

## **Scalable, High Performance S3 Object Storage on Tape**

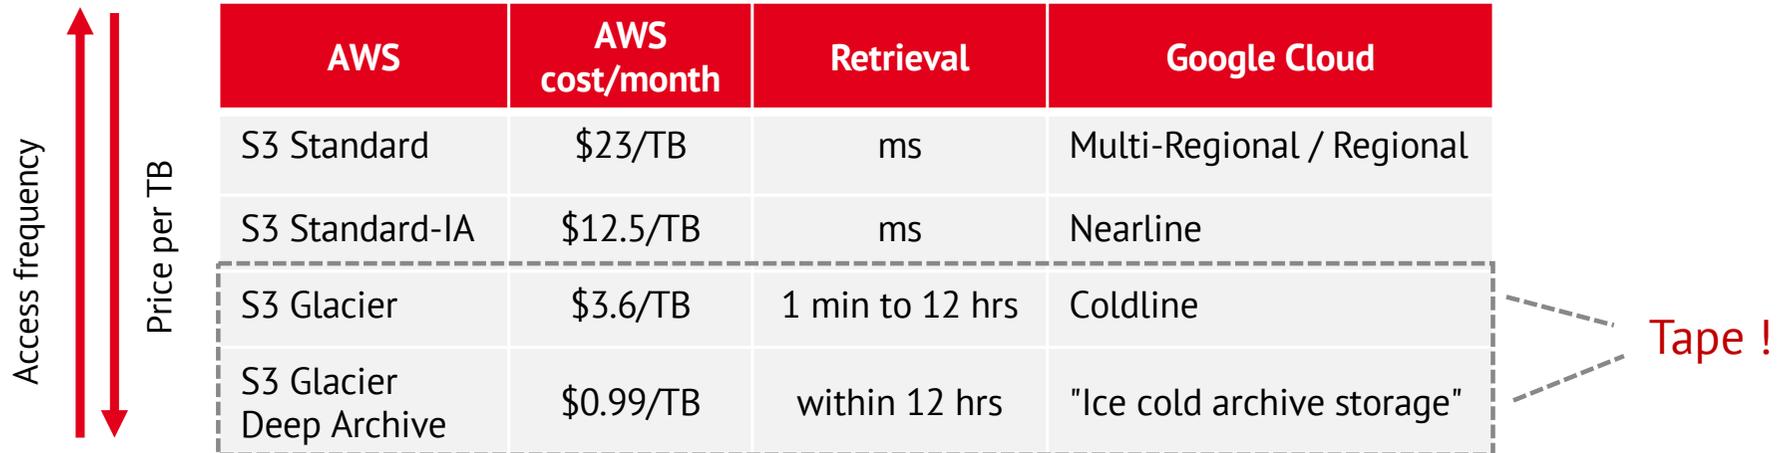
Thomas Thalmann  
PoINT Software & Systems GmbH

**„HSM & Cluster FS” - Konferenz 2023**  
**Deutsches Zentrum für Luft- und Raumfahrt e.V. Neustrelitz**  
**14. - 16.6.2023 „Gut Ulrichshusen” - Tressow**

## Tape-based Object Storage

Motivation

Storage classes at AWS, Google etc.



AWS	AWS cost/month	Retrieval	Google Cloud
S3 Standard	\$23/TB	ms	Multi-Regional / Regional
S3 Standard-IA	\$12.5/TB	ms	Nearline
S3 Glacier	\$3.6/TB	1 min to 12 hrs	Coldline
S3 Glacier Deep Archive	\$0.99/TB	within 12 hrs	"Ice cold archive storage"

**Mark Russinovich, Microsoft Azure CTO:**

*"After evaluating various technologies including Blu-ray and magnetic disk, Microsoft came to the conclusion that **tape** was still the way to go for Azure's archival storage tier"*

## Advantages of Object Storage Approach for Tape

- Designed for massive amounts of data without decrease of efficiency
- Ingest of complete objects - perfectly supports sequential character of tape
- Scalability and system security by adding nodes
- Rich custom metadata (in contrast to file systems)
- Immutability (Versioning, Object Lock) – fits well with tape
- Redundancy by Erasure Coding – can be realized with multiple tape media
- Standardized S3/HTTP application interface
- HTTP can easily be routed and connected to any networked system
- HTTP concept simplifies handling of tape-specific challenges like high access times and timeouts



- **Don't tier or replicate data from object storage to file system**
- **Don't put S3 service on top of a file system**

## Tape-based Object Storage

### PoINT Archival Gateway

Scales by Node



High Availability



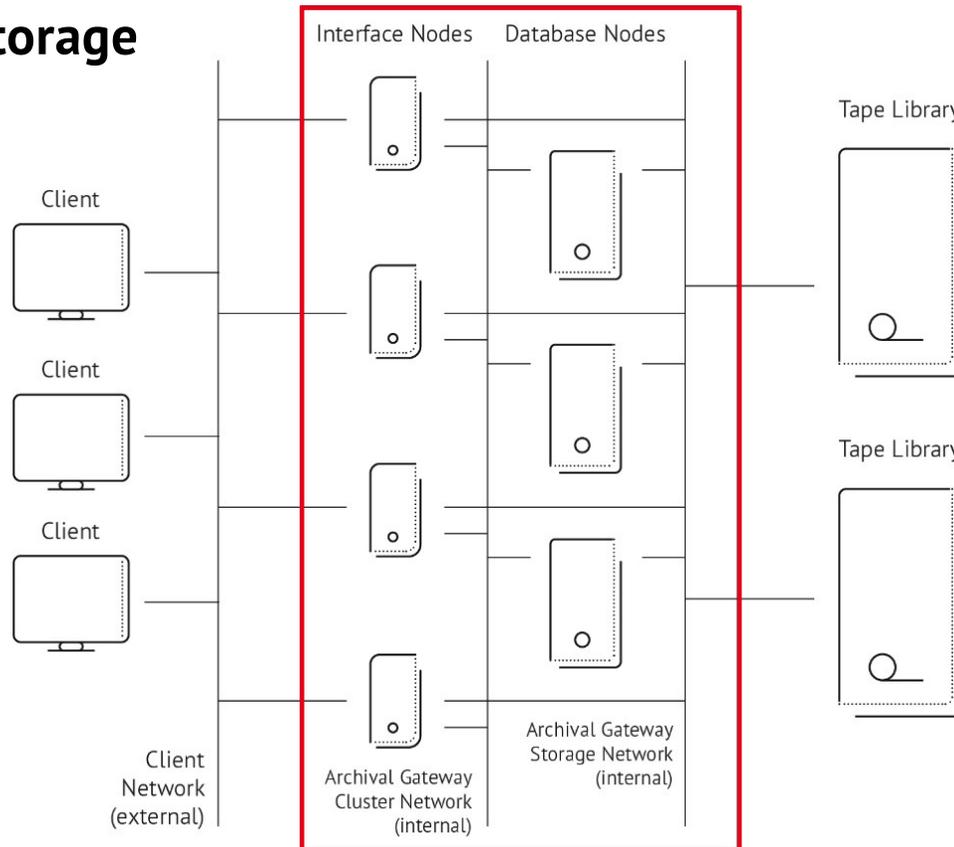
Native S3 API  
incl. Glacier



Erasure Coding  
for tape storage



Object Lock &  
Versioning



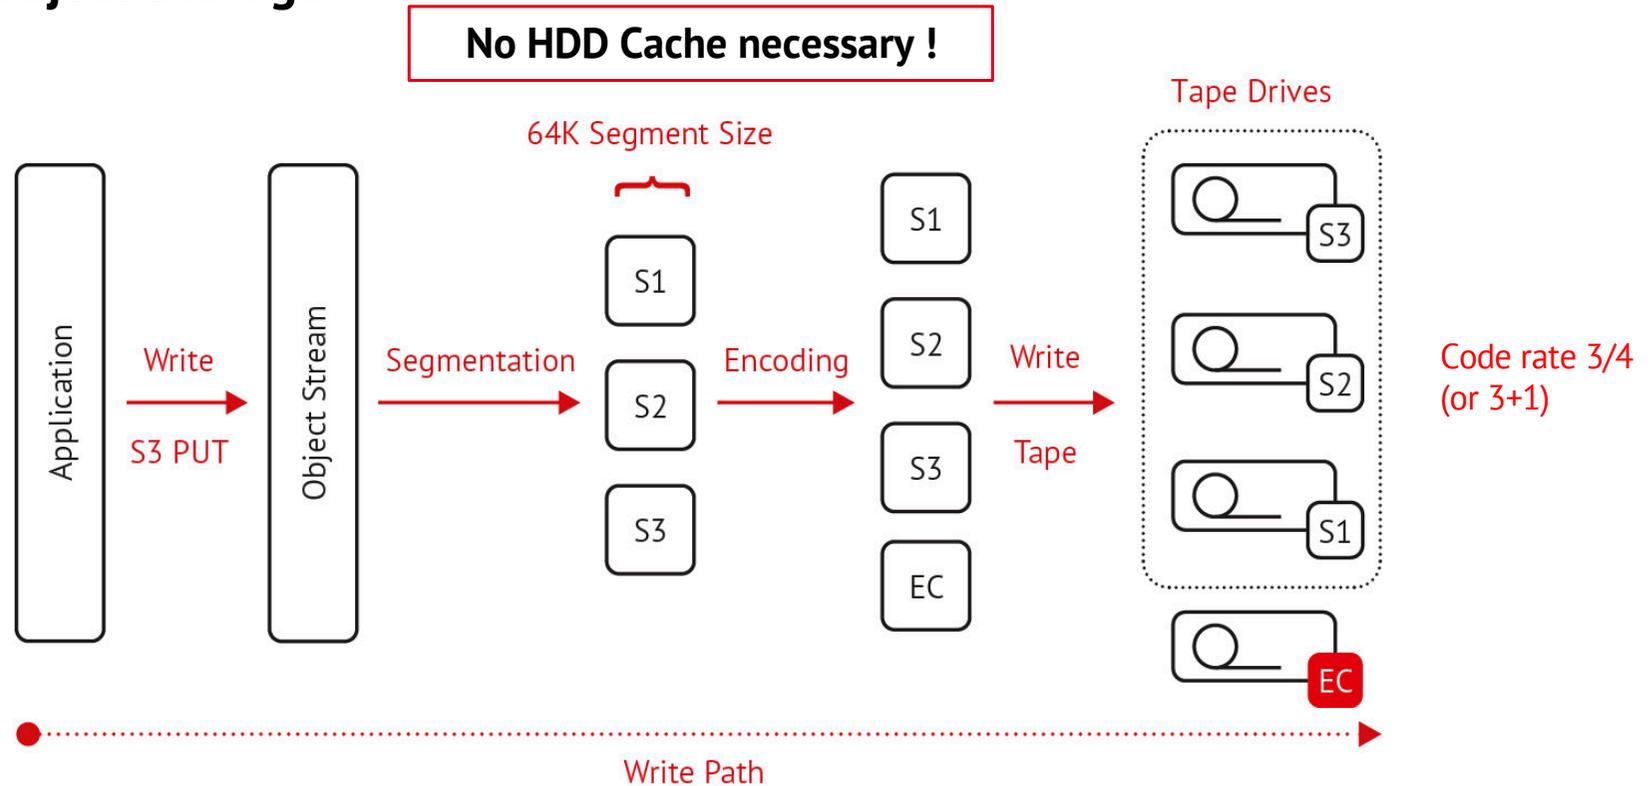
### Interface Nodes

- S3 REST API
- Library drives control
- Up to 8 drives per node
- Up to 230 GB/s throughput
- Multiple active nodes
- Up to 265 tape drives

### Database Nodes

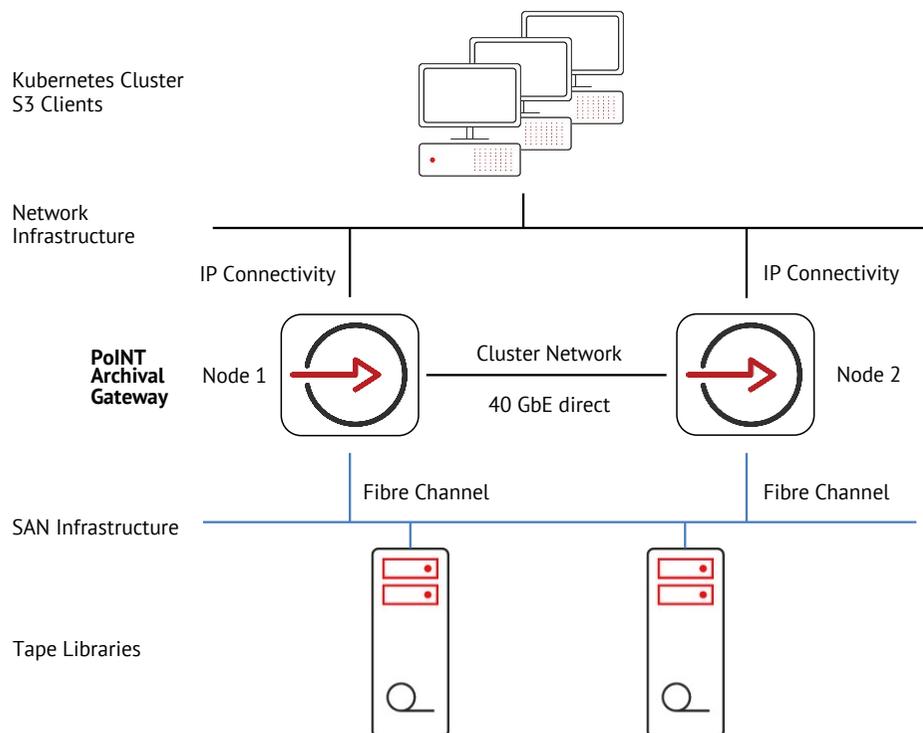
- Databases, logs, monitoring incl. SYSLOG, SMTP
- HTML, API admin interface
- Library robotics control
- Up to 4 nodes possible
- Up to 50 billion objects per S3 bucket

## Tape-based Object Storage Erasure Coding



## Tape-based Object Storage

### Case Study – EMBL European Bioinformatics Institute



### The challenges

- Archiving workloads from Kubernetes Cluster via S3
- Storage capacity in high three-digit PB range
- Long-term preservation and high availability
- Multiple vendor support incl. support for LTO and 3592 drives
- High performance for read and write access (1 PB per week)
- Low TCO

### The solution – PoINT Archival Gateway

- Scale-out architecture and unlimited storage space
- Native S3 interface including versioning
- Logical partitioning for workload requirements
- WORM and retention management
- Read and write performance up to 3 PB per week

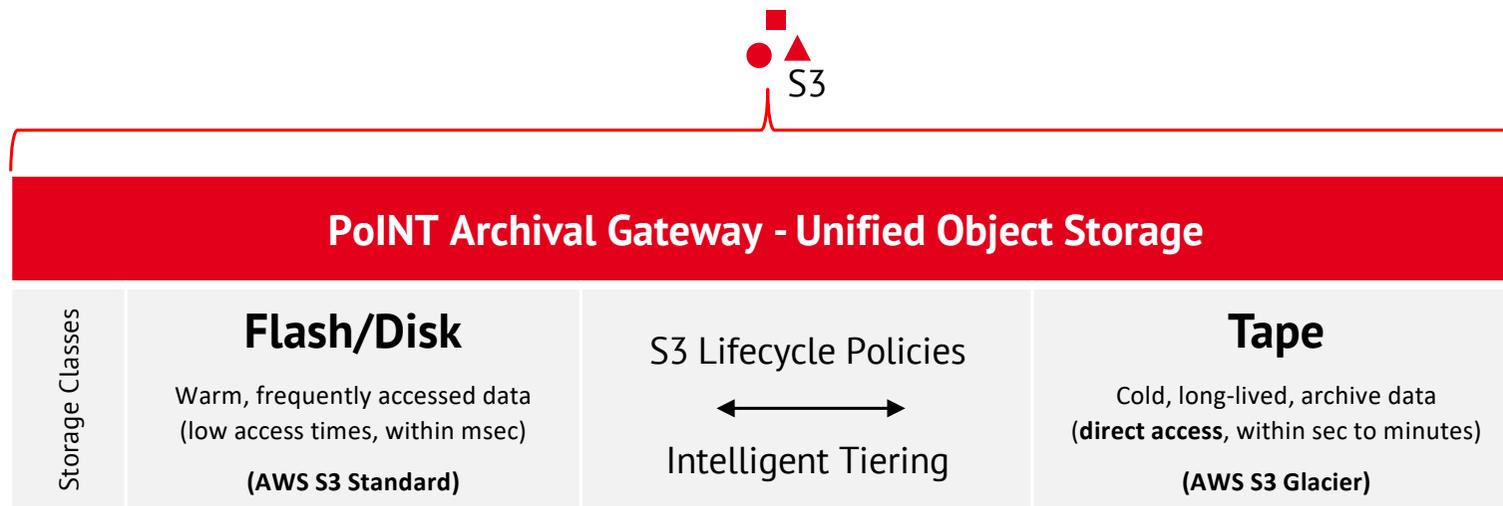
### The benefits

- Compatibility with any S3 client
- Workload specific configuration on bucket or partition level
- Scalable by nodes, drives, slots and media
- Software-defined for independence and sustainable planning
- No vendor and technology lock-in

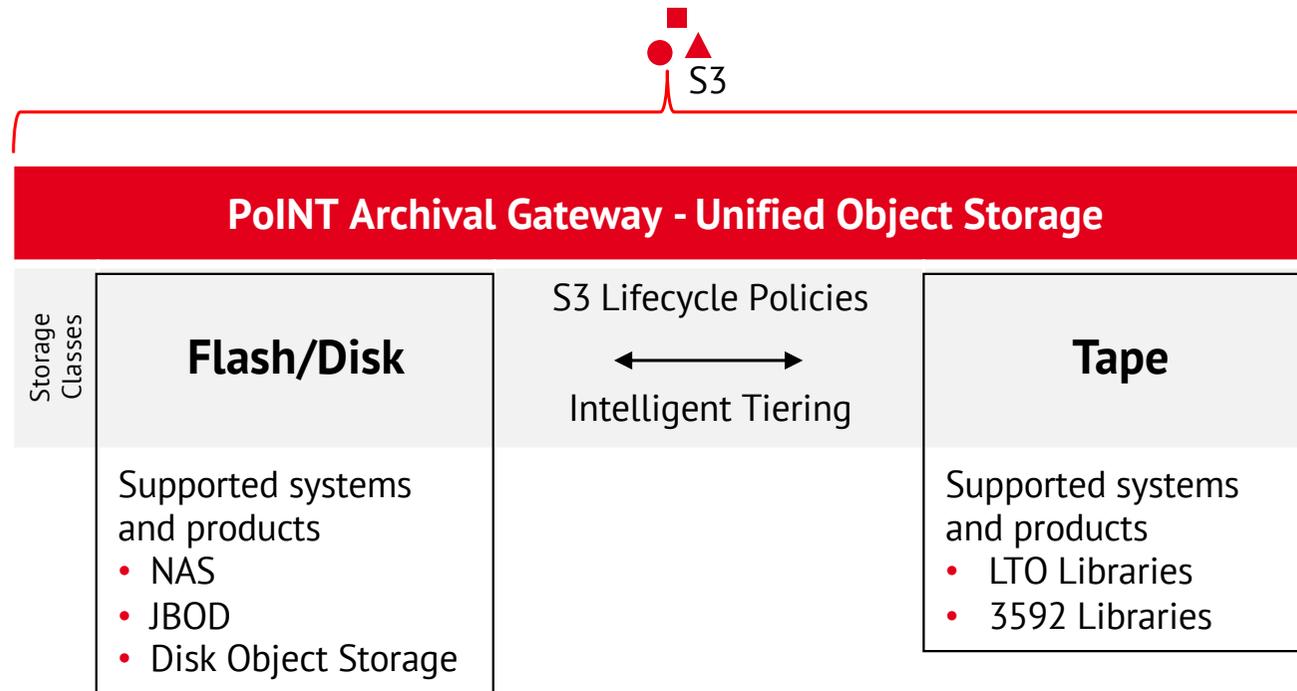
## Unified Object Storage

### Multiple Storage Classes

- Software-defined
- Flexible configurations
- Free choice of storage hardware
- Single namespace across flash/disk and tape
- Transparent access to flash/disk and tape storage class
- AWS S3 compatible
- Automatic replication to disk and tape
- No S3 Object Restore to disk necessary



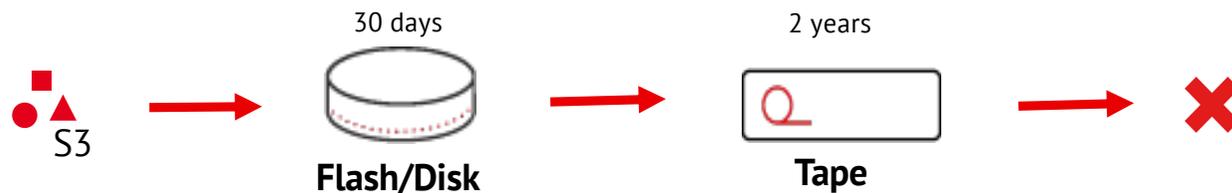
## Unified Object Storage Configurations



## Unified Object Storage

### Lifecycle Policies

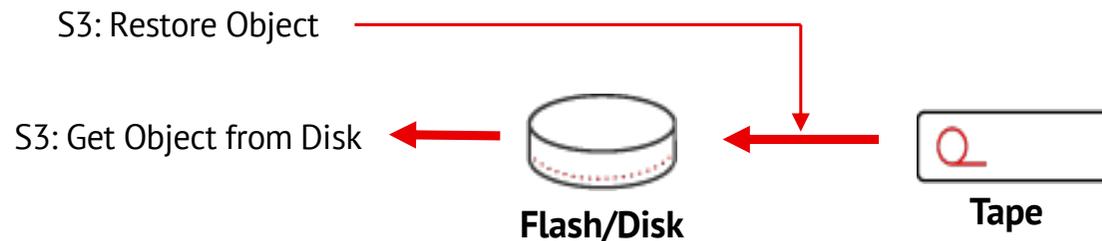
- AWS S3 compatible
- Transition policies
  - E. g. migrate objects from disk to tape 30 days after creation
- Expiration policies
  - E. g. delete objects one year after creation
- Example
  - Storage of periodic logs in disk storage class for one month for analysis purposes
  - After one month migrate to tape storage class for archiving purposes
  - After further 24 months delete objects



## Unified Object Storage

### Restoring Objects from Tape to Disk

- AWS S3 compatible command: Restore Object (Glacier API)
- Objects can temporarily be restored on disk
- Number of days how long a restored object is stored on disk can be specified



## Use Cases



HSM / File Tiering

HSM / File Tiering  
& Archiving to S3



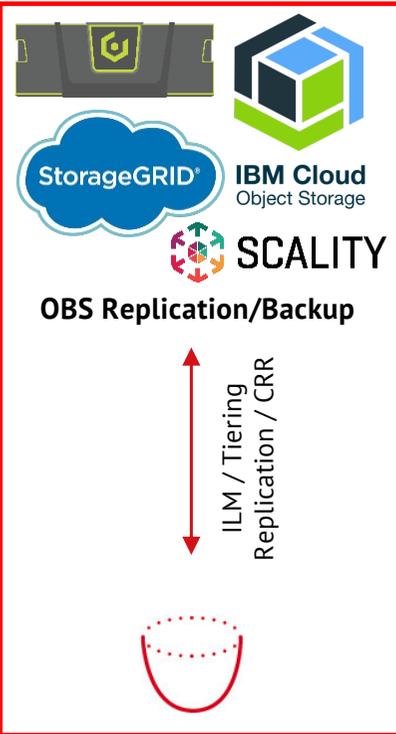
Cloud/Object Backup

S3-to-S3  
Replication



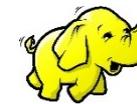
HPC File Systems

S3 API



OBS Replication/Backup

ILM / Tiering  
Replication / CRR



Hadoop

HDFS S3a  
backend



Backup

Backup to S3



Any S3  
compatible App

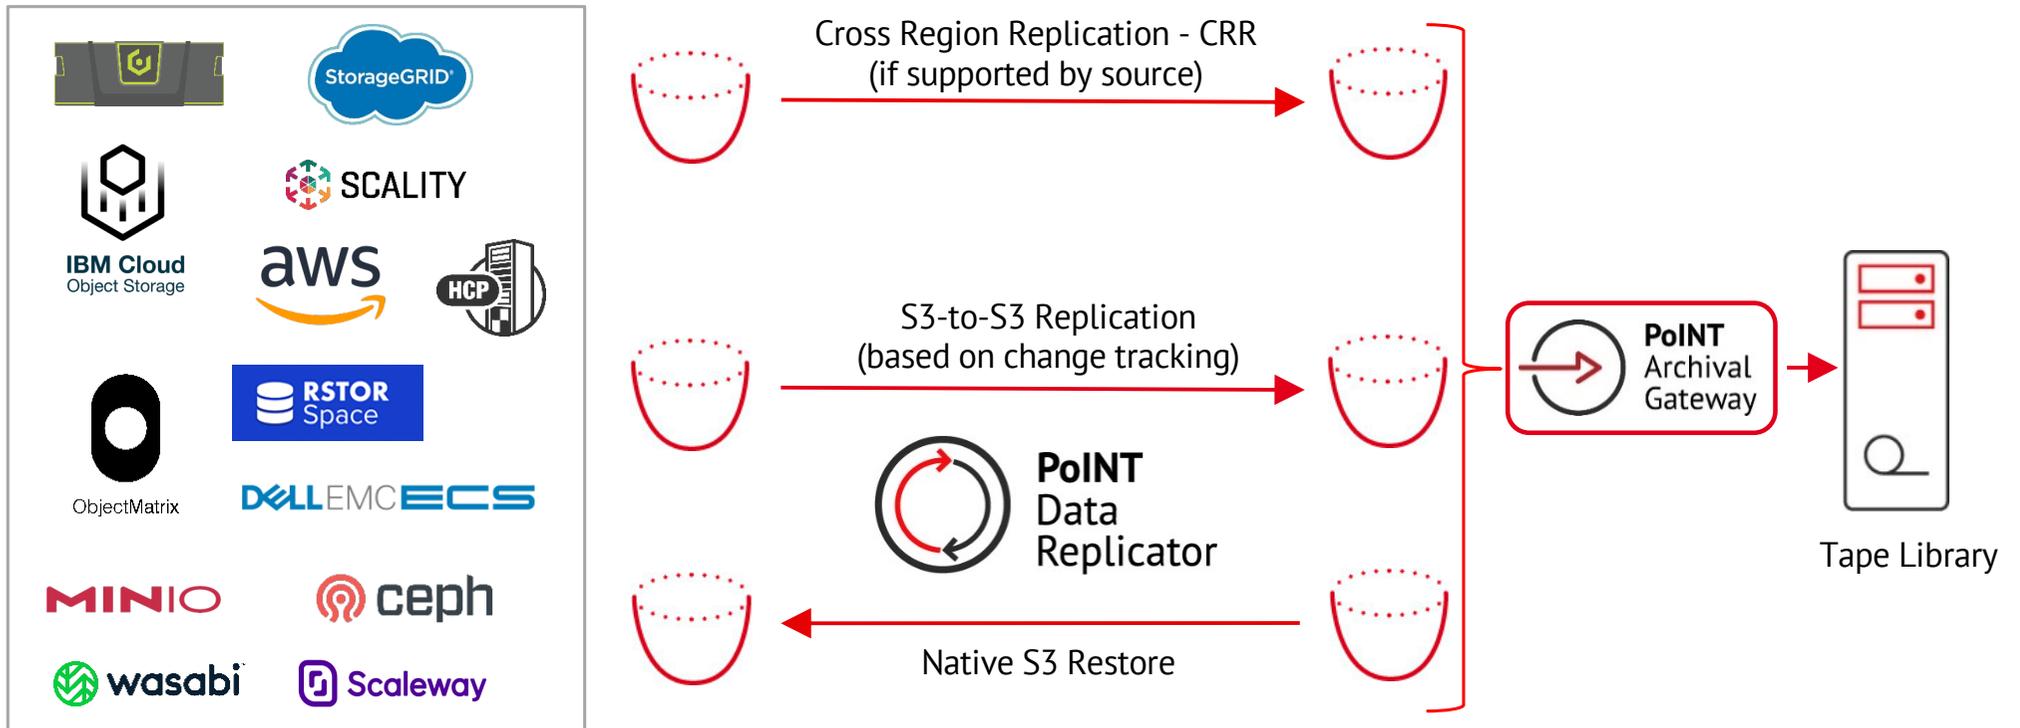
S3 API



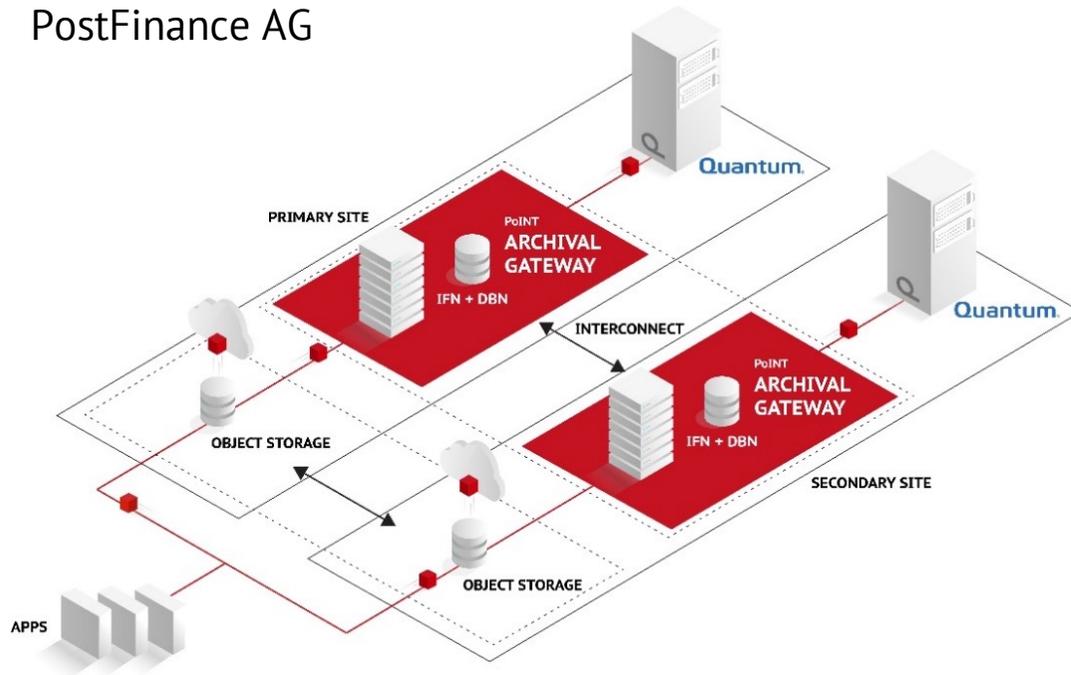
PoINT Archival Gateway

## Use Case

### Object/Cloud Backup to Tape



## Case Study PostFinance AG



### The challenges

- Minimize risk of data loss
- Increase data security through independent replication
- Cost-efficient storage

### The solution – PoINT Archival Gateway

- Introduction of an additional S3 storage class
- Automatic asynchronous data replication on object level
- Access to backup via native S3 API
- Scale capacity, performance and protection as needed

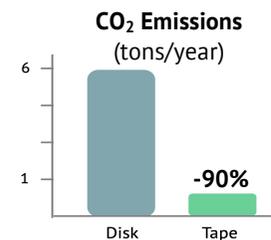
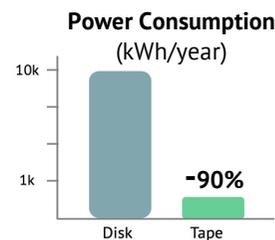
### The benefits

- Minimization of data loss risks
- Cost-effective storage on tape technology
- Flexible and simple expansion options
- Investment protection through manufacturer-independent software solution

## PoINT Archival Gateway

### Benefits

- Savings in storage costs (compared to all data on HDD)
- No ingest, egress and transaction costs (compared to public cloud)
- No vendor lock-in regarding storage hardware
- Data sovereignty (on-premises private cloud)
- Fulfillment of archiving and compliance requirements (by long-term data preservation)
- Protection from ransomware and cybercrime (“Air-gap” by tape)
- Investment protection by standardized interface (AWS S3)
- Sustainability by less energy consumption and lower CO<sub>2</sub> emissions (compared to all data on HDD)



---

**[www.point.de](http://www.point.de)**