INTRODUCING:
UNIVERSAL STORAGE

DEMOCRATIZING FAST ACCESS TO ALL DATA

Petr Plodík, Sales Director, petr.plodik@mcomputers.cz

VAST

Agenda

představení společnosti VAST data
VAST Universal Storage
technologie redukce dat
řešení dostupnosti
správa VAST úložiště
příklady nasazení

zkušenosti z České republiky

kvíz o ceny

VAST DATA LEADERSHIP SINCE 2016

A PROVEN TEAM

RENEN HALLAK

CEO & FOUNDER

XtremIO

MICHAEL WING

PRESIDENT

DELL EMC

JEFF DENWORTH

CMO & CO-FOUNDER

ddn

SHACHAR FIENBLIT

VP R&D & CO-FOUNDER

kaminario

AVERY PHAM

VP OPERATIONS

PURESTORAGE

INVESTORS

83NORTH · commonfund · DELL Technologies CAPITAL · Goldman Sachs · GREENFIELD · Next47 · NORWEST · NVIDIA

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

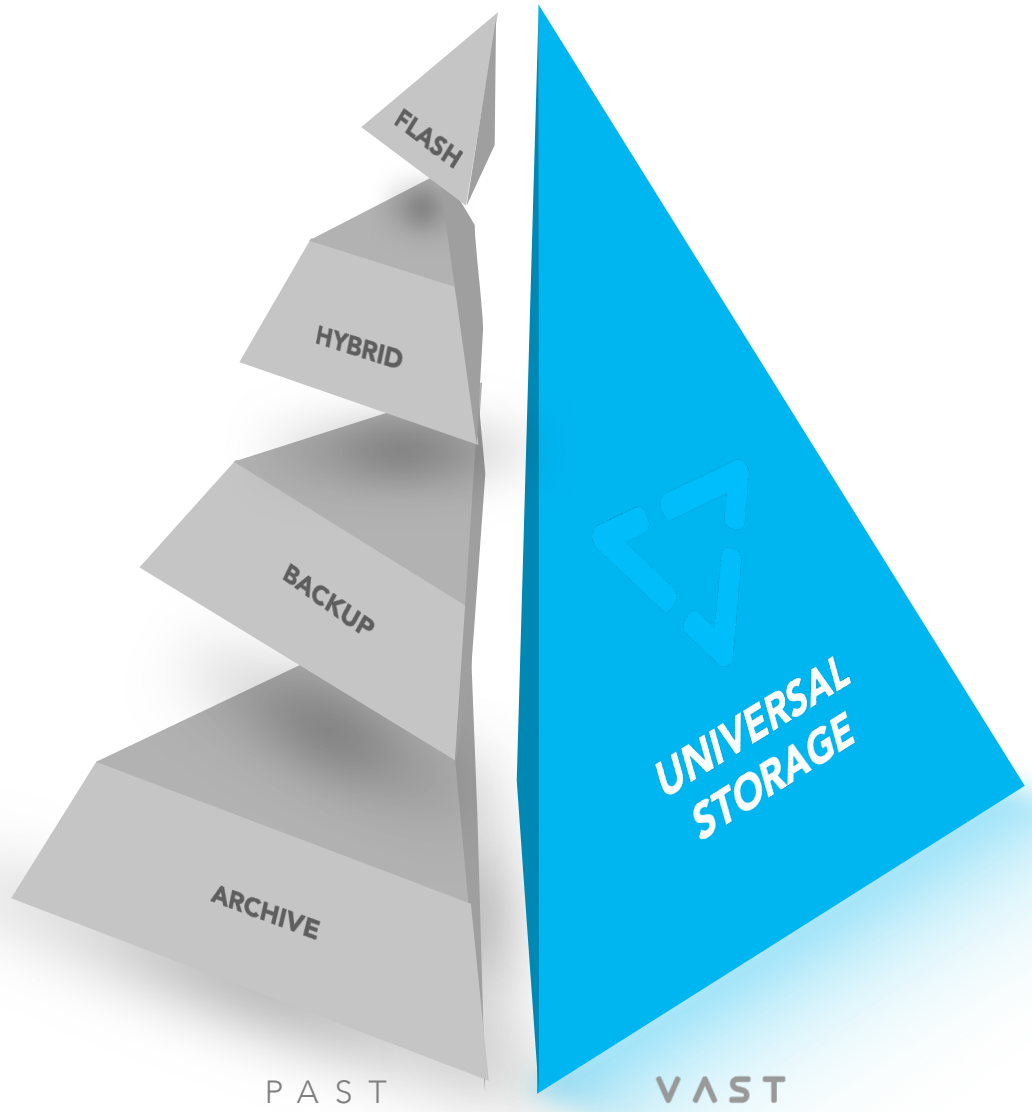

UNIVERSAL STORAGE

ONE *UNIVERSAL* DATA PLATFORM FOR ALL YOUR DATA

ALL-NVMe PERFORMANCE

Eliminate application bottlenecks

EXABYTE SCALE, WITH ARCHIVE ECONOMICS

Eliminate the HDD with a single-tier flash cloud

ENTERPRISE NAS & OBJECT STORAGE

Standard interfaces with game-changing performance

DELIVERED VIA OUR REVOLUTIONARY GEMINI PROGRAM

The world's first disaggregated commercial model

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.

3D XPOINT / SCM FOR LONGEVITY & EFFICIENCY

Enables write shaping to QLC and rich metadata.

USE CASES

PURPOSE BUILT FOR ALL DATA

VERTICAL SOLUTIONS	SECONDARY STORAGE	BIG DATA & AI	ENTERPRISE INFRASTRUCTURE
Life Science	COMMVAULT	splunk>	kubernetes
Finance	veeam	Apache Spark	OPENSHIFT
Content	rubrik	TensorFlow	vmware
HPC			

VAST Universal Storage: One Platform For All Data Center Data

SHARED-NOTHING VS. VAST'S DASE

LEGACY, SHARED-NOTHING CLUSTER ARCHITECTURE ARE LIMITED

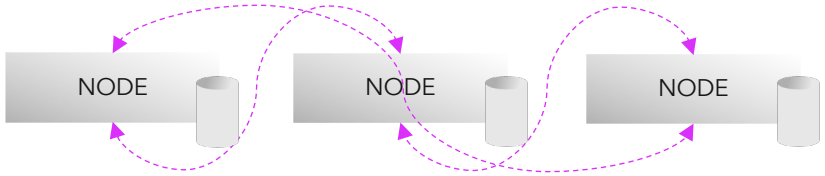

crosstalk, rebuilds and interdependencies increase geometrically with cluster size

DISAGGREGATED, SHARED-EVERYTHING CLUSTER ARCHITECTURE

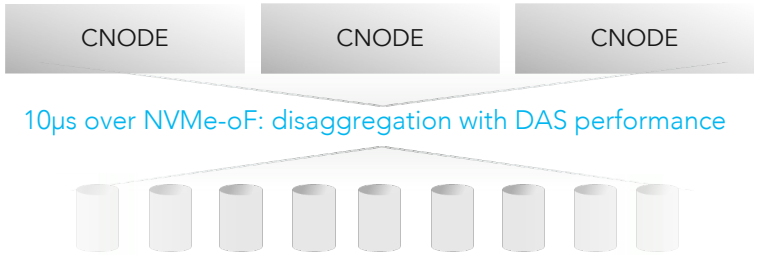

SHARED-NOTHING		VAST DASE
–	Linear performance scaling	✓
–	Independent scaling: CPUs & drives	✓
–	Composable infrastructure	✓
–	Asymmetric scaling	✓
–	Global efficiency codes	✓
–	No rebuilds on server failure	✓

DISAGGREGATED, SHARED-EVERYTHING

REVOLUTIONARY SCALE

- *SCALE TO EXABYTES, ASYMMETRICALLY*
- *LINEARLY SCALE BY ADDING CPUS*
- *SCALE COMPUTE INDEPENDENT OF CAPACITY*
- *ASYMETRIC SCALING ENDS THE INFINITE REFRESH CYCLE*

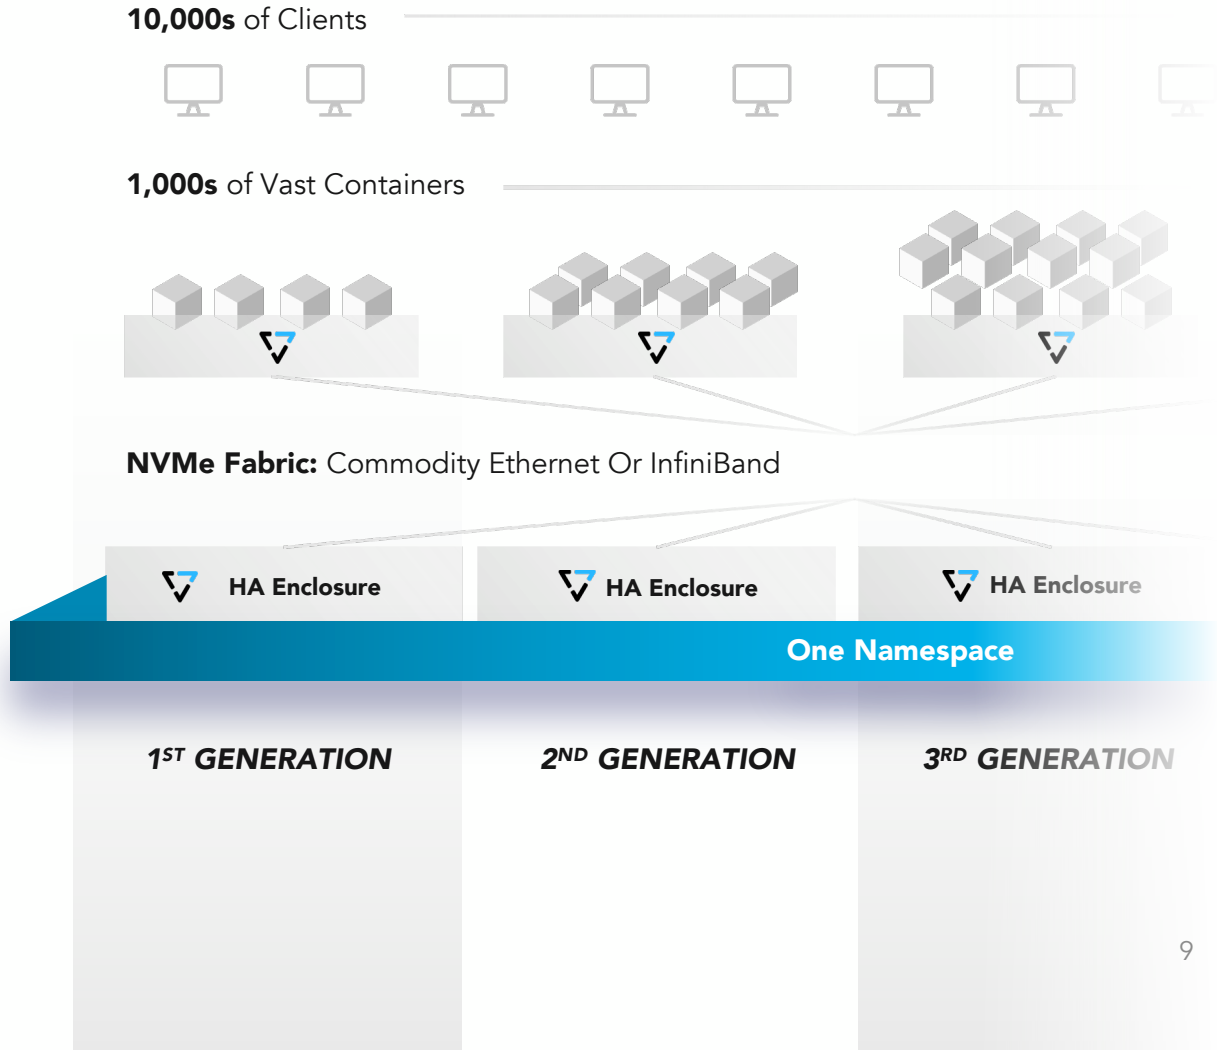

DISAGGREGATED, SHARED-EVERYTHING

A COMPOSABLE FLASH CLOUD

- ALLOCATE CPUS PER TENANT FOR QOS
- LINEARLY SCALE BY ADDING CPUS
- CONNECT SERVERS TO MULTIPLE NETWORKS
- ELIMINATES ISLANDS OF INFRASTRUCTURE

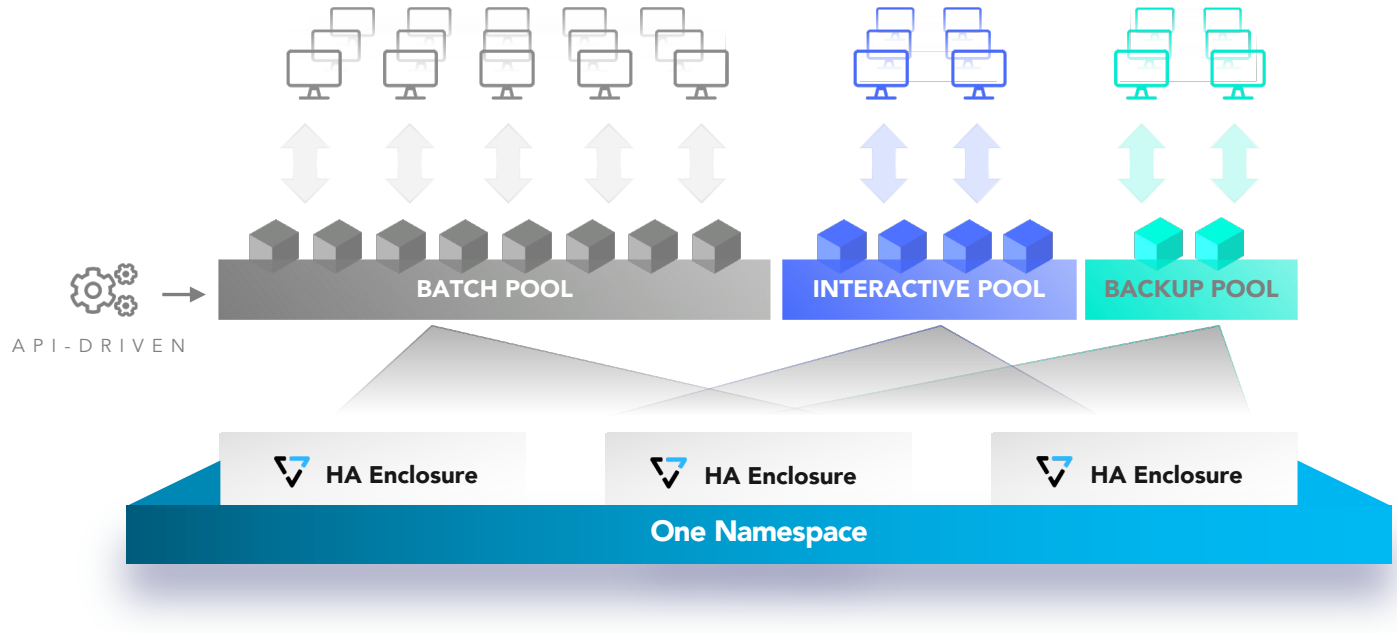

UNIVERSAL STORAGE

FLEXIBLE & EFFICIENT DATA STRUCTURES

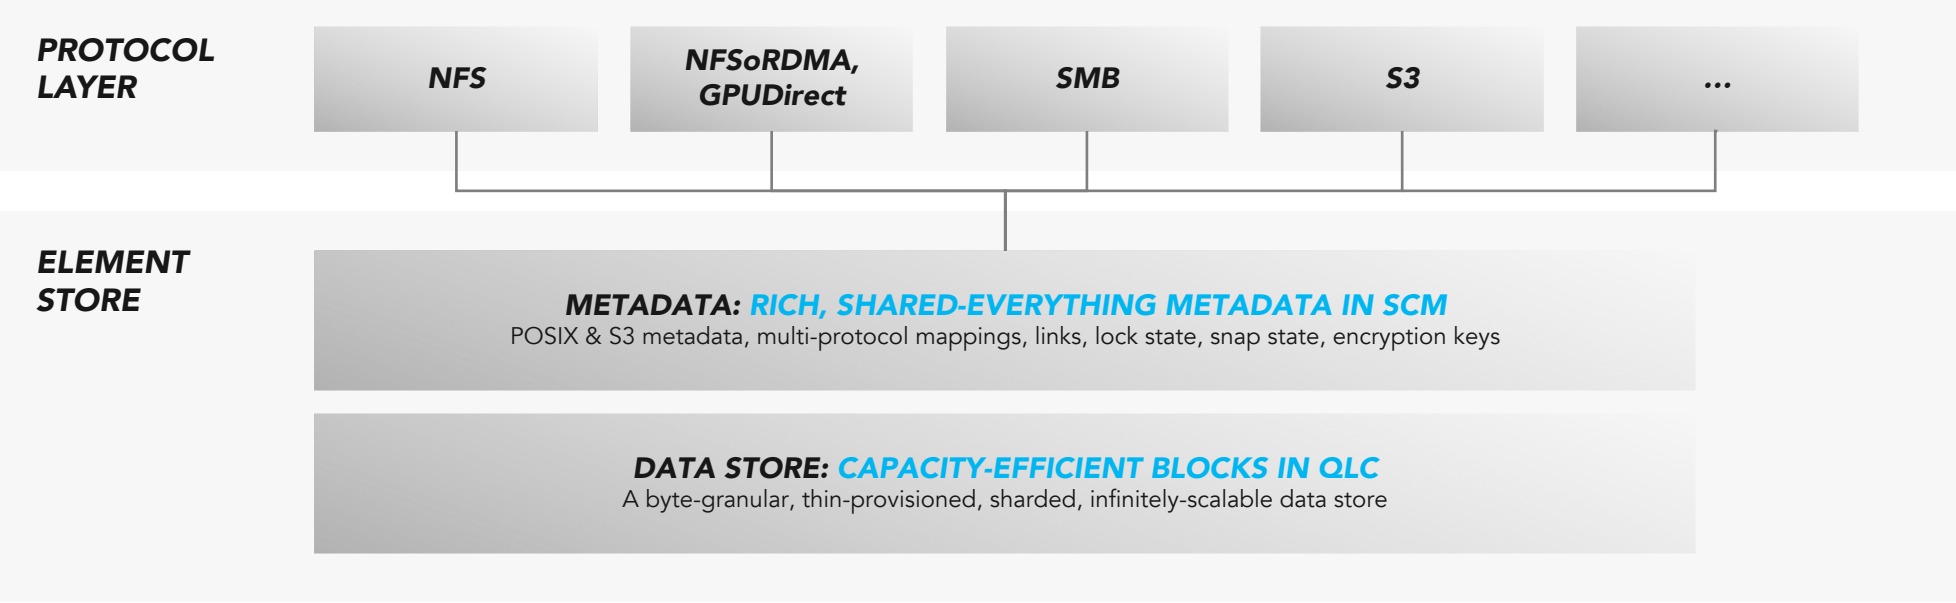

FEATURE SET

NFS V3.0, 4.1
w/ RDMA, GPUDirect, byte-range locks

SMB 2.1 (RESILIENT SMB)
Stateful file services; stateless servers.

S3-COMPATIBLE API, W HTTPS
Stateful file services from containers.

MULTI-PROTOCOL NAMESPACE
Easily traverse between file & object

LOAD BALANCING
Internal DNS services

ENCRYPTION-AT-REST
FIPS-grade security

NO-OVERHEAD SNAPSHOTS
Space and performance efficient

REPLICATION TO CLOUD/S3
Snapshot to an S3 endpoint of choice

NATIVE REPLICATION
Industry-low levels of RPO (15s)

ENTERPRISE AUTHENTICATION
Support for LDAP, AD, NIS

SIMILARITY-BASED DATA REDUCTION
Unprecedented storage efficiency

LOCALLY-DECODABLE DATA PROTECTION
Unrivaled resilience, only 2.5% overhead

AUTOMATION PLUGINS
Kubernetes CSI, Manilla Driver

DIRECTORY QUOTAS

RESILIENT SMB2 FOR WINDOWS & MAC

NEVER LOSE A CLIENT CONNECTION DUE TO NODE FAILURE OR UPGRADE

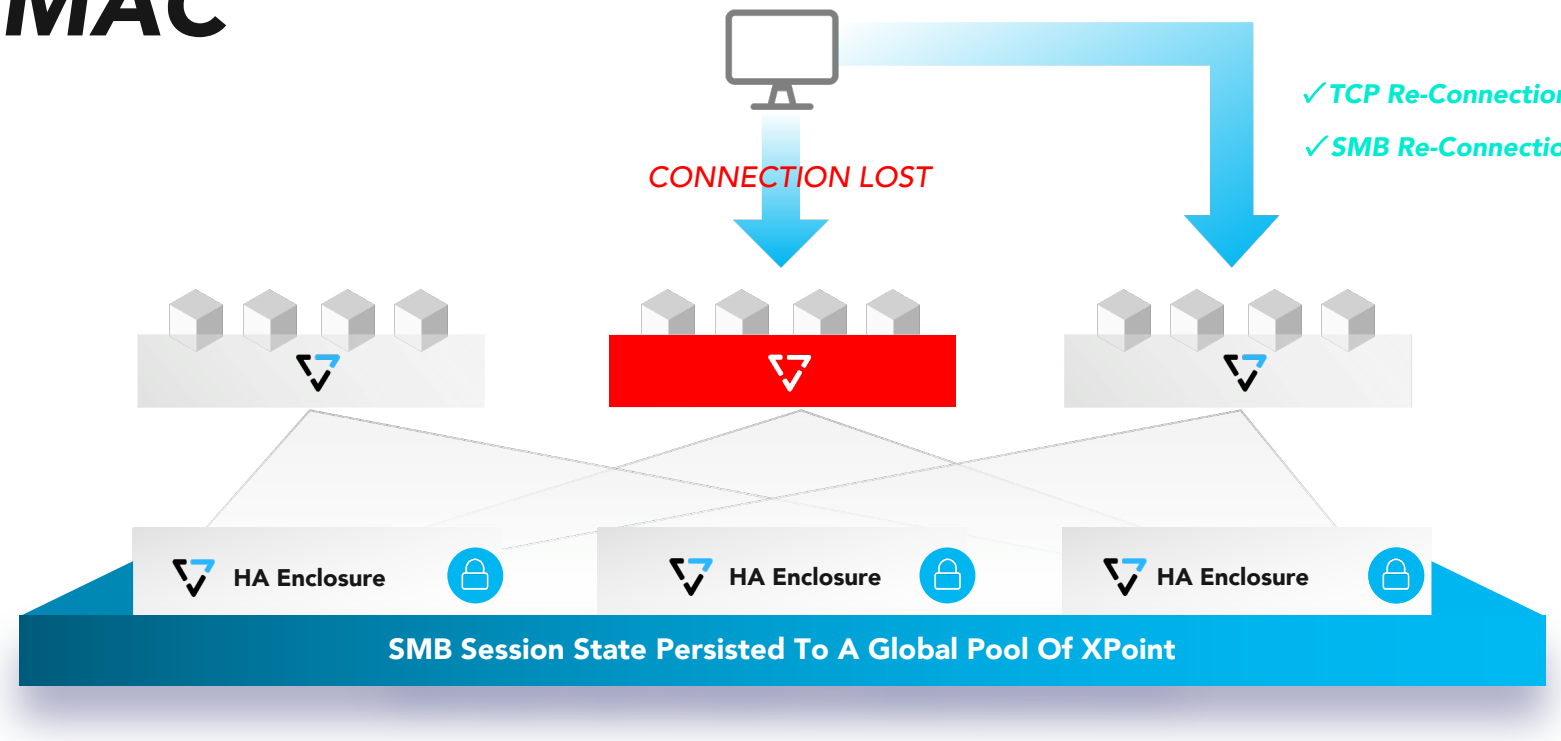

STORAGE EFFICIENCY

TRANSFORMING THE CALCULUS OF FLASH OWNERSHIP

VAST DATA FORESIGHT

QLC + 10-Year Longevity

LOCALLY DECODABLE ERASURE CODES

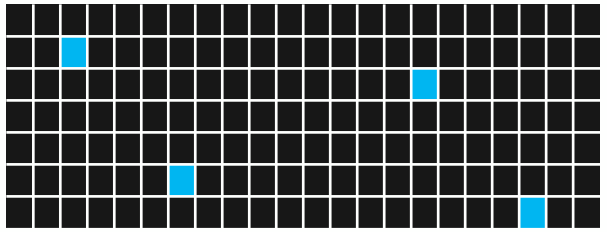

2.5% Overhead Data Protection

SIMILARITY-BASED DATA REDUCTION

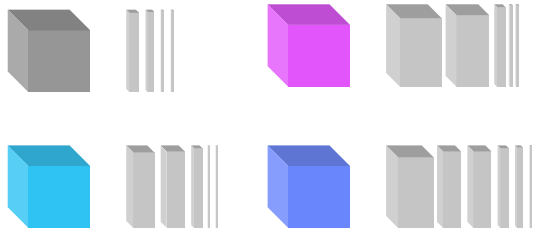

Global, Fine-Grained Compression

QLC flash economics

- **Indirect-on-Write System Always Writes in Full QLC Erase Blocks**
- **3D XPoint Buffering Ensures Full-Stripe Writes**
- **Universal Storage Wear Leveling Amortizes Write Endurance**
- **VAST Predictive Data Placement**

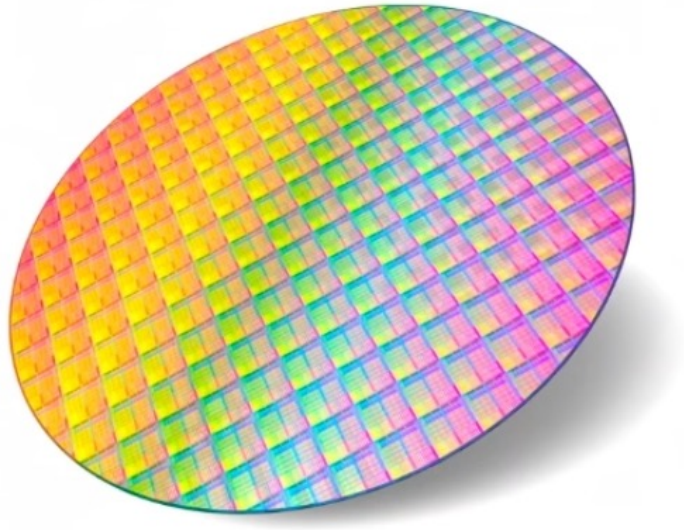

PROS

- 80% Less than enterprise flash ($/GB)
- Reads at the same speed as enterprise flash
- No failure-prone mechanical media

CONS

- <500 Write Cycles
- Not Dual-Ported
- No Cache Power Protection
- Poor Write Speed

Data protection

Historic Barriers to Wide Striping

- Shared Nothing Designs
- Small, Volatile Write Caches
- Reed-Solomon

VAST's Next-Generation Rapid Rebuild Encoding

VAST codes accelerate rebuild speed by using a new type of algorithm that gets faster with more redundancy data. Everything is fail-in-place.

- 150+4: 3x faster than HDD erasure rebuilds, 2.7% overhead
- 500+10: 2x faster than HDD erasure rebuilds, 2% overhead

Additional redundancy enables MTBF of over 100,000 years at scale.

LOCALLY-DECODABLE ERASURE: HOW IT WORKS

vastdata.com/whitepaper#similarity

- SCM/XPoint buffer is a persistent space to build large stripes
- DASE enables large stripes: 36+4 to 146+4 at scale
- +4 codes deliver 60M years of MTTDL, with only 2.7% overhead

Locally-decodable codes reduce rebuild times by 75%

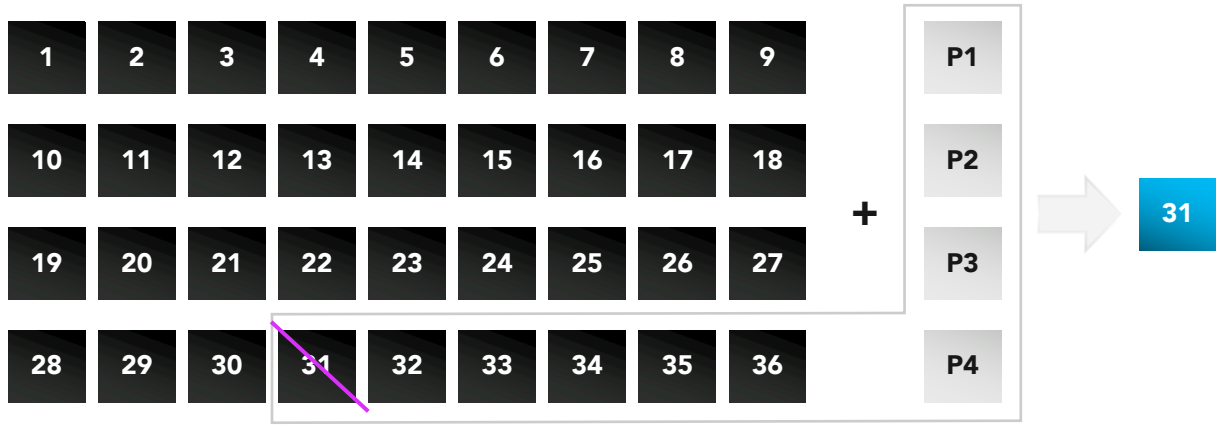

SIMILARITY: HOW IT WORKS

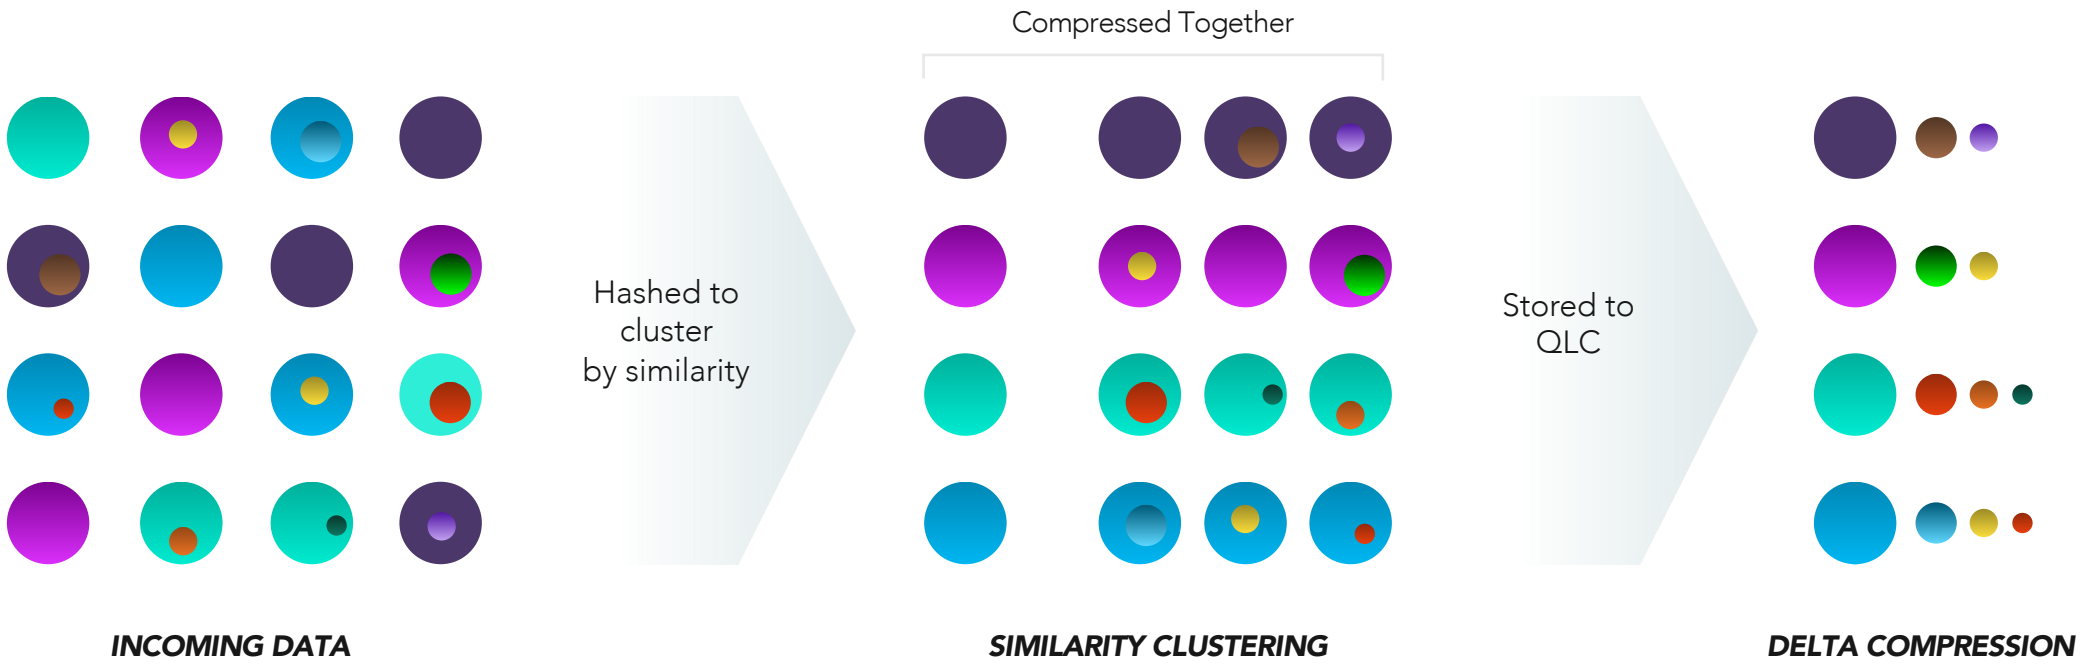

https://vastdata.com/whitepaper/#similarity-reduction-to-the-rescue

STORAGE EFFICIENCY

SIMILARITY IS GAME-CHANGING

COMPRESSION

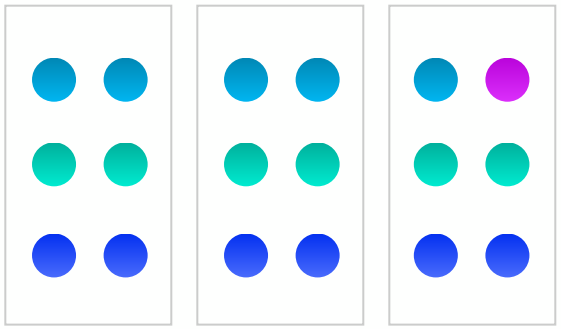

Fine Grained, But Local

DEDUPLICATION

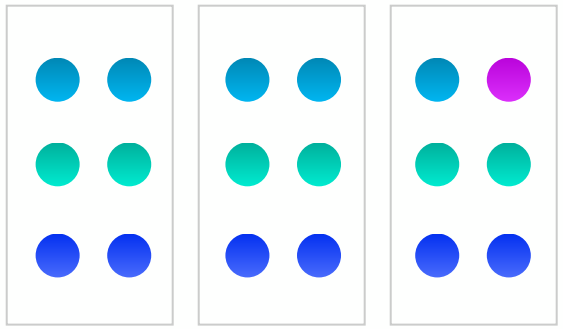

Global, But Coarse

VAST DATA SIMILARITY REDUCTION

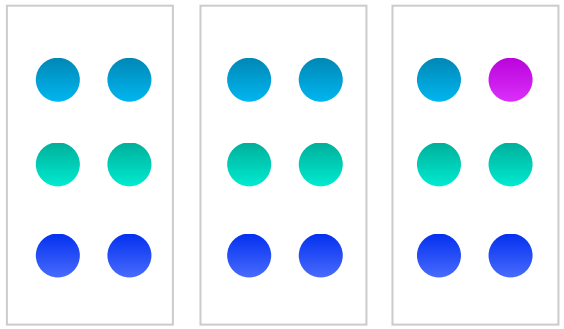

Global & Fine Grained

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

TRANSFORMING THE CALCULUS OF FLASH OWNERSHIP

CUSTOM TCO REPORTS
AVAILABLE AT VASTDATA.COM

VASTDATA.COM/SAVINGS

RADICAL SAVINGS TO BRING AN END TO HDD TIERS

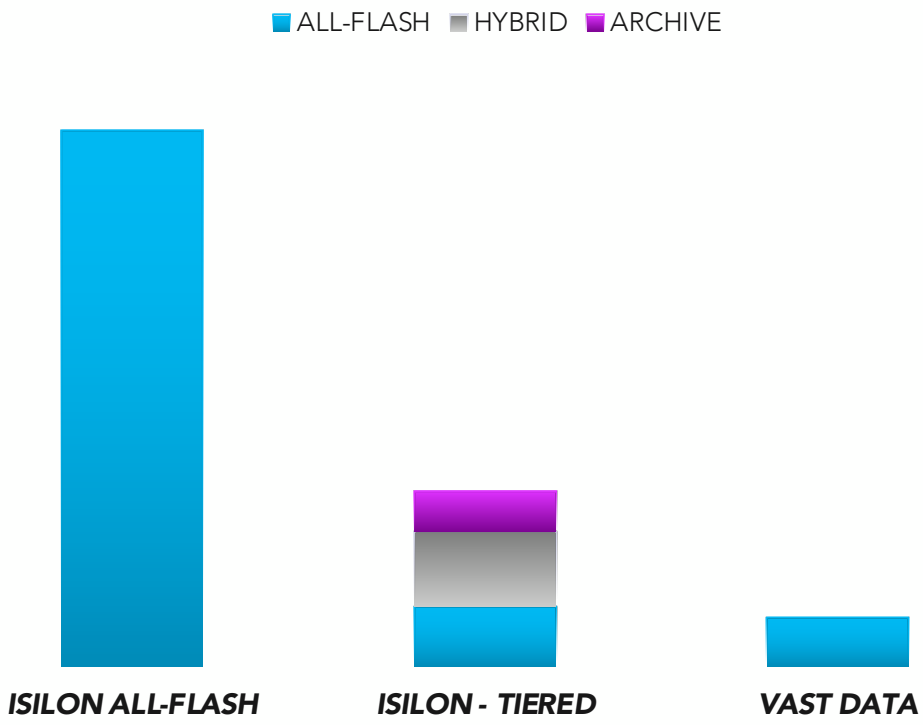

WITH YOU, EVERY STEP OF THE WAY

YOUR CO-PILOT, YOUR SUCCESS

— ***VAST CALL HOME RECEIVES TELEMETRY EVERY 5 MINUTES***

100Ks of metrics collected every day (opt-in)
Enables proactive maintenance, rapid response

— ***LEVEL 3 ENGINEERS PARTNERED WITH EVERY ACCOUNT***

Onboarding & daily operations
Capacity and performance monitoring and planning

— ***ONGOING SUCCESS MANAGEMENT***

Prioritization of feature requests
Advance software access for early testing

— ***INFINITE LIFECYCLE SERVICES***

Ongoing education for you and your team
Cluster expansion and upgrade coordination

SIMPLE, ANALYTICS-RICH MANAGEMENT

"SUPER IMPRESSED

- WEHI

AUTOMATED INFRASTRUCTURE

"YOUR API IS FANTASTIC 😀

- BUFFALO

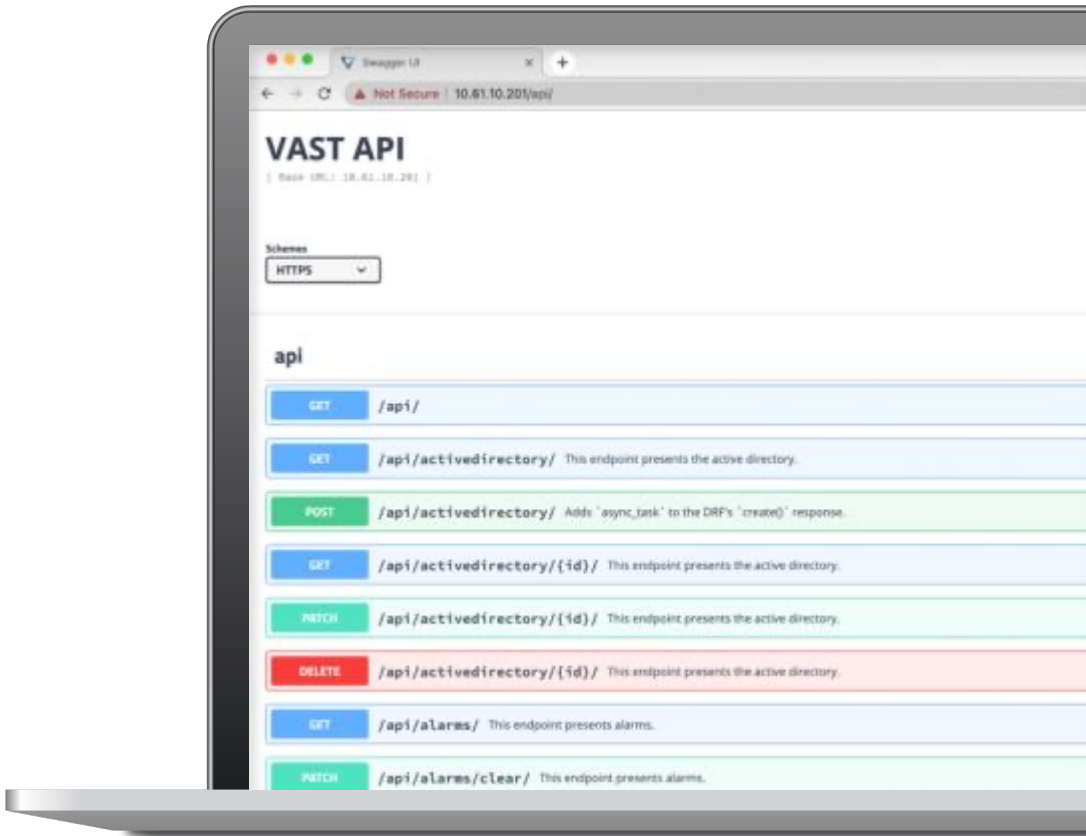

AI PERFORMANCE + NAS SIMPLICITY

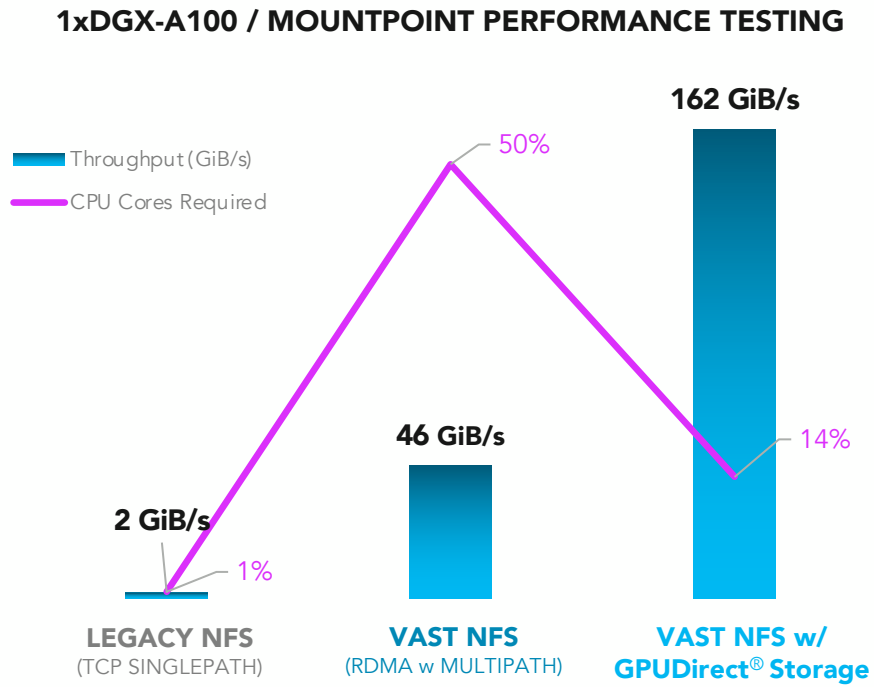

https://vastdata.com/whitepaper/#vast-nfs-parallel-file-system-speed,-nas-simplicity

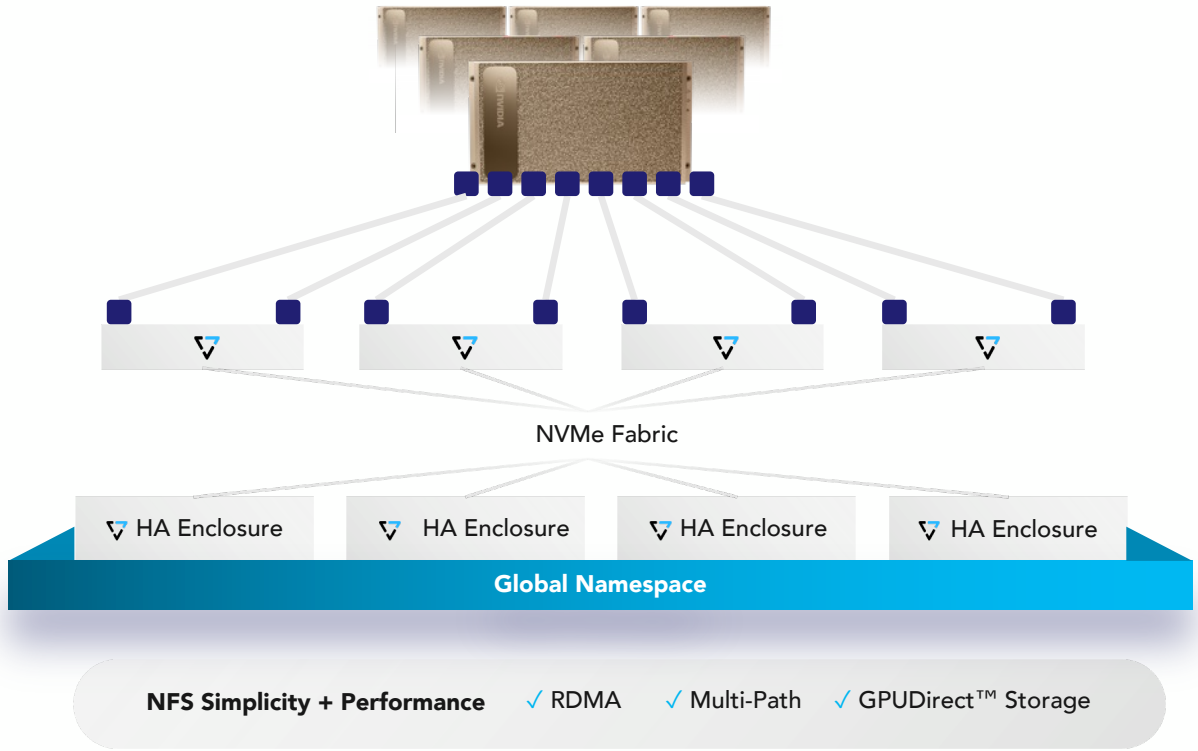

GEMINI
SAAS OFFERING

GEMINI

DISAGGREGATED STORAGE MEETS THE WORLD'S FIRST DISAGGREGATED COMMERCIAL OFFERING.

PURCHASE LIKE A CLOUD PROVIDER,
DEPLOY LIKE AN ENTERPRISE.

VAST

DISAGGREGATED HW SOURCING

Buy Like A Cloud Provider, Deploy Like An Enterprise

AN INFINITE STORAGE LIFECYCLE

Never Pay The Refresh Tax, Never Migrate Again

ZERO-COMPROMISES WARRANTY

Takes The Risk Out Of Adopting Universal Storage

A TRUSTED CO-PILOT

Our Team Is Now Your Team

BUY LIKE A CLOUD VENDOR, DEPLOY LIKE AN ENTERPRISE.

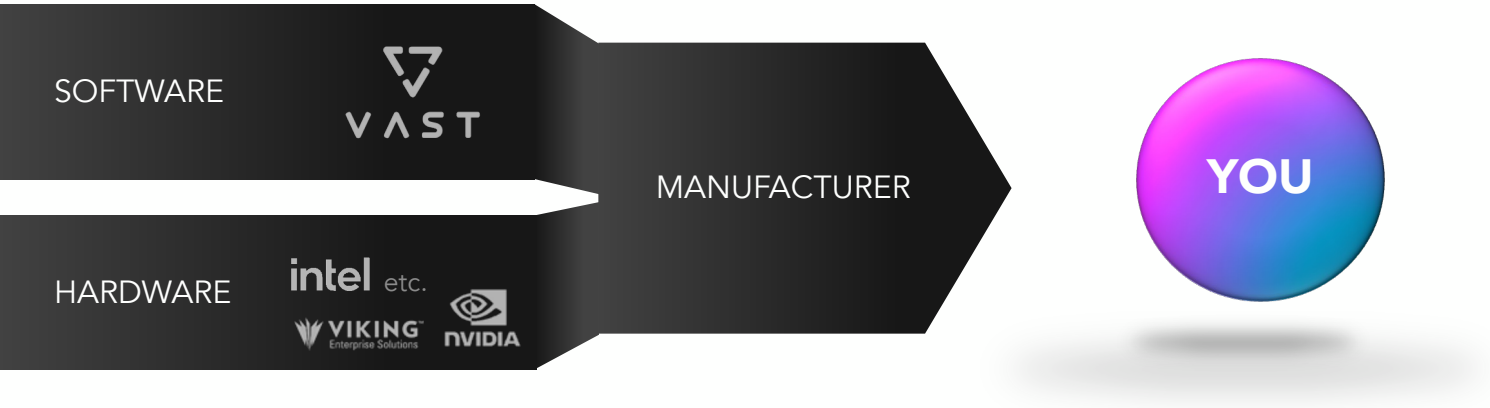

BUY HARDWARE AT COST.

VAST is not in the purchase path

Get full visibility to supply chain costs

MAXIMIZE BUYING POWER.

VAST & AVNET work the supply chain for you

A SIMPLE, TURNKEY APPLIANCE.

Comes to your door ready to turn on

VAST telemetry ensures 24/7 monitoring

VAST processes all HW RMAs

THE 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.

BENEFITS

THE INFINITE REFRESH CYCLE.

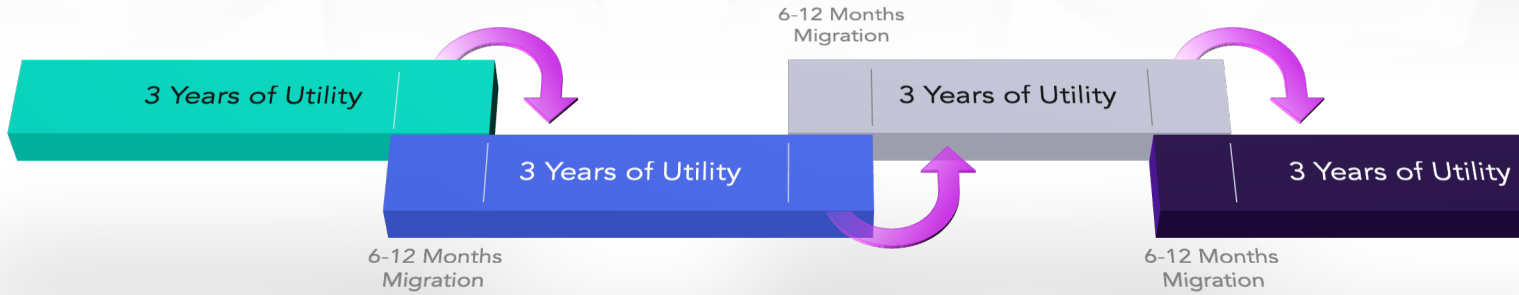

THE INFINITE STORAGE LIFECYCLE.

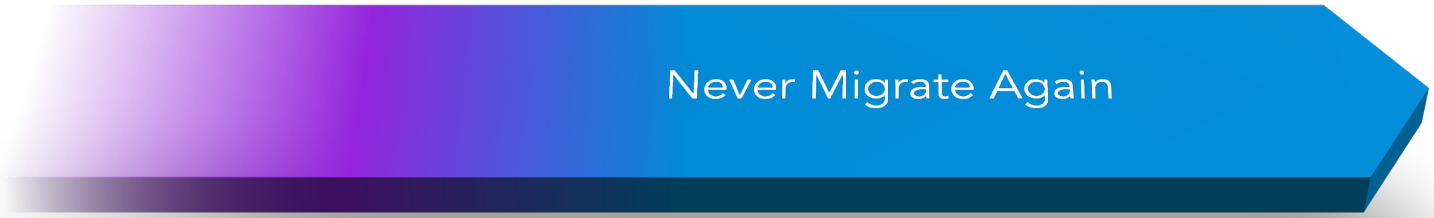

BENEFITS

ZERO COMPROMISES.

REMOVING THE RISK FROM
MODERNIZING YOUR DATA

- Uptime Guaranteed
- Persistence Guaranteed
- Data Reduction Superiority
- All-inclusive Storage Software
- 10 Years Of QLC Longevity
- 60-day Unconditional Right To Return

CUSTOMER STORIES

UNIVERSAL STORAGE USE CASES

VAST DATA FOR HPC & AI

- Blurring the lines between enterprise NAS & HPC storage
- Unprecedented simplicity, and 80x faster NFS performance
- Support for RDMA, Locking, Paralle IO, GPUDirect and more
- Embarrassingly-parallel scale for Exascale Computing
- Powering many world-leading HPC sites:

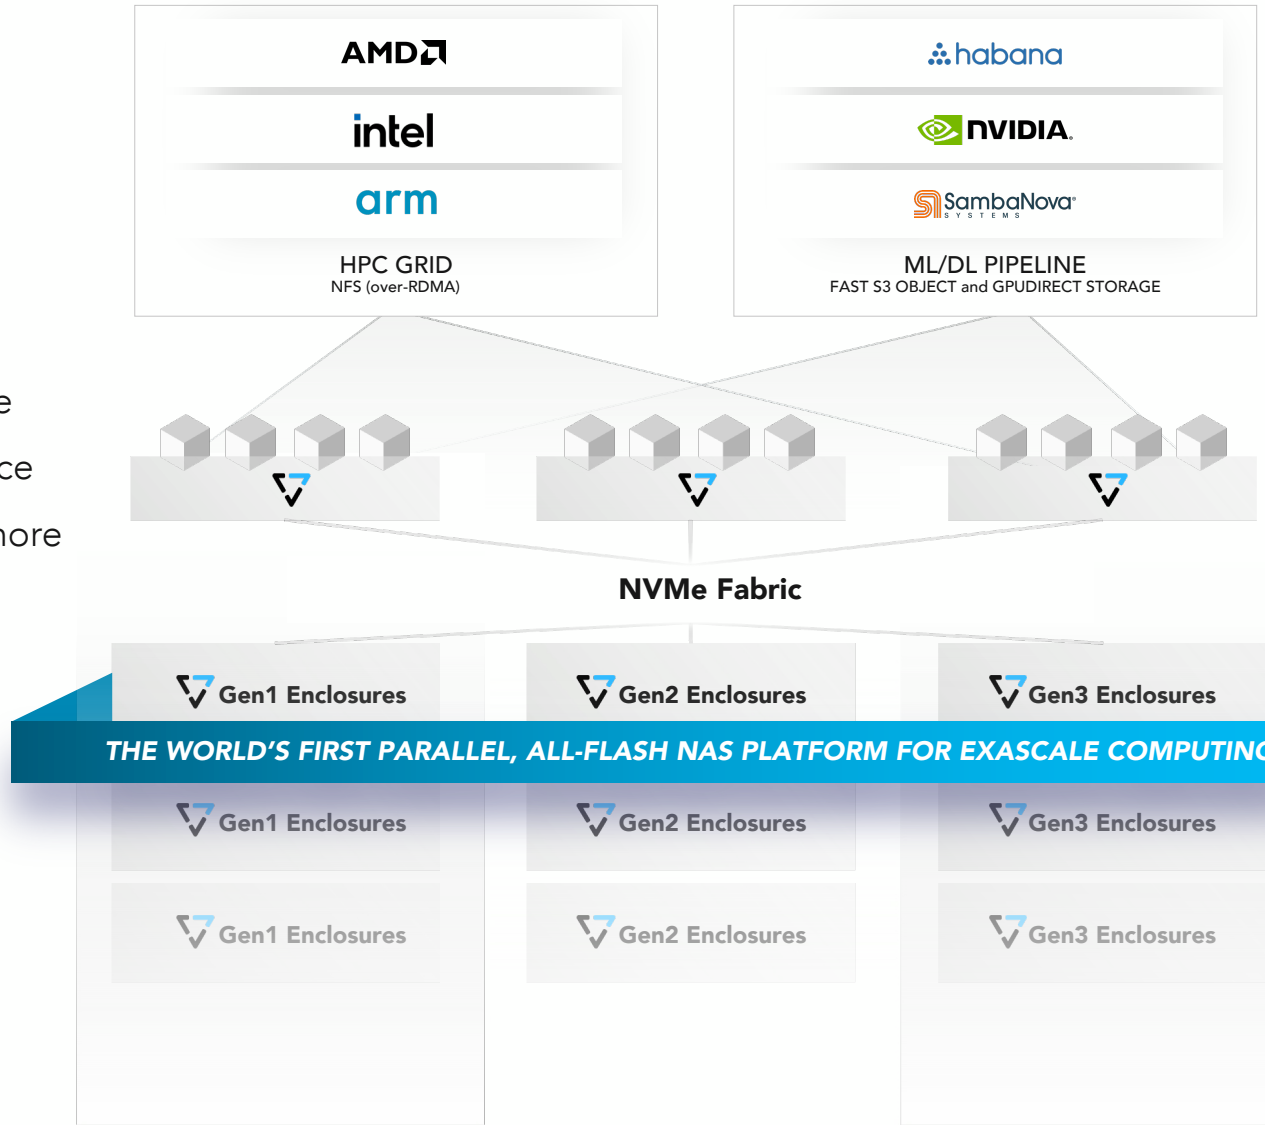

VAST

"VAST DATA IS VERY INTERESTING TO US... IT'S USED FOR ENTERPRISE NAS AND IN PLACE OF PARALLEL FILE SYSTEMS"

SUCCESS STORY

- Support for /home, /scratch, /training – mounted on 22 different clusters
- Rapid access to all data, support for the DOE's hardest challenges
- Embarrassingly parallel architecture for at-scale computing
- 4:1 data reduction on petabytes of file data

Lawrence Livermore National Laboratory

UNIVERSAL STORAGE USE CASES

VAST LIGHTSPEED FOR ML/DL

- 80x faster performance than legacy NAS
- No HPC storage complexity
- Radical flash savings, eliminates tiering complexity

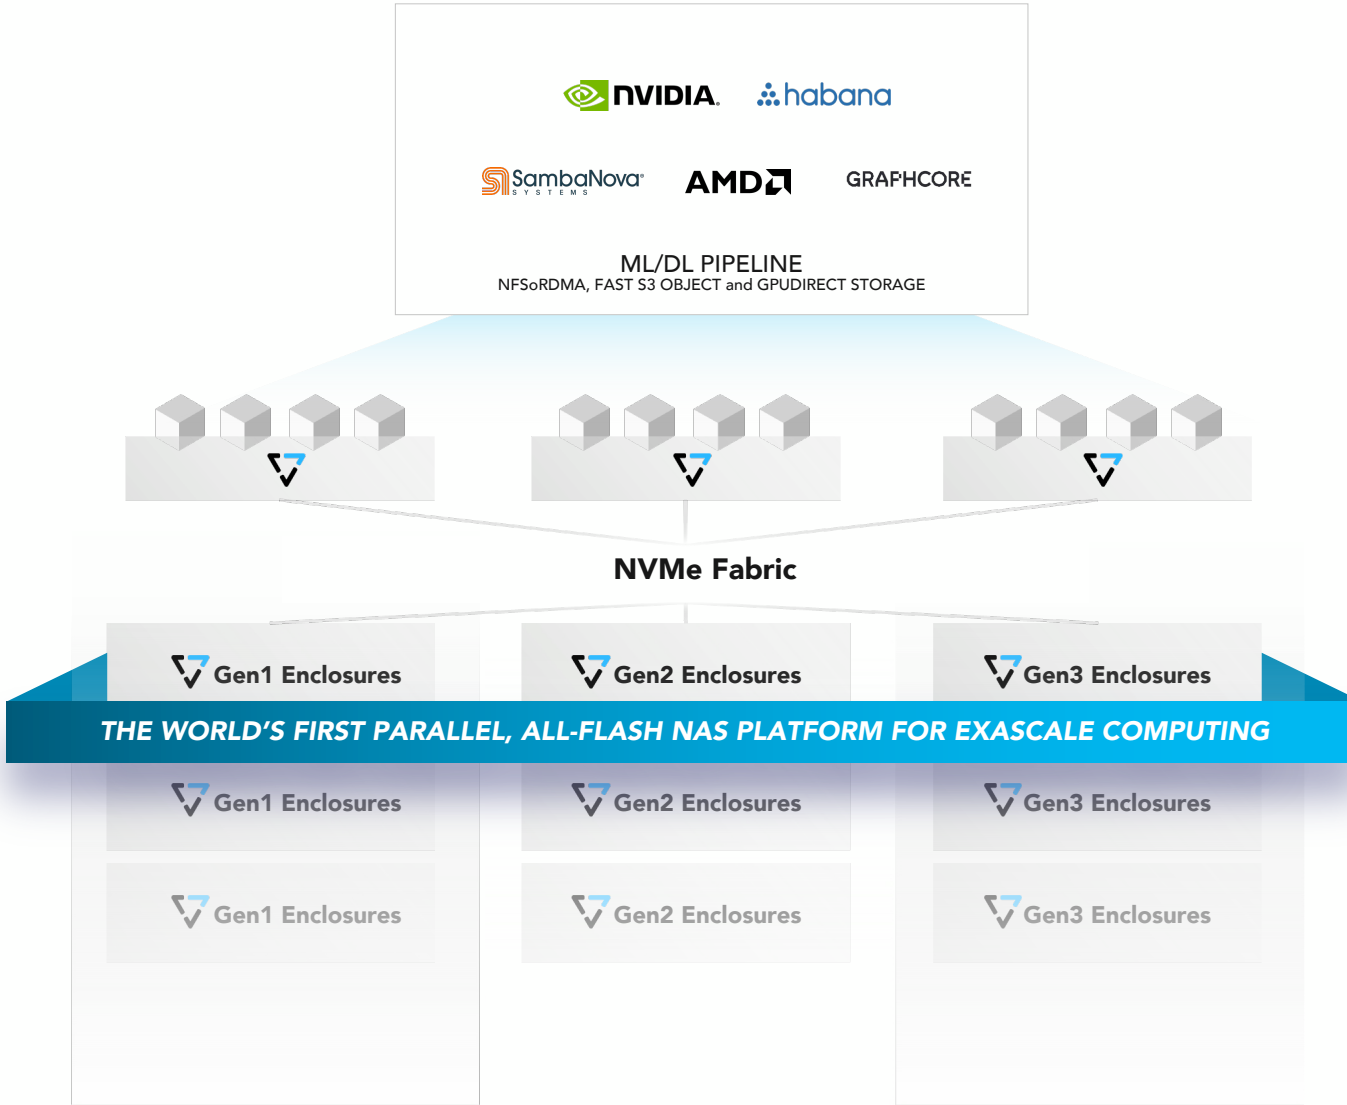

"WE'VE NEVER GOTTEN THIS LEVEL OF SUPPORT AND NEAR-INSTANTANEOUS FEEDBACK FROM ANY VENDOR BEFORE. IT'S AMAZING.

SUCCESS STORY

- All-flash infrastructure for NVIDIA A100s – less than hybrid alternatives
- Takes the guesswork out of deploying data science at scale
- GPUDirect-ready storage

Athinoula A. Martinos Center For Biomedical Imaging

VAST

VAST's new many-port DGX reference architecture delivers superior deep learning throughput when compared to legacy approaches.

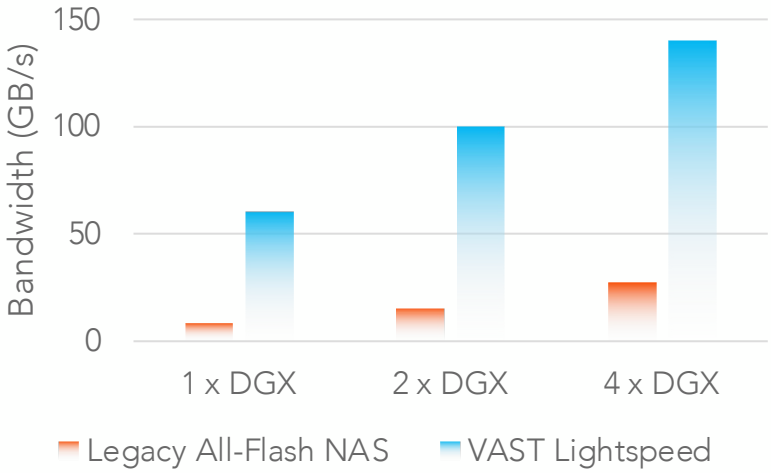

UNIVERSAL STORAGE USE CASES

VAST DATA FOR BIG DATA

- Blurring the lines between DAS and data lake storage
- Faster-than-DAS performance for modern workloads
- Support for RDMA, GPUDirect and more
- No shuffles, every node enjoys direct data access
- Powering many world-leading big data toolkits:

kafka splunk> Apache Spark

presto elastic dremio VERTICA

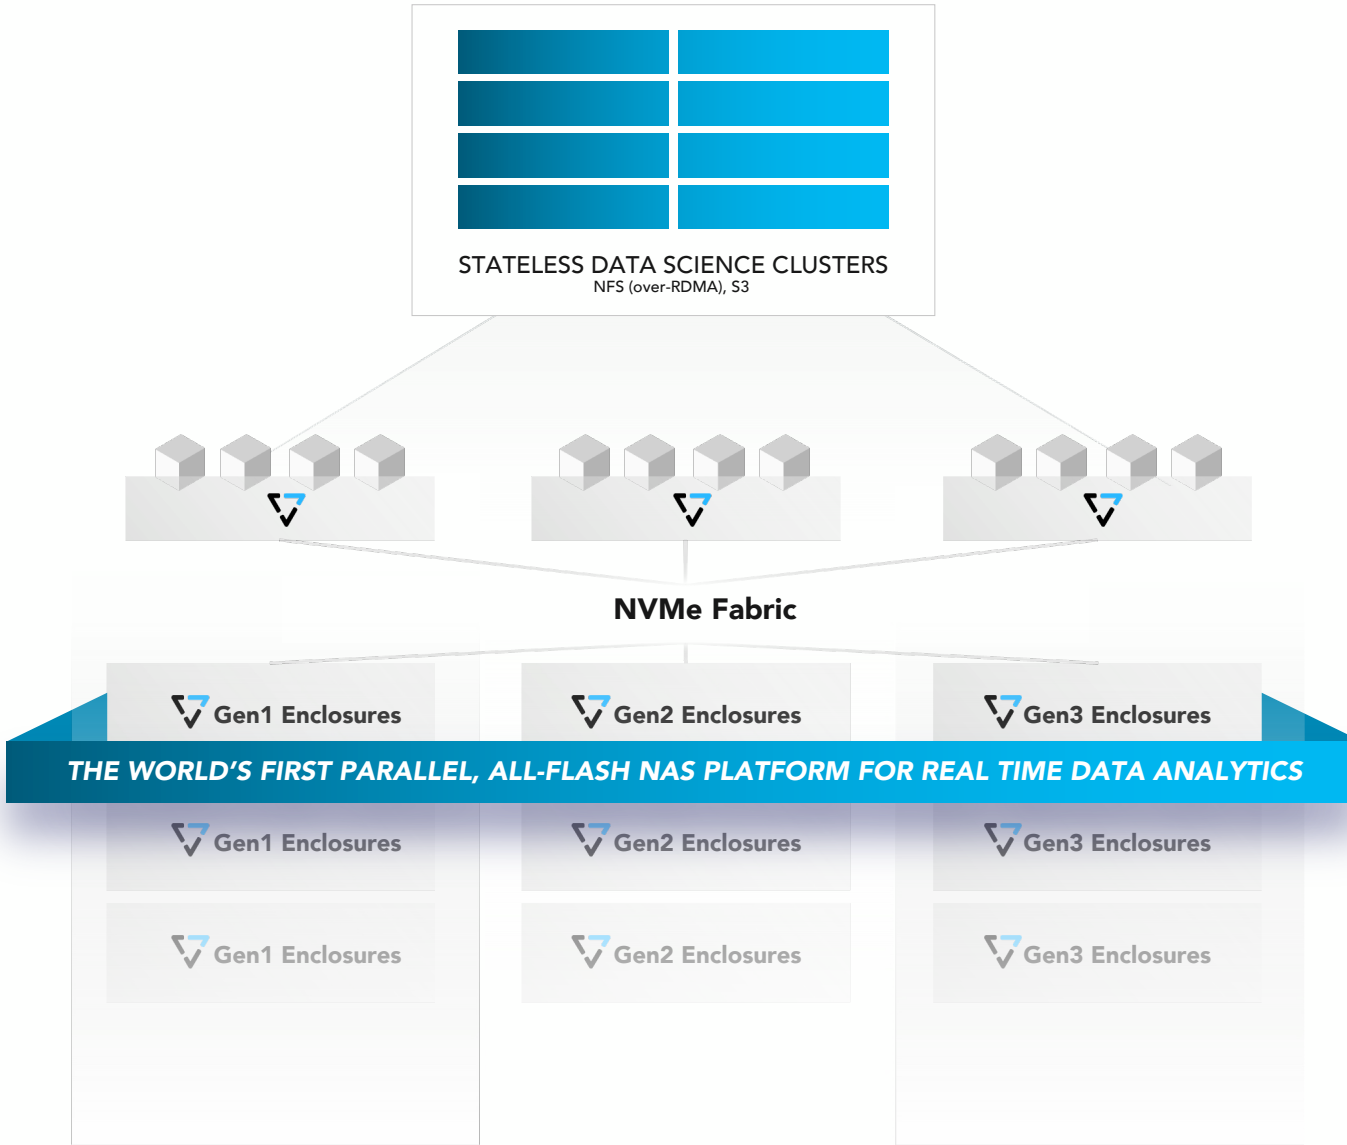

VAST

BIG DATA BENCHMARKING

DASE SCALE IN PRACTICE

PostgreSQL testing at a US Government customer

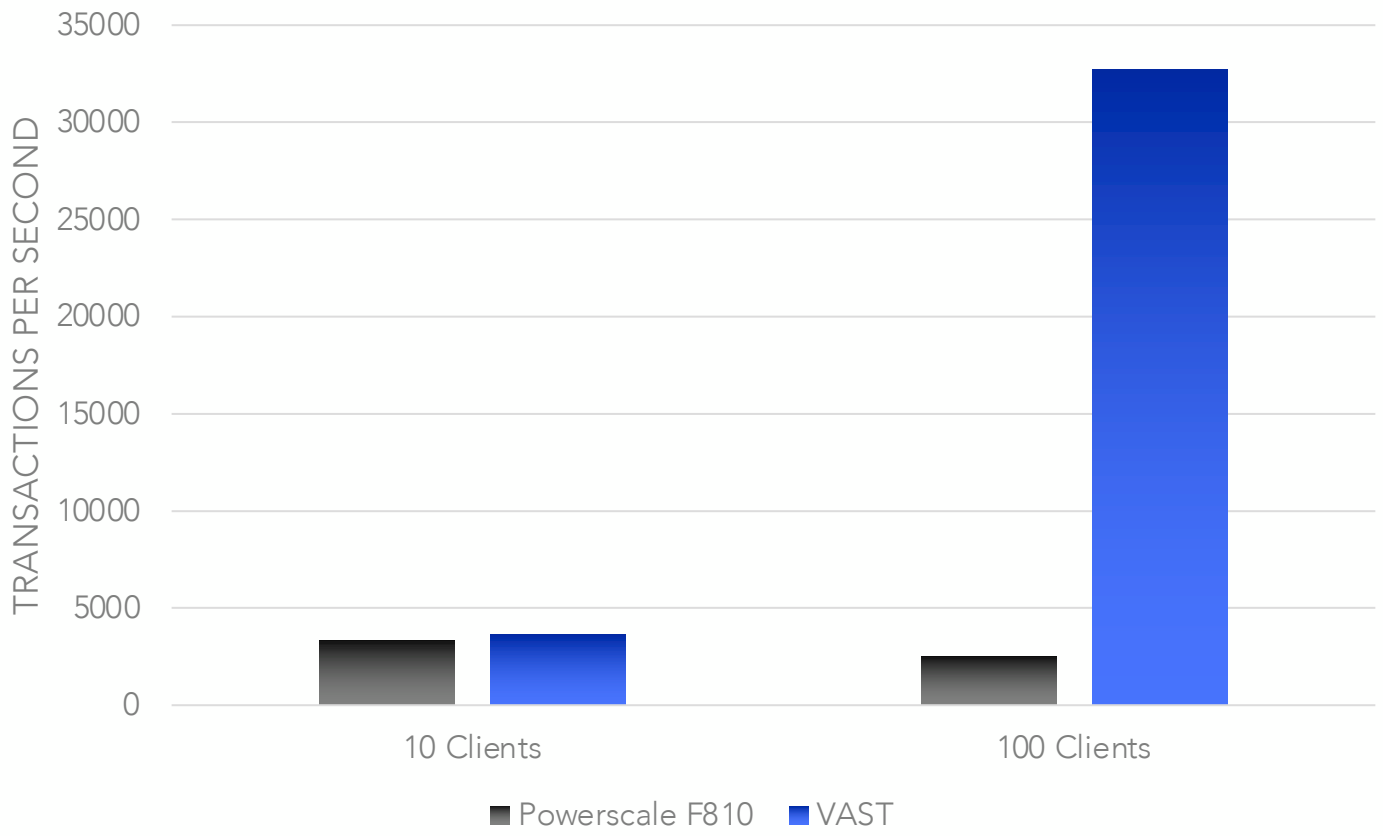

UNIVERSAL STORAGE FOR BIO-IT

THE LIFE SCIENCES DISCOVERY ENGINE.

- All-flash infrastructure to accelerate and simplify genomics, proteomics, brain research, imagery, etc.
- A simple appliance with linear HPC performance
- 10x faster performance than HDD solutions
- The platform of choice for life science leaders:

Boston Children's Hospital — Until every child is well
HudsonAlpha Institute for Biotechnology
NIH Center for Information Technology
Department of Veterans Affairs
Harvard Medical School
Invitae
EMBL-EBI
Ginkgo Bioworks
National Cancer Institute
WEHI — brighter together
Broad Institute

VAST

UNIVERSAL STORAGE USE CASES

AFFORDABLE RAPID RESTORES

- Restore and recover at the speed of flash
- Up to 10X faster restores than disk-based systems
- Archive economics (flash at the cost of HDDs)
- Accelerate and data-reduce backup jobs at any scale
- Certified with leading backup solutions:

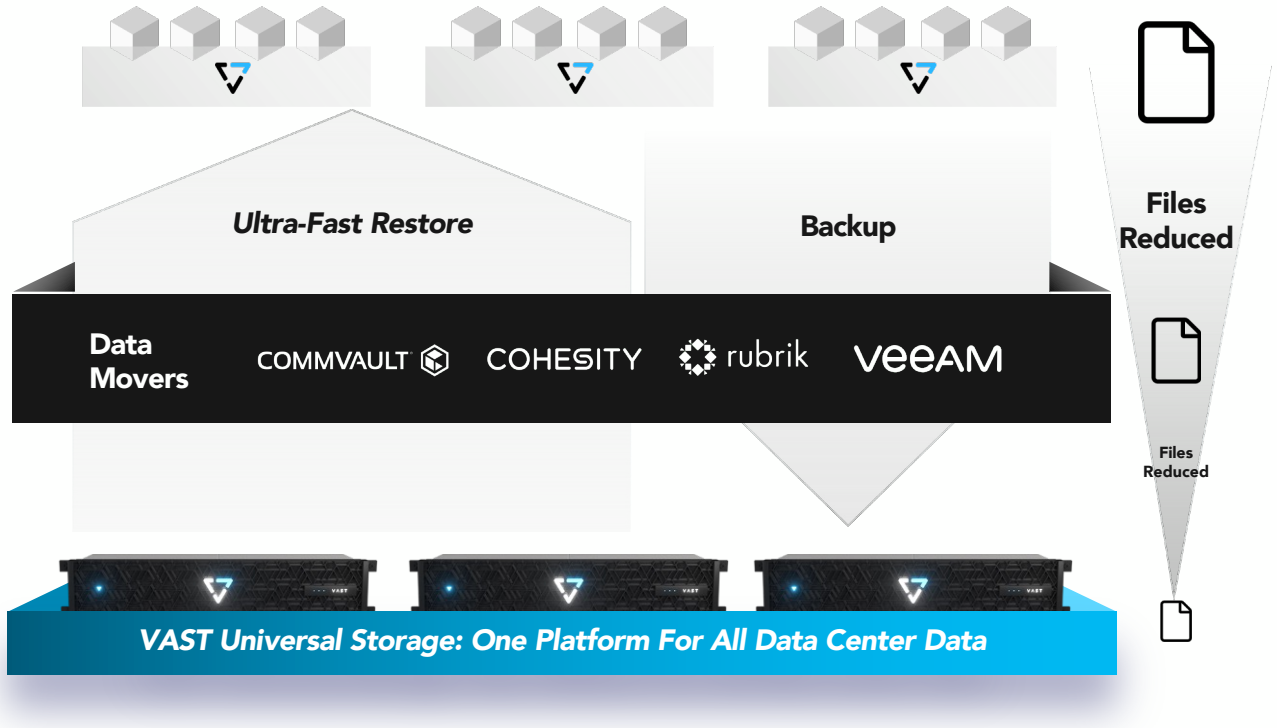

"THE SYSTEM IS STABLE AND SIMPLE...THERE'S LITTLE OVERHEAD FOR THE TEAM TO MANAGE"

SUCCESS STORY

- Veritas NetBackup solution providing flash-level responsiveness
- Hard drive cost economics with a smaller physical footprint
- Data reduction beyond backup software (with encryption enabled)
- System scales to support additional applications (archive, K8S, etc.)

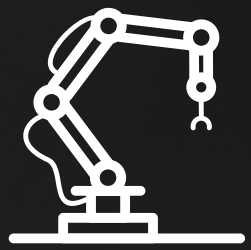

(Confidential Auto Mfgr)

UNIVERSAL STORAGE USE CASES

VAST DATA FOR MEDIA & ENTERTAINMENT

- Edit, render and stream data from all-flash infrastructure
- Flash performance WITH archive economics
- Enterprise NAS simplicity
  - (no file system agents, Fibre Channel, etc.)
- Multi-protocol namespace (NFS, SMB, S3)
- Scale editors, playout servers and render nodes with limitless capacity

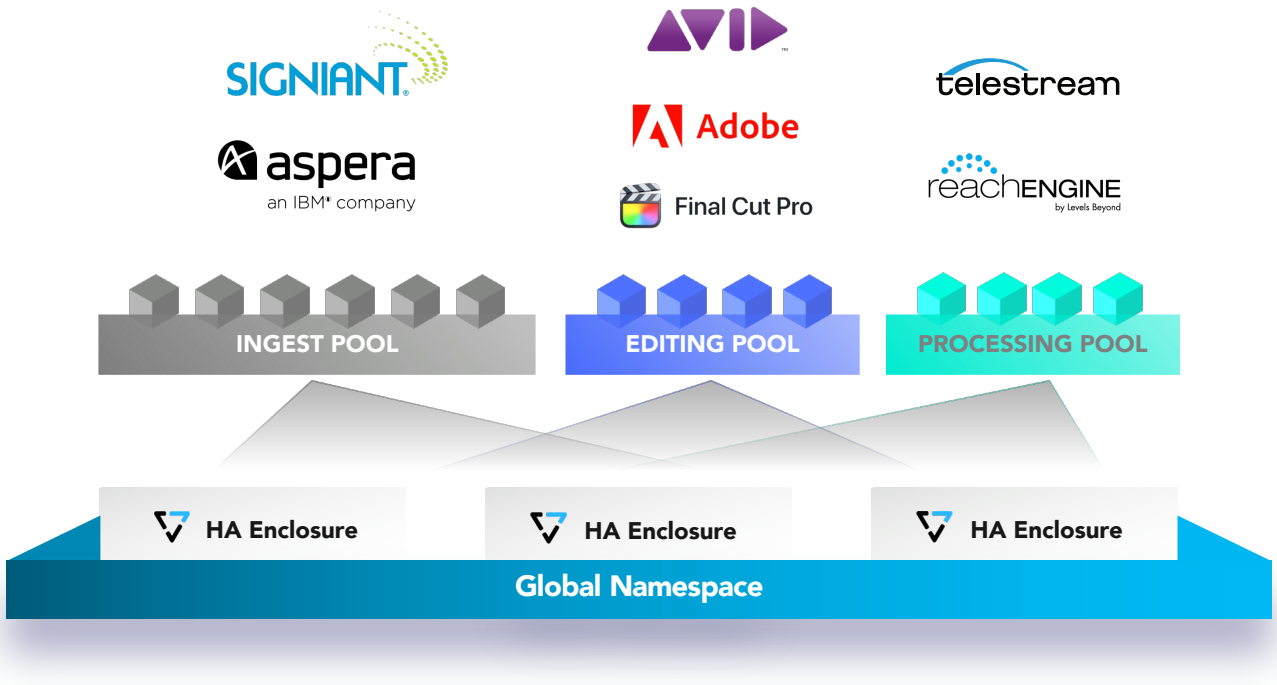

"WHAT WE'RE DOING WOULD BE COST PROHIBITIVE TO SCALE OUT USING ISILON TECHNOLOGY"

SUCCESS STORY

- Scaled CPUs independent of capacity to reduce data access latency
- Pooling CNodes enabled QoS for specific departments/projects
- Data reduction of up to 3:1 for in-flight jobs reducing required investment
- Co-Pilots deliver seamless operational experience

(Confidential Animation Studio)

"WE NEEDED WORKFLOW ISOLATION WITH KILLER PERFORMANCE AND VAST DELIVERED"

SUCCESS STORY

- Ease of use and simplicity that doesn't get more complicated with scale
- Multi-tenancy and workload concurrency accelerated time-to-market
- Flash performance WITH archive economics
- Lower administrative cost

(Confidential Sports Broadcaster)

UNIVERSAL STORAGE USE CASES

VAST DATA FOR NATIONAL SECURITY

- A simple, exabyte-scale storage cluster appliance
- Scale-out file + object for analytics, AI, HPC, Backup, etc
- Access data in real-time from one tier of affordable flash
- De-risks the data processing agenda for any mission
- Already powering many government agencies:

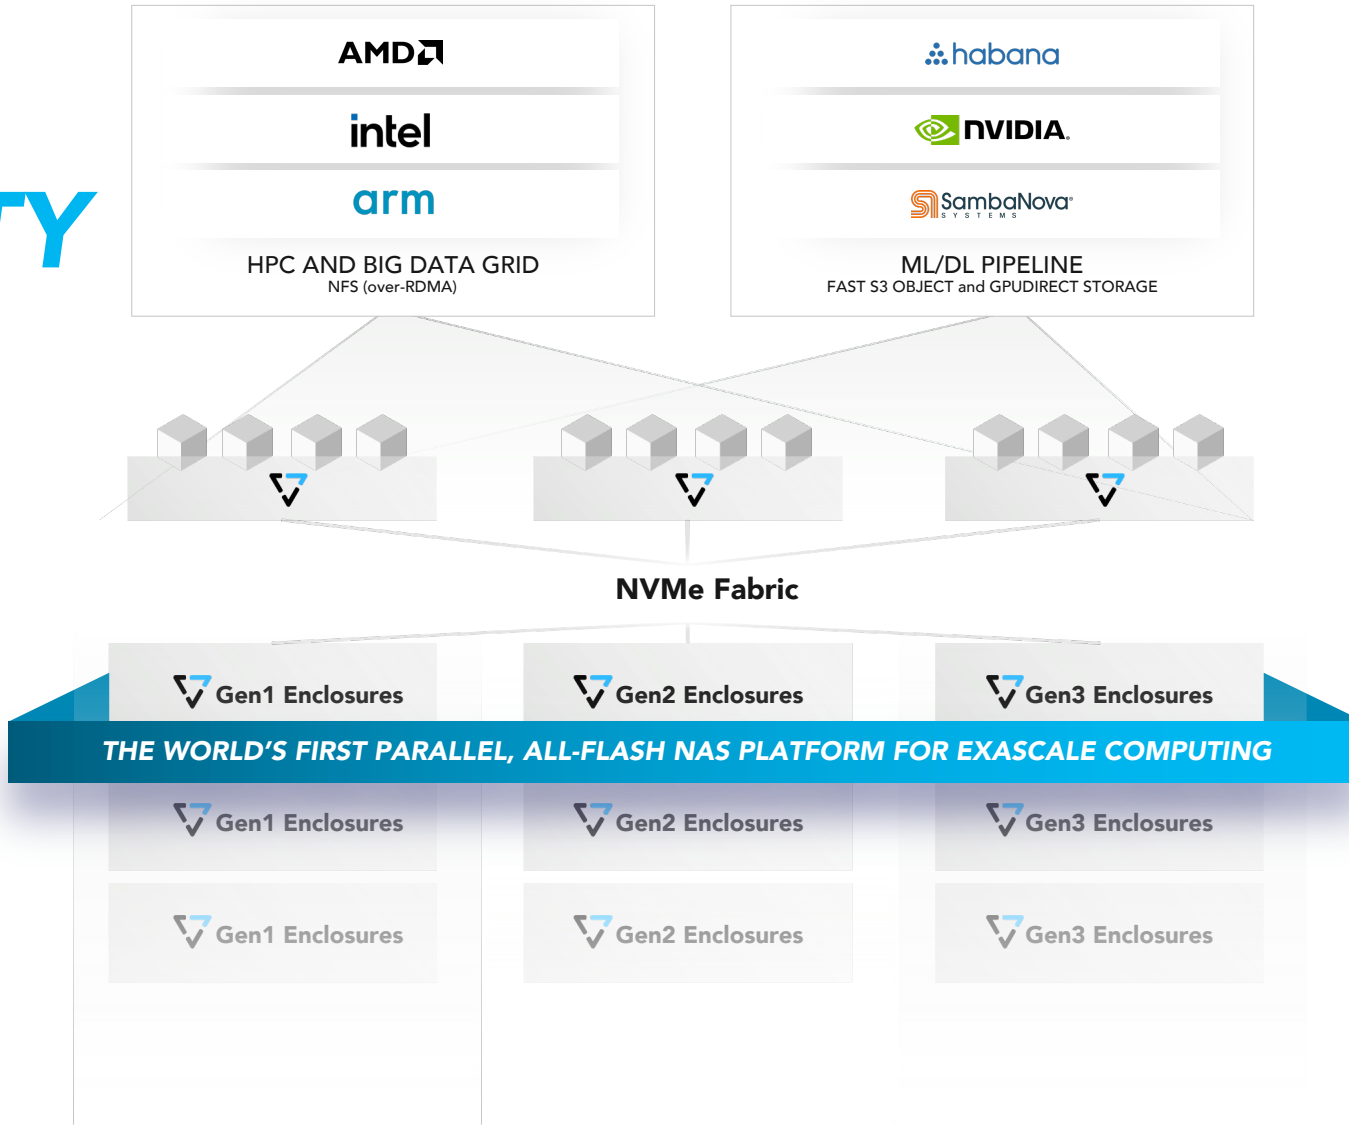

VAST

DISAGGREGATED, SHARED-EVERYTHING

THE STORAGE ARCHITECTURE OF THE FUTURE

- IDC, 2020

- The only enterprise-class storage platform scalable to exabytes
- Groundbreaking economics and data reduction
- 10 year hardware warranty
- Supercompute performance with the ease of a NAS
- Industry leading power efficiency
- Lowest eWaste in the industry
- Customer driven, independent scaling of performance or capacity
- Better than enterprise reliability: 99.9999999% Architecture

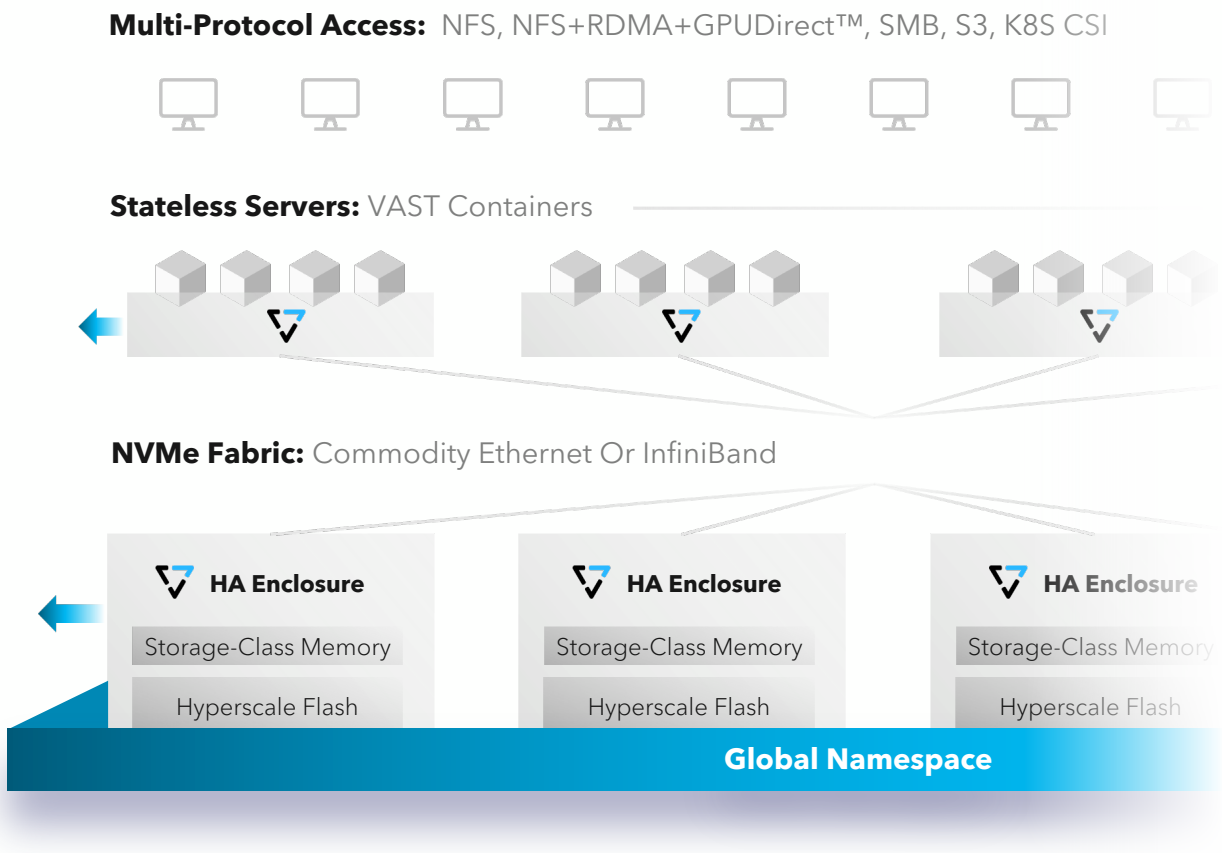

INTRODUCING:
UNIVERSAL STORAGE

BRINGING AN END TO THE HDD ERA

Petr Plodik, Sales Director, petr.plodik@mcomputers.cz

VAST