

# Object Storage 101

Standard file and block storage have their places, but there is an increasing move to the use of object storage. This is not just for technical reasons – object storage meets the governance, risk and compliance needs of organisations, while also providing enhanced information management capabilities.

## Key Points:

Maintaining the integrity and security of information assets is a big issue for organisations today. Large investments are made in layering security and information management systems onto storage systems, ignoring the fact that in the majority of cases, it is the storage system itself that is at fault.

The use of object storage provides a built-in mechanism for creating a platform for information management and security. As data assets are stored as full objects, rather than being broken down into blocks of relatively easily changeable ones and zeros, object storage provides a means of monitoring and auditing all actions taken against a data asset in real time.

Through the use of erasure coding and a global namespace, object storage enables massive scalability across a dispersed storage platform, providing high levels of data availability in a secure and easily manageable environment.

## Report Authors

*Clive Longbottom*

*Tel: +44 118 948 3360*

*Email: [Clive.Longbottom@Quocirca.com](mailto:Clive.Longbottom@Quocirca.com)*

*Marcus Austin*

*Tel: +44 7973 511 045*

*Email: [Marcus.Austin@Quocirca.com](mailto:Marcus.Austin@Quocirca.com)*



## Commissioned by:



## Block and File Storage

**W**e all know what “data” is. It is the digital ones and zeros that make up the zettabytes of stored information around the world. However, the way that this data has historically been stored has a few problems.

Let’s look further into this.

Figure 1 shows a standard, spinning magnetic platter disk drive. It’s a complex, highly engineered item, based on technology that is over 50 years old.

The important part is how data is written to this disk – each platter has a set of tracks that are broken down into blocks which are further broken down into sectors. All those ones and zeros need to be placed onto this storage device by filling up those sectors.




Figure 1: A disk drive

Unfortunately, even though we are now well into the move towards solid state disk drives, the concepts and constraints of spinning magnetic disks are following through – causing the same problems for those working with all flash arrays (AFAs). This need not be the case – flash storage substrates can deal with data in a completely different way to how spinning disks do – and this could lead to a much higher adoption of object storage in the future.

If we examine a data file, it will typically be broken up into lots of different chunks of data and stored across multiple sectors, blocks, tracks and platters.

When you want that file back, all the disparate chunks of data have to be recovered from the various parts of the disk, reconstituted as a single file and delivered to the requesting application or user as required.

Now, not only can this result in high data latency, but it also has some fundamental security issues.

The way that the data has been written is governed by a metadata storage set – a common example of this is logical block addressing (LBA).

LBA is a simple means of logging the tuple of cylinder(platter)/head/sector (known as a CHS address) so that data can be easily identified and retrieved.

However, such a mechanism means that should the underlying data be modified without the LBA knowing about it, it will still be recovered as if it was ‘true’ data. With a simple hex editor, it is possible to rewrite the base data on a block mode disk drive – and so security within such a storage system can be relatively easy to compromise.

One attempt to bypass this problem has been the use of file storage systems. Although this sounds as if it is dealing with data stored as complete items (i.e. files), it is only an abstraction. The two common forms of file storage are network file system (NFS) and server message block (SMB – previously known as CIFS). These were brought about by the need for users to access files that had been stored on network attached storage (NAS) devices.

File storage works by maintaining a hierarchical view of how files are managed. When a remote user makes a file request to a NAS system, it goes through a user interface that abstracts the physical storage from the logical hierarchy. As such, it still suffers from the same problems as block storage – the underlying data can be changed without the overall system being aware of what has happened.

---

*“Object storage is therefore an excellent means of dealing with information storage and management. It has all the capabilities required to create a fully functional platform that understands and deals with different object types and ensures that they are highly available and inherently audited and logged as to any actions taken against the object.”*

---

Much time and effort has gone into attempting to ameliorate this security problem – application and database security has been bolstered through the use of process auditing tools and so on. However, the very fact that block and file storage has this basic flaw should be an area of concern to all organisations.

Object storage has been around for a long time – but has failed to make the breakthrough into widespread usage due to historical technical issues including performance.

As these issues are now dealt with, it may well be time for an organisation to review its storage strategy and consider object storage for multiple different storage workloads.

## Object Storage

**N**ow, what happens if we regard data in the same way as your average, non-technical person does? There is a file – it is a single entity; it has a start, a middle

**and an end. It is globally identifiable as what it is – a single entity of information. We can regard it as an object – something that can be pointed at and described.**

This makes for a different approach to how the underlying data is stored. It now has to be stored as a single entity – and it has to have a series of attributes stored with it to make it identifiable and addressable.

Object storage is built on this concept. Every item is an object; those objects are stored as immutable items on a storage medium. If any attempt is made to change that data, it will be noted and stored as metadata – every action is automatically audited within the system.

Object storage holds a lot of metadata on each object. Some of this metadata will be automatically created as the object is stored; some will be created based on rules defined by the business.

A major part of this automatically created metadata is the globally unique identifier for each object. By holding such a globally unique identifier, this builds in much greater security – and enables greater levels of data availability. Object storage utilises the concept of a namespace: within that namespace, each object uses this unique identifier to ensure that it is always seen as being a different entity – even where the content may be 100% the same. Therefore, the namespace can operate across a globally dispersed set of object stores with the objects distributed accordingly while maintaining each of their globally unique identities.

Rather than storing multiple copies of a file for data protection, objects are broken down into fragments known as ‘shards’ that contain redundant data which are then distributed across the object store. For example, in an 18/5 Erasure Code implementation, the object is broken down into 18 shards, with the shards being stored across different disk drives across multiple storage arrays.

On any disruption to the data, any 13 of these shards are required to rebuild the object (see figure 2). Should a given shard (or even five) become corrupted, the remaining shards are sufficient to rebuild the object: availability is built in without the need for standard RAID.

The use of metadata also enables extra functionality to be built into a storage system. For example, the security level of an object (e.g. public, confidential, for your eyes only) can be stored as an artefact of the document.

As this is tied into the object itself through the global namespace, it can be used in many other areas. It can be indexed, so that those in charge of an organisation's security can search for all objects that are tagged as being 'public'.

The metadata can be used in business processes – for example, when an email is being sent, if the attachment has a security identity of 'for your eyes only', raise an exception via a message to the sender or place the email message into quarantine until it has been checked.

Object storage is also media agnostic: it can be used against spinning magnetic disk drives as well as solid state flash-based systems. It is therefore not tied in to any particular approach – decisions made today on adopting object-based storage will remain valid for the future.

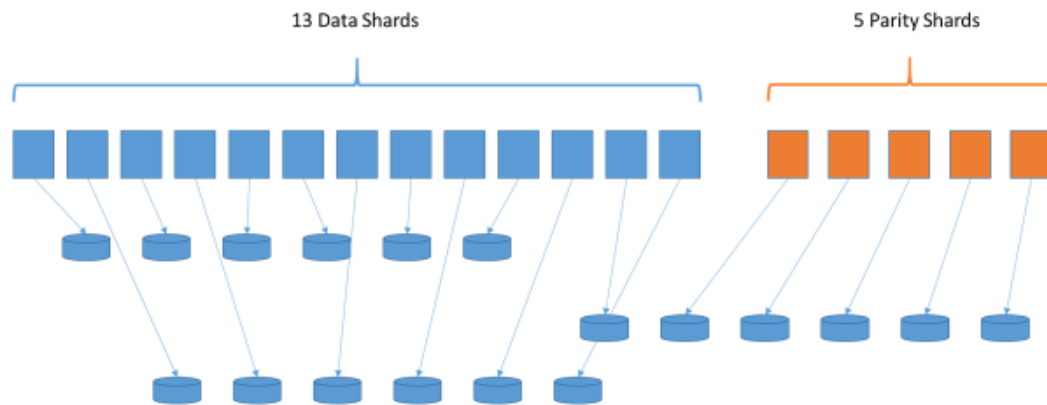


Figure 2

Many people remain worried that object storage is 'the new kid on the block' – it has, however, been around for 20 years now, and underpins cloud services such as AWS, Microsoft Azure, Google and many OpenStack implementations (using Swift). In addition, some vendors are selling object storage software for DIY solutions with others offering purpose-built integrated object storage solutions.

---

*"Rather than storing multiple copies of a file for data protection, objects are broken down into fragments known as 'shards' that contain redundant data which are then distributed across the object store. For example, in an 18/5 Erasure Code implementation, the object is broken down into 18 shards, from which any 13 are required to rebuild the object (see figure 2). Should a given shard (or even five) become corrupted, the remaining shards are sufficient to rebuild the object: availability is built in without the need for standard RAID."*

---

Alongside this, object storage underpins the majority of the world's supercomputers, due to its scalability and proven performance capabilities.

The majority of the Global 2000 organisations are also using object storage – either intentionally or through taking a storage system that just happens to use object storage.

What object storage has shown is its incredible scalability: the large cloud providers are now looking at hundreds of trillions of objects being stored in their object stores around the globe – all as if on a single logical storage platform.

Object storage is therefore an excellent means of dealing with information storage and management. It has all the capabilities required to create a fully functional platform that

understands and deals with different object types and ensures that they are highly available and inherently audited and logged as to any actions taken against the object.

---

*“Consider within your organisation – just what can be considered as objects, rather than just data? Files are the obvious one – but how about virtual desktop infrastructure (VDI) images? How about Docker and Mezos containers? Inventory items in the ERP system; customers in the CRM system? Transactions?”*

---

Look within your organisation – just what can be considered as objects, rather than just data? Files are the obvious one – but how about virtual desktop infrastructure (VDI) images? How about Docker and Mezos containers? Inventory items in the ERP system; customers in the CRM system? Transactions?

Indeed, pretty much everything can be regarded as an object – and using object storage along with business-defined metadata enables so much more value to be extracted from the underlying data.

It is clear that object storage is a proven and useful means of storing information – maybe it is time for you to look at it far more closely?

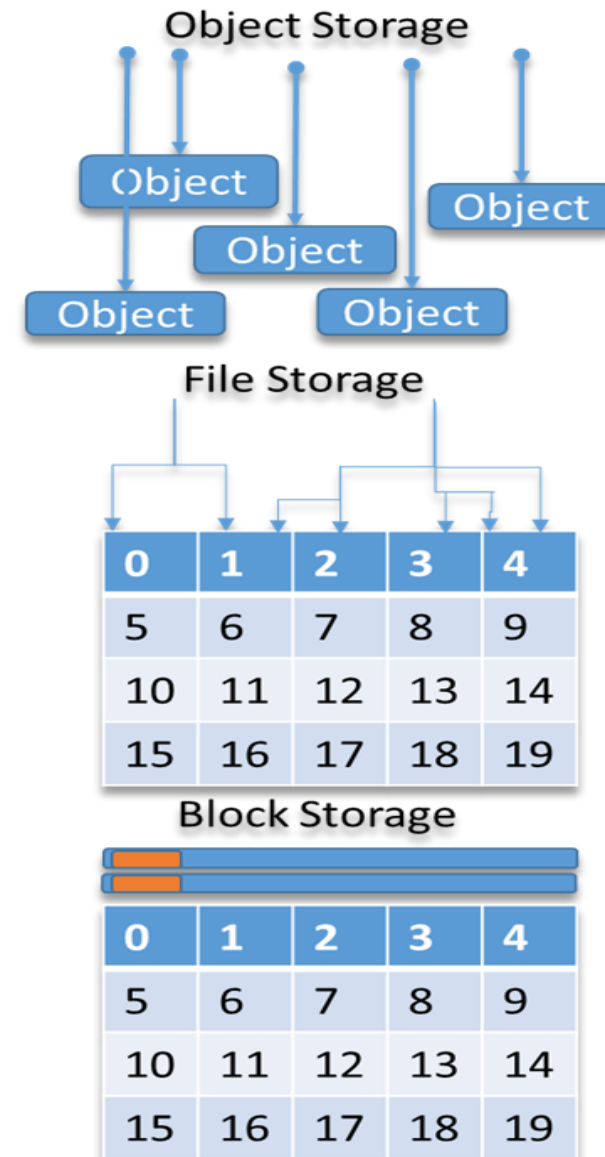


Figure 3



### About Western Digital

Western Digital Corporation (NASDAQ: WDC) is an industry-leading provider of storage technologies and solutions that enable people to create, leverage, experience and preserve data. The company addresses ever-changing market needs by providing a full portfolio of compelling, high-quality storage solutions with customer-focused innovation, high efficiency, flexibility and speed. Our products are marketed under the HGST, SanDisk and WD brands to OEMs, distributors, resellers, cloud infrastructure providers and consumers.

### HGST

HGST (@HGSTStorage), a Western Digital Corporation (NASDAQ: WDC) brand, helps the world harness the power of data. Our smarter storage solutions power the markets and companies that shape our lives—enabling possibilities for the cloud, enterprise and sophisticated infrastructures everywhere. For more information, please visit [www.hgst.com](http://www.hgst.com) #LongLiveData

### About Quocirca

Quocirca is a research and analysis company with a primary focus on the European market. Quocirca produces free to market content aimed at IT decision makers and those that influence them in business of all sizes and public sector organisations. Much of the content Quocirca produces is based on its own primary research. For this primary research, Quocirca has native language telephone interviewing capabilities across Europe and is also able to cover North America and the Asia Pacific region. Research is conducted one-to-one with individuals in target job roles to ensure the right questions are being asked of the right people. Comparative results are reported by geography, industry, size of business, job role and other parameters as required. The research is sponsored by a broad spectrum of IT vendors, service providers and channel organisations. However, all Quocirca content is written from an independent standpoint and addresses the issues with regard to the use of IT within the context of an organisation, rather than specific products. Therefore, Quocirca's advice is free from vendor bias and is based purely on the insight gained through research, combined with the broad knowledge and analytical capabilities of Quocirca's analysts who focus on the "big picture". Quocirca is widely regarded as one of the most influential analyst companies in Europe. Through its close relationships with the media, Quocirca articles and reports reach millions of influencers and decision makers. Quocirca reports are made available through many [media](#) partners.

To see more about Quocirca's analysts, click [here](#)

To see a list of some of Quocirca's customers, click [here](#)

To contact Quocirca, please click [here](#).