OpenIO

Object Storage
for BigData & HPC

OpenIO Object Storage

We think different

Directory with indirections

Track containers and not objects

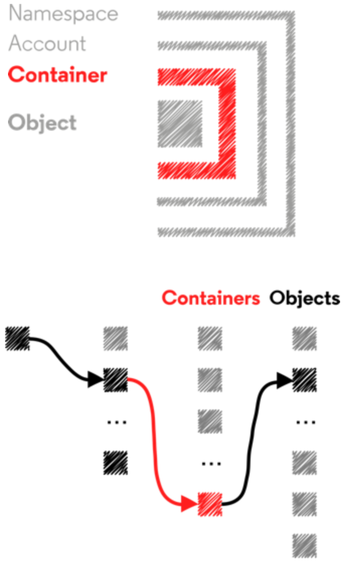

ConsciousGrid technology

Real time LB for optimal data placement.
Never rebalance

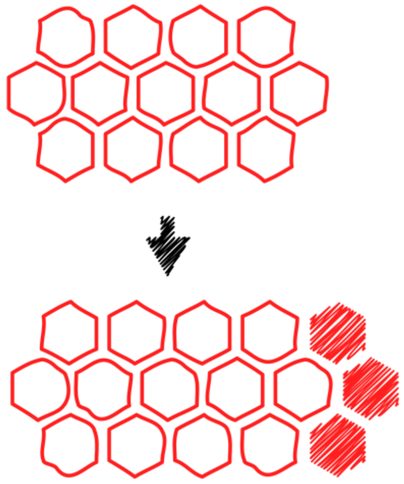

Open Source

Avoid vendor lock-in and Improve freedom of choice.

Keep control of your data

Flexibility!

Heterogeneous platform

Multi-tiers

Data Access

Amazon S3

OpenStack Swift

Filesystem connector (FUSE-based)

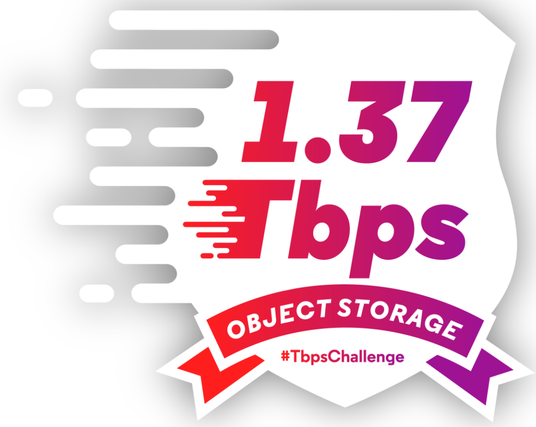

Storage Landscape

Object Storage is the Sweet Spot

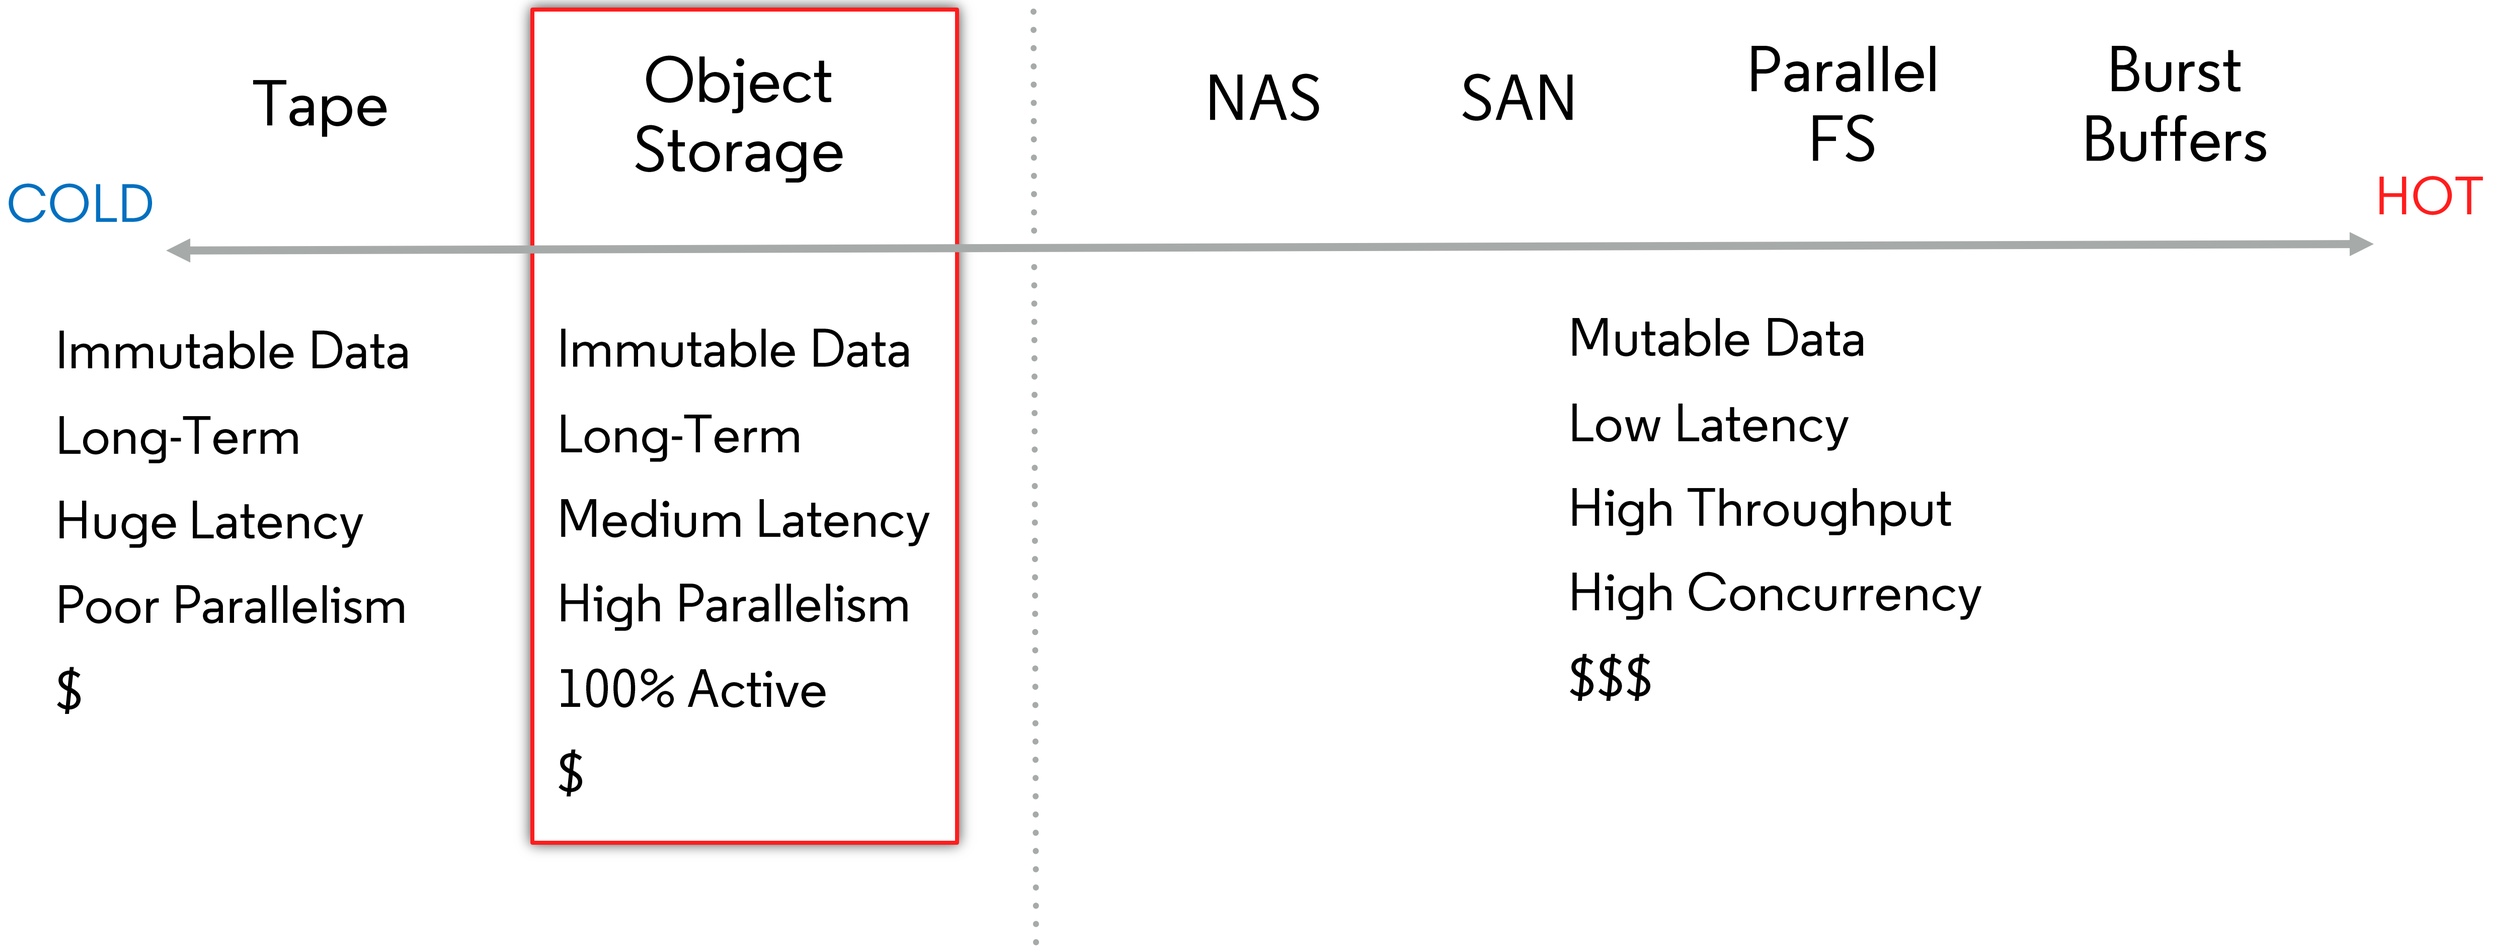

Working with HDF Kita

Because budgets are finite!

The HDF5 community has the same TCO problem

How do you manage the "chasm of data immutability" to let HDF5 benefit from Object Storage?

So HDF Kita was born to map HDF5 files to Objects
1. Start prototyping on a Public Cloud
2. Use the same code on premises

We made the stack efficient with a tight integration...

Working with HDF Kita

Tight Integration

Both stacks merged

Same deployment tools

Same scalability patterns, redundancies removed

Special attention to the efficiency of network aspects

Kita+OpenIO is the most direct way to a TCO-efficient HDF5 usage at scale.