Ceph News at 10

Do you know where your data is?

Ceph Cloud Transition and Beyond

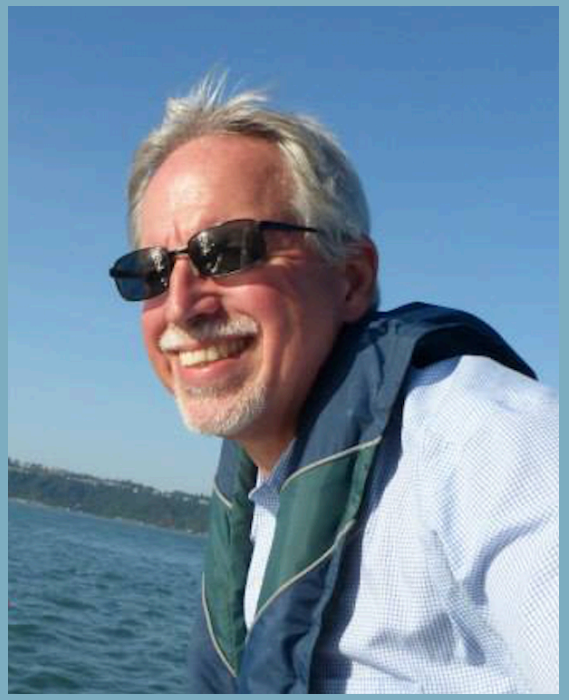

John Shubeck

jrshubeck@gmail.com

LinkedIn: jshubeck

Meet the presenter

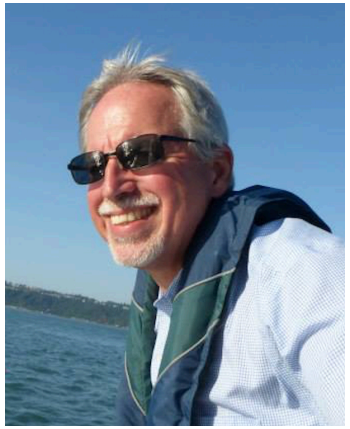

John Shubeck is an information technology professional with over 44 years of industry experience spanning both the customer and technology provider experience. John is currently serving as a Senior Storage Technical Specialist on IBM Object Storage platforms across all market segments in the Americas.

Email address: jshubeck@us.ibm.com
Mobile phone: 253-363-7827

Solid state drive (SSD) inflation curve – last 12 months[1]

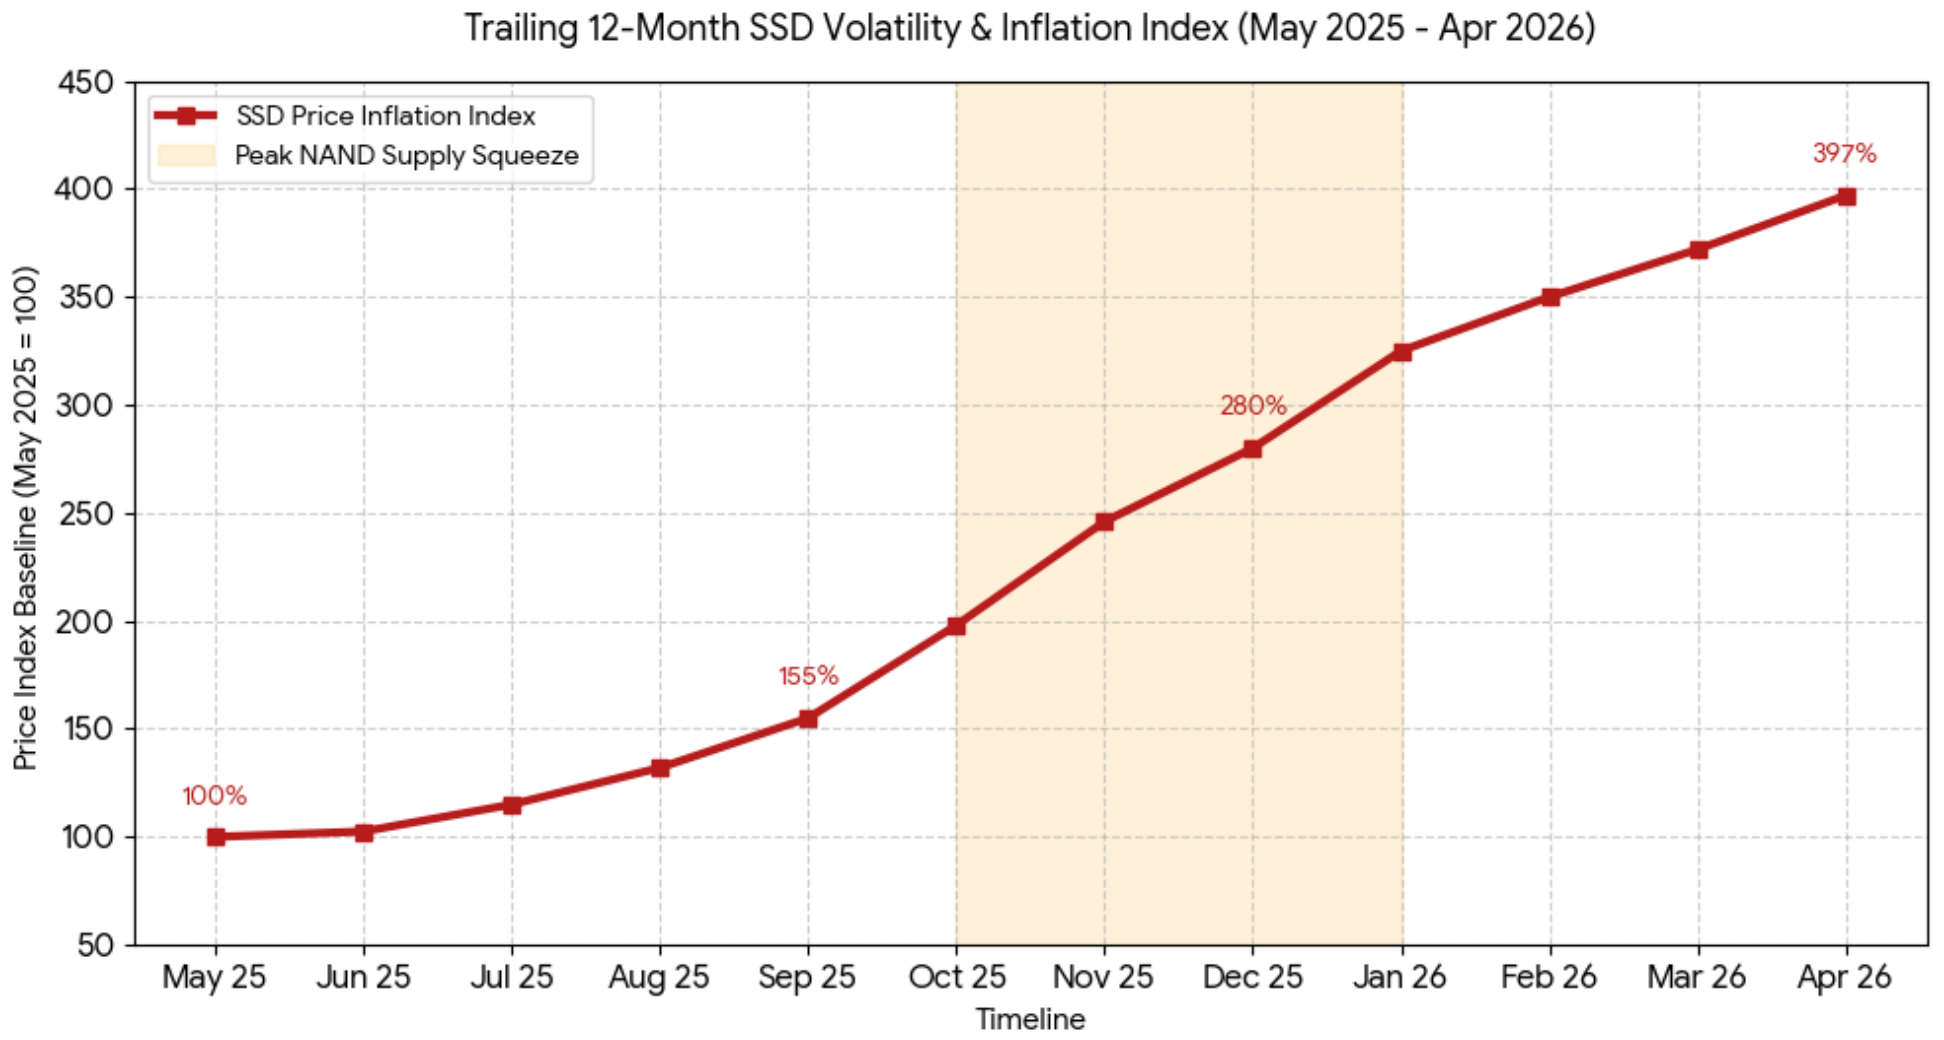

1. Source: Google Gemini AI Forbes (Overall increase since May 2026 = 472%)

Hard disk drive (HDD) inflation curve – last 12 months[1]

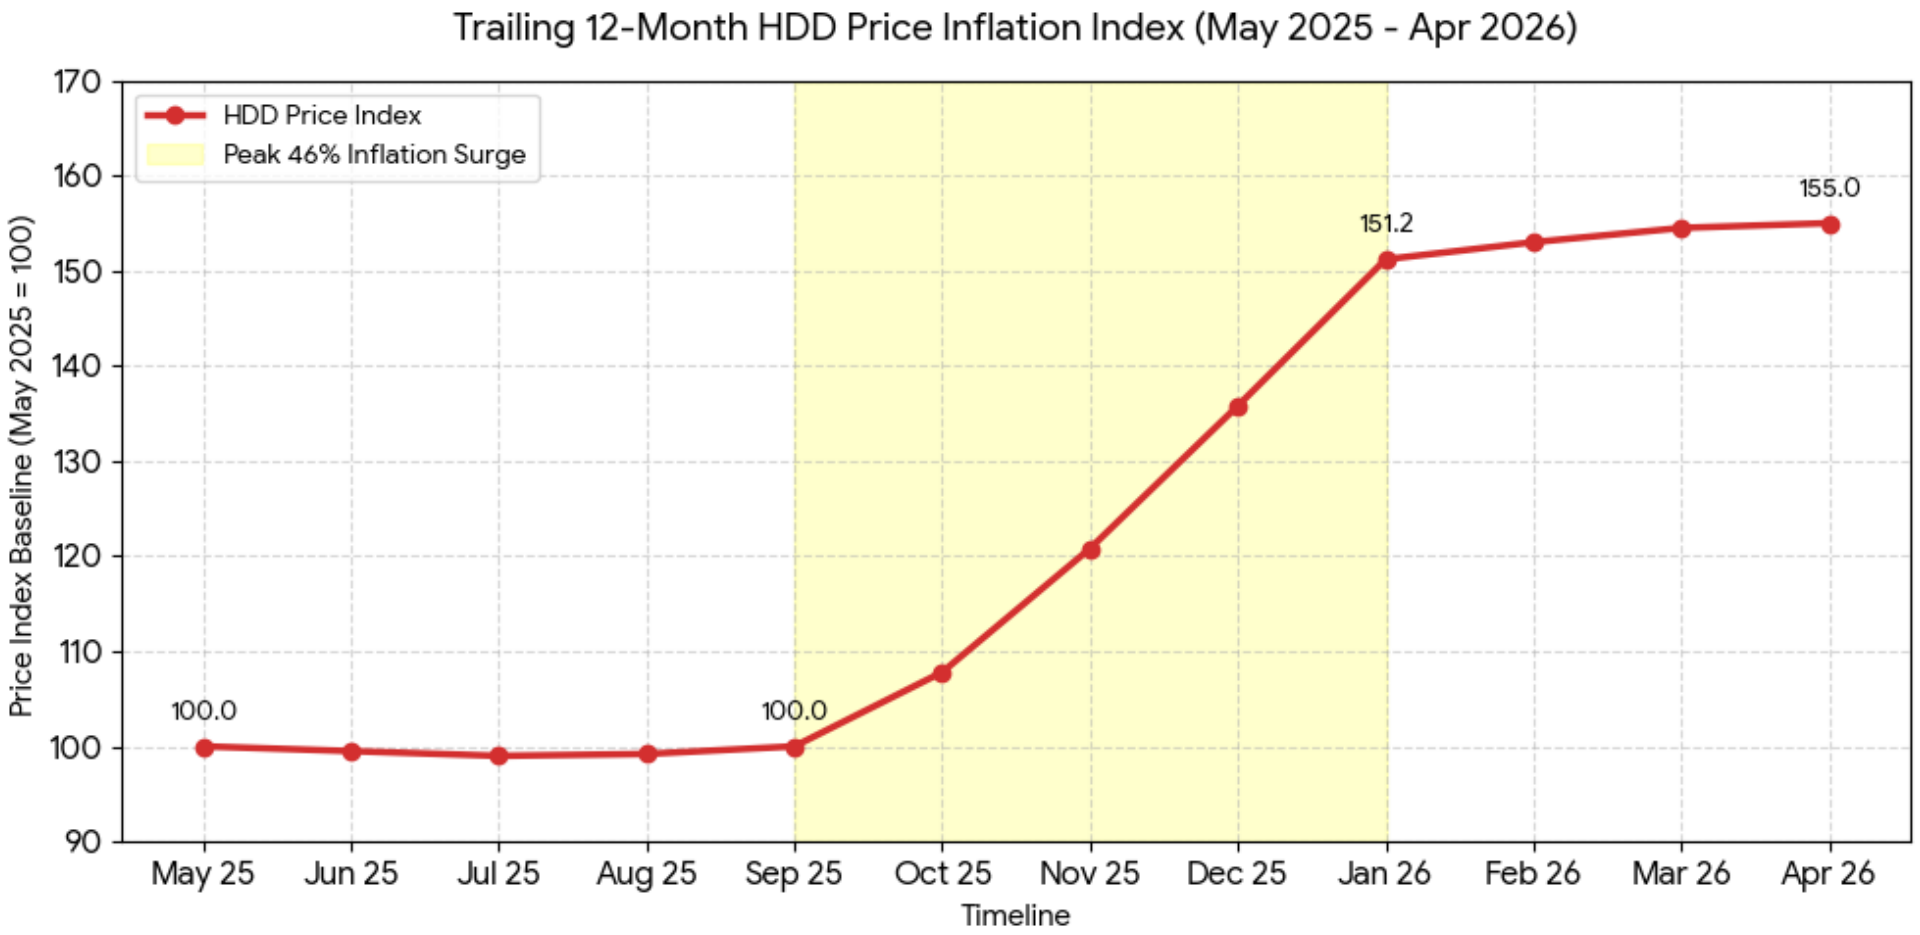

1. Source: Google Gemini AI Forbes (Overall increase since May 2026 = 55%)

Magnetic LTO tape inflation curve – last 12 months[1]

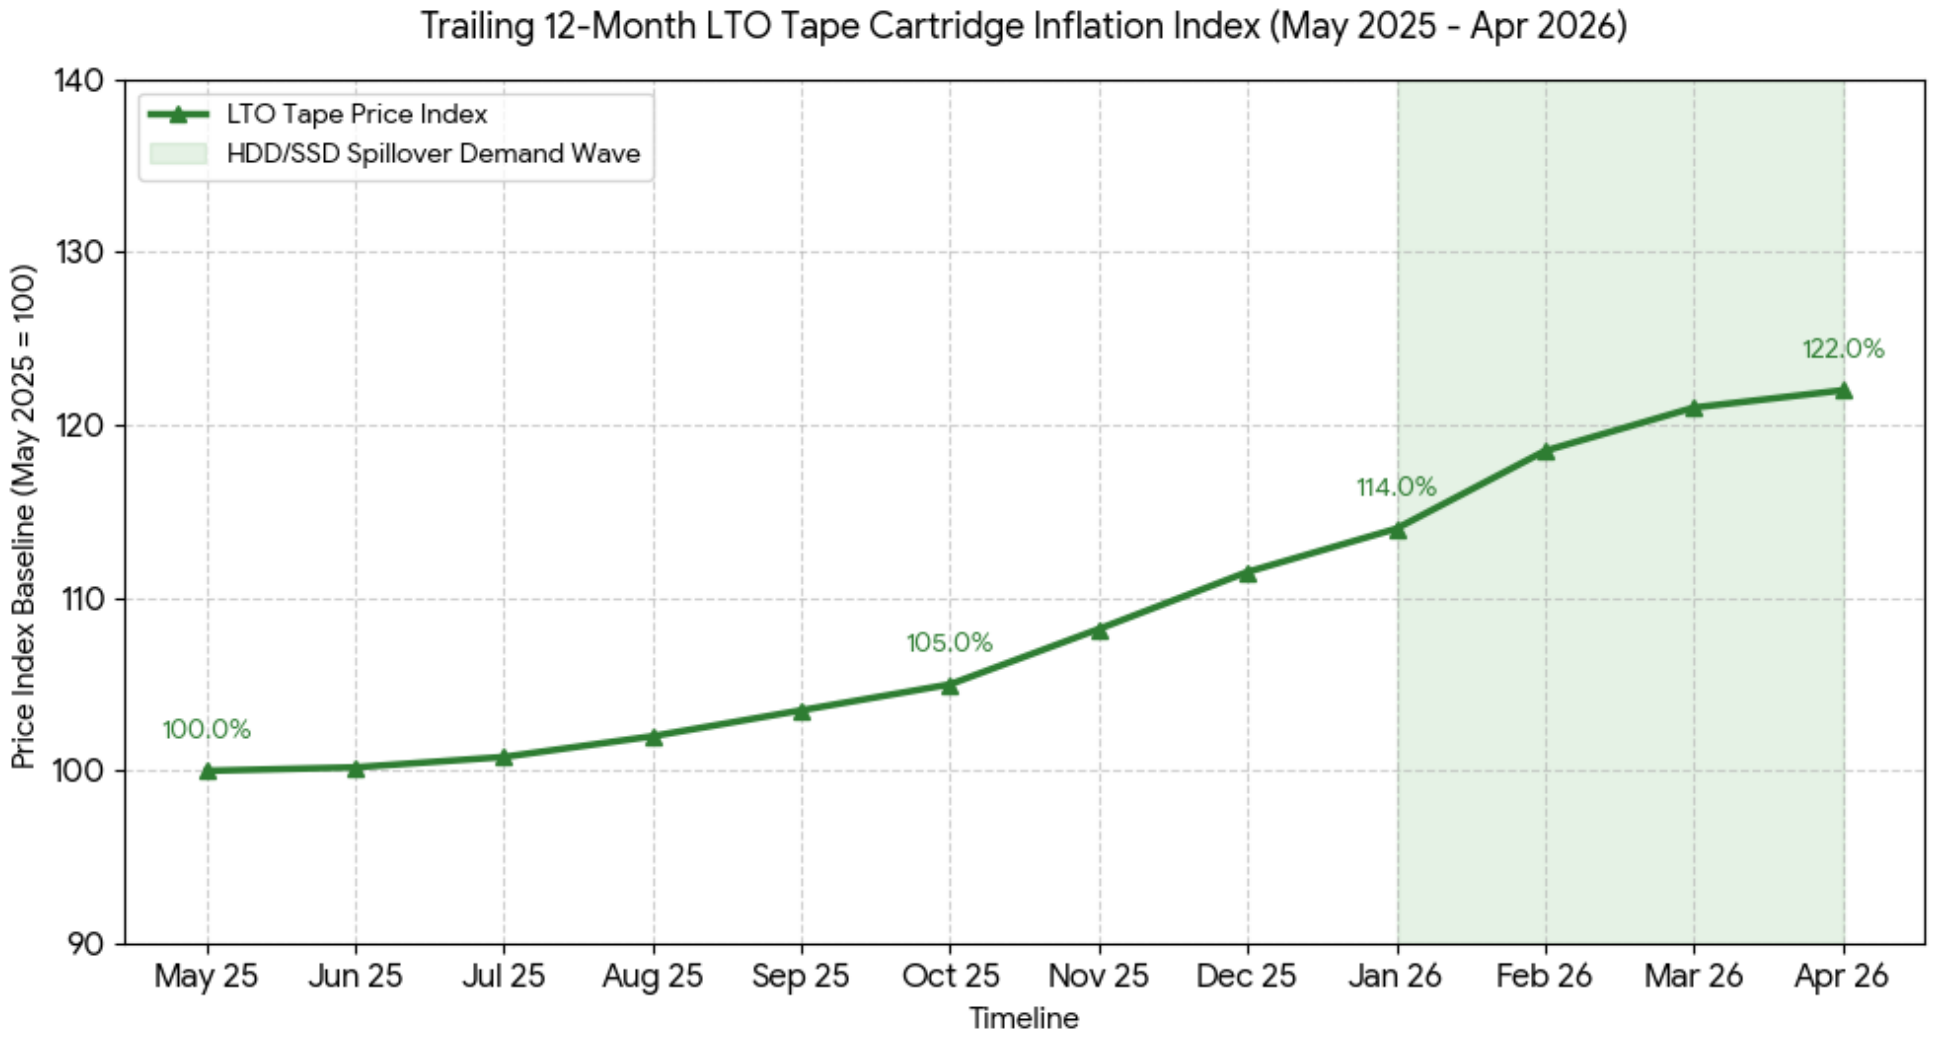

1. Source: Google Gemini AI *Forbes* (Overall increase since May 2026 = 22%)

Inflation rate comparison between SSD vs. HDD vs. LTO[1]

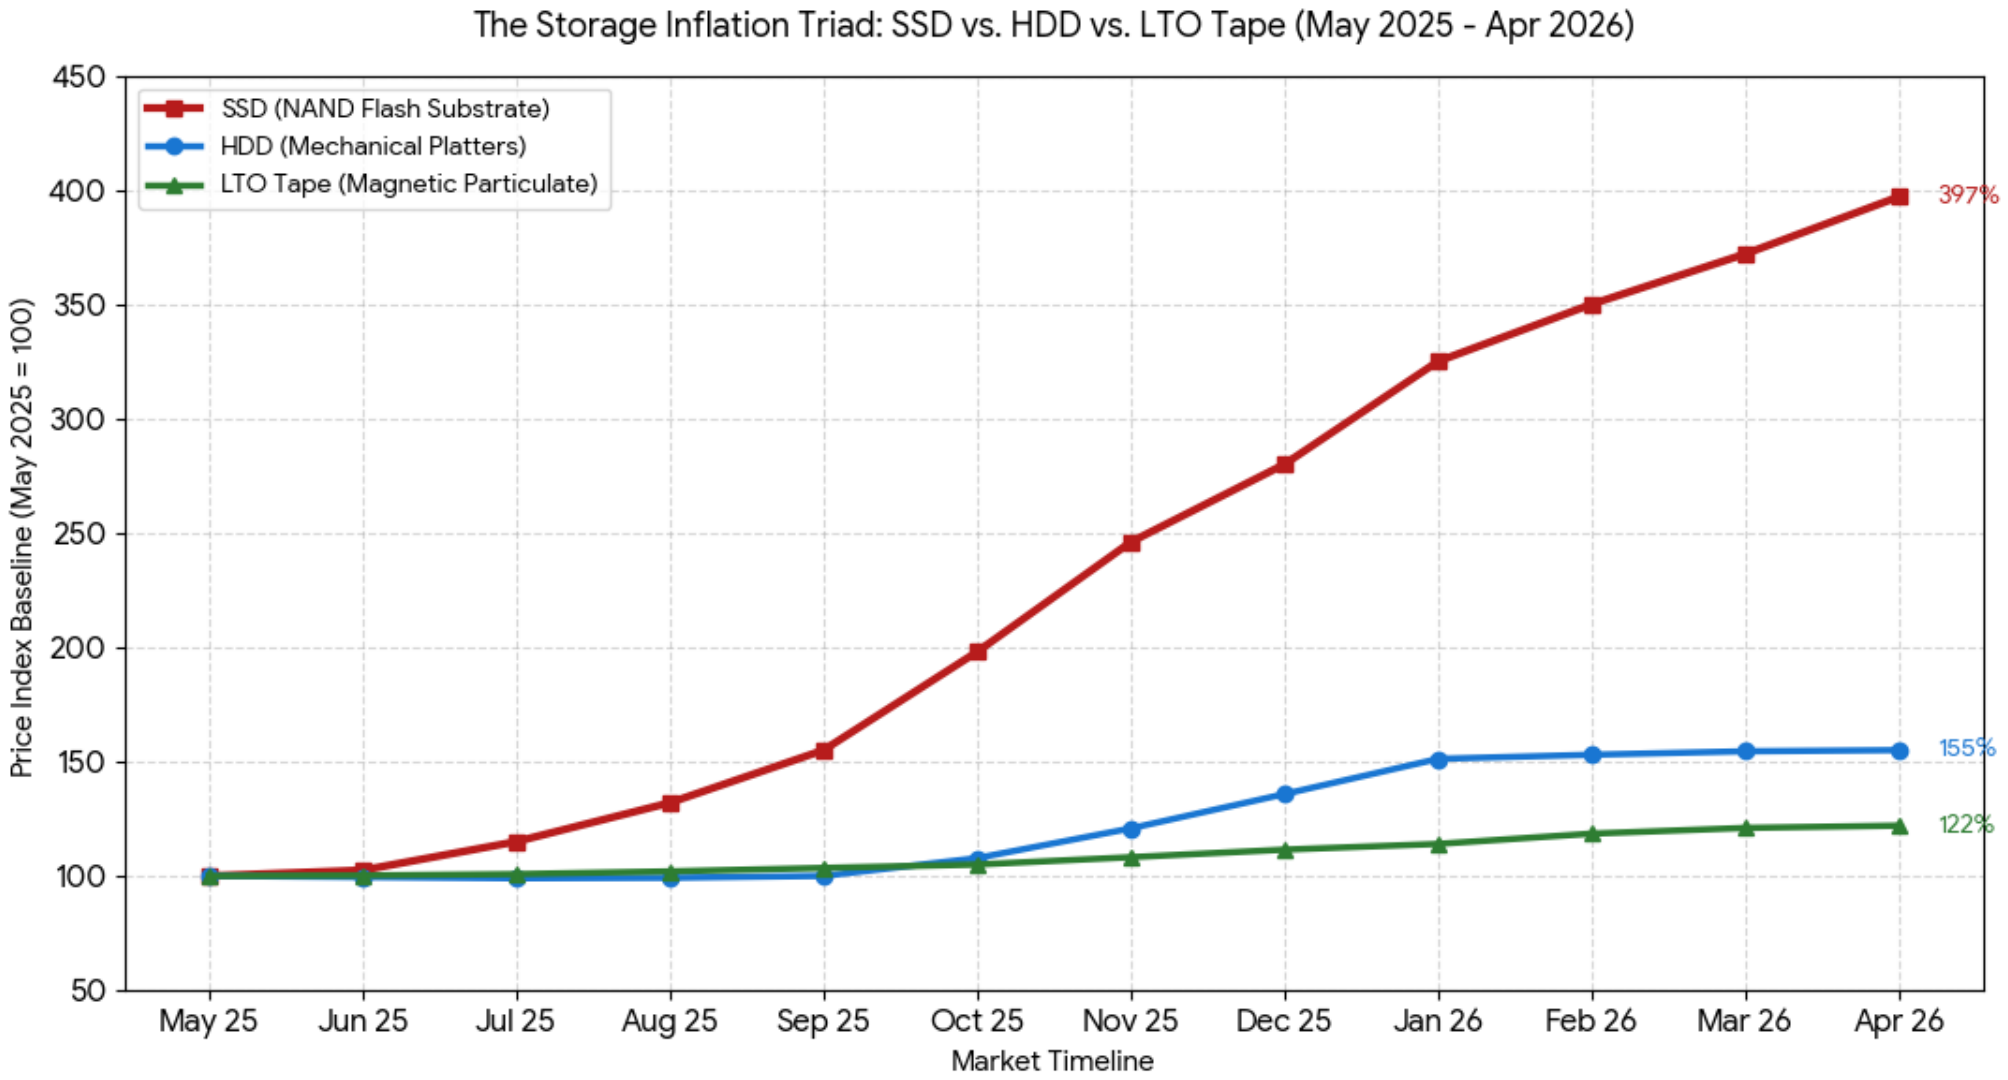

1. Source: Google Gemini AI *Forbes*

Inflation rate comparison between HDD vs. LTO[1]

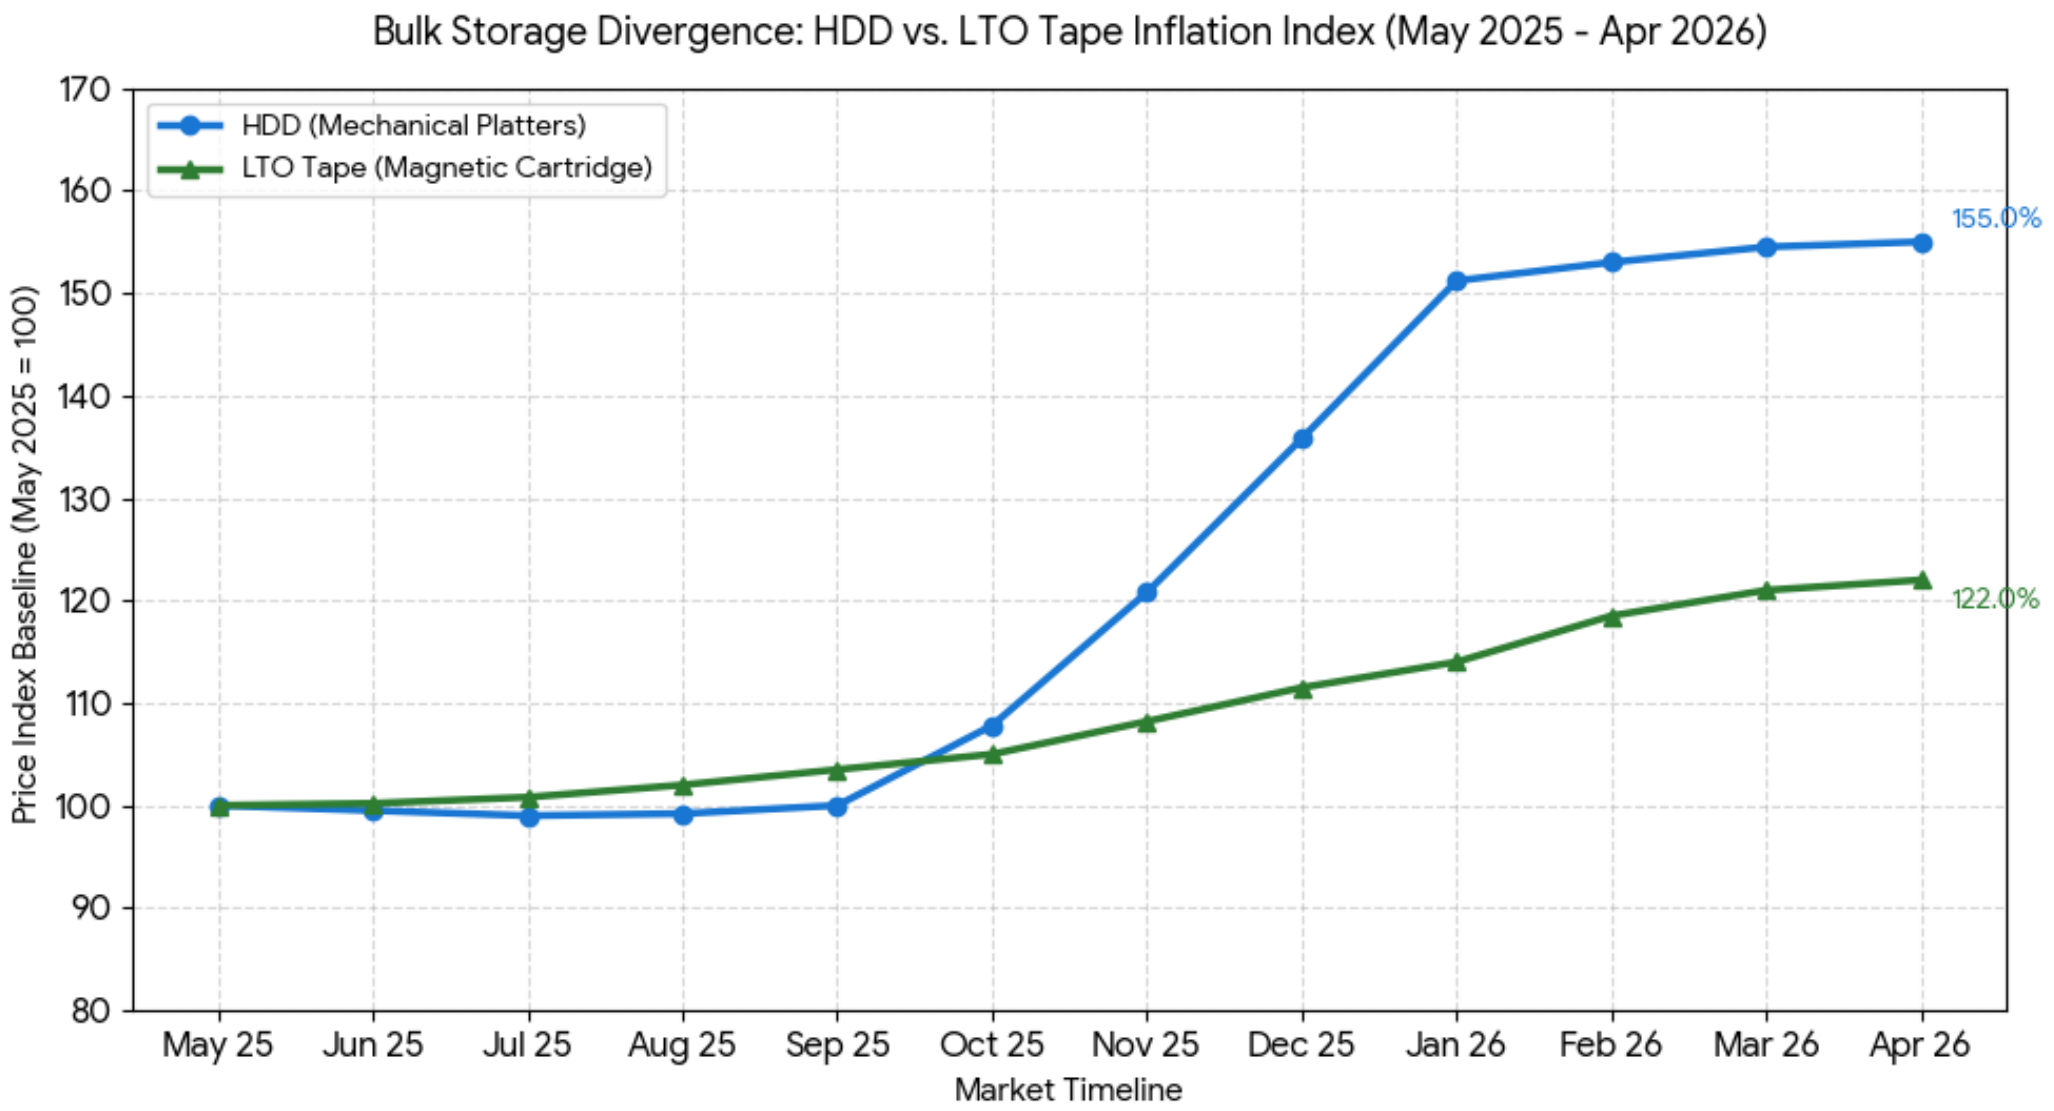

What can be done? Does anybody care?

Sometimes desperate measures are taken

But . . . Ceph object storage (RGW) has an “app” for that

- Storage class
- Protection rules
- Fast EC
- *Tiering*
- Compression
- *Cloud Transition!*
- Transparent recall
- *TAPE!*

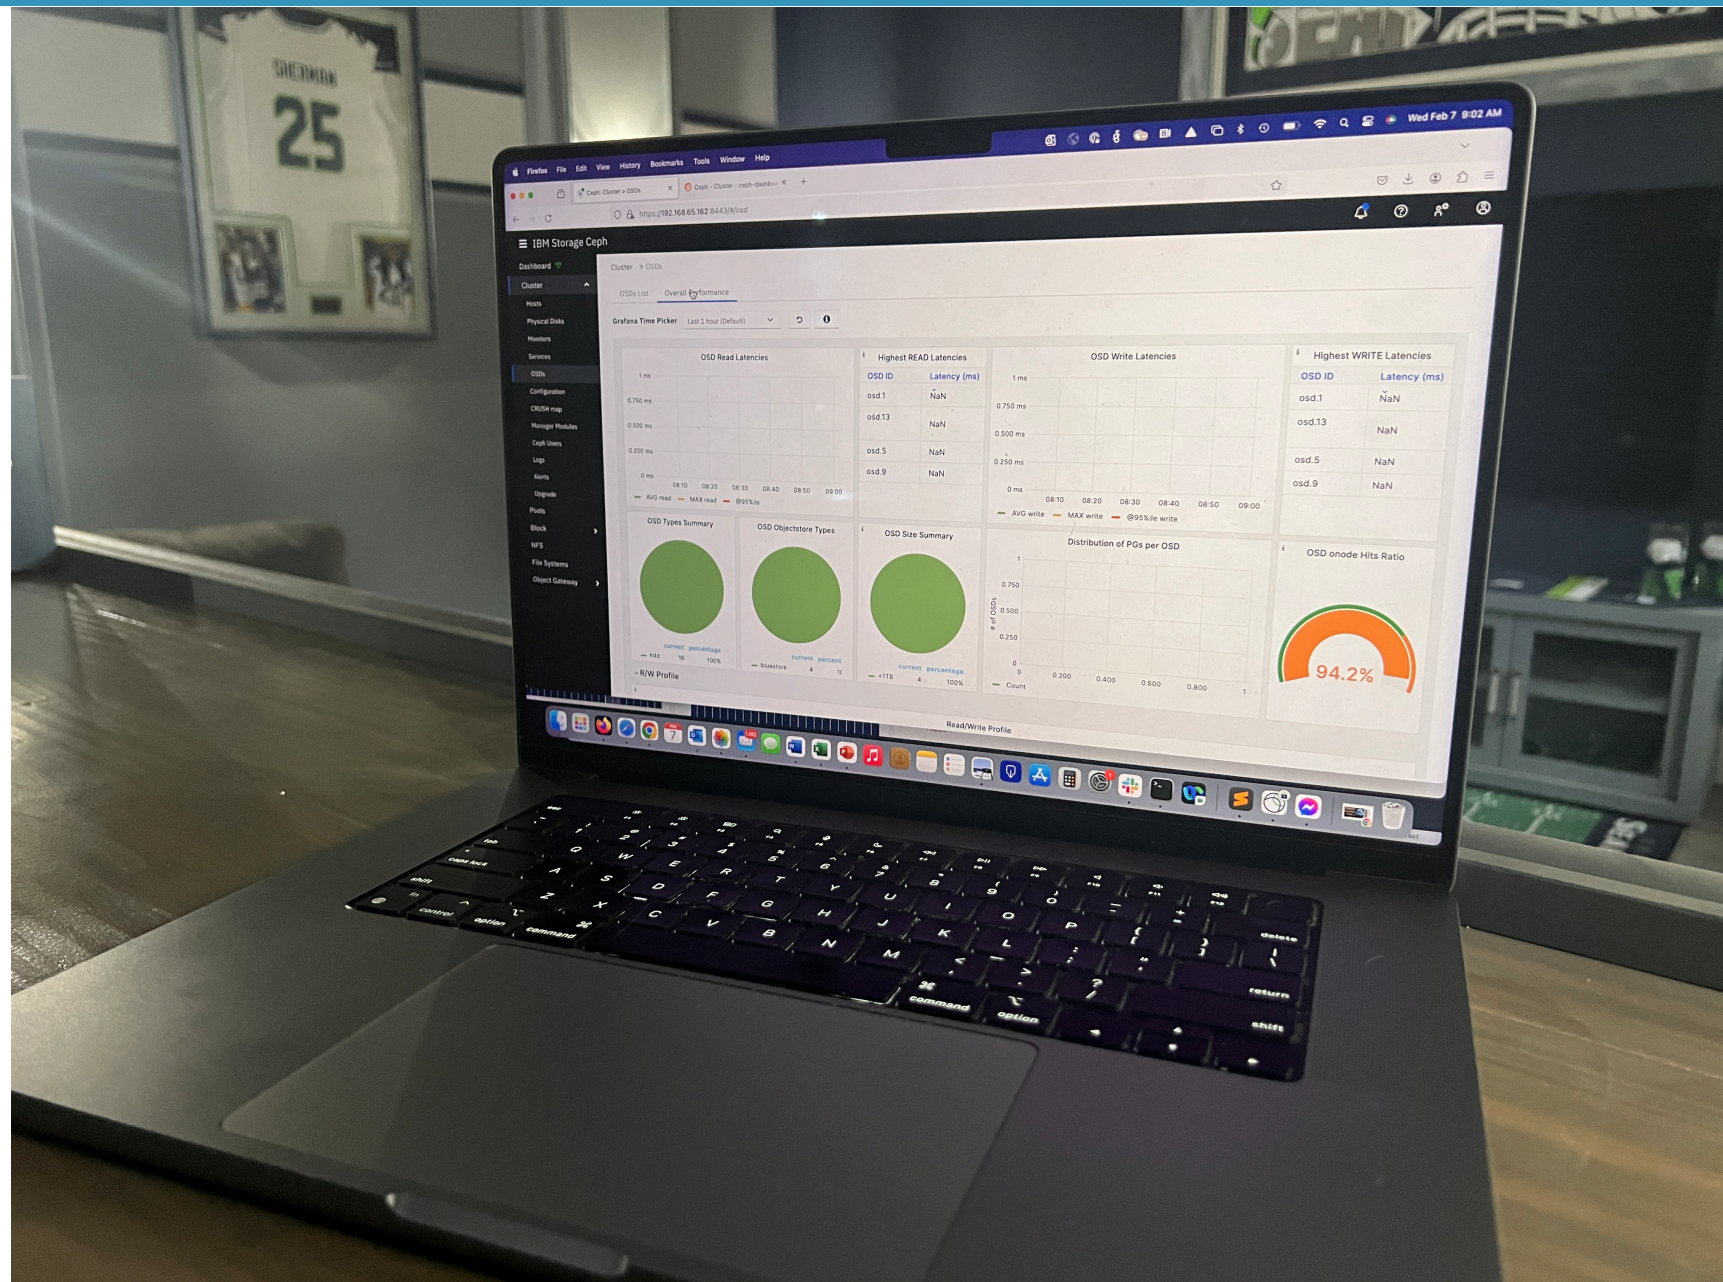

Cloud transition is an extension of bucket lifecycle

Bucket Tiering

- S3 Lifecycle generally means transitions across different storage tiers within the same cluster

Cloud Transition

- Cloud Transition moves the data out of the cluster to a cloud service

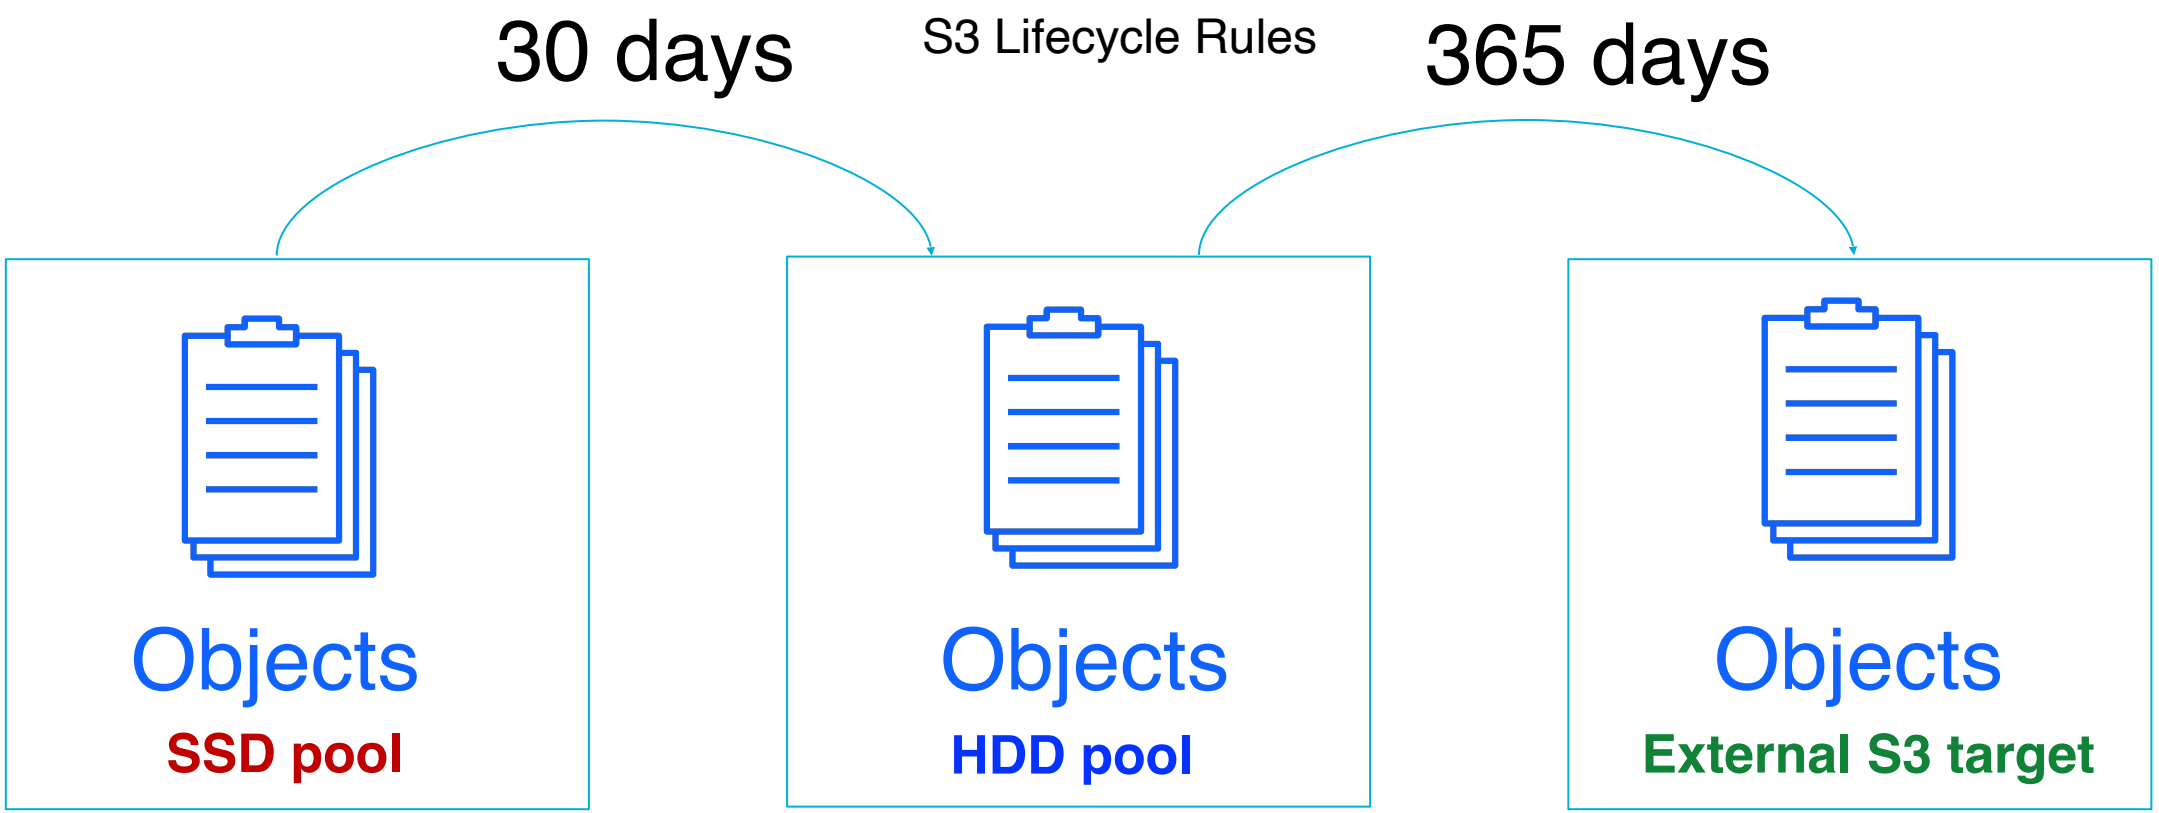

Guidelines for bucket tiering and cloud transition

Guidelines for bucket data storage tiers . . .

- Use Replica x3 for bucket index and bucket metadata
- Use a higher performant Erasure Code profile for bucket data (e.g. Rep x3, EC 2+2 or EC 4+2)
- Use a higher efficiency Erasure Code profile for cold data (e.g. EC 8+3)
- Use a cloud transition for Exabyte scale or GLACIER storage class

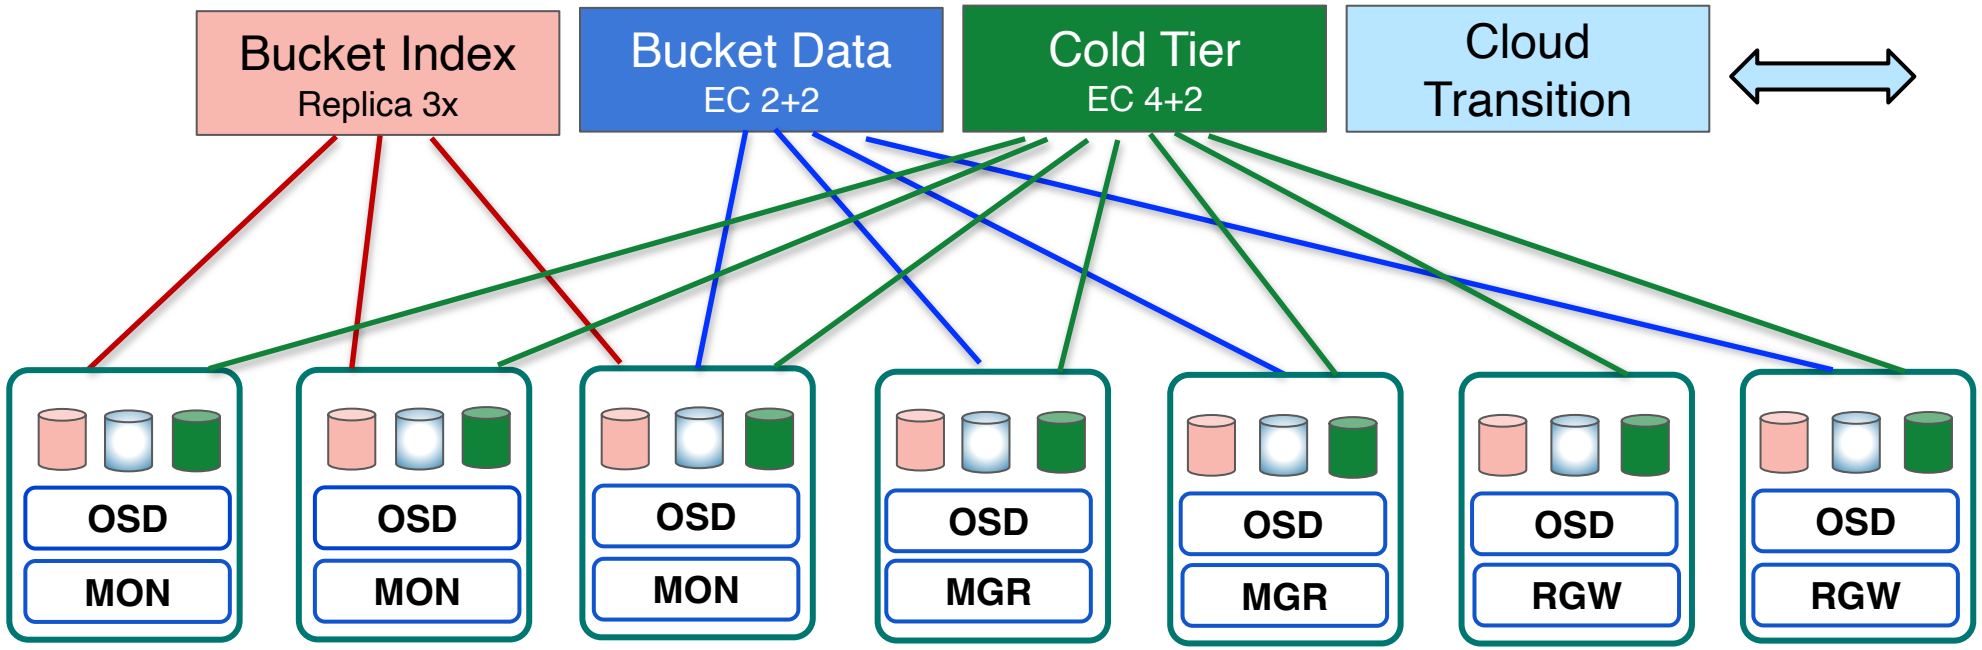

Note: For illustrative purposes. Nodes and full disks are not dedicated to any single pool

The what and why of Ceph cloud transition

This feature enables data movement to a remote cloud service as part of S3 Lifecycle Configuration via Ceph Storage Classes.

Features:

- The goal of this feature is to enable data transition to multiple cloud providers
- The currently supported cloud providers are those that are compatible with AWS (S3)

Benefit:

- IBM Storage Ceph + Cloud Transition is . . .
  . . . a virtually INFINITE object storage capacity

Why use cloud transition?

Ceph S3 object tiering to cloud is a bucket lifecycle rule

IBM Storage Ceph

- Software defined
- Hardware independent
- Highest S3 fidelity
- Block, file, and object services
- Multiprotocol S3 and NFS single namespace
- Multiple tiers internally (i.e. NVMe -> HDD)
- Transitions to multiple clouds including tape

Other on-prem dense S3

- Software defined
- Hardware optimized
- Exabyte scale object store
- High density storage nodes
- Geo-dispersed Erasure Coding
- Hundreds of petabytes per admin
- Extreme immutability

Cloud S3 (AWS, Google, IBM)

- Software defined
- Hardware independent
- Subscription (pay as you go)
- Data access from other cloud services
- Tiering within the realm of the cloud provider

IBM Deep Archive

- Software defined
- Hardware specific (purpose built tape library)
- High density
- Pseudo air gap
- Lowest carbon footprint
- Not (today) facing a global supply chain shortage

Wait, what? Did you say . . . TAPE?

For illustrative purposes. Not this kind of tape . . .

This kind of tape!

IBM Deep Archive[1]

Quantum Scalar i7[2]

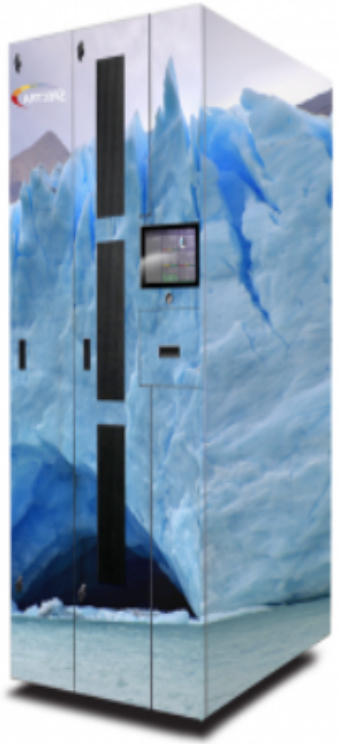

Quantum Scalar i7[3]

1. https://www.ibm.com/products/deep-archive
2. https://www.quantum.com/en/products/tape-storage/scalar-i7/
3. https://spectralogic.com/solutions/object-based-tape/

A high altitude view of Ceph Cloud Transition

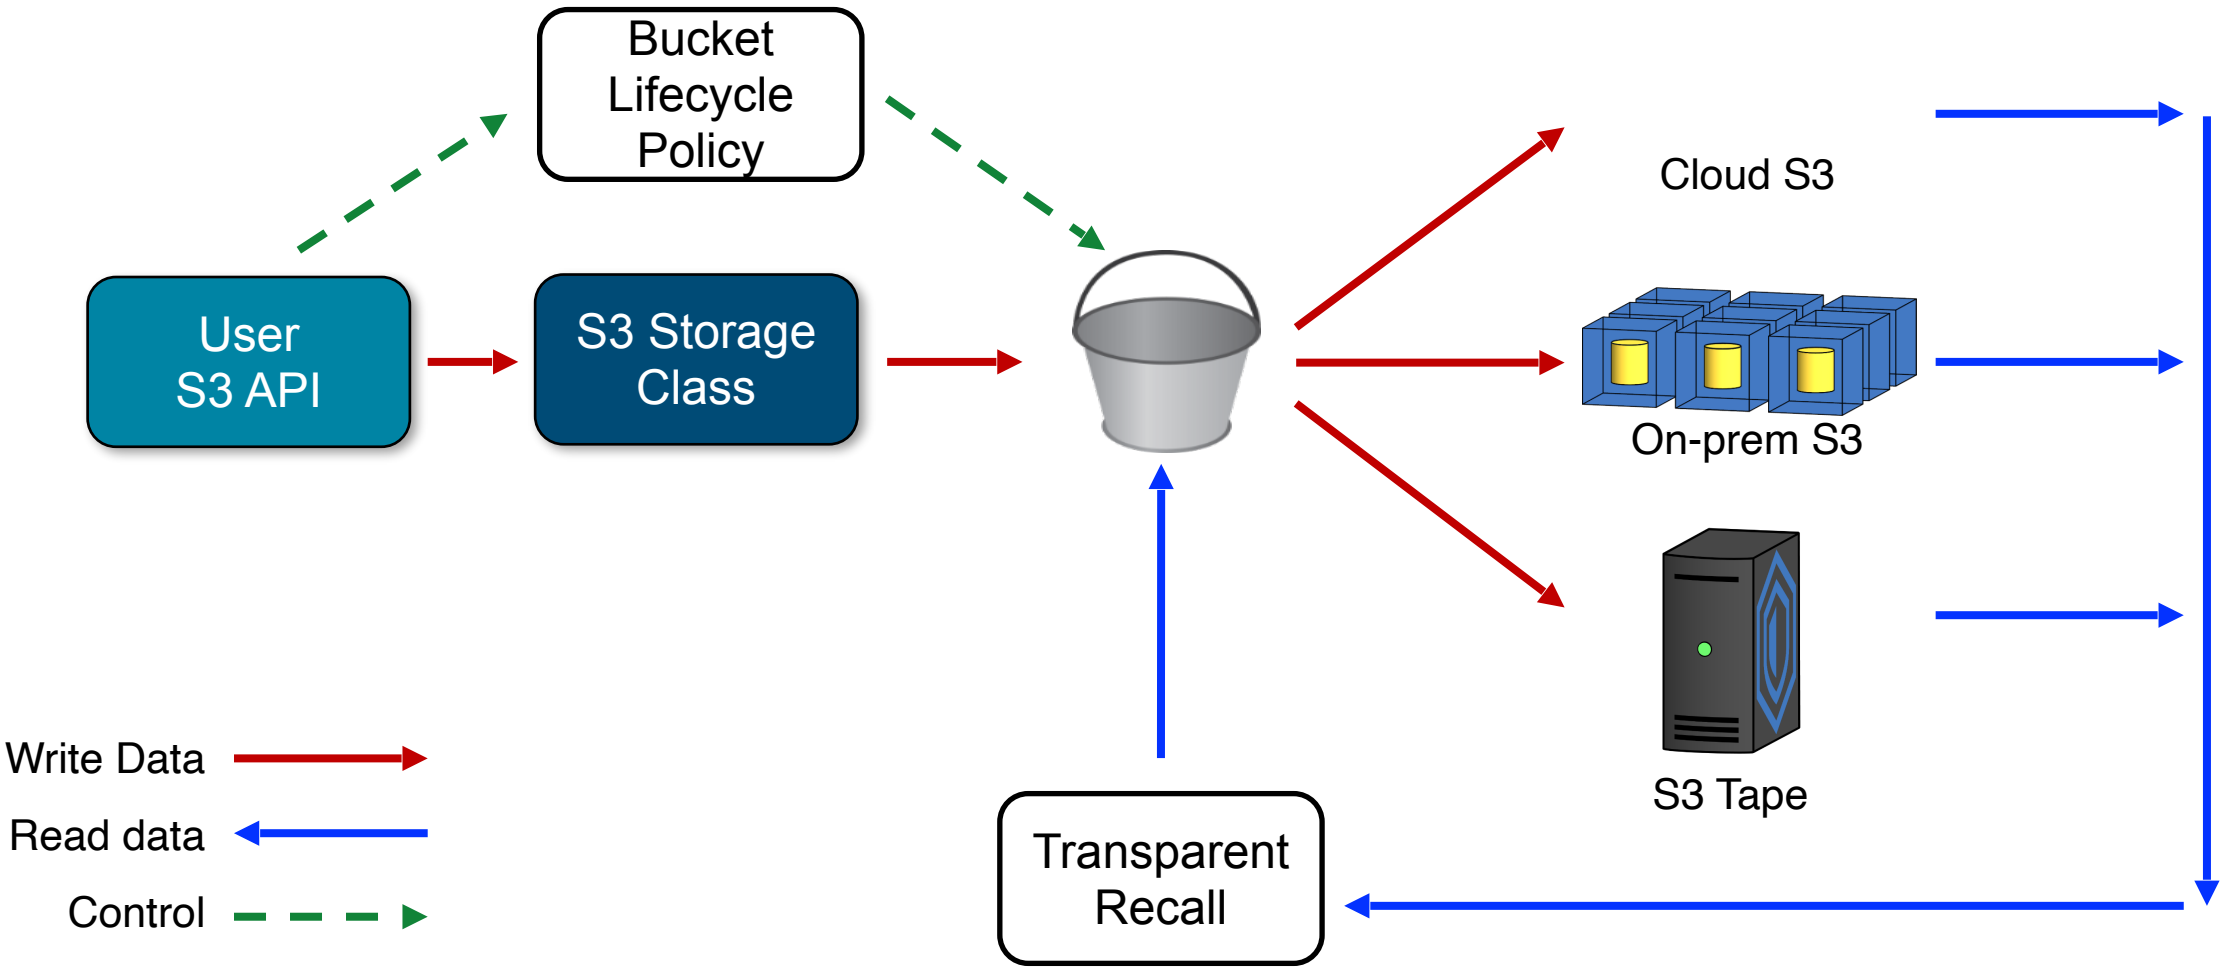

Ceph cloud transition phases and tasks

Administrator tasks

Add a new storage class "**GLACIER**" in the Zonegroup

Map the Storage Class "GLACIER" to **IBM Deep Archive**

Granular settings in the Cephadm CLI

S3 client tasks

Upload objects to the S3 bucket

Apply a lifecycle rule to the S3 bucket

Validate capacity space relief and object meta data

S3 client verification tasks

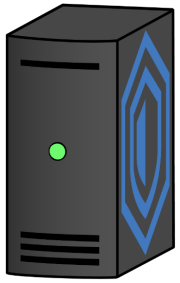

Transparent recall (**RESTORE, HEAD, GET**)

Restored!

A FEW “DAYS” LATER . . .

RGW cloud transition commands for object storage

```
# Display the default Zonegroup storage class listing

# radosgw-admin zonegroup get --rgw-zonegroup=default \
| jq .placement_targets[0].storage_classes

[
  "STANDARD"
]

# Define a new storage class in the Zone Group
# radosgw-admin zonegroup placement add --rgw-zonegroup default  \
                                --placement-id default-placement \
                                --storage-class AWS \
                                --tier-type=cloud-s3

. . . Output omitted . . .
```

RGW cloud transition commands for object storage

```
# Display the default Zonegroup storage class listing

# radosgw-admin zonegroup get --rgw-zonegroup=default \
| jq .placement_targets[0].storage_classes

[
  "STANDARD"
]

# Define a new storage class in the Zone Group
# radosgw-admin zonegroup placement add --rgw-zonegroup default  \
                                --placement-id default-placement \
                                --storage-class AWS \
                                --tier-type=cloud-s3

. . . Output omitted . . .
```

RGW cloud transition commands for object storage

```
# Display the default Zonegroup storage class listing

# radosgw-admin zonegroup get --rgw-zonegroup=default \
| jq .placement_targets[0].storage_classes

[
 "AWS",
 "STANDARD"
]


# Display the Storage Class and placement settings

# radosgw-admin zonegroup get

. . . Output omitted . . .
```

RGW cloud transition commands for object storage

```
# Zonegroup placement tier configuration settings

# radosgw-admin zonegroup placement modify --rgw-zonegroup default  \
--placement-id default-placement \
--storage-class AWS \
--tier-config='endpoint=https://s3.amazonaws.com' \
--tier-config='access_key=AKIAJXABCDEHPSPFIO3Q' \
--tier-config='secret=hl13577MtPKFtjhZNtg1XYZZYEMCDR6OCUhdjoa13' \
--tier-config='target_path=jrs-ceph-archive' \
--tier-config='retain_head_object=true' \
--tier-config='allow_read_through=true' \
--tier-config='read_through_restore_days=1’ \
--tier-config='host_style=virtual ' \
--tier-config='region=us-east-1 ' \
--tier-config='target_storage_class=STANDARD'


. . . Output omitted . . .
```