# RAID for Broadcast Applications

Extracted and revised from Broadcast Engineering January 2006
This article has been edited and extended by mediaproxy Pty Ltd, Australia

Most broadcasters have heard or used the term RAID

**RAID - Redundant Array of Inexpensive Disks**

Typically, RAID systems work by storing extra information along with the regular data. If a disk fails, these extra pieces of information allow the array to rebuild the original stored information in order to provide redundancy or protection.

This article focuses on key features and considerations for RAID when used in the broadcast environment. There are several features that are important to consider in generic RAID applications.

The following list of features and considerations are especially important for broadcasters who are using RAID in their facilities for video and audio content.

**RAID--1 / RAID-5 arrays**

While there are many types of RAID available – each offering different levels of performance, protection and cost, the two variants discussed in this paper are RAID-1 (mirroring) and RAID-5 (single drive parity protection).

RAID-1 operates by maintaining two identical copies of a physical disk drive – in the event that one fails, the second drive continues operation independently until a replacement 'mirror' drive is available and rebuilt in the background. RAID-1 capacity is limited by the physical disk capacity of any single disk drive.

RAID-5 requires 'n+1' drives to provide a dynamic protection capability. The disk capacity is based on the first 'n' drives, while the extra disk maintains protection information - able to replace the functionality of any one of those 'n' drives instantly in the event one should fail. RAID-5 capacity is limited by the number of drives that can be assigned to a single array (within practical and hardware limitations).

**Hot-Swappable Drives**

Many RAID arrays are designed *so* that drives can be replaced while the unit is still in service. If the objective of having a RAID solution is to avoid downtime, then this feature can be fundamental.

Hot-swappable drives have special connectors that are designed to break and make connections in a specific order so that components are not destroyed as the drive is removed or installed. Many people think that hot swappable also means that the drive will automatically rebuild; however, this is not necessarily the case.

**Hot Spare / Cold Spare Drive**

Within a RAID array – there is the option of assigning an available disk drive as a 'hot spare' – which is put in service automatically when one of the *active* drives goes offline unexpectedly. A cold spare drive is physically available, but not brought online automatically.

Note, that when a spare drive is brought into a RAID array – the logical drive presentation is unchanged to the operating system and applications.
However the physical disk drive positions within the RAID chassis may appear in any order – especially after several drive replacements. Strict attention to correct disk-connector, drive-bay and cable-numbering is critical to avoid removing the wrong drive accidentally during drive replacement.

**Rebuild in Background**

This feature allows the data on an individual disk to be reconstructed from the data on other good drives in the array. As noted above, RAID arrays can recreate missing data on the fly so the application never notices that a single drive has failed.

RAID arrays also can put this capability to work - to recreate the data and write it to a new drive installed in the array. Of course, you would expect this functionality to be available. But be careful: You may or may not be able to rebuild the data on the new drive without having to shut down the array. This used to be an exotic feature, but it is beaming much more common.

When a single drive fails, the extra drive can be put online - in many cases, automatically. This can be a valuable option for critical RAID arrays. Remember, however, that the drive will still have to be rebuilt. Because there is no way to know in advance which drive will fail, it is impossible to have the extra drive ready to go at a moment's notice.

**SNMP and other Remote Monitoring**

This is one topic I would encourage you to think about carefully If one of the characteristics you requires of a RAID array is to keep on working even when a drive fails… how will you know when a drive fails?

If you lose a drive in a RAID array and then lose another drive in the same array before the first is replaced and rebuilt, will the RAID array keep working?

The answer to the first question is that you will not have any notice that a drive has failed if you are not monitoring the RAID'S status. The answer to the second is that the RAID will not keep working if a second drive fails during the critical period.

It is extremely important that you monitor the health of the RAID array on a regular basis. In some cases, RAID monitoring is built into the application. If a drive quits, the application lets the operator know.

But many times, especially when using generic applications, there is no notification to the user that a drive has failed. With that said, every RAID array I have ever seen has provisions for monitoring drive status, and almost all of them have remote monitoring provisions.  These include e-mail notification, SNMP and other proprietary techniques.

Be sure to incorporate monitoring of RAID arrays into your overall maintenance strategy. If you fail to do this, you are wasting the value of RAID. Eventually, two drives will fail.

When they do, you will be in the same position you would be in with a single large drive.

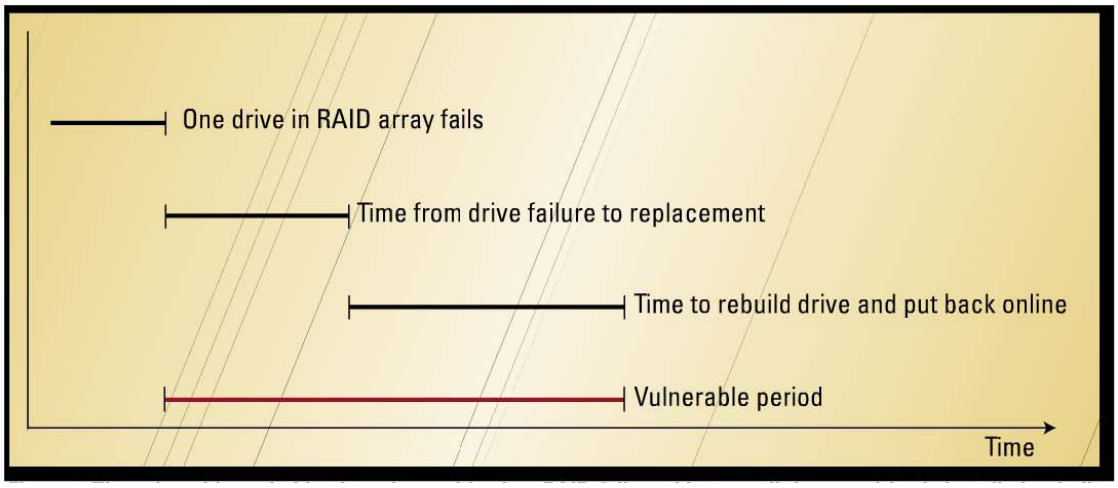


As Figure 1 shows, the vulnerable period begins when a drive in an array fails and lasts until the drive has been replaced, rebuilt and is back online. If, during this vulnerable period, another drive fails, then the RAID array will fail and data will be lost. If you are not monitoring the health of your RAID systems on a regular basis, the time from drive failure to replacement may become lengthy, exposing you to a complete RAID failure.

**Dual Power Supplies and Redundant Cooling**

Dual power supplies and redundant cooling are common in RAID arrays. If you have an important service that you want to protect with RAID, be sure your array has these features. The array should continue to work with a failed cooling fan. I have seen several RAID arrays that incorporate fan speeds and temperature alarms into their remote monitoring facilities. Because drives can fail when they get too hot, it can be important to have remote monitoring on cooling systems.  Remember to monitor the health of these items along with the drives themselves!

**Rebuild Time**

It's important to know how long it takes your system to rebuild after a drive replacement. The more time it takes to rebuild and place a new drive online, the longer you're exposed.

If during the rebuild time you lose another drive, the entire RAID system will fail. While the chances of experiencing two drive failures within a few hours of each other are remote, it is important to know that during the rebuild time, your system is vulnerable.

**Testing**

The best way to know that your RAID system supports the features you expect is to test it. Assuming your system supports hot swapping and rebuilding in background, try the following:

1. Pull out a drive while the system is in operation (hopefully during commissioning tests, and not when it is actually on-air!). Does the system keep running normally? Do you see any indication that the drive has failed either in the application, or on your monitoring system?

2. Plug the drive back in. Do you see any indication that the drive has been installed either in the application, or on your monitoring system? Does the drive begin rebuilding by itself, or do you have to do something to initiate a rebuild? How long does it take before the drive is online again and ready to use?

RAID systems are great examples of how broadcasters can leverage IT technology to create more reliable platforms for their applications. But it is important to understand the features available and how they work to be sure that you get the expected performance from your array.

---

*Original article attributed to Brad Gilmer, President* of Gilmer *& Associates, Executive D*irector of *the AAF Association and E*xecutive D*irector* of *the* Video *Services Forum.*