VAST Data

Rethinking your Data Strategy to improve better Medical Comes

- ✓ Improving Bio-medical Imaging
- ✓ Scaling Alpha Fold and Cyro-EM
- ✓ Speeding up Genome Mapping
- ✓ Powering AI to deliver accurate diagnosis

Howard Fyffe : Managing Director Australia & New Zealand
Email: howard.fyffe@vastdata.com
Mobile: +61435962224

About VAST Data

Fastest Growing Enterprise SW Company In Infrastructure History

Annual Recurring Revenue

$100M+

Elite ”Centaur” Status, 2.5x Y/Y Growth

100%

Recommended

Source: VAST Customer Surveys

Gartner peerinsights

Ranked Among The World’s Fastest-Growing & Innovative Companies

#5

Deloitte.

Fastest Growing Company

Top 50

FAST COMPANY

“Most Innovative”: Data Science

#15

Forbes

Best Startup Employers

$3.7B

Series D Valuation
April 2021

The most valuable private enterprise infrastructure SW company in history

Cashflow Positive For Last 5 Quarters

Over Eight Exabytes of VAST Data

Fastest Growing Enterprise SW Company In Infrastructure History

DATA PLATFORMS FOREVER, HAVE BEEN COMPLICATED.

Re-thinking your Data Strategy

NO MORE TIERS

BREAKING TRADEOFFS TO MAKE ACCESSING AND USING YOUR DATA SIMPLE

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

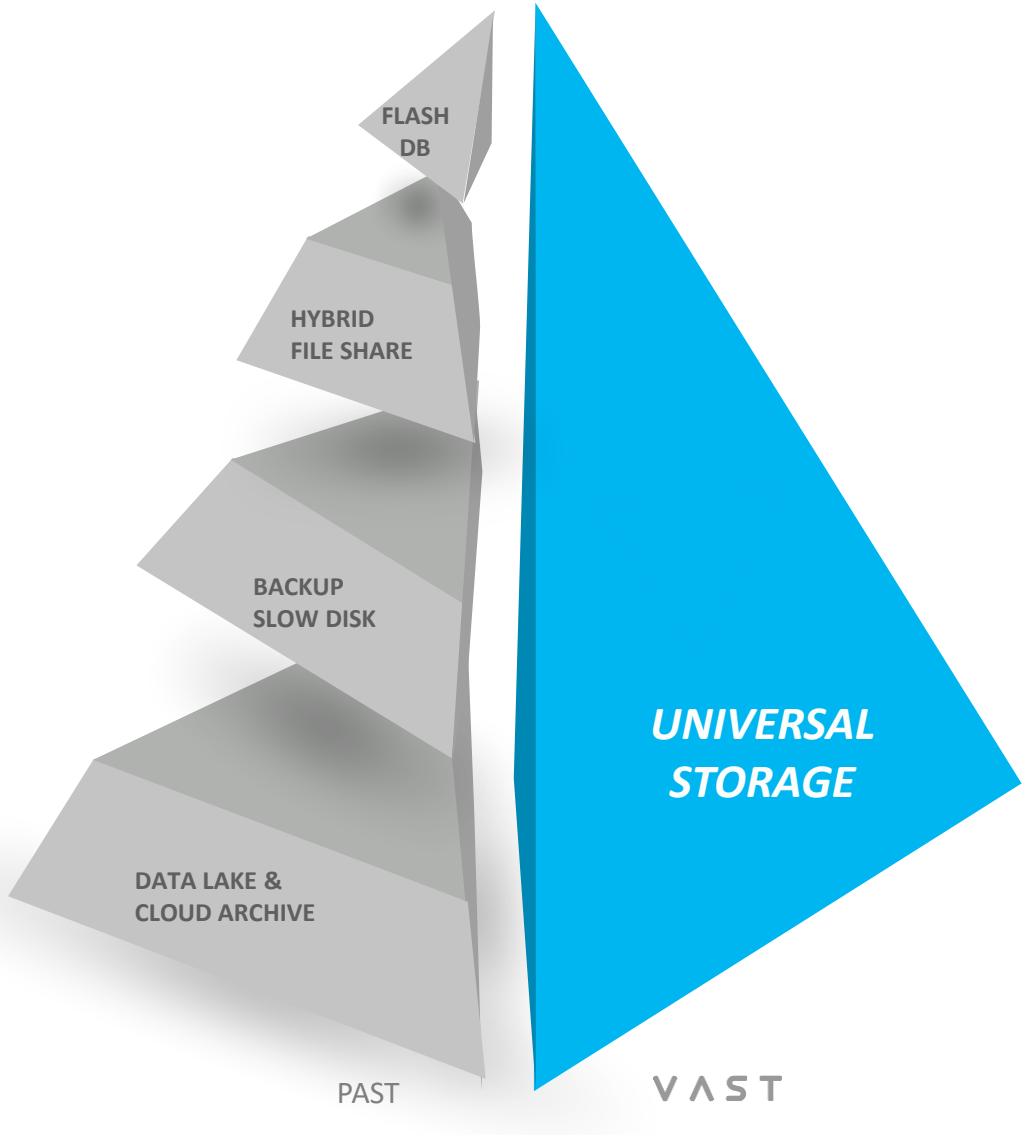

VAST DATA ELIMINATES DATA SILOS

FAST ACCESS TO ALL DATA

DEEP LEARNING & AI TRANSFORMS HOW WE VALUE AND ACCESS DATA

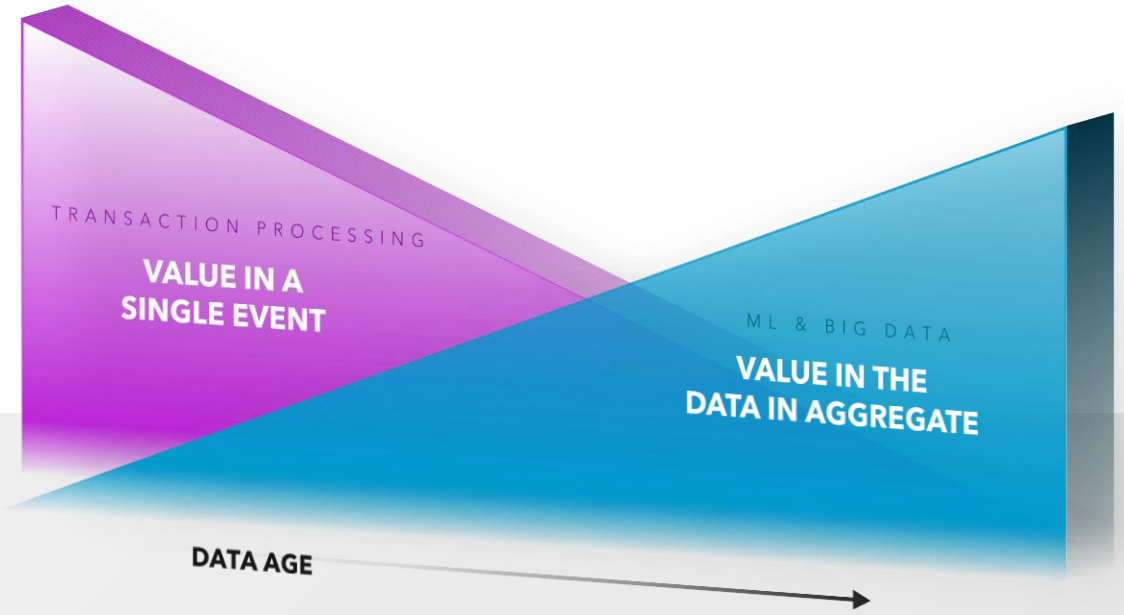

Deep Learning and AI Will create more Economic Value than the Internet

However, legacy data tiering is will not work in this new era

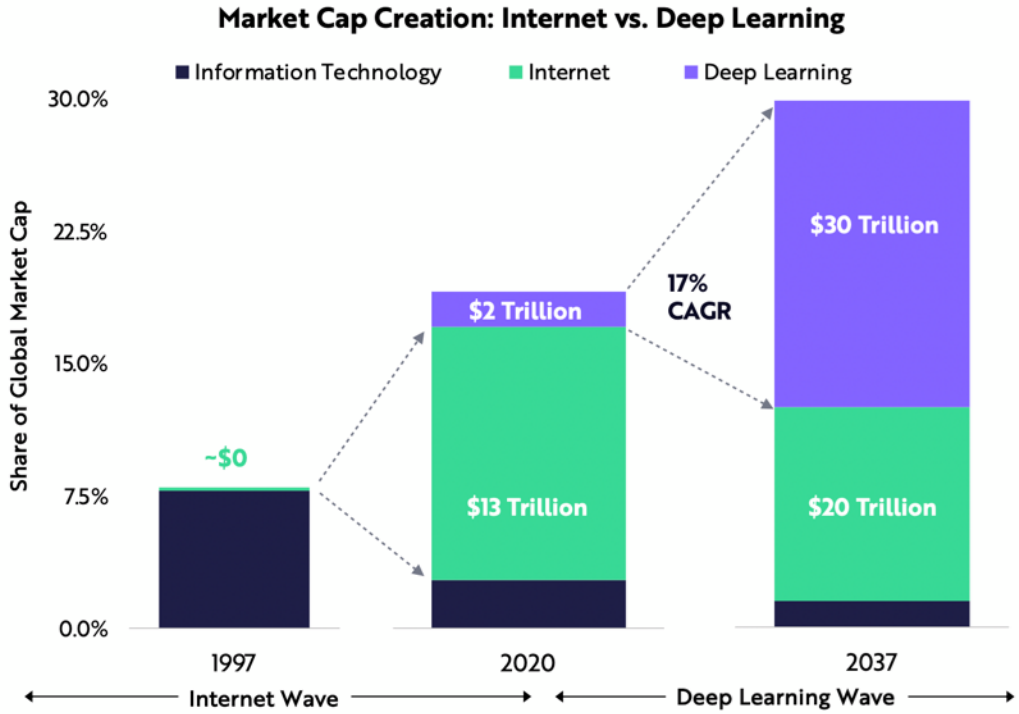

"The accelerator should replace the CPU as the main server compute engine by 2030."

Source: ARK Capital

So Why is VAST Different?

We saw in 2016, 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.

INDUSTRY RECOGNITION

IDC ANALYZE THE FUTURE

STORAGE ARCHITECTURE OF THE FUTURE

Democratizing Flash for Much Broader Workload Consolidation in Enterprise Environments (2020)

Gartner

2022 File & Object Magic Quadrant

The highest placement of any vendor since the file & object Magic Quadrant introduction.

HPCwire 2022 EDITORS' CHOICE AWARDS

Storage Product Of The Year

A breakthrough after 20 years

VMware vSAN

FlashBlade FB

COHESITY

NUTANIX

ISILON SYSTEMS

SAP HANA

ceph

Google

Direct-Attached, Shared Nothing

elastic

hadoop

redis

aws S3

rubrik

VAST DASE

Disaggregated & Shared Everything

A Fundamentally-New Approach to Scale-Out

Disaggregated, Shared-Everything (DASE)

- Scale To Exabytes, Asymmetrically
- Archive Economics
- A Composable Flash Cloud
- 99.9999999% Architecture

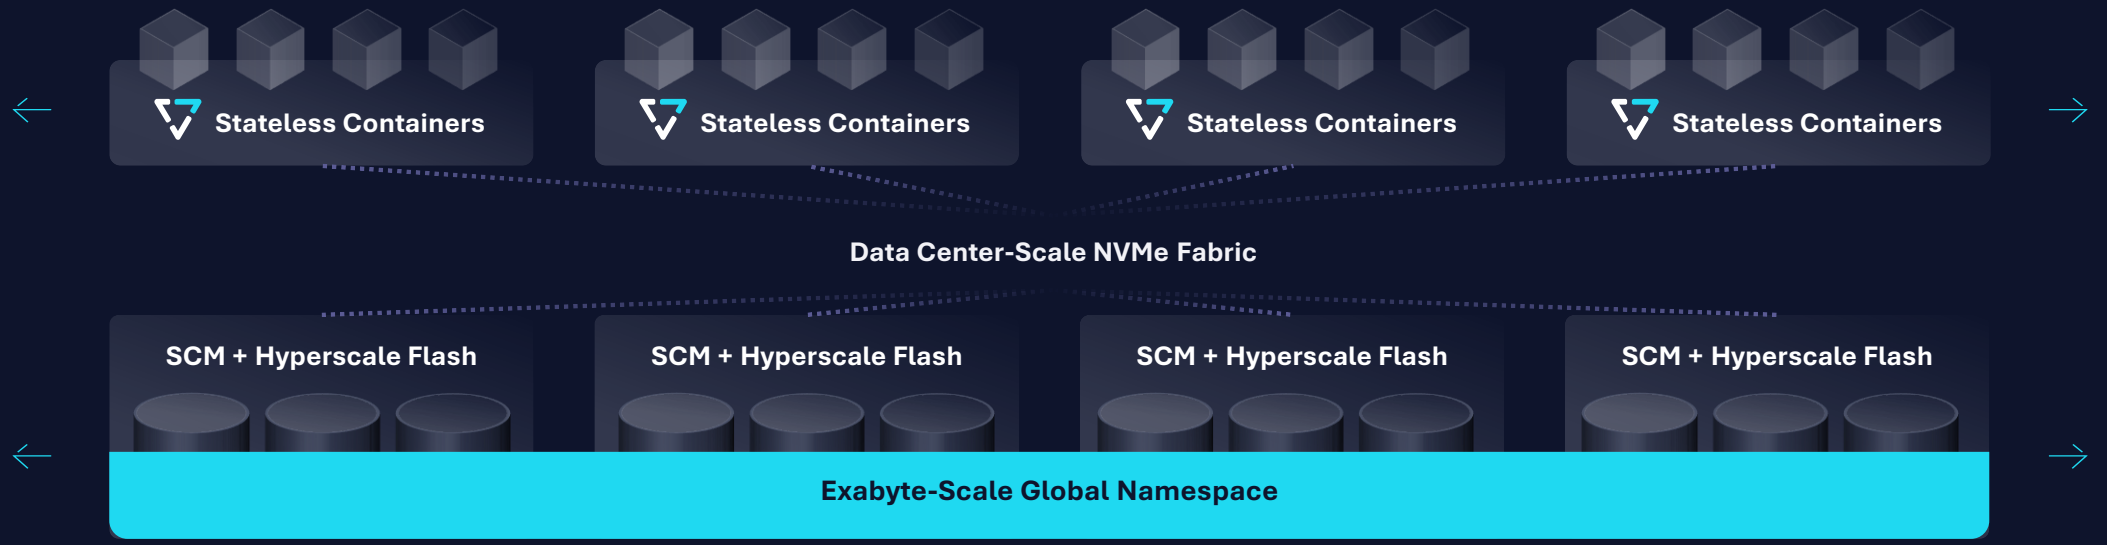

"The Architecture Of The Future" - IDC

Disaggregated

Stateless Nodes Are Composable, Data Always Redundant

Shared-everything

No Need For Cache Coherence, Massively Scalable Writes, Reads, Metadata

Radically Efficient

New Global Efficiency Codes To Revolutionize NVMe Economics

The VAST Data Platform

All of Your Data, from Edge-to-Cloud. Enriched, Fused & AI-Ready.

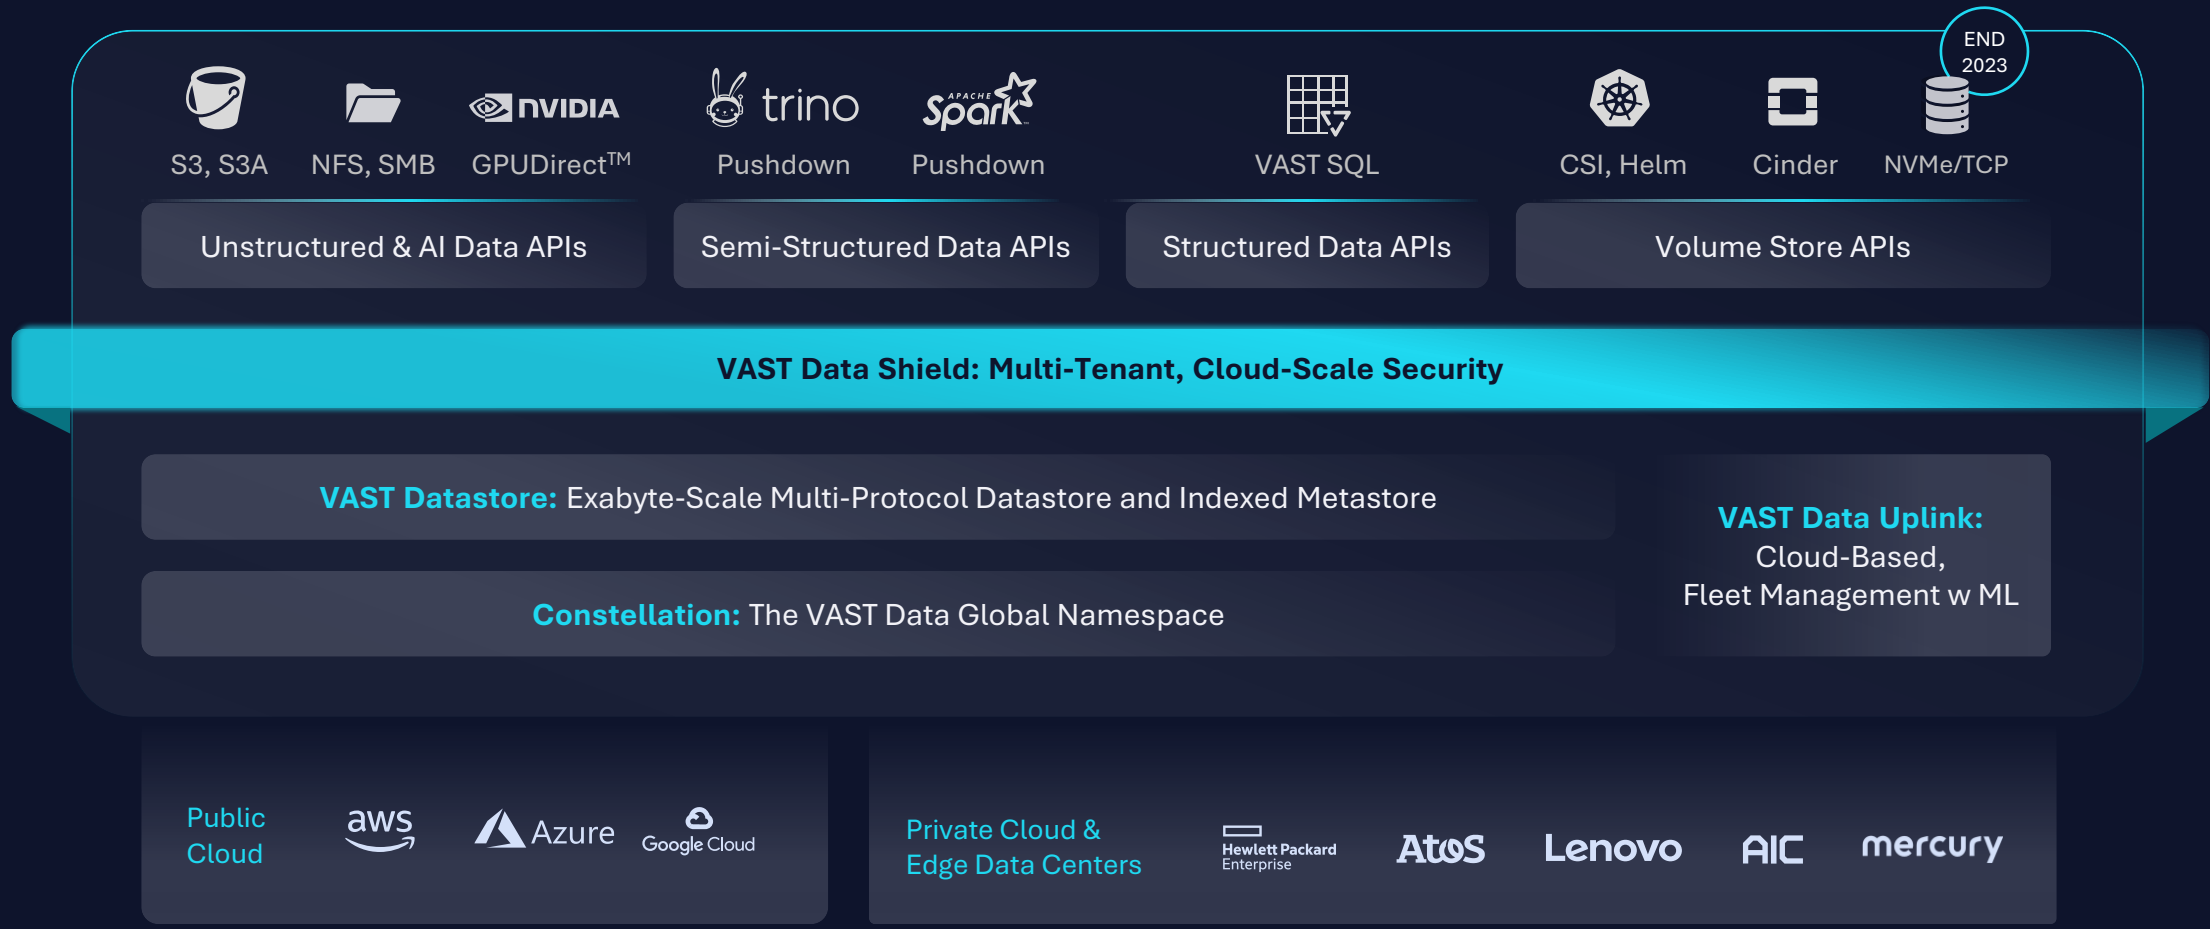

HOW IT WORKS

VAST PROVIDES AN INFINITE STORAGE LIFECYCLE.

Never Migrate Again.

Asymmetric architecture

Forward & backward SW compatibility

Decoupled HW & SW.

Never pay SW tax on a HW refresh

Always leverage HW investments for their useful life

Refresh, at cost, Replace at your convenience.

THE INFINITE
REFRESH CYCLE.

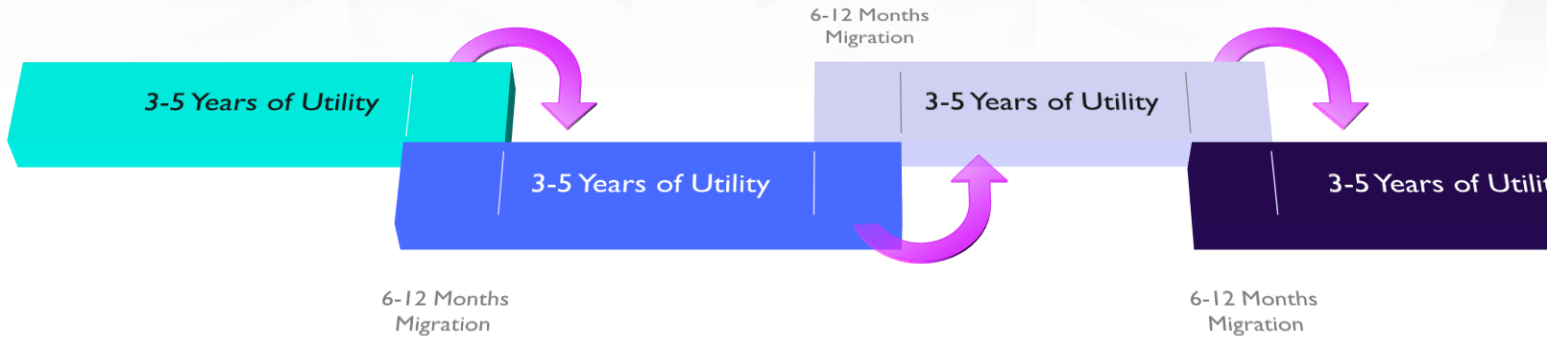

THE INFINITE
STORAGE LIFECYCLE.

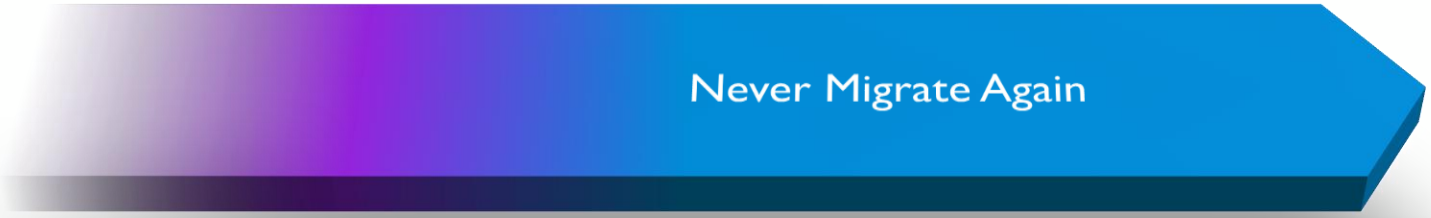

VAST Data Lake Moderization

AGODA LEVERAGES VAST'S S3 COMPATIBLE FAST OBJECT STORAGE TO BUILD COST-EFFECTIVE COMPUTING ENVIRONMENT FOR APACHE SPARK AND APACHE IMPALA PROCESSING

agoda

THE PROBLEM

- Needed to move away from Cloudera HDFS and proprietary toolsets (HIVE and Impala)
- Needed an on-prem performant and scalable S3 platform for PRESTO DB.

THE SOLUTION

- Commercial flexibility that allows scale beyond 15PB of effective capacity, at half the cost of alternative tiering solution
- Multi-protocol support (NFS and S3) with massive scale

YESTERDAY'S BIG DATA INFRASTRUCTURE

- Complex architecture driven by legacy HDFS
- Inability to scale storage & performance independently
- High overhead: 3X replication
- Costly rebuild on node failure

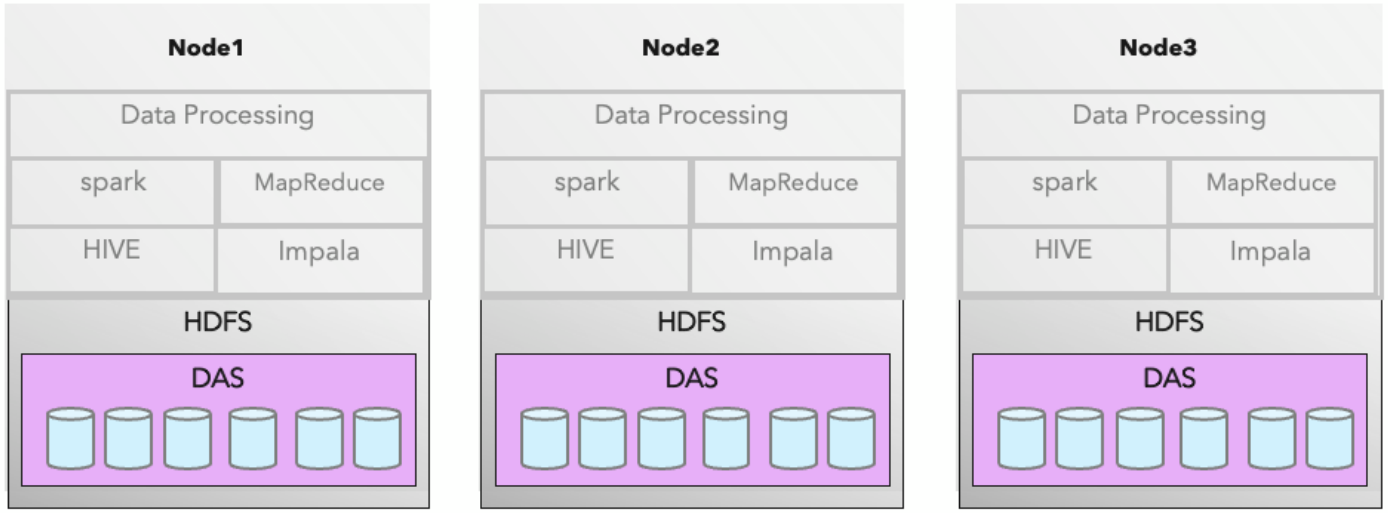

Modernise your Big Data Platform

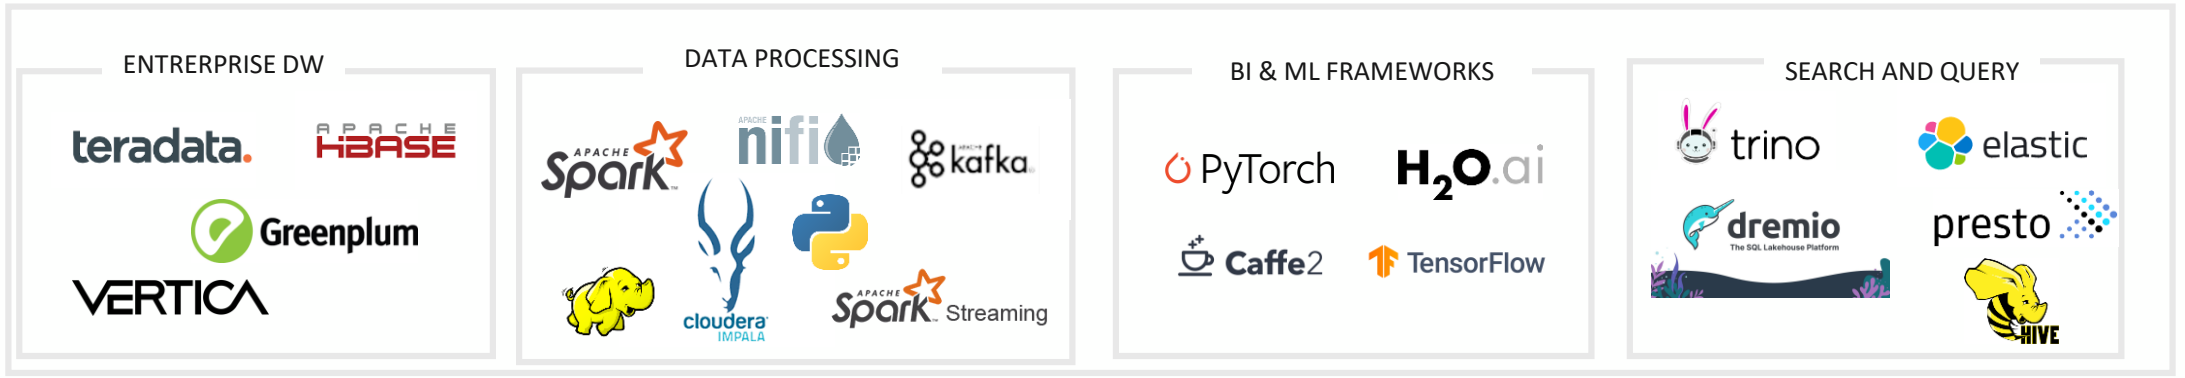

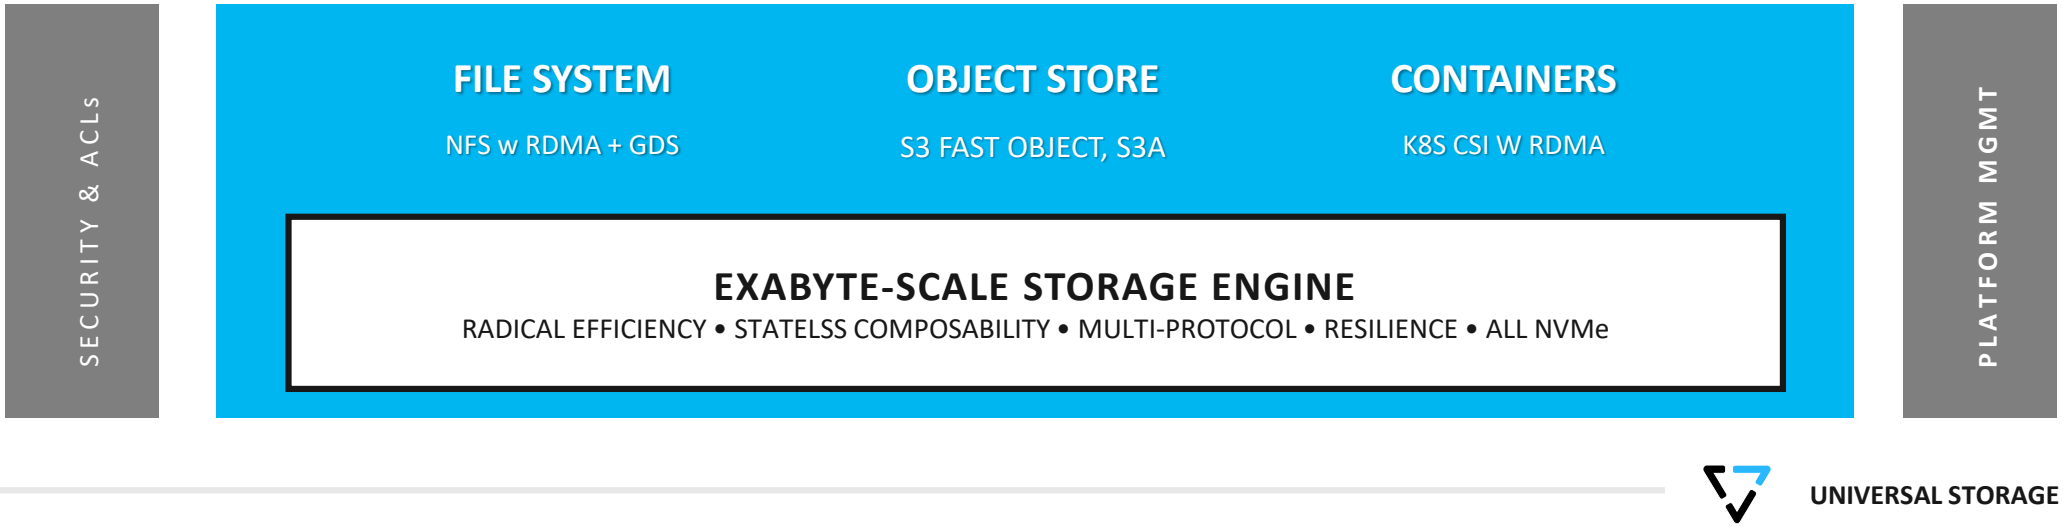

VAST DATA vs HDFS

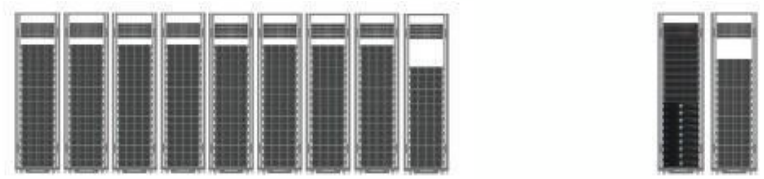

	HDFS	VAST
TOTAL RACK UNITS (COMPUTE+STORAGE)	336	80
RAW CAPACITY	30PB	6PB
AVAILABLE CAPACITY	10PB	10.7PB*
STORAGE THROUGHPUT (READ BANDWIDTH)	180GB/s	360GB/s
RACKS	9	1
CLOUDERA SERVER NODES	150	36

THE VAST ADVANTAGE

80%
LESS CAPACITY NEEDED

2X
FASTER PERFORMANCE

88%
LESS RACK SPACE

80%
LESS SW LICENSE COST

VAST Life Sciences

INVITAE

“Our COVID work took 24 hours on legacy scale-out NAS. It now finishes in 15 minutes.

Success Story

- Went from processing a genome during COVID from 24 hrs to 15 minutes
- 50% more performance than other flash solutions, at 70% less cost
- GPUs + VAST enable the Invitae pipeline
- Now consolidating all workloads onto VAST

VAST

Athinoula A. Martinos Center For Biomedical Imaging

“We’ve never gotten this level of support and near-instantaneous feedback from any vendor before.

PAST

- Struggled to affordably capture and store enormous volumes of raw image data while enabling fast data access to hundreds of researchers’ AI initiatives

VAST

- All-flash infrastructure for NVIDIA A100s – less than hybrid alternatives
- Faster image analysis using affordable all-flash storage capable of scaling to support massive data growth
- Uptime challenges inherent in legacy HDD systems have been eliminated

VAST

THERAPIXEL

“VAST is a game-changer, providing the simplicity, scalability, and capacity needed to run our training algorithms.

PAST

- Legacy HDD-based storage could not provide random read performance required to accelerate the detection of breast cancer mutations at scale

VAST

- Eliminates capacity and performance bottlenecks
- Highly reliable breast cancer detection at first reading
- Faster access to train and iterate algorithms to interpret mammograms

VAST

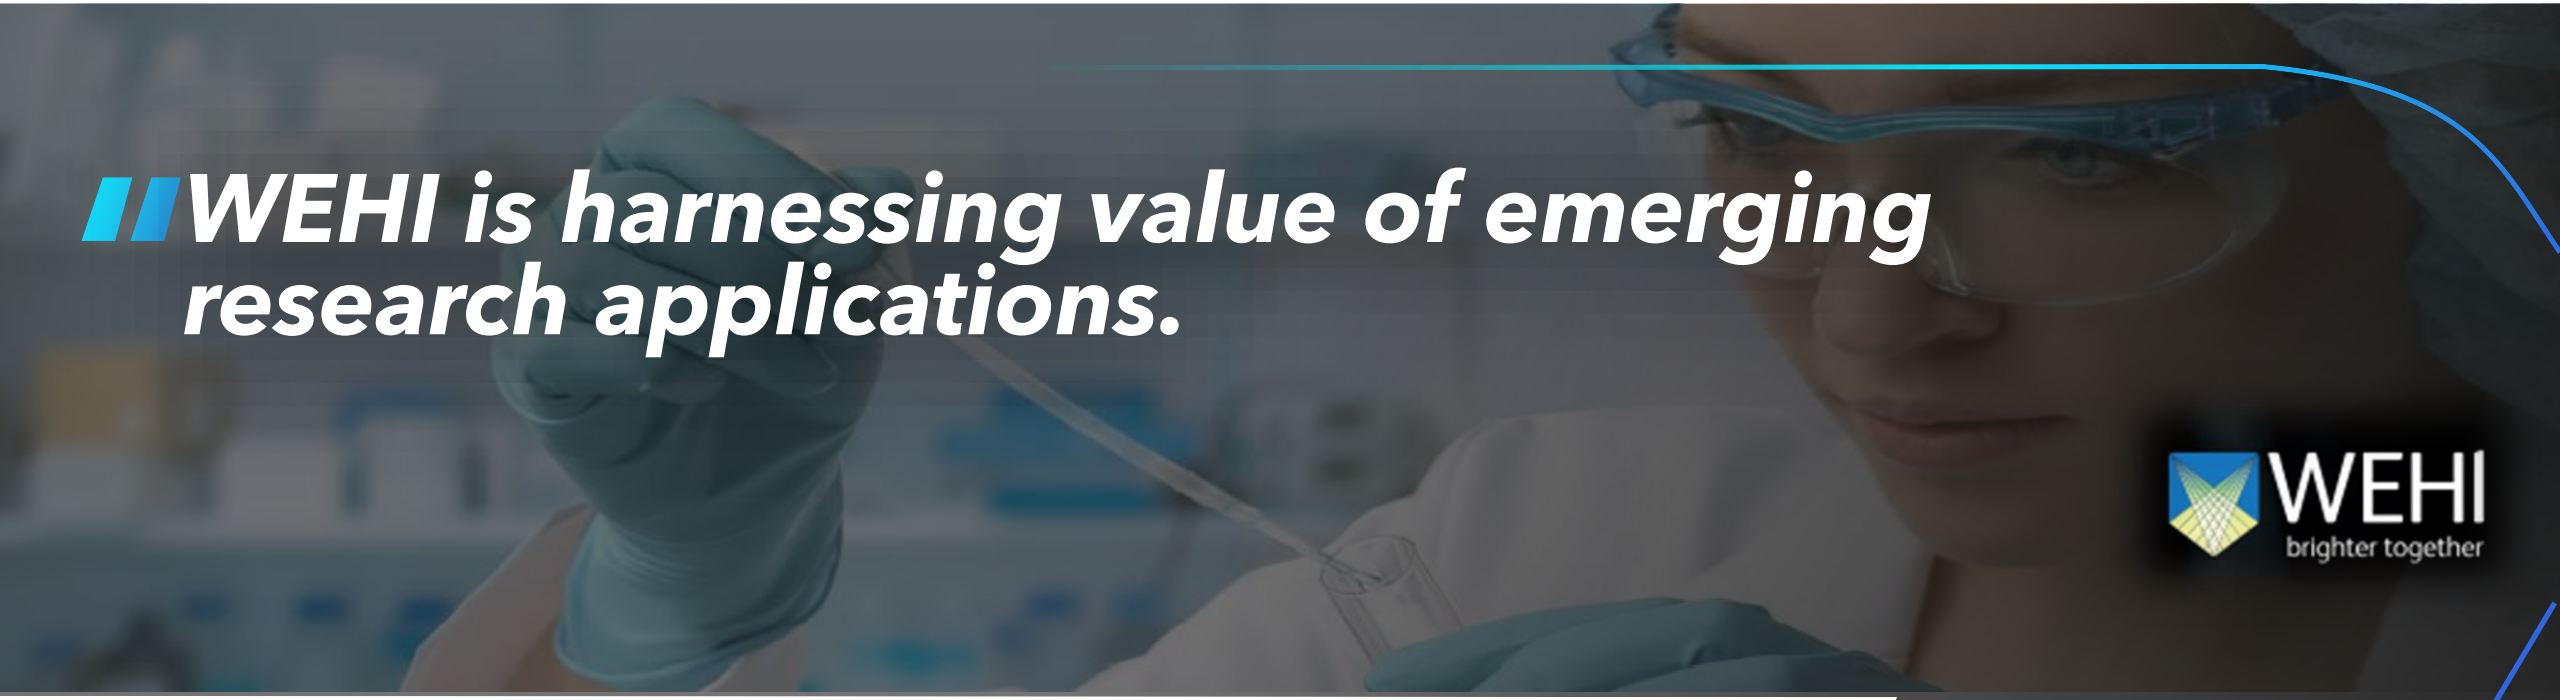

WEHI is harnessing value of emerging research applications.

WEHI
brighter together

THE PROBLEM

- Next-generation techniques, such as Cryo-EM, are particularly demanding on storage
- Random IO often the bottleneck with inconsistent performance on spinning disk

THE SOLUTION

- VAST has enabled us to tackle the challenge of rapid growth in dataset sizes
- Disk IO no longer the bottleneck – even when our batch system is running near capacity
- No downtime required for maintenance, solid API and comprehensive dashboards
- VAST just works! Less time troubleshooting storage, more time for supporting research

VAST

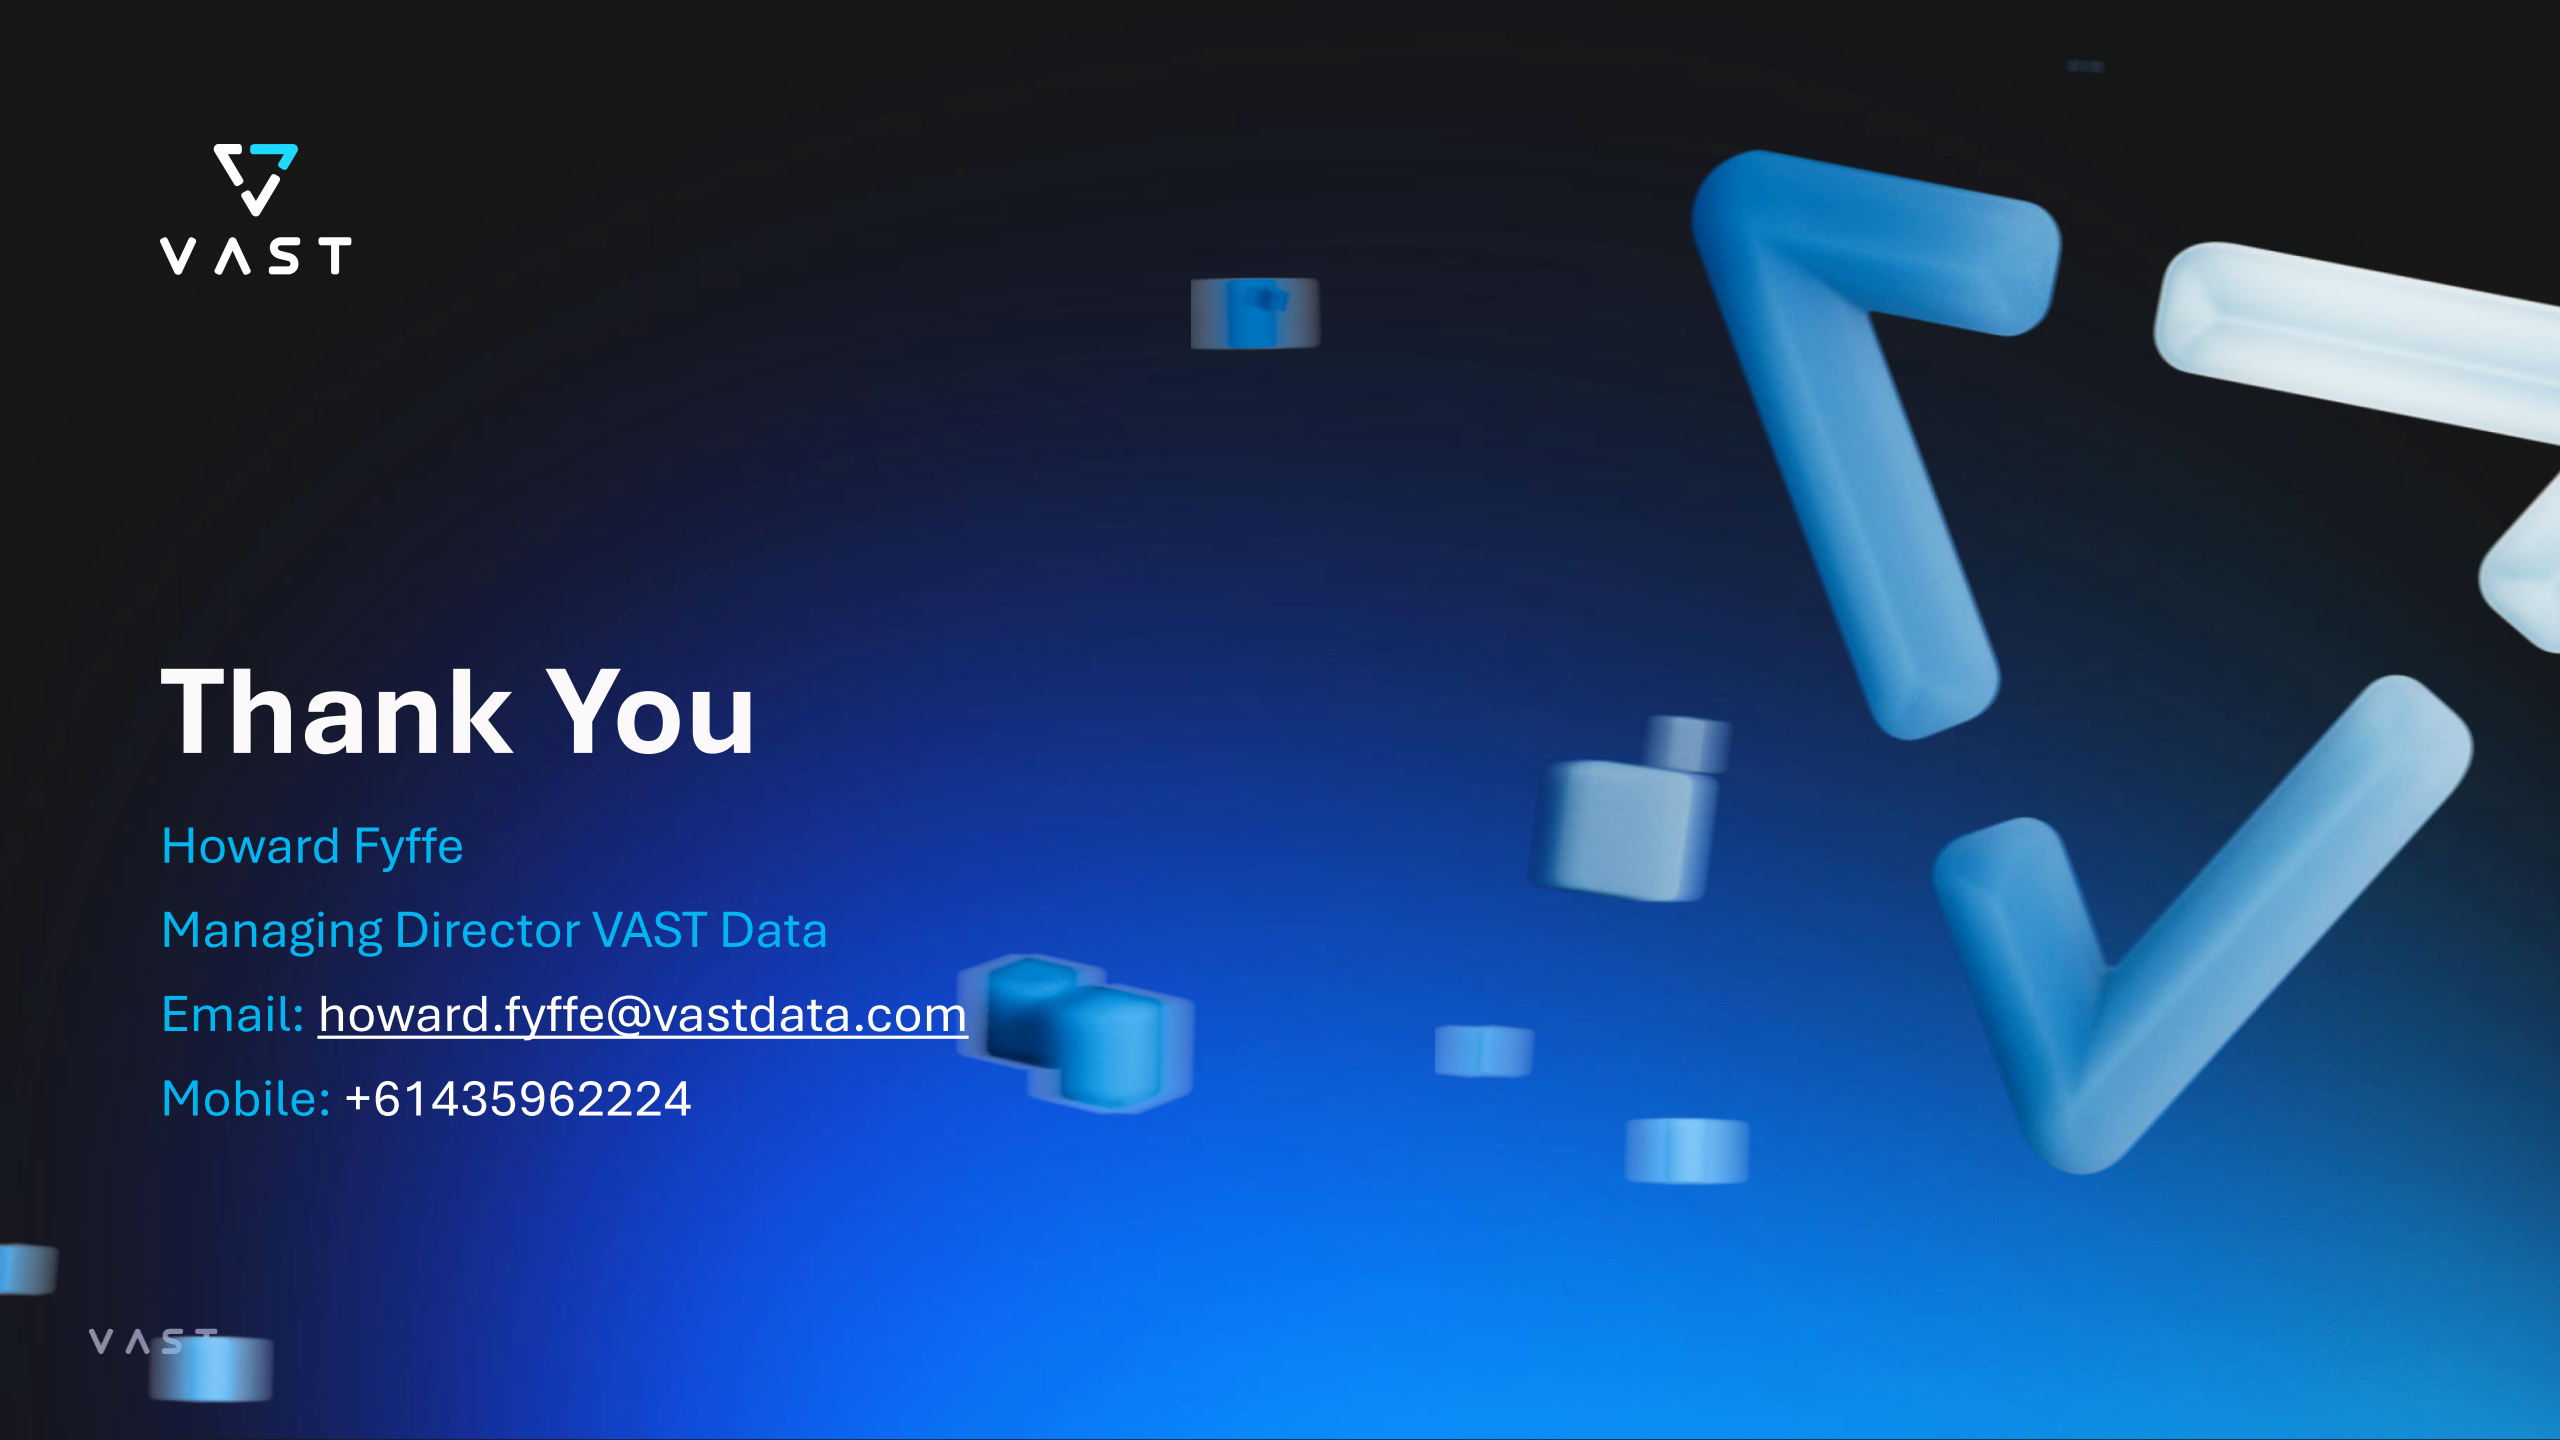

Thank You

Howard Fyffe

Managing Director VAST Data

Email: howard.fyffe@vastdata.com

Mobile: +61435962224

Alternative use cases

VAST Ransomware Protection

"The University of Auckland modernised data protection with immutable backups and all-flash restore in the event of a cyber attack.

UNIVERSITY OF AUCKLAND

THE PROBLEM

- Ransomware threat meant large data restoral from Hard Drive Disk (HDD) storage could take months

- Legacy HDD storage did not offer immutability to protect from a ransomware attack

THE SOLUTION

- All flash data platform delivered as a per TB consumption model offer Cloud Commercial flexibility on premise

- All-flash performance and data immutability critical for the university.