IBM Cloud Object Storage System[™] Version 3.14.3

Definitive Guide to Dispersed Storage



This edition applies to IBM Cloud Object Storage System *Definitive Guide to Dispersed Storage* and is valid until replaced by new editions.

© **Copyright IBM Corporation ..** US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Chapter 1. Why should I read this
guide?
Is your current storage system petabyte-ready? 1
Chapter 2. Information dispersal 3
Dispersed storage defined
The math behind information dispersal 4
Chapter 3. How dispersed storage works 5
Chapter 3. How dispersed storage works 5 How dispersed storage works: Step-by-step 5
Chapter 3. How dispersed storage works5How dispersed storage works: Step-by-stepSystem components
Chapter 3. How dispersed storage works5How dispersed storage works: Step-by-stepSystem componentsObject Storage foundations
Chapter 3. How dispersed storage works5How dispersed storage works: Step-by-step.System components.Object Storage foundations.Access methods
Chapter 3. How dispersed storage works5How dispersed storage works: Step-by-step.System components.Object Storage foundations.Access methods.Security features.10
Chapter 3. How dispersed storage works5How dispersed storage works: Step-by-step.System components.Object Storage foundations.Access methods.Security features.10Availability features.

Chapter 4. The benefits of dispersed

storage		•		•	-	-	-	•	•	•	•	•	•	17
Scalability I	ben	efit	s.											17

Availability benefits						. 17
Security benefits						. 18
Economic benefits .						. 18

Chapter 5. Solutions for dispersed

storage				-					. 21
Use cases									. 21
Content repository storage									. 21
Storage-as-a-service									. 22
Enterprise collaboration .									. 22
Backup									. 23
Active archive	•	•	•	•	•	•	•		. 24
Chapter 6. Conclusio	n	-							. 27
Is dispersed storage right f	or	you	r o	rga	niz	ati	oní	?.	. 27
Notices								-	. 29
Trademarks									. 31
Homologation statement .									. 31

Figures

1.	Data served on multiple devices	
2.	Difficulties with traditional storage approaches 1	
3.	Mathematics of the information dispersal	
	algorithm 4	
4.	How dispersed storage works	
5.	How multi-site failures are managed 6	
6.	How single/multi-site failures are managed 7	
7.	REST API storage interfaces	
8.	Write operation	
9.	Read operation	

10.	AONT encoding
11.	Network security
12.	Device registration
13.	Integrity checks
14.	Continuous error correction
15.	SmartWrite technology
16.	SmartRead technology
17.	Comparison of total cost of ownership of IBM
	Cloud Object Storage System versus NAS 2015. 19

Chapter 1. Why should I read this guide?

Is your current storage system petabyte-ready?

This information describes a new approach to solving the challenges associated with storing large volumes of unstructured data. They are challenges that cannot be met by traditional storage technologies.

The unstructured data explosion



Figure 1. Data served on multiple devices

According to IDC, the total amount of digital information that is created and replicated surpassed 4.4 zettabytes (4,400 exabytes) in 2013. The size of the digital universe is more than doubling every two years and is expected to grow to almost 44 zettabytes in 2020.

Although individuals generate most of this data, IDC estimates that enterprises are responsible for 85 percent of the information in the digital universe at some point in its lifecycle. That means organizations take on the responsibility for designing, delivering, and maintaining information technology systems and data storage systems to meet the demand.

Traditional storage approaches don't work



Figure 2. Difficulties with traditional storage approaches

The figure shows difficulties with storage at the petabyte level such as data errors, disk failures, and security breaches. Addressing these issues with traditional systems means skyrocketing costs with no end in sight.

Technological advances are helping with this challenge to some degree. Computing is getting faster and cheaper. Virtualization is driving up efficiency and usage. Storage devices are growing in terms of capacity while declining in price (more bits per device at a lower cost). They are getting faster with the advent of solid-state technologies (although not currently at a suitable price point for all workloads). Delivery mechanisms such as cloud computing are also helping to lower costs and drive efficiencies. In some cases, the advances in technology (specifically the capacity expansion of storage devices) are putting a strain on traditional methods of protecting and preserving digital information.

Traditional storage protection technologies such as RAID are inadequate when it comes to protecting digital information from data loss at petabyte-scale and beyond. Traditional storage architectures are not designed to scale to the petabyte range. They are less secure, less reliable, and more expensive.

Consider the following challenges that traditional storage systems face when they reach petabyte-scale:

- Data integrity suffers when system size is 10 billion times larger than the bit error rate of a hard disk drive.
- Data availability suffers when hundreds of drives fail every day and require a week to rebuild.
- Data security suffers with millions of devices and multiple copies in multiple locations.

Enterprises that need to store large volumes of unstructured data must look beyond their current storage solutions and evaluate new approaches. *The Definitive Guide to Dispersed Storage* helps you understand how dispersed storage works and how its unique benefits helped other organizations achieve high levels of scalability, availability, and security while still controlling storage costs.

Chapter 2. Information dispersal

Dispersed storage defined

The IBM Cloud Object Storage System[™] software platform uses an innovative approach for cost-effectively storing large volumes of unstructured data while still ensuring security, availability, and reliability.

The IBM Cloud Object Storage System[™] storage technology uses Information Dispersal Algorithms to separate data into unrecognizable "slices" that are distributed via network connections to storage nodes locally or across the world. The collection of dispersed storage appliances creates what is called a dispersed storage network. With dispersed storage technology, transmission and storage of data are inherently private and secure. No complete copy of the data exists in any single storage node. Only a subset of nodes needs to be available to fully retrieve the data on the network.

Background

By taking the methods that the internet uses for data networking and applying them to data storage, dispersed storage allows companies to store massive amounts of content (video, audio, photo, text) securely and reliably.

Much like how the internet uses an open protocol (TCP/IP) based on the improved design of packet switching in comparison to the established telephony protocols used in older circuit-switched networks, dispersed storage is a commercial-grade implementation of a technology for data storage called Information Dispersal Algorithms (IDAs).

IDA technology transforms data into slices by using equations such that a subset of the slices can be used to re-create the original data. These slices, which are like packets but are for data storage, are then stored across multiple storage appliances (or storage nodes). Slices are created by using a combination of erasure coding, encryption, and sophisticated dispersal algorithms.

Dispersed storage systems are well-suited for storing unstructured data like digital media of all types, documents that are produced by desktop productivity applications, and server log files, which are typically larger files. Dispersal is not optimized for transaction-oriented primary storage for databases and similar high IOP workloads because of the extra processing associated with slicing and dispersing.

What is information dispersal?

At the foundation of the IBM Cloud Object Storage System[™] storage platform is a technology that is called Information Dispersal. Information Dispersal is the practice of using Erasure Codes as a means to create redundancy for transferring and storing data.

An Erasure Code is a Forward Error Correction (FEC) code that transforms a message of k symbols into a longer message with n symbols such that the original message can be recovered from a subset of the n symbols (k symbols).

Erasure Codes use advanced math to create "extra data" that allows a user to need only a subset of the data to re-create it.

An IDA can be made from any Forward Error Correction code. The additional step of the IDA is to split the coded data into multiple segments, which can then be stored on different devices or media to attain a high degree of failure independence. Using forward FEC alone on files on your computer does not do much to help if your hard disk drive fails. If you use an IDA to separate pieces across machines, you can now tolerate multiple failures without losing the ability to reassemble that data.

The math behind information dispersal

How do these IDAs work?



Figure 3. Mathematics of the information dispersal algorithm

The answer is simple. If you can remember back to high school algebra class, you might recall that when you have a system of equations with, for example, five variables, you can solve for those variables when you have at least five outputs from different equations by using those variables.

The figure shows five variables (a through e) and eight different equations that use these variables, with each yielding a different output. To understand how information dispersal works, imagine the five variables are bytes. Following the eight equations, you can compute eight results, each of which is a byte. To solve for the original 5 bytes, you can use any five of the resulting 8 bytes.

This example shows how information dispersal can support any value for k where n-k is the number of variables, and n is the number of equations.

Chapter 3. How dispersed storage works

How dispersed storage works: Step-by-step

At a basic level, the IBM Cloud Object Storage System[™] platform uses three steps for slicing, dispersing, and retrieving data.

- 1. Data is virtualized, transformed, sliced, and dispersed by using IDAs. In the first figure, the data is separated into 12 slices. So the "width" (n) of the system is 12.
- 2. Slices are distributed to some combination of separate disks, storage nodes, and geographic locations. In this example, the slices are distributed to three different sites.
- **3**. The data is retrieved from a subset of slices. In this example, the number of slices that are needed to retrieve the data is 7. So the "threshold" (k) of the system is 7.

Given a width of 12 and a threshold of 7, you can refer to this example as a "7 of 12" (k of n) configuration.

The configuration of a system is determined by the level of reliability needed. In a "7 of 12" configuration," five slices can be lost or unavailable and the data can still be retrieved because the threshold of seven slices is met. With a "5 of 8" configuration, only three slices can be lost, so the level of reliability is lower. Conversely, with a "20 of 32" configuration, 12 slices can be lost, so the level of reliability is higher.



Figure 4. How dispersed storage works

Multi-site failure

With dispersed storage, only a subset of slices is needed to retrieve the data. A dispersed storage system can tolerate appliance failures both within a single site and across multiple sites, as shown in the following figure.

1. Data is virtualized, transformed, sliced, and dispersed by using Information Dispersal Algorithm (IDAs). The "width" (n) of the system in this example is 12.

- 2. Slices are distributed to separate disks, storage nodes, and geographic locations. In this example, the slices are distributed to four geographically dispersed sites.
- **3**. The data is retrieved from a subset of slices. In this example, the number of slices that are needed to retrieve the data is 7. So even though failures are occurring across all three sites, the data is still available to be retrieved because the "threshold" of seven available slices is reached.



Figure 5. How multi-site failures are managed

Single-site, multiple-device failure

A dispersed storage system can also be deployed in a single site with the ability to tolerate the failure of multiple appliances within that site, as shown in the following figure.

- 1. Data is virtualized, transformed, sliced, and dispersed by using IDAs. The "width" (n) of the system in this example is 12.
- 2. Slices are distributed to separate disks, storage nodes, and geographic locations. In this example, the slices are distributed to four different racks within a single site.
- **3**. The data is retrieved from a subset of slices. In this example, the number of slices that are needed to retrieve the data is 7. So even though each rack experienced one or more device failures, the data can be retrieved because the "threshold" of seven slices is met. Even with five slices unavailable, the data can be bit-perfectly re-created.



Figure 6. How single/multi-site failures are managed

System components

You can use the IBM Cloud Object Storage System[™] platform to create storage systems that have three software components: the IBM Cloud Object Storage Manager[™] software, IBM Cloud Object Storage Accesser[®] software and IBM Cloud Object Storage Slicestor[®] software.

The software components can be deployed on a wide range of compatible industry-standard hardware platforms, as virtual machines, and in the case of the IBM Cloud Object Storage Accesser[®] software, as a software application that is running on a Linux operating system. Physical and virtual deployment can be combined in a single system by using virtual machine deployment for the IBM Cloud Object Storage Manager[™] and IBM Cloud Object Storage Accesser[®] software and physical servers for the IBM Cloud Object Storage Slicestor[®] software as an example.

Each of the three software components serves a specific function.

- The IBM Cloud Object Storage Manager[™] software is responsible for monitoring the health and performance of the system, configuring the system and provisioning storage, managing faults, and other administrative and operational functions.
- The IBM Cloud Object Storage Accesser[®] software is responsible for encrypting/encoding data on ingest and decoding/decrypting it when read and managing the dispersal of slices of data resulting from this process across a set of IBM Cloud Object Storage Slicestor[®] nodes.
- The IBM Cloud Object Storage Slicestor[®] software is responsible for the storage of slices.

When the IBM Cloud Object Storage Manager[™] software, IBM Cloud Object Storage Accesser[®] software, and IBM Cloud Object Storage Slicestor[®] software are deployed on a hardware platform that is certified by IBM[®], the following benefits result:

- Minimum time to production on initial deployment because hardware and software compatibility and configuration are predefined and validated by IBM.
- Hardware configuration optimized to maximize value of the IBM Cloud Object Storage System[™].

- Increased system reliability due to low-level monitoring and management of hardware component health.
- Access to IBM support staff that are familiar with both the hardware and software components of the system.

Object Storage foundations

The IBM Cloud Object Storage SystemTM is based on a simple Object Storage approach that efficiently stores billions of data objects in a single, flat, namespace. It allows the data to be accessed through a REST interface by using the HTTP-based protocol.

The old file-based way

Traditional storage systems organize data in a hierarchical file system and the data can be accessed via NAS-based protocols like NFS and SMB. The file system approach is ideal for human users but not for big data applications that must manage billions of objects. File systems allow humans to organize content in an understandable hierarchy where access speed is not critical and directory contention is minimized. When you are dealing with billions of objects, a file system approach does not scale and incurs performance breakdowns and bottlenecks.

The new object-based way

The dynamic data addressing capabilities of object-based storage lead to a number of advantages over traditional storage. Among them are massive scalability, improved storage efficiency and ease of data migration and movement. Object-based storage also allows for more metadata than traditional storage, making it easier to manage data tiering, security, and migration. [Forrester Research: "Prepare for Object Storage in the Enterprise," by Andrew Reichman. November 29, 2010.]

The benefits of object-based storage.

Scalability

Reach new cost, capacity, and accessibility milestones with petabyte scalability and beyond.

Security

Protect mission-critical data with zero-touch encryption and built-in carrier-grade security.

Availability

Make sure that your data is always available, regardless of planned or unplanned downtime.

Efficiency

Simplify management with an intuitive platform that is 15 times more efficient than earlier storage operations.

Economics

Dramatically reduce long-term total cost of ownership with a premium software-based solution that runs on commodity hardware.

With file system storage, data is closely tied to its location. Object-based storage overcomes this limitation by decoupling data from its physical placement in the storage system. Analyst firm Forrester Research cites valet parking as an apt analogy for object-based storage. [Forrester Research: "Prepare for Object Storage in the Enterprise," by Andrew Reichman. November 29, 2010.] When you valet park your car, the attendant gives you a claim ticket that allows you to retrieve your car when ready. While the attendant has your car, they might move it around as needed to optimize space in the lot or garage. The claim ticket identifies your car, but not a particular parking space. With object-based storage, an object ID identifies a particular piece of data, but not its specific location in the system. Data can be moved around the system as needed, and the object ID is the "claim ticket" needed to retrieve it, wherever it is located.

Access methods

The underlying storage pool of a dispersed storage system can be shared and is jointly accessible by multiple access protocols.

Object-based access methods

The Simple Object interface is accessed with a HTTP/REST API. Simple PUT, GET, DELETE, and LIST commands allow applications to access digital content, and the resulting object ID is stored directly within the application. With the IBM Cloud Object Storage Accesser[®] application, no IBM Cloud Object Storage Accesser[®] application is needed since the application server can talk directly to IBM Cloud Object Storage Slicestor[®] storage nodes.



REST API access to storage

Figure 7. REST API storage interfaces

REST is a style of software architecture for distributed hypermedia information retrieval systems such as the World Wide Web. REST-style architectures consist of clients and servers. Clients initiate requests to servers. Servers process requests and return associated responses. Requests and responses are built around the transfer of various representations of the resources.

The REST API works in way that is similar to retrieving a Universal Resource Locator (URL). But instead of requesting a webpage, the application is referencing an object.

REST API access to storage offers several advantages:

- Tolerates internet latency
- Provides for "programmable" storage
- Provides efficient global access to large amounts of data

File-based access methods

Dispersed storage can also support the traditional NAS protocols SMB/CIFS and NFS through integration with third-party gateway appliances. Users and storage administrators are able to easily transfer, access, and preserve data assets over standard file protocols.

Security features

SecureSlice^{\mathbb{M}} is the technology that is used to guarantee confidentiality, integrity, and availability of data stored on the system.

Data security

SecureSlice combines two algorithms: an Information Dispersal Algorithm (IDA) and an All-or-Nothing Transform (AONT). AONT is a mode of encryption in which the information can be deciphered only if all the information is known. The diagrams illustrate basic write and read operations by using SecureSlice.







2. With the complete AONT package, the data is decrypted.

1. Any threshold number of slices are put through the IDA to recover the original input, in this case the AONT package.



Internals of AONT encoding



Figure 10. AONT encoding

When a segment of data is to be stored in a dispersed storage system, an integrity check value is first appended to the data. The integrity check value can be any well-known constant value, if its length is sufficient. This value will be checked after decoding to ensure that no corruption occurred.

If any slice used in the reassembly of the data segment is corrupted, a high probability that the integrity check value is also corrupted exists. The dispersed storage system notices this corrupted value and prevents the invalid data from reaching the user. If the integrity check value is corrupted on a slice, the dispersed storage system attempts to find a valid combination of slices to retrieve the complete data segment.

Network security

All network traffic that is flowing into or out of appliances in a dispersed storage system is encrypted by using TLS with AES. Storage nodes can be placed anywhere without complex firewall or VPN setup, as shown in the following figure.



Figure 11. Network security

Device registration - certificate authority (CA) trust

It is not enough to say devices use TLS. These protocols do not prevent "man in the middle" attacks unless the connections are authenticated. Authentication with TLS requires the use of digital certificates, and these certificates must be verifiable as belonging to a valid node in the storage network. To accomplish it, nodes are given a signed digital certificate at the time they are approved into the storage network.

Such approval requires an administrator to log in to the management interface, view the request, and authorize it. The administrator can see the IP address, MAC address, and fingerprint of the device that is making the request, and verify that each is valid before the device is accepted into the system. When approved, the node is granted a certificate that is signed by the certificate authority (CA) for the storage network. All devices in the storage network trust this certificate authority, and by extension, any node that owns a valid certificate that is signed by this CA. Appliances can, at some future time, be retired or become compromised. Now, the IBM Cloud Object Storage Manager[™] can revoke the device's certificate, by adding it to the Certificate Revocation List (CRL), which is periodically polled by every node in the system.



Figure 12. Device registration

The device registration process includes the following steps:

- 1. The device registers with the Manager certificate authority (CA).
- 2. The certificate is sent back to the device.
- 3. A separate, secure connection is created to each device verified against the Manager CA.
- 4. The client is authenticated by using either username/password or PKI.
- 5. The password is checked against the authentication service.

Availability features

The availability features of a dispersed storage system provide continuous error detection and correction, ensuring bit-perfect data availability.

Integrity check on all slices and files

A dispersed storage system checks for data integrity through an intelligent background process that proactively scans and corrects errors. It scans data slices for integrity, rebuilds any corrupted slices, and checks for both slice integrity and file data integrity before delivery. This process guarantees bit-perfect data delivery through proactive correction of bit errors and correction of latent soft errors that might occur during normal read/write operations. It also ensures that data cannot be modified without authorization and that malicious threats are detected.



Figure 13. Integrity checks

Continuous error correction

If a slice is determined to be corrupted, meaning that the integrity check value is invalid, the IBM Cloud Object Storage Slicestor[®] appliance starts the distributed rebuilder technology to replace the slice with a valid slice. If the slice is missing, the distributed rebuilder technology re-creates a valid slice. Continuous error correction increases system availability because it is not waiting for data to be read to detect errors. It is crucial with long-term archives and massive digital stores where information isn't as frequently read. The distributed rebuilder model allows for predictability because the rebuilder is "always on" at a moderated rate, making I/O performance much more predictable, and scalable, as the rebuilder grows with storage.



Figure 14. Continuous error correction

Performance optimization features

Dispersed storage uses SmartWrite[™] and SmartRead[™] technology to optimize writes and reads of slices, resulting in improved throughput and efficiency.

Improved performance of k of n writes

SmartWrite enables a successful write operation even if the full width of slices cannot be written. For example, when a failure condition at a node or within the network exists. SmartWrite optimistically attempts to write all slices. When the write threshold of slices is achieved, SmartWrite considers the write successful. The remaining slices continue to attempt to write asynchronously. If a write operation times out, it is detected and rebuilt.



Figure 15. SmartWrite technology

Improved performance of k of n reads

SmartRead predicts the optimal network routes and storage nodes to most efficiently retrieve data. Data is reassembled in segments, and for each segment, thousands, if not millions, of combinations of slices are examined to determine the best delivery path. SmartRead ranks storage nodes by real-time performance and requests the optimal combination of slices to re-create the data. If a slice request is not performing, SmartRead requests a slice from another node.



Figure 16. SmartRead technology

Chapter 4. The benefits of dispersed storage

Scalability benefits

Dispersed storage provides massive scalability with minimal administrative tasks. Systems can grow easily from terabytes to petabytes to exabytes.

Multiple drivers of scalability

Storage and management of data at the level of petabytes, exabytes, and beyond requires an architecture that can scale. With no centralized servers, the capacity and performance of a dispersed storage system can be scaled independently.

The Object Storage foundation of dispersed storage enables the data mobility, scalability, and storage efficiency crucial for limitless scale storage. Traditional file systems do not scale beyond the petabyte level without bolt-on measures.

Dispersed storage delivers a single addressable global namespace that virtualizes all individual storage nodes, providing a single point of management. More benefits of using a global namespace approach include the ability to open up more storage pools for larger working pools of disks, migrate data transparently, and reduce the number of mount points/shares in an environment.

The dispersed storage protocol can be used within application servers or devices, each independently accessing storage nodes. This protocol enables massive parallel writes and content distribution to be achieved. And it avoids the choke points of a gateway, improving performance in a distributed environment.

Availability benefits

Dispersed storage maintains 100 percent data integrity even as millions of physical bit errors occur or as multiple drives, servers, containers, or locations change or are replaced. The same information that goes in is the same that comes out, guaranteed, and the information is accessible from anywhere, anytime. Data is always available with an architecture that can tolerate simultaneous failures.

Configurable availability and zero-downtime upgrades

Dispersed technology provides superior data protection and availability. Dispersal is significantly better than other storage solutions because it does not replicate data to overcome the shortfalls of RAID implementations. It does not suffer the significant risk of data loss as does RAID during the rebuild process, which might take many hours, or even days, for even a single hard disk drive. It is configurable to provide higher levels of fault tolerance (k of n) when compared to RAID 5 (1 of N) and RAID 6 (2 of n). By using IDAs and storing the resulting slices on independent hardware that can be either in a single site or geographically dispersed, the IBM Cloud Object Storage System[™] drives reliability and availability without replication.

The IBM Cloud Object Storage System[™] software platform allows enterprises to tolerate entire site failures and still have seamless access to data without expensive copies. So taking a data center offline for routine maintenance does not change availability. With dispersed storage, zero-downtime upgrades are possible. Rolling upgrades enable the system to remain operable with data accessible throughout the process. No scheduled maintenance window is needed.

Security benefits

Dispersal ensures data confidentiality even when multiple drives, servers, containers, or locations are compromised. Data in motion and data at rest is encrypted to make it unrecognizable and inherently secure to eliminate opportunities for security breaches.

Superior security for data at rest and data in motion

Dispersed storage technology uses encryption, all-or-nothing transformation (AONT), an integrity check value, and IDAs to effectively split data into inherently secure slices. Each slice is encrypted, but no external key management is needed. All data is computationally secure unless a "threshold" of slices is available to decrypt it. The result is slices that do not contain any representation of the data and that require a threshold number to bit-perfectly re-create the data.

These slices are stored on independent hardware. No full copy of the data exists on any storage volume. Individual servers that contain slices are useless without possessing a threshold number of them, which must be taken from many different physical locations. It means that the likelihood of a security breach is significantly lower and in some cases eliminated. Further, since the slices are created before traversing the network, the slices are secure against on-the-wire security breaches. The result is superior data security for both data over the network and data at rest.

Economic benefits

The IBM Cloud Object Storage System[™] delivers significantly lower total cost of ownership for storage systems at the petabyte level and beyond by eliminating expensive replication and associated incremental costs. Hardware, electricity, floor space, support, and management costs are also reduced.

No copies, lower costs

Dispersed technology eliminates the need for costly replication. The IBM Cloud Object Storage System[™] delivers the equivalent availability of 3 - 4 replicated copies of data while reducing storage requirements by up to five times when compared to traditional approaches.

Figure 17 on page 19 shows the cost of ownership in dollars per terabyte, based on a customer analysis. In this example, NAS requires:

- Over three times the raw storage of information dispersal
- Three times the power and cooling costs
- Four times the floor space
- Added costs for hardware, software, and a tape copy

The cost of the IBM Cloud Object Storage System[™] is more than 80 percent lower.



Figure 17. Comparison of total cost of ownership of IBM Cloud Object Storage System[™] versus NAS 2015

Chapter 5. Solutions for dispersed storage

These case studies make a powerful argument for IBM Cloud Object Storage System[™].

Use cases

The IBM Cloud Object Storage System[™] platform is ideal whenever enterprises need to securely store large volumes of unstructured data with high availability and where latency is not a primary consideration.

Unstructured data delivery

With the unprecedented growth in new digital information, use cases have emerged that enable organizations to store and distribute limitless data. A distributed and decentralized storage architecture along with an Object Storage interface enables enterprises to deliver data to their users across the globe as never before. These use cases include content repository, storage-as-a-service, enterprise collaboration, backup, and archive.

Content repository storage

The IBM Cloud Object Storage System[™] provides the most reliable, scalable platform for your business critical data.

Effectively store and secure valuable content

Consumers access content from different locations worldwide, making it a business priority to secure irreplaceable originals at scale. IBM Cloud Object Storage System[™] content repository solutions deliver data availability at petabyte and beyond scalability. Its easily scalable system delivers carrier-grade security for a single copy of original content before dispersing it geographically. The IBM Cloud Object Storage System[™] technology ensures data integrity from start to finish. Organizations need content storage that can be distributed across their infrastructure to effectively store and distribute content. IBM Cloud Object Storage System[™] software storage provides a high-availability environment, long-term file integrity and access, and authentication enforcement. Whether organizations have less active, fixed, or frequently accessed content that users are collaborating on, IBM Cloud Object Storage System[™] offers a secure, reliable, and cost-effective approach. A shared content storage repository can be securely accessed by people inside or outside of an organization, enabling collaboration across geographies.

Case Study of the Miami Marlins

Problem

Like most Major League Baseball teams, video means practically everything to the Miami Marlins. Video photographers capture every pitch and every hit of every player at every game all year long. Those recordings prove critical for coaches. As a result, the team collects a large amount of unstructured video data to store, secure, and make available to their coaches and employees.

Solution

The IBM Cloud Object Storage SystemTM is ideal for cloud storage, especially for archiving and backup. Its benefits include massive scalability, geographic independence, multi-tenant features, and the ability to use proprietary, off-the-shelf technology, which provides more cost savings. The Marlins placed IBM Cloud Object Storage SystemTM software in the back end of their system, with a controller as the local interface to the global file system for seamless data access. They immediately noticed the difference that it made to the IT staff. The IBM Cloud Object Storage

System[™] solution eliminated 30 percent of IT professionals' time to serve as a backup administrator and coordinate all the replication routines.

Results

IBM Cloud Object Storage System[™] technology gave the Marlins the ability to have cross-site access to data from each of its data sites without sacrificing security. The IT security chief can keep all data up to date, backed up and replicated across all of their sites. The coaches, scouts, and trainers have access to the content via the cloud by logging in and finding what they need to retrieve. In addition to enhanced data security and seamless data colocation, the Marlins value the systems' transparency to the user. Based on the successes the team experienced with IBM Cloud Object Storage System[™] technology, the Marlins are expanding the cloud storage system into their affiliates in the minor leagues.

Storage-as-a-service

The IBM Cloud Object Storage System[™] software helps deliver new levels of storage capacity and availability with carrier-grade data security to a company's user base.

Delivering storage capacity and availability

Being able to sell capacity to customers on a centralized infrastructure is a must for service providers and large enterprises. The IBM Cloud Object Storage System[™] helps these organizations implement storage-as-a-service solutions that consolidate users and customers onto a single platform. The IBM Cloud Object Storage System[™] helps streamline management and efficiently scale storage to meet customer demands. With secure multi-tenancy, zero-touch encryption and robust management APIs, IT can build a storage offering that is as scalable and reliable as it is easy to manage and cost-efficient.

Case study of Solid-State Solutions (S3)

Problem

S3, a leading data storage integrator, decided to expand their managed and hosted data storage services. Aware that many of their potential customers were operating in hyper-growth environments, they realized that the platform they were reselling fell short. To serve these customers and expand their business, S3 needed a new storage-as-a-service offering.

Solution

The IBM Cloud Object Storage System[™] reliably stores a large amount of data at low cost, making it suited for the cloud. After a thorough evaluation, S3 selected the IBM Cloud Object Storage System[™] platform to provide a competitive data storage-as-a service solution that is more cost efficient and easier to deploy than public cloud providers.

Results

The IBM Cloud Object Storage System[™] platform allowed S3 to provide storage services with data reliability that is greater than 9 nines (99.9999999%) while still competing with public cloud pricing. The simple pricing model is attractive to potential S3 customers in hyper-growth environments. The easy-to-manage IBM Cloud Object Storage System[™] platform also enables S3 to provide customers with rich data and insights such as performance reports, capacity consumption reports, and technology updates that are not typically available from public cloud offerings.

Enterprise collaboration

The IBM Cloud Object Storage System[™] provides secure, distributed access to valuable content, making it easier to enable workplace productivity across the globe.

Collaboration and productivity

Today's workforce is constantly on the move, with global businesses. To be successful, employees need seamless access to mission-critical data from anywhere, at any time. The IBM Cloud Object Storage System[™] delivers a data hub that allows business to provide global access to data. By simultaneously securing it on-premises with zero-touch encryption, they provide safe, distributed data access that enables enterprise collaboration and improves productivity.

Case study of Hogarth

Problem

Hogarth, a leading global marketing and implementation agency, produces advertising and marketing communications for clients across all media and all languages. The company was growing rapidly, with each campaign consuming much storage space in the production environment. Their IT team was looking for a storage solution that would cost-effectively address their expanding amount of unstructured data. They also needed a solution that would enable their international workforce to collaborate globally and without interruption.

Solution

Optimized for storing high volumes of data-driven content, the IBM Cloud Object Storage System[™] solution met Hogarth's requirements and was implemented in all of their main studios. With a single addressable global namespace, the IBM Cloud Object Storage System[™] delivers a unified, single point of management and access that can scale beyond the limits of traditional centralized metadata servers. All employees can now write to it and any of the end offices can pull the data back up to their systems as needed. With most of the Hogarth production studios connected, productivity increased across the company.

Results

The IBM Cloud Object Storage System[™] enables the Hogarth staff to archive much more aggressively and limits the amount of expensive production storage they use. The system also allows Hogarth staff to collaborate more easily between offices and gives their clients easy, safe, reliable access to the assets they need to ensure regulatory compliance.

Backup

The IBM Cloud Object Storage System[™] provides scalable backup and always-on data availability for dependable recovery and security at an 80% lower infrastructure cost.

Cost-effective, secure, and accessible storage

It is a challenging task for IT to collect and back up data from diverse application servers and user machines. Storing this data long term and at scale is even more difficult. The IBM Cloud Object Storage System[™] backup solutions provide a cost-effective, easily accessible storage platform for long-term data protection. SecureSlice[™] zero-touch encryption secures data before cost-effectively distributing it across multiple sites, ensuring long-term, bit-perfect protection at scale. The IBM Cloud Object Storage System[™] software platform enables fast access to data after it is backed up, speeding business recovery time after a disaster.

Case study of a major retailer

Problem

For a major retailer, the steady production of unstructured data, including videos, photos, and more, was increasing at a phenomenal rate. At the same time, IT storage platforms were struggling to scale without dramatically increasing in price and decreasing in reliability.

Solution

The IBM Cloud Object Storage System[™] enables business to efficiently store, manage, and access data at petabyte scale and beyond. Using erasure coding, a type of forward error correction, the

IBM Cloud Object Storage System[™] solution offers far higher data resilience than RAID and requires far less storage capacity than standard Object Storage solutions. IBM Cloud Object Storage System[™] demonstrated the scalability, efficiency, security and simplicity of its object software storage, and the retailer's upper management quickly got on board.

Results

The retailer is on track to hit roughly one-half petabyte, or 500 terabytes, of storage by the end of 2015. The IBM Cloud Object Storage System[™] backup solution stores mountains of video and visuals easily and economically. The major driver for the move to the IBM Cloud Object Storage System[™] was cost efficiency, and the savings were immediate. The retailer realized a 50 percent savings (per byte) compared to its previous storage vendor. The IBM Cloud Object Storage System[™] method of slicing and structuring data makes it impossible to damage or steal critical information, ensuring the data is secure.

Active archive

The IBM Cloud Object Storage System[™] keeps content accessible with a scalable, reliable, and secure long-term data archive.

Scalable, reliable, and secure archive storage

Many organizations are seeking an archival solution that provides their users with immediate access to their data. The IBM Cloud Object Storage System[™] offers an archive storage solution that combines limitless availability with the highest levels of data integrity and confidentiality.

Proactive error correction is crucial to keeping a long-term archive working correctly, since information isn't as frequently read. The IBM Cloud Object Storage System[™] employs an intelligent background process that scans storage nodes, checking for and correcting errors.

Deployments of IBM Cloud Object Storage System[™] solution can span multiple data centers. An archive that is distributed across multiple offsite locations protects data against a single location failure or catastrophic disaster, making it securely accessible for long-term retention. Because it is not tied to a specific server or storage device, the data is automatically reconstituted as new storage nodes are installed in the system. The IBM Cloud Object Storage System[™] archive solution enables organizations to meet their compliance requirements and long-term preservation goals.

Case study of Shutterfly

Problem

With an active archive of billions of photos in constant motion, and continuously growing, Shutterfly faced significant challenges to make sure it could keep pace with its customers' needs and maintain the same levels of performance, availability, and reliability. With petabytes of raw storage and a double-digit growth rate, the cost to store this data was growing rapidly too. Shutterfly had to find alternative ways to drive down the cost of storage and make it easier to manage.

Solution

The IBM Cloud Object Storage System[™] software storage approach to ingesting and storing data inherently solved Shutterfly's most critical issues (eliminate single points of failure and deliver high levels of fault tolerance). An IBM Cloud Object Storage System[™] solution also had a number of properties that made the management of Shutterfly's image archives much simpler. Advanced erasure coding techniques disassociate the performance and reliability of individual components from application level performance and reliability. It allows the company to have continuous availability of its data, making it far less susceptible to hardware and software problems in its storage tier.

Results

Presently Shutterfly has over 150 PB of storage in production and is growing rapidly with a

limitless capacity to scale. The company is realizing significant power consumption and management cost savings across the board. Shutterfly is now able to manage the entire storage platform containing billions of objects and over 150 PB of capacity with only three part-time storage administrators.

Chapter 6. Conclusion

Enterprises that need to store large volumes of unstructured data must look beyond their current storage solutions and evaluate new approaches. Dispersed storage is one such innovative approach for cost-effectively storing large volumes of unstructured data while ensuring security, availability, and reliability.

The Definitive Guide to Dispersed Storage described the features and benefits of dispersed storage in five critical areas:

Availability

Data is always available, regardless of planned, or unplanned downtime.

Scalability

Systems can easily grow from terabytes to petabytes to exabytes.

Security

Data confidentiality is ensured even when multiple drives, servers, containers, or locations are compromised.

Economics

The need for costly replication is eliminated, significantly lowering the total cost of ownership for storage systems at the petabyte level and beyond.

Efficiency

Easily manage tens of petabytes of storage per administrator.

For these reasons, dispersed storage is ideal whenever enterprises need to store large volumes of unstructured data and where latency is not a primary consideration. Enterprises with content storage, active archive, or content distribution needs should evaluate dispersed storage as a new technology option.

Is dispersed storage right for your organization?

Use this checklist to determine whether your organization might benefit from dispersed storage.

- Do you have applications that require long-term retention of data?
- Does the data consist of large, unstructured objects such as images, movies, and documents?
- Do you have 80 usable terabytes or more of this data?
- Do you have requirements for data security, availability, scalability, and cost-effectiveness?
- Do you have the infrastructure that is needed to support dispersed storage?
 - Do you have network connectivity?
 - Do you have high-quality bandwidth (if geo-dispersed)?

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing IBM Corporation North Castle Drive Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing Legal and Intellectual Property Law IBM Japan, Ltd. 19-21, Nihonbashi-Hakozakicho, Chuo-ku Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing IBM Corporation North Castle Drive, MD-NC119 Armonk, NY 10504-1785 US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

Accesser[®], Cleversafe[®], ClevOS[™], Dispersed Storage[®], dsNet[®], IBM Cloud Object Storage Accesser[®], IBM Cloud Object Storage Dedicated[™], IBM Cloud Object Storage Insight[™], IBM Cloud Object Storage Manager[™], IBM Cloud Object Storage Slicestor[®], IBM Cloud Object Storage Standard[™], IBM Cloud Object Storage System[™], IBM Cloud Object Storage Vault[™], SecureSlice[™], and Slicestor[®] are trademarks or registered trademarks of Cleversafe, an IBM Company and/or International Business Machines Corp.

Other product and service names might be trademarks of IBM or other companies.

Homologation statement

This product may not be certified in your country for connection by any means whatsoever to interfaces of public telecommunications networks. Further certification may be required by law prior to making any such connection. Contact an IBM representative or reseller for any questions.

IBM.®

Printed in USA