

# Chapter 2: Virtual Devices

This chapter discusses virtual devices, which are devices that exist only in software. These devices are used to provide a virtualized environment for applications and services. The most common virtual devices are virtual disks and virtual network interfaces. These devices are used to provide a virtualized environment for applications and services. The most common virtual devices are virtual disks and virtual network interfaces.

ZFS is a file system that is designed to be used on virtual devices. ZFS provides a number of features that make it well-suited for use on virtual devices, including support for large file sizes and high performance.

## Virtual Disks and Other Storage Media (Disks and Other Storage Media)

ZFS can be used on a variety of storage media, including virtual disks and other storage media. ZFS provides a number of features that make it well-suited for use on virtual disks and other storage media, including support for large file sizes and high performance.

## Raw Disk Storage (Raw Disk Storage)

Raw disk storage is a type of storage that is not formatted with a file system. Raw disk storage is used to store data in a binary format. ZFS can be used on raw disk storage, but it is not recommended.

FreeBSD provides a number of tools for managing raw disk storage. These tools include `dd`, `disk`, and `zfs`. These tools are used to create and manage raw disk storage.

6TB is a common size for raw disk storage. ZFS can be used on raw disk storage, but it is not recommended. ZFS provides a number of features that make it well-suited for use on virtual disks and other storage media, including support for large file sizes and high performance.

Virtual Disks



GEOM 是一個用於管理磁碟的軟體層。HAST 是一個用於高可用性配置的軟體層。HAST 和 GEOM 可以一起使用，以提供高可用性和數據冗余。HAST 使用 ZFS 來實現數據冗余，而 GEOM 則用於管理磁碟。HAST 和 GEOM 的組合可以為您的系統提供高可用性和數據冗余。HAST 和 GEOM 的組合可以為您的系統提供高可用性和數據冗余。

GEOM 提供了一個用於管理磁碟的軟體層。disk ident, gptid, GPT 是 GEOM-specific label 的示例。GEOM 提供了一個用於管理磁碟的軟體層。disk ident, gptid, GPT 是 GEOM-specific label 的示例。GEOM 提供了一個用於管理磁碟的軟體層。disk ident, gptid, GPT 是 GEOM-specific label 的示例。

GEOM 提供了一個用於管理磁碟的軟體層。HBA(RAID 控制器) 是一個用於管理磁碟的硬體層。HBA(RAID 控制器) 是一個用於管理磁碟的硬體層。HBA(RAID 控制器) 是一個用於管理磁碟的硬體層。HBA(RAID 控制器) 是一個用於管理磁碟的硬體層。HBA(RAID 控制器) 是一個用於管理磁碟的硬體層。HBA(RAID 控制器) 是一個用於管理磁碟的硬體層。

GEOM 提供了一個用於管理磁碟的軟體層。I/O 是一個用於管理磁碟的硬體層。I/O 是一個用於管理磁碟的硬體層。I/O 是一個用於管理磁碟的硬體層。I/O 是一個用於管理磁碟的硬體層。I/O 是一個用於管理磁碟的硬體層。I/O 是一個用於管理磁碟的硬體層。

## File-Backed Storage (File-Backed Storage)

ZFS 是一個用於管理磁碟的軟體層。ZFS 是一個用於管理磁碟的軟體層。ZFS 是一個用於管理磁碟的軟體層。ZFS 是一個用於管理磁碟的軟體層。ZFS 是一個用於管理磁碟的軟體層。ZFS 是一個用於管理磁碟的軟體層。

## Providers vs. Disks (Providers vs. Disks)

FreeBSD 提供了一個用於管理磁碟的軟體層。GEOM 提供了一個用於管理磁碟的軟體層。GEOM 提供了一個用於管理磁碟的軟體層。GEOM 提供了一個用於管理磁碟的軟體層。GEOM 提供了一個用於管理磁碟的軟體層。GEOM 提供了一個用於管理磁碟的軟體層。

FreeBSD 提供了一個用於管理磁碟的軟體層。GEOM 提供了一個用於管理磁碟的軟體層。GEOM 提供了一個用於管理磁碟的軟體層。GEOM 提供了一個用於管理磁碟的軟體層。GEOM 提供了一個用於管理磁碟的軟體層。GEOM 提供了一個用於管理磁碟的軟體層。

ZFS 是一個用於管理磁碟的軟體層。ZFS 是一個用於管理磁碟的軟體層。ZFS 是一個用於管理磁碟的軟體層。ZFS 是一個用於管理磁碟的軟體層。ZFS 是一個用於管理磁碟的軟體層。ZFS 是一個用於管理磁碟的軟體層。



RAID-Z1

RAID-Z1 is a RAID configuration that uses three disks to store data and parity. It is similar to RAID-5, but it uses a different parity scheme. RAID-Z1 is designed for high performance and reliability. It is suitable for use in environments where data integrity and availability are critical.

## RAID-Z1(3 disks)

ZFS RAID-Z1 uses three disks to store data and parity. It is similar to RAID-5, but it uses a different parity scheme. RAID-Z1 is designed for high performance and reliability. It is suitable for use in environments where data integrity and availability are critical.

RAID-Z1 is a RAID configuration that uses three disks to store data and parity. It is similar to RAID-5, but it uses a different parity scheme. RAID-Z1 is designed for high performance and reliability. It is suitable for use in environments where data integrity and availability are critical.

RAID-Z1 VDEV uses three disks to store data and parity. It is similar to RAID-5, but it uses a different parity scheme. RAID-Z1 is designed for high performance and reliability. It is suitable for use in environments where data integrity and availability are critical.

## RAID-Z2(4 disks)

RAID-Z2 is a RAID configuration that uses four disks to store data and parity. It is similar to RAID-6, but it uses a different parity scheme. RAID-Z2 is designed for high performance and reliability. It is suitable for use in environments where data integrity and availability are critical.

## RAID-Z3(5 disks)

RAID-Z3 is a RAID configuration that uses five disks to store data and parity. It is similar to RAID-7, but it uses a different parity scheme. RAID-Z3 is designed for high performance and reliability. It is suitable for use in environments where data integrity and availability are critical.

## RAID-Z Disk Configurations (RAID-Z Disk Configurations)

RAID-Z disk configurations are used to optimize performance and reliability. They are designed for use in environments where data integrity and availability are critical. RAID-Z disk configurations are used to optimize performance and reliability.

20 disks are used in RAID-Z2 configurations. RAID-Z2 uses four disks to store data and parity. It is similar to RAID-6, but it uses a different parity scheme. RAID-Z2 is designed for high performance and reliability. It is suitable for use in environments where data integrity and availability are critical.



RAID 2 的 寫入 延遲 "寫入 洞 (write hole)" 問題 較 嚴重 . RAID 5 的 6 個 磁碟 中 有 5 個 磁碟 在 寫入 時 必須 等待 其他 磁碟 寫入 完畢 . 而 RAID 6 的 6 個 磁碟 中 有 5 個 磁碟 在 寫入 時 必須 等待 其他 磁碟 寫入 完畢 . RAID 6 的 6 個 磁碟 中 有 5 個 磁碟 在 寫入 時 必須 等待 其他 磁碟 寫入 完畢 . RAID 6 的 6 個 磁碟 中 有 5 個 磁碟 在 寫入 時 必須 等待 其他 磁碟 寫入 完畢 .

RAID 6 的 寫入 延遲 問題 較 嚴重 . RAID 6 的 6 個 磁碟 中 有 5 個 磁碟 在 寫入 時 必須 等待 其他 磁碟 寫入 完畢 . RAID 6 的 6 個 磁碟 中 有 5 個 磁碟 在 寫入 時 必須 等待 其他 磁碟 寫入 完畢 .

ZFS 的 寫入 延遲 問題 較 嚴重 . ZFS 的 寫入 延遲 問題 較 嚴重 . ZFS 的 寫入 延遲 問題 較 嚴重 . ZFS 的 寫入 延遲 問題 較 嚴重 . ZFS 的 寫入 延遲 問題 較 嚴重 . ZFS 的 寫入 延遲 問題 較 嚴重 . ZFS 的 寫入 延遲 問題 較 嚴重 . ZFS 的 寫入 延遲 問題 較 嚴重 .

## VDEV (Special VDEVs)

VDEV 的 寫入 延遲 問題 較 嚴重 . VDEV 的 寫入 延遲 問題 較 嚴重 .

## Separate Intent Log; SLOG, ZIL

ZFS 的 寫入 延遲 問題 較 嚴重 . ZFS 的 寫入 延遲 問題 較 嚴重 .

ZFS 的 寫入 延遲 問題 較 嚴重 . ZFS 的 寫入 延遲 問題 較 嚴重 . ZFS 的 寫入 延遲 問題 較 嚴重 . ZFS 的 寫入 延遲 問題 較 嚴重 . ZFS 的 寫入 延遲 問題 較 嚴重 .

ZFS 的 寫入 延遲 問題 較 嚴重 . ZFS 的 寫入 延遲 問題 較 嚴重 . ZFS 的 寫入 延遲 問題 較 嚴重 . ZFS 的 寫入 延遲 問題 較 嚴重 . ZFS 的 寫入 延遲 問題 較 嚴重 .

ZFS 的 寫入 延遲 問題 較 嚴重 . ZFS 的 寫入 延遲 問題 較 嚴重 . ZFS 的 寫入 延遲 問題 較 嚴重 . ZFS 的 寫入 延遲 問題 較 嚴重 . ZFS 的 寫入 延遲 問題 較 嚴重 .

ZIL 的 寫入 延遲 問題 較 嚴重 . ZIL 的 寫入 延遲 問題 較 嚴重 . ZIL 的 寫入 延遲 問題 較 嚴重 . ZIL 的 寫入 延遲 問題 較 嚴重 . ZIL 的 寫入 延遲 問題 較 嚴重 .

## Cache: L2ARC

Cache is a small amount of memory that stores frequently accessed data. In the context of ZFS, the L2ARC (Level 2 Adaptive Replacement Cache) is a cache of data that is stored in RAM. It is used to store data that is frequently accessed, so that it can be retrieved quickly. The L2ARC is managed by the ZFS kernel, and its size is determined by the amount of RAM available. The L2ARC is a cache of data that is stored in RAM, and it is used to store data that is frequently accessed, so that it can be retrieved quickly. The L2ARC is managed by the ZFS kernel, and its size is determined by the amount of RAM available.

The L2ARC is a cache of data that is stored in RAM, and it is used to store data that is frequently accessed, so that it can be retrieved quickly. The L2ARC is managed by the ZFS kernel, and its size is determined by the amount of RAM available. The L2ARC is a cache of data that is stored in RAM, and it is used to store data that is frequently accessed, so that it can be retrieved quickly. The L2ARC is managed by the ZFS kernel, and its size is determined by the amount of RAM available.

The L2ARC is a cache of data that is stored in RAM, and it is used to store data that is frequently accessed, so that it can be retrieved quickly. The L2ARC is managed by the ZFS kernel, and its size is determined by the amount of RAM available. The L2ARC is a cache of data that is stored in RAM, and it is used to store data that is frequently accessed, so that it can be retrieved quickly. The L2ARC is managed by the ZFS kernel, and its size is determined by the amount of RAM available.

The L2ARC is a cache of data that is stored in RAM, and it is used to store data that is frequently accessed, so that it can be retrieved quickly. The L2ARC is managed by the ZFS kernel, and its size is determined by the amount of RAM available. The L2ARC is a cache of data that is stored in RAM, and it is used to store data that is frequently accessed, so that it can be retrieved quickly. The L2ARC is managed by the ZFS kernel, and its size is determined by the amount of RAM available.

## VDEVs (How VDEVs Affect Performance)

VDEVs (Virtual Devices) are the building blocks of a ZFS storage pool. They are used to store data, and their performance characteristics can significantly affect the overall performance of the pool. The number of VDEVs, their size, and the type of storage they use (e.g., SSDs, HDDs) all play a role in determining the pool's performance.

The performance of a ZFS storage pool is determined by the performance of its VDEVs. The number of VDEVs, their size, and the type of storage they use (e.g., SSDs, HDDs) all play a role in determining the pool's performance. The performance of a ZFS storage pool is determined by the performance of its VDEVs. The number of VDEVs, their size, and the type of storage they use (e.g., SSDs, HDDs) all play a role in determining the pool's performance.

The performance of a ZFS storage pool is determined by the performance of its VDEVs. The number of VDEVs, their size, and the type of storage they use (e.g., SSDs, HDDs) all play a role in determining the pool's performance. The performance of a ZFS storage pool is determined by the performance of its VDEVs. The number of VDEVs, their size, and the type of storage they use (e.g., SSDs, HDDs) all play a role in determining the pool's performance.

The performance of a ZFS storage pool is determined by the performance of its VDEVs. The number of VDEVs, their size, and the type of storage they use (e.g., SSDs, HDDs) all play a role in determining the pool's performance.

The performance of a ZFS storage pool is determined by the performance of its VDEVs. The number of VDEVs, their size, and the type of storage they use (e.g., SSDs, HDDs) all play a role in determining the pool's performance.



### 3 (Three Disks)

3 disks can be configured as RAID-Z1 or RAID-Z2. RAID-Z1 provides a single parity disk, while RAID-Z2 provides two parity disks. RAID-Z1 uses 33% of the disks for parity, while RAID-Z2 uses 66%.

RAID-Z1 is a good choice for applications that require high read performance and can tolerate a single disk failure. RAID-Z2 is a good choice for applications that require high write performance and can tolerate two disk failures. RAID-Z1 provides a single parity disk, while RAID-Z2 provides two parity disks. RAID-Z1 uses 33% of the disks for parity, while RAID-Z2 uses 66%.

RAID-Z1 is a good choice for applications that require high read performance and can tolerate a single disk failure. RAID-Z2 is a good choice for applications that require high write performance and can tolerate two disk failures. RAID-Z1 provides a single parity disk, while RAID-Z2 provides two parity disks. RAID-Z1 uses 33% of the disks for parity, while RAID-Z2 uses 66%.

Table 3: Three-Disk Virtual Device Configurations

Disks	Config	Read IOPS	Write IOPS	Read MB/s	Write MB/s	Usable Space	Fault Tolerance
3	1 x 3 disk Mirror	750	250	300	100	1 TB (33%)	2
3	1 x 3 disk RAID-Z1	250	250	200	200	2 TB (66%)	1

RAID-Z1 provides a single parity disk, while RAID-Z2 provides two parity disks. RAID-Z1 uses 33% of the disks for parity, while RAID-Z2 uses 66%. RAID-Z1 is a good choice for applications that require high read performance and can tolerate a single disk failure. RAID-Z2 is a good choice for applications that require high write performance and can tolerate two disk failures. RAID-Z1 provides a single parity disk, while RAID-Z2 provides two parity disks. RAID-Z1 uses 33% of the disks for parity, while RAID-Z2 uses 66%.

### 4 or 5 (Four or Five Disks)

4 or 5 disks can be configured as RAID-Z1 or RAID-Z2. RAID-Z1 provides a single parity disk, while RAID-Z2 provides two parity disks. RAID-Z1 uses 25% of the disks for parity, while RAID-Z2 uses 40%.

RAID-Z1 is a good choice for applications that require high read performance and can tolerate a single disk failure. RAID-Z2 is a good choice for applications that require high write performance and can tolerate two disk failures. RAID-Z1 provides a single parity disk, while RAID-Z2 provides two parity disks. RAID-Z1 uses 25% of the disks for parity, while RAID-Z2 uses 40%.

RAID-Z1 is a good choice for applications that require high read performance and can tolerate a single disk failure. RAID-Z2 is a good choice for applications that require high write performance and can tolerate two disk failures. RAID-Z1 provides a single parity disk, while RAID-Z2 provides two parity disks. RAID-Z1 uses 25% of the disks for parity, while RAID-Z2 uses 40%.

RAID-Z1 is a good choice for applications that require high read performance and can tolerate a single disk failure. RAID-Z2 is a good choice for applications that require high write performance and can tolerate two disk failures. RAID-Z1 provides a single parity disk, while RAID-Z2 provides two parity disks. RAID-Z1 uses 25% of the disks for parity, while RAID-Z2 uses 40%.

RAID-Z1 (MB/s) .

Table 4: Four- or Five-Disk Virtual Device Configurations

Disks	Config	Read IOPS	Write IOPS	Read MB/s	Write MB/s	Usable Space	Fault Tolerance
4	2 x 2 disk Mirror	1000	500	400	200	2 TB (50%)	2 (1/VDEV)
4	1 x 4 disk RAIDZ-Z1	250	250	300	300	3 TB (75%)	1
4	1 x 4 disk RAIDZ-Z2	250	250	200	200	2 TB (50%)	2
5	1 x 5 disk RAIDZ-Z1	250	250	400	400	4 TB (80%)	1
5	1 x 5 disk RAIDZ-Z2	250	250	300	300	3 TB (60%)	2
5	1 x 5 disk RAIDZ-Z3	250	250	200	200	2 TB (40%)	3

RAID-Z1 (MB/s) RAID-Z2 , RAID-Z3 .

VDEV n - 1 VDEV .

### 6~12 (Six to Twelve Disks)

RAID-Z VDEV .

6 3 2 VDEV . 3- VDEV 3 2 RAID-Z VDEV .

6 RAID-Z VDEV . 12 VDEV .

Table 5: Six- to Twelve-Disk Virtual Device Configurations

Disks	Config	Read IOPS	Write IOPS	Read MB/s	Write MB/s	Usable Space	Fault Tolerance
6	3 x 2 disk Mirror	1500	750	600	300	3 TB (50%)	3 (1/VDEV)



