



OBJECT STORAGE AS A PLATFORM FOR EXPONENTIAL DATA GROWTH





DROWNING IN DATA

In 2014, EMC and IDC predicted that 90% of all data generated in the next decade would be unstructured*; a prediction that is quickly becoming a reality.

Whether it's the Internet of Things (IoT), the Connected Everything, the use of 8K video, the digitisation of assets or the proliferation of mobile devices – the rate at which unstructured data is growing is phenomenal. Not only is the volume of data being generated vast, but it also has intrinsic value that organisations are trying to harness.

In addition to the rapid data growth, regulated industries such as legal and financial services or project/customer driven industries such as Engineering, Manufacturing or Advertising, demand

that data is retained for long periods of time, in some cases forever.

Collaboration both internally within the business, and with external third parties is also key, and in some industry sectors such as Architecture, Engineering and Construction (AEC) it's almost mandatory.

Historically, when organisations have wanted to collaborate, it has led to the storing of multiple silos of data, or the use of unsanctioned file sharing applications.

What does this all mean? It means that the amount of data being generated is rapidly growing, it needs to be stored for longer and shared repeatedly. Traditional storage methods struggle to cope with these demands, particularly at large scale

How Does This Change Today's Storage Requirements?

Before we consider how these demands are changing the requirements for enterprise storage systems, first we need to understand the challenges that traditional SAN and NAS platforms can't overcome.

- 1. Physical Restrictions:** traditional platforms are built on proprietary hardware with finite capacity based on the number of disks and shelves that they contain; limiting their ability to scale and forcing costly forklift upgrades
- 2. Layers of Complexity:** to provide an end-to-end solution for data storage there are typically a number of layers e.g. primary storage, D2D/D2T and off site tape storage. In a distributed enterprise the cost and complexity of scaling this are nearly impossible to contain
- 3. Spiraling Cost:** due to it's layers of complexity, containing the cost of these systems is nearly impossible as you attempt to scale to meet data growth
- 4. Cold Data Is Cold:** many organisations have a choice: either store data on a cold tier (archive, often tape based) with slow and often inconsistent restore times, or leave archive data on expensive SAN/NAS storage. This is often a cost based decision but todays business needs to have it's data more readily available

INEFFICIENCY IN NUMBERS



Data Generated In The Next Decade Will Be Unstructured



The DU Will Contain Valuable Info If Analysed by 2020



Every TB Of Storage Needed, Firms Buy 5TB Of Capacity

It's these challenges, coupled with sheer volume of data being stored and used, that is causing organisations to consider new ways of storing their unstructured data. The overall goal being to:

- Reduce the cost and complexity of their storage infrastructure
- Enable internal and external collaboration
- Create a platform that can scale, quickly and with predictable cost
- Make data more readily available to be accessed, i.e. create a more nearline archive

Object Storage For A Data Driven World

So how do we create a platform to store this mass of unstructured data that can be made available globally, supports multiple use cases and scales with predictable costs?

The only way to accomplish this is to centralise your data storage, and use an object storage platform to enable simplification and scale - it's no coincidence that the worlds largest Cloud storage infrastructures are based on object storage.

Object storage differs from traditional file and block storage, storing data as HTTP objects instead of files or blocks; whilst it's less performant (so unsuited to application workloads) it is equally beneficial due to its granular metadata and near-infinite scalability.

Object storage usually consists of low-cost storage nodes with directly attached commodity drives with an object-based storage stack that manages the distribution of content across nodes; providing a series of benefits:

1. Reduction in physical footprint and environmental resources

By removing the need for RAID and replication, and by leveraging standard X86 server infrastructure as the underlying hardware platform, object storage enables significant reductions in datacenter footprint, in some cases up to 75%.

2. Simplified and improved data protection with less duplication

By using distributed erasure coding to protect your data, object storage systems enable you to remove the use of RAID and replication, which in turn reduces the amount of data that is duplicated.

3. Reduced total cost of ownership

By reducing replication of data, the physical footprint of your infrastructure (and therefore the environmental impact) and by leveraging object storage as a platform for multiple use cases, organisations dramatically reduce the total cost of ownership of their storage.

4. Easier to scale with more predictable cost

Object storage is typically deployed on lower cost, commodity hardware, making the cost to scale systems more predictable and manageable. Operationally they become easier to scale with the ability to add more X86 storage nodes to extend capacity as required.

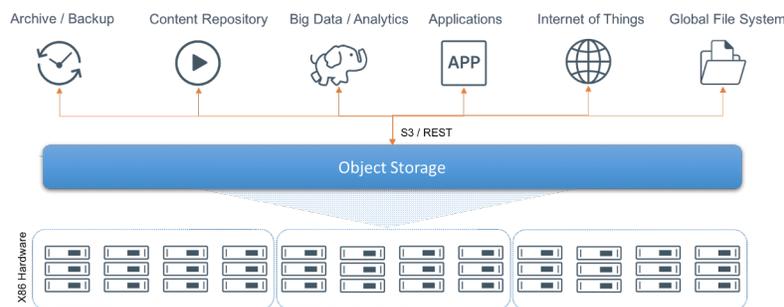
The Importance of An Ecosystem

Whilst they are intelligent in the way that they store and protect your data, used in isolation object storage systems have a limited scope, i.e. they deliver one thing (large scale, low cost capacity storage) extremely well.

To drive the most value from it though it must be combined with third party ISV (Independent Software Vendor) solutions to enable you to use it as a single horizontally scalable platform that can deliver multiple use cases. By using a common communication protocol (typically an S3 interface) ISV partners have a common interface to integrate with to enable a simple plugin to an object storage system.

Object Storage Use Cases

Because of the common S3 interface that object storage systems support, it's possible to support multiple use cases either by writing applications natively to an object store, or by using an ISV ecosystem to provide specific features as illustrated below



Outside of native application integration, some of the more common use cases for object storage have been:

- **Long Term Archive:** The use of tape for long term data retention/archive has long been problematic due to the management overhead, cost and reliability. This as well as the fact that organisations want more readily available access to their archive makes it a prime use case for object storage due to its scalability, cost and active nature
- **Global File System:** traditional infrastructure approaches have led to silos of infrastructure and duplicated data. Object storage enables the centralisation of data, that when coupled with Cloud storage gateways enable the delivery of a single global unified file system, reducing cost, duplicate data and enabling global collaboration
- **Big Data Analytics:** by their nature Big Data workloads demand large data sets, that drive larger capacity storage requirements that scales to meet the demand. Object storage can be used as the centralised data repository, moving data between the object store to a performance tier for analysis.
- **IoT Data:** as we enter a world of the connected everything data from sensors and other IoT devices are being generated, collected and stored at growing volumes. This data needs to be stored and stored at scale, with native application integration.

Real World Benefits

When used to its full capability, and when partnered with ISV partners, object storage offers a horizontally scalable storage system that reduces the cost and complexity of traditional approaches.

The key to recognising these benefits are often held in the understanding of your data - once you understand the data that you hold, the workflows that surround it and the protection that it requires you are then able to build a strategy to leverage object storage as a platform.

Even at a basic level for a single use case such as global file systems, object storage instantly demonstrates value, as highlighted in the example below:

	Capacity	Reliability	Management & Environmental
1PB Usable Storage 3 Copies, RAID 6 + Tape Backup	<ul style="list-style-type: none">• 5,050 TB Raw• 1,350 * 3TB Disks	<ul style="list-style-type: none">• 5 * 9's With Disk-to-Tape	<ul style="list-style-type: none">• 3 * FTE's• ~4 * Rack footprint• ~4 * power, cooling etc.
1PB Usable Object Store Over 3 Sites	<ul style="list-style-type: none">• 1,600 TB Raw• 534 * 3TB Disks	<ul style="list-style-type: none">• 12 * 9's Total• 5 * 9 Per Site	<ul style="list-style-type: none">• 0.5 FTE• 1.6 * rack footprint• 1.6 * power, cooling etc.
Object Storage Savings	60%	100 million * More Reliable	80%

These operational and environmental reductions, as well as the removal of individual storage silos, translates to real world cost savings. This cost reduction can be anywhere between 50-80% reduction in total ownership costs. Not only this but it simultaneously reduces the management overhead, environmental impact and risk of data loss.

How Can Nephos Help You Leverage Object Storage?

With over 15PB of object storage infrastructure in customer environments, and a focus on delivering Next Generation Datacenter infrastructure, Nephos Technologies are uniquely positioned to deliver an object storage platform that delivers value to your business.

Through Nephos' SMARTService offerings, Nephos works with our customers to understand their data landscape, technical and commercial requirements, to design and deliver an object storage strategy that ensures best value.

Nephos Technologies Ltd

International House, St Katharines Way, London, E1W 1UN

T. +44 (0)8453 104 105 | WWW.NEPHOSTECHNOLOGIES.COM | SALES@NEPHOSTECHNOLOGIES.COM