

ADVANCED TECHNOLOGY GROUP (ATG)



Accelerate with ATG Webinar

IBM Cloud Object Storage System

Why IBM Cloud Object Storage System and, ... why now?

John Shubeck – ATG Storage Technical Specialist

Date: June 20, 2024



About the Presenter



John Shubeck is an information technology professional with over 42 years of industry experience spanning both the customer and technology provider experience. John is currently serving as a Senior Storage Technical Specialist on IBM Object Storage platforms across all market segments in the Americas.

Introducing our panelists




Shu Mookerjee is a Level 2 Certified Technical Specialist with over twenty years at IBM, working in a variety of roles including sales, management and technology. For the last decade, he has focused exclusively on storage and has been the co-author of four (4) Redbooks. Currently, Shu is part of the Advanced Technology Group where he provides education, technical guidance, Proofs of Concept and Proofs of Technology to IBMers, business partners and clients.

Introducing our panelists



Jerrod Carr is an IBM Principal Storage Technical Specialist in IBM Storage Solutions. Jerrod Carr has been in the Storage industry for over 21 years selling hardware and software for various large technology companies. With beginnings in the Cleversafe IBM team for 8 years providing expertise in Cloud Object Storage, the last 3 years working on the Americas SWAT team as a Senior Storage Specialist providing unstructured data experience to the various markets.



Todd Johnston is a proven applied technologist bringing decades in customer advocacy, technology implementations, service provider enablement, sales engineering, and solution architecture. Todd is especially focused on Ceph innovation and early adoption, technical agility, and disruptive technologies. Todd's expertise in lies storage software, virtualization, and container and hybrid cloud infrastructures. Todd joins us from the IBM SWAT practice in our cross functional team.

Summary of topics



- Overview and review of IBM Cloud Object Storage
- What are customers asking for with object storage
- Key reasons to deploy IBM Cloud Object Storage
- IBM Storage Ready Node update
- IBM Cloud Object Storage use cases
- IBM Cloud Object Storage differentiators
- What about swim lanes?

What is IBM Cloud Object Storage?

The new cloud native architecture on-premises holds the promise of traditional data center security and reliability with cloud agility and scalability. IBM Cloud Object Storage System is a framework that powers the world's largest data repositories, delivering data security, availability, and durability at significantly lower cost than traditional storage and cloud storage. Join the Advanced Technology Group for a closer look at IBM Cloud Object Storage System with emphases on the recent innovations, use cases, and solutions for the cloud native on-premises architecture that answers the question, “Why now?”.

Why IBM Cloud Object Storage?

“IBM Cloud Object Storage is unequaled in the industry when it comes to data resiliency and large deployments. So, customers with multiple data centers, three or more data centers, that want geographically dispersed erasure coding, IBM Cloud Object Storage is absolutely the right answer.”

IBM Cloud Object Storage offerings

What is the difference?



IBM Cloud Object Storage System *Object storage on premises*

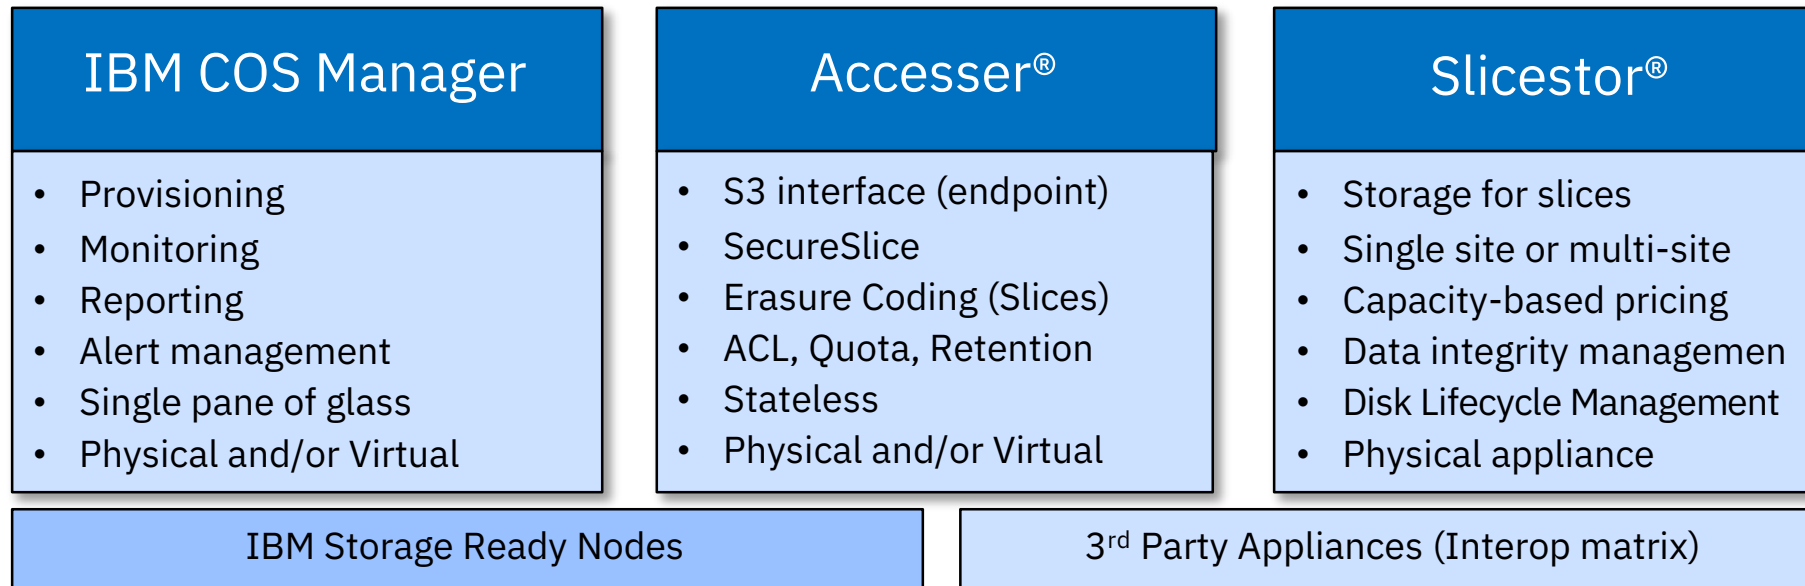
- Customer managed
- Exabyte capacity
- Data center networking
- Private enterprise cloud
- Corporate IAM
- Bucket explorer on Github
- Data center security perimeter



IBM Cloud Object Storage Service *Object storage in the cloud*

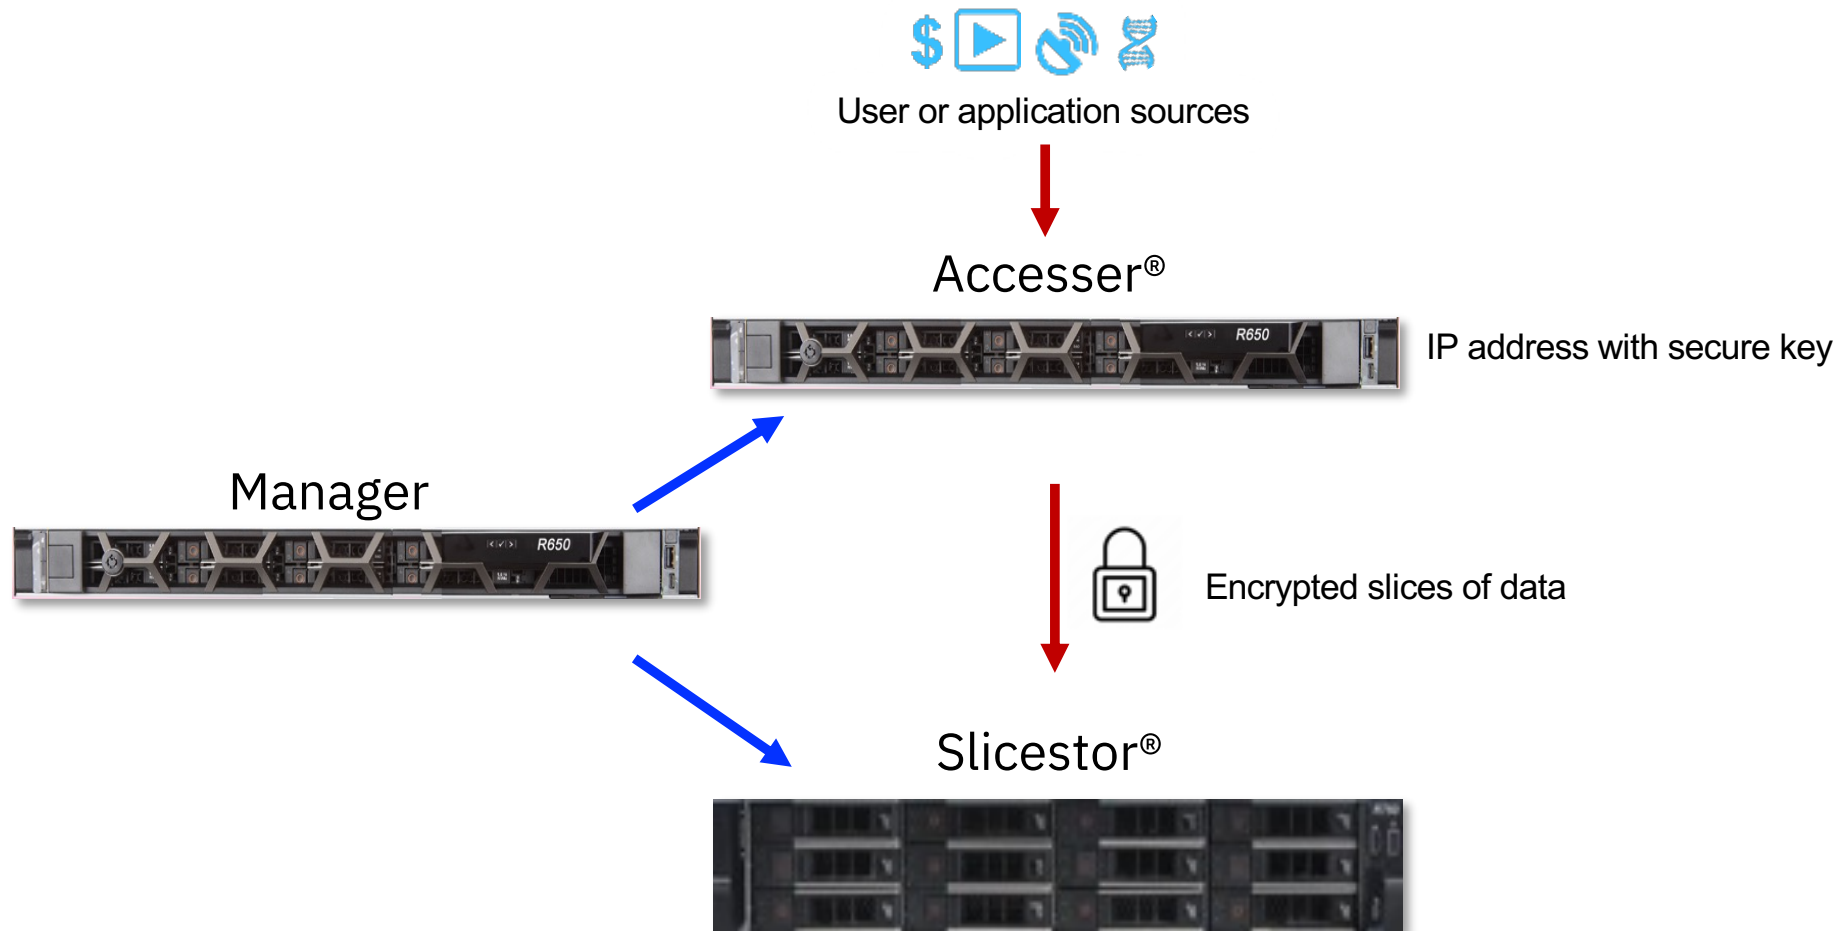
- IBM managed
- Unlimited capacity
- Unlimited bandwidth
- Cloud IAM integration
- Cross region replication
- Bucket explorer
- GLACIER storage class
- Automatic Smart Tiering
- Cloud key management (KMS)

IBM COS system components

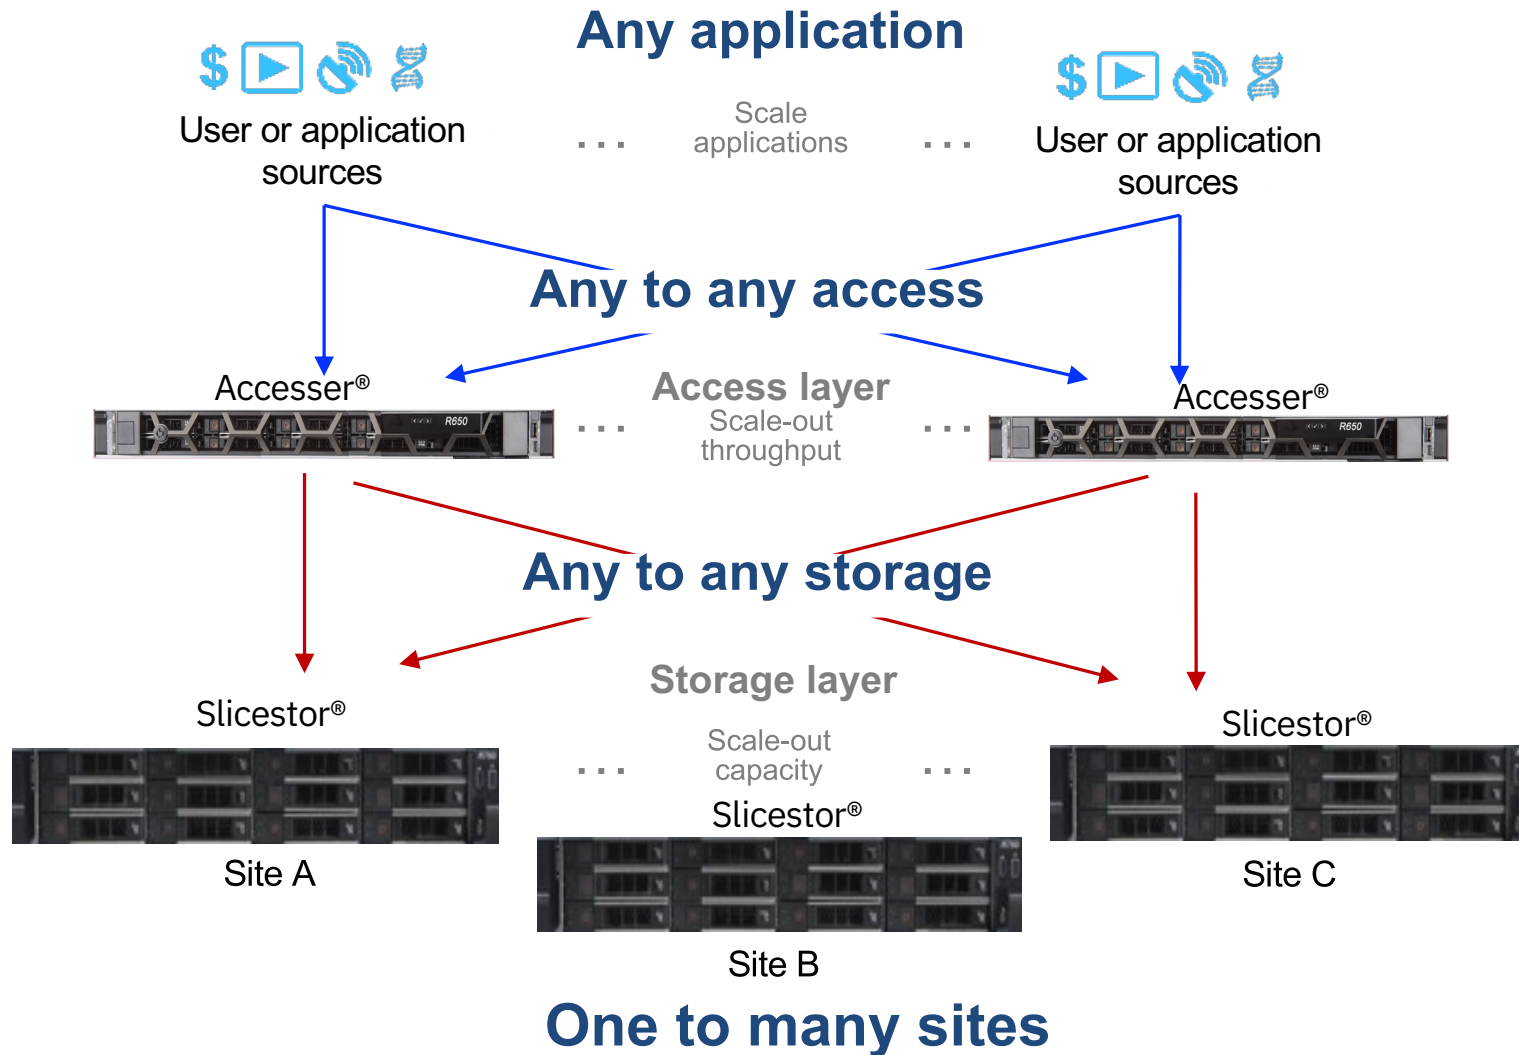


 Cloud Object Storage

IBM COS system components



IBM COS system deployment



IBM COS fundamental – The economy of Geo-dispersed data protection

Traditional RAID Storage




IBM Cloud Object Storage requires less than half the storage and 70% lower TCO.

You can lose some number of slices due to failure or disaster, and still quickly recover 100% of your data.

Slices are distributed geographically for durability and availability.

IBM Cloud Object Storage



Align with customer requirements



Compliance Archive (Retention Vault or S3 Object Lock)



Extreme availability (Multiple Managers)



Security (Multi Factor Authentication)



Data and AI (S3 Select)

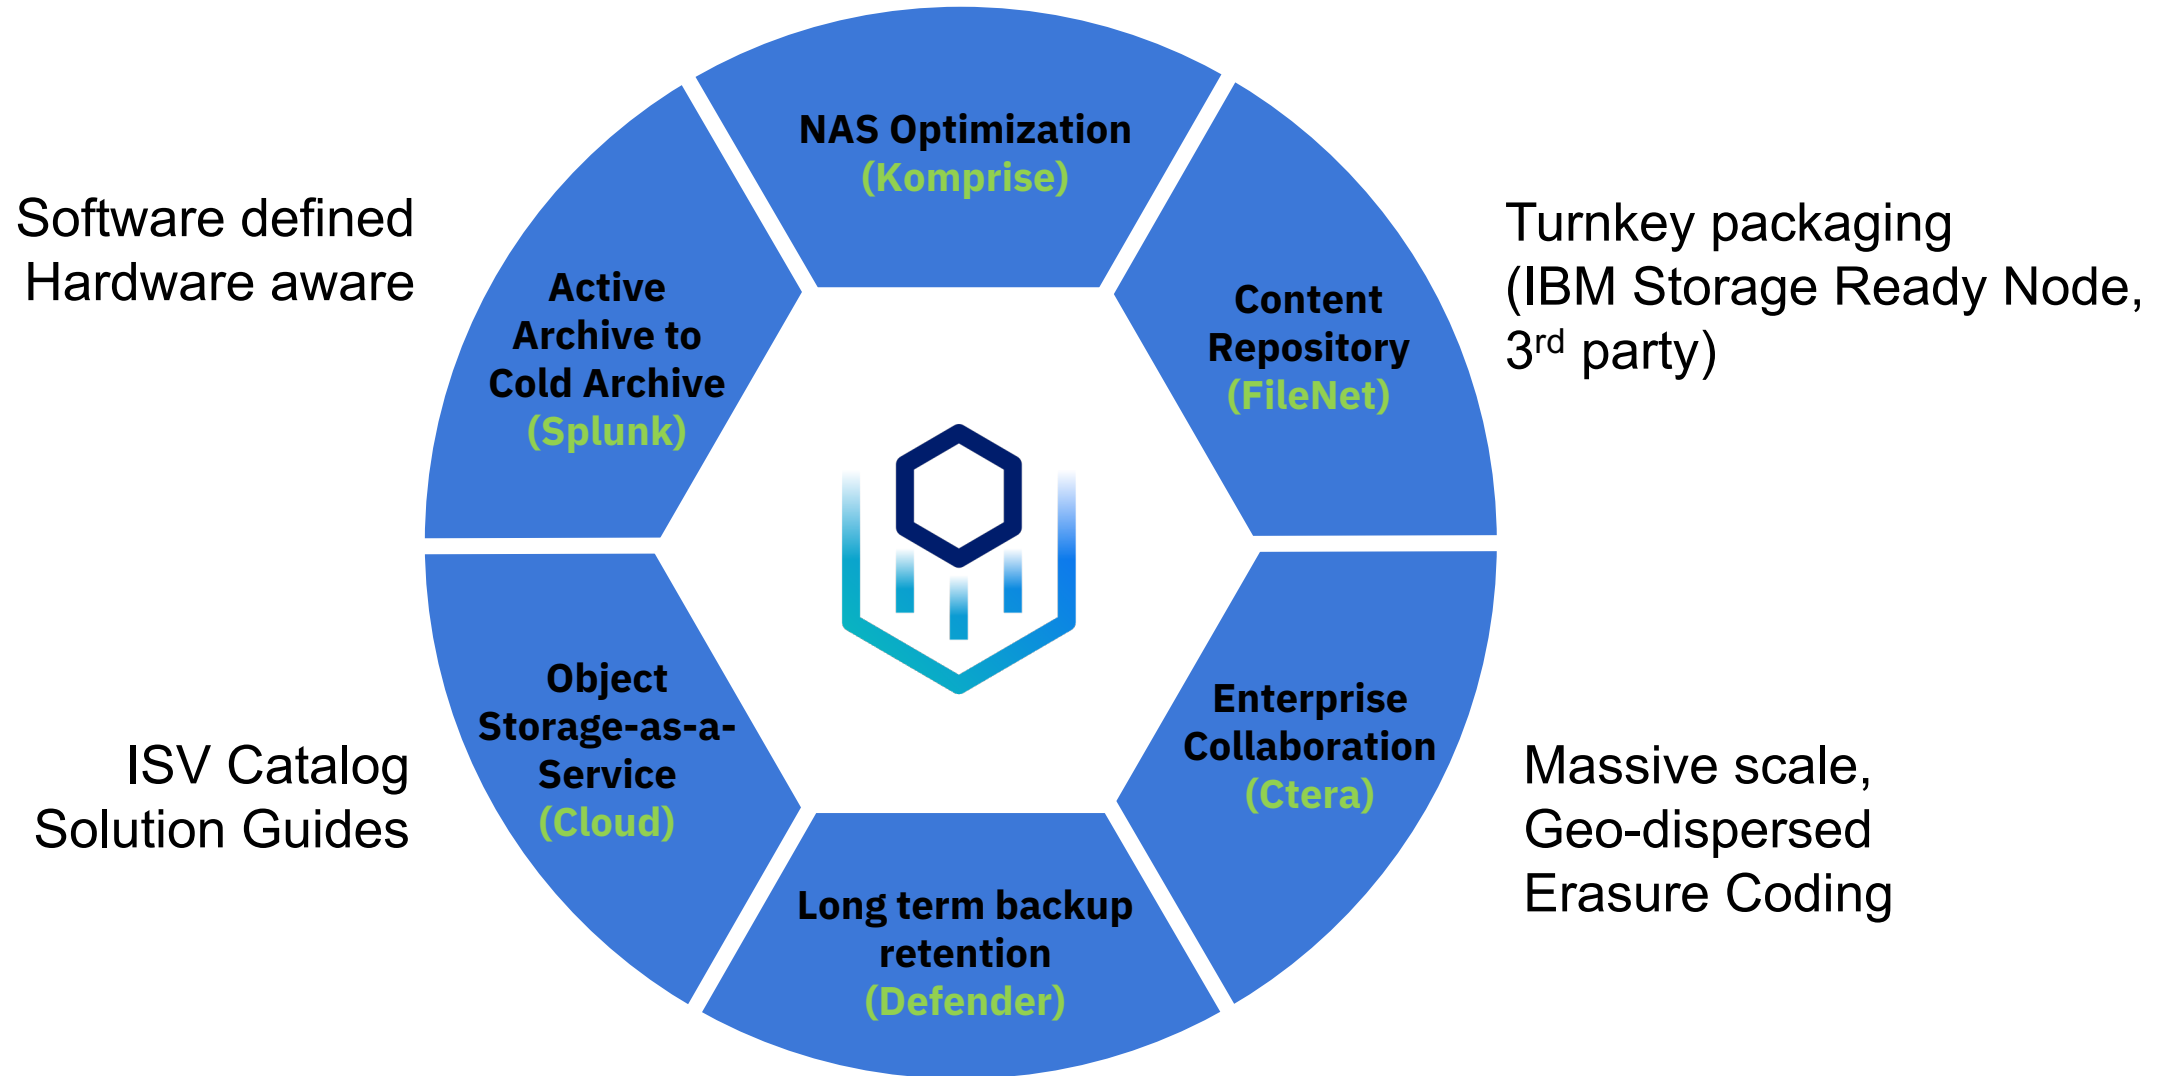


Cyber resiliency (S3 Versioning, S3 Object Lock)



Simplicity (IBM Storage Ready Nodes)

IBM Cloud Object Storage attributes use cases



Reason no. 1 - ease of deployment and turnkey packaging




IBM Storage COS buying and delivery options


1 Software only



IBM Storage COS



2 Ready Node HDD




IBM Storage COS System

3 Ready Node NVMe

Future - *exploring...*

4 Appliance



IBM Storage Ready Node highlights



Storage Ready nodes

IBM Storage Ready Nodes are a simple, cost-effective way for clients to deploy IBM Storage software on industry-standard server hardware.

Building blocks

A simple building block approach to software defined storage and infrastructure, all supported by IBM.

Scalability

Start with a cluster configuration that meets current needs and then scale with additional nodes when required.

Online expansion!

Use cases

Suitable for all IBM Cloud Object Storage use cases.

- NAS optimization
- Content repository
- Enterprise collaboration
- Backup retention
- SaaS
- Active archive

IBM tested software

IBM qualified and supported hardware for running IBM's object storage platforms.

Configuration options tested and certified.

IBM product integration

Investment protection, IBM Storage Ready Nodes can be integrated into existing IBM object storage configurations.

Hardware overview summary matrix

IBM Storage Ready Node for Cloud Object Storage

	Manager / Accesser / Controller	Slicestor	
	R650	R750	ME484
CPU	2 x Intel® Xeon® Silver 4314 2.4G	2 x Intel® Xeon® Silver 4314 2.4G	
Memory	16 x 16GB RDIMM (256GB)	16 x 16GB RDIMM (256GB)	
OS Disks	2 x 960GB SSD SATA 2.5"	2 x 240GB SSD SATA 2.5"	
Data Disks Quantities		12	28, 56, 84
Data Disk Sizes (TB)		8, 12, 16, 20	8, 12, 16, 20
RAW Capacity (TB)		96, 144, 192, 240	672, 1008, 1344, 1680
Network	2 x 10GbE (SFP+)	2 x 10GbE (SFP+)	
Dimensions	1.68" H x 18.97" W x 28.9" D (1U Rack Height)	3.41" H x 18.97" W x 29.85" D (2U Rack Height)	8.75" H x 19.01" W x 38.31" D (5U Rack Height)
Software Support	ClevOS 3.17 Release	ClevOS 3.17 Release	ClevOS 3.17 Release
Additional Notes		Integrated Controller	R650 Controller Required

Hardware overview summary matrix - continued

IBM Storage Ready Node with NVMe			
	Slicestor Nodes	OSD Nodes	
	Ready Node 750	Ready Node with NVMe	Ready Node next
CPU	2 x Intel® Xeon® Silver 4314 2.4G	2 x Intel® Xeon® Gold 6438N Turbo	
Memory	16 x 16GB RDIMM (256GB)	16 x 32GB RDIMM (512GB)	
OS Disks	2 x 240GB SSD SATA 2.5"	2 x 480GB BOSS-N1	
Device Count	12 HDD	8, 16, 24 NVMe	
Device Sizes (TB)	8, 12, 16, 20	3.84TB, 7.68, 15.36TB	
RAW Capacity (TB)	96, 144, 192, 240	122, 244, 366 (24 devices)	
Network	2 x 10GbE (SFP+)	2 x 100GbE (QSFP28)	
Dimensions	3.41" H x 18.97" W x 29.85" D (2U Rack Height)	3.41" H x 18.97" W x 29.85" D (2U Rack Height)	
Software Support	IBM Storage Ceph 6.1	IBM Storage Ceph 7.0	

R650 Manager / Accesser / Controller

Descriptions	Specifications
CPU	2 x Intel® Xeon® Silver 4314 2.4G
Memory	16 x 16GB RDIMM (256GB)
OS Disks	2 x 960GB SSD SATA 2.5"
Data Disks	N/A
Data Disk Quantities	N/A
Network	2 x 10GbE (SFP+ Optical Transceivers Included)
Dimensions	1.7" H x 18.97" W x 31.85" D (1U Rack Height)
Software Support	ClevOS 3.17 Release



R750 Slicestor

Descriptions	Specifications
CPU	2 x Intel® Xeon® Silver 4314 2.4G
Memory	16 x 16GB RDIMM (256GB)
OS Disks	2 x 240GB SSD SATA 2.5"
Data Disks	8TB / 12TB / 16TB / 20TB SATA
Data Disk Quantities	12 Disks Fully Populated Only
Network	2 x 10GbE (SFP+ Optical Transceivers Included)
Dimensions	3.41" H x 18.97" W x 29.85" D (2U Rack Height)
Software Support	ClevOS 3.17 Release



ME484 Slicestor

Descriptions	Specifications
CPU	2 x Intel® Xeon® Silver 4314 2.4G (in the R650 Controller)
Memory	16 x 16GB RDIMM (256GB) (in the R650 Controller)
OS Disks	2 x 960GB SSD SATA 2.5" (in the R650 Controller)
Data Disks	8TB / 12TB / 16TB / 20TB SATA
Data Disk Quantities	28 / 56 / 84
Network	IPMI
Dimensions	8.75" H x 19.01" W x 38.31" D (5U Rack Height)
Software Support	ClevOS 3.17 Release



IBM Storage Ready Node with NVMe

Descriptions	Specifications
CPU	2 x Intel® Xeon® Gold 6438N 2G, 32C/64T, 16GT/s, 60M Cache, Turbo, HT (205W) DDR5-4800
Memory	16 x 32GB RDIMM (512GB)
OS Disks	BOSS-N1 controller card + with 2 M.2 480GB (RAID 1)
NVMe Device Density	3.84TB, 7.68, 15.36TB NVMe
NVMe Device Quantity	8, 16, or 24 Disks
Network	2 x 10GbE (SFP+ Optical Transceivers Included)
Network	2 x 100GbE (QSFP28 Optical Transceivers Included)
Dimensions	3.41" H x 18.97" W x 29.85" D (2U Rack Height)
Software Support	ClevOS 3.18.2 and later



Reason no. 2 – Massive scale, Geo-dispersed Erasure Coding



IBM COS fundamental – The economy of Geo-dispersed data protection

Traditional RAID Storage




IBM Cloud Object Storage requires less than half the storage and 70% lower TCO.

You can lose some number of slices due to failure or disaster, and still quickly recover 100% of your data.

Slices are distributed geographically for durability and availability.

IBM Cloud Object Storage



Deployment options across one, two or multiple sites (What if I need to store 1 PB)

Single site most cost effective



Site A

1.7 PB raw

Traditional two site with mirroring or replication



Site A

1.7 PB raw

Site B

1.7 PB raw

3.4 PB raw

Geo-dispersed multiple sites and most deployed by customers



Site A



Site B

1.7 PB raw



Site C

Reason no. 3 – ISV catalog



ISV solutions (Validation, qualification, solution guides)

Big Data / IoT / Object Buckets / Distribution



Remote file access / collaboration



Data protection



Active archive



Reason no. 4 – Software defined and hardware aware




Which would you prefer?

- Without hardware awareness, disks are nothing more than . . . “dumb” disks in a device table

Drive List

Capacity	Device Name	Disk Status	Disk Actions
17.19 GB	sda	ONLINE	
8.58 GB	sdb	ONLINE	
8.58 GB	sdc	ONLINE	
8.58 GB	sdd	ONLINE	
8.58 GB	sde	ONLINE	
8.58 GB	sdf	ONLINE	
8.58 GB	sdg	ONLINE	
8.58 GB	sdh	ONLINE	
8.58 GB	sdi	ONLINE	
8.58 GB	sdj	OFFLINE	Dispose
8.58 GB	sdk	ONLINE	
8.58 GB	sdl	ONLINE	
8.58 GB	sdm	ONLINE	

Device Capacity



Which would you prefer?

- With hardware awareness, the storage layout, geometry, details, and disk state are visible!

Device Diagram

Front View


1	4	7	10
2	5	8	11
3	6	9	12

Bottom

Drive Information and Actions
Drive: main:front:10
Model: ST10000NM0016-1T
Firmware: SNE0
Serial Number: ZA20SZX0
Capacity: 10 TB
Device Name:
Status: UNUSABLE
Drive Action:

[close](#)

Device Capacity

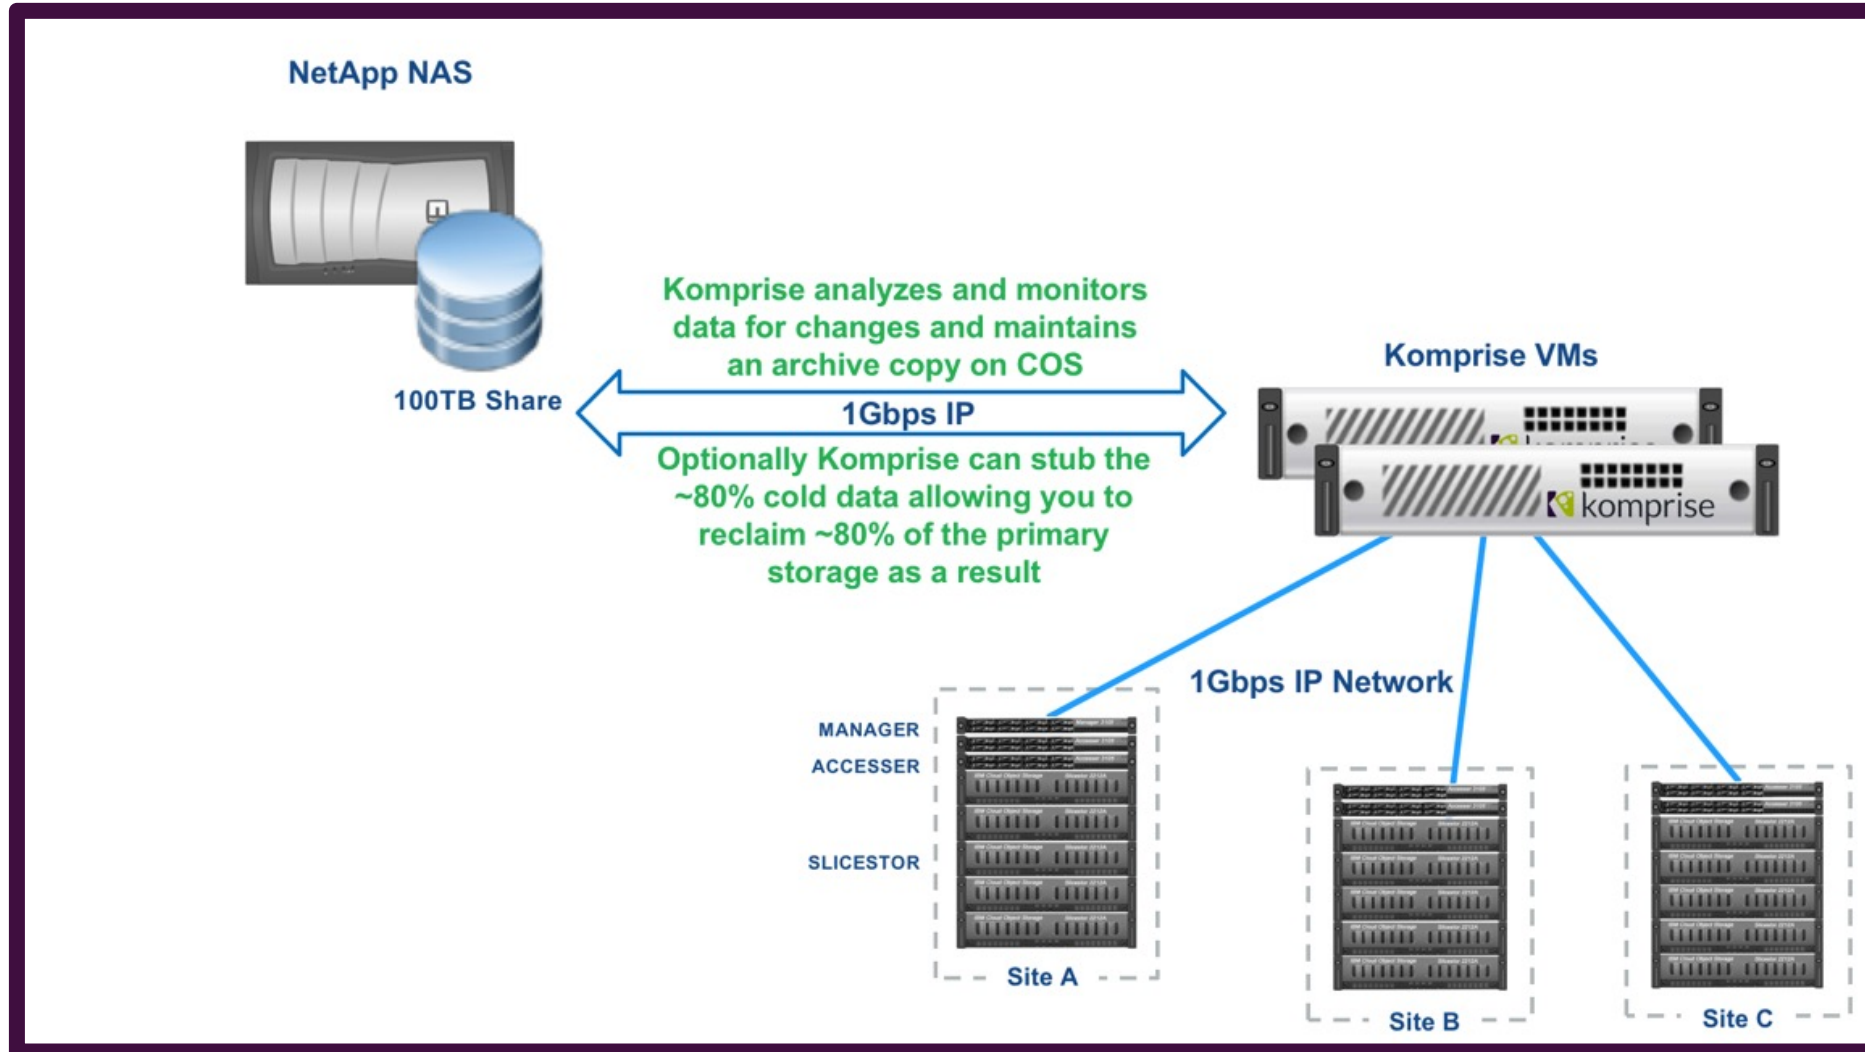


Open Incidents (1)

Object storage use case no. 1 – NAS Optimization



NAS Optimization / AKA “Drain the NAS” . . .

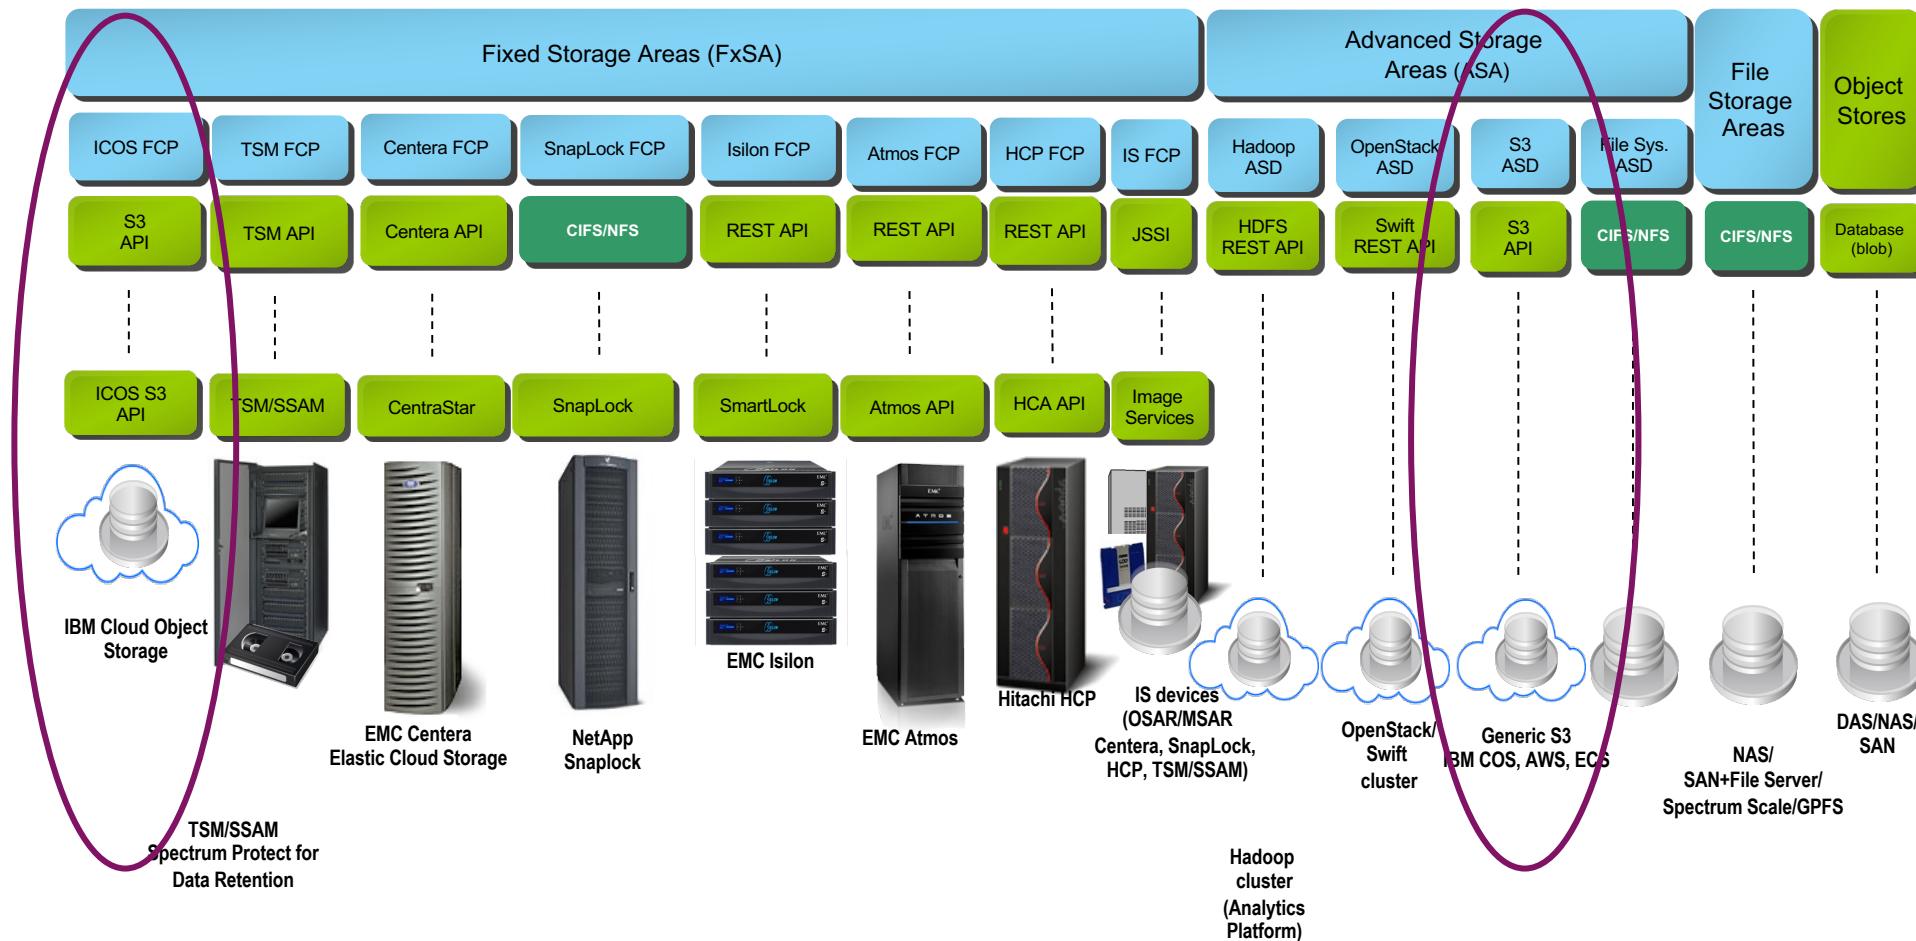


Object storage use case no. 2 – Content repository

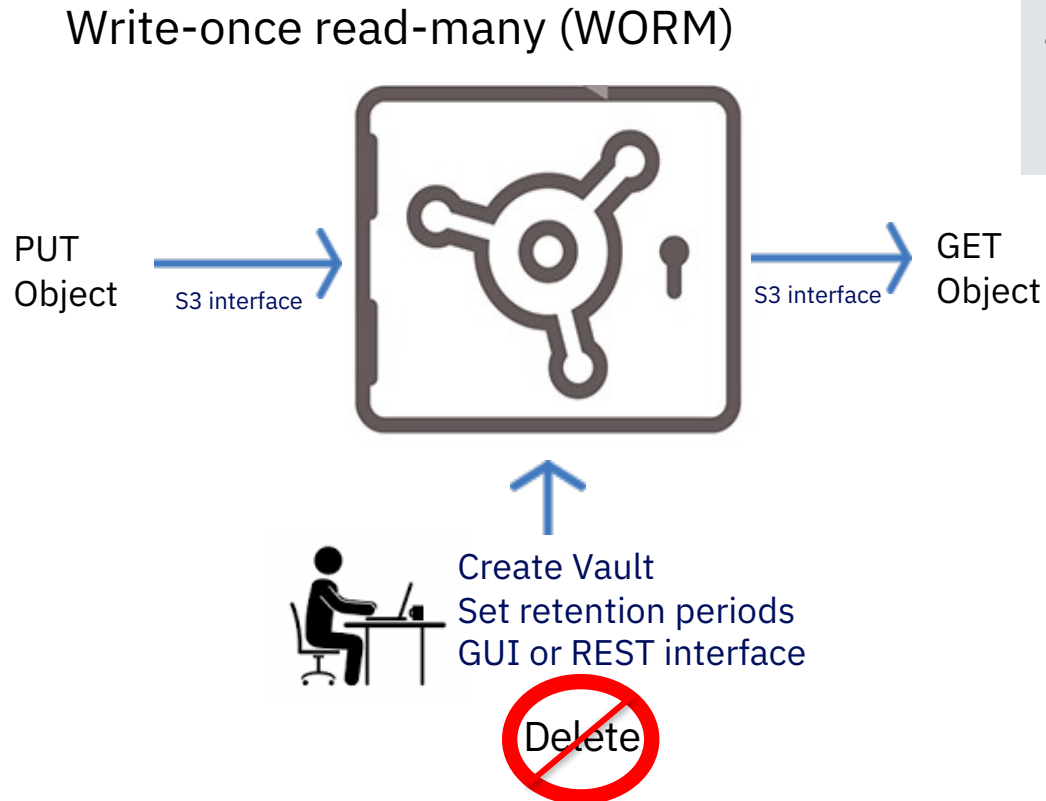


FileNet Content Manager integration

IBM COS is tightly integrated with the FileNet application engine and can be used as an Advanced Storage Area or Fixed Storage Area which supports WORM retention and compliance requirements



Locking down content with retention



- FileNet P8 writes the content and applies retention to each object individually and are write time based on its own retention setting, which then gets propagated to the device.
- FileNet "owns" the content it writes and wants the device to just "enforce" the retention setting it dictates.

- ✓ **Easy to create**
- ✓ **Easy to access**
- ✓ **Unable to delete or modify**

Protection

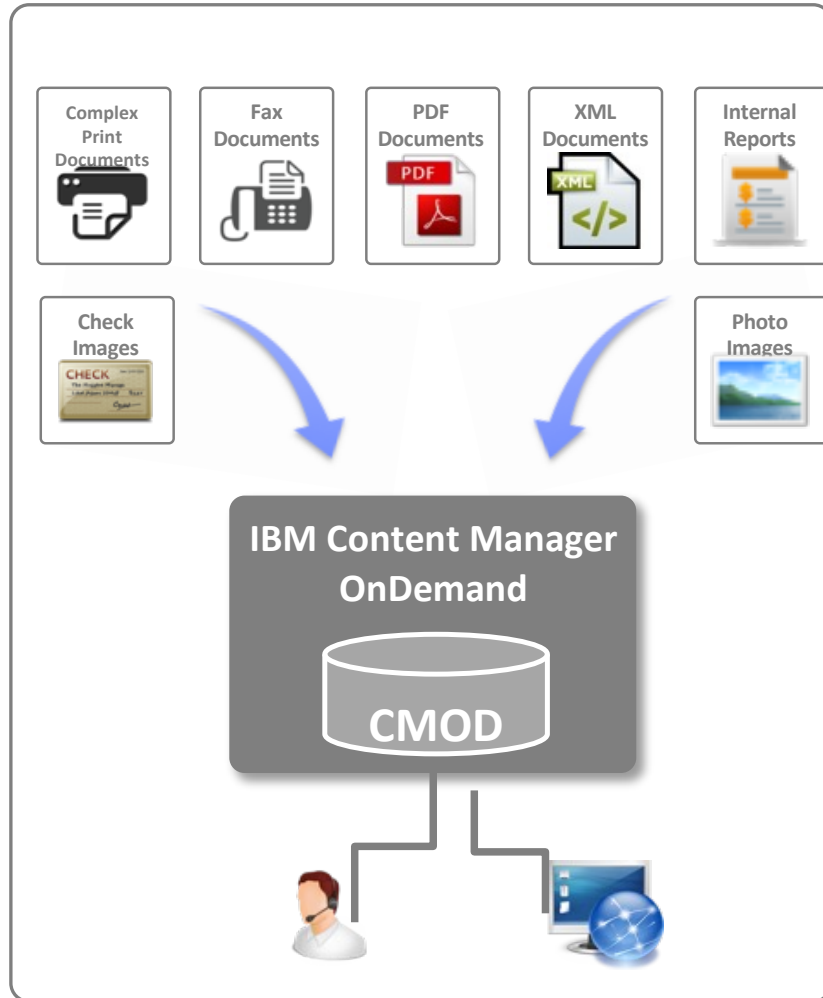
Level:

Data is retained for the default retention duration unless you specify a duration

Data Retention Durations

Default Duration:	<input type="text" value="2"/>	<input type="text" value="Years"/>	<input type="text" value="730"/>
Minimum Duration:	<input type="text" value="1"/>	<input type="text" value="Years"/>	<input type="text" value="365"/>
Maximum Duration:	<input type="text" value="20"/>	<input type="text" value="Years"/>	<input type="text" value="7,305"/>

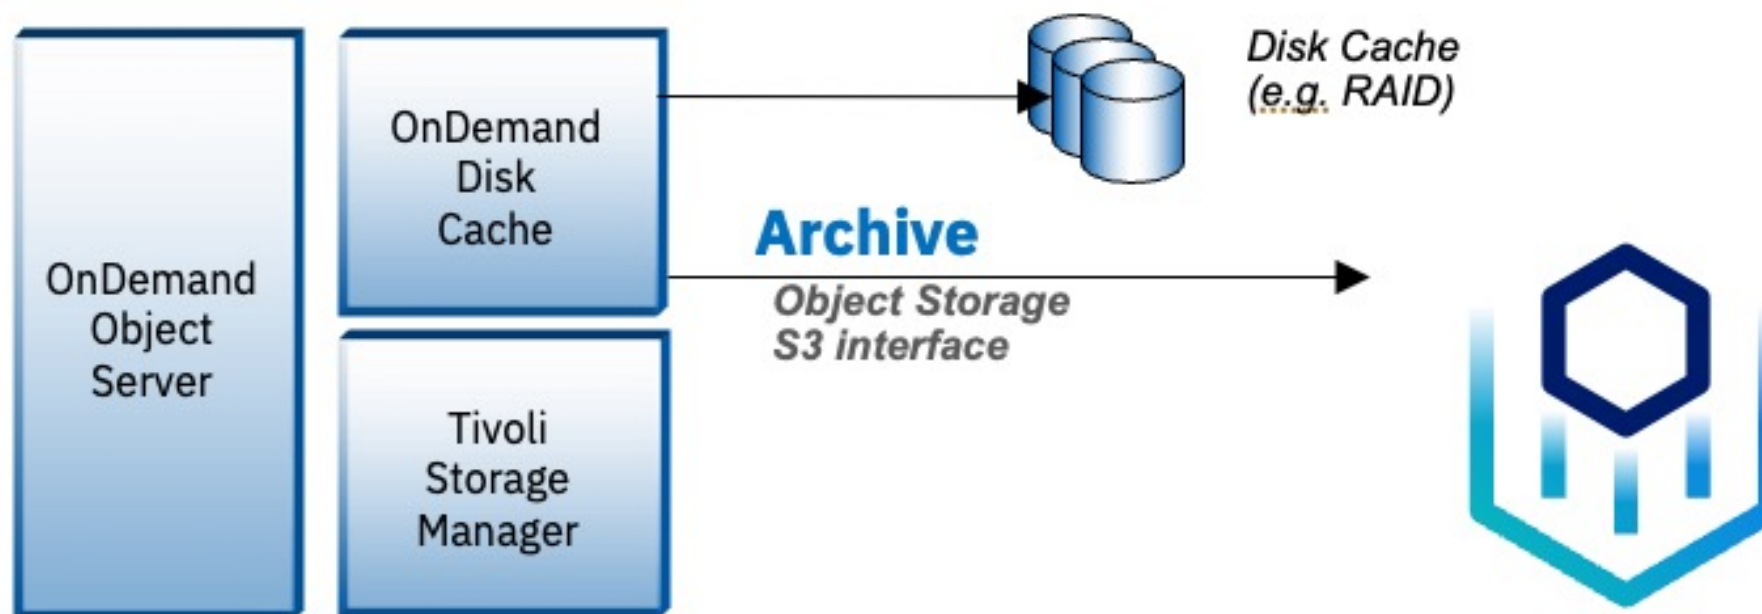
Content Manager on Demand (CMOD) repository



- Enterprise Archive and Customer Communication Management
- Archives customer correspondence like statements, invoices, policies, confirmations, internal reports, all types of data
- Provides immediate access to statements by customers through web portal
- Up to 90% storage savings through advanced compression
- Highly scalable and very efficient
- Easy to install and configure, highly reliable, lower cost to operate than most other repositories
- Available on Premise, On Cloud or Hybrid

IBM COS for CMOD storage management

- COS has direct S3 integration with CMOD and supports retention compliance
- No 3rd party application is required



Object storage use case no. 3 – Enterprise collaboration

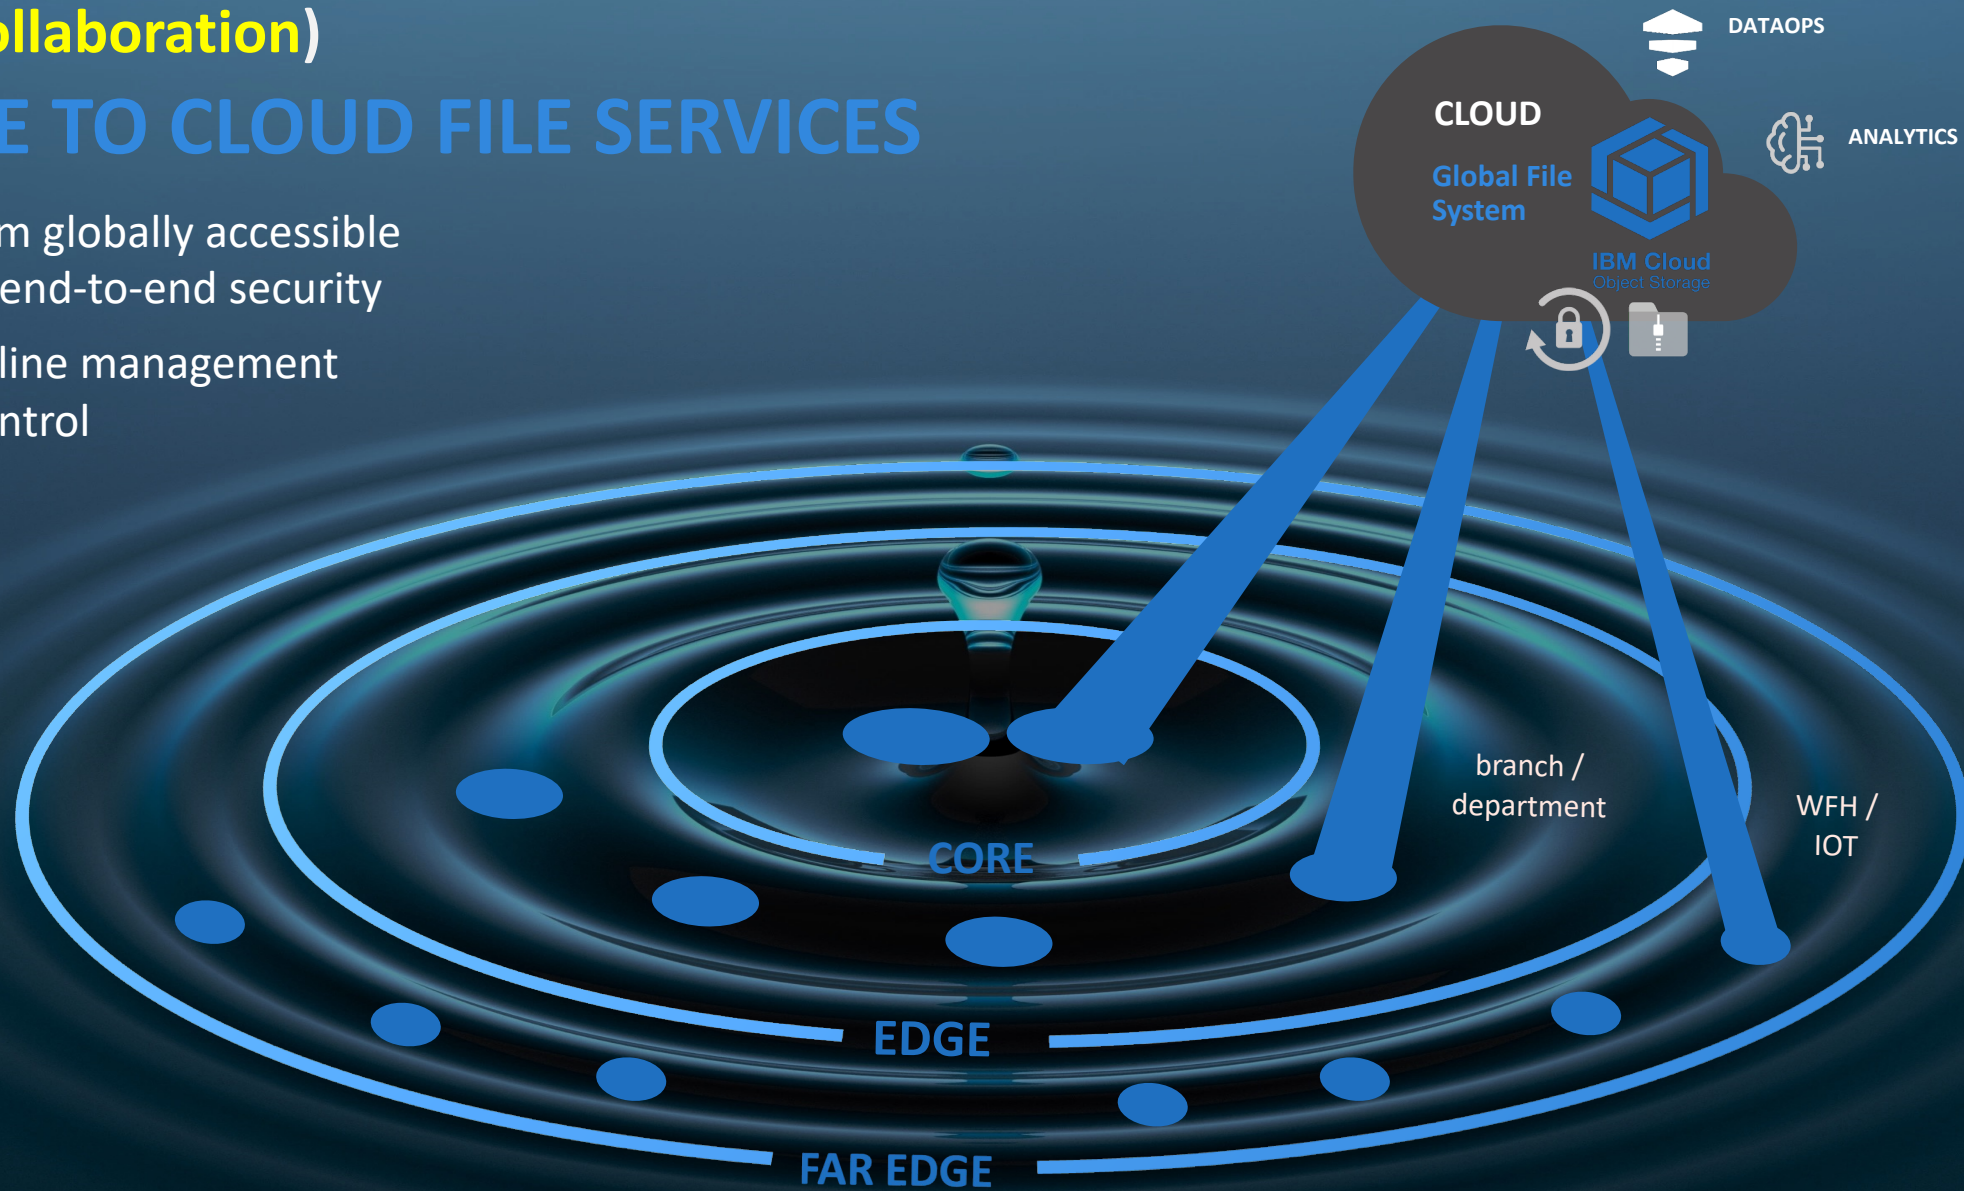


CTERA (Global Collaboration)

SECURE EDGE TO CLOUD FILE SERVICES

Object-based file system globally accessible
with local caching and end-to-end security

Cloud-native data pipeline management
with full insight and control

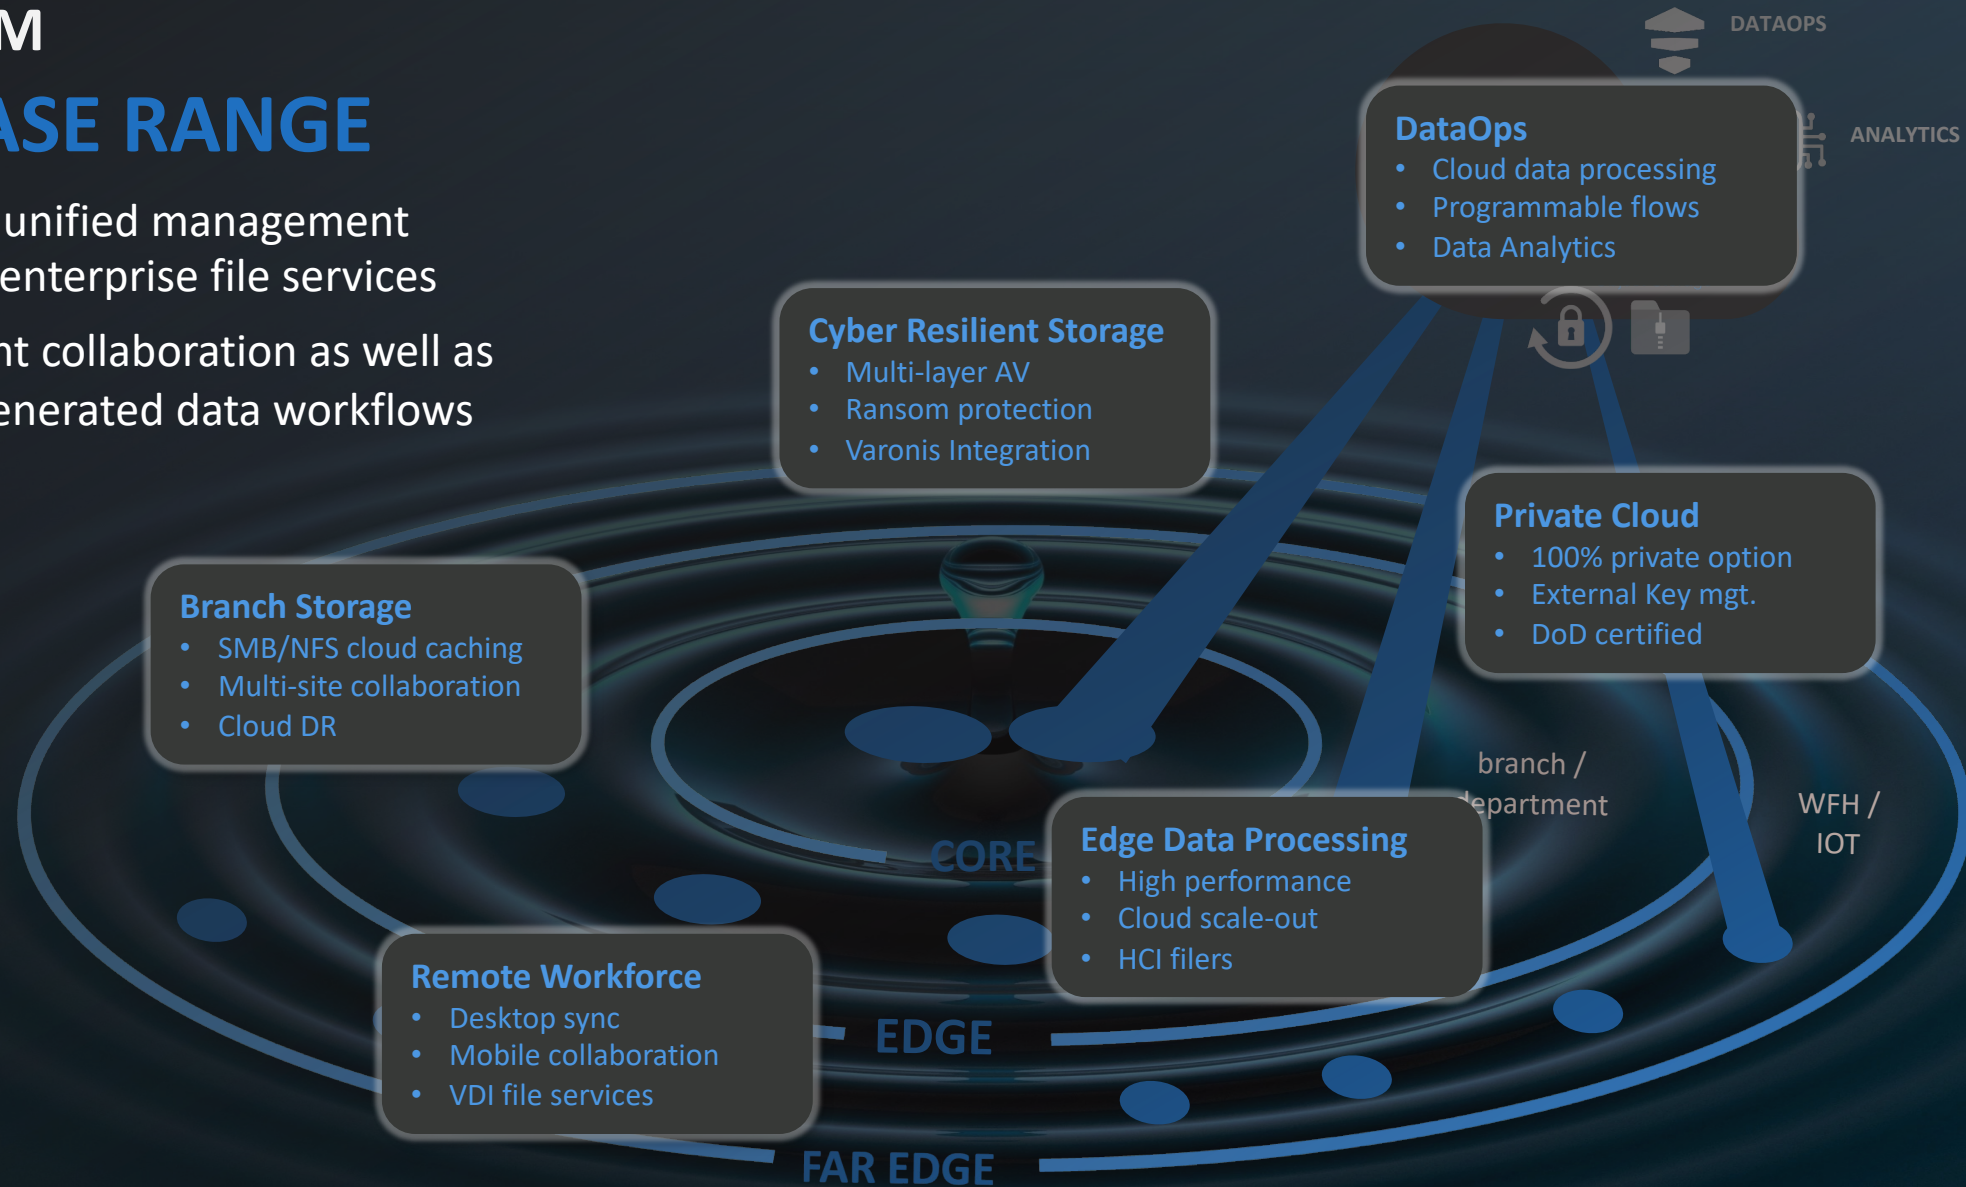


SINGLE PLATFORM

WIDE USE CASE RANGE

Single namespace and unified management across a continuum of enterprise file services

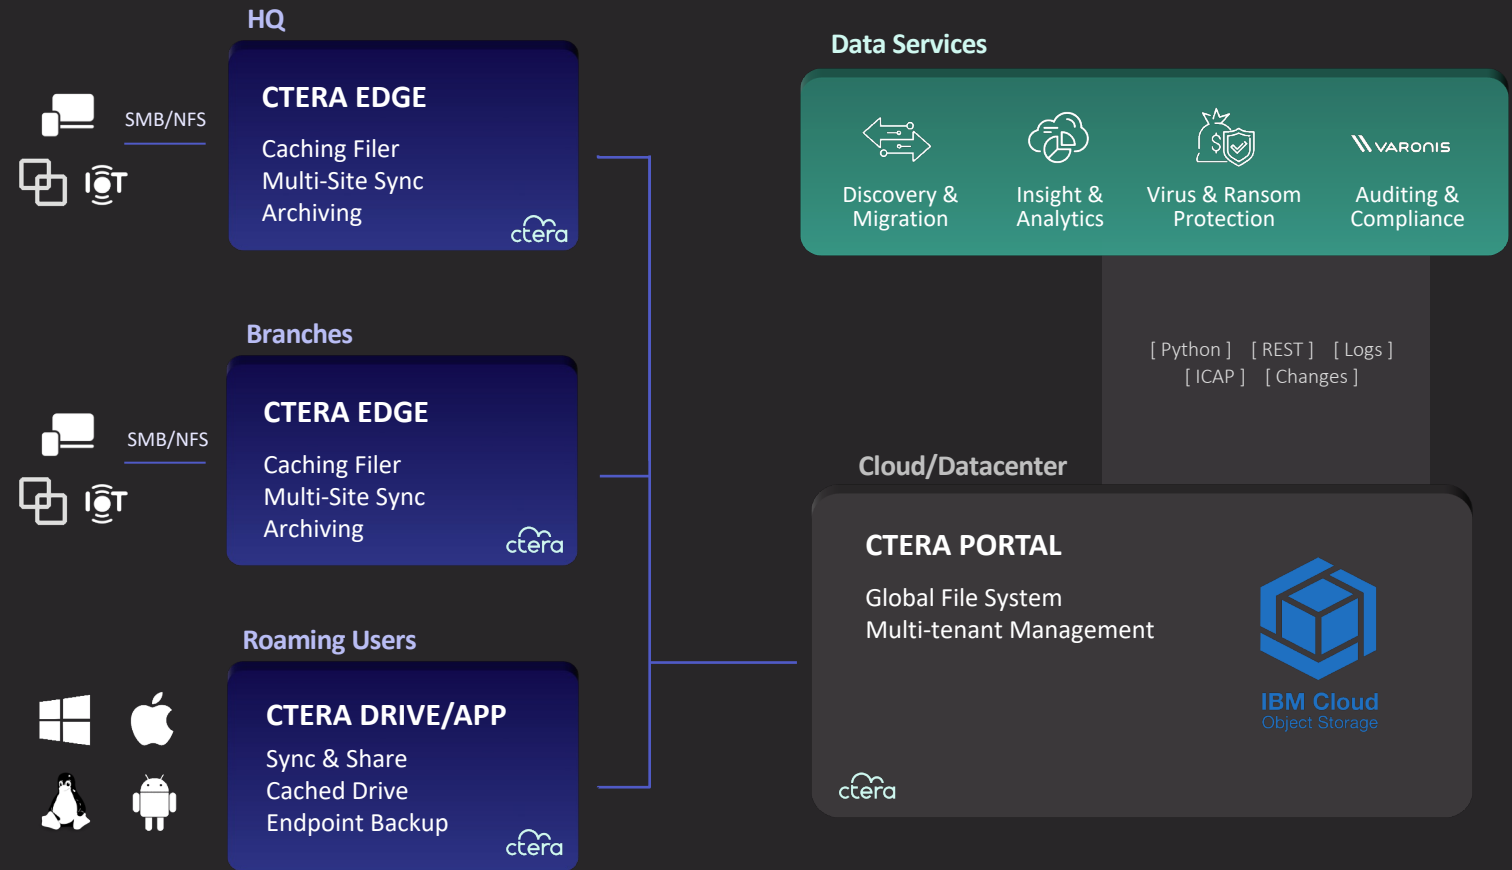
Addressing user content collaboration as well as automated machine-generated data workflows





GLOBAL FILE SYSTEM ARCHITECTURE

- NAS Replacement and Modernization using Single namespace over cloud, core and edge for distributed user content and machine-generated data
- Accelerated remote cache-based with built-in data protection and rapid DR
- Object-native backend using customer's choice of private or public infrastructure
- Seamless migration from existing NAS systems with full ACL support
- Programmable API for data pipeline automation and analytics
- Military-certified security with built-in virus and ransomware protection



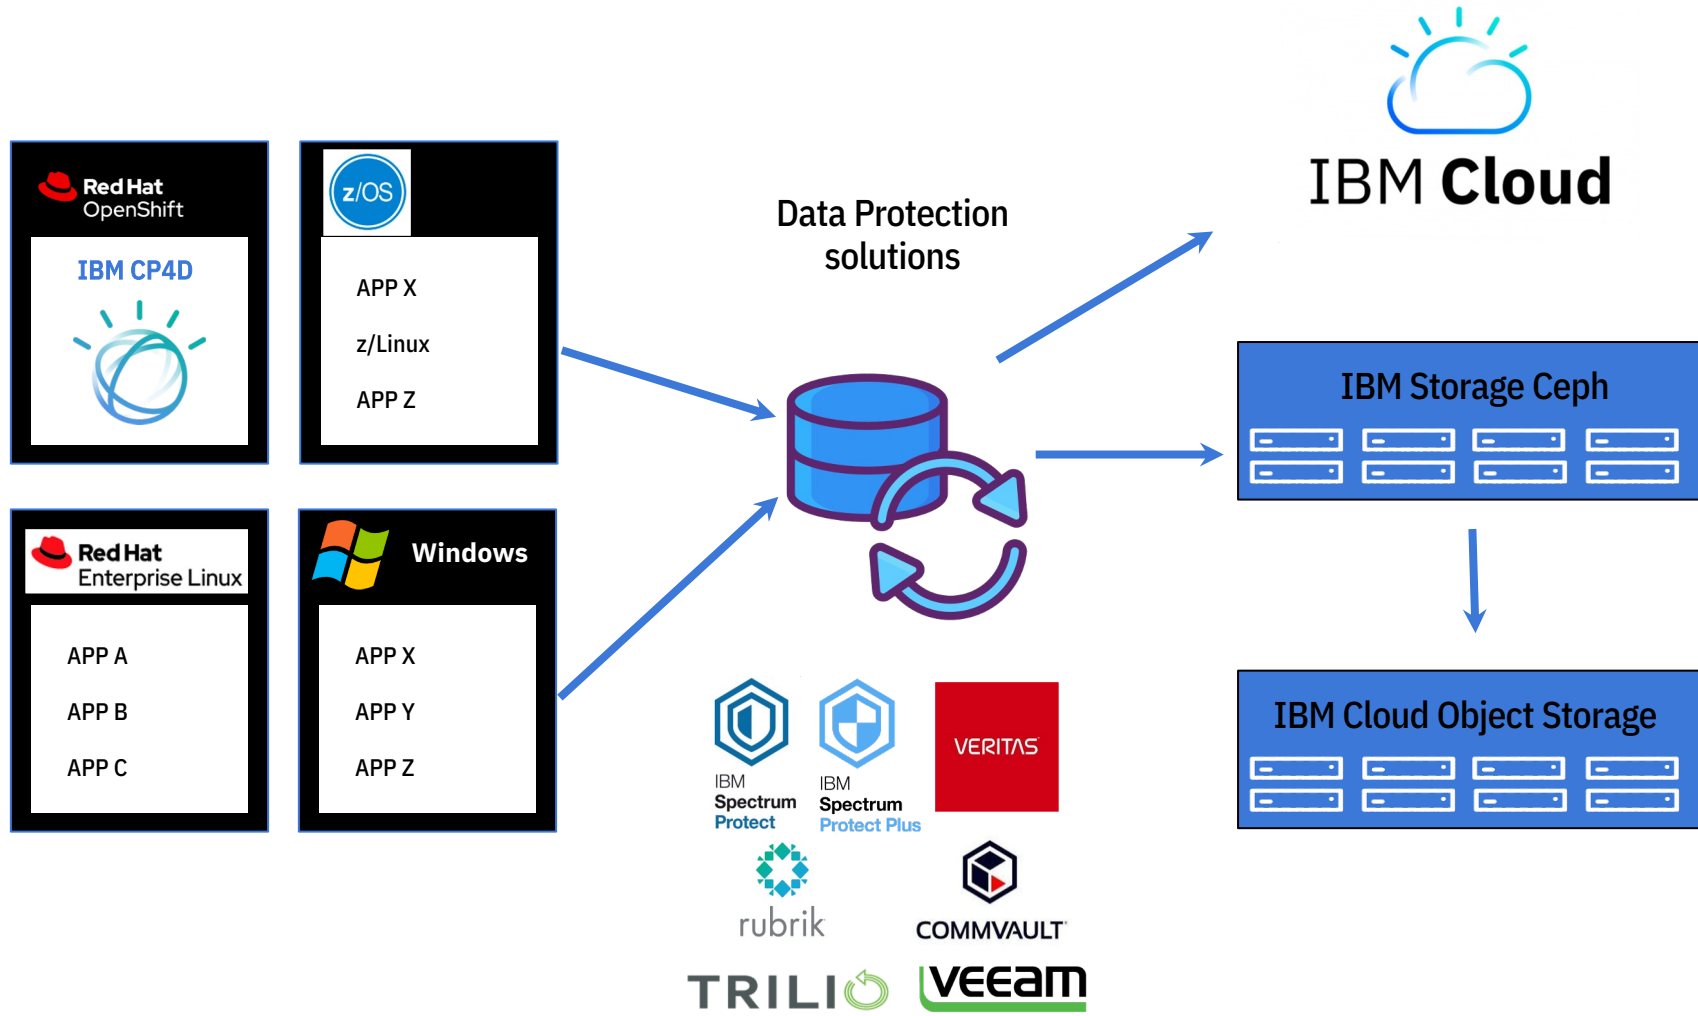
Cyber Resiliency

- Continuous replication to an immutable data pool with instant roll-back
- Live file operations analysis for early alert
- Distributed trust with central control, resilient to site or endpoint compromise
- Self-defending against virus and ransomware attacks
- Data classification for auditing compliance and governance

Object storage use case no. 4 – Cyber resiliency



IBM COS as a backup target (operational, long term, retention)

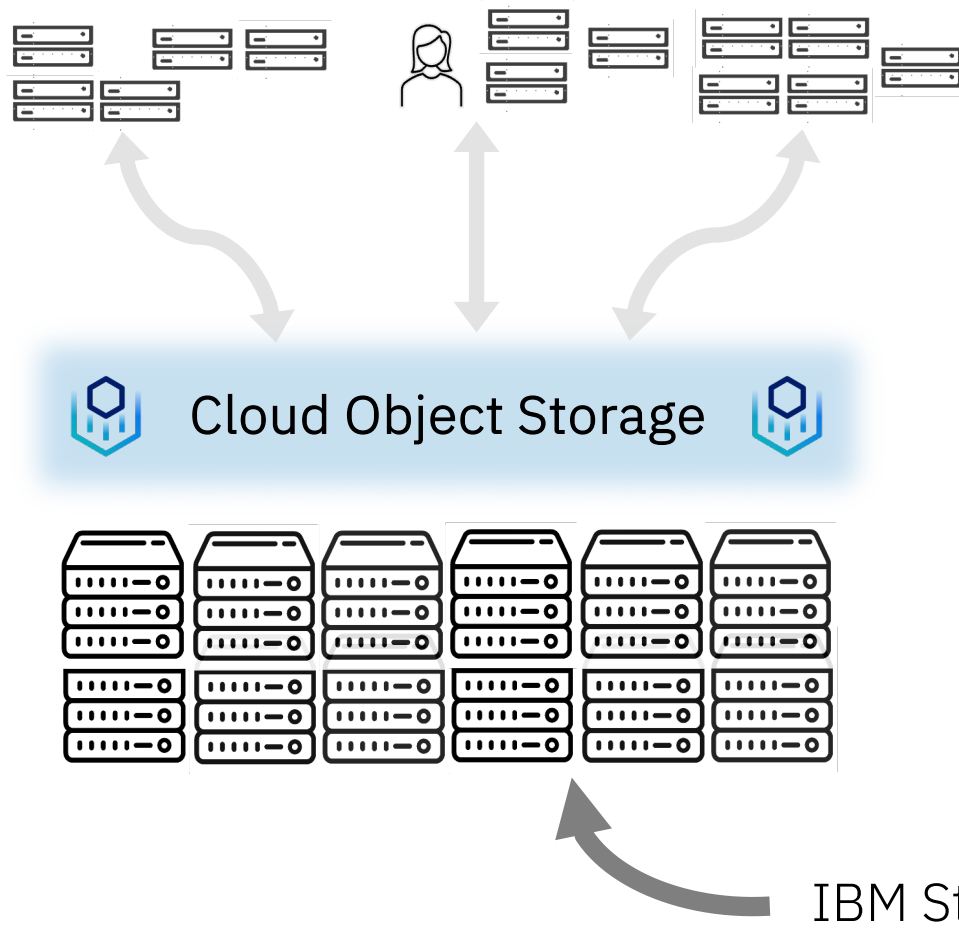


Object storage use case no. 5 – Storage as a Service of IT



Storage as a Service of Enterprise IT

Enterprise Storage w/public cloud like operational efficiency....



- ✓ On demand, horizontal scaling to exabyte capacity with single point of administration
- ✓ Fully automated, self healing rebalancing resiliency
- ✓ API driven architecture for automation and the ability to increase capacity without adding operational staff
- ✓ Runs on low cost, industry standard servers for flexibility and cost efficiency
- ✓ Open source based software flexibility and agility
- ✓ Reduced storage TCO by 40%

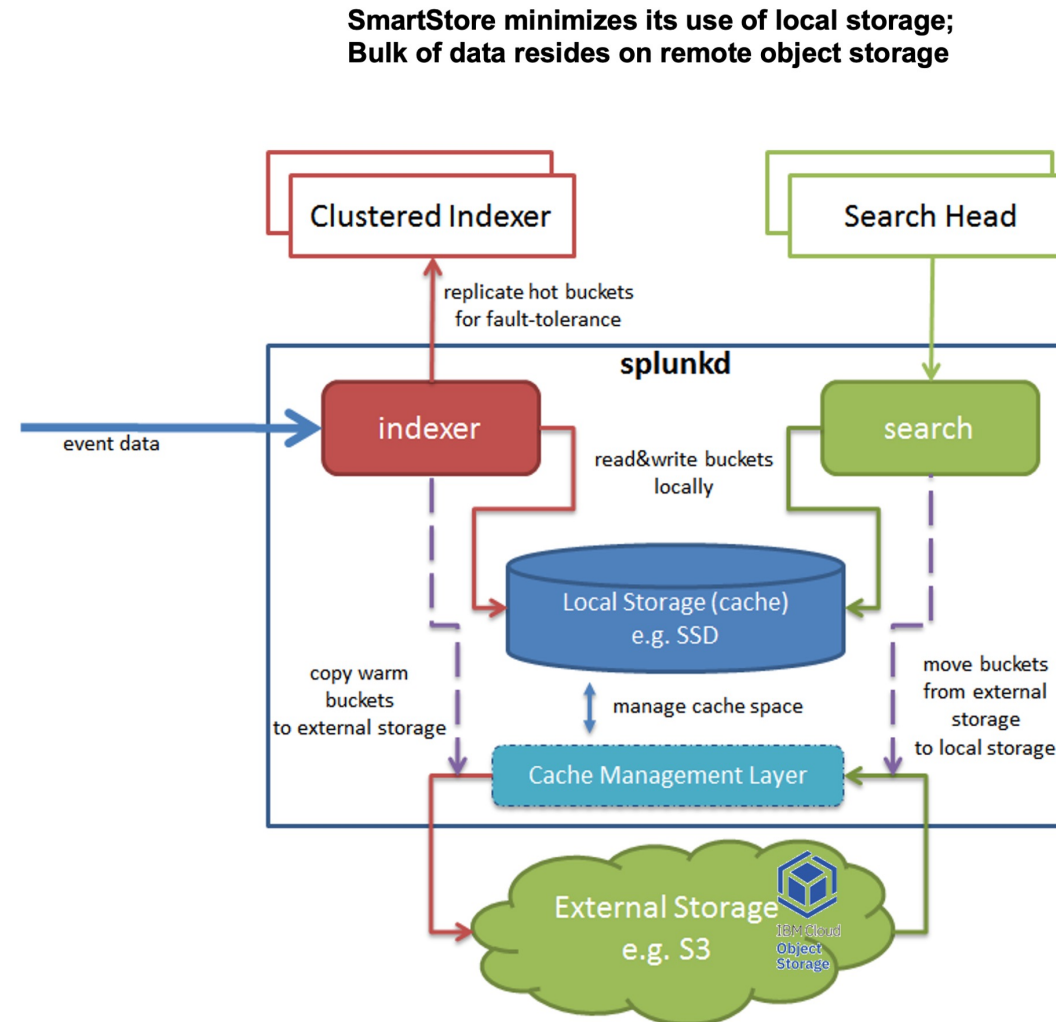
Object storage use case no. 6 – Active archive



Active archive – Splunk SmartStore

SmartStore data flow

- Data streams to Source Indexer that saves it locally in a hot bucket; also replicates it to target indexers
- When the hot bucket rolls to warm, the source indexer copies the warm bucket to the object store while leaving the existing copy in cache
- Target indexer deletes their copy per policy
- Master copy resides in Object Store – **all searchable by Splunk**
- Ensures High Availability

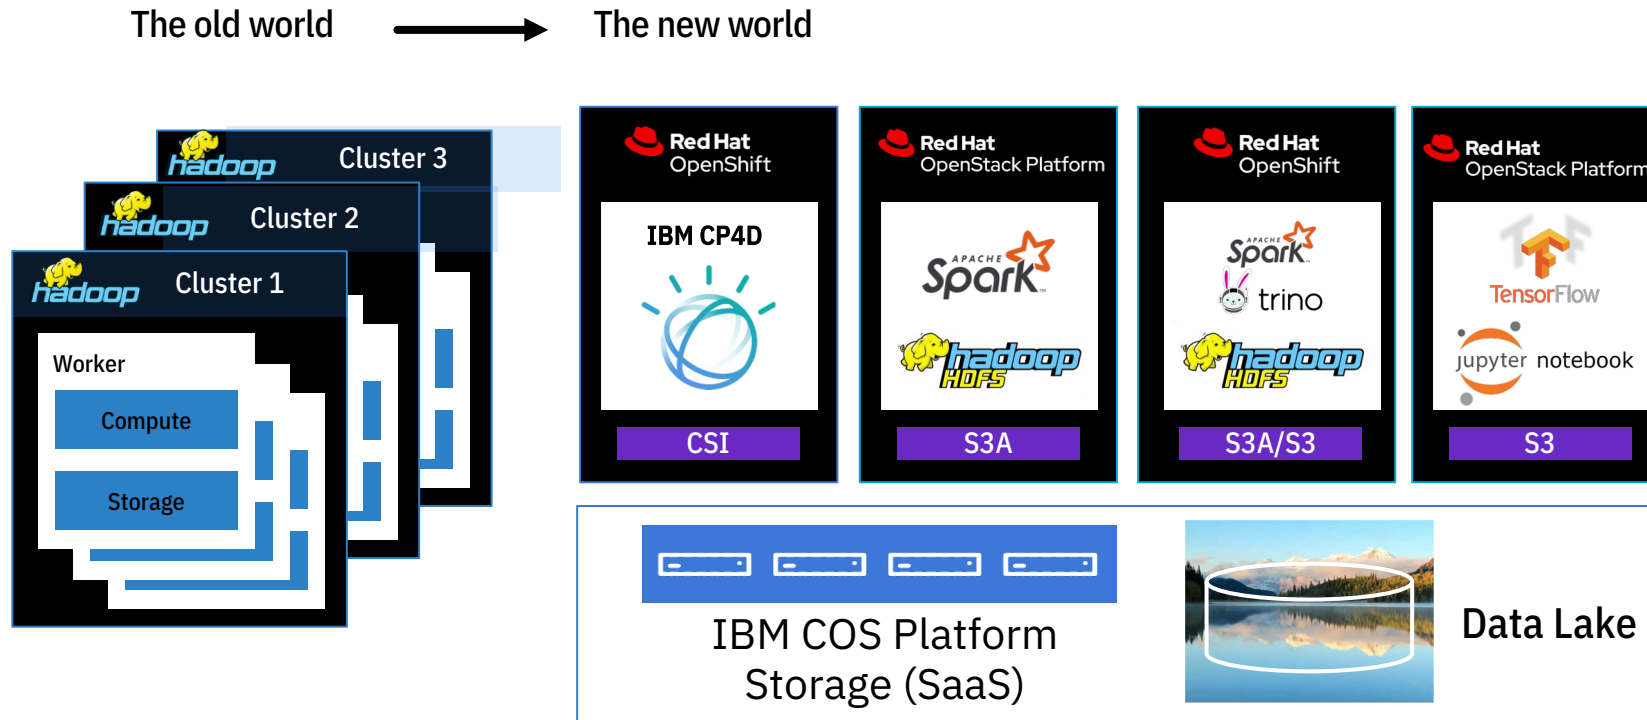


Object storage BONUS use case – Analytics repository



Object storage for live Data Analytics and AI/ML Infrastructure

Multi-tenant workload isolation with shared data context



IBM COS Supports S3 Select (CSV, JSON, Parquet) in ClevOS 3.18.0 (Long Term Support Release)

IBM Cloud Object Storage differentiators

EASE OF USE

Ease of deployment. Ease of understanding. Ease of software update. (Knowledge transfer and functional overview requires only two (2) hours.)

SUPPORT

IBM Cloud Object Storage Team (i.e. Pre Sales, Lab Services, Advanced Technology Group, Customer Support) remains actively engaged with all deployments.

PARTNERSHIP

Object use case adoption partnership. (i.e. Identifying new use cases, developer tools, ISV solutions, Solution assurance)

SELECTABLE DATA PROTECTION

Selectable data protection and relative storage efficiency across multiple sites. Erasure Coding is inline. Data is fully protected upon write commit.

EASE OF TECHNOLOGY UPDATES

Most customers update their own software without assistance. Many customers add their own hardware, particularly x86 commodity users.

HARDWARE LIFECYCLE MANAGEMENT

No vendor lock-in. 3rd Party hardware qualifications ongoing. Density at 1840 TB in a 4u storage enclosure. Constantly adding new hardware to the portfolio


COMMODITY HARDWARE SUPPORT

More than a dozen x86 SKUs supported for bare metal IBM COS Slicestor appliances with Disk Lifecycle Management (DLM). Plus VMware, plus Docker.

PROVEN AT MASSIVE SCALE

IBM COS System, IBM COS Service, private and public implementations of the same ClevOS software. Multiple EB customers in production. Single global object store.

What about swim lanes?



Where we're going, we don't need swim lanes!

Storage!



The other bonus reason – The strength of the IBM object storage portfolio

Ceph	COS	Scale
Native Object		Native File
Open source (Downstream)	Simple to deploy and manage	High performance – POSIX parallel file system
Flexible – granular controls and customization	Proven at very high capacity -> 1 EB +	Rich feature set and enterprise functionality
Rich Amazon S3 Feature Set and Compatibility	Strongest when 3 sites available	Widely recognized leader in HPC storage
Offered as Software Defined Storage or as “appliance”	Offered as an “appliance” (IBM Storage Ready Node)	Single namespace access to all other storage platforms and vendors on prem and in the cloud

Learning resources



IBM Redbooks for IBM Cloud Object Storage System

Learn more about IBM Cloud Object Storage

IBM Cloud Object Storage

Enable AI workloads and consolidate primary and secondary big data storage with industry leading object storage.

Read the Gartner MQ

Talk to an expert



www.ibm.com/products/cloud-object-storage



SCAN ME

Updated for 2023



<http://www.redbooks.ibm.com/redbooks/pdfs/sg248439.pdf>

<http://www.redbooks.ibm.com/redpapers/pdfs/redp5537.pdf>

Accelerate with ATG Survey

Please take a moment to share your feedback with our team!

You can access this 6-question survey via [Menti.com](#) with code 1708 6924 or

Direct link <https://www.menti.com/alwhyze7z1gz>

Or

QR Code



IBM