

Foto: © Sergey Nivens | gettyimages.com



WHITE PAPER | OBJECT STORAGE & S3

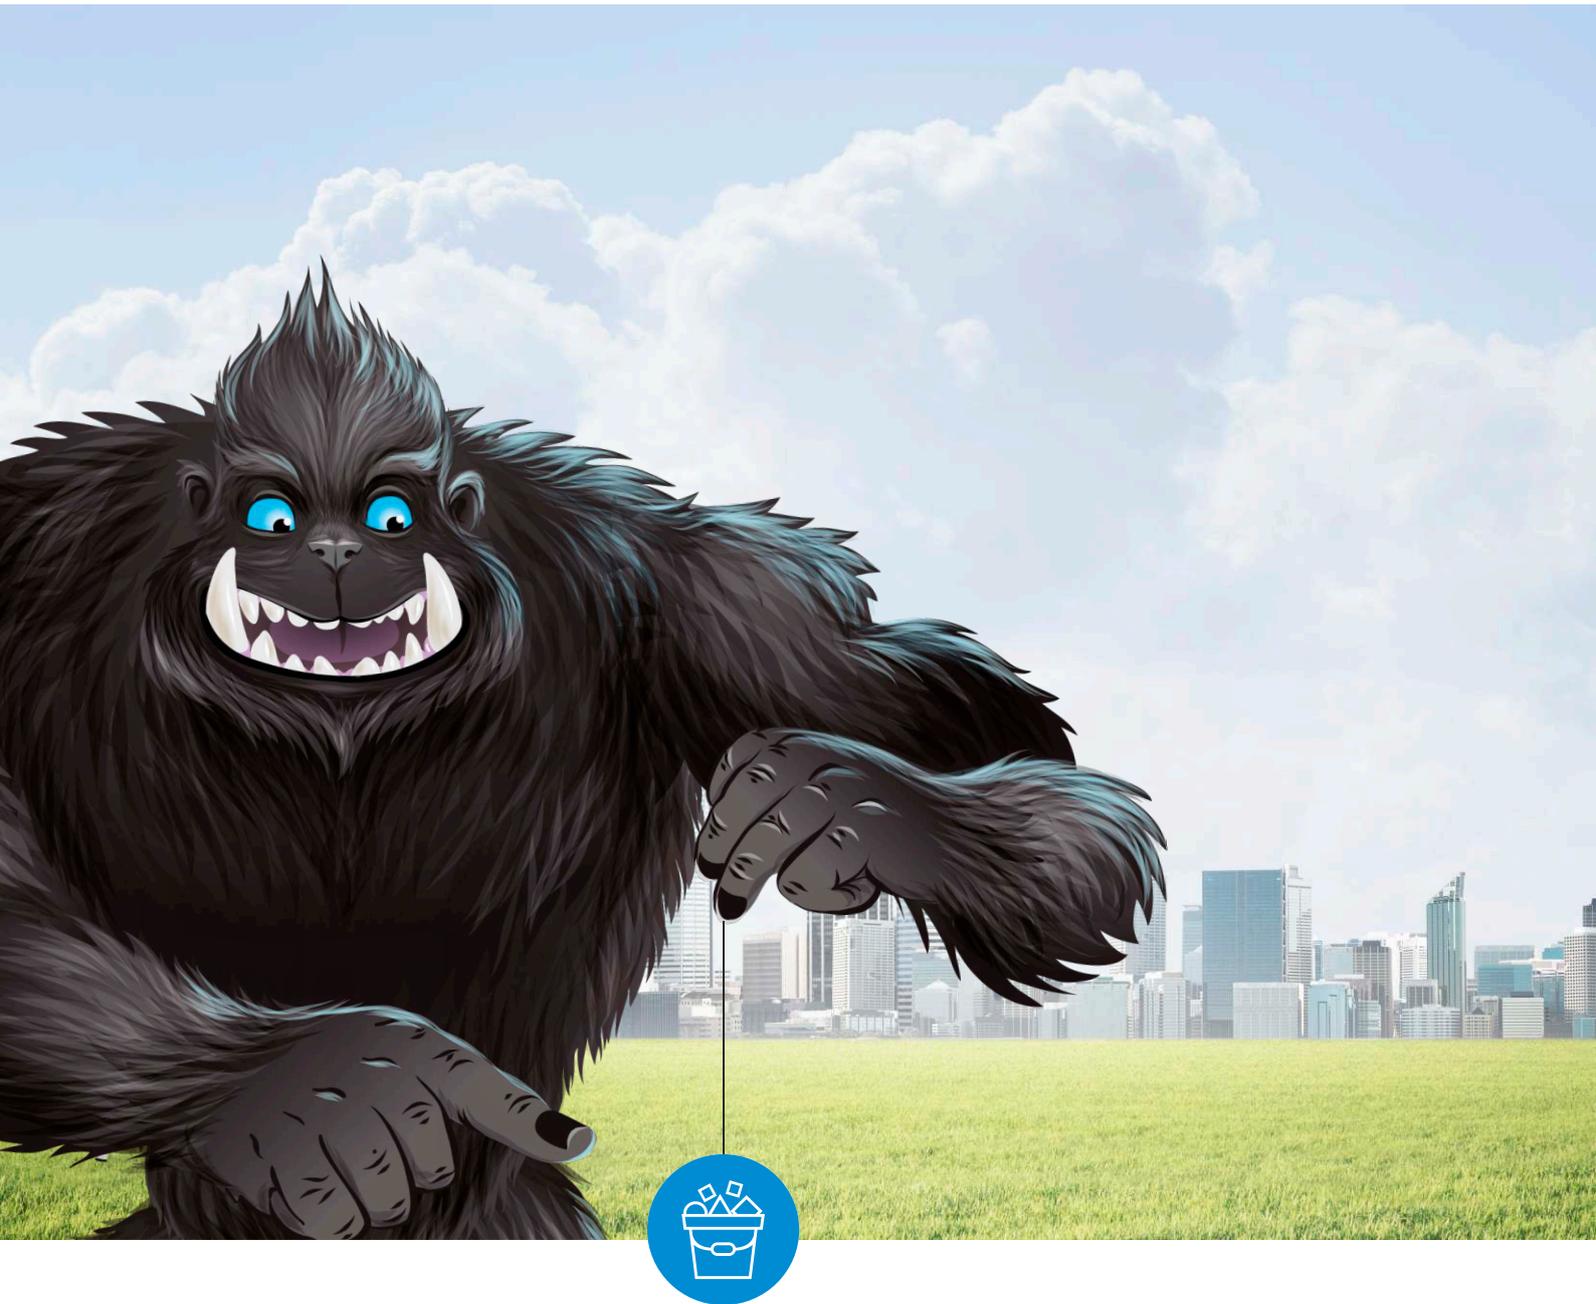


# Object Storage and S3 Not just for the big guys!



RNT Rausch. Making IT possible.







## Object Storage and S3

### Not just for the big guys!

*Object storage, Amazon S3, cloud storage, hybrid cloud, so many terms are floating around, but what is this all about? Is this just big enterprise technology or could it make a difference for every business? What are the benefits especially for data volumes in the terabyte (TB) rather than the petabyte (PB) space?*

**Let's take a closer look.....**

### Introduction

Information is the most valuable business asset. As a result, data is stored 'just in case' even though its current business value is unknown, because it might be useful in the future. The data volumes, management burden and associated costs for managing, securing, analysing and storing are growing alongside the demand to access data when and wherever needed.

Computer technology has allowed businesses to grow faster and achieve far more than during the days when everything was paper based. IT has brought huge benefits for businesses on one hand. On the other hand, the technology of the time was often less able and couldn't deliver more capability, capacity, and flexibility at a reasonable cost. In the last few years, digital transformation has driven a major growth in compute power and software suites are removing many of the technological restrictions of the past.

In our digital age, far more data is generated and reprocessed which is often not just text-based but far 'richer' than a few years ago, such as images, presentations, audio, video, etc. This type of data is known as 'unstructured data' as it cannot be arranged into a traditional relational database. Typically, unstructured data represents 80-90%+ of all the data in a business and various analysts predict that we'll globally create more data by 2025 than all the data created over the past 30 years.



## The data storage challenge

Data storage has been a fundamental of computing since Tommy Flowers built the first programmable electronic computer in 1943. Today, all businesses are more reliant than ever on IT and quick data analysis to function, compete and grow - making data the most valuable company asset.

Data faces a hurdle, it requires storage. While subjects such as server performance, hyper-convergence, virtualisation and ML/AI are seen as hip subjects, storage has always languished near the bottom of the hip pile. This is very surprising because storage is the foundation of every IT system, without data and timely access no computer system, however hip, can function.

The huge advancement in compute technology in recent years combined with virtualisation has driven an explosion in flexible server and application deployment. This is not so true of data storage. It has definitely got faster, but it's still heavily based on traditional block and file solutions which are very good for certain use cases. However, they have limits and restrictions because they are relatively old designs, and adapting them to store our modern unstructured data sets in the collaborative cloud era is neither easy nor ideal.



## How S3 became the storage language of the Cloud

In the mid 1990s, there were a number of developments that laid the foundations for a new method to store data with the aim to eliminate many limitations that block and file storage impose. By 1999, the first standards emerged, covering object storage, scalable computing, platform independence, and storage management; the Cloud Data Management Interface (CDMI) published by SNIA (Storage Networking Industry Association).

In 2000, Amazon.com began to focus on their version of object storage to be able to deliver cloud-based infrastructure-as-a-service. This led to the introduction of Amazon Simple Storage Service (S3) in 2006, which has become the object storage quasi-standard protocol in the meantime.

Amazon's own object storage is clearly 100% S3, and software vendors such as Cloudian are also purely S3. However, other vendors have their own object storage protocols/languages but most support enough S3 API (application programming interface commands for varying levels of interoperability).

There are many choices for object storage, but they all deliver key values compared to block and file. Object storage has been designed to simplify the storage, management, and retrieval of unstructured data in an easier way and thus improves flexibility at a reduced cost.

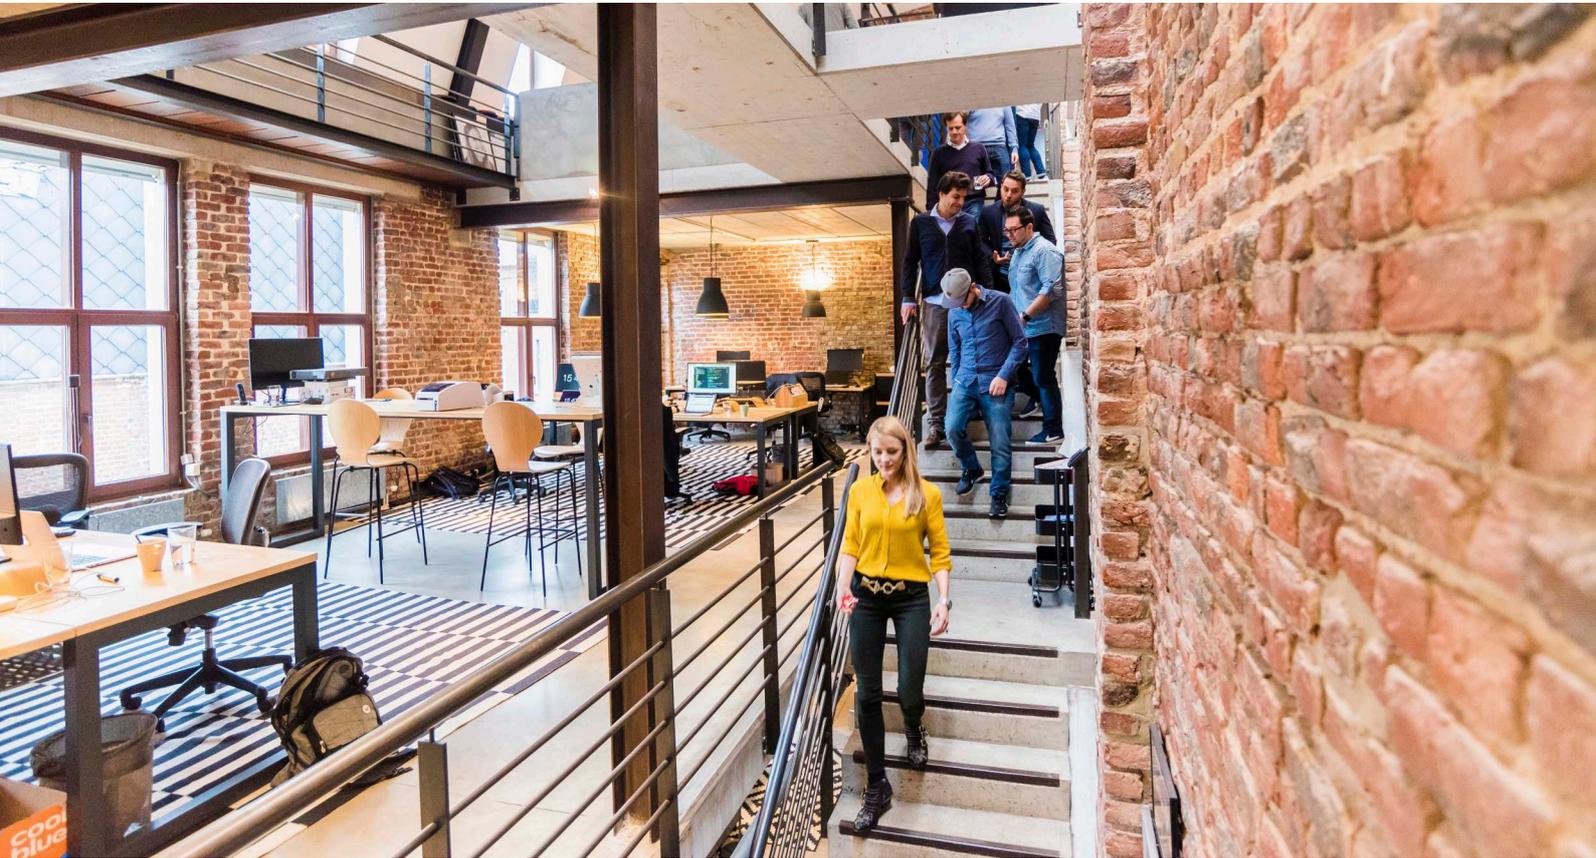


## What's so different in Object Storage?

Each item held in object storage is called an object and is given a unique ID. An object also has metadata associated but unlike file and block storage this metadata can be edited and expanded by humans, AI, or other software. The metadata could have key words to categorise any object such as:

- ⊕ who created it
- ⊕ target audience description (i.e. end user, internal, technical, sales, and more)
- ⊕ content key words
- ⊕ etc.

With such editable and rich metadata, the search for objects/files is simplified. Additionally, the unique ID makes it simple to directly retrieve and access any object from any location and without the need to use layers of folders (hierarchical directory structures). There is far less requirement for data management; just put it in a bucket and call it back out directly whenever needed.



## What's in it for me

and why should I care?

*Object storage was designed with the cloud and changing (digital) business practices in mind. Allowing multiple differing operating systems and applications to all jointly connect without the traditional technical connectivity challenges and storage siloes; you just need a web browser.*

***Let's have a closer look at some core values of Object Storage and the S3 protocol:***

1

## Handling data growth quickly and efficiently

Much is spoken about the explosion in data volumes and invariably we are presented with graphics about the growth in internet videos or music for some enterprise businesses. For smaller businesses, that are not seeing petabyte (PB)+ per year growth, the information is all very interesting but is the message applicable? Well, yes. An explosion in data volume doesn't need to be in the hundreds of TBs or PBs; it is always relative. For example, data growth from <5TB to >15TB is still a data explosion as the volume triples. The simplicity to handle this growth quickly and efficiently is one of the great values that object storage brings: simple scalability whether that is on premise, in the cloud, or both as a hybrid solution.



2

## Flexibility of access

Object Storage is making information securely accessible from anywhere, barriers due to network limitations are in the past. This makes accessing, collating, or analysing data so much simpler, easier, and more effective for businesses of any size.

Every day we benefit from this flexibility when we browse the internet as most of the pictures, videos, and downloads, etc. are stored on object storage somewhere in the world. All we care about is to quickly access what we need on any device and when we want it.



3

## Secure data access

File and block storage typically needs to be accessed on a local area network (LAN) or virtual private network (VPN), because data must be accessed securely, and their protocols were not designed with the cloud era in mind.

The ability to use cloud storage, whether dedicated or hybrid, highlights security as another value. Object storage is designed to access data wherever needed using the standard highly secure protocols that underpin the internet

such as https. This means all data in flight is highly encrypted and as these standards are updated all object storage traffic will automatically be covered. Data access is no longer limited to one network but available anywhere, anytime, easily, securely and with total confidence.

This ease of secure access adds another value, the simplicity with which software can talk directly to object storage which does not need to be on the same network or location. This allows a whole new slew of opportunities and efficient yet secure architectures that were difficult or even impossible to deploy in the past.



The benefits of object storage data access and management are clear. The broad adoption of S3 as a quasi-standard has resulted in large number of software suites integrating S3 compatibility. This makes it simple for companies to adopt object storage as part of their strategy. A good example is Microsoft SQL Server which now supports S3, even though their own cloud storage offerings do not have S3 support.

Accessing data stored as block and file is just a read and write process. However, accessing data via the S3 API gives the applications complete access to all the API's features. Object storage

certainly brings new capabilities that allow far more intelligent and flexible data management but most importantly allow applications to have complete control over their data and its management. This brings about a wholesale change; any application can now use the S3 API to tier its data for example or create new buckets itself. With block and file the application is reliant on an independent system, such as the storage controller or people to manage its data which, even with some proprietary interface, is not as efficient or cost effective.



## 4

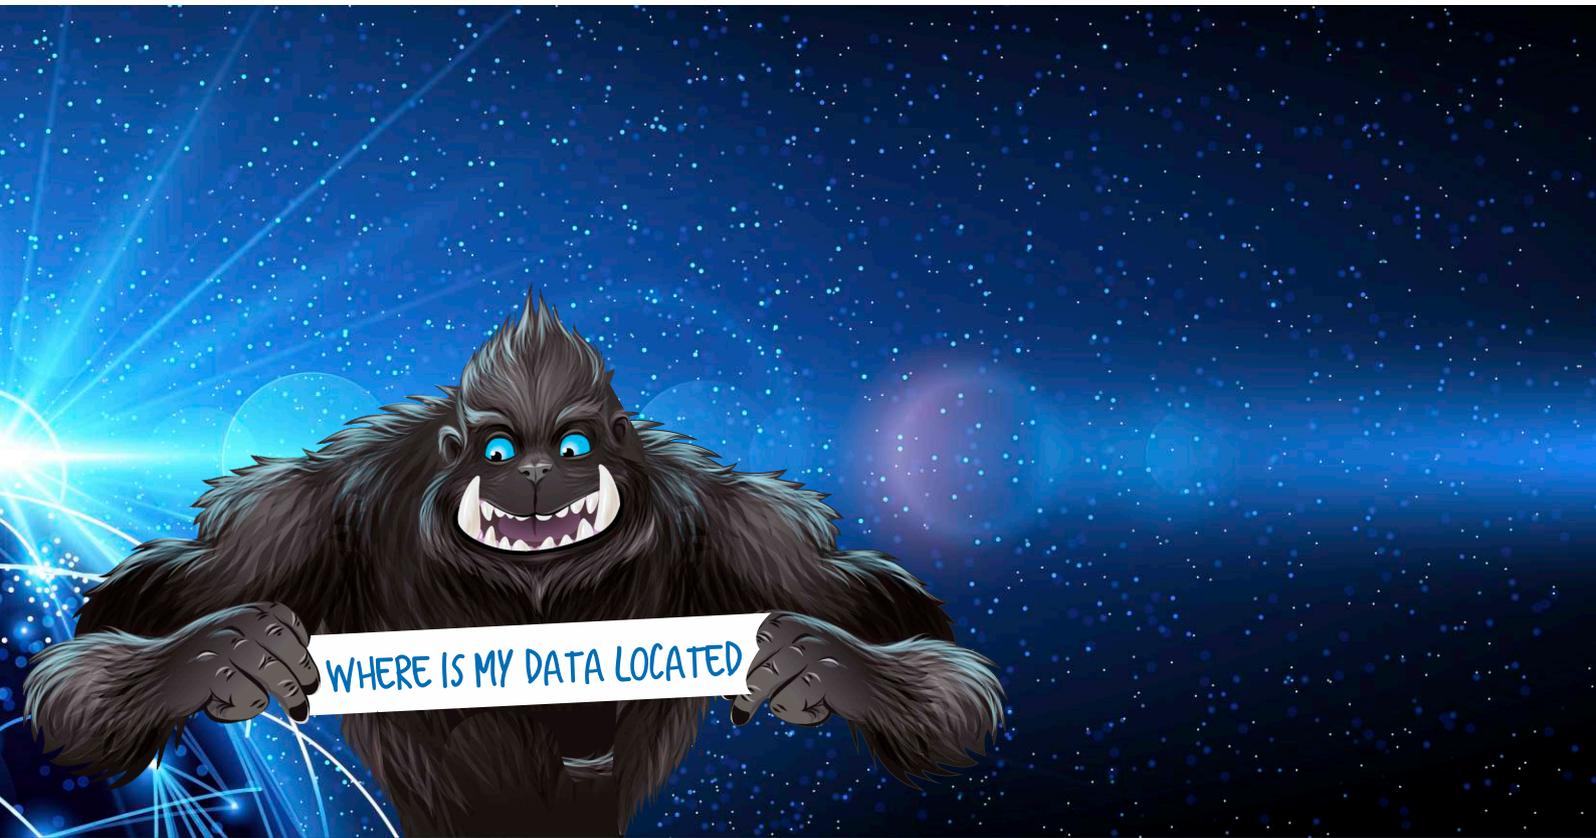
### Data sovereignty – put the cloud where you need it

#### 'It's in the cloud',

wonderful! But physically, where is it and who has legal access? Who has access? Organisations and Governments around the globe are waking up to the value of data and creating legal mechanisms to extricate it from foreign countries for their own use. At the same time, many of those same administrations are passing regulations to force businesses to 'defend' their data by ensuring it is kept inside certain regions or geographies. This can easily create a couple of serious conflicts and legal implications which is making data sovereignty a business imperative.

An excellent example of such a conflict are the 2018 U.S. Cloud Act or to give its full name, The [U.S. Clarifying Lawful Overseas Use of Data Act](#), and the [EU General Data Protection Regulation](#), in short GDPR.

In case of an ongoing criminal investigation, the U.S. Cloud Act gives United States federal law enforcement authorities the power to request access to all data stored by electronic communication or computing service providers subject to US jurisdiction. This also applies if the appropriate data is stored in a data centre outside the US territory.



Disclosure of all data is therefore required under US law although it violates GDPR. As a result, affected EU citizens have no legal protection against requests under the CLOUD Act and since the EU-US Privacy Shield was declared invalid by the [Court of Justice of the European Union \(CJEU\)](#) on July 16th, 2020, it's almost impossible for companies subject to EU jurisdiction to stay GDPR compliant if their data is stored with providers headquartered in the US. Additionally, GDPR does not provide any protections against requests for company related data that could be highly confidential, such as pending patents, research or test results.

Thus, data sovereignty is imperative for all companies that don't want to face such legal, business, and financial problems. Two key steps are vital

- ⊕ Ensuring the (Cloud) provider is headquartered in a legally safe jurisdiction and/or;
- ⊕ the data is physically within a specified geography as required by local laws.

Some of object storage's fundamental features make it invaluable for data sovereignty situations, especially if budgets are tight. Principally, object storage's cloud friendly design makes its deployment swift and easy without the need for private networks which add more complexity to the entire architecture. With Object Storage companies have full control of where the data resides while still enabling very cost efficient, flexible, simple, and secure access, whether delivered as an on-prem solution, private, hybrid or public cloud.

5

## Business continuity – the goal of data replication

Copying and storing data in multiple locations creates a distributed environment that improves:

- + data availability
- + data access for local users;
- + data durability

Replication enables local users to access data faster, and without disrupting other users. Inbuilt replication between object storage systems is something many object storage vendors offer. For example, the S3 API provides a suite of replication features enabling object storage systems and applications to offer higher data protection and more flexible architectures for data access and sharing.

The ease with which replication can be configured and managed (especially compared with many legacy storage systems) makes offsite copies not just simple but very compelling and provides much greater security for any company's most valuable asset: its data.



Data replication can be kept within a company's infrastructure/private cloud. However, the flexibility of object storage also allows some, or indeed all the data, to be replicated to public or 3rd party cloud providers creating a hybrid cloud environment.

The ability to easily manage and replicate data geographically provides speed of access to businesses with multiple offices or remote workplaces. Only a few, typically larger companies, benefit from huge pipes to the internet. As such, the often-required immediate access to relevant data can be far less than instant, resulting in commercial impacts.

Having data that is accessible locally is usually the best option for speed but multiple copies across different sites often lead to data management nightmares. Object storage puts an end to these challenges by not only providing the services to deliver local storage but also replicate the information to another site and/or cloud vendor.

6

## The data protection challenge

Data forms the foundation of business today and is every business' most valuable (although often undervalued) commercial asset. Accessing, processing and quickly analysing data for commercial gain is now more important than ever to ensure business success and keep customers happy. There is one key aspect about data that is still a hot topic: the malicious deletion or encryption of data to bring a business operation to a complete halt and/or cause severe damage to its reputation.

Cyberattacks in general, and ransomware in particular, became the digital plague and both are still on the rise. Deliberate or accidental deletion of data is not new, which is why we should have firewalls, access control, trusted and educated users and locks on the doors of data centres amongst a variety of measures that need to be part of the complete data protection strategy.



The rise of ransomware has probably become one, if not the most dangerous global threat to businesses of any size. Our need to be connected 24x7 and to utilise so much more data and applications gives cybercriminals far more ways to attack and encrypt entire systems.

Additionally, there are more rules and regulations around what kind of data must be retained and for how long, from within business as well as official regulations and laws.

Storage is often the last line of defence in any given data protection strategy and should at least store data using a write once, read many (WORM) model. To protect any company's most valuable asset, use Object Lock, the only cloud object storage WORM capability in the market today.

7

## Object Lock – immutable data for everyone!

Whatever the reason for an attempt to delete or encrypt data, the key is to prevent all damage and so minimise the business impact. Some object storage systems have a feature called Object Lock that provides the ability to apply individual immutability (retention) settings for objects and/or default immutability settings for all objects within a bucket.

The exact features of vendors' object lock functionalities and protection levels vary from supplier to supplier but fundamentally they all provide the ability to prevent data sets from being accidentally or deliberately deleted or encrypted.

Object Lock should be considered a must-have in the data protection arsenal. Immutable data is the ultimate final defence against attacks from in-house or cybercriminals from the outside. Either way, companies do not have to worry about the ransom costs or whether they'll even be provided with the de-encryption key(s) on payment.

Data can be manually locked or managed through scripts or applications that can work with the Object Storage Object Lock API (the interface for control of the locking process). A much overlooked but crucial application is backup or, even more importantly, recovery. By necessity, it's there for the bad days to get the data and systems back ASAP. While ransomware may have encrypted systems from servers to laptops, object lock compatible backup applications such as Veritas, Commvault, Veeam, amongst others, ensure nothing is lost and operations are restored as swiftly as possible from a clean backup.

Vendors, such as Cloudian, provide object lock at an additional level for enhanced data security.

- ① They protect the data from deletion at the S3 level so an external hacker cannot use S3 storage commands to modify or delete the data
- ② The data is protected at the disk level, so someone with local console access and full admin rights cannot delete the data either

This is vital to defend against hackers with full access to company's systems.





## 8

## Data recovery - where should I put my data?

There is a final and very important concern, where should the storage be situated physically, on-premises or in the cloud or both? If data sovereignty is applicable, the physical storage holding that data clearly must be either on-prem or in the relevant physical geography. However, if the location is not a factor, with object storage the data could physically be anywhere in the world. Perhaps with the cheapest cloud provider the web could suggest if budgets are tight. This may work but the bandwidth and speed of light can often prove to become enemies of such a strategy.

There is another important factor, recovery time in case a disaster strikes. Backup windows are rarely an issue and are often discussed but the recovery time always matters and is often not considered properly in a system design. There's a lot to consider when designing a recovery plan but the two key questions are:

- ⊕ In the event of an outage what are the key systems that are needed for the business to function?
- ⊕ How fast must they be back online?

The answer to those two questions will very clearly determine the bandwidth and performance of both storage and network and whether one backup copy should be kept on-prem or solely in the cloud.

Object storage can be used to provide the most efficient commercial solution for any recovery scenario. Relevant backups can be stored on the local site where they are needed, with copies stored offsite as an extra level of security. Once configured, these replicated copies are managed by the object storage whether they are at another office or in the cloud. Such a strategy delivers a swift restore even if a site is taken offline, for example in response to a ransomware attack to prevent the attackers from continuing their activities. The off-site copies provide another level of security but obviously the restore time is often far slower due to the available bandwidth.





## Summary

Is object storage just for the big guys? Absolutely not. Object storage is not the answer to every storage use case but it's a game changer for all businesses dealing with vast amounts of burgeoning unstructured data volumes. Although it's not a certified protocol, S3 is, by far, the most well adopted object storage protocol and delivers all the services needed for a modern digital world. Data is invaluable and every business needs to be able to access and utilize it quickly to grow and compete successfully. Cost, security and ease of access are three key aspects that should govern data storage. Object storage makes access and management of unstructured data more flexible while applying highest security standards for data protection and reducing the overall storage cost. Clearly, every business, small and big, should evaluate how object storage can help them grow, drive innovation, and meet digital transformation goals.

What used to be a challenge, was finding a cost-effective and all-in-one enterprise class object storage solution designed for smaller data volumes while delivering all the enterprise grade values of both the hardware and software features such as immutable storage.

The [Yowie Appliance portfolio](#) from RNT Rausch puts an end to the search. Powered by [Cloudian Hyperstore](#) the installation is done via a simple wizard including installing object lock which is certified by 25 countries including USA, Germany, and Switzerland. The Yowie Appliance family is clearly challenging the status quo of 'SMB Storage' and brings the values of object storage and S3 from 2 TB useable capacity to smaller companies with Simplicity-as-a-Service in mind.

RNT Rausch GmbH  
WHITE PAPER  
OBJECT STORAGE & S3

# RNT Rausch. Making IT possible.

RNT Rausch is a Germany based technology pioneer with 20+ years of experience in the high-tech server and storage industry and on a mission to always be ahead of technology trends. RNT offers tailor-made products and services that help customers and partners to enhance security, flexibility, scalability and sustainability and become fit for tomorrow's technical revolution.

## Let's talk!

The RNT channel team is looking forward to getting in touch with you now!

Mail: [channel@rnt.de](mailto:channel@rnt.de)