PhonePe

# The Rise of Computational Storage

May 2019

# Vision

**India’s largest transaction platform, built on payments**

# Strategy – Follow the money journey


| | | | |
|---|---|---|---|
| **KEY USE CASES** | P2P Money transfer | Utility Bills & Mobile Recharge | Banking as a Platform |
| | Gifting | Offline Merchants | Gold |
| | Remittances | Online Merchants | Insurance |
| | Group Payments | In-App: AppStore | Mutual Funds |
| | | | Loans |
| | | | Credit |

# 2nd-largest payments company within 2 years of launch

**100+ M**
Registered Users

**40 M**
Monthly Active Users

**25 M**
Monthly Active Customers

**1 M**
Merchant Network

**180 M**
Monthly Transactions

**48 Billion**
Annualized Transaction Processing Value

**Flipkart accounts for <0.5% of our transaction volume**

# Fastest growing merchant network

- Live on 100 top online merchants
- Live at 1+ Million merchants across India

# Innovation examples (1/4) – Open ecosystem approach

Payment container

- All payment instruments supported
  - UPI
  - Debit Card
  - Credit Card
  - Wallet

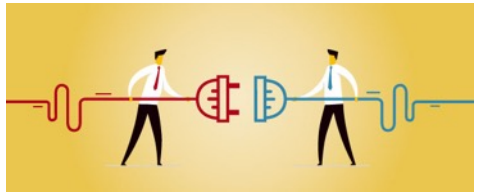

Wallet interoperability

- PhonePe wallet integrated with other wallets

Jio freecharge airtel money


Use cases

- Partner with leading players to offer consumer use cases

OLA redBus goibibo.com MMRC

# Innovation examples (2/4) – In app platform

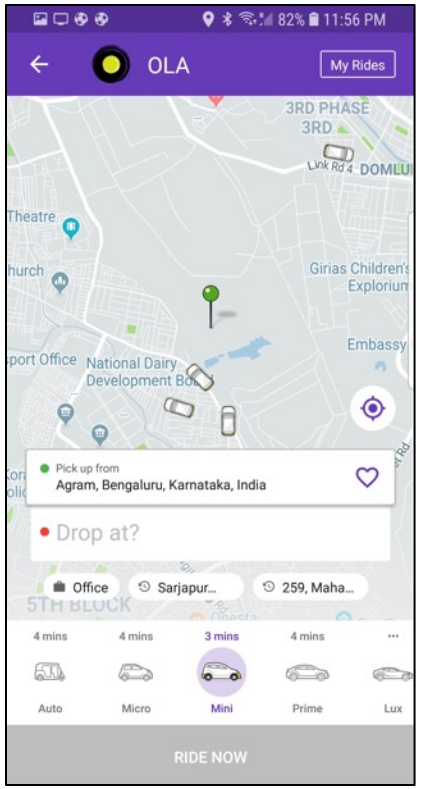


- **Autopay**: Initiate payments without second factor authentication
- Works with both Credit and Debit Cards

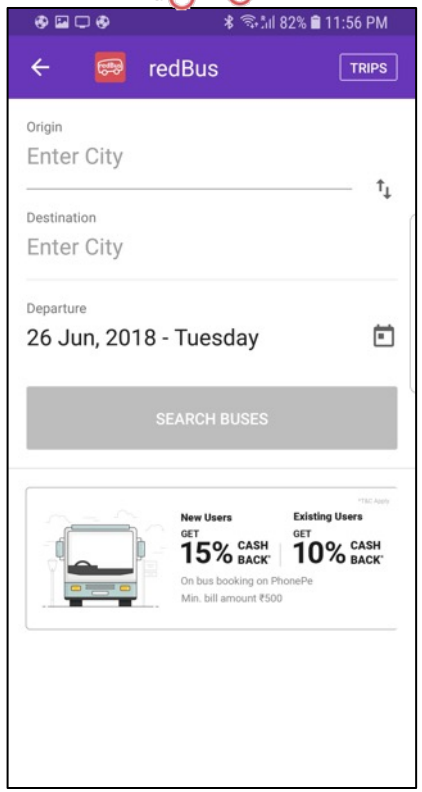


- **10+%** of the consumer driven Redbus bookings happening through PhonePe

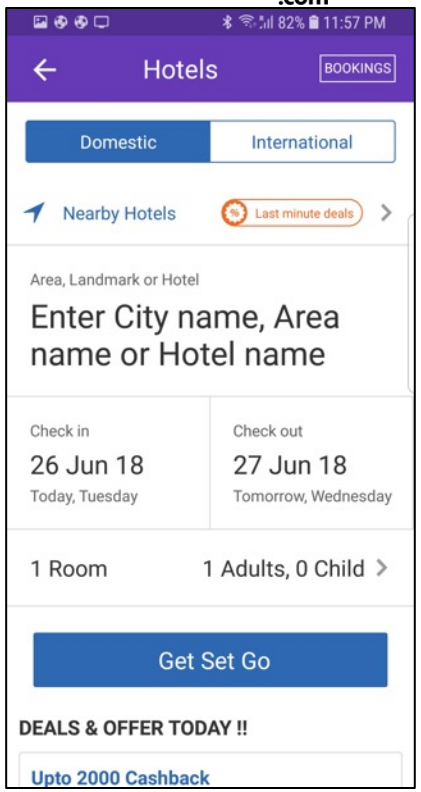


- Make hotel bookings within the PhonePe app
- Users can **use goCash** on the PhonePe platform

# Innovation examples (3/4) - Made for India POS device


- **Launched in Nov 2017**

  First of its kind innovation basis merchant insights

- **World's cheapest POS ($10)**

  All payment methods supported, no data requirement from merchant side

- **First 5,000 POS distributed in 3 weeks**

  Initial launch of devices has been well received – plan to deploy 1 Mil devices across top 60 cities

- **Won NPCI RFP for proximity**

# Innovation examples (4/4) – Gold marketplace

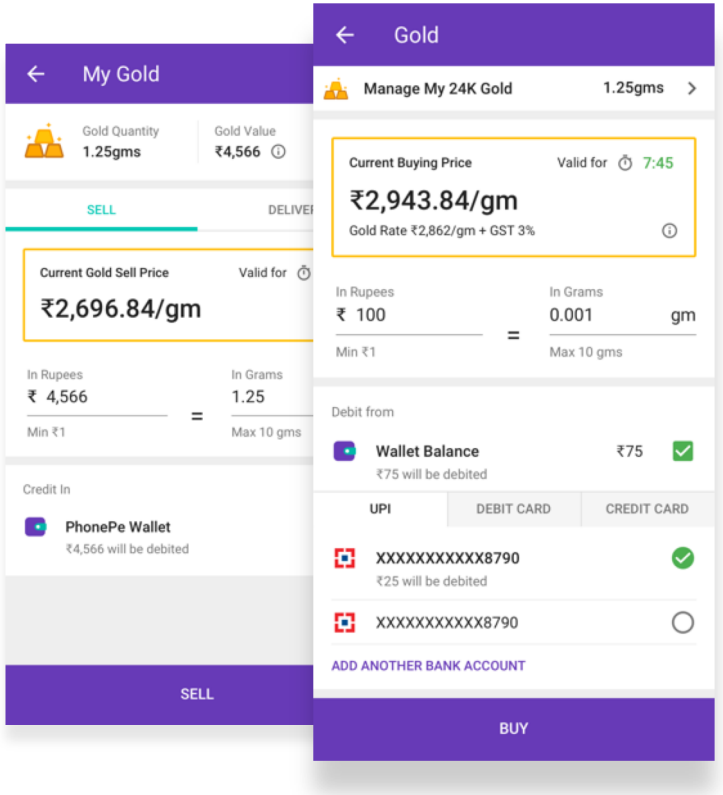


- First to offer a marketplace model
- Market leader with 60% share
- More than 1 Tonne worth of gold transactions

# Data Center Trends

## STORAGE

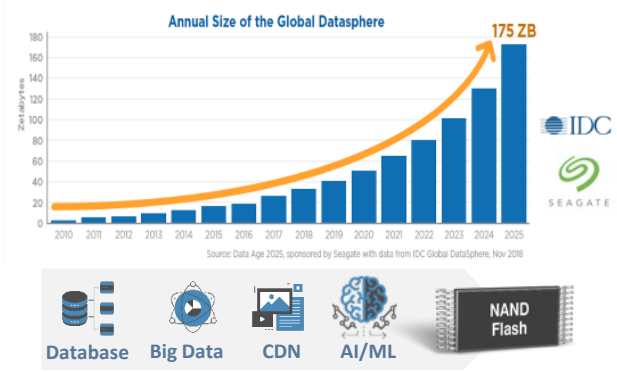


**Massive, Fast Data**

## NETWORKING

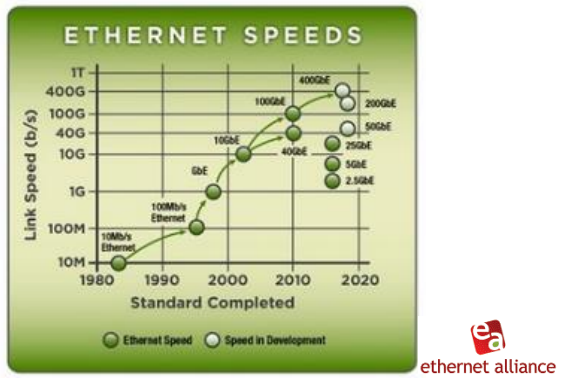


10 → 40 → 100Gb+

## COMPUTE

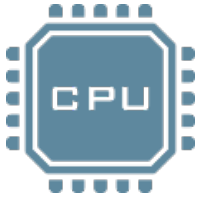


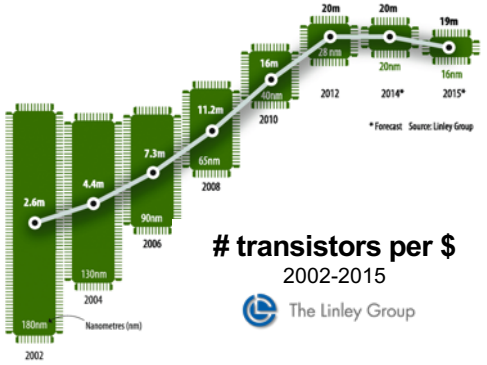


End of Moore's Law

## Domain Specific Compute Solutions



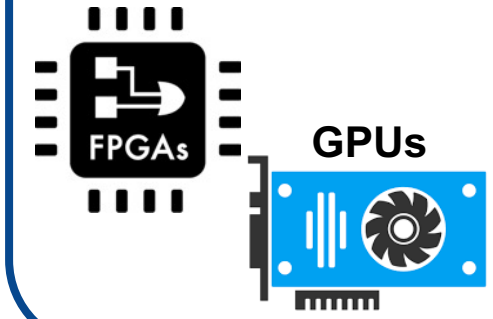

# What Is Computational Storage?

## Traditional Infrastructure

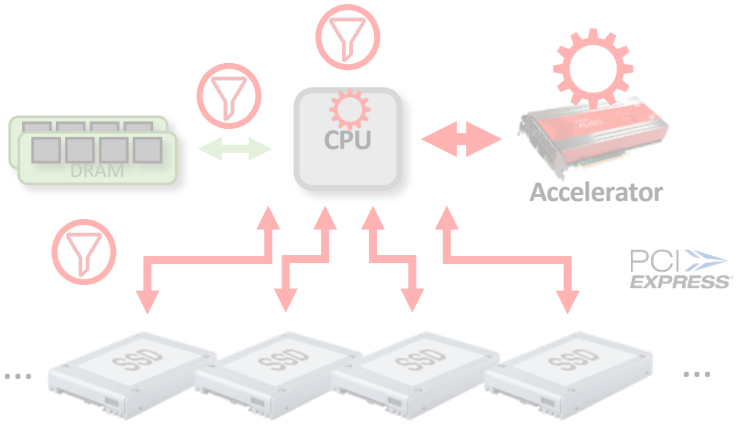


PCIe increases storage I/O 6X+ vs. SAS/SATA

All data moves to host for processing

Host CPU / memory bottlenecks

No compute parallelism

Data-driven application performance challenges

## Computational Storage

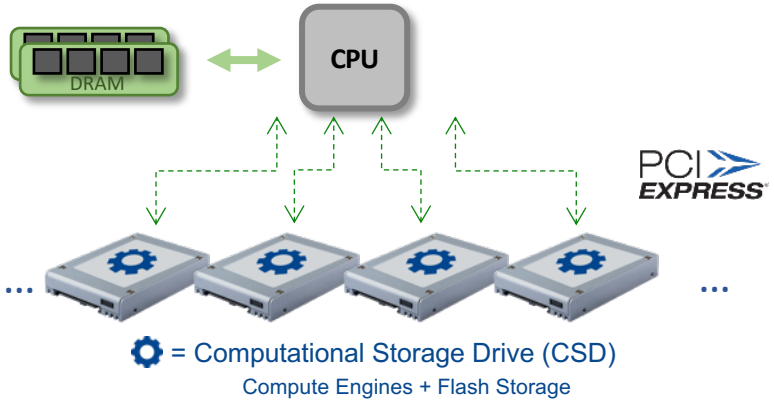


Balanced compute resource & storage I/O

Minimize data movement

Multiple FPGAs, easily plug-in via storage

Maximum compute parallelism

Fits into existing, standard PCIe Storage slots

Standardization (in process) for easy app integration

**A New Paradigm to Scale Compute Resources with Storage Capacity**

# Parallelizing Workloads With Computational Storage

## GZIP Compression

(Raw Performance CPU zlib vs. CS zlib, corpus.cantebury files, E5-2667v4)
GZIP uses Huffman Encoding which is inefficient on x86

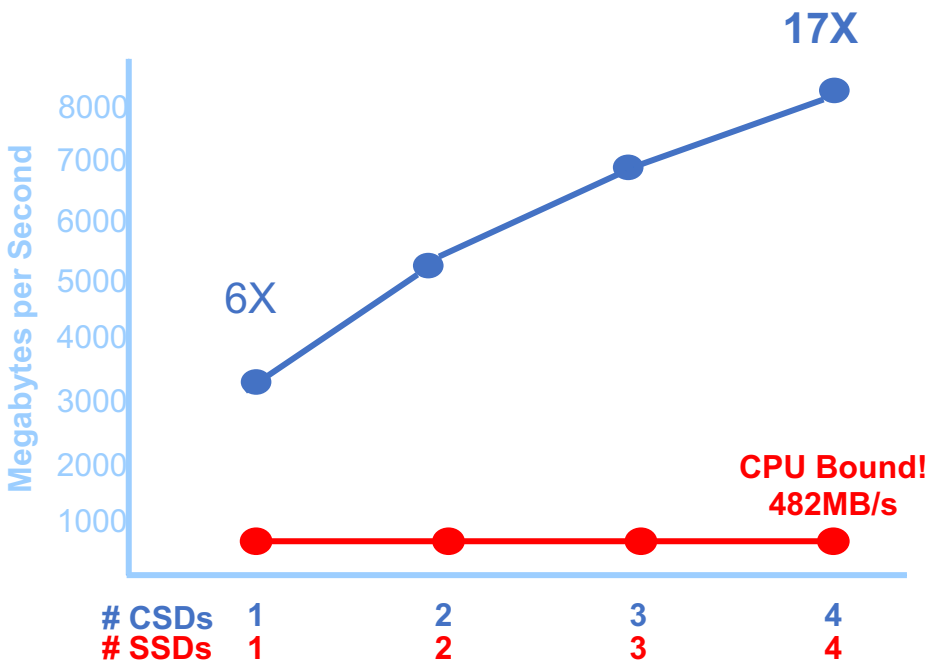


## Fuzzy Text Search (POC)

(CPU vs. CSD agrep, Unindexed Text Data, Edit Distance = 8, E5-2637v3)
agrep = approximate grep which is inefficient on x86 with high edit distance

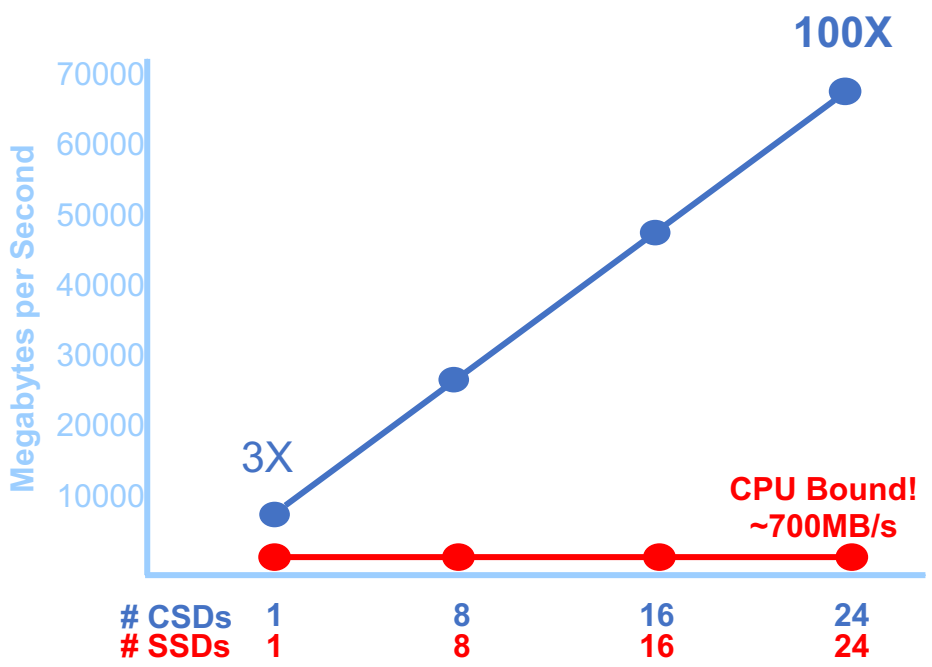


Cohesively Unlock Compute & Storage I/O Bottlenecks

# Computational Storage Services

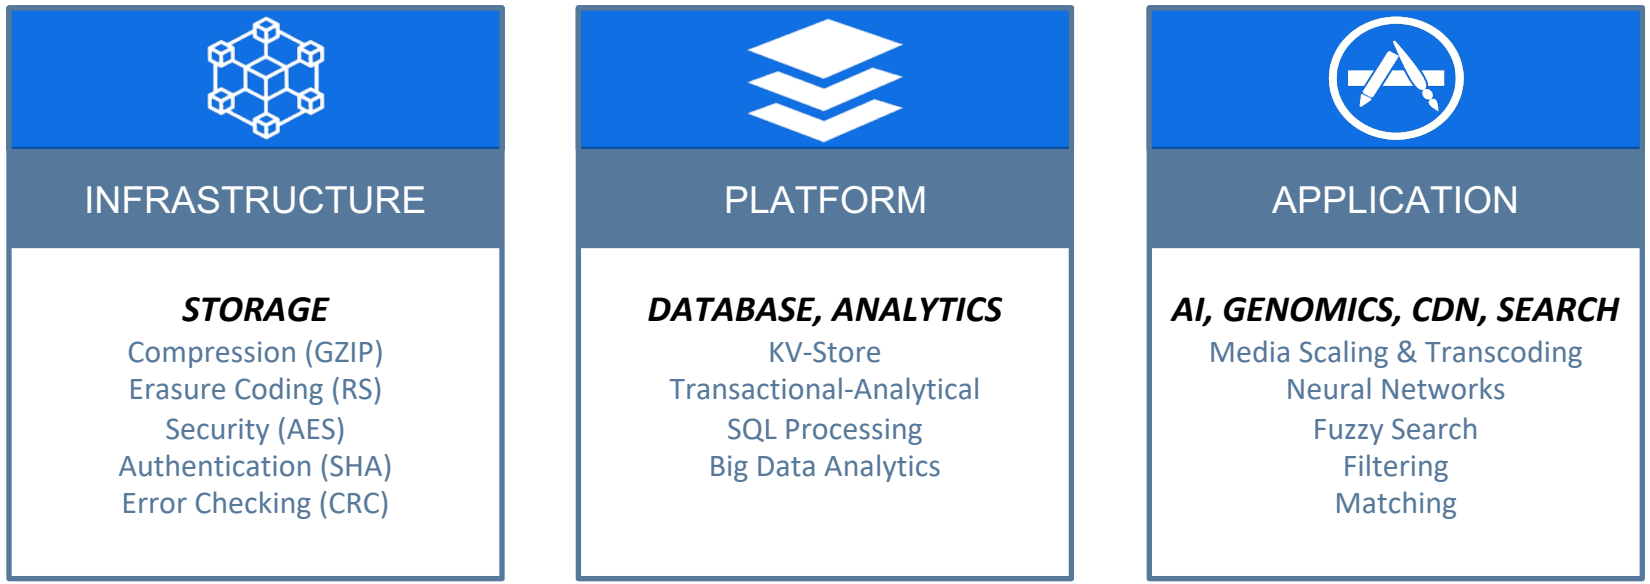


**Ideal Targets:**

- Fixed algorithms
- Bit-wise data comparison/manipulation (slow on x86)
- Compute required across the entire storage data set

***Evolutionary*** integration through Flash storage (standard hardware)

***Revolutionary*** impact through parallelization of workloads

# Computational Storage Use Models

## Accelerator + Flash

Data Starts in Host DRAM

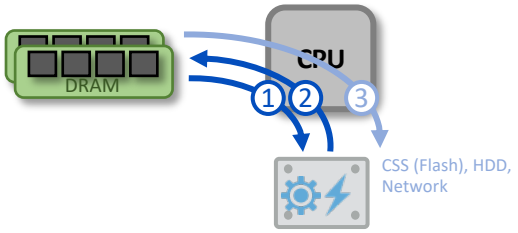


- Examples: GZIP, EC (RS), AES, SHA, transcoding…

## Data Path Processing

Data Starts in Host DRAM (on Write) or CSD (on Read)

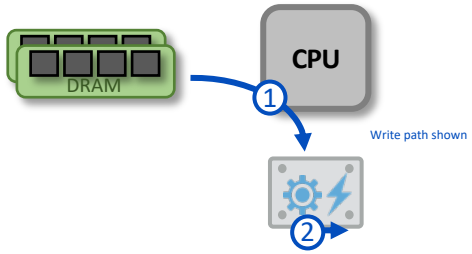


- Examples: in-line GZIP, AES, SHA, transcoding, AI inference
- Compute while data write to CSD or read from CSD

## In-Storage Processing

Data Starts in CSD

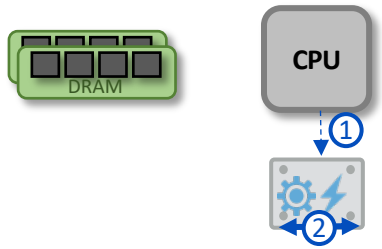


- Examples: Database queries, fuzzy search, pattern matching…
- Compute locally on each CSD, return only return results back to CPU
- Save massive data movement

**Common Computational Storage Benefits:**

PCIe slot consolidation (Compute + Flash)

Easily parallelize computation across multiple CSDs

Same hardware can be used regardless of model

Multiple models can be utilized simultaneously

# Solutions That Simply Run Faster on Computational Storage

AEROSPIKE

**200%**

**Latency Consistency**

**Host FM & Tunable Performance**

Customer Specific Workload vs. NVMe

MySQL

**200%**

**Queries Per Second**

**Atomic Write Support**

Sysbench OLTP write-only vs. NVMe

OpenZFS

**Up to 5X**

**GZIP Write Throughput**

**GZIP Compute Engines**

FIO CPU GZIP vs. CSS GZIP

hadoop MapReduce 3.0

**161%**

**Jobs Completed**

**GZIP & EC Compute Engines**

Teragen, Terasort vs. HDD only

*Applies to HDFS & Spark

APACHE HBASE

**260%**

**GZIP Write Throughput**

**GZIP Compute Engines**

YCSB Load Benchmark vs. NVMe

# Example Computational Storage: MyRocks Database Query

## TPC-H Database Query Benchmark Results

Push down DB queries (compute intensive) to Computational Storage through MyRocks

Up to 71% latency reduction in query time vs. using Host x86 for queries

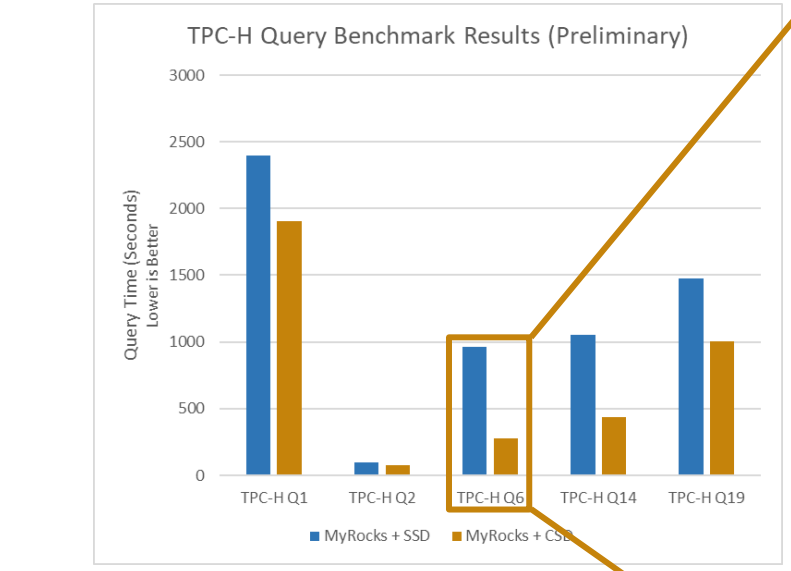


TPC-H Benchmark, 4K Block Size
CentOS v. 7.2.1511
100GB Database

Dual
E5 2620 V3
@ 2.2GHZ

128GB
DRAM

MySQL Server

MySQL v5.6, 1Thread

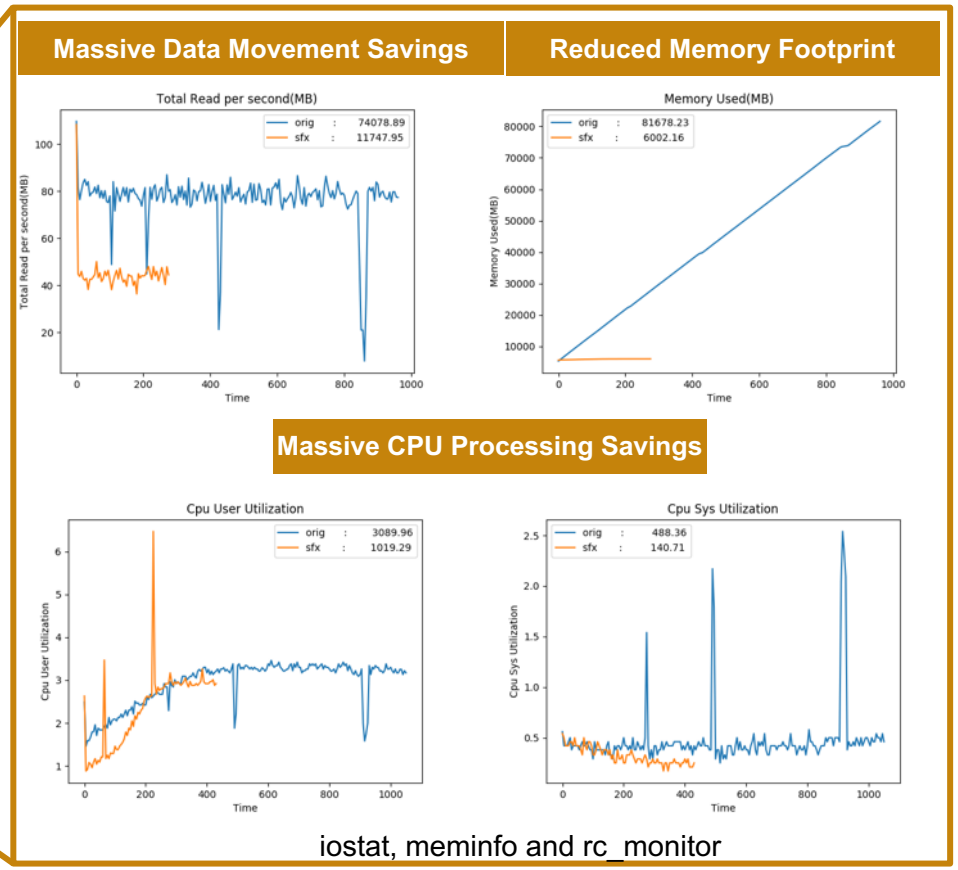

# Standardization of Computational Storage

**Definition of product types:**

- CSD = Computational Storage Drive
- CSP = Computational Storage Processor
- CSA = Computational Storage Array

**Management:** Discovery, Security

**Operation on data types:** LBA (Logical Block Address), KV (Key-Value), File, Object, Persistent Memory

SNIA
Storage Networking Industry Association

Founders: EIDETICOM, NGD systems, ScaleFlux

Initial Supporters: arm, Calypso Systems, GigaIO, inspur, Kalray, Lenovo, Marvell, Micron, NetApp, NetInt, Nyriad, SK hynix, Samsung, Toshiba, Western Digital, Xilinx

Broadcom, Dell EMC, Hitachi, Huawei, IBM, intel, Mellanox, Microsemi, NEC, Oracle, Supermicro, Seagate, vmware

** New since February


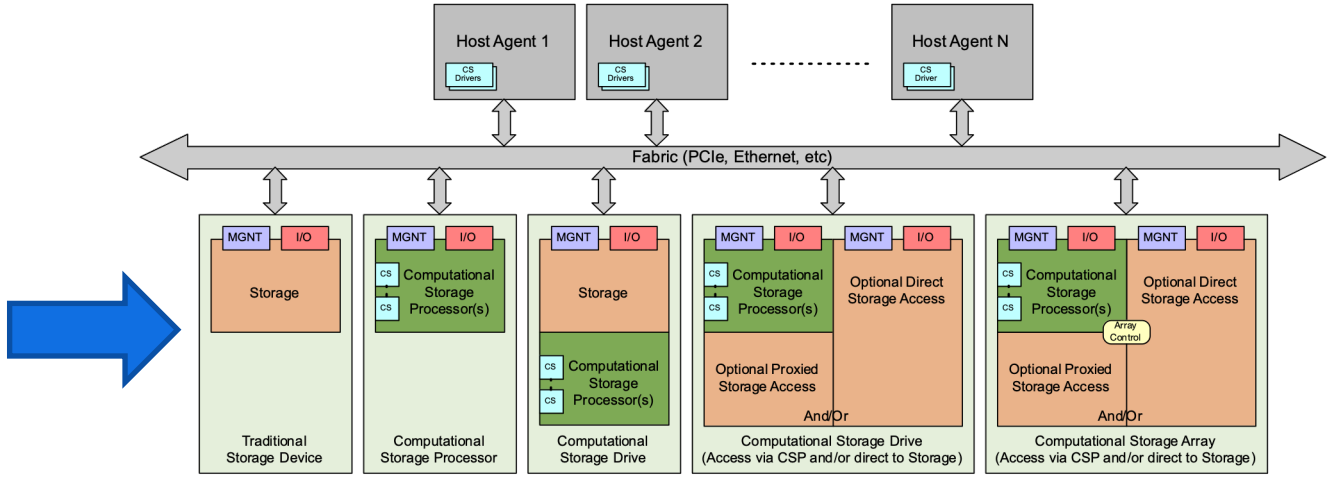


**Easily Consume Computational Storage From Multiple Vendors**

# Summary: Industry Embracing Computational Storage

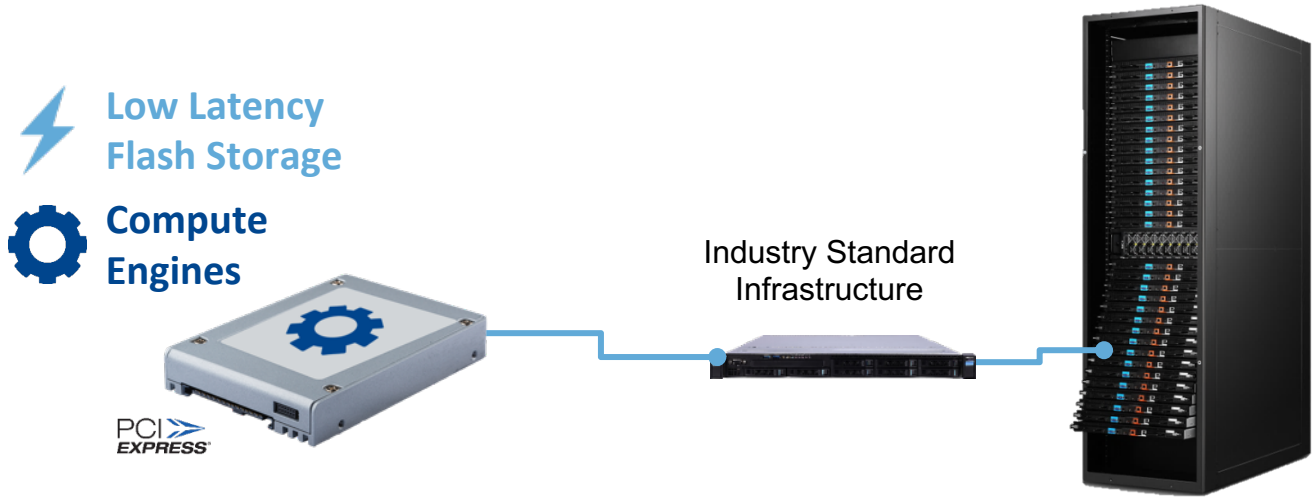


Computational Storage easily fits into existing PCIe SSD storage slots

Delivers immediate application level benefits with simple integration

Save expensive data movement and optimize server compute & memory resources

Industry-wide effort to standardize integration amongst multiple vendors

Bringing Compute to Data for Modern Data Driven Applications

# Thank You

Shameless plug: We are hiring, if you are passionate about solving large scale problems, love opensource and working with cool people, please reach out to us!

# Large Version of Pictures from Trends Slide from Info Only