Disk Redundancy (RAID)

A Primer for Business

Dvana's Primers for Business series are a set of short papers or guides intended for business decision makers, who feel they are being bombarded with terms and want to understand a complex topic. This guide will break down the technology and explain the important points without jargon. The intention is to give you the knowledge needed to make business decisions, enabling you can use the right tool for the job.

We understand that not everyone can be an expert with the intricacies of modern computing and IT technologies, unless it is your passion of course!

This guide will explain the concept of a Redundant Array of Independent Disks (RAID). It covers the different levels, their advantages and disadvantages, along with a brief discussion on where to use each level.

Introduction

It is important to understand that the techniques presented in this report are intended to protect the data on the system from the failure of disk drives. They do not provide any backup of the data, nor do they facilitate archiving of the data. For both of those requirements, other techniques need to be employed, which are beyond the scope of this report. Please see other reports in the series for information on backups and archiving. Please note that all of these techniques do not protect against accidentally deleting data or protect against a virus affecting the data.

The different Redundant Array of Independent Disks (RAID) levels attempt to solve a particular problem, namely the inherent lack of reliability in hard disk drives. RAID combines two or more physical hard disks into a single logical unit. A set of disk drives configured to work together as a group is referred to as a “RAID Array” or more simply “Array”. There are two basic methods of achieving this; the first is software RAID, while the second is hardware RAID. There are advantages and disadvantages to both.

Software implementations are typically implemented in the operating system and usually present the RAID volume as a single drive to applications running within the operating system.

Hardware implementations are typically designed to present themselves to the attached system as a single hard drive, so that the operating system would be unaware that the drive was not a single disk drive.

There are three key basic concepts in RAID:

1. Mirroring is the writing of the same data to more than one disk
2. Striping is the splitting of the data across more than one disk
3. Error uses redundant information called "parity" which is stored on the disk drives to provide a mechanism to recover data should a drive failure occur. This is known as fault tolerance

In this report we will consider the standard levels and consider the different uses for the different options. The purpose of these techniques is to increase the reliability of data storage and to make the data more available. This protects the stored data from hardware failure, while often increasing the performance of the system.

As additional hard disks are added to a system, the probability that one of the disk drives will fail increases, this is the reason for redundancy. This redundancy however, comes at a cost, in that the redundant data must be stored somewhere. It is typically stored on additional hard disk drives.

The RAID techniques presented here are those most commonly used, there are other variants which have either been superseded or are no longer commonly utilized for one reason or another. Each of the different techniques has an impact on the performance and the reliability of the system, with their own quirks, advantages and disadvantages as we shall see in the following sections.

Striping or RAID 0

Description

Striping in its most basic form is used to improve the speed of data access, both reading and writing. It works by writing the data equally to all the disks at once. In a “Striped Set” there can be two or more disks all

working together. Since this method spreads the data across all disk drives, in the event of a disk drive failure, the whole striped set will be lost. The more disks are in the striped set, the higher the risk of data loss.

Configuration

- Absolute Minimum
  - Two data disks
- Practical
  - Two data disks
  - Three data disks
- Practical Maximum
  - Five data disks
  - Six data disks

Advantages

- Faster than a single drive
- Utilizes the full capacity of all drives
- Best price to performance ratio

Disadvantages

- Increased risk of data loss
- All disks must be the same size to eliminate wasted space

Ideal Uses

- Workstation
  - Video editing scratch disks
  - Photoshop scratch disks
  - Temporary users space
- Server
  - Temporary file storage
  - Backup snapshots prior to archiving
- General
  - Temporary storage for files or intermediate results

Striping is used where the performance is of benefit, but the significant disadvantage of 100% data loss in the event of a single disk drive failure is not an issue. This limits its use to niche applications, as most data storage situations require the data to be available in the event of a failure.

Mirroring or RAID 1

Description

The mirroring of data between two or more disk drives means that should a single drive fail, all the data is available on the other drive. The mirroring process writes the duplicated data to all the drives simultaneously. Writing of data to the disk drives is close to the same speed as for a single disk drive. However, all the drives must write the data before it is written and this can take a little longer than for a single disk drive. Data is read from all the disk drives at once, significantly increasing the read performance, since files can be read in parallel.

Configuration

- Absolute Minimum
  - Two data disks
- Practical
  - Two data disks
- Practical Maximum
  - Two data disks

Note: It is possible to use more than two disks in theory, but in practice this is rarely implemented in hardware, which is the performance option.

Advantages

- Multiple exact duplicates of the data
- Very fast and easy recovery of data
- Potentially faster than a single disk drive
- Read performance is increased
- Easy to upgrade, both for disks or the system

Disadvantages

- Uses twice the number of disk drives
- All disks must be the same size to eliminate wasted space
- A small write penalty might be incurred
- Most expensive method for the capacity

Ideal Uses

- Workstation
  - System drive
- Server
  - System drive
  - Database servers
- General
  - Any situation in which the data must not be lost, and read performance is required

Mirroring is the most basic of the redundant methods; it provides good performance and enables easy access to the data in the event of a failure. It is most ideally suited for use where performance is important, but data cannot be lost.

Redundant Striping or RAID 5

Description

Redundant striping of data is one of the most common forms of protecting data against loss. In this system, an additional disk drive is used to add capacity which is then used for storing “parity” or redundant data. The data is striped over all the disks, which gives the advantages of simple striping and the parity is rotated between all disks, this improves performance. All but one of the disk drives are required for the data to be recovered, so a single failure can be tolerated.

It is imperative that in the event of a single disk failure, the defective disk drive be replaced as soon a possible and the array allowed to rebuild itself. If a second disk drive fails, all the data on the array is lost. During a rebuild of the additional redundant data, the performance of the system will be degraded.

Configuration

- Absolute Minimum
  - Three disks, two for data and one for parity
- Practical
  - Any number of data disks and a single parity drive
- Practical Maximum
  - Fifteen data disks and one parity

Advantages

- Good compromise between capacity and performance
- Good compromise on cost
- Read performance is increased

Disadvantages

- Small writes are slower than a single disk
- All disks must be the same size to eliminate wasted space
- A rebuild can be very slow for large arrays

Ideal Uses

- Workstation
  - Not commonly needed or used in workstations
- Server
  - User shared data
  - General files
- General

- Any situation a large capacity is required
- Read performance is more important than write.

Redundant Striping is a good general purpose solution for reasonable performance and capacity which has a degree of protection against hardware failures.

It is possible to have a “Hot Spare” disk associated with the array; in this case, a disk drive is designated to be used to rebuild the redundant information automatically in the event of a disk drive failure. With the advent of very large disks, it would be generally better to have Multi-Redundant Striping (RAID 6) instead of Redundant Striping (RAID 5) with Hot Spare.

Multi-Redundant Striping or RAID 6

Multi-Redundant Striping is very similar to Redundant Striping (Raid 5), but differs in the use of two or more disks for redundancy. This means that for each additional parity disk drive, it is possible to loose that number of drives without loosing all the data.

This is particularly important when the number of disks in an array passes 15, since the extra protection of the additional parity drive or drives is needed.

Configuration

- Absolute Minimum
  - Four disks, two for data and two for parity
- Practical
  - Any number of data disks and two or more parity
- Practical Maximum
  - Any number of data disks and two or more parity

Advantages

- As for Redundant Striping (RAID 5)
- Ability to increase array size without practical limit
- Protection against a second disk drive failure

Disadvantages

- As for Redundant Striping (RAID 5)
- The use of a second parity drive affects write performance
- Additional cost as compared to Redundant Striping (RAID 5)
-

Ideal Uses

- Workstation
  - Not commonly needed or used in workstations
- Server

    - User shared data
    - General files
- General
    - Any situation where a large capacity is required
    - Read performance is more important than write

Multi-Redundant Striping (RAID 6) is suited to situations where more disks are needed for capacity than can be provided by Redundant Striping (RAID 5) or where the additional protection is needed.

Advanced

It is possible to combine the various techniques presented in this report to create a whole host of other combinations.

For example, it is possible to have two Redundant Stripes (RAID 5) and Mirror (RAID 1) them together so that all the data written to one is present on the other. In this case, it would be possible to lose all of one array and the data would still be present and safe. However, it would mean doubling the total equipment requirements. Utilising this arrangement, it is possible to have the two Redundant Stripe arrays at different locations and update in parallel both copies of the data, this can provide protection against the loss of data at a single site.

It is also possible to Stripe (RAID 0) data over Redundant Stripes (RAID 5) to gain the benefits of writing the data to different storage units, but at the same time eliminating the vulnerability inherent in Striping (RAID 0). This is of course at the cost of doubling the hardware requirements.

Everything applicable to Redundant Striping is also applicable to Multi-Redundant Striping.

Conclusion

There are various mechanisms to ensure that data can be protected against the loss of a single disk drive, with a little design thought, it is possible to make the system highly available.

By using redundancy, it is possible to sacrifice a little capacity to ensure that data will not be lost in the event of a hardware failure. As the need for storage grows and the reliability of the storage devices remains almost constant, the use of the techniques outlined here are of particular interest to businesses who are ever more reliant on their data being “Just There” when they need it.

The use of specialized hardware and storage devices can provide the performance in terms of speed and capacity a business needs. Scalability and upgradability should also be considered in provisioning such a system.

Spending a little more time and money from the outset will save large amounts of both in the future, when the systems need upgrading as your business need grow.

About Dvana

Dvana is a Management Consultancy, **Boosting Business Productivity** for all our clients.

This is accomplished by deeply understanding the current position of your business and its future direction.

Dvana run Professional Courses aimed at business professionals wanting or needing that extra edge over their competitors. The Professional Courses focus on getting results now, while building skills for the future.

Talk with Dvana today and see how we **Build Better Businesses**, so you become more successful and more profitable.

Contact Information

Web	www.dvana.com	
Phone	(01492) 55 63 84	UK
	+44 1492 55 63 84	International
Email	Info@dvana.com	

Terms of Use