Computer Back-up Systems

Computer Failures

1. **Oops** = I just deleted a file I needed and emptied the trash so it REALLY gone. HELP!!!

2. **GR-I-I-I-I-ND** = What's that strange noise the disk is making. And the computer won't start up.

3. **STOP, THIEF** = hardware is gone or destroyed. Fire, flood, thief, etc.

1

No perfect solution to preventing data loss

Jerry Lange
June 14, 2010

Data Loss Prevention

Manual solutions - Will you really do it?

Back-up CD's or DVD's -
and how many do you need for the 300 GB of data? Not workable in the multi-gigabyte world of videos, audio tracks and photos.

Backup hard drives - works fine, but will you actually do it?

Automatic solutions

In House

Time Machine - great for data loss from the system drive. No good for system files back-up.

SuperDuper, et al - Can do automatic back-ups of just data or entire system disk (i.e. System clones that can boot the system) Can be automatic or manual.

BUT - What about the destruction of your house? I keep a copy in my detached garage

Out House

Carbonite, Mozy, etc. OK for data, but maybe not for huge amounts of data, 100's of gigabytes for instance. Costs some $ but not outrageous, Around $55/year the ad says for Carbonite unlimited. Can't recreate system disk.

Another tool - RAID Arrays - *Redundant Array of Inexpensive Disks*
Make with hardware (buy it) or software (Snow Leopard Disk Utility)
Common array types - RAID 1 has redundant disk copies

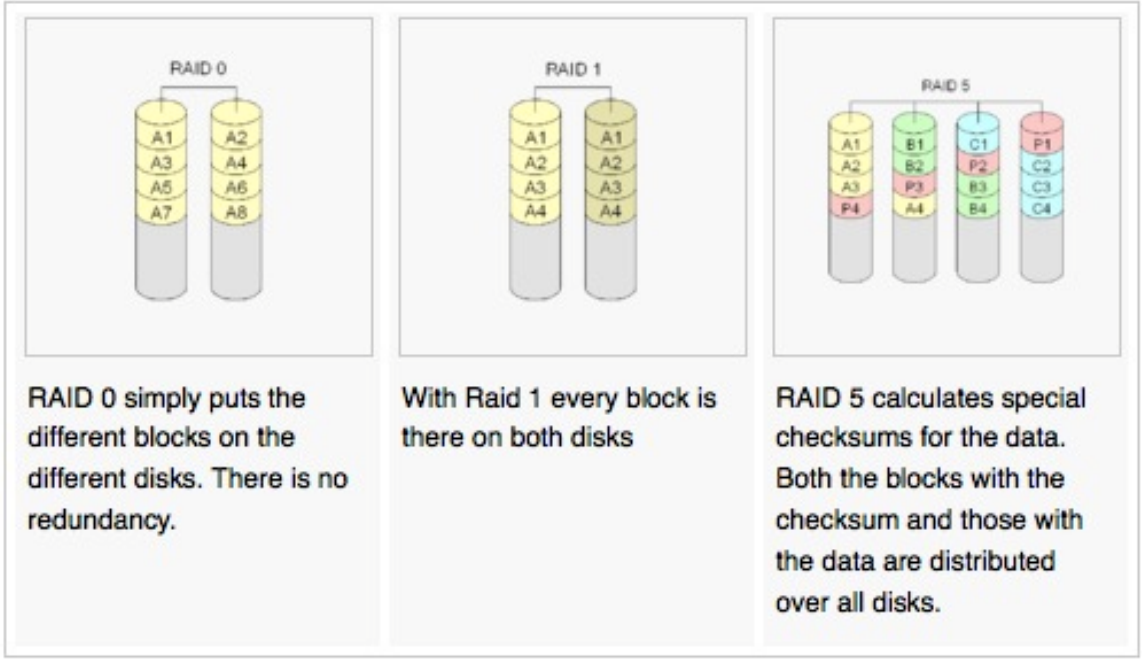

Some not so common types

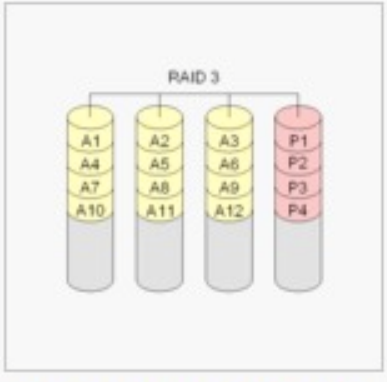

RAID 3 is much like RAID level 1. An extra disk is added that will hold a checksum for each block of data.

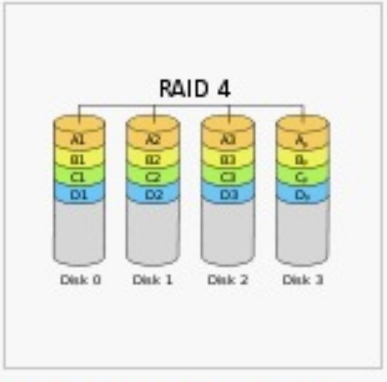

RAID 4 is similar to RAID level 3, but calculates parity over larger blocks of data

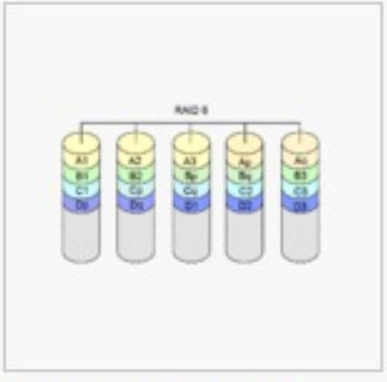

RAID 6 is similar to RAID 5, but it calculates two different checksums. This allows for two disks to fail, without data loss.

And you can go nuts with combos

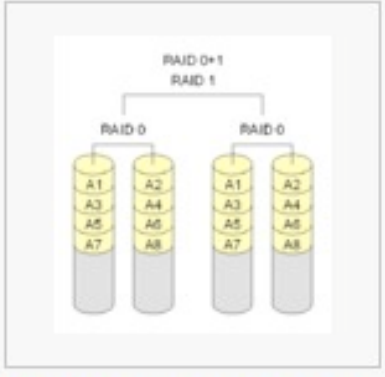

RAID 0+1: Several RAID 0 arrays are combined with a RAID 1

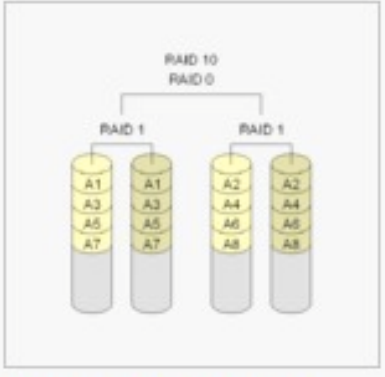

RAID 1+0: More robust than RAID 0+1; supports multiple drive failures, as long as no two drives that make a mirror fail.

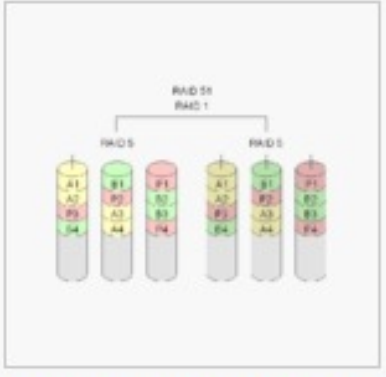

RAID 5+1: Any three drives of this can fail, without data loss.

Some definitions

Mirroring: More than one copy of the data [change]

When talking about a mirror, this is a very simple idea. Instead of the data being in only one place, there are several copies of the data. These copies usually are on different hard disks (or disk partitions). If there are two copies, one of them can fail without the data being affected (as it still is on the other copy). Mirroring can also give a boost when reading data. It will always be taken from the fastest disk that responds. Writing data is slower though, because all disks need to be updated.

Striping: Part of the data is on another disk [change]

With striping, the data is split into different parts. These parts then end up on different disks (or disk partitions). This means that writing data is faster, as it can be done in parallel. This does not mean that there will not be faults, as each block of data is only found on one disk.

Error correction and faults [change]

It is possible to calculate different kinds of checksums. Some methods of calculating checksums allow finding a mistake. Most RAID levels that use redundancy can do this. Some methods are more difficult to do, but they allow to not only detect the error, but to fix it.

Hot spares: using more disks than needed [change]

Many of the ways to have RAID support something is called a *hot spare*. A hot spare is an empty disk that is not used in normal operation. When a disk fails, data can directly be copied onto the hot spare disk. That way, the failed disk needs to be replaced by a new empty drive to become the hot spare.

Stripe size and chunk size: spreading the data over several disks [change]

RAID works by spreading the data over several disks. Two of the terms often used in this context are *stripe size* and *chunk size*.

The *chunk size* is the smallest data block that is written to a single disk of the array. The *stripe size* is the size of a block of data that will be spread over all disks. That way, with four disks, and a stripe size of 64 kilobytes (kB), 16 kB will be written to each disk. The chunk size in this example is therefore 16 kB. Making the stripe size bigger will mean a faster data transfer rate, but also a bigger maximum latency. In this case, this is the time needed to get a block of data.

And for Tom, who has a bunch of old disks laying around

Putting disk together: JBOD, concatenation or spanning [change]

Many controllers (and also software) can put disks together in the following way: Take the first disk, till it ends, then they take the second, and so on. In that way, several smaller disks look like a larger one. This is not really RAID, as there is no redundancy. Also, spanning can combine disks where RAID 0 cannot do anything. Generally, this is called *just a bunch of disks* (JBOD).

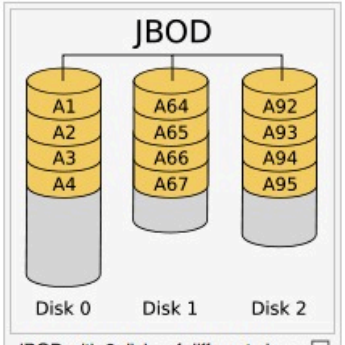

JBOD with 3 disks of different sizes

This is like a distant relative of RAID because the logical drive is made of different physical drives. Concatenation is sometimes used to turn several small drives into one larger useful drive. This can not be done with RAID 0. For example, JBOD could combine 3 GB, 15 GB, 5.5 GB, and 12 GB drives into a logical drive at 35.5 GB, which is often more useful than the drives alone.

In the diagram to the right, data are concatenated from the end of disk 0 (block A63) to the beginning of disk 1 (block A64); end of disk 1 (block A91) to the beginning of disk 2 (block A92). If RAID 0 were used, then disk 0 and disk 2 would be truncated to 28 blocks, the size of the smallest disk in the array (disk 1) for a total size of 84 blocks.

Some RAID controllers use JBOD to talk about working on drives without RAID features. Each drive shows up separately in the operating system. This JBOD is not the same as concatenation.

Many Linux systems use the terms "linear mode" or "append mode". The Mac OS X 10.4 implementation — called a "Concatenated Disk Set" — does not leave the user with any usable data on the remaining drives if one drive fails in a concatenated disk set, although the disks otherwise operate as described above.

Concatenation is one of the uses of the Logical Volume Manager in Linux. It can be used to create virtual drives.

What do I do?

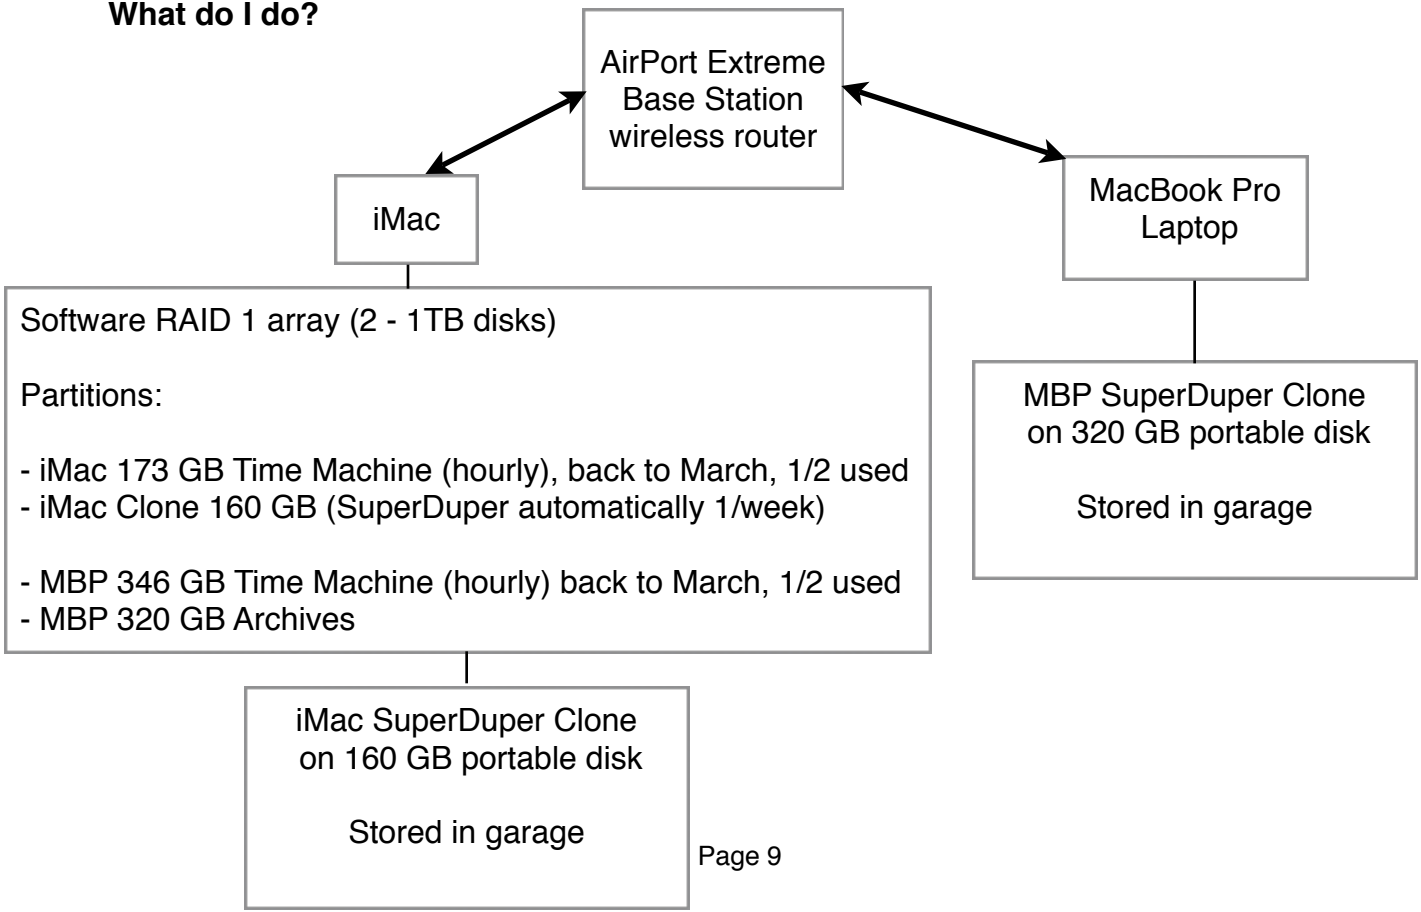
AirPort Extreme
Base Station
wireless router
iMac
MacBook Pro
Laptop
Software RAID 1 array (2 - 1TB disks)
Partitions:
- iMac 173 GB Time Machine (hourly), back to March, 1/2 used
- iMac Clone 160 GB (SuperDuper automatically 1/week)
- MBP 346 GB Time Machine (hourly) back to March, 1/2 used
- MBP 320 GB Archives
MBP SuperDuper Clone
on 320 GB portable disk
Stored in garage
iMac SuperDuper Clone
on 160 GB portable disk
Stored in garage