Eagle Eye Application Note – AN012

Understanding RAID and Building Resilience with Eagle Eye CMVRs

2022-06-07 Revision 1.2

Target Audience

This Application Note is intended for installers and technicians that will be performing CMVR and camera installations in conjunction with the Eagle Eye Cloud VMS and is intended to share best practices in regards to RAID.

Feature Overview

RAID stands for Redundant Array of Inexpensive Disks; it is a way of logically putting multiple disks together into a single array. The idea is that these disks working together will have the speed and/or reliability of a more expensive disk.

Background

A RAID system consists of two or more drives working in parallel. These can be standard hard discs, but there is also a trend to use the technology of Solid-State Drives (SSD). Sometimes, disks in a storage system are defined as **JBOD**, which stands for *Just a Bunch of Disks*. This means that those disks do not use a specific RAID level and simply act as a stand-alone disk.

Functionality

There are different RAID levels, each optimized for a specific situation. There is no standardization by any industry group or committee at present. Here are the following common RAID levels:

- RAID 0 – striping

- RAID 1 – mirroring
- RAID 5 – striping with parity
- RAID 6 – striping with double parity
- RAID 10 – combining RAID 1 mirroring and RAID 0 striping

In the Eagle Eye VMS, CMVR 520 and 620 both come with RAID 5, while CMVR 820 uses RAID 6.

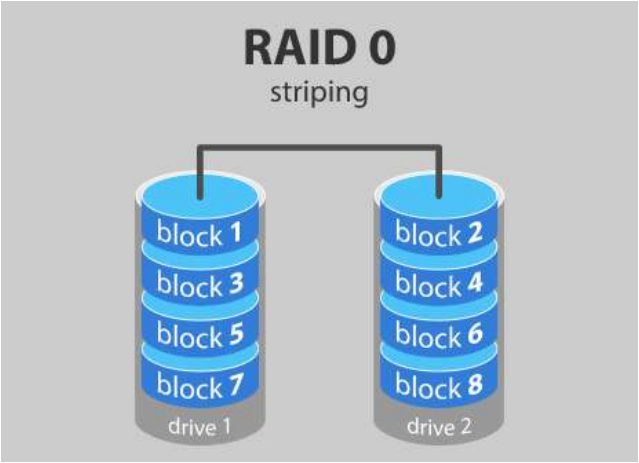	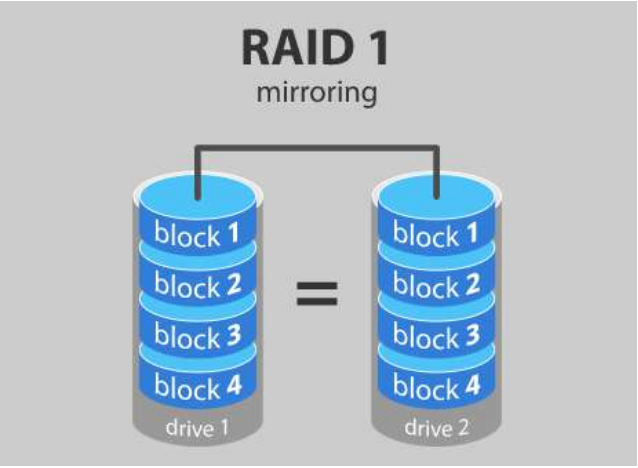
System data split up into blocks that get written across all the drives in the array. By using multiple disks (at least 2) at the same time, this offers superior Input/Output operations per second (IOPS) performance. **Advantages**: • Great performance both in read and write operations. • No overhead caused by parity controls and all storage capacity is used. • Easy to implement. **Disadvantages**: • It is not fault tolerant. • If one drive fails, all data in the RAID 0 array will be lost. • It should not be used for any mission-critical systems.	Data are stored twice by writing them to both data drives (or a set of data drives) and a mirror drive (or a set of drives). If a drive fails, the controller uses either the data drive or the mirror drive for data recovery and continuous operation. **Advantages:** • Excellent read and write-speed that is comparable to a single drive. • In the case where a drive fails, data does not get rebuilt, it is just copied to the replacement drive. • It is simple technology. **Disadvantages:** • The effective storage capacity is only half of the total drive capacity due to all the data being written twice. • A failed drive can only be replaced after powering down the computer that it's attached to.

Table 1a: Background and advantages and disadvantages of RAID 0 and RAID 1

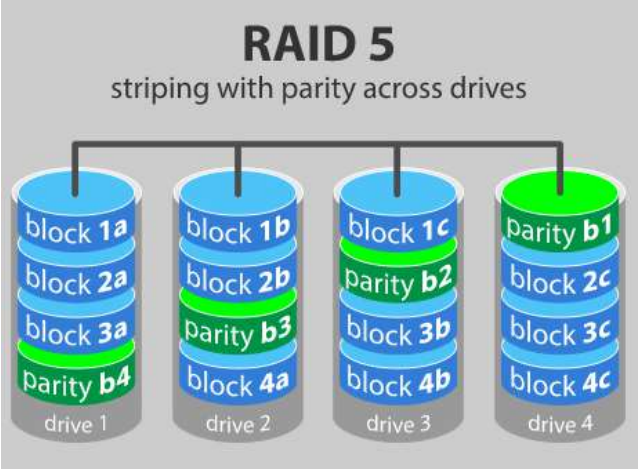	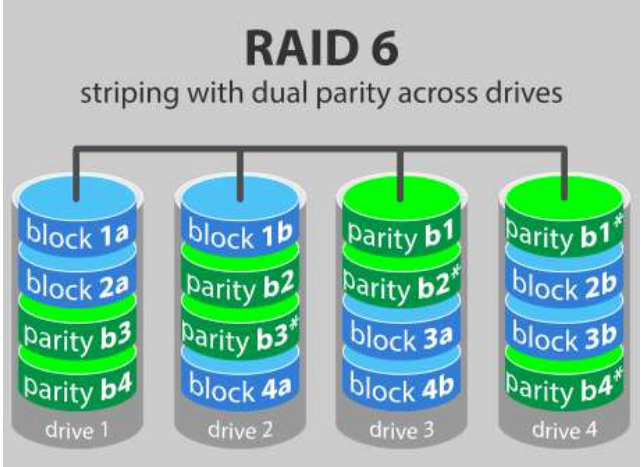
This is the most common secure RAID level. It requires at least 3 drives but can work with up to 16. Data blocks are striped across the drives and on one drive with a parity checksum. **Advantages**: • Read data transactions are very fast while write data transactions are somewhat slower due to parity that has to be calculated. • If a drive fails, data will still be accessible, even when the failed drive is being replaced and the storage controller rebuilds data on a new drive. **Disadvantages**: • Drive failures influence throughput, although this is still acceptable. • It is a complex technology. • In drives with higher capacity, if the disk fails and is replaced, restoring the data (rebuild time) may take longer depending on the load on the array and the speed of the controller. • In the event of a second drive failure, the data will be lost forever.	RAID 6 is like RAID 5, but the parity data are written to two drives. That means it requires at least 4 drives and can withstand 2 drives failing simultaneously. **Advantages**: • Like with RAID 5, read data transactions are very fast. • If two drives fail, data will still be accessible, even when the failed drive is being replaced. • More secure than RAID 5. **Disadvantages**: • Drive failures influence throughput, although this is still acceptable. • Write data transactions are 20% slower than RAID 5 due to the additional parity data that must be calculated. • It is a complex technology. • Rebuilding an array in which one drive failed can take a long time.

Table 1b: Background and advantages and disadvantages of RAID 5 and RAID 6

It is possible to combine the advantages (and disadvantages) of RAID 0 and RAID 1 in one system. This is a nested or hybrid RAID configuration. It provides security by mirroring all data on secondary drives while using striping across each set of drives to speed up data transfers.

Advantages:

- If something goes wrong with one of the disks in a RAID 10 configuration, the rebuild time is amazingly fast since all that is needed is copying all the data from the surviving mirror to a new drive. This can take as little as 30 minutes for drives of 1 TB.

Disadvantages:

- Half of the storage capacity goes to mirroring, so compared to large RAID 5 or RAID 6 arrays, this is an expensive way to implement redundancy

Table 1c: Background and advantages and disadvantages of combining RAID 1 and RAID 0

There are different RAID levels 2, 3, 4 and 7 but they are not commonly used due to the parity data always writing to the same drive.

Application

RAID 0 for non-critical data. In this scenario, carefully consider if the data truly matters. This may sound like an odd decision to consider but it is an especially important factor to determine usage of RAID 0.

RAID 1 for a two-disk array. It is suitable for small application usage where only two drives will be used. It requires high availability and cost is not typically a constraint.

RAID 5 is ideal for file and application servers such as data storage that has a limited number of data drives. It is a good all-around system that combines efficient storage with excellent security and decent performance.

RAID 6 is similar to RAID 5 in usage where more read than write performance. The difference is the recovery disk failures, the cost and slower write performance.

RAID 10 combines the performance benefits of RAID 0 with the redundancy benefits of RAID 1. RAID 10 is ideal in environments where capacity overhead is affordable. It offers the mix of performance and protection available from the traditional RAID.

Notes and Other Helpful Details

Features	RAID 0	RAID 1	RAID 5	RAID 6	RAID 10
Minimum # of drives	2	2	3	4	4
Data Protection	No Protection	Single-drive failure	Single-drive failure	Two-drive failure	Up to one disk failure in each sub-array
Read Performance	High	High	High	High	High
Write Performance	High	Medium	Low	Low	Medium
Capacity Utilization	100%	50%	67%-94%	50%-88%	50%