



Understanding Usable versus Raw Storage Capacity



Terascale, Inc.
145 Bodwell St
Avon, MA
www.terascale.com

Summary

When planning for deployment of a storage solution, understanding the difference between raw and usable storage capacity is a key consideration. The Terascale RTS 1000 offers excellent capacity delivering over 85% of raw capacity to end users as usable storage space, making it the leader in this area among Lustre based solutions.

The Problem

There are two key capacity metrics that need to be considered when specifying storage systems. The first metric is “raw disk capacity” which refers to the amount of raw space available on a particular storage system. The second metric is “usable capacity” which refers to the amount of space available for end user data after the system has been fully setup.

Raw capacity is typically calculated by adding together the capacity of the individual disk drives embedded in the system. These capacities are specified by the disk drive manufacturers and come in standardized sizes, such as 250 GB, 500 GB, or 1 TB. Typically storage systems are specified as having “40 TB of capacity”. This usually refers to the raw capacity available.

(Note: The technical definition of a gigabyte is 2 raised to the 30th power, or 1,073,741,824 bytes. However, gigabyte has been commonly rounded to refer to 1 billion bytes. Most disk drives manufactured now specify their drives with gigabyte being equal to 1 billion bytes.)

Usable capacity can be more complicated to understand and specify. In a storage solution, there are components that consume disk capacity. Below are some of the major ones.

The operating / file system that runs the storage solution is one consumer of capacity. Depending on where the files system resides, this will consume disk capacity that will no longer be available for user data.

Journaling is a feature that consumes disk capacity. It helps protect the file system against unexpected power outages or hardware component failures. With journaling enabled, a continuous record of changes to files on the disk is maintained in a journal file. If the system stops because of a power failure or some other issue, the journal is used to restore the disk to a known-good state when the system restarts.

Parity is the technique of checking whether data has been lost or written over when it's moved from one place in storage to another. The parity data is stored on the disks and reduces usable storage.

Formatting loss also impacts usable capacity. When the raw disks are formatted, a portion of the disk ends up as unusable from a user data storage perspective.

Impact on Lustre based environments

For solutions leveraging the Lustre parallel file system, the most important capacity metric is the amount of usable storage space per Object Store Target (OST). Every Lustre system will have a Metadata server solution which will have its own storage capacity, as well as

numerous OSTs with their own raw capacity. It is the OST usable capacity that defines the true usable storage capacity for a Lustre solution.

Below is a table that highlights the Terascale usable capacity per OST compared with two other alternatives. As this table shows, how disk space is allocated to the various aspects of a full solution can have a large impact on a capacity available for data storage.

One element that sets Terascale apart is its unique journaling subsystem that enhances overall system performance and delivers more usable capacity.

Usable Capacity				
	Terascale	Alt. 1	Alt. 2	
Total Gross Space (GB)	8000	12000	8000	
Per Disk Capacity (GB)	1000	1000	1000	
Disks Per OST	8	12	8	
Less Journaling (GB)	0	2000	2000	
Less Parity (GB)	1000	1000	500	
Less Parity 2 (GB)	0	1000	0	
Less Hot Spare (GB)	0	0	1000	
Less OS & Lustre (GB)	24	0	0	
Less Formatting Loss (GB)	54	54	54	
Usable Space Per OST (GB)	6922	7946	4446	
Disk Utilization for Objects	86.5%	66.2%	55.6%	

Conclusion

Terascale unique approach in its Lustre implementation combined with its specialized journaling system leaves over 85% of system raw capacity for use of user data.

About Terascale

Terascale develops high throughput, high capacity and cost-effective storage solutions. Its unique storage appliance approach is changing the dynamics of the performance driven computing market, enabling existing users to do more for less while enabling new users to maximize the capabilities of their processing infrastructure. Founded in 2005, Terascale is based in Avon, MA.



Terascale, Inc.
 145 Bodwell St.
 Avon, MA 02322
 Tel: 508-588-1501
 Email: sales@terascalacom
www.terascalacom