



Flash Memory Summit

# Computational Storage: Acceleration Through Intelligence & Agility

Dr. Hao Zhong  
CEO & Co-Founder, ScaleFlux™

Flash Memory Summit 2018  
Santa Clara, CA





# What's the Big Deal?







## High Cost

- Exorbitant licensing fees
- Expensive migrations

**70%**  
Data Islanding  
of Alibaba Cloud's clients encounter:

- 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

MARKET PROBLEM



High Cost



OPPORTUNITY

Cloud Scale Economics




Data Islanding  
Unified Storage

Real-time Analytics



SOLUTION



Computational Storage

**HTAP: HYBRID**  
**Transactional / Analytical Processing**

**No lag for analytics**  
**Low cost, on storage**



**Alibaba Cloud**

**POLARDB**


**Hardware Acceleration**  
**Compute at data**  
**Embracing new hardware**

**High**  
**Performance &**  
**Scalability**

**High**  
**Compatibility**

**High Reliability**  
**& Availability**

# Computational Storage Provides the Solution



Real-time analytical processing from transactional data

- Intelligent data management
- Parallelize Compute at Data
- Programmable hardware



## 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

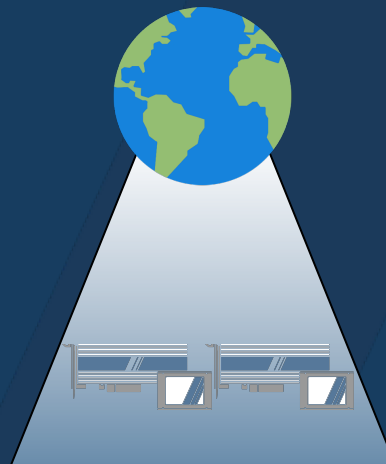
# Open Channel (Host) Flash Management

## Data Placement




- ✓ Controllable
- ✓ Application Awareness

## Global View



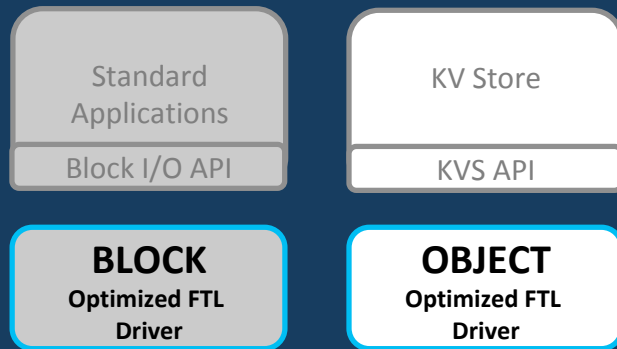
- ✓ Reduce Overprovision
- ✓ Minimize Write Amplification

## Multi-Tenant



- ✓ Isolation
- ✓ Consistent

## Value Example: KV Store



KV Tailored FTL

Compatible Hardware

Fast Integration

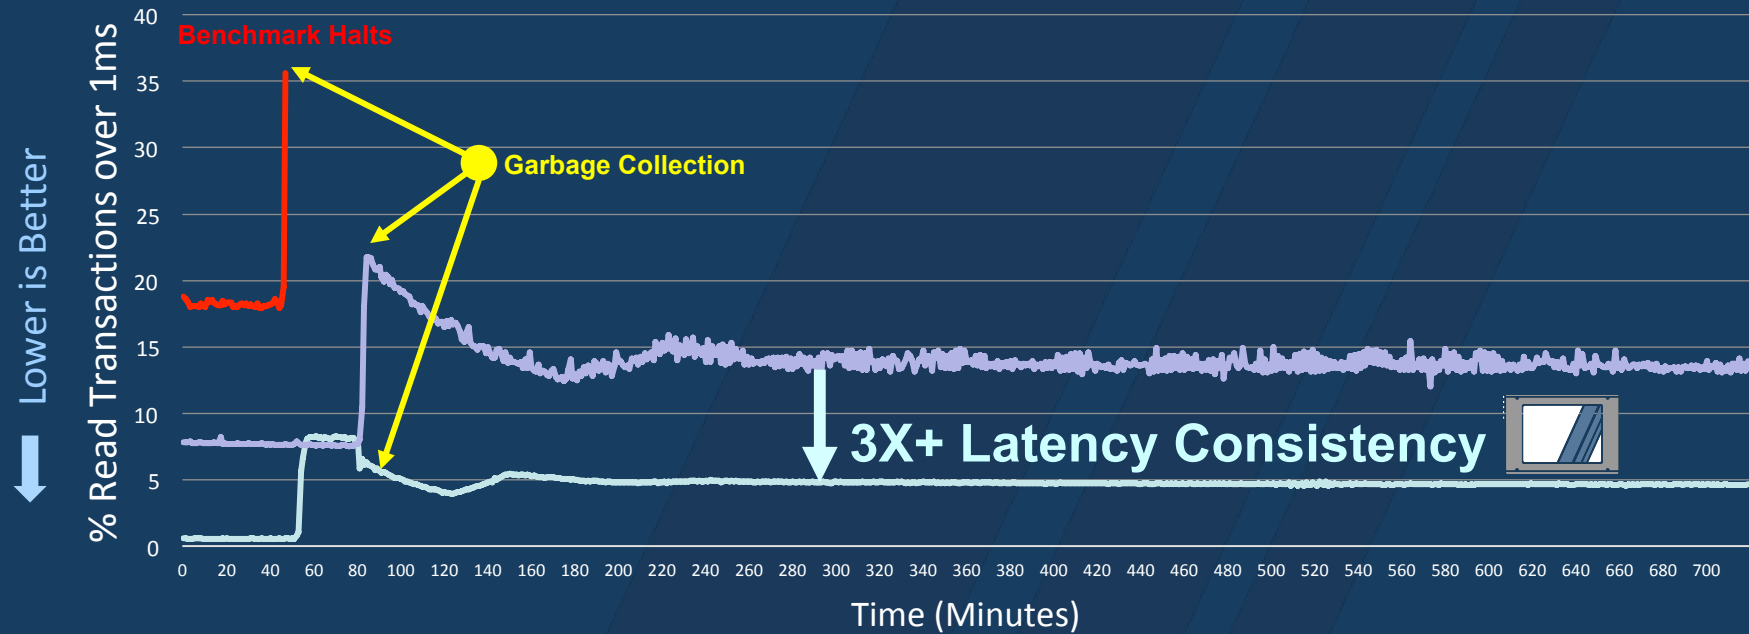
# Optimize Application QoS, not just at SSD Level

## Aerospike Certification Tool (ACT) v4

300K Transactions per second (100K/s 128KB writes with 200K/s 1536B reads)

Dual Intel 6126 CPU, 256GB DRAM

AEROSPIKE



Vendor A 3.2TB U.2

Vendor B 3.2TB U.2

ScaleFlux CSS 3.2TB U.2




## Other Open Channel Management Values

Easily Tunable

3D NAND TLC to QLC+

Simple NVM integration

**Next: Industry standardization**



# Cohesive Application to Storage Acceleration



Open Channel  
Flash Management

Computation  
Acceleration

Solution  
Agility



# Acceleration

Performance & Scalability



**Intense Compute**  
(compression, fuzzy search)



SLOW



SLOW



SLOW

Computational Storage Subsystem (CSS)

**Limited I/O and  
Memory Capacity**



- ...
- Reduce Data Movement
- Accelerate Computation
- Parallelize Processing

# Tradeoffs and Design Consideration



## Compute Functions

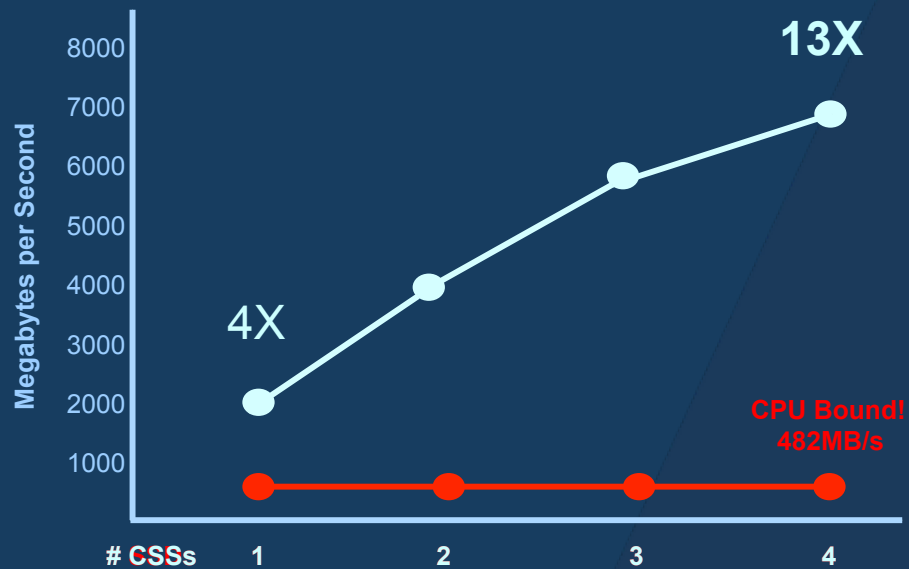
- Data intensive, fixed function
- 5-100x speed up vs. CPU



# 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)



# Identify 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

# Cohesive Application to Storage Acceleration



Open Channel  
Flash Management

Computation  
Acceleration

Solution  
Agility

# Agility is Important



## ENGAGEMENT

Demand to POC < 6 months



## FLASH LIFECYCLES


Reducing to 12 months



## HW AGILITY


Update after deployment

# Solution Agility Across Whole Stack



Compute Libraries / APIs

Simple Interface  
Easy Integration



Controllable Data Placement  
Performance/QoS (latency)

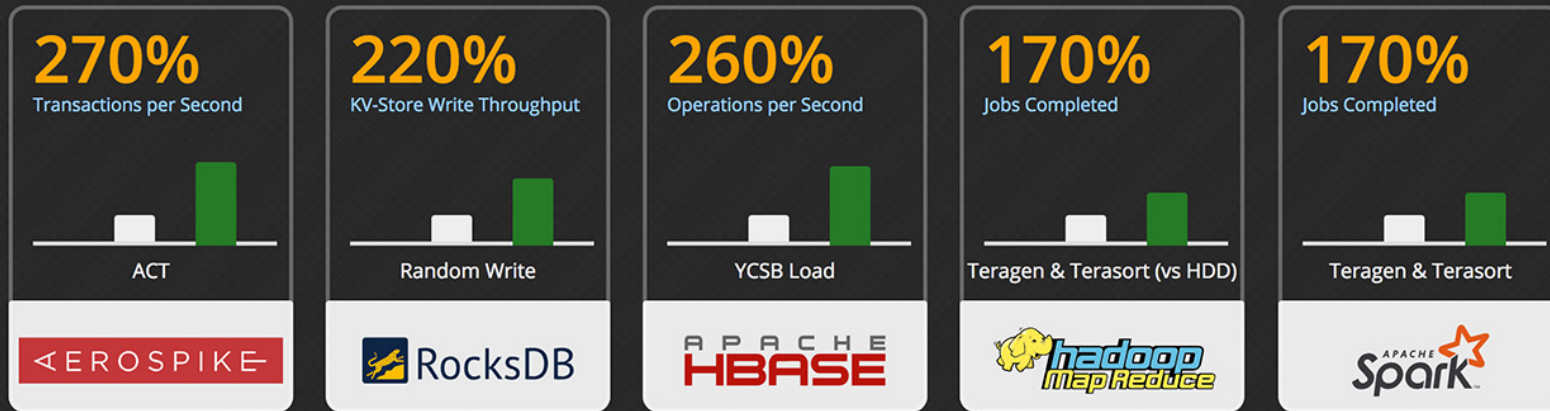
PCIe

AIC or U.2



Programmable HW engines

# Fast TTM for Turnkey Apps



vs. NVMe

Available Through:

DELL EMC

inspur

packet

ScaleFlux™



# Delighted Customers

A blue circle containing the word "FAST" in white capital letters. The background of the circle features faint, semi-transparent logos of various companies, including Playtika.

**FAST**

“...delivering fantastic **OPERATIONS PER SECOND** for our latest NoSQL database...”

A blue circle containing the word "EASY" in white capital letters. The background of the circle features faint, semi-transparent logos of various companies, including Google.

**EASY**

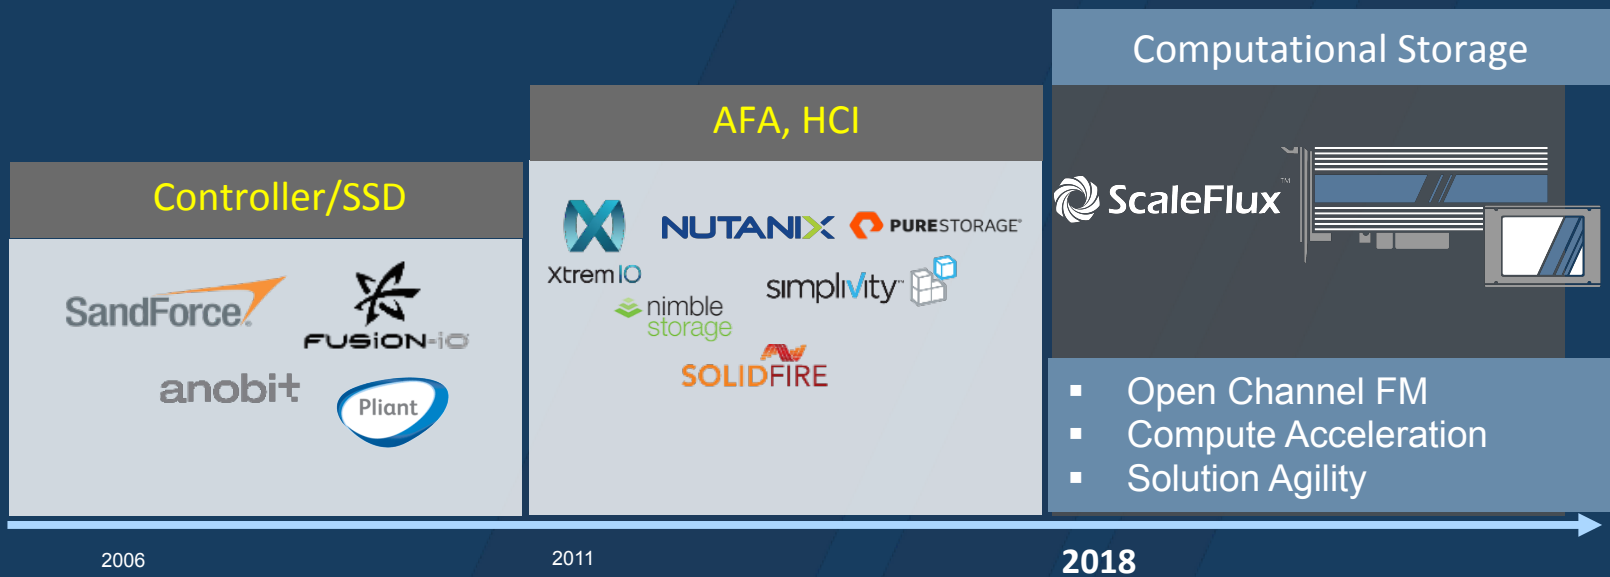
“...**INSTANTLY** saw how this can help us **COST-EFFECTIVELY** scale our infrastructure ...”

A blue circle containing the word "AGILE" in white capital letters. The background of the circle features faint, semi-transparent logos of various companies, including Flipkart and PhonePe.

**AGILE**

“...accelerating **MULTIPLE, BUSINESS-CRITICAL APPLICATIONS** for us...”

# Ride the Computational Storage Wave!





**The pioneer in deploying Computational Storage at scale**

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

**Thank You!**

**Come visit us at Booth #113**

**[www.scaleflux.com](http://www.scaleflux.com)**



BACKUP



Flash Memory Summit

Thank You!

Come visit us at Booth #113  
[www.scaleflux.com](http://www.scaleflux.com)

