

RAID and Storage Solutions

For the free video please see
<http://itfreetraining.com/server#storage-solutions>

This video will look at a number of different storage solutions. These include software and hardware based systems. A storage solution is a system that allows more drives to be combined together for performance or redundancy reasons.

What's in this video

- Software vs Hardware
- JBOD
- Spanning
- RAID
- Storage solutions
 - Windows Storage Spaces
 - Logical Volume Manager

What's in this video

The following storage solutions will be looked at in this video.

Software vs Hardware: The advantages to using hardware over a software solution.

JBOD: Just a bunch of disks. Allows different sized drives to be combined together to form the one drive.

Spanning: Allows multiple drives to be combined of different sizes.

RAID: Redundant Array of Inexpensive Disks is a system that allows multiple drives to be combined to form the one drive.

Windows Storage Spaces: This is a new system implemented in Windows Server 2012 that allows multiple drives to be combined together.

Logical Volume Manager: Is an alternative storage system used by operating systems like Linux.

Software vs Hardware


- Hardware
 - Costs more
 - Faster and less load on the system
 - Supports booting the OS
- Software
 - May not be able to boot the OS
 - May offer additional features

Software vs Hardware

Hardware based systems typically cost more than a software solution as software solutions usually come with the operating system free of charge. A lot of motherboards now come with free hardware based solutions. You will find that if you purchase a server this may come with some hardware based solutions. Some servers may require additional hardware in the server or a higher model may need to be purchased to gain access to some hardware based solutions. The biggest advantage of a hardware solutions is that the operating system sees the drive as a single physical drive. This means the operating system can be booted from this drive. Some software based solutions do not support booting of the operating system. Software solutions may also support some additional features not supported by hardware. For example a software based solutions may allow for multiple files containing the same data to use the same physical space on the drive. Enterprise hardware solutions will often offer additional features as well but do cost more. For example, enterprise hardware solutions will have a web interface allowing access to additional features.

JBOD

- Just a bunch of disks
- Combines different size drives together
- No redundancy (If one drive fails all data is lost)
- No performance improvements

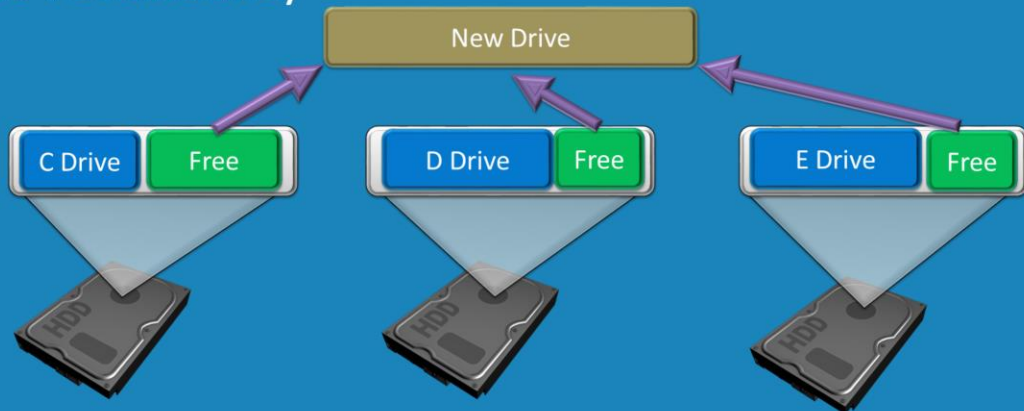


JBOD

Just a Bunch of Disks allows multiple drives to be combined together. This includes different sized drives and different types. For example you could combine solid state drives and mechanical drives together. JBOD does not offer any performance increase and if one of the drives was to fail you would lose all the data on all the drives.

Spanning

- Combines free space on multiple drives together
- No performance increase
- No redundancy



Spanning

Spanning is similar to JBOD however it combines free space on multiple drives together into the one drive. The advantage of spanning is that it allows space that may have otherwise been lost to be used. Spanning does not provide any speed advantages and also does not offer any redundancy. If a drive that is used in spanning was lost, then all the data in the spanned set would be lost.

RAID


- Redundant array of independent disks
- Combines multiple drives together
 - Drives need to be the same size
 - Should be similar speeds
- May not be able to add additional drives later

RAID

Redundant Array of Inexpensive Disks is a system which allows multiple drives to be combined together to form the one drive. The drives need to be the same size in order to be used. If one drive is larger than the others, typically it still can be used, however the extra space will be left unused. Depending on which RAID solution is used will determine if there are any redundancy or speed advantages. The more expensive RAID solutions may allow drives to be added to the RAID, increasing the amount of space in the RAID. A lot of RAID solutions do not offer this feature and thus if you want to change the size of the RAID you need to destroy the RAID and recreate it.

RAID 0 (Striping)

- Requires 2 or more disks
- Improves read and write performance
- No redundancy




RAID 0 (Striping)

RAID 0, otherwise known as striping, requires 2 or more disks in order to operate.

Data is divided between the drives in blocks. This offers fast read and write performance but no redundancy. If you have a file that is smaller than a block it will be stored on the one drive. Files that are larger than a block will be divided up and stored on different drives. If one drive is lost in the set, all data in the set is lost. If you have difficulty remembering which RAID level is striping, remember that RAID 0 offers zero redundancy.

RAID 1 (Mirroring)

- Requires 2 disks
- Halves useable space but adds redundancy
- No performance increase




RAID 1 (Mirroring)

RAID 1 requires two drives. It works by storing two copies of the data. One copy on each drive. This means that if one drive was to fail no data would be lost. This gives redundancy; however, it does not offer any performance increases. It also effectively doubles the cost of storage. RAID 1 is often used for the operating system drive.

RAID 5 (Striped With Parity)

- Requires 3 or more drives one lost to parity
- Can still operate with one drive missing
- Fast reads, slow writes




RAID 5 (Striped with Parity)

RAID 5 divides the data up onto multiple drives using one drive as parity. The parity information allows for a single drive in the system to fail while allowing the RAID to still operate. Let's consider the following $1 + 2 = 3$. If you did not know the value 2, that is $1 + x = 3$, you could work out that the missing value is 2. This is the same principle of parity, adding additional information so that if a drive is lost, the additional information can be used to work out what this information was. RAID 5 requires 3 or more drives to operate. It offers excellent read performance as multiple drives can be read at once. Write performance is slow however when data is written to the drive as the parity information must also be updated. Depending on how much data is being written to the drive, updating the parity may involve a read before the write can be performed.

RAID 1+0


- Excellent read and write performance
- Can still operate after multiple failures
- Increase cost to implement



RAID 1+0


This is sometimes referred to as RAID 10. RAID levels like this combine two RAID levels together. In this case drives are stored using RAID 1. All the RAID 1 drives are then striped, thus RAID 1+0 is also known as a stripe of mirrors. This solutions offers excellent read and write performance as there is no parity drive that needs to be updated. The disadvantage is the cost of storage doubles. The advantage of RAID1+0 is that is can still operate after multiple failures.

RAID 1+0



RAID 1+0

RAID 1+0 uses multiple drives to store the same data on and thus can still operate after multiple failures. Since the system does use RAID 1 sets, it is possible for the system to fail after losing 2 drives. This will occur if two drives were to fail that were holding the same data. If multiple drives were to fail that are not holding the same data the RAID 1+0 will still work. Business solutions will often have a large number of drives in them so the chance of two drives holding the same data to fail is very low. Thus RAID 1+0 offers a performance increase over RAID 5 but does not have the disadvantages of RAID 5.

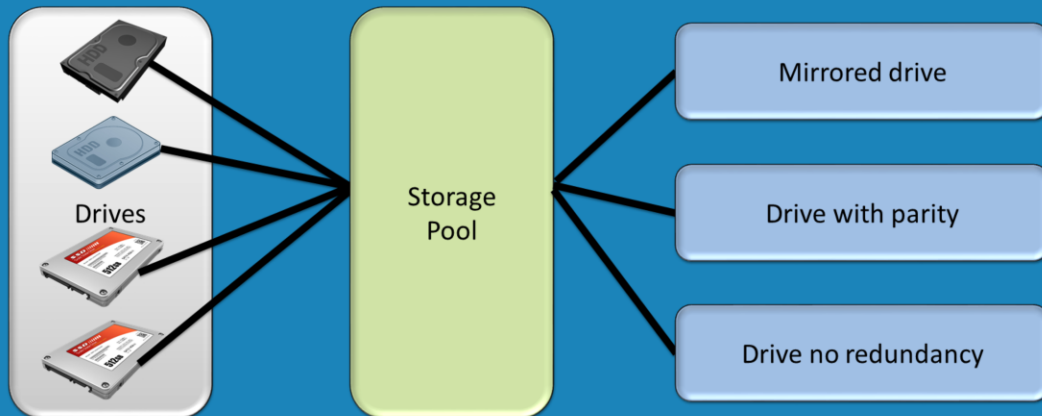


RAID 0+1

There is also RAID 0+1 which is essentially a mirror of a striped set. Like RAID 1+0 it offers excellent read and write performance and can still operate after multiple drive failures. However, if 2 drives were to fail with the same information on them then all data would be lost.

Storage Systems

- Can add/remove drives as required
- Manages available drives into storage pool



Storage Systems

A storage system is a system that manages the physical drives in a system for the administrator. The administrator will generally create drives using the free space. For example they could create a mirrored drive or a drive with parity. The storage system is responsible for keeping the required level of redundancy. If a physical drive was to fail and later was replaced, the storage system would take care of making sure the redundancy level is still kept. For example, if a large drive was to fail and was replaced by two smaller drives, the storage system would reorganize the data so the required redundancy level is still kept. Storage systems can be hardware or software based. A storage system in hardware is usually found in high level solutions like SAN solutions. Software based storage solutions are often included as part of the operating system.

Windows Storage Spaces

- Requires Windows Server 2012/Windows 8
- Replaces Logical Disk Manager
- Supports redundancy
 - Mirroring, parity
- Allows thin provisioning
- Can use ReFS
- Does not allow booting

Windows Storage Spaces

Windows Storage Spaces was added in Windows Server 2012 and Windows 8. It is designed as a replacement for the existing system, Logical Disk Manager, however it lacks some features of Logical Disk Manager, most noticeable is that the operating system cannot be booted from a Windows Storage Spaces drive. Windows Storage Spaces support the creation of multiple drives from a storage pool. The drives can have mirroring, parity or no redundancy. If you are using Windows Server 2012 R2, there is the option for 3 way mirror and extra redundancy. Windows Storage Spaces adds the option for thin provisioning. This allows a virtual drive to be created that only uses space as required. This is the same system that is used for virtual machines. Windows Storage Spaces was also designed to work with ReFS. ReFS is a new file system that was introduced in Windows Server 2012 and Windows 8.

Logical Volume Manager (LVM)

- Linux based solution
- Add remove drives without downtime
- Create volumes with redundancy
 - Or without
- Allows volumes to be resized
- Operating system can be booted from an LVM
- Support on most Linux Distributions

Logical Volume Manager (LVM)

This is a system used in a lot of Linux based systems. It can create drives with or without redundancy and allows for features like the volume to be resized as required. The advantage of LVM is that the operating system can be booted from a LVM drive.

See <http://YouTube.com/ITFreeTraining> or <http://itfreetraining.com> for our always free training videos. This is only one video from the many free courses available on YouTube.

References

“Installing and Configuring Windows Server 2012 Exam Ref 70-410” pg 49-55

“RAID” <http://en.wikipedia.org/wiki/RAID>

“Nested RAID levels” http://en.wikipedia.org/wiki/Nested_RAID_levels

“Step-by-step for Storage Spaces Tiering in Windows Server 2012 R2”

<http://blogs.technet.com/b/josebda/archive/2013/08/28/step-by-step-for-storage-spaces-tiering-in-windows-server-2012-r2.aspx>

“Logical Volume Manager (Linux)”

[http://en.wikipedia.org/wiki/Logical_Volume_Manager_\(Linux\)](http://en.wikipedia.org/wiki/Logical_Volume_Manager_(Linux))