ADVANCED TECHNOLOGY GROUP (ATG)

Accelerate with ATG Webinar:
Revolutionizing Data Archiving – IBM Storage Deep Archive on Diamondback

Jason Peipelman - UX Architect IBM Storage, Master Inventor

Edgar Su – IBM Storage Tape Product Manager

Accelerate with ATG Technical Webinar Series

Advanced Technology Group experts cover a variety of technical topics.

Audience: Clients who have or are considering acquiring IBM Storage solutions. Business Partners and IBMers are also welcome.

To automatically receive announcements of upcoming Accelerate with IBM Storage webinars, Clients, Business Partners and IBMers are welcome to send an email request to accelerate-join@hursley.ibm.com.

2024 Upcoming Webinars – Register Here!

IBM Fusion and RedHat OpenShift Virtualization – August 6th, 2024

Important Links to bookmark:

ATG Accelerate Site: https://ibm.biz/BdSUFN

ATG MediaCenter Channel: https://ibm.biz/BdfEgQ

Offerings

Client Technical Workshops

- **IBM Storage Scale System & Storage Scale Workshop: August 8th in San Jose, CA**
- **IBM Fusion & Ceph: A Deep Dive into Next Gen Storage: August 21st – 22nd in Atlanta, GA**
- **IBM FlashSystem Deep Dive & Advanced Functions: September 18th – 19th in Paramus, NJ**
- IBM DS8900F Advanced Functions
- IBM Cyber Resiliency with IBM Storage Defender

TechZone Test Drive / Demo's

- IBM Storage Scale and Storage Scale System GUI
- IBM Storage Virtualize Test Drive
- IBM DS8900F Storage Management Test Drive
- Managing Copy Services on the DS8000 Using IBM Copy Services Manager Test Drive
- IBM DS8900F Safeguarded Copy (SGC) Test Drive
- IBM Cloud Object Storage Test Drive - (Appliance based)
- IBM Cloud Object Storage Test Drive - (VMware based)
- IBM Storage Protect Live Test Drive
- IBM Storage Ceph Test Drive - (VMware based)

Please reach out to your IBM Representative or Business Partner for more information.

***IMPORTANT* The ATG team serves clients and Business Partners in the Americas, concentrating on North America.**

Registration Open!

Storage @ IBM TechXchange Conference 2024

October 21-24, 2024
Mandalay Bay | Las Vegas
#IBMTechXchange

Key Learnings

- Practical how-to advice
- Patterns and best practices
- Success stories, IBM PoV, proven techniques

Featured Products

IBM Storage Defender
IBM Storage Fusion
IBM Storage Scale + IBM Storage Ceph
IBM Tape + IBM SAN
IBM Storage FlashSystem + IBM Storage DS8000

Collaborate. Learn. Play.

Community

IBM Champions
User Groups
Tech Peers
Business Partners

Sandbox

Network
Learn
Collaborate
Play

Roadmaps

Go deep with people in the know and set the stage for where IBM is going in the future

Breakout Sessions

Trends and Directions
User Groups
Product Deep Dives
Meet the Expert
Professional Development
Show the Code
Birds of a Feather
Academic/Research

Accelerate your Career

Labs (Instructor-Led, Self-paced)
IBM Certification Testing
Earn up to 25 hours in CPE credits

https://www.ibm.com/community/ibm-techxchange-conference/

Game On! →

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

ADVANCED TECHNOLOGY GROUP (ATG)

Accelerate with ATG Webinar:
Revolutionizing Data Archiving – IBM Storage Deep Archive on Diamondback

Jason Peipelman - UX Architect IBM Storage, Master Inventor

Edgar Su – IBM Storage Tape Product Manager

Meet the Speakers - Example

Jason Peipelman is the lead designer and UX architect for the Tape Storage and Mainframe Storage teams. He has over 20 years of storage experience in both development and design and helps ensure our products are usable, accessible, and intuitive.

Edgar Su is the Storage Tape product manager, responsible for the entry-mid range physical library and Tape Software. He has over 10 years of experience in storage products in both supply chain and product management.

Unstructured data growth and cost of storing it in the cloud

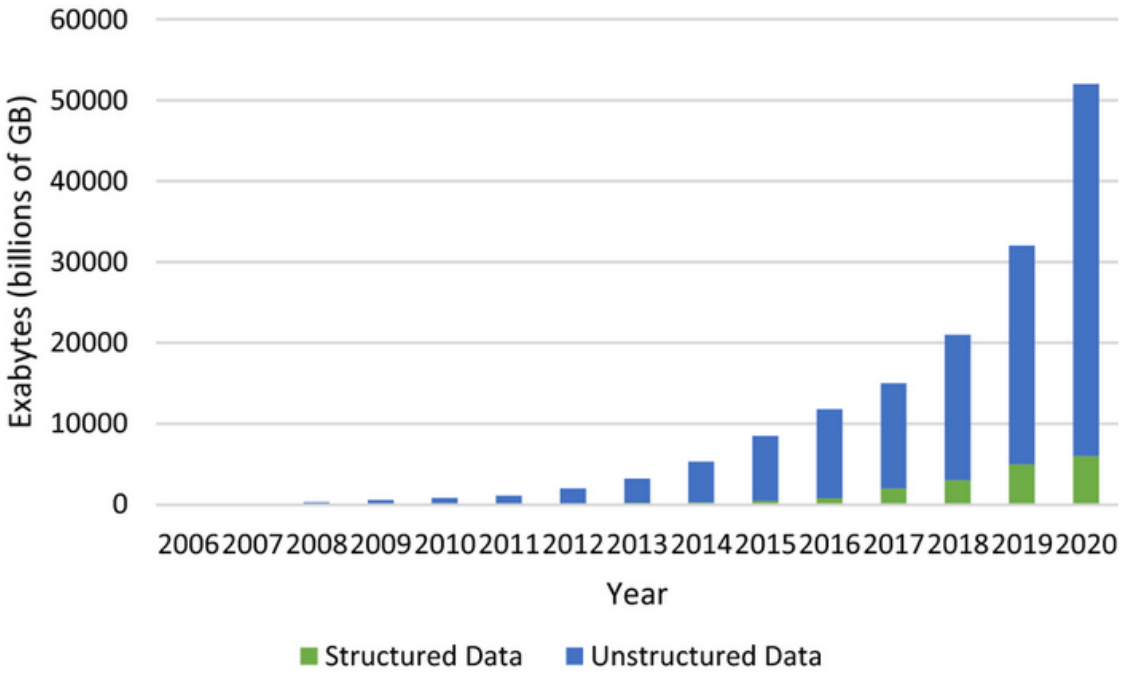

- **80-90%** of Data is Unstructured
- Unstructured Data is Growing **55-65%** Annually
- **43% of IT Decision Makers** fear their IT infrastructure won't handle future data demands

Unstructured data growth and cost of storing it in the cloud

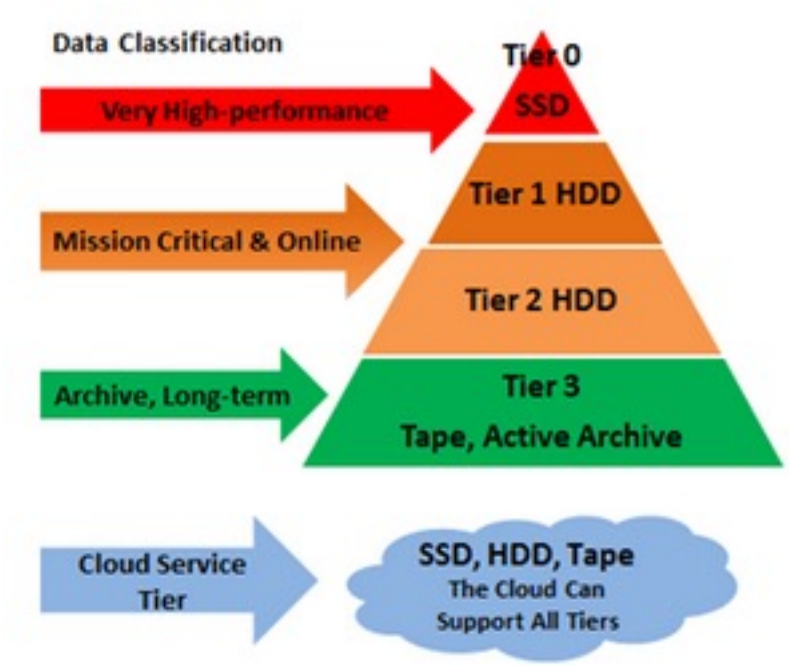

- **Active archive**
  - Required for compliance & security
  - Integrates with data governance policies
  - Allow backup & archive for high availability and disaster recovery
  - S3 Glacier storage classes are being widely adopted as the industry standard interface for cloud based archival solutions.

Unstructured data growth and cost of storing it in the cloud

Cloud archival storage

- 90% of IT Decision Makers say cloud is a “cornerstone” of their digital strategy.
- 94% say their cloud storage costs are rising, and 54% say they are growing faster than other costs.
- 37% of business owners reported being taken aback by cloud charges with 82% incurring unnecessary and unpredictable cloud costs due to things like data egress fees.
- 82% of IT Decision Makers stated their biggest cloud challenge was managing cloud spending - even beating out security at 79%.

What people want in an Active Archive solution

Survey in 2023 showed the following value statements were prioritized:

1.	**Integration**	Standard, S3 compliant interface allows easy integration into existing infrastructure and software.
2.	**Cost**	Cost benefits that ensure stability in pricing and no unexpected fees.
3.	**Black box**	A self-contained solution with a simple set of options that comes with everything you need.
4.	**Resiliency**	Data resiliency that ensures the protection of data is prioritized over data access.
5.	**Ease of use**	A system that doesn't require advanced tape knowledge to run and can be configured operated by an S3 administrator.
6.	**Secure**	Security by design is implemented from the start including encryption, system monitoring, and air-gapped access control.
7.	**Sustainability**	Must help to meet reduced energy consumption, carbon footprint, and cooling requirements

What is Deep Archive?

IBM Deep Archive

For organizations looking for cost-effective long-term data storage in an integrated solution, **IBM Deep Archive** delivers the benefits of tape storage in an easy-to-use, easy-to-integrate solution in the footprint of a single 19" rack with up to 27 PB of capacity and up to 19.1 TB/hour throughput.

This solution can be used directly by your existing S3 Glacier Flexible Retrieval storage class compatible software to reduce the cost of your cold data by up to 80%, with no data egress fees.

A black-box experience ensures complexity is hidden and internally optimized so that the entire solution is supportable by S3 administrators with no technical tape experience.

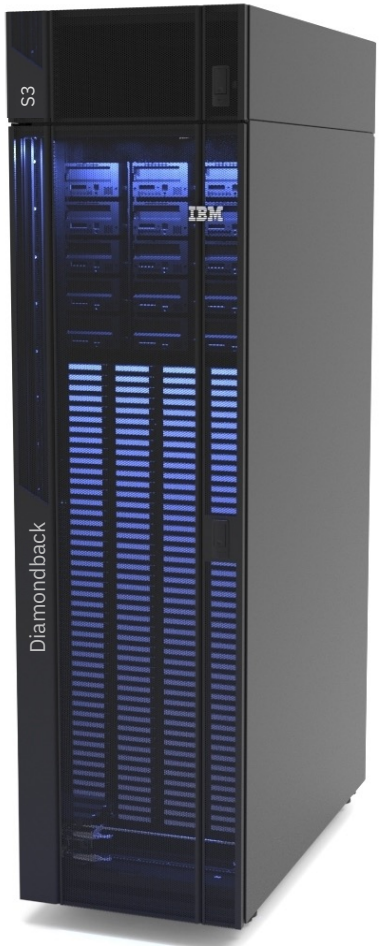

Benefits

Cost benefits

- Up to **80% lower TCO** than comparable cloud storage.
- No data egress fees.

Standard interface

- Industry standard S3 API to tape allows leveraging of existing data movers and S3 software.
- Compatible with the **S3 Glacier Flexible Retrieval storage class**.

Simplified management

- S3 administrators require **no specialized knowledge of tape storage**.
- Installed in less than a day.
- No tuning required, plug-and-play.

Black-box solution

- Comes with everything you need from a single IBM seller
- Storage Expert Care provided for the entire solution.
- No FC or SAS external to the system.

Sustainability

- Object storage energy consumption reduced by up to 97%
- 85% lower carbon footprint than comparable hard disk solutions.

Security

- Keep your data your data by ensuring it is kept behind your firewall.
- Consistently apply security, encryption, system monitoring, and air-gapped control access.

Node Configuration

Entry Configuration

Server Hardware
1 Server Node

LTO 9 Tape Drives
4 drives
Up to 4.6 TB/hour.

LTO 9 Tape Cartridges
Up to 27 PB of storage

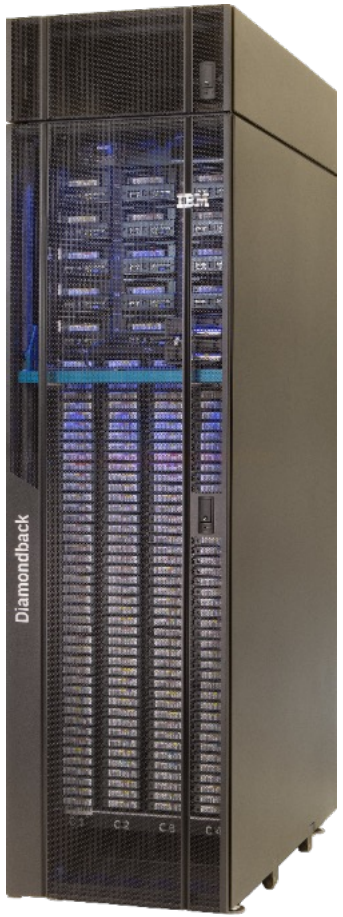

4.6 TB / hr
23.34 watts/PB
Deep archives
Air gap
Second/Third copy
Regulatory compliance

Tape Library configuration

Simplified architecture

- Singe-frame scale-out infrastructure
- Single robot, high availability grippers
- 1 service magazine with 10 I/O slots
- Up to 27 PB uncompressed capacity
- Up to 1548 Tape cartridges
- Up to 14 LTO-9 Tape drives

Enhanced capabilities

- Cost kept to a minimum
- Fast deployment, <1-hr install
- Pre-loaded media option
- Storage pod design no regular media removal
- Increased data density and energy savings
- Quick Mean Time to Repair

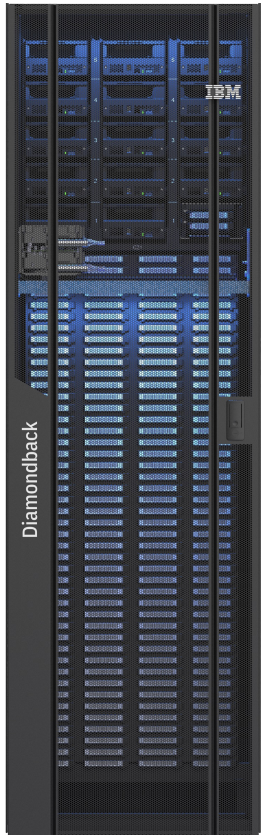

Capacity options

- **14 PB** preloaded media

800 cartridges

- **27 PB** preloaded media

1548 cartridges

- **Supply your own** media.

Integration

Customer provided S3 orchestration

Deep Archive	Deep Archive	Primary Tier	Off-prem Cloud
S3 Glacier	S3 Glacier	File or S3 Standard	

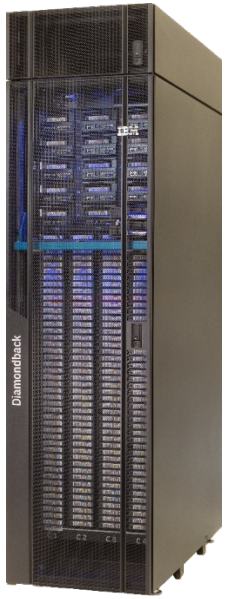

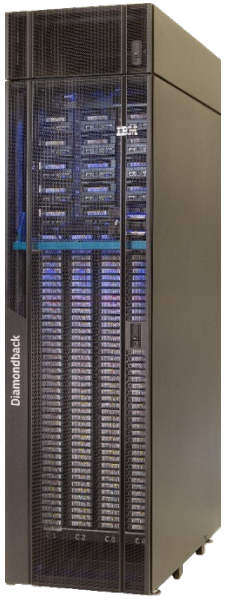

Cloud Storage

aws

Microsoft Azure

Install and Configure

- IBM install →
  1. **Install and configure the tape library**
  2. **Install and cable the S3 nodes**
  3. **Update code on the nodes**
  4. **Update firmware on library and drives**
  5. **Configure network settings**
- Cabling →

1x - Management network
2x - S3 network (10G)
2x - Power

1x - Library connection
1x - iDRAC connection
1x - TSSC connection

Optional

- Customer configuration →
  6. **Setup user accounts, create a bucket**

S3 Glacier Flexible Retrieval

S3 applications must communicate using the S3 Glacier Flexible Retrieval storage class to read and write data via a 25G ethernet connection.

Writing data is first staged on the Deep Archive nodes. The response for the S3 API is returned and within 10 minutes the data is automatically migrated from the nodes to the tape library. Multipart upload is recommended with object sizes from 100 MiB to 5 TiB.

Restoring data in an S3 Glacier Flexible Retrieval storage class are not accessible in real time. The AWS RestoreObject request is required to create a temporary copy of the object available on the nodes. Restore times include the mount and seek times for the tape cartridges that contain the data and the restore state can be monitored using the Head command.

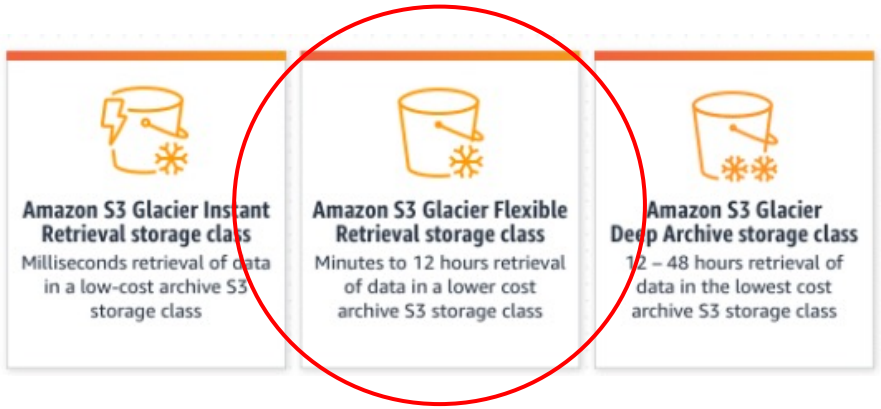

Supported maximums:

• Object size	5 TiB
• S3 accounts	100
• S3 buckets	1500
• Objects	1.5 billion
• Objects per S3 bucket	1,000,000
• Recommended average object size	100 MiB

TapeCloud Manager

Monitoring

- System heath
- Capacity metrics
- Throughput metrics
- Event notifications
- TapeCloud nodes
- Tape libraries
- Tape drives
- Data cartridges
- Cleaning cartridges
- Node hardware

Management

- System updates
- Library/drive firmware
- Certificates
- User accounts
- Library user accounts
- Buckets
- Software logs
- Library/drive logs
- Node logs
- Encryption settings
- NTP/time settings
- Network settings

Service actions

- Start/stop/restart nodes
- Power off/reboot nodes
- Rescan library
- Disable/enable drives
- Replace data cartridges
- Replace cleaning cartridges
- Validate cartridges

TapeCloud Manager CLI
Tape GUI
iDRAC
Linux

Pricing

- This shows a comparison against HDD and AWS Deep Archive.
- **AWS S3 Glacier Flexible Retrieval** is **3.6x** more expensive with even greater retrieval costs
- …so, it would be off the charts.

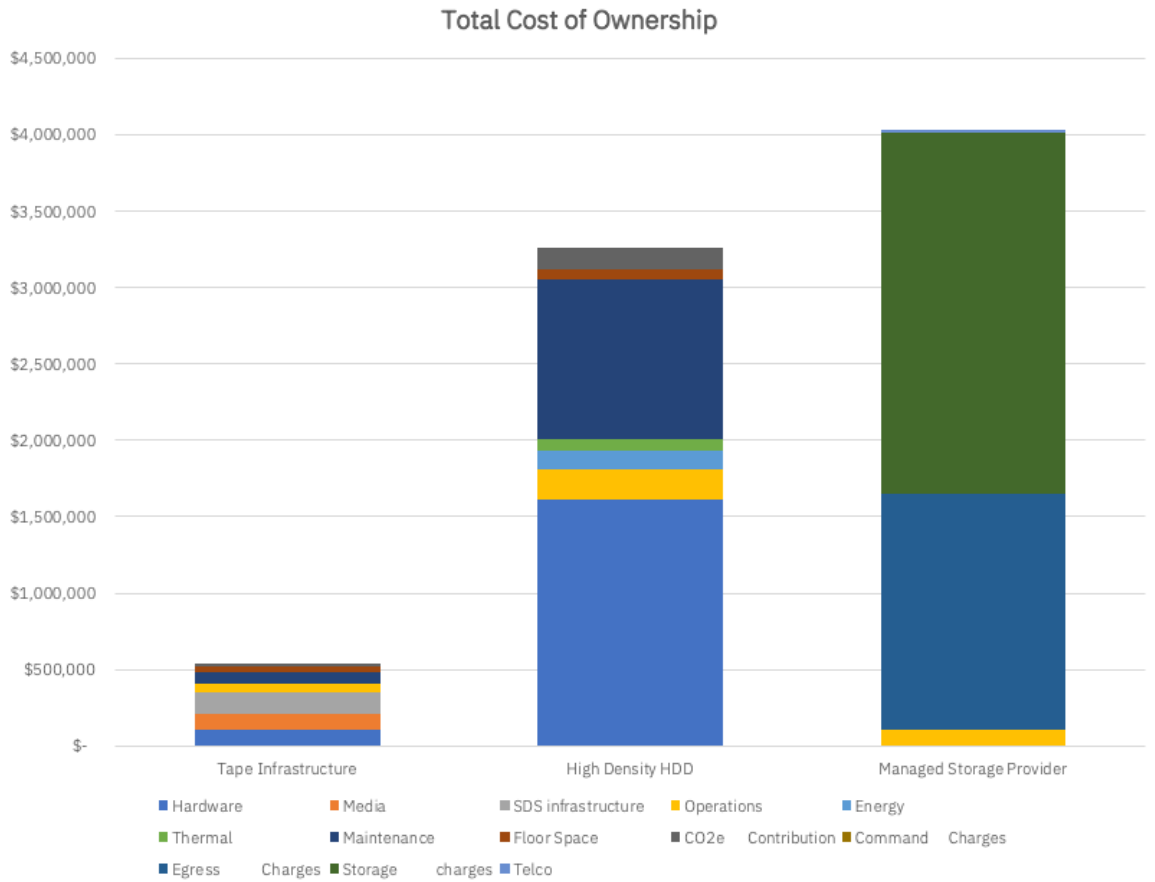

IBM Cost Compare Tool: 10-year retention, 27PB data, 16TB HDD commercial S3, AWS Glacier Deep Archive, 1% recall/month

Comparison

	IBM Storage Deep Archive	AWS Glacier Instant Retrieval	AWS Glacier Flexible Retrieval	AWS Glacier Deep Archive
10-year TCO Saving	-	98%	97%	85%
Cost ($) per 1000 commands	$0	$0.01-$0.02	$0.004-$10	$0.004-$0.10
Per GB Cost for recall to premise	N/A	.05 average up to 1PB in 30-day charge period*		
Retrieval time - Expedited	1-5 minutes	milliseconds	1-5 minutes	N/A
Retrieval time – Standard	1-5 minutes	milliseconds	3-5 hours	5-12 hours
Retrieval time – Bulk	1-5 minutes	milliseconds	Within 12 hours	Within 48 hours
$/GB Retrieval cost - Expedited	$0	$0.03	$0.03	$0
$/GB Retrieval cost – Standard	$0	$0.03	$0.01	$0.02
$/GB Retrieval cost – Bulk	$0	$0.03	$0	$0.0025

*See https://aws.amazon.com/s3/pricing/ for pricing at each TB tier.

Use Cases

A low-cost, on-premises cloud providing secure and durable storage for data archiving and online backup.

- Standard S3 Glacier commands
- On-Premise data control
- Data encryption
- No cost for retrieval priority
- Air gap storage
- Integrated system monitoring

NAS Cold Data offload

StorageGRID

komprise

Cloud pools * with orchestration

VERITAS Enterprise Vault

Object Archive/Repatriation

Orchestra

MoonwalkSoft

<API data mover/orchestration>
rclone move with –min-age

Back-up Retention Archive

veeam

COHESITY

rubrik

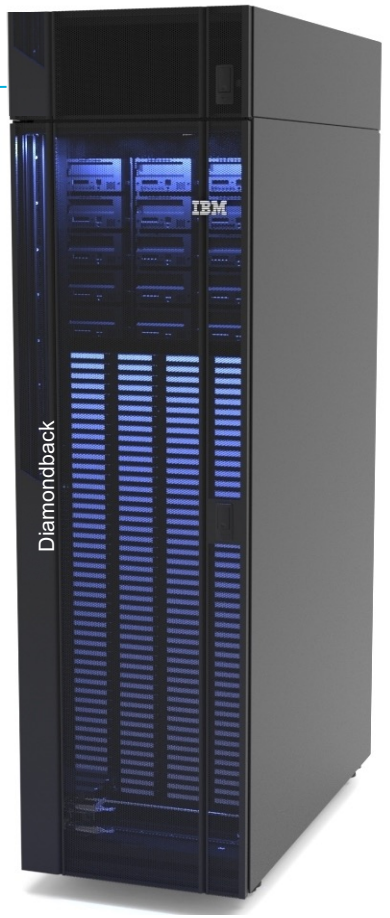

Positioning

Deep Archive

- Active archive for infrequent access
- Object only with S3 Glacier storage class
- User/application accessible data archive
- Low-cost, high-capacity
- 3 PB minimum, 9+ PB optimal

IBM Storage Scale + CES

- Transactional storage for AI and data analytics
- File and object (S3 Standard storage class) storage
- Tiered storage for user object/file frequent accessibility
- Full feature information lifecycle management platform
- Fully automated data durability platform

What's next?

Want to more about IBM Storage Deep Archive?

IBM Deep Archive webpage

- https://www.ibm.com/products/s3-deep-archive

Thank you!

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