

# Storage Capacity 101 to Post-Grad in a Flash

Clive Longbottom, Service Director



## Take a nibble at a bite

OK – we all know that data is made up from bits of information – 8 bits to a byte (4 to a nibble, for those who care). Kilobytes, Megabytes, Gigabytes, Terabytes and Petabytes all used to be based on the ‘power of 2’ system, so 1kB was 1,024 bytes (2 to the power 10), but has now come in line with standard metric notation, and is now 1,000 bytes.

This immediately has an impact on storage for anyone who was brought up on the power of 2 notation. As 1MB would be 1024 bytes squared (1,048,576 bytes), 1GB 1024 cubed, and 1TB 1024 to the power 4, we start to see how the quoted storage of a disk is less than many of us expect. It may only lead to a 1% error at the kilobyte level, but it’s getting close to 10% at the terabyte level; that is capacity that many think they are getting that they are not. To this end, the IEC has renamed old-style units as kibibytes (KiB), mebibytes (MiB) and gibibytes (GiB) (I kid thee not).

## The terror of non-standard terabytes.

So, let’s take that 1TB capacity and see how big it really is. This capacity, as provided by a manufacturer, has a certain amount of space taken up by the vendor’s own systems – only a few percent, but it still takes some more of that raw storage away from you. Formatting so that it can be recognised by an operating system takes up space. This depends on the operating system, sure – but you are still not going to get the raw storage space you were hoping for. Configuring the raw capacity into RAID groups up will take more space – for example a RAID 5 configuration across 5 disks will lower actual storage by 20% - at least, as the formatting of RAID tends to take up more space than a simple formatting, and the data indexing for RAID itself takes up yet more space.

We are now getting to the point where we should be able to see that there is a major difference between ‘raw’ and ‘usable’ storage. That 1TB has come down quite a bit – but we’re not finished yet.

How have you formatted that capacity and at what level? Have you used basic data compression? How about in-line deduplication?

All of these will have an impact on data storage capacity achieved; some in a negative and some in a positive way. After all this formatting is complete, now you have the ‘usable’ capacity of the storage. But alas, with all-flash storage you do get to actually use much more of your available capacity than with a disk-based solution.

*The move from a kilobyte being 1024 bytes to 1000 bytes means that 1 terabyte is nearly 10% less than it used to be.*

# Storage Capacity 101 to Post-Grad in a Flash

Clive Longbottom, Service Director



## The death of the '60% rule'

*“As [flash] is a direct access system, the standard, received wisdom approach of ordering more capacity when storage hits 60% capacity no longer applies, as you are not using magnetic disks and thus not subject to their limitations. You can stuff a flash-based system as full as you want – and performance will still be excellent.”*

With flash-based arrays, there is no need to defragment the 'disks' – as direct access systems, the data is pulled from the platform at full speed – every time. As it is a direct access system, the standard, received wisdom approach of ordering more capacity when storage hits 60% capacity no longer applies, as you are not using magnetic disks and thus not subject to their limitations. You can stuff a flash-based system as full as you want – and performance will still be excellent.

There is also a geeky aspect around sector size. When floppy disks came about as the first spinning storage devices, a 512 byte sector size was chosen. This is the smallest possible data size that can be written to a disk – but it also has its own payload of various bits of extra data required by the disk controller. By upping that sector size to 4 kbytes, the smallest amount of data that can be written is upped by a factor of 8 – while that non-productive payload overhead remains the same. This leads to an immediate increase in usable storage of 12.5%.

Now, compression. Most data contains a lot of 'space'. By looking at the data as a string of bits, patterns can be recognised and the data remapped. A bit string such as 00001011111110001111 can be compressed to say that it is four zeroes followed by a one, then a zero, then eight ones, three zeroes and four ones (see figure 1). In this way, more data can be written to the available storage. However, again, it takes up resources – and for magnetic disks, this means a noticeable impact on performance. However, for flash-based arrays, that performance impact is negligible.

# Storage Capacity 101 to Post-Grad in a Flash

Clive Longbottom, Service Director





Figure 1

# Storage Capacity 101 to Post-Grad in a Flash

Clive Longbottom, Service Director



## Deduplication deduplication deduplication...

The same with data deduplication. At a high level, this can be done by identifying what files are the same across a storage platform – and many organisations have been pretty amazed to find out how much storage they are paying for to store the same files over and over again. This is too simple, though. Along with compression, looking at data at the block level can lead to far greater understanding of what bits and bytes are being stored redundantly. By creating an index of specific patterns, data can be massively compressed with the data being ‘rehydrated’ as it is retrieved back to the original data format.

This is heavily resource dependent – and most spinning disk array vendors have stuck to using it only for target-based systems – backup and archive. However, with flash-based arrays, the raw power in the systems means that compression and advanced inline data dedupe can be used for most any primary storage workload.

This is the land of confusion

So, are you sufficiently confused at this point? Let’s sum up what we’ve learned by considering the three common ways storage vendors express capacity:

- **RAW** – is the actual physical capacity purchased, i.e. what’s shipped in the box
- **USABLE** – is the available capacity after RAID, formatting, etc., i.e. what you can actually use
- **EFFECTIVE** – is how much data you can effectively store once data reduction (deduplication and compression) are brought into play, i.e. the equivalent amount of usable capacity without data reduction

But, which measure is really important? That depends upon your workload(s) and if you are attempting an apples-to-apples comparison across different solutions.

- **RAW** is only important as a measure of physical capacity; it is rarely helpful in real life.
- **USABLE** is important for performance-oriented workloads. The overhead of data reduction affects latency and IOPS, and can have other detrimental impacts on certain applications such as databases.
- **EFFECTIVE** is important for primary storage where a balance between performance and cost effectiveness is desired.

# Storage Capacity 101 to Post-Grad in a Flash

Clive Longbottom, Service Director

Note that USABLE and EFFECTIVE are not the same measure. Some sneaky vendors might try to convince you to the contrary, especially if their data reduction capabilities cannot be turned off. Performance sensitive workloads can be negatively affected by data reduction, and their needs should be addressed through true USABLE capacity.

So, what we have are several ways to describe capacity, each relevant within a given context, but subject to misuse that can result in the expressed figures becoming a complete red herring that cannot be trusted in any way. What really counts is how much data can you store on any given storage platform based upon your specific workloads. Due to the way that flash-based systems operate, they can store data far more effectively and efficiently than magnetic-based systems.

And this is why flash-based storage can beat magnetic-based systems in a per real byte storage competition.



violin<sup>®</sup>  
MEMORY

# Storage Capacity 101 to Post-Grad in a Flash

Clive Longbottom, Service Director



## About Violin Memory

Violin Memory transforms the speed of business with continuous data protection through high performance, always available, low cost management of critical business information and applications.

Violin's All-Flash optimized solutions accelerate breakthrough CAPEX and OPEX savings for building the next generation data centre. Violin's Flash Fabric Architecture (FFA) speeds data delivery with chip-to-chassis performance optimization that achieves lower consistent latency and cost per transaction for Cloud, Enterprise and Virtualized mission-critical applications. Violin's All-Flash Arrays empowered by our enterprise data management software solutions enhance agility and mobility while revolutionizing data centre economics.

Founded in 2005, Violin Memory is headquartered in Santa Clara, California.

Further details are available at <http://www.violin-memory.com>



# Storage Capacity 101 to Post-Grad in a Flash

Clive Longbottom, Service Director



violin  
MEMORY

## About Quocirca

Quocirca is a primary research and analysis company specialising in the business impact of information technology and communications (ITC). With world-wide, native language reach, Quocirca provides in-depth insights into the views of buyers and influencers in large, mid-sized and small organisations. Its analyst team is made up of real-world practitioners with first-hand experience of ITC delivery who continuously research and track the industry and its real usage in the markets.

Through researching perceptions, Quocirca uncovers the real hurdles to technology adoption – the personal and political aspects of an organisation's environment and the pressures of the need for demonstrable business value in any implementation. This capability to uncover and report back on the end-user perceptions in the market enables Quocirca to advise on the realities of technology adoption, not the promises.

Quocirca research is always pragmatic, business orientated and conducted in the context of the bigger picture. ITC has the ability to transform businesses and the processes that drive them, but often fails to do so. Quocirca's mission is to help organisations improve their success rate in process enablement through better levels of understanding and the adoption of the correct technologies at the correct time.

Quocirca has a pro-active primary research programme, regularly surveying users, purchasers and resellers of ITC products and services on emerging, evolving and maturing technologies. Over time, Quocirca has built a picture of long term investment trends, providing invaluable information for the whole of the ITC community.

Quocirca works with global and local providers of ITC products and services to help them deliver on the promise that ITC holds for business. Quocirca's clients include Oracle, IBM, CA, O2, T-Mobile, HP, Xerox, Ricoh and Symantec, along with other large and medium sized vendors, service providers and more specialist firms.

Full access to all of Quocirca's public output (reports, articles, presentations, blogs and videos) can be made at <http://www.quocirca.com>