

# *Towards an exabyte-scale unified data platform for research data*

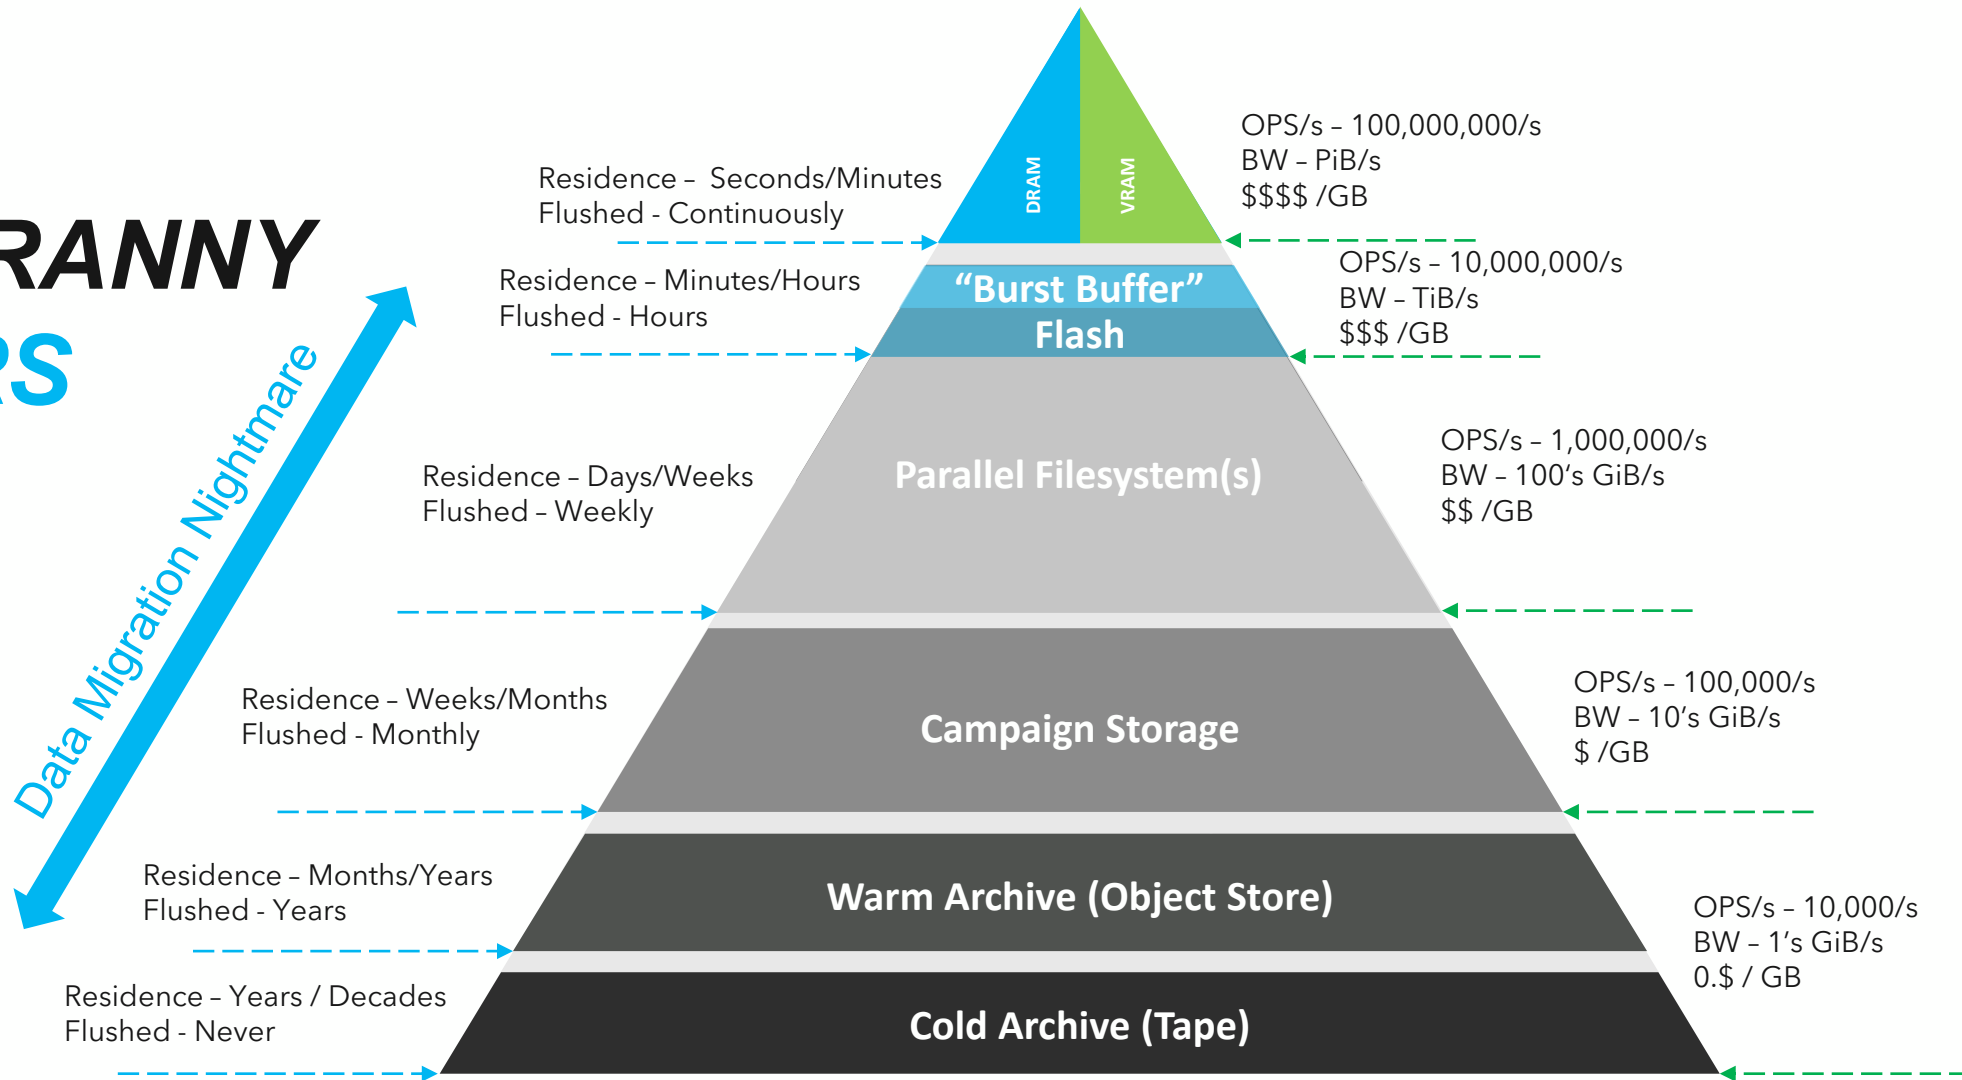
KQ Codes Technical Social | Wed 18th May 2022

**Ray Coetzee – Solution Architect Lead & “Recovering Parallel Filesystem Engineer”**


**Ross Cooper-Smith – Senior Solution Architect & “Recovering Enterprise Storage Engineer”**



# THE TYRANNY OF TIERS




# LOW-COST FLASH WILL BENEFIT FROM CONTINUED INNOVATION...




# NEW AI WORKLOADS ARE **OPPOSITE** FROM HPC WORKLOADS

Classic HPC Simulation I/O Pattern



Numerically-Intensive Scientific Computing

New Artificial Intelligence I/O Pattern



Data-Intensive Deep Learning Training

CREDIT: DR. ENG LIM GOH (HPE)

# THE COST OF DEEP LEARNING


“While advances in hardware and software have been driving down AI training costs by 37% per year, the size of AI models is growing much faster, 10x per year. As a result, total AI training costs continue to climb. We believe that state-of-the-art AI training model costs are likely to increase 100-fold, **from roughly \$1 million today to more than \$100 million by 2025.**”

ARK CAPITAL  
BIG IDEAS 2021 REPORT



**10,000X  
GROWTH  
IN 4 YEARS**

# TRAINING COSTS ARE EXPLODING







**1MB**



Feature  
extraction  
algorithm



**100KB**





VS

Nvidia H100



@1MB Sequential: **140MB/s**  
@100KB Random Read: **10.5MB/s**

**128,000 MB/s**  
**Bi-Directional**



VS

Nvidia H100



**8,000 HDDs REQUIRED FOR A SINGLE MODERN AI PROCESSOR**

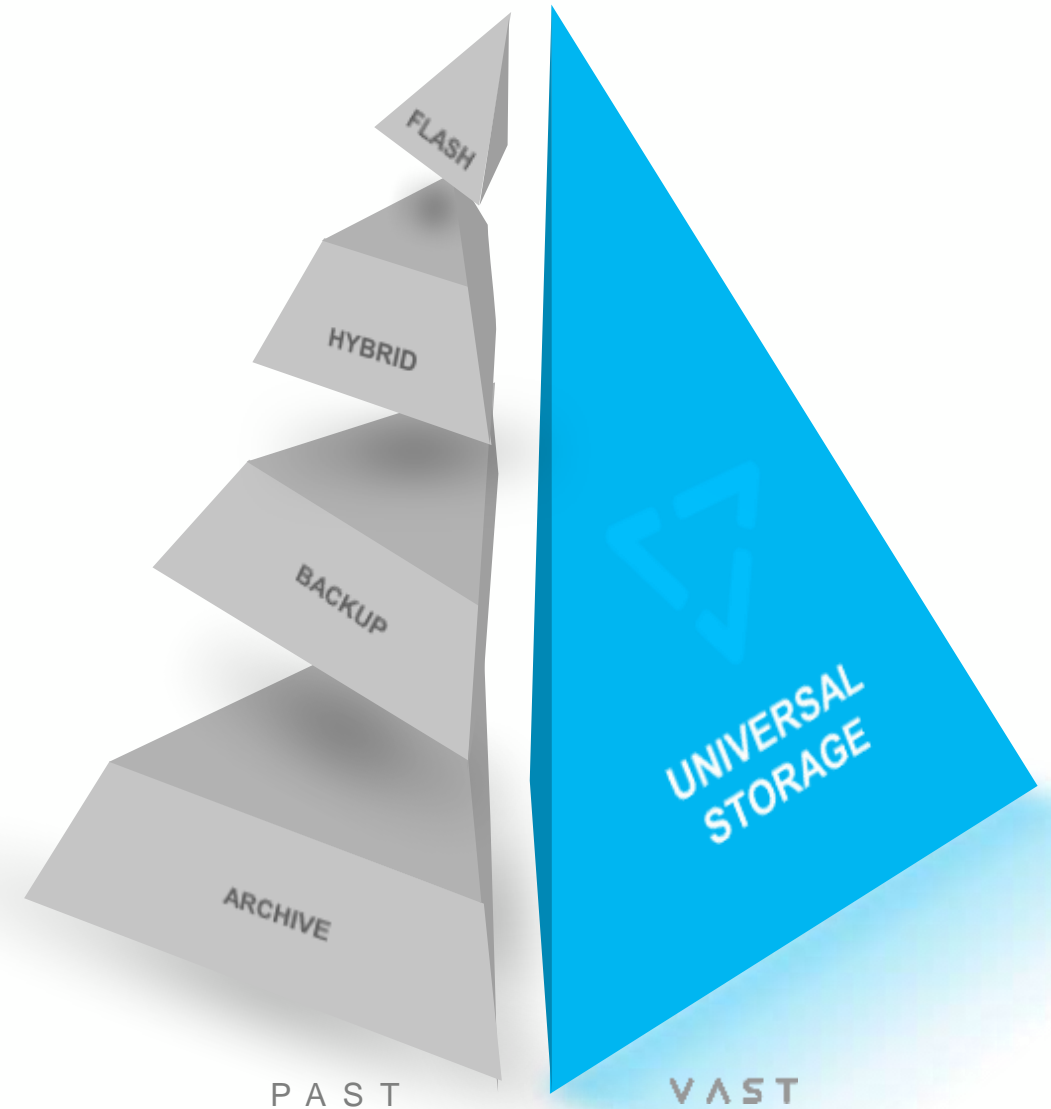
OUR MISSION

# ***NO MORE TIERS***

## **BREAKING TRADEOFFS TO MAKE STORAGE SIMPLE**

- Extinction-level Event for The Hard Drive
- An End To 30 Years of Complex Storage Tiering
- Unleash Big Data and ML Insights

VAST



UNIVERSAL STORAGE

# **STARTING FROM A RENAISSANCE IN HARDWARE**



## **NVME OVER FABRICS FOR DISAGGREGATION**

The latency of DAS, over switched commodity networks.



## **QLC FLASH FOR COST SAVINGS**

Low-cost, lower-endurance hyperscale flash.



## **STORAGE CLASS MEMORY LONGEVITY & EFFICIENCY**

Enables write shaping to QLC and rich metadata.

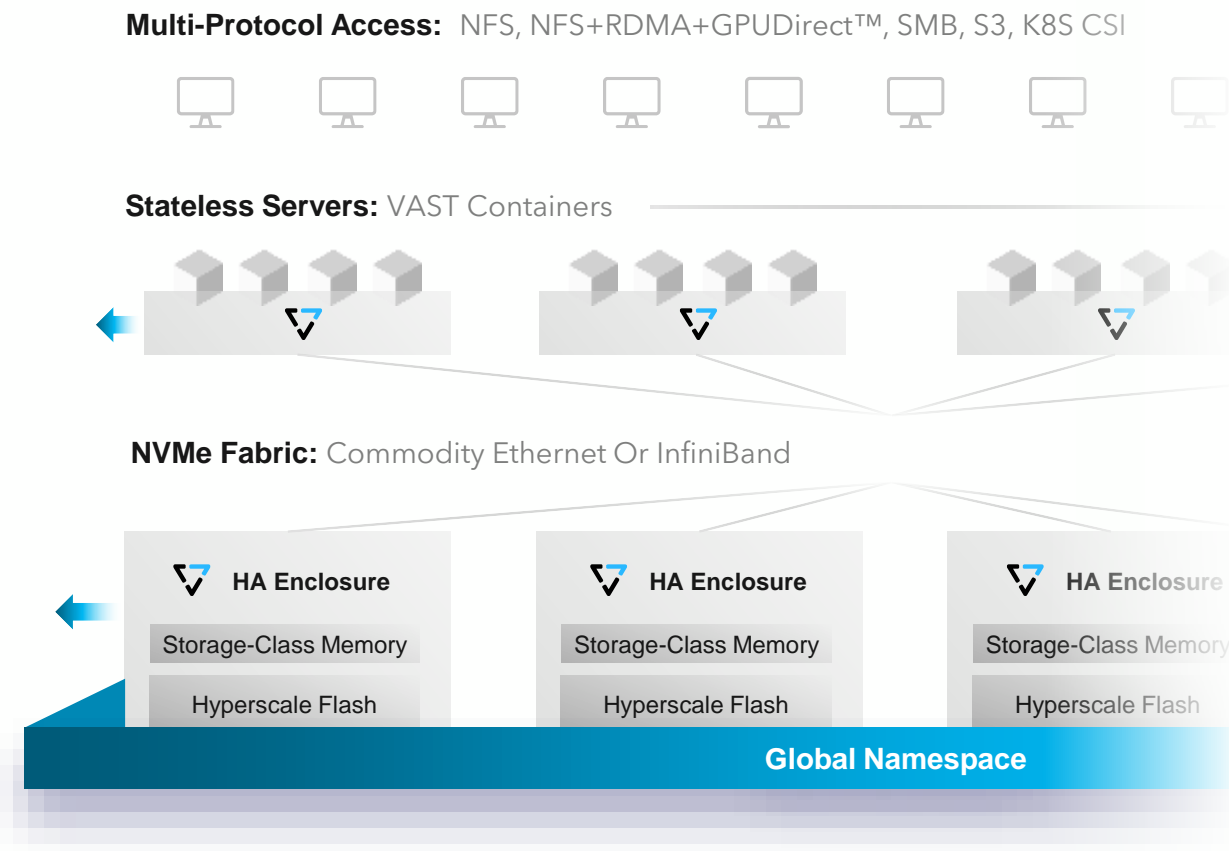
DISAGGREGATED, SHARED-EVERYTHING

# ““ THE STORAGE ARCHITECTURE OF THE FUTURE

- IDC, 2020

- Scale to exabytes, asymmetrically
- A composable flash cloud
- Archive economics
- 99.9999999% Architecture

VAST



STORAGE EFFICIENCY


# TRANSFORMING *THE CALCULUS* OF FLASH OWNERSHIP

## VAST DATA FORESIGHT




QLC + 10-Year Longevity

## LOCALLY DECODABLE ERASURE CODES



2.7% Overhead Data Protection

## SIMILARITY-BASED DATA REDUCTION




Global, Fine-Grained Compression

STORAGE EFFICIENCY


# *SIMILARITY IS GAME-CHANGING*

COMPRESSION




Fine Grained, But Local

DEDUPLICATION



Global, But Coarse

VAST DATA SIMILARITY REDUCTION



Global & Fine Grained

EXAMPLE SAVINGS FROM SIMILARITY

**3:1** Pre-Reduced Backups

**3:1** Pre-Compressed Log Files

**2:1** Life Science Data

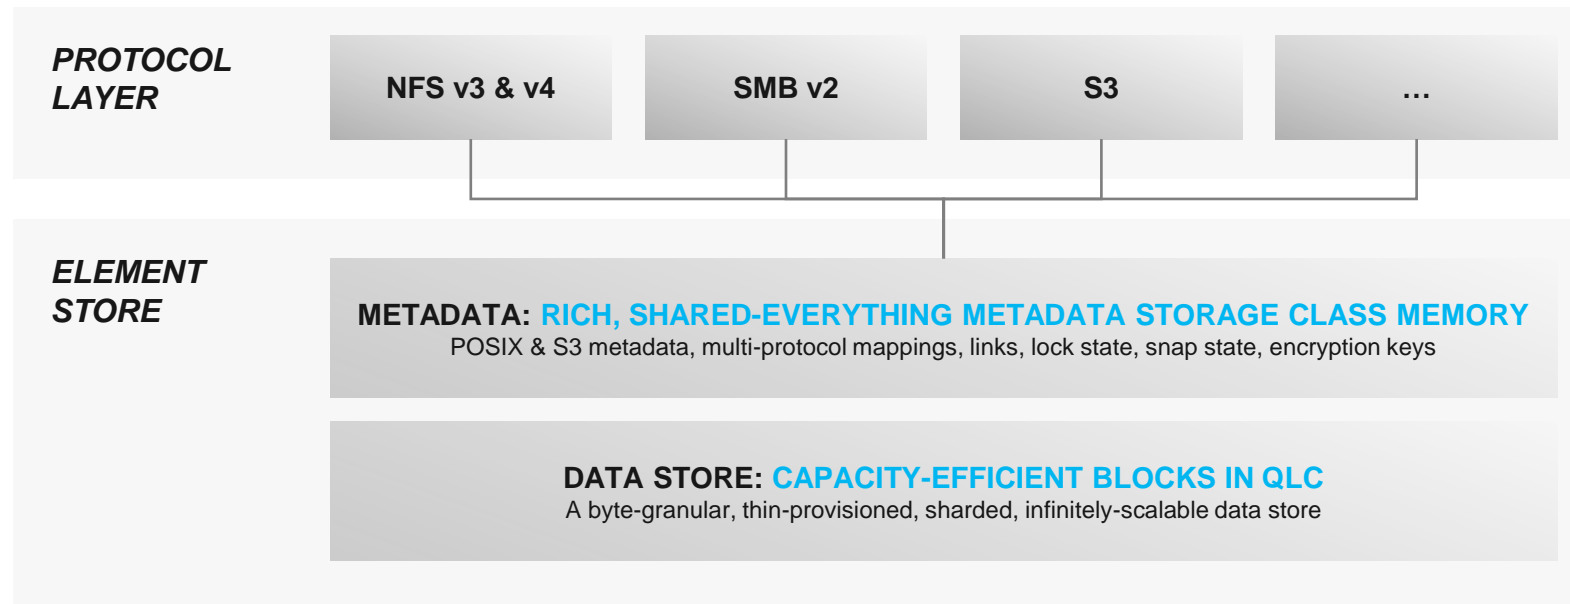
**3:1** HPC Data

**3:1** Animation Data

**8:1** Uncompressed Time-series Data

UNIVERSAL STORAGE

# ***FLEXIBLE & EFFICIENT DATA STRUCTURES***





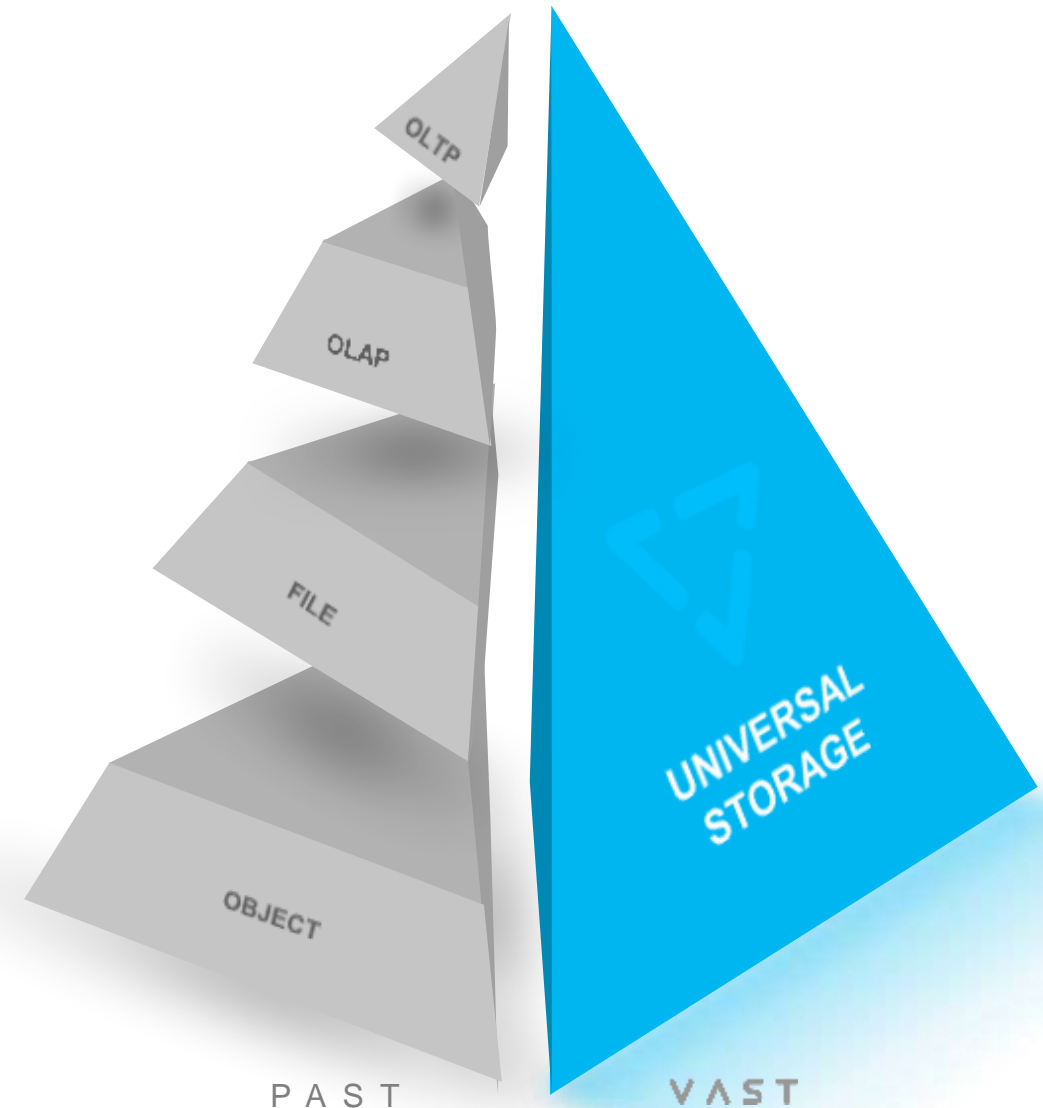
OUR MISSION

# ***NO MORE TIERS***

## **BREAKING TRADEOFFS TO MAKE STORAGE SIMPLE**

- Extinction-level Event for The Hard Drive
- An End To 30 Years of Complex Storage Tiering
- Unleash Big Data and ML Insights

VAST



*QUESTIONS?*

