

---

# RAID by Sight and Sound

---

By Scott A. Moulton  
Of  
My Hard Drive Died &  
Forensic Strategy Services LLC

# Topics for Talk

---

- Brief Coverage of NOT RAID
- Brief Intro to RAID
- RAID 0
- RAID 5
- Demo of Rebuild

# Assumptions

---

- We are assuming you have already done what I previously described in previous videos and classes to repair the damaged drive.
- You now have a running drive and have imaged it in some fashion and will use that for reassembly.

# RAID Array's

---

- Redundant Array of (Inexpensive or Independent) Disks.
- Some arrays are not “Redundant”
- Different types of arrays may need different numbers of drives in the array and may give you different results in free space available by how they are setup.

# JBOD Drives

---


- Means “Just a Bunch of Disks” and they are just linked logically together end to end.
- These drives usually have no fan, get very hot and contain several drives. Sometimes the cables are melted together.
- I have found that some recoveries that are using these JBOD will work once repaired, even if the Lacie board is burnt out, just by placing them in a G5 and connecting them and booting on a Mac external disk. This is great for when you cannot repair the Lacie board or power supply.

# LaCie and other JBOD Drives

---



# LaCie and other JBOD Drives



# LaCie NAS Boxes









# Dynamic Disks

---

- Dynamic disks do not use partition tables, they use LDM which is at the end of the disk and needs to be done backwards. It uses one single partition occupying the entire disk minus one cylinder. When volumes are added or deleted the partition table is not updated.
- This will be noticed right away by some data recovery software like R-Studio.

# Types of RAID Arrays Overview


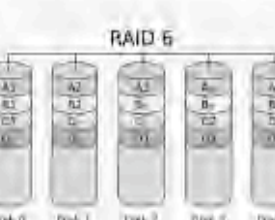
RAID 0 is one of the types of RAID that are often in for recovery

Level	Description	Minimum # of disks	Image
RAID 0	<b>Striped set without parity.</b> Provides improved performance and additional storage but no fault tolerance. Any disk failure destroys the array, which becomes more likely with more disks in the array. A single disk failure destroys the entire array because when data is written to a RAID 0 drive, the data is broken into fragments. The number of fragments is dictated by the number of disks in the drive. The fragments are written to their respective disks simultaneously on the same sector. This allows smaller sections of the entire chunk of data to be read off the drive in parallel, giving this type of arrangement huge bandwidth. RAID 0 does not implement error checking so any error is unrecoverable. More disks in the array means higher bandwidth, but greater risk of data loss. <a href="#">SNIA definition</a>  .	2	 <p>RAID 0</p> <p>Disk 0    Disk 1</p>
RAID 1	<b>Mirrored set without parity.</b> Provides fault tolerance from disk errors and single disk failure. Increased read performance occurs when using a multi-threaded operating system that supports split seeks, very small performance reduction when writing. Array continues to operate so long as at least one drive is functioning. <a href="#">SNIA definition</a>  .	2	 <p>RAID 1</p> <p>Disk 0    Disk 1</p>

From Wikipedia.org

# Types of RAID Arrays Overview

RAID 5 is one of the most common types of RAID that are seen for recovery

RAID 5	<p><b>Striped set with distributed parity.</b> Distributed parity requires all drives but one to be present to operate; drive failure requires replacement, but the array is not destroyed by a single drive failure. Upon drive failure, any subsequent reads can be calculated from the distributed parity such that the drive failure is masked from the end user. The array will have data loss in the event of a second drive failure and is vulnerable until the data that was on the failed drive is rebuilt onto a replacement drive. <a href="#">SNIA definition</a></p>	3	 <p>RAID 5</p> <p>Drive 0 Drive 1 Drive 2 Drive 3</p>
RAID 6	<p><b>Striped set with dual parity.</b> Provides fault tolerance from two drive failures; array continues to operate with up to two failed drives. This makes larger RAID groups more practical, especially for high availability systems. This becomes increasingly important because large-capacity drives lengthen the time needed to recover from the failure of a single drive. Single parity RAID levels are vulnerable to data loss until the failed drive is rebuilt: the larger the drive, the longer the rebuild will take. Dual parity gives time to rebuild the array without the data being at risk if one drive, but no more, fails before the rebuild is complete. <a href="#">SNIA definition</a></p>	4	 <p>RAID 6</p> <p>Drive 0 Drive 1 Drive 2 Drive 3 Drive 4</p>

From Wikipedia.org

# RAID 0: How it works

---


- RAID 0 has no redundancy for protecting data and if one drive fails you loose the array.
- IN short RAID 0 should be called AI DS.
  - Array of Inexpensive Drives that Suck

# RAID 0: How it works

---

- RAID 0 is a very simple type of raid that is used mainly to gain speed and performance by spreading your reads and write cycles over two drives.
- RAID 0 can, depending on the controller and the choice of the manufacture, vary the size and location of the strip and it is helpful to know the brand and type of controller.

# RAID 0 strips with only two drives



From Wikipedia.org

# RAID 0 Two Drives +

---

- There are RAID 0 arrays with more than two drives, however there is no sequencing numbers and the order of the array is determined by the data's location there is no way to determine the order of the drives in the array.
- With two drives, the array is either stored in one order, or the other, but if there are 4 drives in the array, there are at least 72 different combinations and no indication of what way is correct. This will make recovery exponentially difficult and in many cases a disaster.

# RAID 0

---

- In most cases you can determine the first drive in the array, depending on the slice size.
- How?
- If the slice size is larger than 32k, at sector 63 you will see the active boot partition, in most cases...

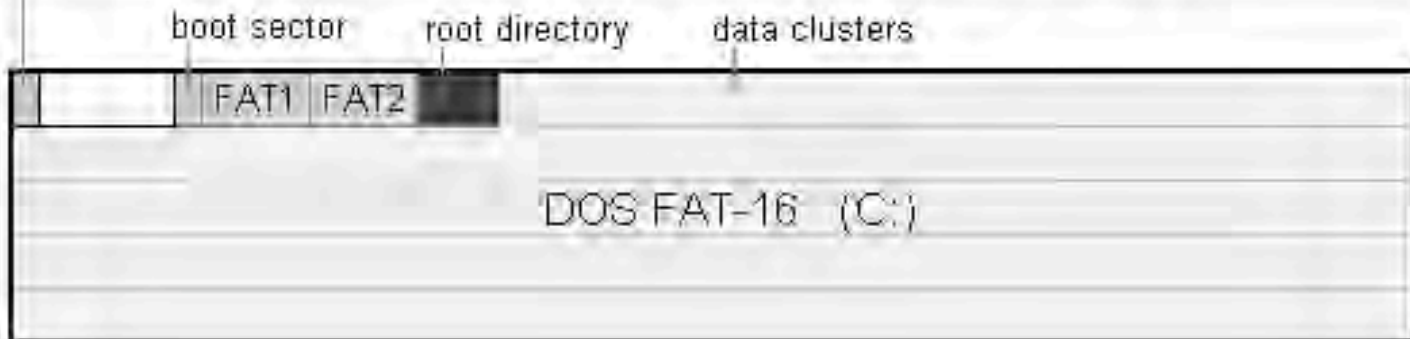


# Partition Example

**Example 1** Hard disk 340M [ 665 cyl x 16 heads x 63 sects ]

## MBR

#	Partition Type	Starting			Ending			Size[K]
		Cyl	Side	Sect	Cyl	Side	Sect	
1	DOS FAT-16	0	1	1	664	15	63	335,128
2	Unused	0	0	0	0	0	0	
3	Unused	0	0	0	0	0	0	
4	Unused	0	0	0	0	0	0	



From <http://www.ranish.com/part/primer.htm>

# NTFS Boot Sectors

```
Physical Sector: Cyl 0, Side 1, Sector 1
00000000: EB 52 90 4E 54 46 53 20 - 20 20 20 00 02 08 00 00 .R,NTFS .....
00000010: 00 00 00 00 00 F8 00 00 - 3F 00 FF 00 3F 00 00 00 .....?..?..
00000020: 00 00 00 00 80 00 80 00 - 1C 91 11 01 00 00 00 00 .....
00000030: 00 00 04 00 00 00 00 00 - 11 19 11 00 00 00 00 00 .....
00000040: F6 00 00 00 01 00 00 00 - 3A B2 7B 82 CD 7B 82 14 .....!..{..{..
00000050: 00 00 00 00 FA 33 C0 8E - D0 BC 00 7C FB B8 C0 07 .....3.....|....
```

From Microsoft.com

# RAID 0

---

- Put the first drive in the first slot of whatever software you are using..
- Put the other drives in their slots
- Set your size of your slice to your guess.... Usually 64 is the defaults (unless some tech messed with it)
- Scan for Pictures (JPG,JPEG,GIF) or MP3s.
- Stop, extract, view, listen...

# What to Extract

---

- Extract between the boundaries of the controller...if you don't know, usually guess:
  - » 32k
  - » 64k
  - » 128k
  - » 256k
  - » 512k
  - » 1024k

---

---

How do you know when you  
are wrong??

---


# Large Files in Megs

---




# Recognizable 140k File..

---



# Small Files under 32k Intact

---






# Small Files under 64k Intact

---



# Files Just Over 64k

---




# Files Just Over 64k

---



# Files Over 2 Megs

---



# Large RAW Files



# Once you get it Right – 700k

---



---

---

# 2 Meg MP3 Sound File

---

Sample



# RAID 5






# RAID 5: How it works

---

- RAID 5 Array protects the server from “down time.” It will allow for a drive to fail and your system to continue running without the result of lost data.
- RAID 5 does this by storing parity data on all the hard drives. Parity is a formula that calculates error correction data.
- By distributing parity across all drives it creates a safety net for the data when a drive fails.

# RAID 5: How it works



# RAID 5: Controllers

---








- There are two kinds of controllers for RAID, Host Based and Discrete controllers.
- Host controllers use the processor power in your computer to process the calculations for the array.
- Discrete controllers have a processor that will do the mathematics for calculating the array values. A discrete controller is much faster and leaves your processor to do other tasks, however, they are much more expensive.
- You are going to try to do this in software!

# RAID 5: How it works

---

- There have also been times where RAID 5 arrays have failed a single drive, but no one noticed before a second one failed.
- If two drives fail and the array goes down, which drive do you need to repair???

# RAID 5 Simplist Sample



# RAID5 For Real??

---

- Parity is calculated by XORing the data with the number of slices in the row -1. For 4 drives it looks like this:

SliceA xor SliceB xor SliceC xor SliceD = Parity

# Arrangements: Left Async

RAID-5 Segments  
Left Asynchronous

Drive 0	Drive 1	Drive 2	Drive 3	Drive 4
0	1	2	3	Parity
4	5	6	Parity	7
8	9	Parity	10	11
12	Parity	13	14	15
Parity	16	17	18	19
20	21	22	23	Parity

# Arrangements: Left Sync

RAID-5 Segments  
Left Synchronous

Drive 0	Drive 1	Drive 2	Drive 3	Drive 4
0	1	2	3	Parity
5	6	7	Parity	4
10	11	Parity	8	9
15	Parity	12	13	14
Parity	16	17	18	19
20	21	22	23	Parity



# Arrangements: Right Async

RAID-5 Segments  
Right Asynchronous

Drive 0	Drive 1	Drive 2	Drive 3	Drive 4
Parity	0	1	2	3
4	Parity	5	6	7
8	9	Parity	10	11
12	13	14	Parity	15
16	17	18	19	Parity
Parity	20	21	22	23

# Arrangements: Right Sync

RAID-5 Segments  
Right Synchronous

Drive 0	Drive 1	Drive 2	Drive 3	Drive 4
Parity	0	1	2	3
7	Parity	4	5	6
10	11	Parity	8	9
13	14	15	Parity	12
16	17	18	19	Parity
Parity	20	21	22	23

# Steps to rebuild RAID 5 array

---

1. Repair all necessary BAD drives.
2. Image the damaged drive(s) and recover as many sectors as possible.
3. Image all the good drives.
4. Use software to analyse and re-weave the images back together virtually. Test data!
5. Write the newly weaved image back to a hard drive to start the logical recovery (follow the logical recovery section for the type of format).

# RAID 5

---

- Put the first drive in the first slot of whatever software you are using..
- Put the other drives in their slots
- Set your size of your slice to your guess.... And your ARRANGMENT to the order.
- Scan for Pictures (JPG, JPEG, GIF) or MP3s.
- Stop, extract, view, listen...

# Code to do it for you...

- ```
#!/usr/bin/perl -w
#
# raid5 perl utility
# Copyright (C) 2005 Mike Hardy <mike [at] mikehardy.net>
#
# This script understands the default linux raid5 disk layout,
# and can be used to check parity in an array stripe, or to calculate
# the data that should be present in a chunk with a read error. my [at] array_components.
= (
"/dev/loop0",
"/dev/loop1",
"/dev/loop2",
"/dev/loop3",
"/dev/loop4",
"/dev/loop5",
"/dev/loop6",
"/dev/loop7"
);
my $chunk_size = 64 * 1024; # chunk size is 64K
my $sectors_per_chunk = $chunk_size / 512;
```

**[http://www.freesoftwaremagazine.com/articles/recovery\\_raid](http://www.freesoftwaremagazine.com/articles/recovery_raid)**

---

---

# RAID Live Demo

---

Using R-Studios

# The End

---

