


BEGINNER'S GUIDE TO STORAGE



Chetan Gowda

Associate Sales Engineer Analyst

Dell Technologies

Chetan_gowda@dell.com



The Dell Technologies Proven Professional Certification program validates a wide range of skills and competencies across multiple technologies and products.

From Associate, entry-level courses to Expert-level, experience-based exams, all professionals in or looking to begin a career in IT benefit from industry-leading training and certification paths from one of the world's most trusted technology partners.

Proven Professional certifications include:

- Cloud
- Converged/Hyperconverged Infrastructure
- Data Protection
- Data Science
- Networking
- Security
- Servers
- Storage
- Enterprise Architect

Courses are offered to meet different learning styles and schedules, including self-paced On Demand, remote-based Virtual Instructor-Led and in-person Classrooms.

Whether you are an experienced IT professional or just getting started, Dell Technologies Proven Professional certifications are designed to clearly signal proficiency to colleagues and employers.

[Learn more at www.dell.com/certification](http://www.dell.com/certification)

Table of Contents

Overview	4
Marketed capacity vs Real capacity	4
RAID.....	4
Hot Spare & Distributed Sparing.....	5
Actual Usable Capacity.....	5
Efficiency	6
Conclusion.....	7

Disclaimer: The views, processes or methodologies published in this article are those of the author. They do not necessarily reflect Dell Technologies' views, processes or methodologies.

Overview

This article explains the basics of storage capacity in a Storage Area Network (SAN) environment. It will help those who are non-technical as well as employees working as an analyst in I5 Grade understand the basic concepts of capacity in a SAN world, which will aid their design of technical storage solutions.

Marketed capacity vs Real capacity

When sizing a storage array, you need to carefully consider the capacity unit that you are sizing for; Drive manufacturers market the capacity in “Base-10” system of measurement, but the Operating system will see the capacity in “Base-2” system of measurement; that is the difference between TB (Base-10 system of measurement) vs TiB (Base-2 system of measurement). For example, if a customer requires 200 Terabyte of raw capacity for their environment and you may size a storage solution of 200 Terabyte keeping Base-10 system of measurement in mind. However, when the customer uses this capacity in their environment, they will probably see about 180 Terabyte of capacity. This is because the Operating system will see the capacity in Base-2 system, and the solution will be under-sized. Hence, it is very important to consider the unity of measurement when sizing a storage solution.

RAID

Next thing to take into consideration when sizing a storage solution is the level of protection and performance that the customer is expecting. Redundant Array of Independent Disks (RAID) is the technology that is going to offer protection against the drive failure. RAID is going to combine multiple drives in a set, and it will be called as RAID Set (In Traditional Storage Pool). Data will be stored across multiple drives in RAID, while adding protection to the data in the form of redundancy; Except RAID 0, all other RAID level offer protection in form of redundancy. Since data is distributed across multiple drives, RAID can improve performance by serving the data using multiple drives simultaneously.

There are three RAID techniques which are going to define the RAID Levels:

- a) Striping
- b) Mirroring
- c) Parity

Striping: Striping is the technique where the data is split across multiple drives; for example, if we have 4 drives then data ‘A’ can be split into, $A = A1 + A2 + A3 + A4$; and when the data ‘A’ is retrieved, then all the four drives will work simultaneously to rebuild the data ‘A’. This parallel operation will improve the performance.

Mirroring: Mirroring is the Technique where the same data is copied on different drives. For example, if there are two drives, then data ‘A’ will be copied on both the drives. Even if one of the drives fail, we have the remaining drive which will serve the data. Mirroring does incur IO overhead, since the same data must be written on different drive

Parity: Parity is the technique which is going to protect the Striped data. Depending on the level of protection, most commonly we use single parity or dual parity. For example, If we are striping data ‘A’ into five drives using single parity protection then, $A = A1 + A2 + A3 + A4 + AP$; AP is the parity data, which will be used to rebuild data of any failed drive. In case of dual Parity, data ‘A’ will be stripped

across six drives, $A = A1 + A2 + A3 + A4 + AP1 + AP2$. Single parity will protect against one drive failure, and dual parity will protect against two drive failure in the same set (Traditional Storage Pool)

Most commonly used RAID levels are:

- RAID 0 – Striping but No-Fault Tolerance
- RAID 1 – Mirroring
- RAID 10 – Striping and Mirroring
- RAID 5 – Striped and distributed parity
- RAID 6 – Striped and dual distributed parity

RAID incurs capacity overhead. For example, if there is 100 Terabyte of capacity and we use RAID 1, which is Mirroring, 50 Terabytes will be lost to redundancy. RAID also incurs performance penalty during writes. For example, if using RAID 1, the same data needs to be written to different drives. The following table shows the capacity overhead and penalty incurred on writes.

RAID Level	Available Storage Capacity (%)	Write Penalty
RAID 1	50	2
RAID 5	$[(n-1)/n]*100$	4
RAID 6	$[(n-2)/n]*100$	6

Here n is the number of drives. Keep the capacity and performance overhead into consideration while sizing. Consider the Cost and latency requirement of the customer to arrive at the right capacity.

Hot Spare & Distributed Sparing

Hot spare is a drive, whose purpose is to replace the failed drive in the array. If Mirrored RAID is used, then we rebuild the data from the surviving drive on to the Spare drive. The number of spare drives will depend on the number of drives configured and the type of drive configured. Check the design of the array you are sizing to get the correct number of Hot Spares.

Recently, more arrays have Distributed sparing, where instead of having an idle drive as a hot spare, certain capacity is reserved in all the drives, this capacity is equal to the hot spare capacity; when a drive fails, the data is rebuilt onto the reserved capacity. Distributed sparing improves the performance, since there are no drives which is sitting idle and during the rebuild, all the drives take part in the operation, hence the rebuild time is faster.

Account for the capacity overhead caused by this spare capacity.

Actual Usable Capacity

Let's look at how the raw capacity translates to usable capacity considering all the factors mentioned in this article thus far. Consider a Storage array with the following characteristics:

- 50 Drives of 3.84 Terabyte
- Total Raw capacity at the beginning: 192 Terabyte
- 2 Hot Spares
- RAID 5 (4+1)

- 20% Vendor hold back (all Storage vendors reserve 15 to 20% of capacity for performance and metadata)
- Total Array cost with 192 Terabyte raw capacity

	Remaining Capacity
Starting Capacity	192 Terabyte
Capacity after hot spares	184 Terabyte
Capacity that OS will see (Base-2)	167 Terabyte
Capacity after RAID 5(4+1) overhead	133 Terabyte
Capacity after 20% Vendor hold back	106 Terabyte

	Array Cost/Capacity	\$/Terabyte
Starting Cost	\$800,000 / 192 Terabyte	\$4,166
Cost after hot spare	\$800,000 / 184 Terabyte	\$4,347
Cost after OS (Base-2)	\$800,000 / 167 Terabyte	\$4,790
Cost after RAID 5(4+1)	\$800,000 / 133 Terabyte	\$6,015
Cost after 20% Vendor Hold back	\$800,000 / 106 Terabyte	\$7,547

For this example, we considered the total cost of the array including the support and services as \$800,000. What started as \$4,166 / Terabyte with 192 Terabyte of raw capacity, ended with \$7,547 / Terabyte with 106 Terabyte of usable capacity.

Efficiency

To drive down the \$ / Terabyte cost, vendors started to implement a technology called Data Reduction Ratio (DRR); expressed as X:1, DRR typically X varies between 2 to 4 for most common workloads; X can even go up to 65 for backups. DRR is a capacity optimization technique. After applying DRR to usable capacity, we will get effective capacity. Usable Capacity + DRR = Effective Capacity.

DRR = Data deduplication + Compression.

Data Deduplication: This technique avoids storing the duplicate blocks of data in the storage array. It uses pattern recognition software to identify the duplicate data.

Compression: Compression algorithms reduces the amount of data that needs to be stored onto the physical drive.

For example, if we say that the array has 300 Terabyte effective capacity at 3:1 DRR, this means that the array can store 300 Terabyte of capacity onto 100 Terabyte of physical usable capacity. DRR is going to help drive down the \$ / Terabyte cost.

Conclusion

The journey from Raw capacity to Usable capacity is important to keep in mind whenever discussing capacity with the customer or sizing a solution, and the impact it has on \$ / Terabyte. Have a clear understanding of the workload that customers want to implement, and if that workload qualifies to achieve the DRR that you used to size the storage array.

Dell Technologies believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." DELL TECHNOLOGIES MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying and distribution of any Dell Technologies software described in this publication requires an applicable software license.

Copyright © 2021 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.