

1: ZFS

ZFS 是 FreeBSD 的默认文件系统，也是目前最流行的文件系统之一。ZFS 是 FreeBSD 10.1 引入的，它结合了 RAID-Z 和 ZIL 等技术，提供了强大的数据保护和性能。ZFS 支持快照、克隆、压缩、加密等功能，并且具有自我修复的能力。

ZFS 的架构非常复杂，它由多个层组成。ZFS 的底层是 UFS，而 UFS 的底层是 BSD 的 UFS2。Linux 的 extfs 也是基于 UFS2 的。ZFS 的架构使得它能够支持各种硬件配置，并且能够在不同的操作系统上运行。ZFS 的快照功能非常强大，它允许用户创建数据的副本，并在需要时恢复到之前的状态。ZFS 的克隆功能也非常有用，它允许用户创建数据的副本，而不需要复制整个数据集。ZFS 的压缩功能可以减少数据的存储空间，并且不会影响性能。ZFS 的加密功能可以保护数据的安全，防止未经授权的访问。ZFS 的自我修复功能可以检测和修复数据损坏，确保数据的完整性和可用性。

ZFS 的快照功能非常强大，它允许用户创建数据的副本，并在需要时恢复到之前的状态。ZFS 的克隆功能也非常有用，它允许用户创建数据的副本，而不需要复制整个数据集。ZFS 的压缩功能可以减少数据的存储空间，并且不会影响性能。ZFS 的加密功能可以保护数据的安全，防止未经授权的访问。ZFS 的自我修复功能可以检测和修复数据损坏，确保数据的完整性和可用性。

ZFS 数据集 (ZFS Datasets)

ZFS 数据集是 ZFS 文件系统的基本单位。每个数据集都是一个 UFS 文件系统，并且可以包含子数据集。ZFS 数据集的创建和删除可以通过 `df(1)`, `newfs(8)`, `mount(8)`, `umount(8)`, `dump(8)`, `restore(8)` 等命令来完成。ZFS 数据集的命名规则是：父数据集名/子数据集名。例如，`zroot` 是根数据集，`zroot/ROOT` 是根数据集的子数据集，`zroot/ROOT/default` 是 `zroot/ROOT` 的子数据集。ZFS 数据集的命名规则是：父数据集名/子数据集名。例如，`zroot` 是根数据集，`zroot/ROOT` 是根数据集的子数据集，`zroot/ROOT/default` 是 `zroot/ROOT` 的子数据集。

`zfslist` 命令用于列出 ZFS 数据集的详细信息。

```
$ zfs list
NAME                USED AVAIL REFER MOUNTPOINT
zroot                429M 13.0G  96K none
zroot/ROOT           428M 13.0G  96K none
zroot/ROOT/default  428M 13.0G  428M /
zroot/tmp            104K 13.0G  104K
/tmp zroot/usr       428K 13.0G  96K /usr
...
```

mount(8) df(1) 命令用于挂载和检查数据集。UFS 和 extfs 是 ZFS 数据集的底层文件系统。ZFS 数据集的命名规则是：父数据集名/子数据集名。例如，`zroot` 是根数据集，`zroot/ROOT` 是根数据集的子数据集，`zroot/ROOT/default` 是 `zroot/ROOT` 的子数据集。

ZFS 的快照功能非常强大，它允许用户创建数据的副本，并在需要时恢复到之前的状态。ZFS 的克隆功能也非常有用，它允许用户创建数据的副本，而不需要复制整个数据集。ZFS 的压缩功能可以减少数据的存储空间，并且不会影响性能。ZFS 的加密功能可以保护数据的安全，防止未经授权的访问。ZFS 的自我修复功能可以检测和修复数据损坏，确保数据的完整性和可用性。

zroot 429MB 13GB

REFER ZFS root 429MB " " 96KB 13GB

zroot ZFS

zroot/ROOT 96KB

zroot/ROOT/default 428MB

ZFS FreeBSD 10.1-p1 zroot/ROOT/10.1-p1 /oldroot

zroot/tmp /tmp

ZFS (ZFS partitions and properties)

ZFS LBA () ()

ZFS

ZFS

zfs set quota=2G zroot/var/log

zfs get quota zroot/var/log

```
$ zfs set quota=2G zroot/var/log
```

zfs get quota zroot/var/log

```
$ zfs get quota zroot/var/log
NAME          PROPERTY VALUE SOURCE
zroot/var/log quota      2G local
```

zfs get all ZFS properties and their values.

ZFS Limits (ZFS Limits)

File systems have various limits. FAT file systems have a 32MB limit, while ZFS has a 2TB limit. UFS and ext2/3/4fs have different limits.

ZFS has a 2TB limit. ZFS also has a 128KB limit for the number of files. ZFS also has a 16TB limit for the number of files. ZFS also has a 256TB limit for the number of files.

ZFS also has a 264TB limit for the number of files. ZFS also has a 264TB limit for the number of files. ZFS also has a 264TB limit for the number of files.

zpool(8)

zpool(8) is a manual page for ZFS pools. It describes how to create and manage ZFS pools. It also describes the zpool(8) command.

```
$ zpool status
```

```
pool: zroot
```

```
state: ONLINE
```

```
scan: none requested
```

```
config:
```

NAME	STATE	READ	WRITE	CKSUM
zroot	ONLINE	0	0	0
gpt/zfs0	ONLINE	0	0	0

```
errors: No known data errors
```

zroot is a ZFS pool on a RAIDZ2 configuration. It consists of 2 disks in a RAIDZ2 configuration. The pool is currently online and has no known data errors.

ZFS scan is a command used to scan the pool for data errors. It can be used to scan the entire pool or a specific dataset. The scan process can be interrupted at any time.

zpool status is a command used to check the status of the pool. It provides information about the pool's state, configuration, and any errors.

Virtual Devices (VDEVs)

VDEVs are virtual devices that are used to represent physical disks in a ZFS pool. They can be configured in various ways, such as RAIDZ, RAIDZ2, RAIDZ3, and RAIDZ4. VDEVs can also be configured in a mirrored configuration.

zpool status is a command used to check the status of the pool. It provides information about the pool's state, configuration, and any errors. The output of the command shows the pool's state as ONLINE and the configuration as RAIDZ2.

ZFS is a file system that is designed for high performance and reliability. It is based on the VFS (Virtual File System) layer and provides a wide range of features, including data integrity, snapshots, and compression.

ZFS is a file system that is designed for high performance and reliability. It is based on the VFS (Virtual File System) layer and provides a wide range of features, including data integrity, snapshots, and compression.

zpool status is a command used to check the status of the pool. It provides information about the pool's state, configuration, and any errors. The output of the command shows the pool's state as ONLINE and the configuration as RAIDZ2.

Blocks and Inodes

File systems are categorized into two main types: **block-based** and **inode-based**. Block-based file systems store data in blocks, while inode-based file systems store data in inodes. Common examples include BSD UFS, Linux extfs, and Microsoft FAT.

ZFS is a file system that uses a combination of blocks and inodes. It is designed for high performance and reliability. ZFS uses a **copy-on-write** mechanism, which means that data is only written to a new block when the original block is modified. This allows for efficient storage and recovery of data.

UFS (Unix File System) is a block-based file system. It uses a **block pointer** to store data. ZFS, on the other hand, uses a **copy-on-write** mechanism. This means that ZFS does not use a block pointer to store data. Instead, it uses a **copy-on-write** mechanism to store data. This allows ZFS to store data in a more efficient manner than UFS.

ZFS uses a **copy-on-write** mechanism, which means that data is only written to a new block when the original block is modified. This allows for efficient storage and recovery of data. ZFS also uses a **checksum** to verify the integrity of data. This means that ZFS can detect and correct errors in data.

Revision #2

Created 21 January 2024 03:08:23 by [User] (MeatDumpling)

Updated 21 January 2024 03:37:05 by [User] (MeatDumpling)