



Flash Memory Summit

Data-Driven Applications Gravitate to Computational Storage

Thad Omura

EVP of Marketing & Operations, ScaleFlux™

Flash Memory Summit 2019
Santa Clara, CA




The Answer to All Your Questions is...



Question:

What transformative technology are you going to learn about right now?





- PCIe Standard Form Factors & Power
- Open Channel-like Block Storage
 - High throughput, low latency
 - Workload tuning
 - I/O optimization: Atomic Writes & Streams
- HW *Adaptable* Compute Engines
 - Existing APIs for GZIP & EC



Big Data



Fast Data



Fast Data

First Name	Last Name	Email	Country	IP address
Dalton	Kramer	dalton@email.com	France	211.91.226.108
Gita	Tetterton	gita@email.com	USA	222.153.179.100
Weston	Jurgens	weston@email.com	Spain	203.123.236.1
Brad	Chupp	brad@email.com	France	202.183.111.122
Marybeth	Baumann	marybeth@email.com	Italy	214.132.168.129
Allyson	Feder	allyson@email.com	Italy	182.108.190.85
Lucile	Folks	lucile@email.com	Greece	18.64.161.62
Mickey	Rusk	mickey@email.com	Canada	40.18.115.207
Clarine	Esslinger	clarine@email.com	Greece	185.134.23.86
Kimberly	Penny	kimberly@email.com	France	34.72.165.11
Colleen	Kellough	colleen@email.com	USA	73.51.152.185
Nettie	Edmonds	nettie@email.com	Spain	94.133.138.234
Duncan	Rickenbacker	duncan@email.com	France	211.91.226.108
Marchelle	Diedrich	marchelle@email.com	Italy	222.153.179.100
Mariano	Murrell	mariano@email.com	Italy	203.123.236.1



Low Latency, ~~subset~~ dataset
Transactional Database

Big Data

Last Name	Country	IP address
Kramer	France	211.91.226.108
Tetterton	USA	222.153.179.100
Jurgens	Spain	203.123.236.1
Chupp	France	202.183.111.122
Baumann	Italy	214.132.168.129
Feder	Italy	182.108.190.85
Folks	Greece	18.64.161.62
Rusk	Canada	40.18.115.207
Esslinger	Greece	185.134.23.86
Penny	France	34.72.165.11
Kellough	USA	73.51.152.185
Edmonds	Spain	94.133.138.234
Rickenbacker	France	211.91.226.108
Diedrich	Italy	222.153.179.100
Murrell	Italy	203.123.236.1

Indexes
.
.
.

~~Batch~~, full dataset
Analytical Database


Question:

**What Storage Solution Unifies
Fast & Big Data?**






Computational Storage Unifies Fast & Big Data

Intensive Compute
Extract/Transform
Record Filtering




- Low-Latency **FAST** Data Transactions
- Scale **BIG** Data analytics workloads with capacity
- Commodity infrastructure

Analytical Queries on Transactional Data




768GB DRAM
8 host CPU cores, 100GB, 1 drive per DB instance
TPC-H Benchmark Q6

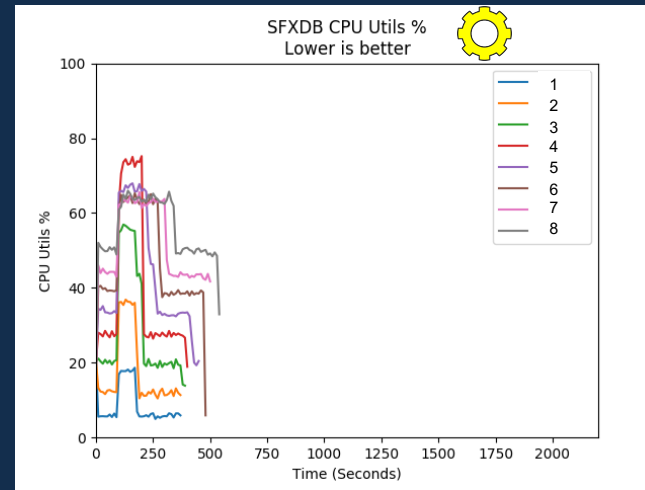


Compute

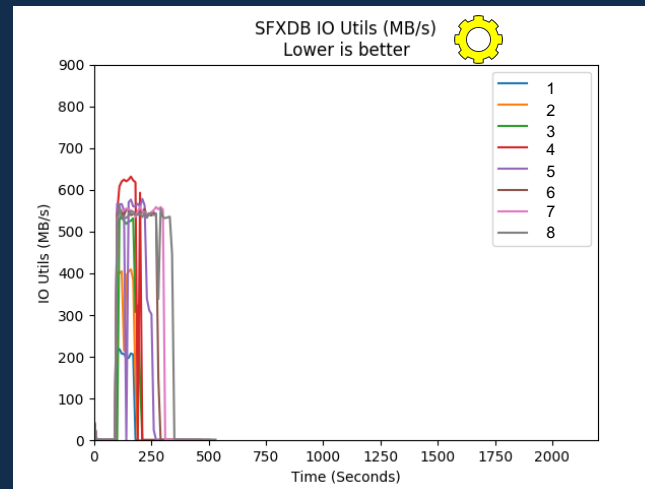

CPU



Computational Storage



I/O






POLARDB

- Cloud native database with distributed storage
- HTAP: Hybrid Transactional-Analytical Processing
 - No lag for analytics
 - Reduced storage capacity required vs. traditional architectures
- Embrace Computational Storage
 - Best price/performance, new class of analytical responsiveness
 - Commodity server/storage infrastructure

Computational Storage Value @ Cloud Scale

Alibaba Cloud **POLARDB**



4 database + 3 storage nodes cluster
Commodity HW w/ Computational Storage

Computational Storage



Scan
Throughput

8GB/s

24GB/s

Query Latency



Up to 70%
Reduction

TCO:
Scale CPU & Memory
3X more Efficiently

Performance:
Best User Experience

“ScaleFlux Computational Storage is required to delivery consistent transactional processing performance with analytical workloads on the same data.” – Alibaba RDS



Question:

What storage solution offers a new paradigm to scale compute resources with workload capacity?



New Paradigm for Infrastructure Scaling

- Scale performance with capacity @ cloud scale
- Enable a new class of service
 - Optimize latency, eliminate stale data analysis
- Infrastructure scaling TCO benefits
- Parallelize intense compute @ data



Compute Acceleration Along the Whole Stack



IN STORAGE
ANALYTICAL DB



DATA PATH
COMPRESSION/
DECOMPRESSION

- ✓ Fixed algorithms
- ✓ Bit-wise compare/manipulate
- ✓ Required across all data

Database, AI/ML, Genomics,

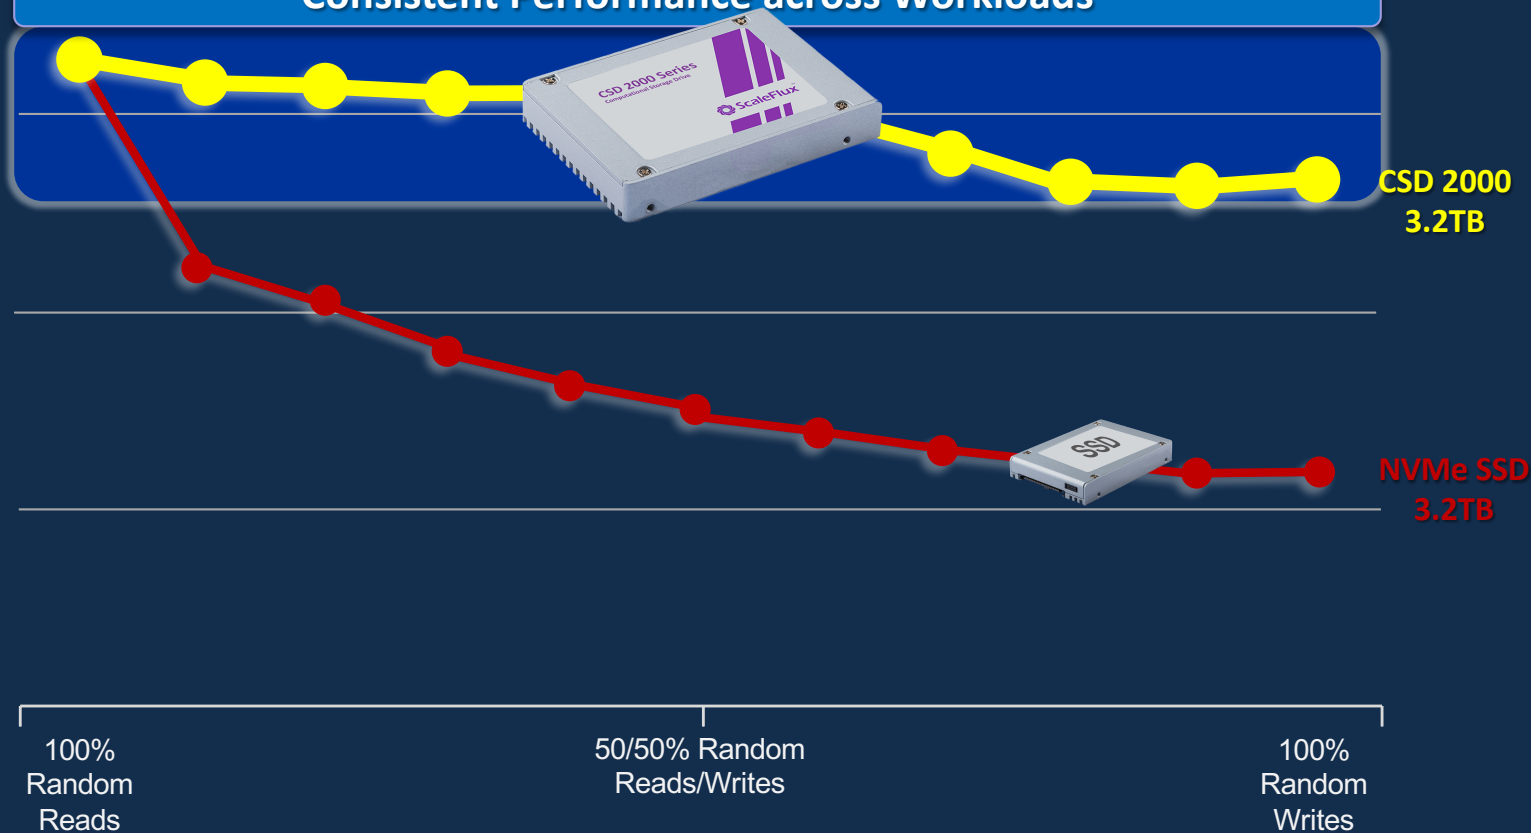


Compression (GZIP)
Erasure Coding (RS)
Security (AES)
Authentication (SHA)
Error Checking (CRC)
Application & Storage Acceleration
Deduplication
...

Industry's Most Advanced Computational Storage Drive

Data Path Compression & Decompression

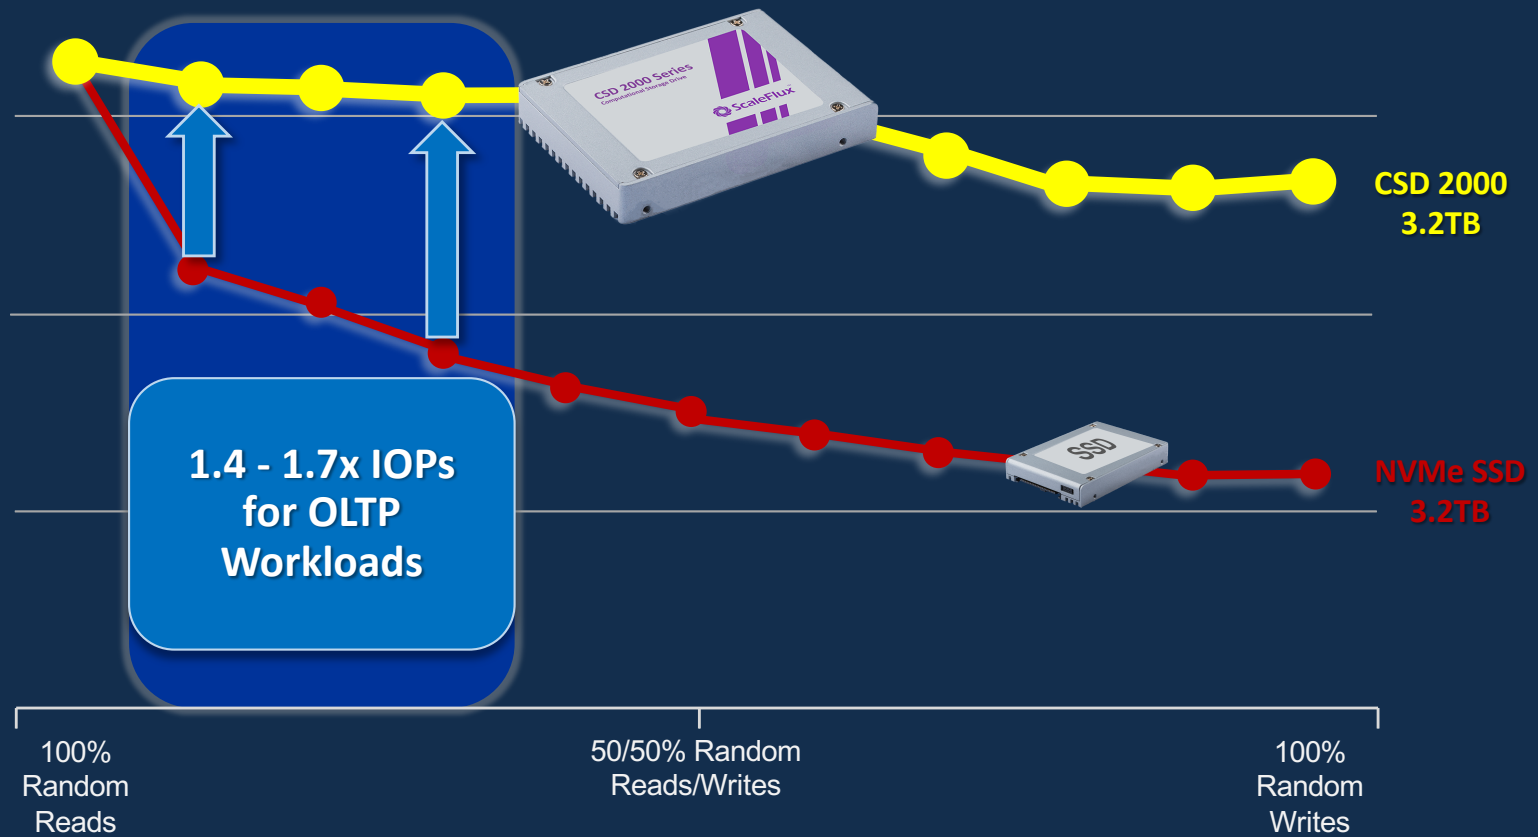
Consistent Performance across Workloads



(Random, 2.5:1 compressible, FIO, 4KB I/Os, 8 Jobs, 32 QD)



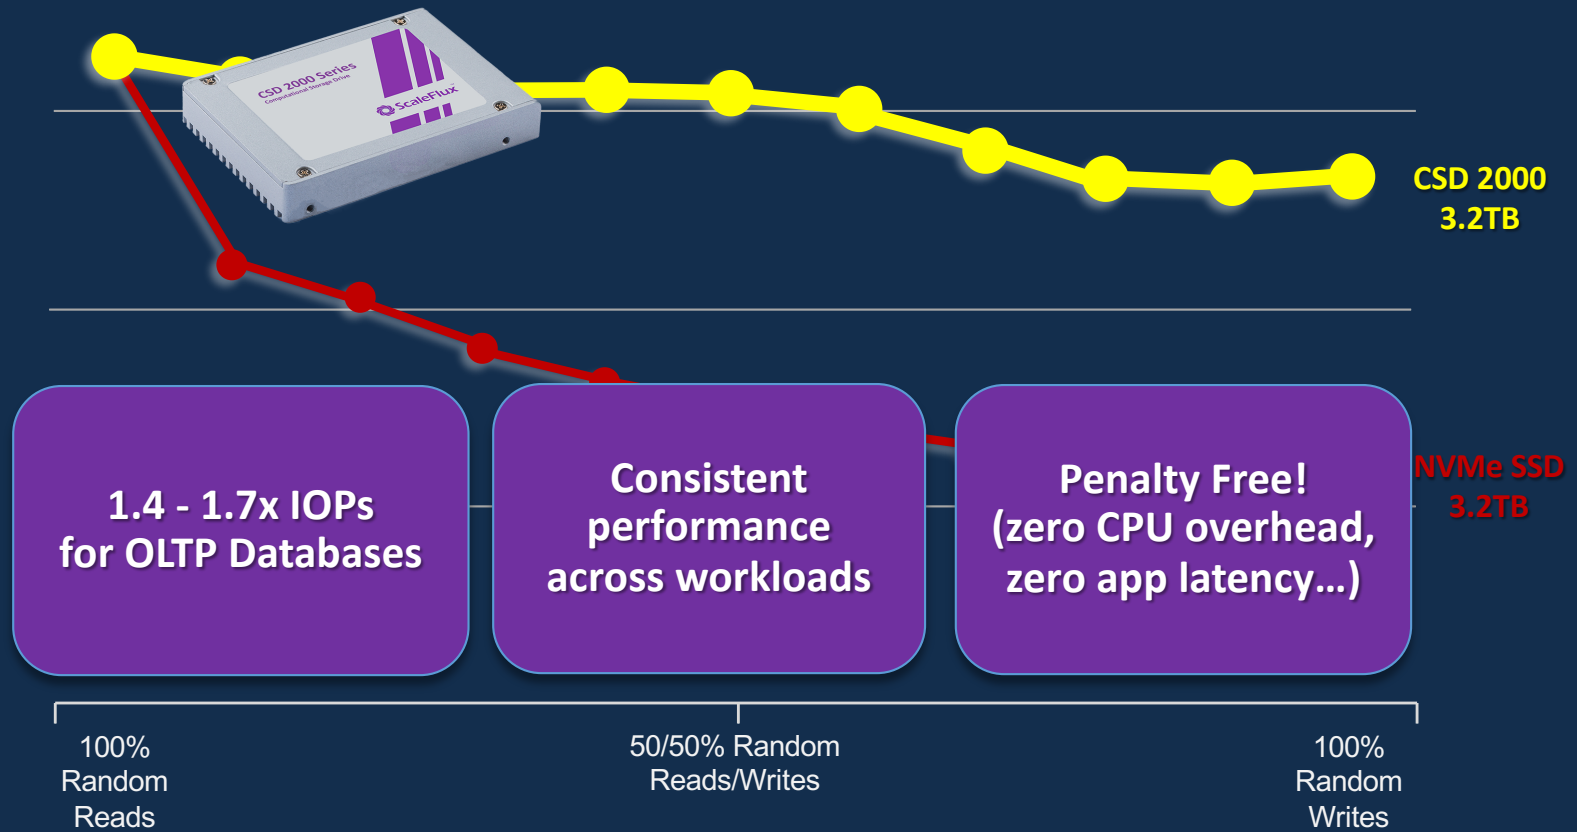
Data Path Compression & Decompression



**1.4 - 1.7x IOPS
for OLTP
Workloads**


(Random, 2.5:1 compressible, FIO, 4KB I/Os, 8 Jobs, 32 QD)

Data Path Compression & Decompression

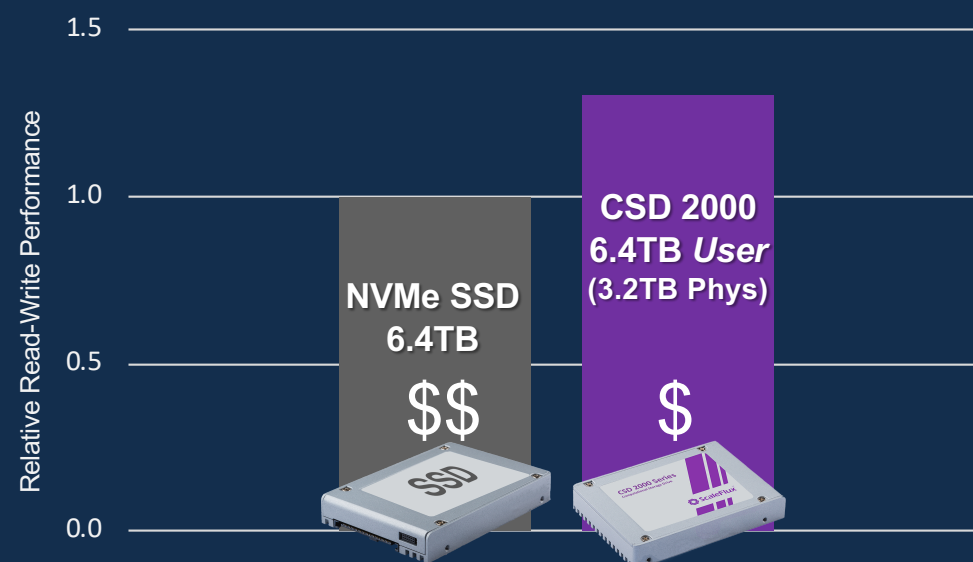


(Random, 2.5:1 compressible, FIO, 4KB I/Os, 8 Jobs, 32 QD)

OLTP Optimized



~50% More Performance



~25% More Performance, Half the Spend

Sysbench with MySQL, InnoDB, 50M records, 64 Threads, 1hr Test run, Read-Write test
Perf Opt Mode: 2.4TB Raw Files Size; 0.9TB Compressed. Cost Opt Mode: 4.8TB Raw File Size, 1.6TB Compressed



Question:

What storage solution cohesively opens **APPLICATION**
and **INFRASTRUCTURE** computing bottlenecks?





Responsive Performance

OLTP and Mixed R/W
Unify Fast & Big Data
Minimize Data Movement



Affordable Scaling

Reduce \$/GB
Performance w/ Capacity
Commodity Infrastructure



Agile Infrastructure

Transparent Benefits
In-System HW Updatable
Application Tuning

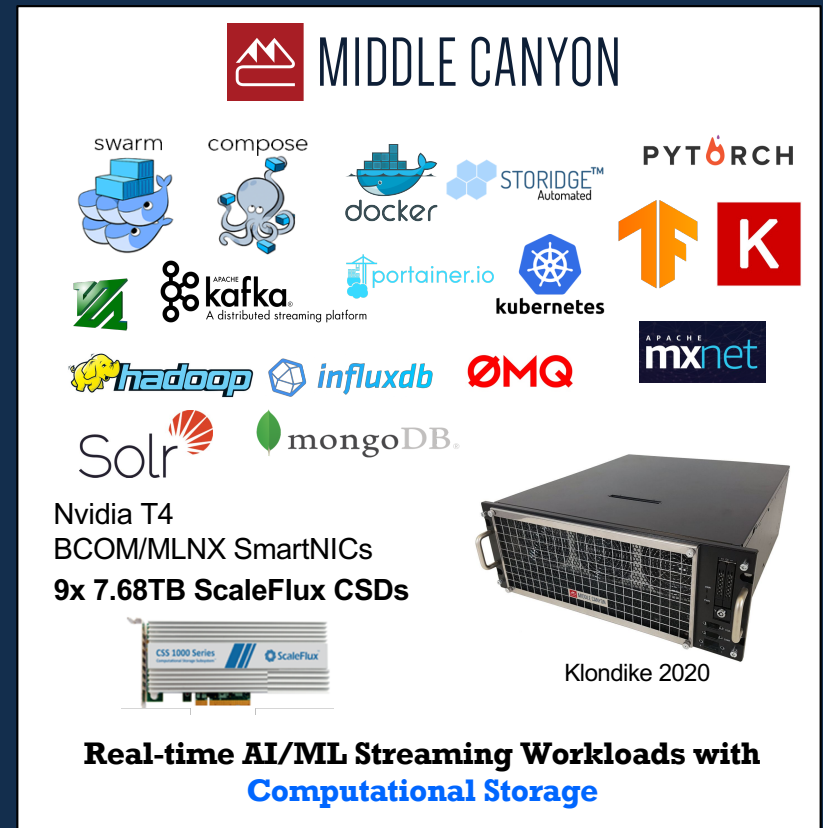
A Growing Ecosystem of Innovative Systems

DELL EMC

Extreme Scale Infrastructure (ESI)
Large Scale Customers



PowerEdge R640 (Intel Skylake)
PowerEdge R6415 (AMD Epyc)



MIDDLE CANYON


swarm compose docker STORIDGE™ Automated PYTORCH

portainer.io kafka A distributed streaming platform kubernetes TF K

hadoop influxdb OMQ APACHE mxnet

Solr mongoDB.

Nvidia T4
BCOM/MLNX SmartNICs
9x 7.68TB ScaleFlux CSDs



Clondike 2020

**Real-time AI/ML Streaming Workloads with
Computational Storage**

Computational Storage Proliferation

Expanding Apps




Transparent Benefit
Ease of Use

Standardization



Identification
Activating Computation

Pipeline Compute



Application + Infrastructure
Multiply benefit

The Rise of Computational Storage



Unify Fast & Big Data

New paradigm to scale performance with capacity

Growing ecosystem & shipping in volume



Deploying Computational Storage at Scale

Thank You!

Come visit us at Booth #113
www.scaleflux.com

