	Controls whether to hide or show the zfs directory at the root of the file system.	
version	Version of this file system on the disk	
volsize	Specifies the logical size of the volume.	
vscan	Controls whether to execute a virus scan on a normal file when the file is opened or closed.	

# <<Reference>> File System Properties 4/4

Property Name	Description	Remarks
xattr	Controls whether extended attributes are valid on the file system.	
zoned	Controls whether to manage the data set from a non-global zone.	
casesensitivity	Specifies whether the file name matching algorithm used by the file system is case-sensitive, case-insensitive, or a combination of both methods of matching.	
normalization	Specifies whether to always perform Unicode normalization of file names and which normalization algorithm to use, when comparing 2 file names in the file system.	
utf8only	Sets whether the file system rejects any file name that contains anything not in the UTF-8 character set.	
encryption	Defines the encryption algorithm and key size used for an encrypted data set.	Solaris 11 11/11 and later
multilevel	Controls whether to individually label with explicit secret label attributes that are automatically generated.	When Trusted Extensions is enabled
keysource	Defines the format and location of the wrapping key for the data set key. This property must be specified when the file system is created.	Solaris 11 11/11 and later
defaultuserquota	Makes the default assignment limit on users.	Solaris 11.3 and later
defaultgroupquota	Makes the default assignment limit on groups.	Solaris 11.3 and later

# <<Reference>> Legacy Mount Setting

- You can also mount a file system in the same way as in UFS.
  - Legacy mount is used to manage mounting with /etc/vfstab and the mount command like in the conventional file system (UFS).
  - Specify "legacy" for the mountpoint property to set legacy mount.
  - Set legacy mount in cases such as specifying a shared file system in a Solaris zone.
- \* If you set legacy mount, use the mount/unmount command instead of the zfs mount/zfs unmount command for manual mounting/unmounting.

## ✓ Legacy mount setting (mountpoint property)

```
# zfs set mountpoint=legacy rz2pool/data4
```

## ✓ Mounting (mount command)

```
# mount -F zfs rz2pool/data4 /zfs/legacy
```

Specify "zfs" for the file system type.

## ✓ Unmounting (umount command)

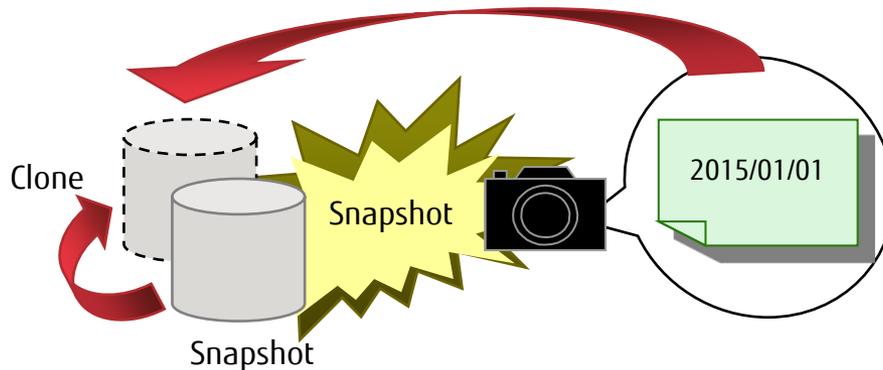
```
# umount /zfs/legacy
```

# 3. Creating Snapshots and Clones

This chapter describes how to create and manipulate snapshots and clones.

## ■ Creating snapshots and clones

This chapter describes how to create and manipulate snapshots and clones.

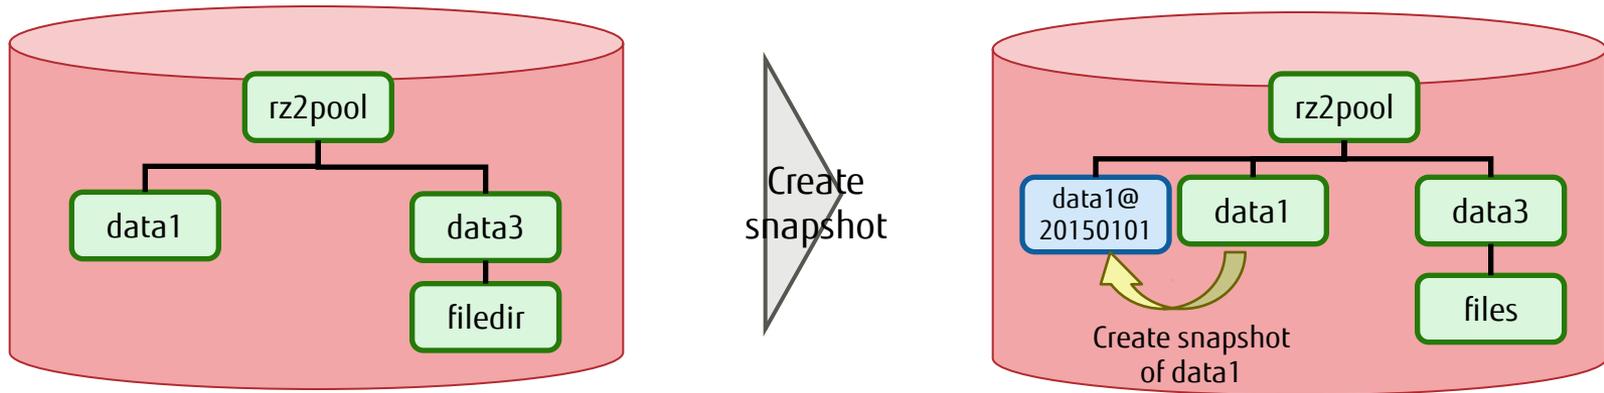


## ■ Snapshot operations and clone operations

- Creating a snapshot
- Changing the snapshot name
- Snapshot differential display
- Creating a clone
- Replacing a file system
- Rollback to a snapshot
- Setting automatic snapshots

# Creating a Snapshot

- Specify a data set (mainly a file system) to create a snapshot.



- ✓ Create a snapshot (zfs snapshot command).

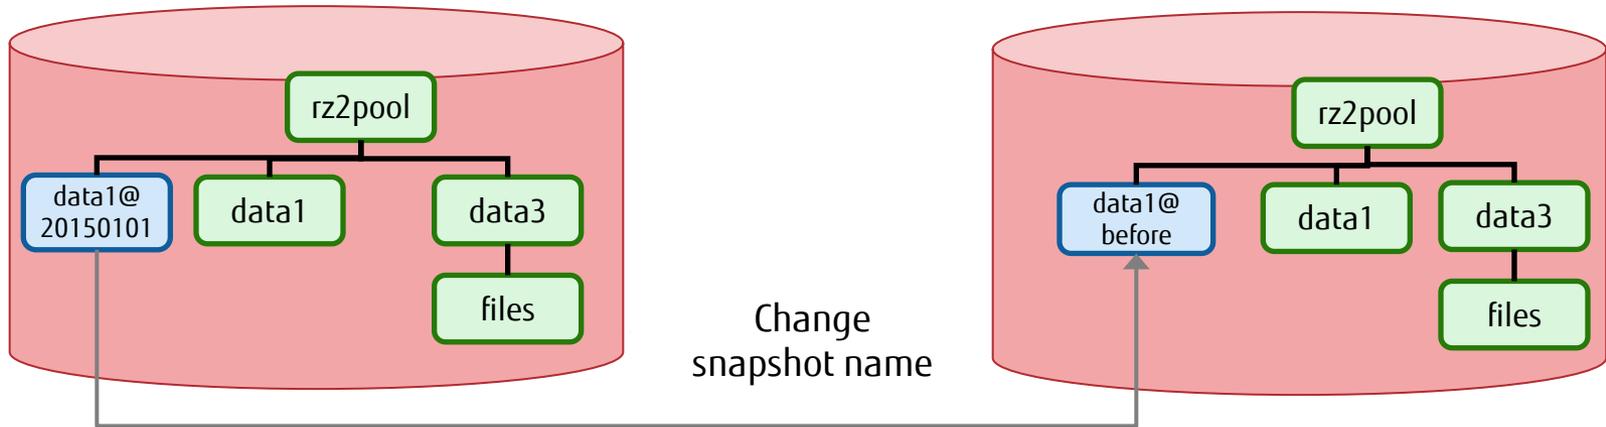
Syntax: `zfs snapshot [file_system_name@snapshot_name]`

```
# zfs snapshot rz2pool/data1@20150101
```

- 💡 - The snapshot is created in the same storage pool as the specified data set.
- By specifying the `-r` option, you can create snapshots of all the data sets under the specified data set simultaneously.  
# `zfs snapshot -r rz2pool/data1@snap1`

# Changing the Snapshot Name

- After creating a snapshot, you can change the snapshot name.



- ✓ Change the snapshot name (zfs rename command).

Syntax: `zfs rename [snapshot_name_before_change][snapshot_name_after_change]`

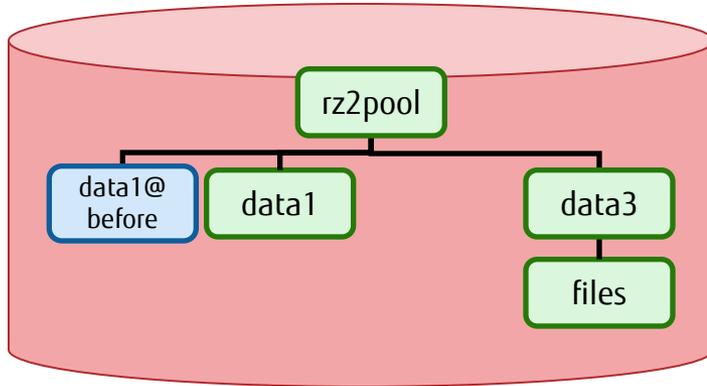
```
# zfs rename rz2pool/data1@20150101 rz2pool/data1@before
```

- A snapshot is created using the format of *data\_set\_name@snapshot\_name*. So, changing the snapshot name changes the snapshot name portion after @.

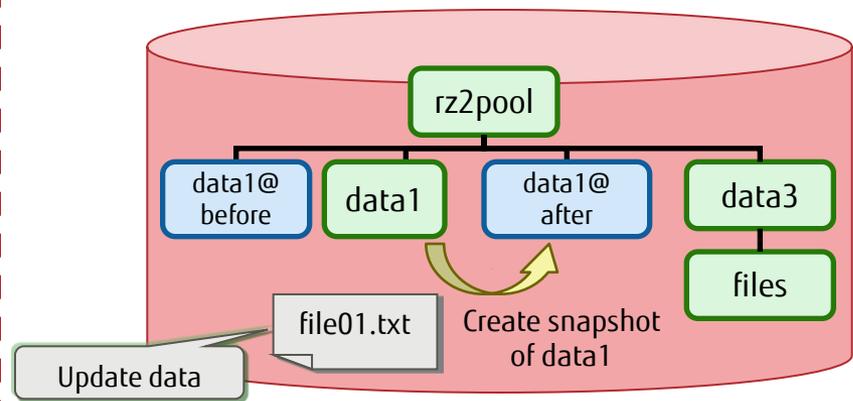
# Snapshot Differential Display

- You can check the differences between two snapshots.

Snapshot A (before data update)  
rz2pool/data1@before



Snapshot B (after data update)  
rz2pool/data1@after



- ✓ Check snapshot differences (zfs diff command).

Syntax: `zfs diff [snapshot_name_1][snapshot_name_2]`

```
$ zfs diff rz2pool/data1@before rz2pool/data1@after
M /rz2pool/data1/
+ /rz2pool/data1/file01.txt
```

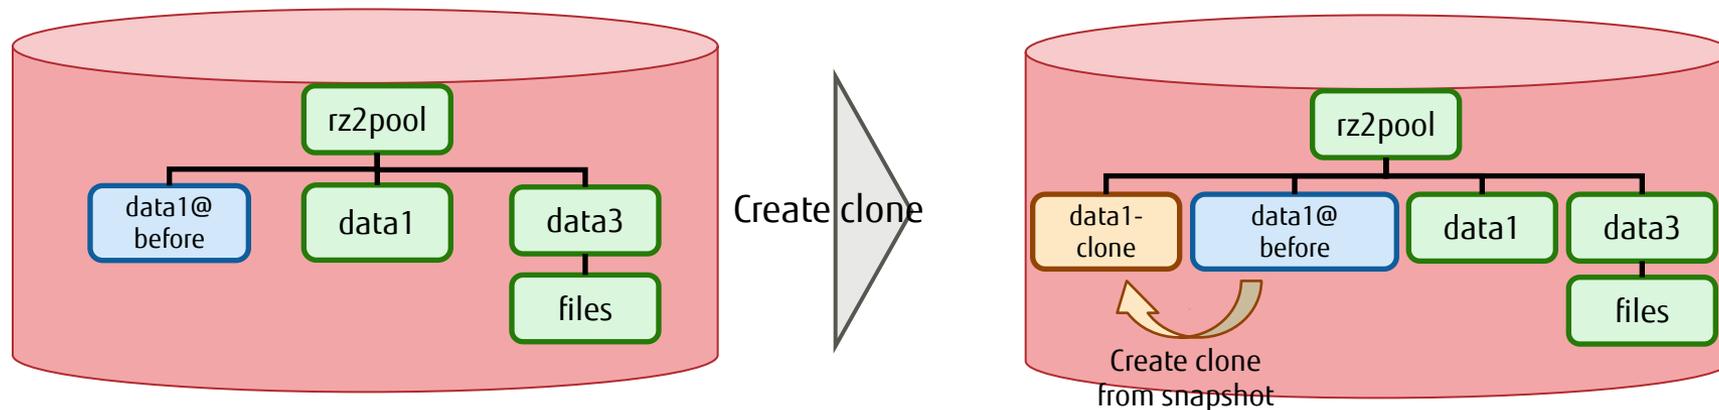
## Snapshot differential indicators

- M: Indicates a change of a file or a directory.
- R: Indicates that a file name or directory name has changed.
- : Indicates that a file or directory resides in an old snapshot but not in a new snapshot.
- +: Indicates that a file or directory resides in a new snapshot but not in an old snapshot.

- 💡 - If the `zfs diff` command is executed with only one snapshot specified, the command displays the differences between the snapshot and the current file system.
- The `-r` option displays the differences between the specified snapshot and all child snapshots.

# Creating a Clone

- Create a clone of a data set from a snapshot.



- ✓ Create a clone (zfs clone command).

Syntax: `zfs clone [snapshot_name][file_system_name_of_clone]`

```
# zfs clone rz2pool/data1@before rz2pool/data1-clone
```

- ✓ Create a clone of an encrypted file system.

```
# zfs snapshot rz2pool/data4@encryption
# zfs clone rz2pool/data4@encryption rz2pool/data4-clone
Enter passphrase for `rz2pool/data4-clone':*****
Enter again:*****
```

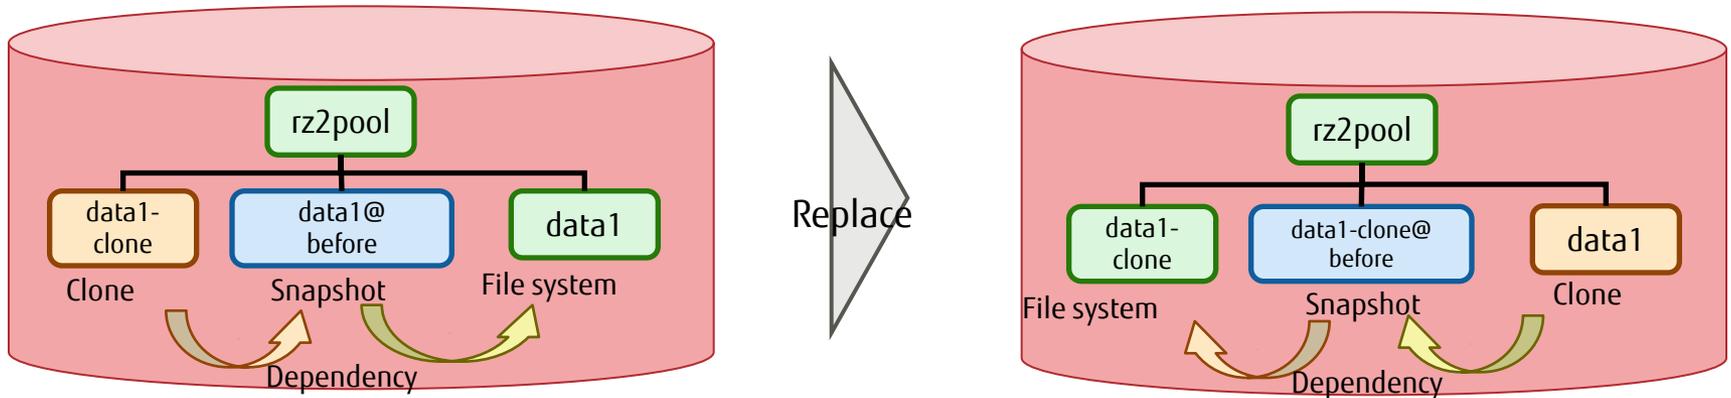
- \* When you create a clone of an encrypted file system, a policy is applied to request a password.  
-> For details on encryption of the file system, see "[Setting a Property 8/9 - Data Encryption](#)."



- You can also create a clone under a different file system within the same specified storage pool.  
Example: `# zfs clone rz2pool/data1@before rz2pool/data3/data1-clone`

# Replacing a File System

- Invert the dependency between the source file system (master) and a clone so that the clone becomes the master file system.



- ✓ Replace the source file system with a clone (zfs promote command).

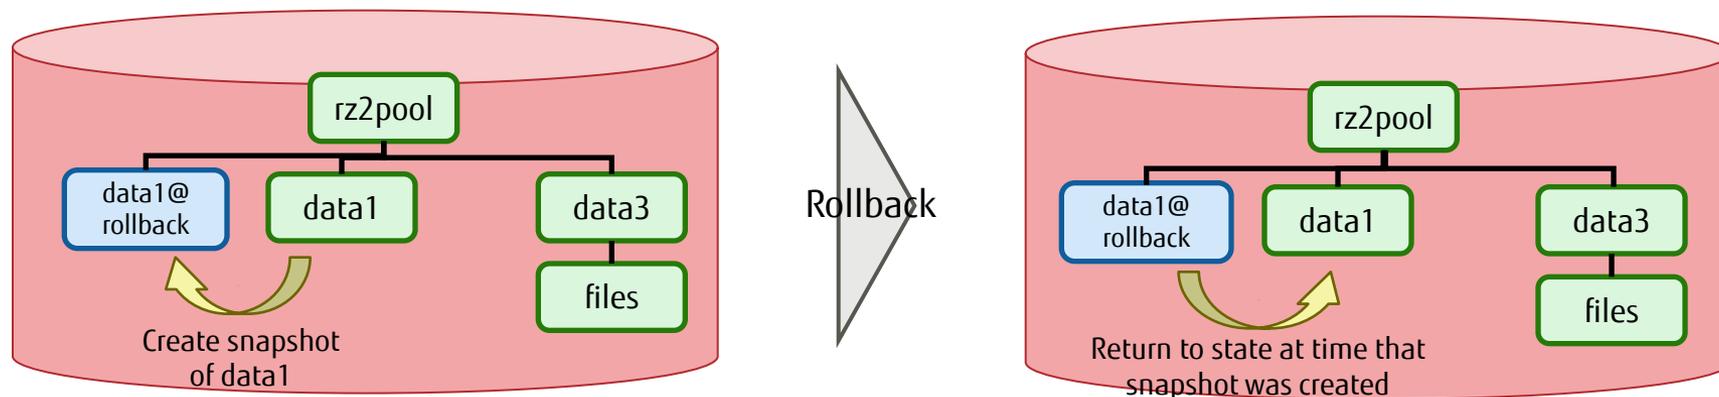
Syntax: `zfs promote [file_system_name_of_clone]`

```
# zfs promote rz2pool/data1-clone
```

- 💡 - With ZFS, you can invert the dependency between the clone and the file system. After replacement, the snapshot name is also changed.
- After replacement, you need to set the file system properties appropriate to the environment because the clone and file system have different properties.

# Rollback From a Snapshot

- A rollback from a snapshot restores the file system to the state of the time that the snapshot was created.



✓ Rollback (zfs rollback command)

Syntax: `zfs rollback [snapshot_name]`

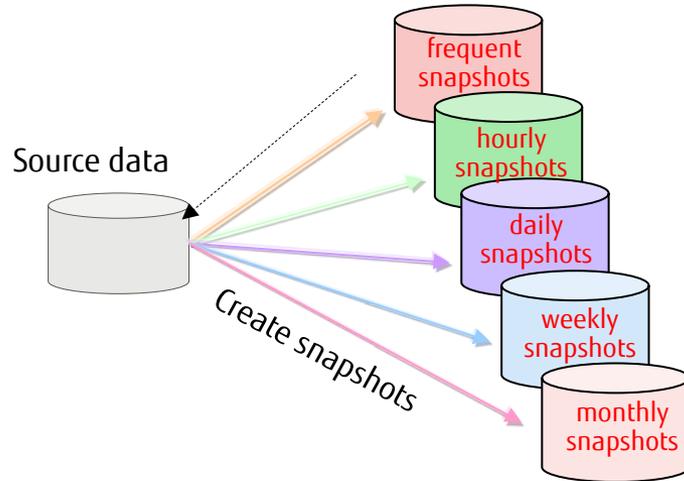
```
# zfs rollback rz2pool/data1@rollback
```



- With ZFS, you can roll back a file system from a snapshot.
- After the rollback, the snapshot remains on the rolled-back system.

# Automatic ZFS Snapshot

- Automatically create snapshots at a certain interval.



	Creation Timing	Quantity That Can be Retained
frequent snapshots	Every 15 minutes	4
hourly snapshots	Every hour	24
daily snapshots	Every day	31
weekly snapshots	Every week	7
monthly snapshots	Every month	12

- ✓ Change the automatic snapshot property setting (com.sun:auto-snapshot property).

Syntax: `zfs set com.sun:auto-snapshot=[true | false][target_file_system_name]`

```
# zfs set com.sun:auto-snapshot=true rz2pool/data1
```

- ✓ Enable the automatic snapshot service.

Syntax: `svcadm enable svc:/system/filesystem/zfs/auto-snapshot:creation_timing`

```
# svcadm enable svc:/system/filesystem/zfs/auto-snapshot:frequent
# svcadm enable svc:/system/filesystem/zfs/auto-snapshot:hourly
```

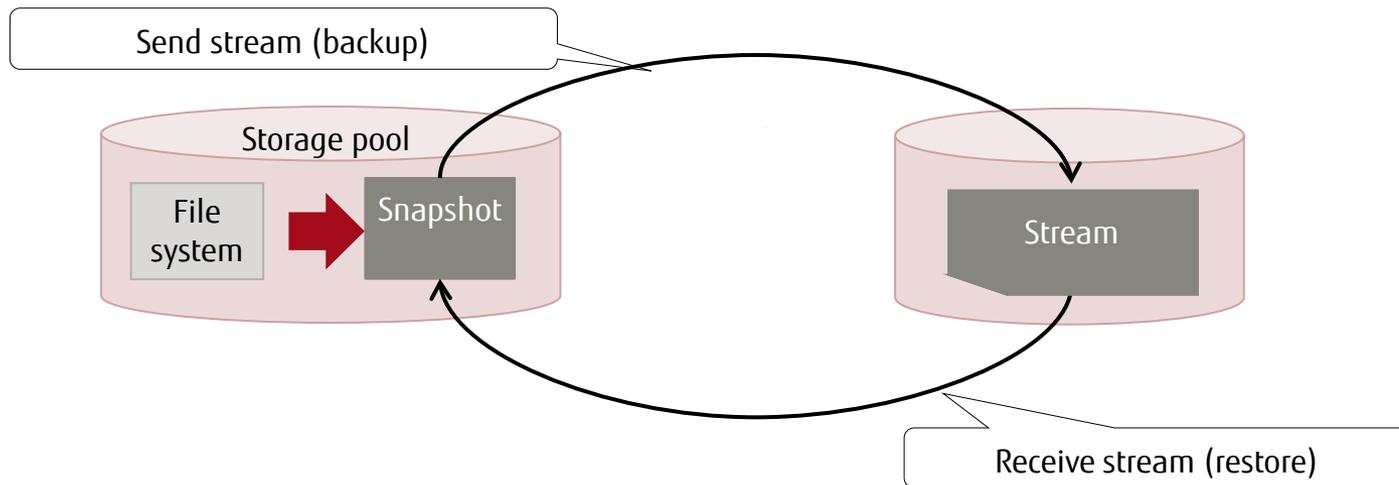
- 💡 - The automatic snapshot function can effectively use disk space since you can specify the applicable data sets and creation timing for the function.

## 4. Backup/Restore

ZFS backs up/restores a file system by sending/receiving a stream.  
This chapter describes how to send/receive a stream.

## ■ Flow of backup/restore

ZFS backs up/restores a file system by sending/receiving a stream.  
This chapter describes how to send/receive a stream.

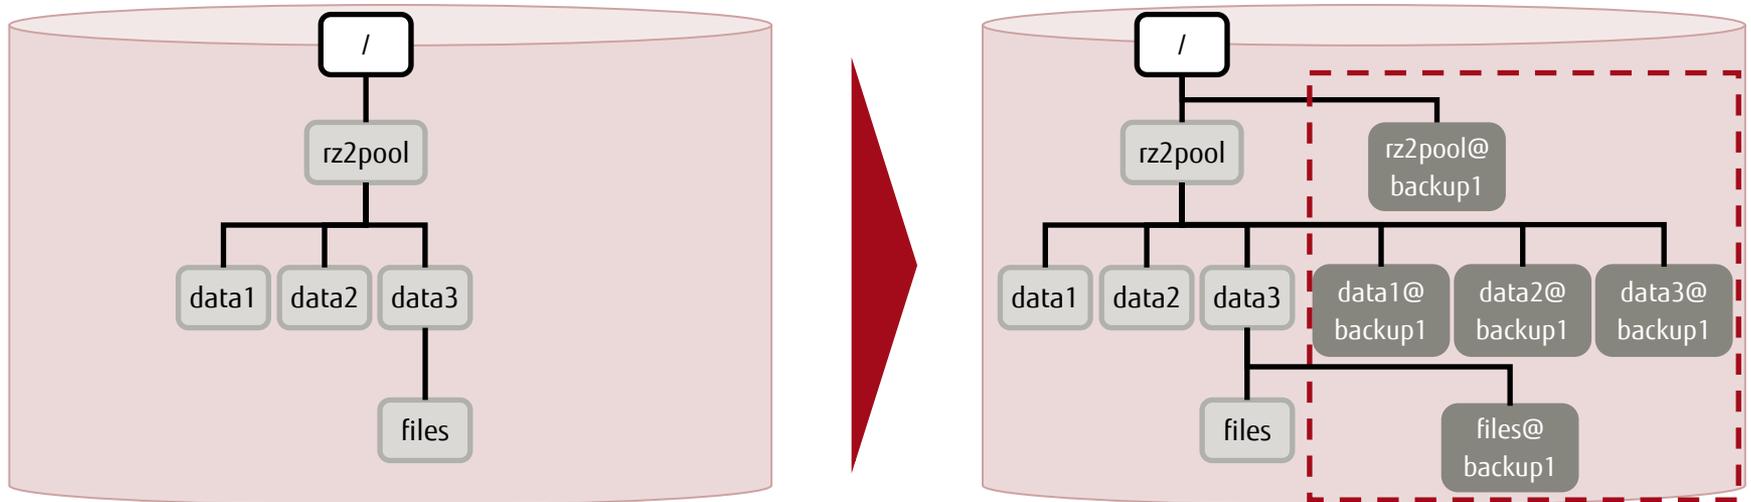


## ■ Backup/Restore methods

- Method 1: Full backup/restore  
Back up/restore all the file systems within a storage pool.
- Method 2: Specific file system backup/restore
- Method 3: Differential backup/restore  
Perform a differential backup/restore of all file systems within a storage pool.

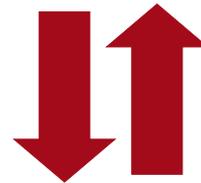
# Method 1: Full Backup/Restore 1/2

- Back up/restore all file systems within a storage pool.



(1) Create snapshot of whole file system

(2) Send stream



(3) Receive stream



## ■ Backup

### (1) Create a snapshot of the whole file system.

Specify the `-r` option to create a snapshot of the specified file system and snapshots of all the lower-level file systems simultaneously.

```
# zfs snapshot -r rz2pool@backup1
```

### (2) Send a stream.

The `-R` option sends a stream of the specified snapshot and of the lower-level subordinate file systems, and then saves it in one file.

```
# zfs send -R rz2pool@backup1 > /mnt/snap2_1.dat
```

\* The stream can be sent to any destination.

## ■ Restore

### (3) Receive a stream.

If the `-R` option was specified when the stream was sent, specify the `-d` option.

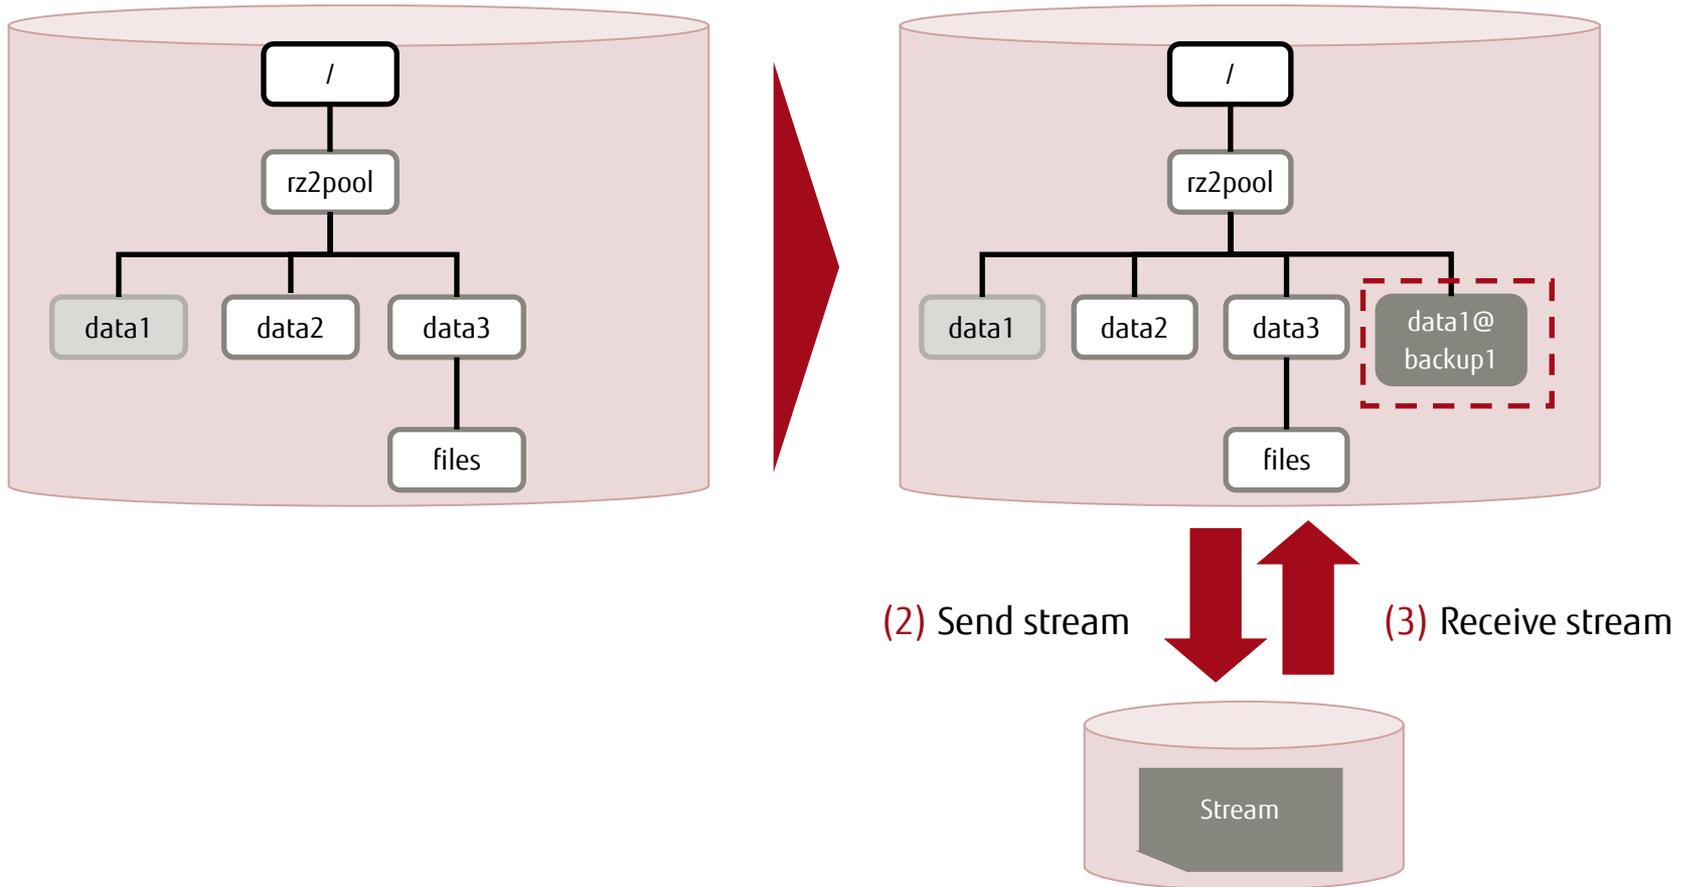
The `-F` option rolls back the file system to the latest state before receiving the stream.

```
# zfs receive -d -F rzpool < /mnt/snap2_1.dat
```

\* Without the `-d` option specified, reception of a stream created with the `-R` option will fail with the output of the following error:  
**cannot receive: must use -d to receive replication (send -R) stream**

## ■ Back up/restore a specific file system.

- (1) Create snapshot of specific file system (rz2pool/data1)



## ■ Backup

(1) Create a snapshot of a specific file system.

```
# zfs snapshot rz2pool/data1@backup1
```

(2) Send a stream.

```
# zfs send rz2pool/data1@backup1 > /mnt/snap3_1.dat
```

## ■ Restore

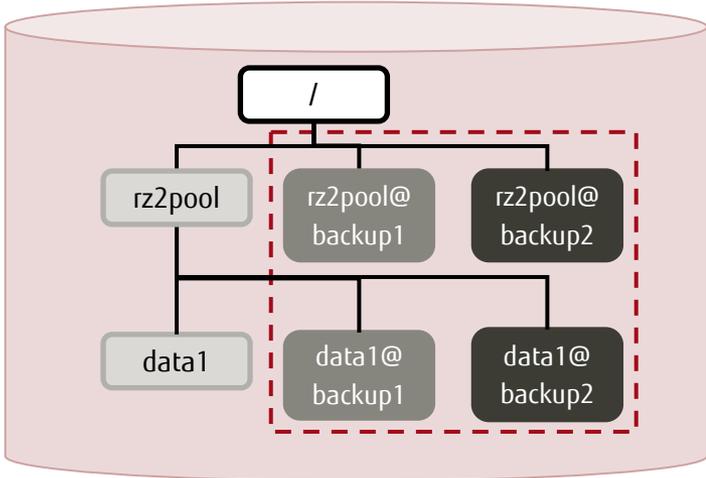
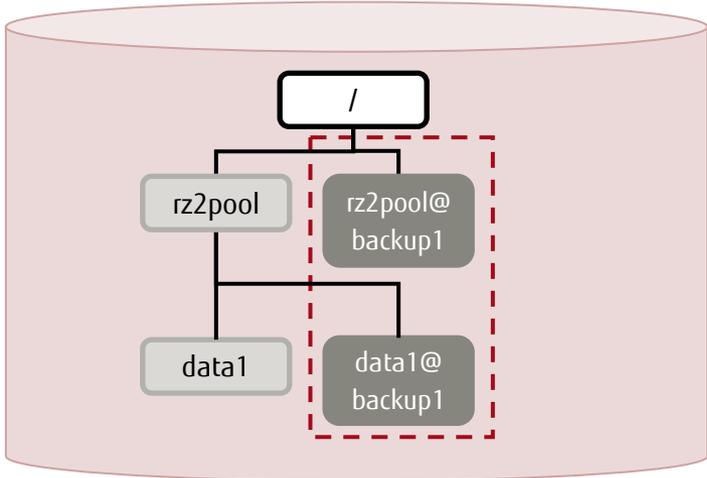
(3) Receive a stream.

```
# zfs receive rzpool/data1 < /mnt/snap3_1.dat
```

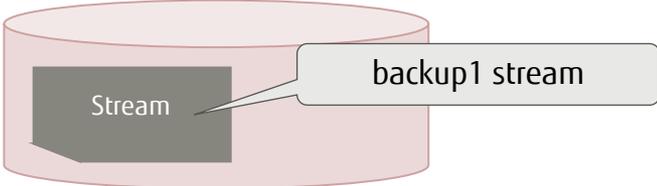
# Method 3: Differential Backup/Restore 1/3

■ Perform a differential backup/restore of all file systems within a storage pool.

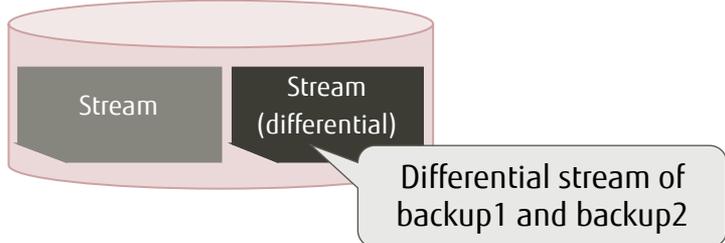
- (1) Create snapshot (backup1)
- (3) Create snapshot (backup2)



(2) Send stream  
(5) Receive stream



(4) Send differential stream  
(6) Receive differential stream



## ■ Backup

(1) Create a snapshot.

```
# zfs snapshot -r rz2pool@backup1
```

(2) Send a stream.

```
# zfs send -R rz2pool@backup1 > /mnt/snap4_1.dat
```

(3) Create a snapshot.

```
# zfs snapshot -r rz2pool@backup2
```

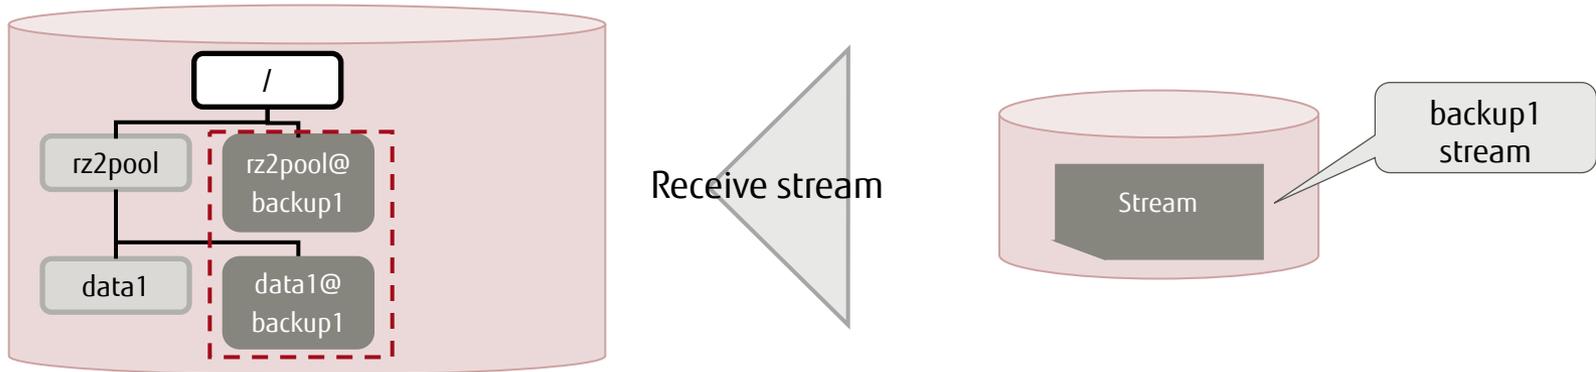
(4) Send a differential stream.

```
# zfs send -R -i rz2pool@backup1 rz2pool@backup2 > /mnt/snap4_2.dat
```

## ■ Restore

### (5) Receive a stream.

```
# zfs receive -d -F rz2pool < /mnt/snap4_1.dat
```



### (6) Receive a differential stream.

```
# zfs receive -d -F rz2pool < /mnt/snap4_2.dat
```



- You can check the progress of the zfs send command. It is useful for improving the accuracy of ZFS save-related schedules and plans.

```
# zfs send -Rv rpool@backup | gzip | pv > /upool/backup/rpool.zfs.gz
sending full stream to rpool@backup
sending full stream to rpool/OVM@backup
sending full stream to rpool/OVM/inst_svr@backup
sending full stream to rpool/OVM/uar2@backup
sending full stream to rpool/OVM/uar1@backup
sending full stream to rpool/VARSHARE@backup
sending full stream to rpool/VARSHARE/pkg@backup
sending full stream to rpool/VARSHARE/pkg/repositories@backup
sending full stream to rpool/VARSHARE/zones@backup
sending full stream to rpool/export@backup
sending full stream to rpool/export/home@backup
sending full stream to rpool/ROOT@backup
sending full stream to rpool/ROOT/solaris@install
sending full stream to rpool/ROOT/solaris@2015-01-05-16:14:20
sending @2015-01-05-16:14:20 to rpool/ROOT/solaris@backup
sending full stream to rpool/ROOT/solaris/var@install
sending @install to rpool/ROOT/solaris/var@2015-01-05-16:14:20
sending @2015-01-05-16:14:20 to rpool/ROOT/solaris/var@backup
estimated stream size:108G
11.6GB 0:23:37 [3.61MB/s] [ <=> ]
```

Displays size of zfs stream file created after command execution (size before compression)

Displays all snapshots of backup targets immediately after command execution

Displays size, elapsed time, and speed while creating



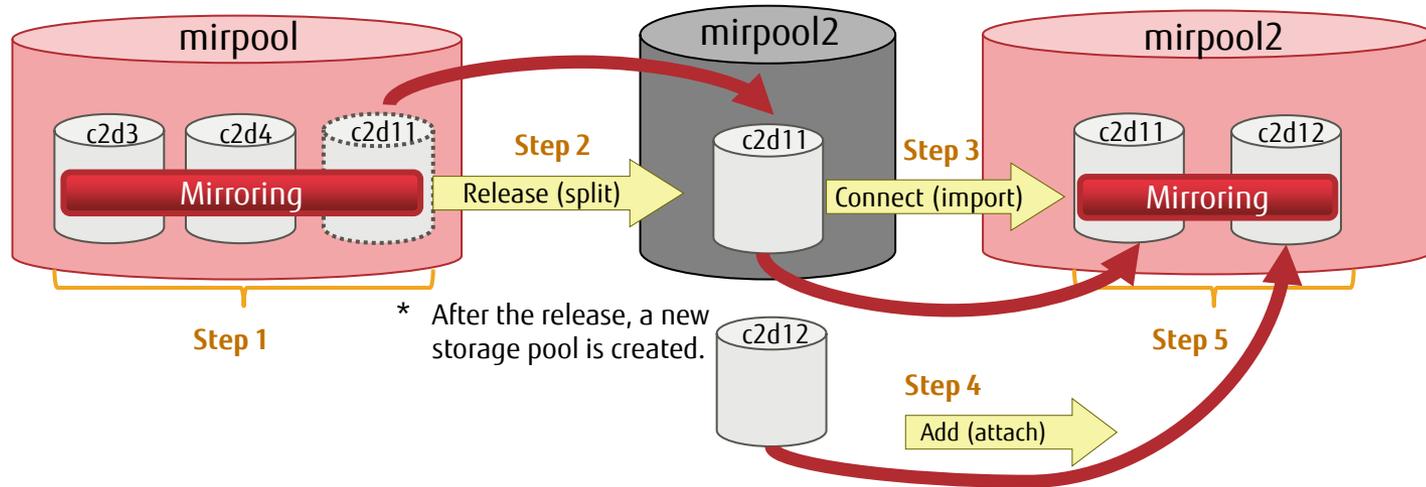
- The added function of checking progress with the zfs send command has been available since Solaris 11.2.

# 5. Releasing a Mirror Disk

This chapter describes how to release a mirror disk and use this disk to make a new storage pool.

## ■ Environment of releasing a mirror disk

This chapter describes how to release a mirror disk and use this disk to make a new storage pool. The following figure shows the configuration to be built.



\* You can release a disk from a storage pool in a mirror configuration and import the released disk to create a new storage pool that holds the same file systems and data as in the original storage pool. After the import, add and incorporate another disk to make a two-way mirror configuration.

## ■ Flow of releasing a mirror disk

- Step 1: Change the mirror configuration (from a two-way mirror configuration to a three-way mirror configuration).
- Step 2: Release a disk from the three-way mirror configuration, and create a new storage pool.
- Step 3: Import the disk to the new storage pool.
- Step 4: Add a disk to the new storage pool to create a two-way mirror configuration.
- Step 5: Check the new storage pool configuration.

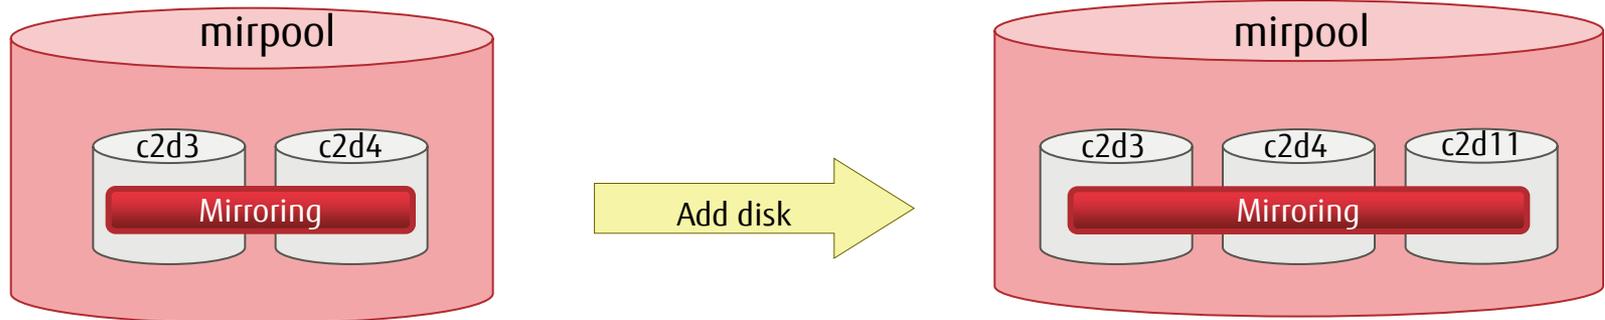


- You can easily move the new storage pool to another server.

# Step 1: Change the Mirror Configuration

- Create a three-way mirror configuration from the two-way mirror configuration to prepare for releasing a mirror disk.

Conceptual illustration of adding a mirror disk



- ✓ Add a disk (zpool attach command).

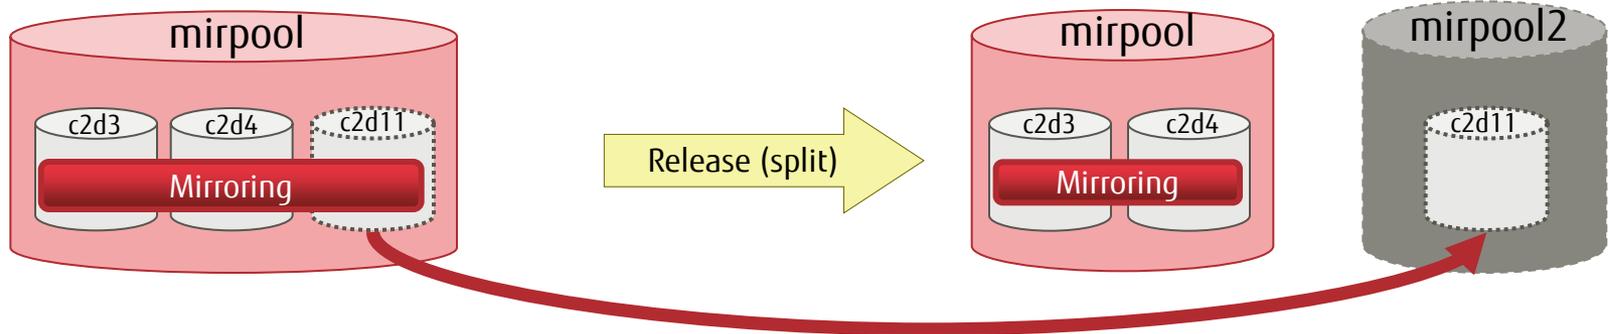
Syntax: `zpool attach pool_name mirror_source_disk mirror_disk`

```
# zpool attach mirpool c2d3 c2d11
```

- 💡 - You can release a mirror disk from a two-way mirror configuration, but after it is released, disk redundancy is lost. If you want to retain disk redundancy, we recommend using a three-way mirror configuration.

# Step 2: Release a Disk From a Three-way Mirror Configuration

- Release a disk from the three-way mirror configuration, and create a new storage pool.
- Conceptual illustration of releasing a mirror disk



- ✓ Release a storage pool (zpool split command).

Syntax: `zpool split storage_pool_name new_storage_pool_name name_of_disk_to_release`

```
# zpool split mirpool mirpool2
```

- \* By default, if no disk name is specified for release, the command releases the last disk connected to the storage pool (c2d11 in this example). The order of display by the zpool status command is the order in which the disks were connected.
- \* The created storage pool (mirpool2) holds the same data as the source storage pool (mirpool).

- ✓ <<Reference>> Specify a disk to release it from a storage pool.

```
# zpool split mirpool mirpool2 c2d11
```



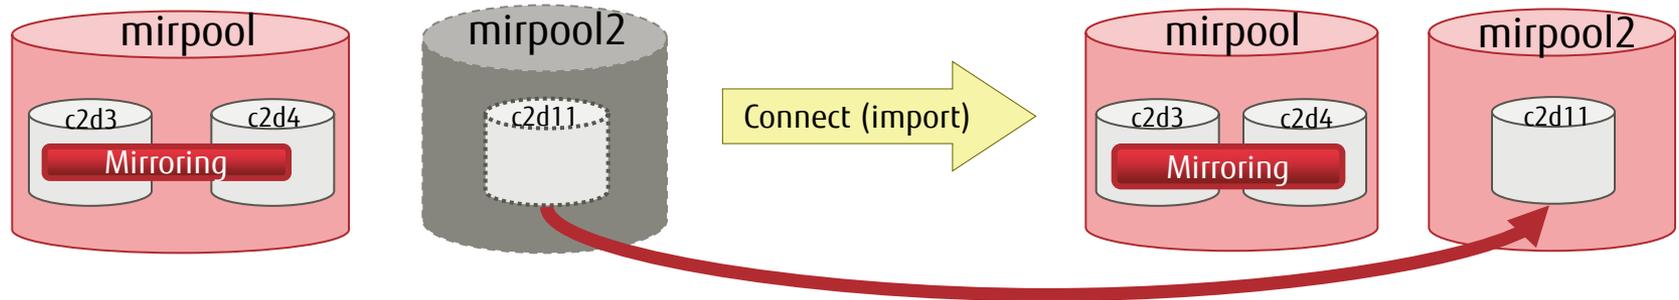
- You can release a disk from a storage pool, with a light load.  
-> For details on the disk load during disk release, see "[<<Reference>> Disk Load Factor in zpool split Execution.](#)"
- A new storage pool is created and exported (export) at the same time that a disk is released.  
\* If resynchronization is in progress, the disk cannot be released.

# Step 3: Import the Released Disk

## ■ Import the released disk to a new storage pool.

- The new storage pool created by the release of a mirror disk is unconnected (exported) when created.
- To use the new storage pool, import it manually.

### Conceptual illustration of recognizing a mirror disk



## ✓ Import a released disk (zpool import command).

Syntax: `zpool import new_storage_pool_name`

```
# zpool import mirpool2
```

\* If no pool name is specified, you can check the storage pools that can be imported.

```
# zpool import
-- <<Omitted>> --
  mirpool2      ONLINE
  c2d11         ONLINE
```

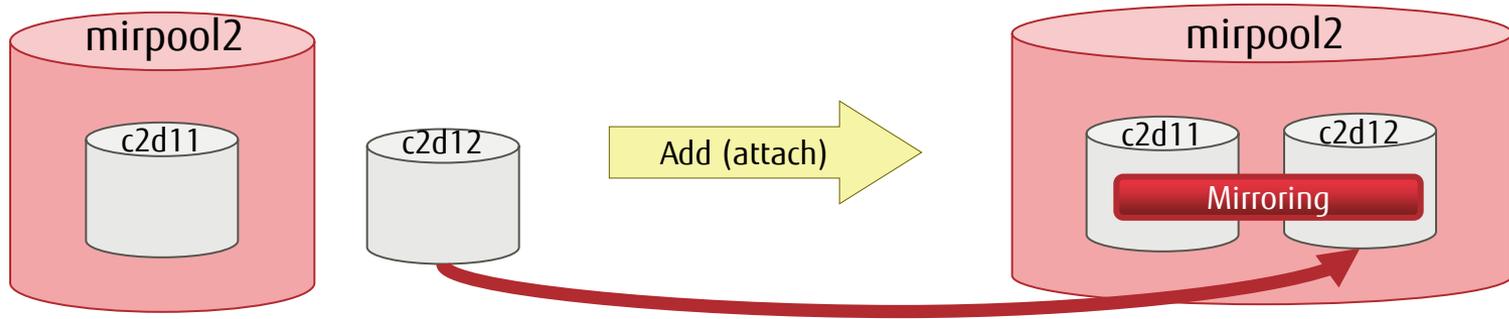


- You can import a disk, with a light load.

-> For details on the disk load during import of a released disk, see "[<<Reference>> Disk Load Factor in zpool import Execution.](#)"

- Add a disk to the new storage pool, and create a two-way mirror configuration.

## Conceptual illustration of creating a 2-way mirror configuration



- ✓ Add a disk (zpool attach command).

Syntax: `zpool attach pool_name mirror_source_disk mirror_disk`

```
# zpool attach mirpool2 c2d11 c2d12
```

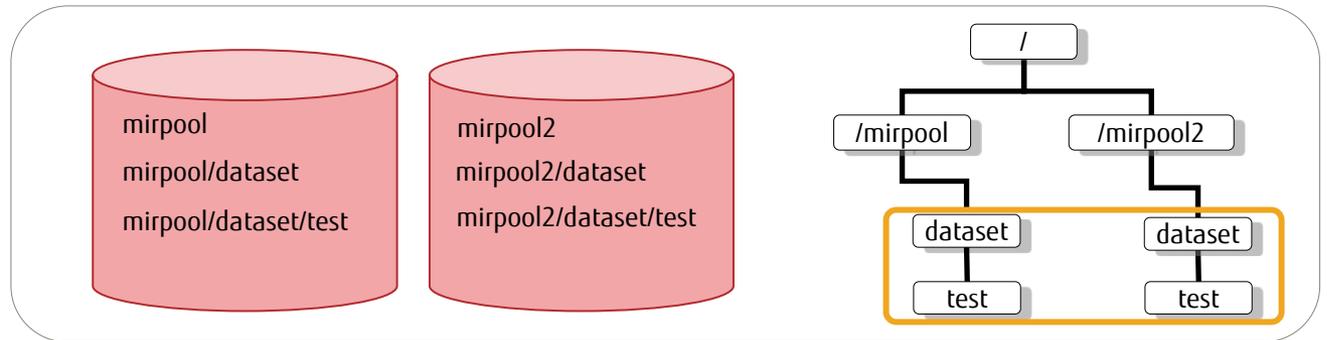
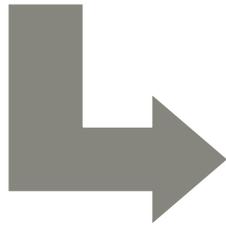
# Step 5: Check the New Storage Pool Configuration

## ■ Check the new storage pool configuration.

- ✓ Check the file systems (zfs list command).

Syntax: zfs list

```
# zfs list
NAME                USED      AVAIL     REFER    MOUNTPOINT
mirpool              173K     9.78G    23K      /mirpool
mirpool/dataset      42K     9.78G    21K      /mirpool/dataset
mirpool/dataset/test 21K     9.78G    21K      /mirpool/dataset/test
mirpool2             175K     9.78G    23K      /mirpool2
mirpool2/dataset     44K     9.78G    23K      /mirpool2/dataset
mirpool2/dataset/test 21K     9.78G    21K      /mirpool2/dataset/test
```



- Confirm that mirpool and mirpool2 have the same file system configuration.

# <<Reference>> Disk Load Factor in zpool split Execution

■ You can release a disk from a storage pool, with a light load.

✓ Disk load factor during five seconds of zpool split execution

iostat -sxn 1 20 command output

```
extended device statistics
r/s   w/s   kr/s  kw/s  wait  actv  wsvc_t  asvc_t  %w  %b  device
0.2   0.2   1.7   2.5   0.0   0.0   0.0     25.7    0  0   c0d3
0.2   0.2   1.6   2.5   0.0   0.0   0.0     26.3    0  0   c0d4
0.5   0.1   5.7   2.0   0.0   0.0   0.0     11.3    0  0   c0d11

4.0   27.9  135.6  77.8  0.0   1.2   0.0     37.8    0  32  c0d3
4.0   27.9  135.6  77.8  0.0   1.4   0.0     44.7    0  37  c0d4
0.0   0.0   0.0   0.0   0.0   0.0   0.0     0.0     0  0   c0d11

0.0   19.0   0.0  473.1  0.0   0.8   0.0     40.4    0  20  c0d3
0.0   19.0   0.0  473.1  0.0   0.8   0.0     41.5    0  22  c0d4
16.0  61.0  922.6  123.5  0.0   2.2   0.0     28.2    0  58  c0d11

0.0   37.1   0.0  524.0  0.0   1.5   0.0     39.8    0  35  c0d3
0.0   37.1   0.0  524.0  0.0   1.7   0.0     45.5    0  41  c0d4
0.0   51.1   0.0  535.6  0.0   1.6   0.0     31.6    0  42  c0d11

0.0   30.9   0.0   60.8  0.0   1.5   0.0     48.7    0  35  c0d3
0.0   30.9   0.0   60.8  0.0   1.5   0.0     49.4    0  35  c0d4
0.0   40.9   0.0  524.1  0.0   1.2   0.0     29.2    0  35  c0d11
```

## Column contents

Item	Description
r/s	Read (times/second)
w/s	Write (times/second)
kr/s	Read (kB/second)
kw/s	Write (kB/second)
wait	Average number of waiting transactions
actv	Average number of transactions in progress
wsvc_t	Average service time (millisecond)
asvc_t	Average service time (millisecond)
%w	Transaction wait time percentage
%b	Disk busy (transaction in progress) percentage
device	Disk used



- A disk load is generated for only several seconds (about four seconds in the above example) during disk release.
- \* The measured values are reference values from a test environment.
- \* The initial output of the iostat command is the cumulative values since the disk went online.

# <<Reference>> Disk Load Factor in zpool import Execution

## ■ You can import a disk, with a light load.

- ✓ Disk load factor during five seconds of zpool import execution

iostat -sxn 1 20 command output

```
extended device statistics
r/s   w/s   kr/s  kw/s  wait  actv  wsvc_t  asvc_t  %w  %b  device
0.2   0.3   2.7   3.2   0.0   0.0   0.0     28.6    0   0   c0d3
0.2   0.3   2.6   3.2   0.0   0.0   0.0     29.8    0   0   c0d4
0.5   0.3   8.1   2.8   0.0   0.0   0.0     14.3    0   0   c0d11

5.0   0.0 1276.0  0.0   0.0   0.0   0.0     8.5    0   4   c0d3
5.0   0.0 1276.0  0.0   0.0   0.1   0.0    10.7    0   5   c0d4
30.9  0.0 2695.6  0.0   0.0   0.1   0.0     3.5    0   8   c0d11

0.0   0.0   0.0   0.0   0.0   0.0   0.0     0.0    0   0   c0d3
0.0   0.0   0.0   0.0   0.0   0.0   0.0     0.0    0   0   c0d4
15.0  92.0  840.0  849.5  0.0   3.2   0.0    29.9    0  73  c0d11

0.0   0.0   0.0   0.0   0.0   0.0   0.0     0.0    0   0   c0d3
0.0   0.0   0.0   0.0   0.0   0.0   0.0     0.0    0   0   c0d4
0.0  33.0   0.0  286.6  0.0   1.0   0.0    30.1    0  26  c0d11

0.0   0.0   0.0   0.0   0.0   0.0   0.0     0.0    0   0   c0d3
0.0   0.0   0.0   0.0   0.0   0.0   0.0     0.0    0   0   c0d4
0.0   0.0   0.0   0.0   0.0   0.0   0.0     0.0    0   0   c0d11
```

## Column contents

Item	Description
r/s	Read (times/second)
w/s	Write (times/second)
kr/s	Read (kB/second)
kw/s	Write (kB/second)
wait	Average number of waiting transactions
actv	Average number of transactions in progress
wsvc_t	Average service time (millisecond)
asvc_t	Average service time (millisecond)
%w	Transaction wait time percentage
%b	Transaction busy (transaction in progress) percentage
device	Disk used

- 💡 - A disk load is generated for only several seconds (about three seconds in the above example) when the disk is being connected.
  - \* The measured values are reference values from a test environment.
  - \* The initial output of the iostat command is the cumulative values since the disk went online.

# Reference Information

# Input/Output Statistical Information for Storage Pools

- You can check the input/output statistical information for individual storage pools or individual devices (disks) composing a storage pool.
  - ZFS has a command that displays data input/output statistical information, which is similar to the standard iostat command in Solaris.

- ✓ Storage pool statistical information (zpool iostat command)

Syntax: `zpool iostat [-v] [pool_name] [seconds]`

```
# zpool iostat rz2pool 5
          capacity          operations
bandwidth
pool      used   avail   read   write   read   write
-----
rz2pool   230K   29.7G    0     13    682   12.3K
```

Specify the display interval. In this example, statistics are displayed every 5 seconds.

- ✓ Device statistical information

The -v option retrieves the layout and input/output statistics of the whole virtual device.

```
# zpool iostat -v rz2pool
          capacity          operations          bandwidth
pool      used   avai   read   write   read   write
-----
rz2pool   230K   29.7G    0     0     23     443
  raidz2   230K   29.7G    0     0     23     443
    c2d7    -     -      0     0     71    4.74K
    c2d8    -     -      0     0     71    4.74K
    c2d9    -     -      0     0    263    4.74K
-----
```

<b>pool</b>	Storage pool name, device name
<b>capacity</b>	
used	Size of data used
avail	Size of available data
<b>operation</b>	
read	Number of read operations
write	Number of write operations
<b>bandwidth</b>	
read	Quantity of read data
write	Quantity of write data

# Migrating a Storage Pool (export/import)

- You can release (export) a storage pool and incorporate (import) it in another system.

- ✓ Release a storage pool (zpool export command).

After being released from a system, the storage pool is not recognized by the system.

```
# zpool export rz2pool
```

- ✓ Incorporate a storage pool (zfs import command).

Execute the zfs import command to check the storage pools that can be incorporated.

```
# zpool import
  pool: rz2pool
    id: 192351259005321084
  state: ONLINE
action: The pool can be imported using its name or numeric identifier.
config:

    rz2pool                ONLINE
```

Incorporate a storage pool that can be incorporated.

```
# zpool import rz2pool
```

## ■ Monitor the operations on storage pools and file systems.

- ✓ Monitor ZFS operation (zpool monitor command).

Syntax: `zpool monitor -t provider [-T d|u] [[-p] -o field[,..]] [pool] ... [interval [count]]`

- `-t provider`: send, receive(recv), destroy, scrub, resilver  
=> Specifies a monitored operation.
- `-o field` : DONE, OTHER, PCTDONE, POOL, PROVIDER, SPEED, STRTTIME, TAG, TIMELEFT, Timestmp, TOTAL, all  
=> Specifies a display field. You can specify multiple fields delimited by a comma. If nothing is specified, POOL, PROVIDER, PCTDONE, TOTAL, SPEED, and TIMELEFT are displayed by default.

```
# zpool monitor -t scrub 5
POOL          PROVIDER  PCTDONE  TOTAL  SPEED
TIMELEFT
rpool         scrub    6.8      20.8G  506M   39s
rpool         scrub    26.9     20.8G  737M   21s
rpool         scrub    63.0     20.8G  1.03G  7s
:
```

- \* Output continues until Ctrl+C is input or processing ends.  
If interval is not specified, the command displays only one line, indicating the status at the command execution time.

- You can monitor the progress of the following operations.

File System Operation (zfs Command)	Storage Pool Operation (zpool Command)
<ul style="list-style-type: none"><li>- Send/Receive ZFS data (send/receive)</li><li>- Destroy snapshot (destroy)</li></ul>	<ul style="list-style-type: none"><li>- Verify file system (scrub)</li><li>- Resynchronize pool (resilver)</li></ul>

- ZFS-related operations are recorded, and you can display the history of ZFS-related commands by using the `zpool history` command.
  - ✓ Display the history of ZFS-related commands.

```
# zpool history
History for 'upool':
2016-11-14.23:07:52 zpool create upool mirror c0t500003942823F558d0 c0t50000394281A9F74d0
2016-11-14.23:08:02 zfs create upool/zfs1
2016-11-14.23:08:13 zfs snapshot
```

# Appendix

*Managing ZFS File Systems in Oracle® Solaris 11.3* (Oracle)  
[https://docs.oracle.com/cd/E53394\\_01/pdf/E54801.pdf](https://docs.oracle.com/cd/E53394_01/pdf/E54801.pdf)

# Revision History

Edition	Date	Description
First	December 2016	First edition created

## ■ Terms of Use

### ■ Copyrights, trademarks, and other intellectual property rights

- The contents (text, images, audio, etc.) are protected by copyrights, trademarks, and other intellectual property rights. The contents can be printed or downloaded for personal use only. The contents may not be otherwise used (reused on a personal webpage, uploaded to another server, etc.) without the prior permission of Fujitsu or the rights holder.

### ■ Limitation of warranties

- Fujitsu does not warrant the accuracy, merchantability, fitness of use, etc. of the contents. Fujitsu assumes no liability for any damages arising from use of the contents. The contents may be modified or deleted without any prior notice.

### ■ Export or provision to another party

- Before exporting this product or providing it to another party, check the applicable export control regulations such as the Foreign Exchange and Foreign Trade Act of Japan and the Export Administration Regulations of the United States, and follow the necessary procedures.

## ■ Trademarks

- UNIX is a registered trademark of The Open Group in the United States and other countries
- SPARC Enterprise, SPARC64, SPARC64 logo, and all other SPARC trademarks are trademarks or registered trademarks of SPARC International, Inc. in the United States and other countries and used under license.
- Oracle and Java are registered trademarks of Oracle Corporation and/or its affiliates in the United States and other countries.
- All other product names mentioned herein may be product names, trademarks, or registered trademarks of their respective owners.



**FUJITSU**

shaping tomorrow with you