XDF XILINX DEVELOPER FORUM

Computational Storage: Acceleration Through Intelligence & Agility

Presented By

ScaleFlux™

Thad Omura
EVP Marketing
December, 2018

XILINX

ScaleFlux™

What's the Big Deal?

ScaleFlux™

Alibaba Cloud

ScaleFlux™

70%

of Alibaba Cloud's clients encounter:

High Cost

- Exorbitant licensing fees
- Expensive migrations

Data Islanding

- Can't store tons of data in the same place
- Difficult for different applications to share

Data Explosion

- Sheer volume of data taxing to process

ScaleFlux™

MARKET PROBLEM

High Cost

Alibaba Cloud OPPORTUNITY

Cloud Scale Economics

Unified Storage

Data Islanding

Real-time Analytics

ScaleFlux SOLUTION

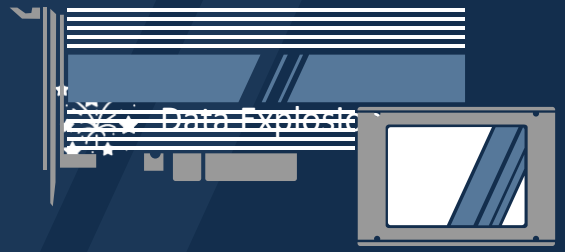

Computational Storage

ScaleFlux

Computational Storage Provides the Solution

Alibaba Cloud

POLARDB

HTAP: Hybrid Transactional-Analytical Database

Real-time analytical processing from transactional data

- Intelligent data management
- **Hardware accelerated, parallelized database compute at data**
- Programmable FPGA hardware to evolve functions with customer demands

ScaleFlux™

10X Transactional-Analytical Processing, **Half** the Flash Capacity

GM at Alibaba Cloud Database

"By bringing compute to the data, ScaleFlux is **transforming** the way we are architecting our **Flash storage infrastructure**.

We're looking to fully utilize the values of Computational Storage in order to cost-effectively scale **real-time analytics** across exploding transactional data sets, all the while delivering the most responsive, cloud-native user experience."

How?

Cohesive Application to Storage Acceleration

Open Channel Flash Management

Computation Acceleration

Solution Agility

ScaleFlux™

Open Channel (Host) Flash Management

Data Placement

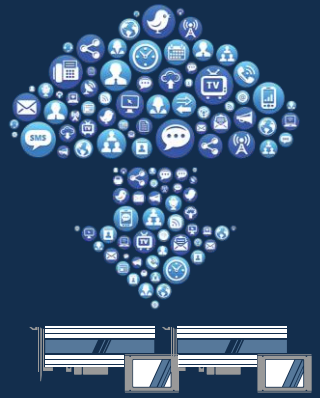

✓Controllable
✓ Application Awareness

Global View

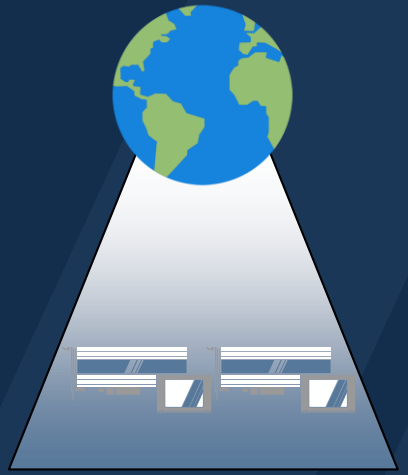

✓ Reduce Overprovision
✓ Minimize Write Amplification

Multi-Tenant

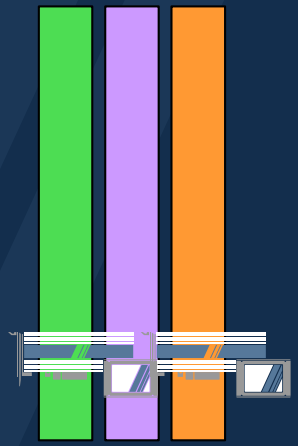

✓ Isolation
✓ Consistent

ScaleFlux™

In-System FPGA (HW) & SW Update for NEW Features

ATOMIC WRITES

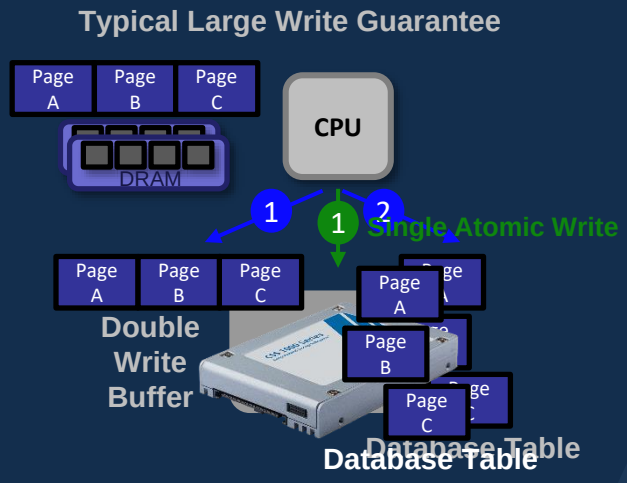

Guarantee larger write units are persisted to memory

2X Flash Endurance!

MULTI-STREAMS

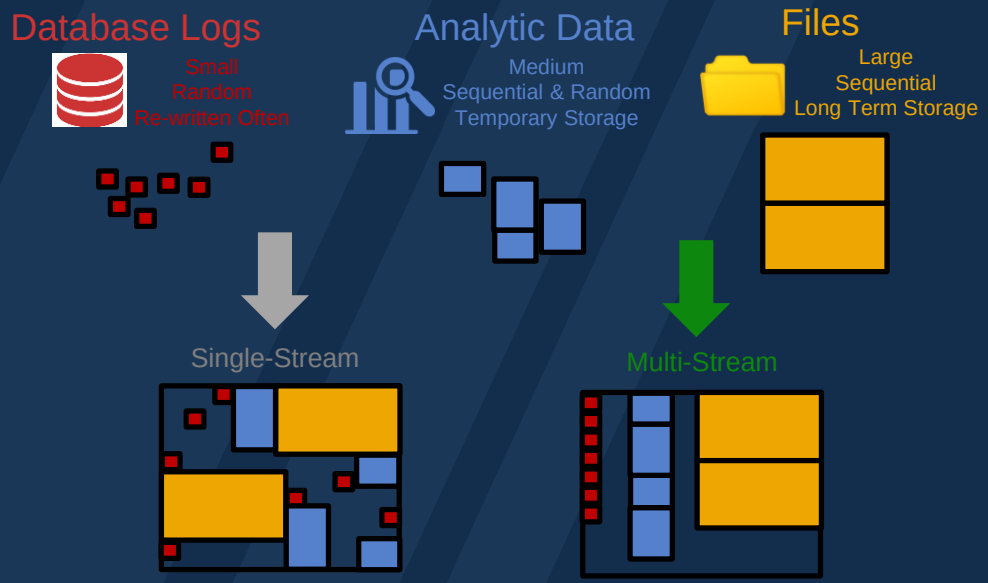

Efficient management of different data types

Performance & Endurance

ScaleFlux™

Production Workload Optimization: Customer Case Study

Orange is ScaleFlux

4 Week Turnaround to tune and test optimized Flash I/O
Minimize both 1ms and 8ms (tail latencies)

Flexible hardware and software platform for application tuning

Significant impact on latency sensitive applications

Customer now using ScaleFlux for lowest latency Flash NoSQL database solution

ScaleFlux™

More Values from Open Channel Model

Fast Transition to Latest 3D NAND TLC to QLC+

Simple Storage Class Memory Integration

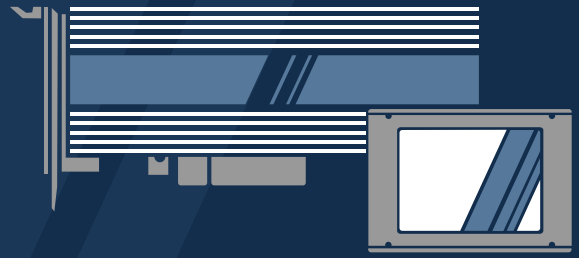

Next: Industry standardization

Cohesive Application to Storage Acceleration

Open Channel
Flash Management

Computation
Acceleration

Solution
Agility

ScaleFlux™

Acceleration

Performance & Scalability

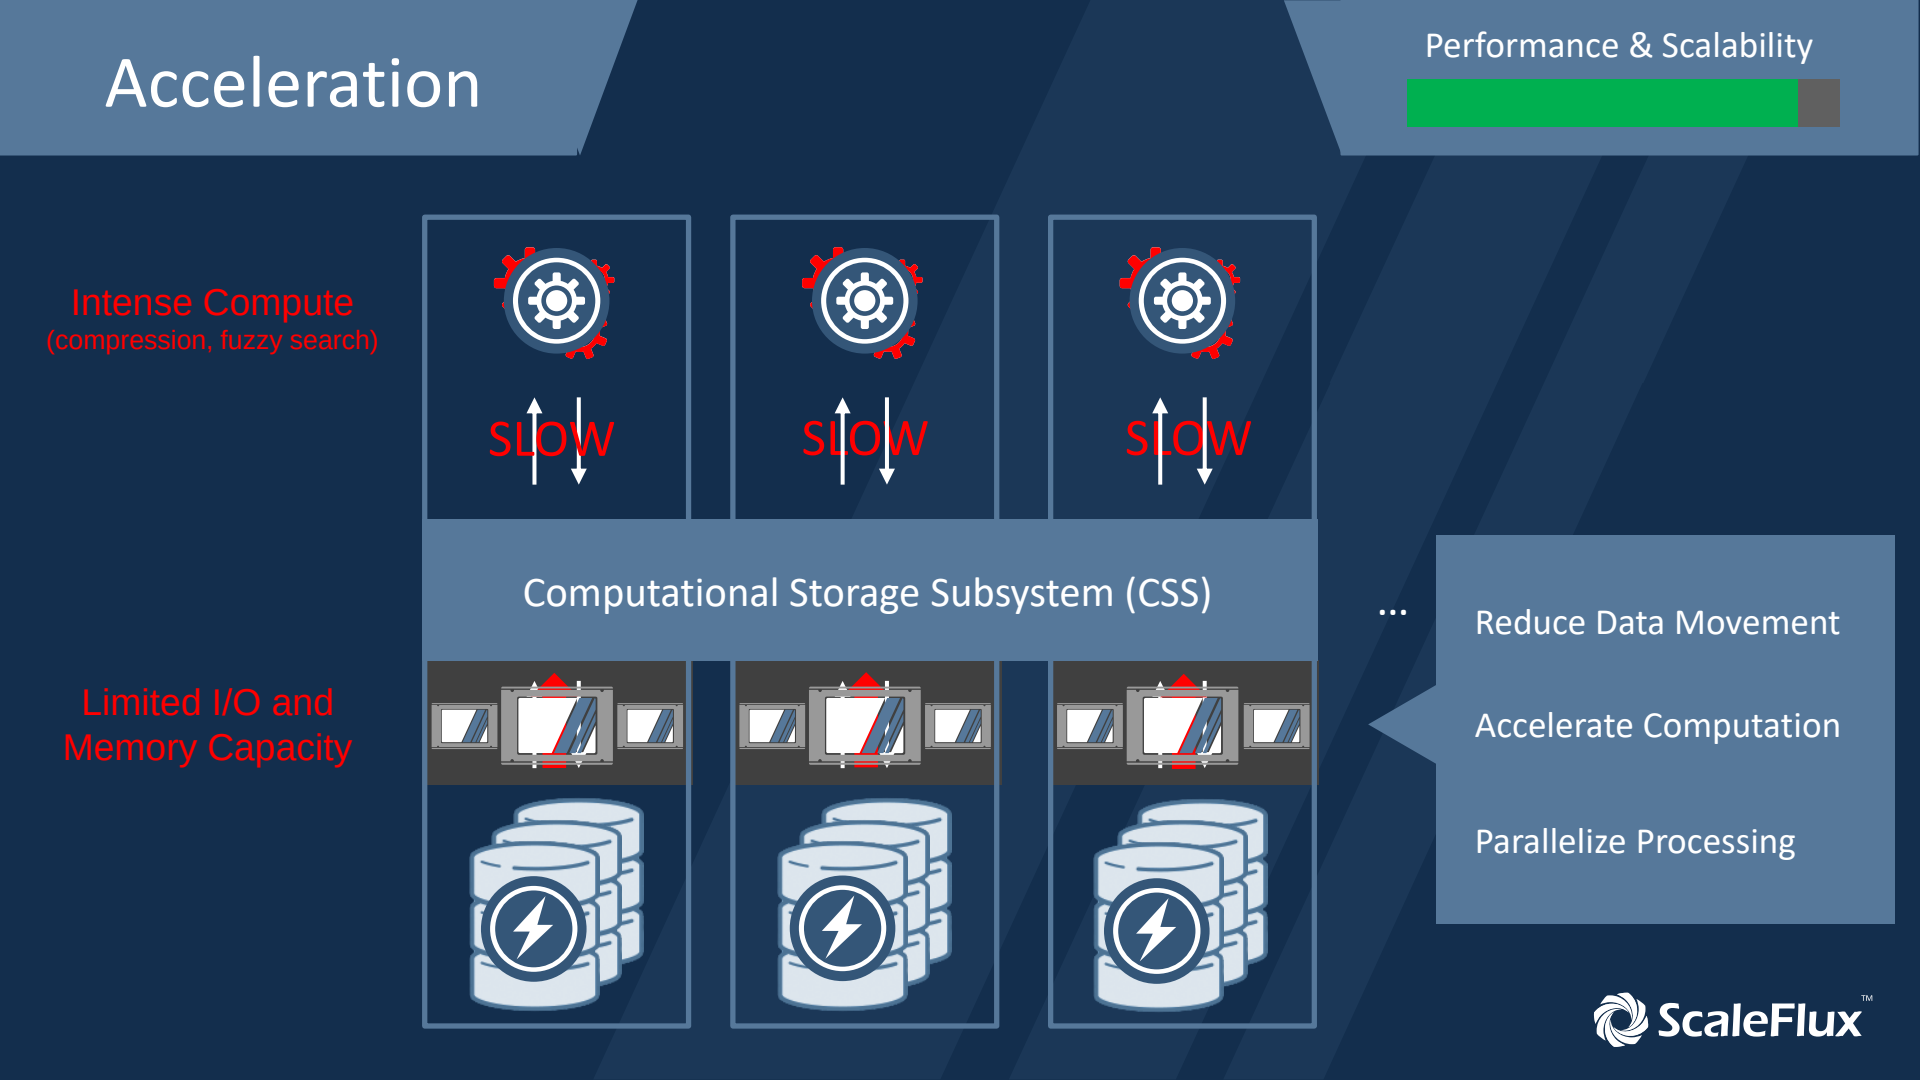

ScaleFlux™

Tradeoffs and Design Considerations

Compute Functions

- Data intensive, fixed function
- 5-100x speed up vs. CPU

ScaleFlux™

Parallelizing Computational Storage

GZIP Compression

(CPU zlib vs. ScaleFlux css_zlib, corpus.cantebury E5-2667v4)

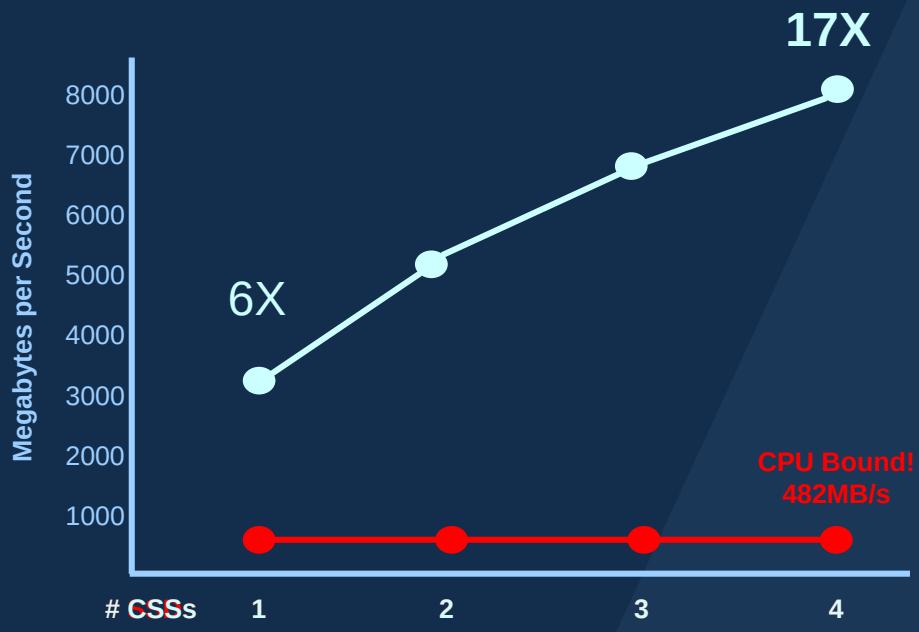

Fuzzy Search

(POC Unindexed Text Data, Edit Distance = 8, E5-2637v3)

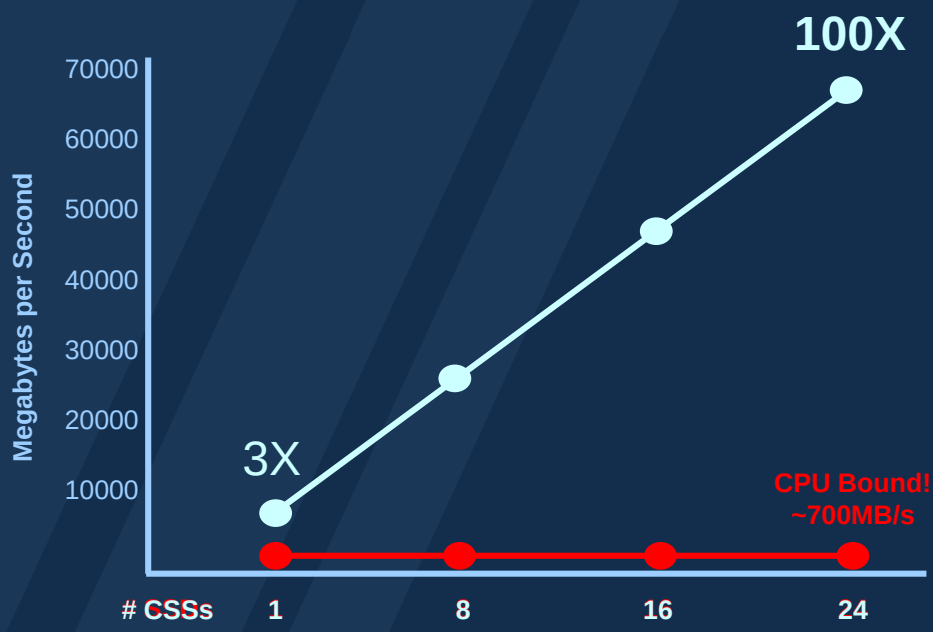

ScaleFlux™

Identifying the Right Workloads

INFRASTRUCTURE

STORAGE
Compression (GZIP)
Erasure Coding (RS)
Security (AES)
Authentication (SHA)
Error Checking (CRC)

PLATFORM

DATABASE, ANALYTICS
KV-Store
Transactional-Analytical
SQL Processing
Big Data Analytics

APPLICATION

AI, Genomics, CDN, Search
Media Scaling & Transcoding
Neural Networks
Fuzzy Search
Filtering, Matching

ScaleFlux™

Cohesive Application to Storage Acceleration

Open Channel
Flash Management

Computation
Acceleration

Solution
Agility

ScaleFlux™

Agility is Important

ENGAGEMENT

Demand to POC < 6 months

+

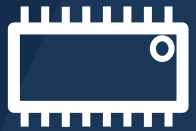

FLASH LIFECYCLES

Reducing to 12 months

+

HW AGILITY

Update after deployment

ScaleFlux™

Solution Agility Across Whole Stack

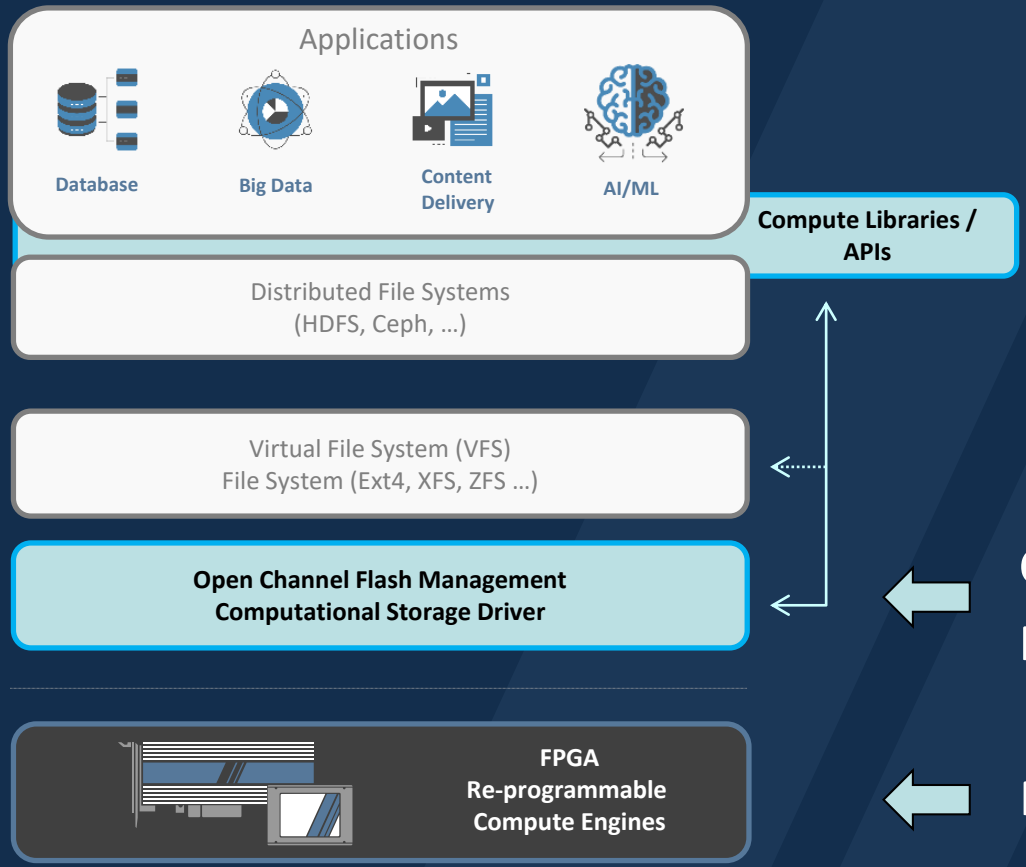

Simple Interface
Easy Integration

Controllable Data Placement
Performance/QoS (latency)

Programmable HW engines

Fast TTM for Turnkey Apps

2.7x
Transactions per Second

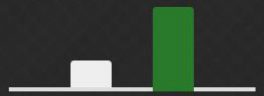

ACT

AEROSPIKE

2.6x
Operations per Second

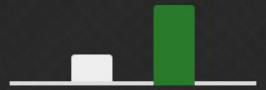

YCSB Load

APACHE HBASE

1.6x
Jobs Completed

Teragen & Terasort (vs HDD)

hadoop MapReduce

Available Through:

DELL EMC

inspur

packet

ScaleFlux™

Turnkey Compute & Storage Acceleration

hadoop 3.1 w/ EC (6+3)

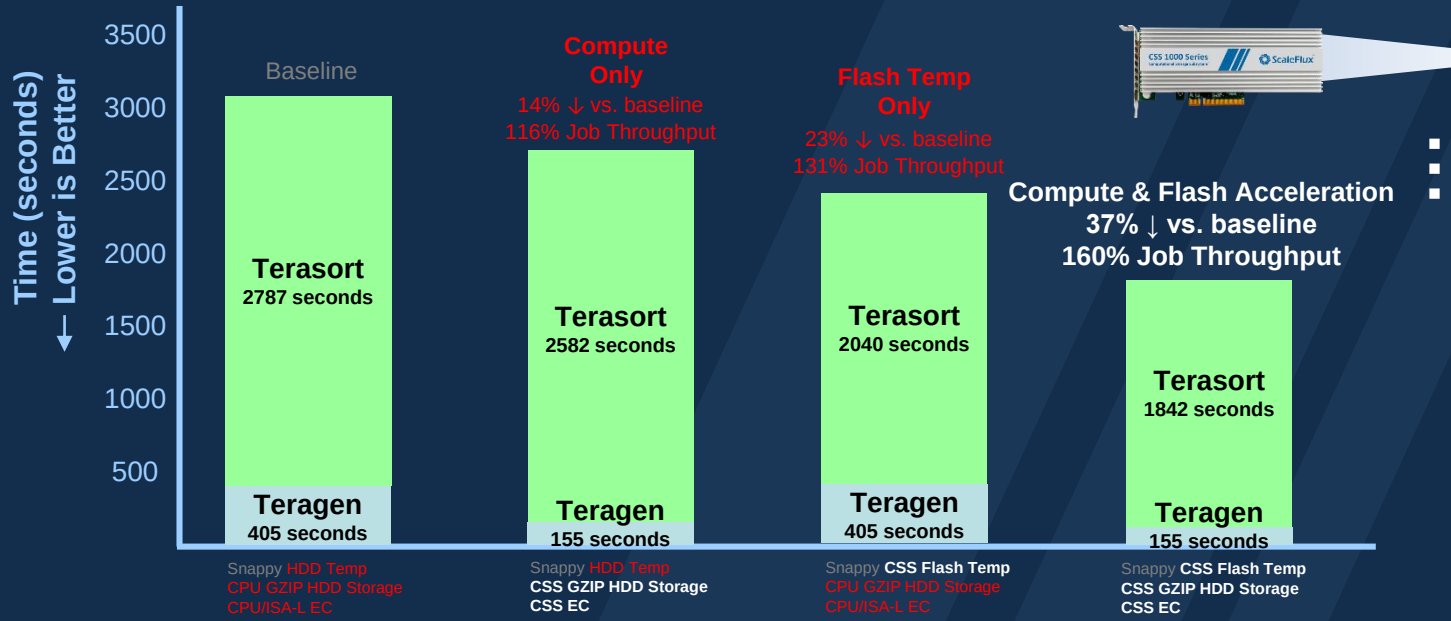

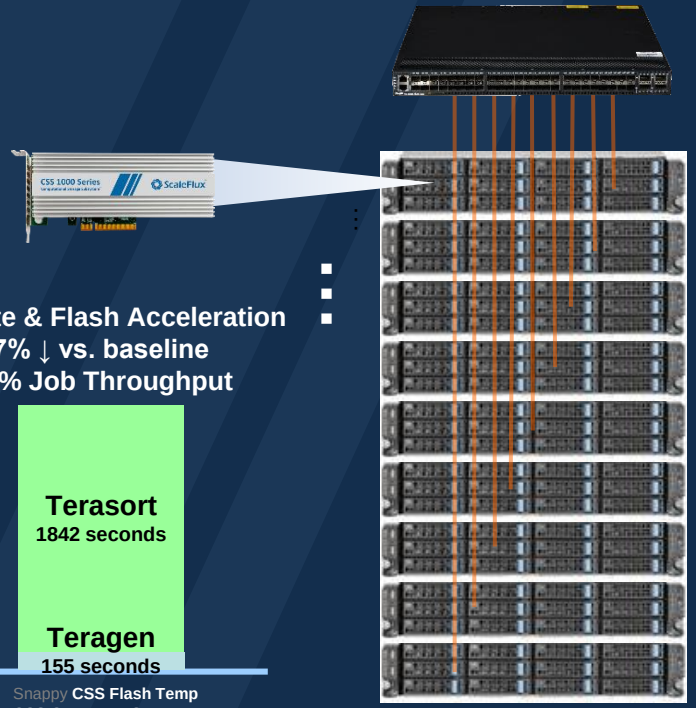

Datanode Config:
Dual E5-2640v3, 128GB DRAM, 12*6TB SAS HDD

24 Mapper/Reducers per Datanode *9 = 216 total
Better performance on CSS reported with lower Mapper/Reducers possible

ScaleFlux™

Delighted Customers

ScaleFlux™

"…delivering fantastic **OPERATIONS PER SECOND** for our latest NoSQL database…"

"…**INSTANTLY** saw how this can help us **COST-EFFECTIVELY** scale our infrastructure …"

"…accelerating **MULTIPLE, BUSINESS-CRITICAL APPLICATIONS** for us…"

ScaleFlux™

ScaleFlux™

The pioneer in deploying Computational Storage at scale

- HQ in San Jose, Offices in China
- Shipping Computational Storage worldwide

Thank You!

Come visit us in the Exhibits!

www.scaleflux.com

ScaleFlux™


XDF
XILINX
DEVELOPER
FORUM