

Raid Explained

What is Raid and the pros and cons of each

Author:

Mark White, The James Comprehensive Cancer Center at The Ohio State University's Wexner Medical Center

Discussion Points:

- Introduction – Who am I?
- Raid Explained (Common types of RAID)
- Radonc at Ohio State University
- Helpful resources (Future of Storage)
- Questions – Hearing from you

Raid Definition

- **Redundant Array of Independent Disks**) A disk subsystem that increases performance or provides fault tolerance or both. RAID uses two or more physical disk drives and a RAID controller, which is plugged into motherboards that do not have RAID circuits. Today, most motherboards have built-in RAID but not necessarily every RAID configuration (see below). In the past, RAID was also accomplished by software only but was much slower. In the late 1980s, the "I" in RAID stood for "inexpensive" but was later changed to "independent."
- RAID is a proven data storage virtualization technology that combines multiple disk drives into a single logical unit. The result is complete data redundancy and/or improved performance.
- In large storage area networks (SANs), floor-standing RAID units are common with terabytes of storage and huge amounts of cache memory. RAID is also used in desktop computers by gamers for speed and by business users for reliability. Following are the various RAID configurations.

Who should use RAID?

- Anyone who needs to keep large quantities of data on hand (System Administrators, Engineers, System Architects)
- Organizations with a need to protect data.

Benefits of RAID?

- Increased storage capacity using a single virtual disk
- Enhanced speed
- Lessening the impact of a disk failure – Data Loss
- Fault tolerance

Raid Types?

RAID



Storage is another very important part of fault tolerance.

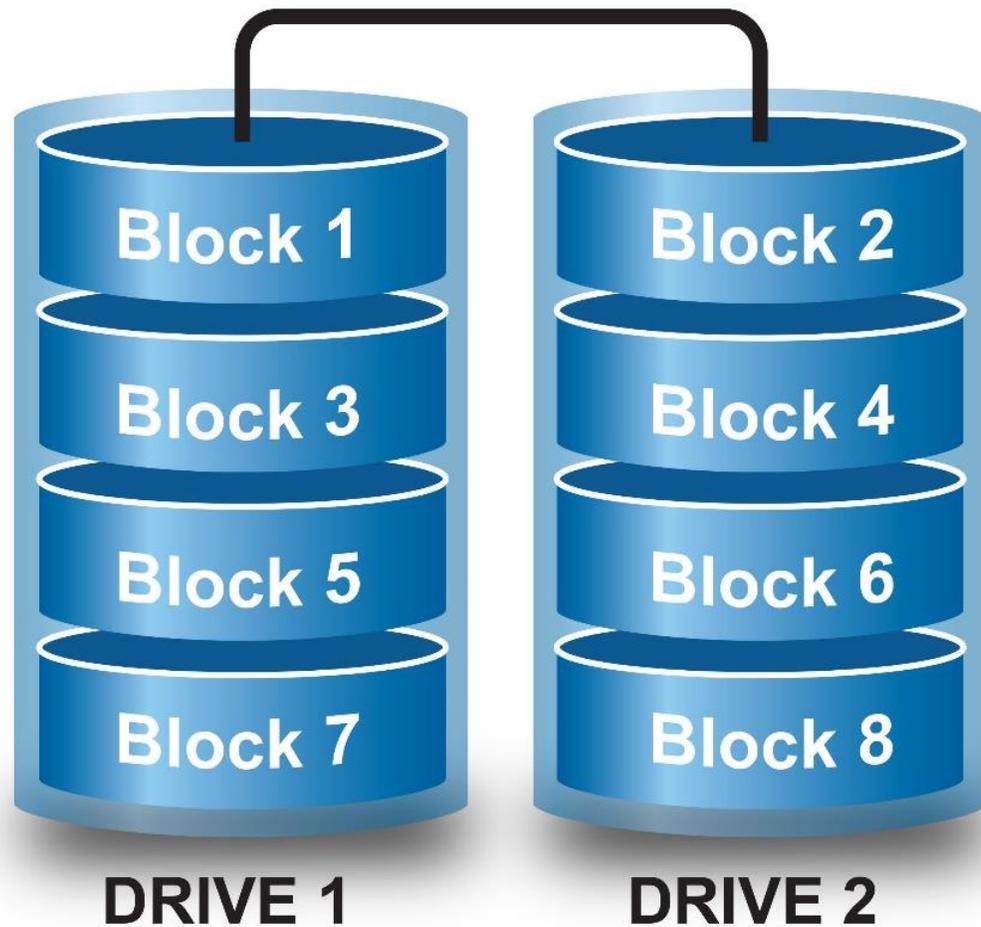
Data loss prevention.

RAID (Redundant Array of Independent Disks)

- RAID 0
- RAID 1
- RAID 5
- RAID 10



RAID 0

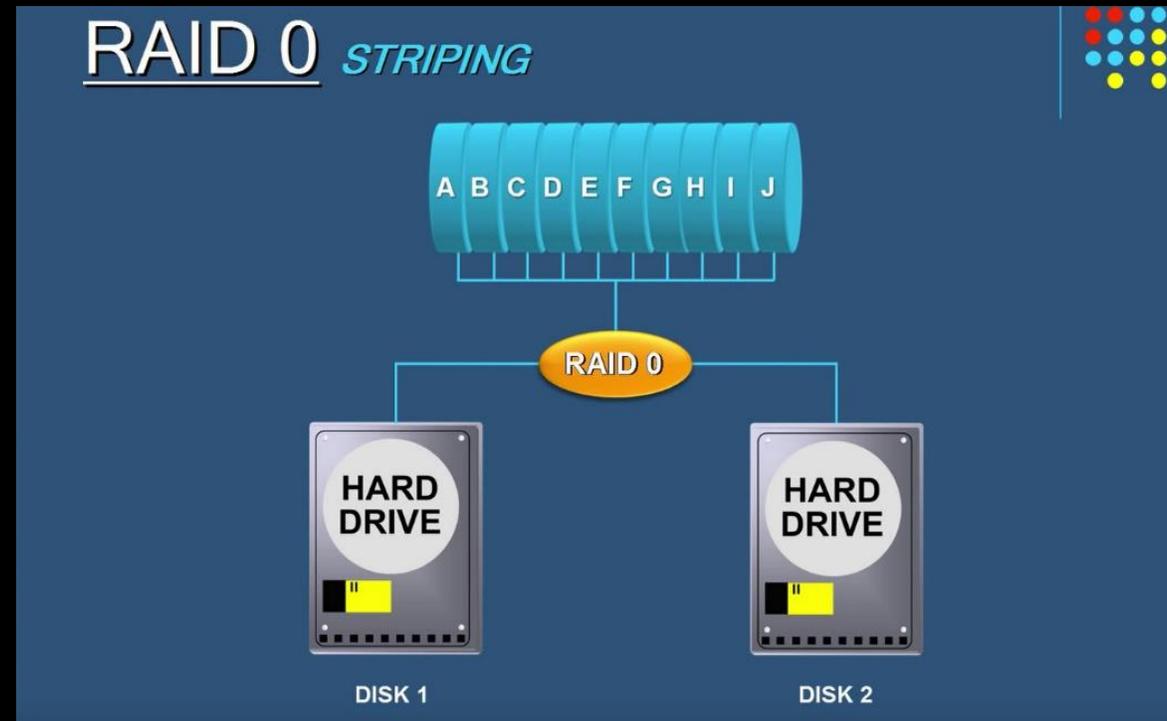


What is Raid 0?

RAID 0 is called disk striping. All the data is spread out in chunks across all the disks in the RAID set. RAID 0 has great performance because you spread the load of storing data onto more physical drives. There is no parity generated for RAID 0. Therefore there is no overhead to write data to RAID 0 disks. RAID 0 is only good for better performance, and not for high availability, since parity is not generated for RAID 0 disks. RAID 0 requires at least two physical disks.

Benefits of Raid 0

- Faster i/o
- Only requires 2 disk
- Least expensive
- Simple configuration

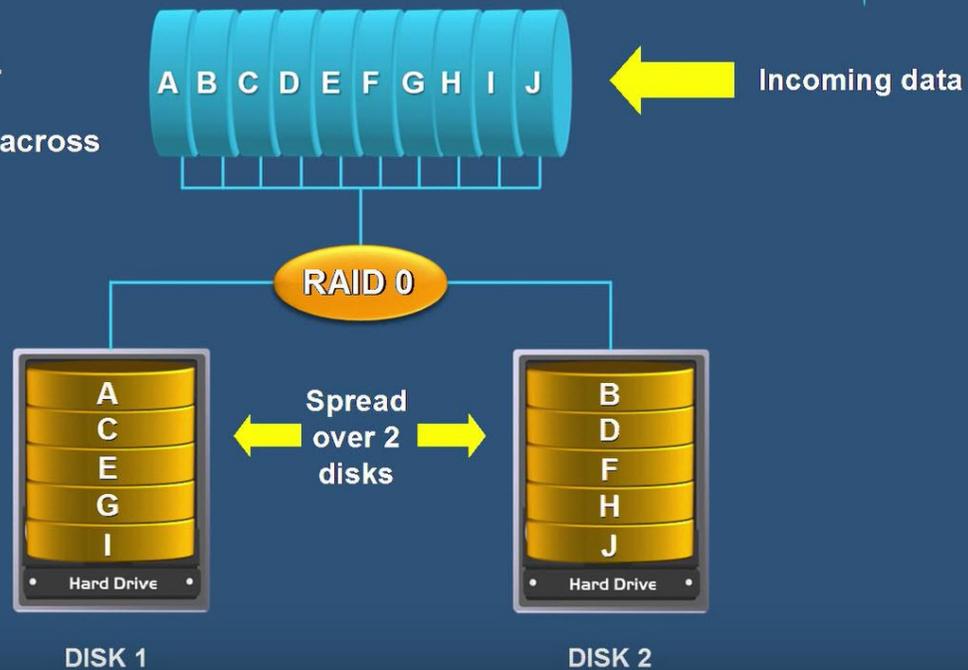


Cons of RAID 0

RAID 0 STRIPING

Not fault tolerant.

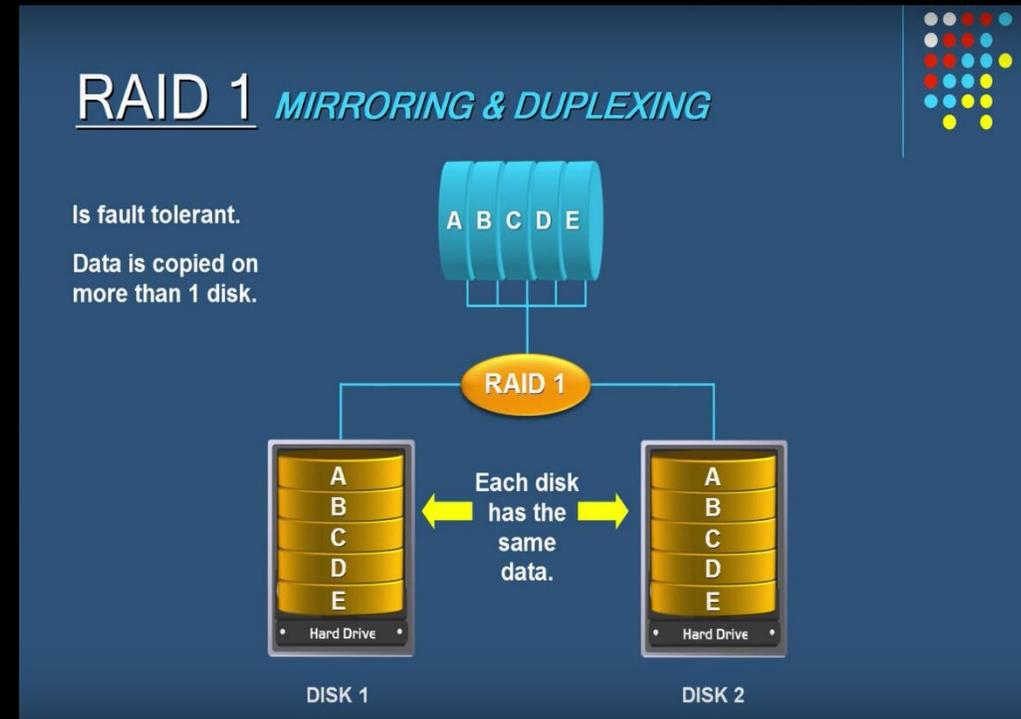
Data is 'striped' across multiple disks.



- Single drive failure
- 2 disk configuration
- Limited Storage Capacity

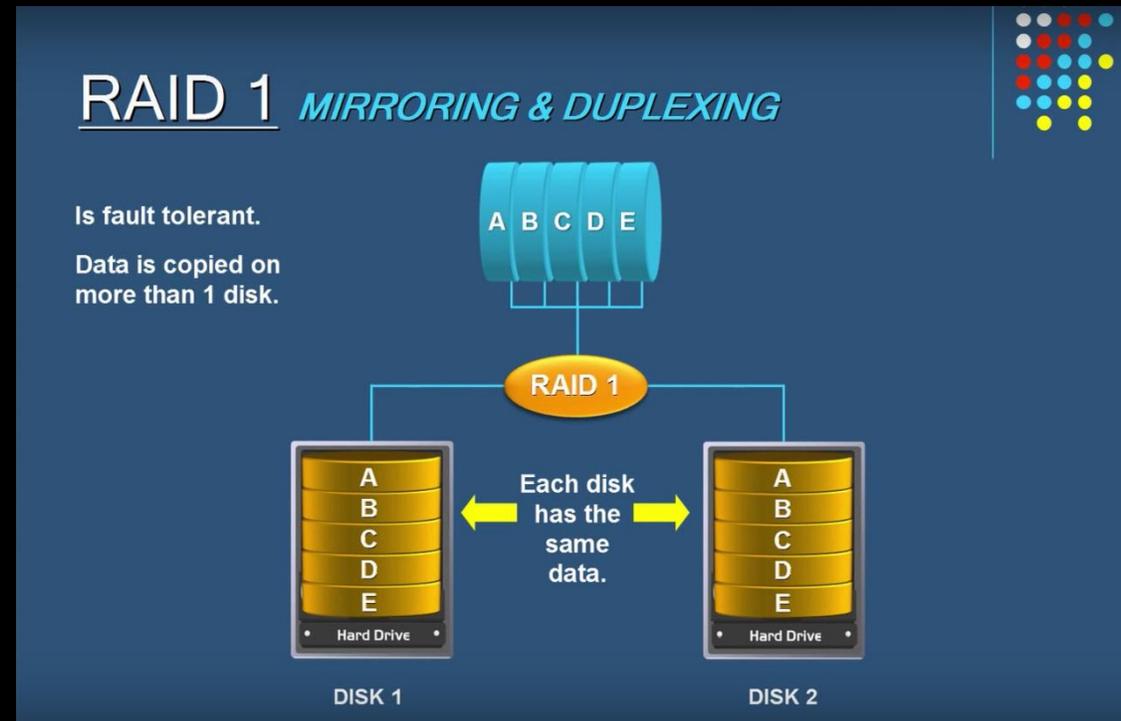
What is Raid 1

- **RAID 1:** RAID 1 is called disk mirroring. All the data is written to at least two separate physical disks. The disks are essentially mirror images of each other. If one of the disks fails, the other can be used to retrieve data.
- Disk mirroring is good for very fast read operations. It's slower when writing to the disks, since the data needs to be written twice. RAID 1 requires at least two physical disks.



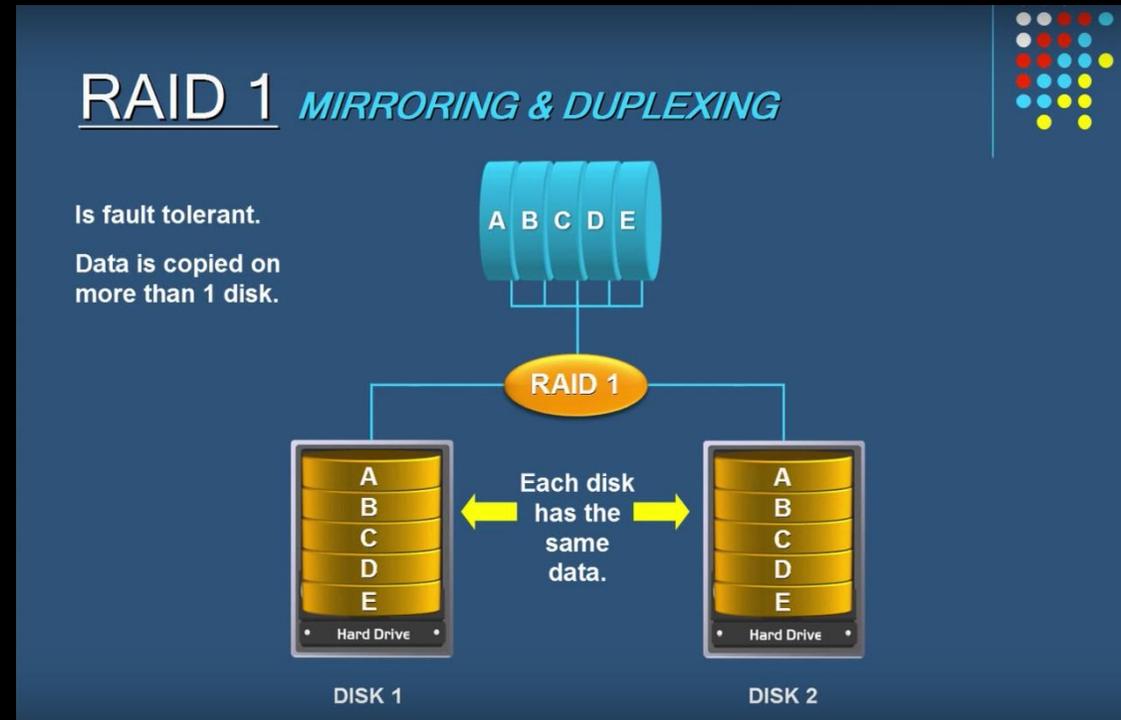
Benefits of Raid 1

- Fault tolerant
- Only requires at least 2 disk
- Simple implementation



Cons of Raid 1

- Not as fast as Raid 0
- Requires at least 2 disks
- Single point of failure on 2nd disk



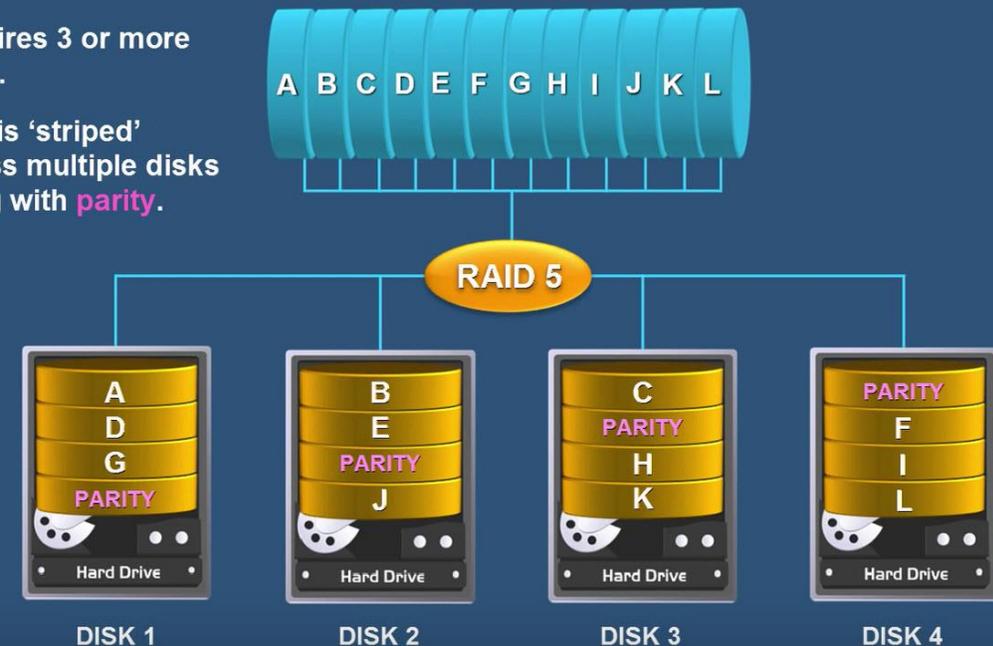
What is Raid 5?

RAID 5: Uses disk striping with parity. The data is striped across all the disks in the RAID set, along with the parity information needed to reconstruct the data in case of disk failure. RAID 5 is the most common method used, since it achieves a good balance between performance and availability. RAID 5 requires at least three physical disks.

RAID 5 *STRIPING WITH PARITY*

Requires 3 or more disks.

Data is 'striped' across multiple disks along with **parity**.



RAID Parity

- RAID Parity
- Parity computations are used in RAID drive arrays for fault tolerance by calculating the data in two drives and storing the results on a third. The parity is computed by XOR'ing a bit from drive 1 with a bit from drive 2 and storing the result on drive 3 (to learn about XOR, see [OR](#)). After a failed drive is replaced, the RAID controller rebuilds the lost data from the other two drives. RAID systems often have a "hot" spare drive ready and waiting to replace a drive that fails.

Benefits of Raid 5

- Most common configuration
- Fault tolerant
- Very fast
- Stores large amounts of data
- Can be rebuilt if 1 drive fails
- Parity (hotspare)



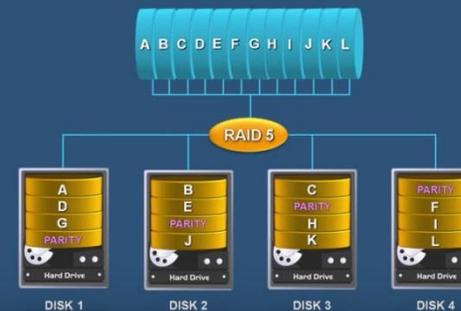
Cons of Raid 5

- Requires more disks (at least 3)
- More expensive
- Entire disk is used in each array
- 25% of disks are used for the array
- More than 1 disk fails then data loss occurs

RAID 5 *STRIPING WITH PARITY*

The equivalent of an entire disk is used to store parity.

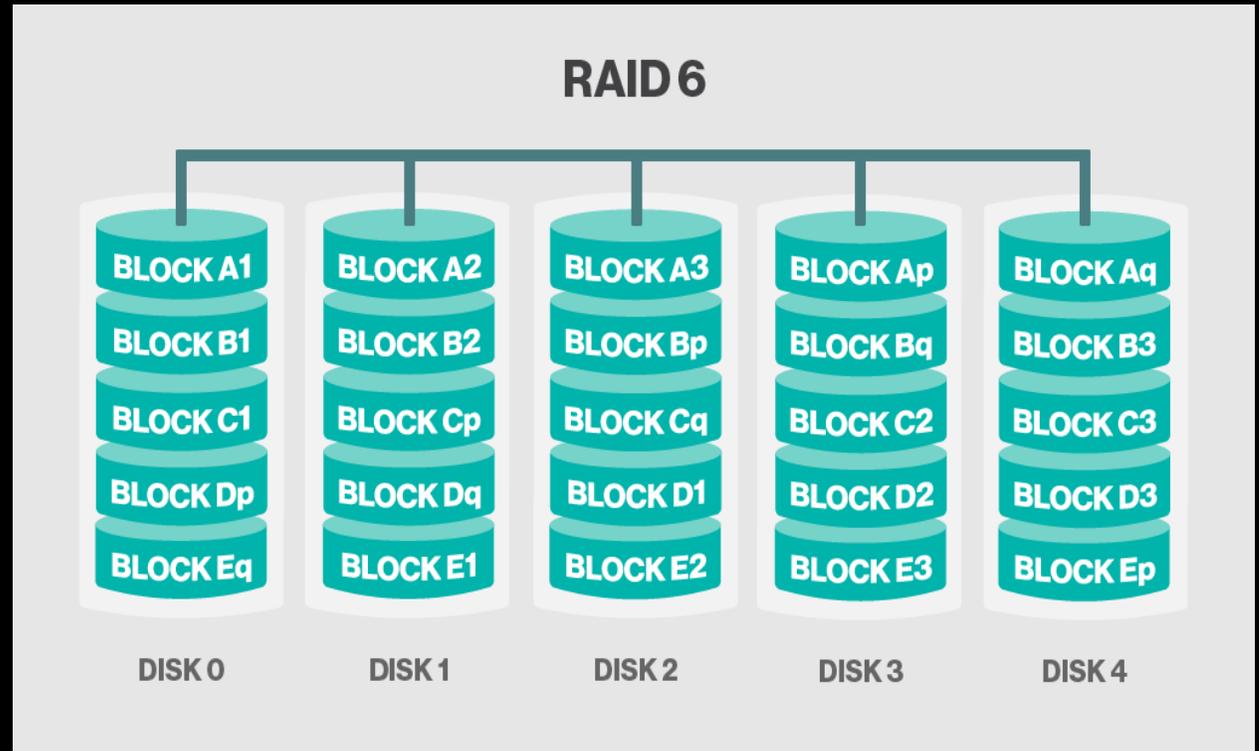
Example: An array of 4 disks totaling 4 terabytes, only 3 terabytes will be used for actual data storage.



What is Raid 6? Double-Parity RAID

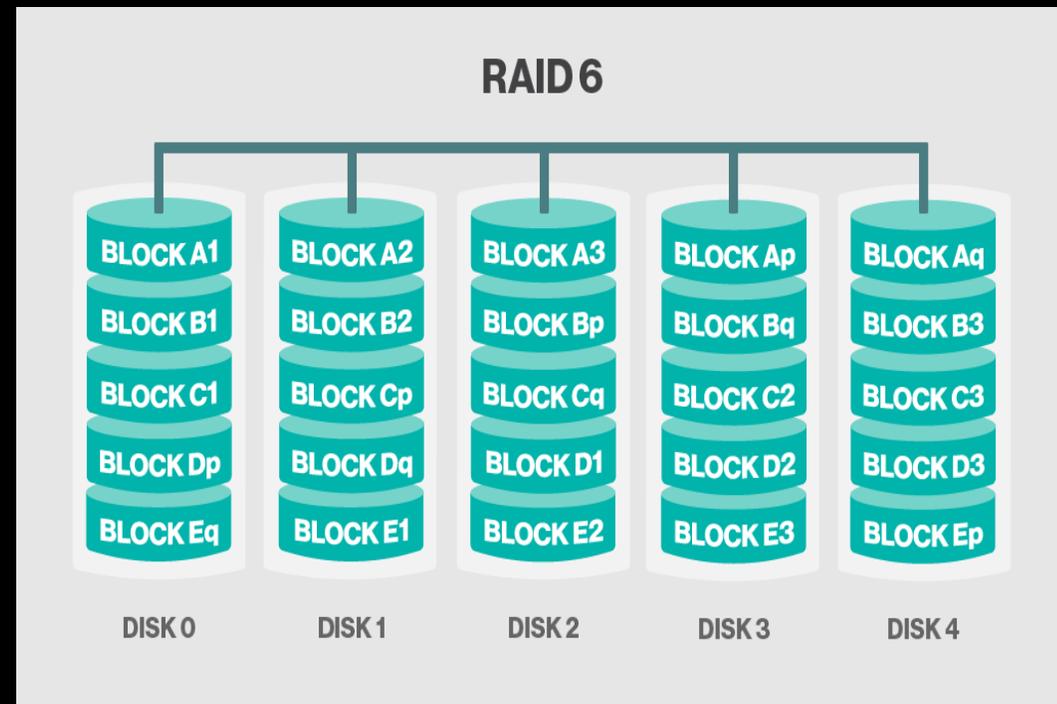
RAID 6 increases reliability by utilizing two parity stripes, which allows for two disk failures within the RAID set before data is lost. RAID 6 is seen in SATA environments, and solutions that require long data retention periods, such as data archiving or disk-based backup.

Raid 6 is rapidly becoming a standard component of modern storage systems. Raid 6 is able to recover from read errors better than Raid 5



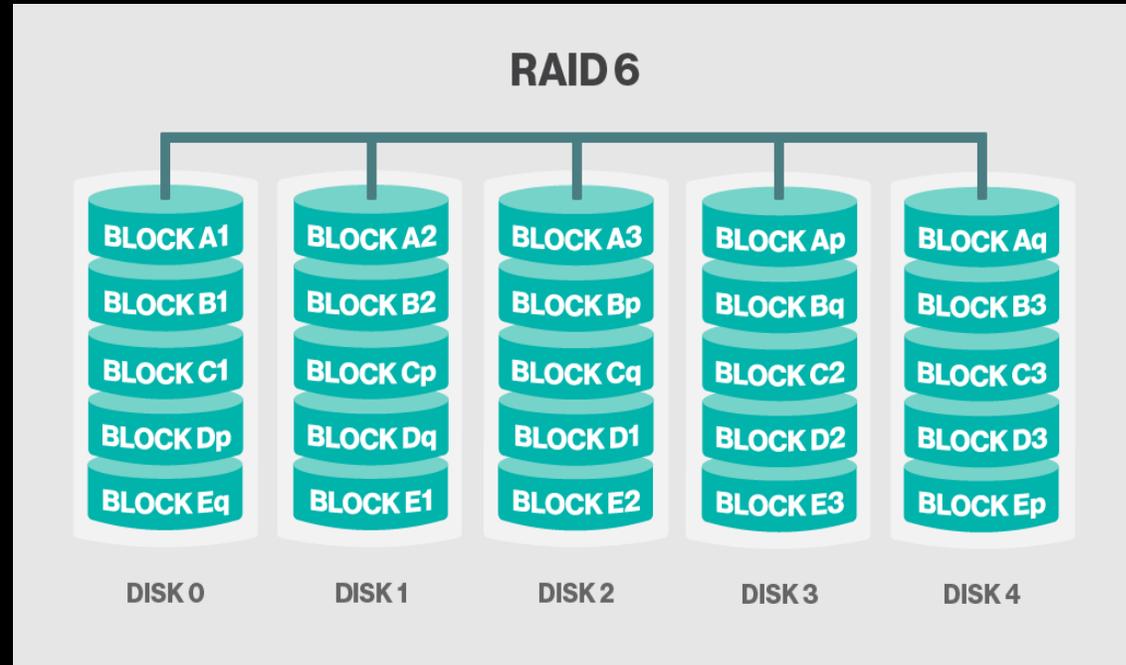
Benefits of Raid 6

- Supports more disk arrays
- More disks can fail
- Dual parity stipes
- Rebuilding is more reliable than Raid 5



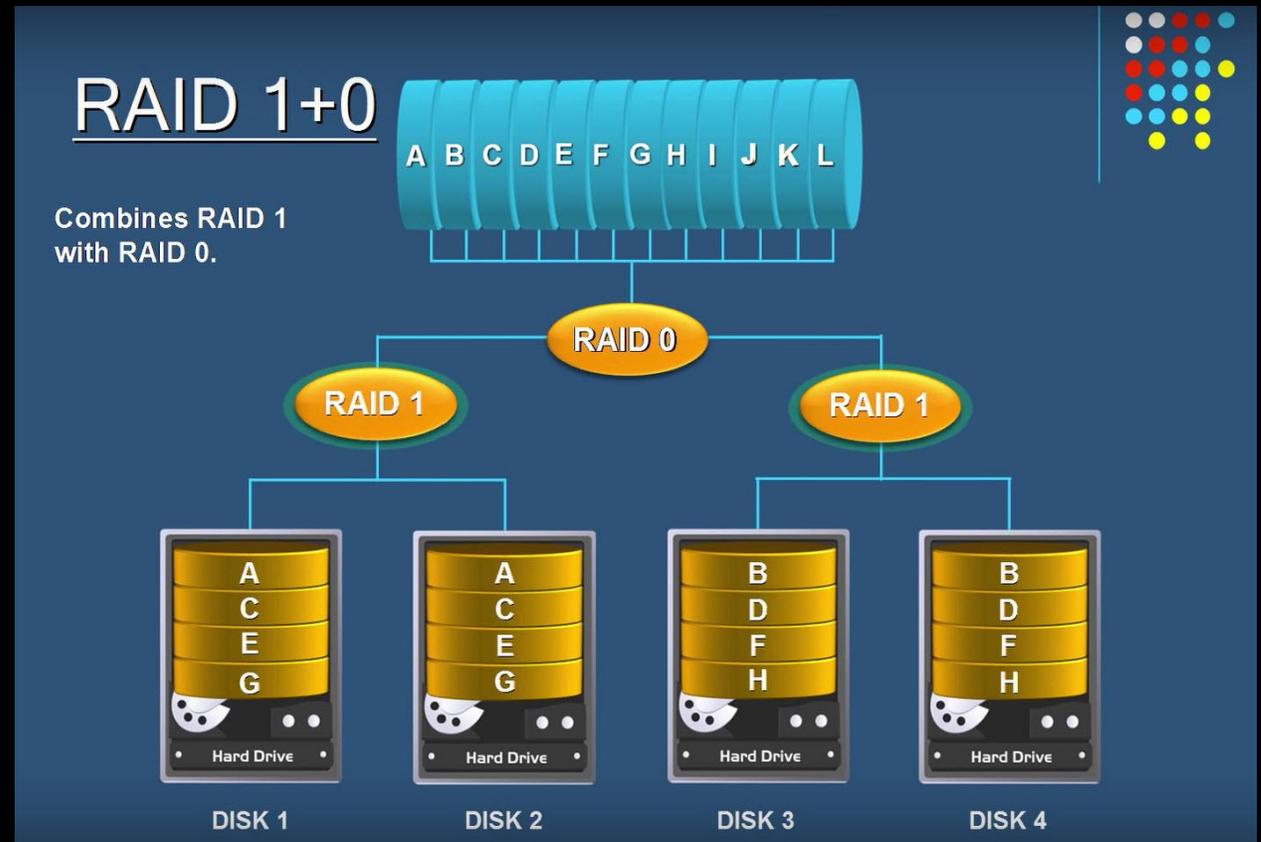
Cons of Raid 6

- More complex setup and configuration
- More expensive to implement
- Lose 2 drives worth of storage



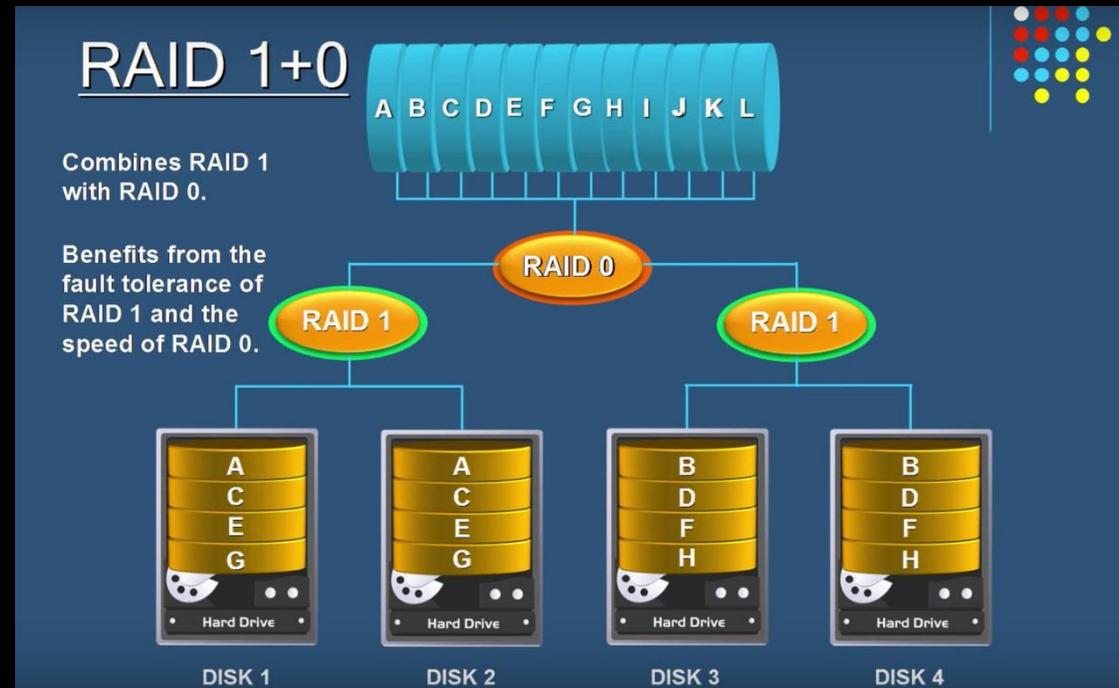
What is Raid 10?

- RAID 10, also known as RAID 1+0, combines disk mirroring and disk striping to protect data
- A RAID 10 configuration requires a minimum of four disks, and stripes data across mirrored pairs. As long as one disk in each mirrored pair is functional, data can be retrieved. If two disks in the same mirrored pair fail, all data will be lost because there is no parity in the striped sets.
- RAID 10 provides redundancy and performance, and is the best option for I/O-intensive applications.



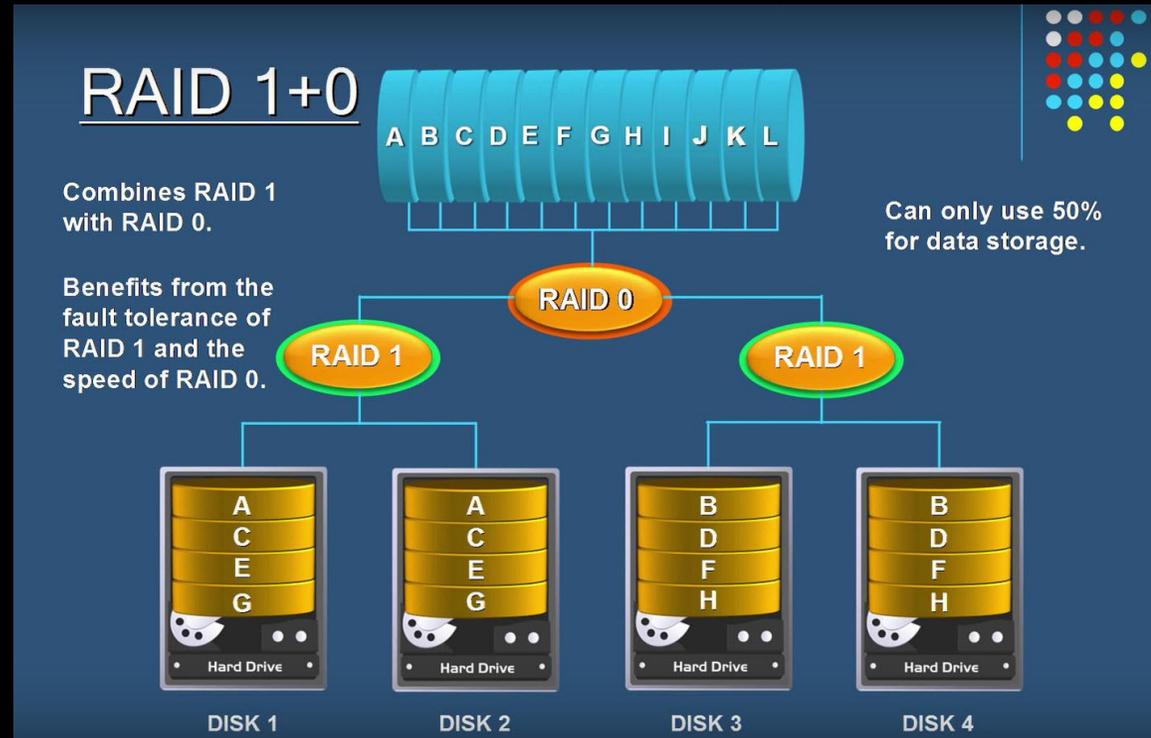
Benefits of Raid 10

- Combines Raid 1 and Raid 0
- Fault tolerant
- Very fast



Cons of Raid 10

- Uses 50% of storage
- More expensive
- Requires more disks



Traditional RAID levels

LEVEL	DESCRIPTION
▶ RAID 0	<ul style="list-style-type: none"> Data striped across hard disk drives (HDDs) for maximum write performance No actual data protection
▶ RAID 1	<ul style="list-style-type: none"> Synchronously mirrors all data from each HDD to an exact duplicate HDD No data lost if HDD faults or fails Typically highest performing RAID level at the expense of lower usable capacity
▶ RAID 2	<ul style="list-style-type: none"> Data protected by error correcting codes (ECCs) Parity HDD requirements proportional to the log of HDD number Somewhat inflexible and less efficient than RAID 5 or RAID 6 with lower performance and reliability Not widely used
▶ RAID 3	<ul style="list-style-type: none"> Data is protected against the failure of any HDD in a group of N+ Similar to RAID 5 but blocks are spread across HDDs Parity is bitwise vs. RAID 5 block Parity resides on a single HDD rather than being distributed among all disks Random write performance is quite poor and random read performance fair at best
▶ RAID 4	<ul style="list-style-type: none"> Similar to RAID 3, stripes data across many HDDs in blocks instead of RAID 3 bytes to improve random access performance Data protection is provided by a dedicated parity HDD Similar to RAID 5 except uses dedicated parity instead of distributed parity Dedicated parity HDD remains a bottleneck, especially for random write performance
▶ RAID 5	<ul style="list-style-type: none"> Most common RAID Provides RAID 0 performance with more economical redundancy Stripes block data across several HDDs while distributing parity among the HDDs Uses HDDs more efficiently, providing overlapped read and write operations Provides more usable storage than RAID 1 or RAID 10 Data protection comes from parity information used to reconstruct data of a failed drive Minimum of three and usually five HDDs per RAID group Rebuilds cause lower storage system performance Potential total RAID group data loss if second drive fails during rebuild Read performance tends to be lower than other RAID types because parity data is distributed on each HDD
▶ RAID 6	<ul style="list-style-type: none"> Similar to RAID 5 but includes a second parity scheme distributed across the HDDs of the RAID group Dual parity protects against data loss if second HDD fails Tends to have lower storage system performance than RAID 5 and can plummet during dual HDD rebuilds
▶ RAID 10 (RAID 1 + RAID 0)	<ul style="list-style-type: none"> RAID 1 striped Improves write performance
▶ RAID 50 (RAID 5 + RAID 0)	<ul style="list-style-type: none"> RAID 5 striped Improves write performance closer to RAID 1
▶ RAID 60 (RAID 6 + RAID 0)	<ul style="list-style-type: none"> RAID 6 striped Improves write performance closer to RAID 1

RADONC IT – Ohio State University

- What is Radonc IT and how does it function at OSU?
- How does Radonc IT support the Oncology Department?
- What are the benefits of communication between Radonc IT and MedCtr IT?
- What are the benefits of communication between Radonc IT and vendors such as Varian?
- How closely does Radonc IT work with the physics department
- What does the OSU Radonc IT Department look like?

Aria/Eclipse Production Environment

Xenapps5-p01 10.127.18.87	Xenapps5-bp07 10.127.18.156	Xenapps5-bp13 10.127.18.162
Xenapps5-p02 10.127.18.88	Xenapps5-bp08 10.127.18.157	
Xenapps5-p03 10.127.18.89	Xenapps5-bp09 10.127.18.158	
Xenapps5-p04 10.127.18.90	Xenapps5-bp10 10.127.18.159	
Xenapps5-p05 10.127.18.91	Xenapps5-bp11 10.127.18.160	
Xenapps5-p06 10.127.18.92	Xenapps5-bp12 10.127.18.161	



RAD-ARIADB-P01
10.81.0.188



Brainlab-p01
10.127.18.123



Brainlab-p02
10.127.18.124



Ariaiem-p01
10.127.25.71



Ariaikoe-p01
10.127.25.75

Ariafas-p14 10.127.26.86	Ariafas-p15 10.127.26.87	Ariafas-p16 10.127.26.88
Ariafas-p17 10.127.26.89	Ariafas-p07 10.127.25.61	Ariafas-p12 10.127.25.66
Ariafas-p03 10.127.25.57	Ariafas-p08 10.127.25.62	Ariafas-p13 10.127.25.67
Ariafas-p04 10.127.25.58	Ariafas-p09 10.127.25.63	
Ariafas-p05 10.127.25.59	Ariafas-p10 10.127.25.64	

VariantBox



ARIAHARP-P01
10.127.25.72

Xenapps5-P14-K 10.127.18.163	Xenapps5-P15-K 10.127.18.164	Xenapps5-P16-K 10.127.18.165
---------------------------------	---------------------------------	---------------------------------

Ackerman/James VM Environment

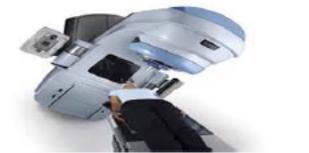
Lantisbk 10.80.3.24	Lantisbackup 10.65.160.183
Rmed-Fs02 10.80.3.25	Rmed-Ps01 10.80.3.26
Radnamer 10.80.3.198	Rmed-Acronis1 10.65.80.3.34
Lantis-X1 10.65.160.230	Lantis-X2 10.65.160.231
Lantis-X3 10.65.160.232	Rmed-Centos1 10.65.160.250
Eswackrmed-p01 10.80.3.20	Eswackrmed-p02 10.80.3.22
Rmed-Dicomsvr1 10.80.3.33	Rmed-Dicompile 10.65.160.93

Non-Clinical Eclipse

Lantis 10.65.160.28	Imagegrid 10.65.160.20	JRADADT 10.65.160.29
Rmed-Upsvr 10.65.160.30	Radvm 10.65.160.31	Radvmctx 10.65.160.56
Rmed-ctxutil 10.65.160.24	Variantbox 10.65.160.69	Exactrac 10.65.160.223
Gamaplan 10.65.160.21	Ct Scanner 10.65.160.73	Nomos 10.65.160.66
Efilm 10.65.160.77	Rmoncologist 10.65.160.86	KodakCr 10.65.160.23
SBRT 10.65.160.39	Oncor 10.65.160.34	Mbe 10.65.160.37
Acuity 10.65.160.237	Brainlab-Iplan 10.65.160.224	Radvarian8 10.65.160.40
Rmed-Redhat 10.65.160.235	Trubeam 10.65.160.225	Radvarian12 10.65.160.62

Training Environment

Trubeam 10.37.216.20	VisionRt 10.37.216.36	Ct Scanner 10.37.216.37
Gating 10.37.216.59	Lap Laser 10.37.216.39	AW Workstation 10.37.216.38



ARIACONNECTSVR



ARIACONNECTSVR
10.216.14.80

STORAGE

MANAGING THE INFORMATION THAT DRIVES THE ENTERPRISE

Hot data storage technologies for 2016

What's hot? What's not? These techs will have an impact on your data centers in the coming year.



SNAPSHOT 1

A third of shops use cloud storage, but more plan to

EDITOR'S NOTE / CASTAGNA

A data-centric view of storage

VDI

Architecting storage for VDI

STORAGE REVOLUTION / TOIGO

The Z(ettabyte) apocalypse

SNAPSHOT 2

Data protection top cloud storage use; security fears remain

HOT SPOTS / BUFFINGTON

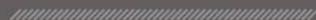
Data protection and the audacity of BYOD

DR 101

DR 101: Get back to the basics

READ-WRITE / TANEJA

Storage performance still an issue



QUESTIONS?

LINKS – Http://

- <http://www.hometoys.com/article/2014/11/the-pros-cons-of-software-vs-hardware-raid/2219/>
- <http://www.pcmag.com/article2/0,2817,2370235,00.asp>
- <http://www.pcmag.com/encyclopedia/term/50148/raid>
- <http://searchstorage.techtarget.com/answer/RAID-types-and-benefits-explained>