Achronix
Data Acceleration

PLIOPS

ScaleFlux™

G2M
RESEARCH

Computational Storage vs Virtualized Computation/ Storage In the Datacenter

Multi-Vendor Webinar
Tuesday July 13, 2021

G2M Research Introduction and Ground Rules

Mike Heumann
Managing Partner, G2M Research

Webinar Agenda

9:00-9:05	Ground Rules and Webinar Topic Introduction (G2M Research)
9:06-9:30	Sponsoring Vendor presentations on topic (10 minute each)
9:31-9:36	Panel Discussion Question #1
9:37-9:37	Audience Survey #1
9:38-9:43	Panel Discussion Question #2
9:44-9:44	Audience Survey #2
9:45-9:50	Panel Discussion Question #3
9:51-9:58	Audience Q&A (8 minutes)
9:59-10:00	Wrap-Up

What is Computational Storage?

Storage Networking Industry Association (www.snia.org) defines computational storage as:

- Architectures that provide Computational Storage Functions coupled to storage, *offloading host processing or reducing data movement*.
- These architectures enable improvements in application performance and/or infrastructure efficiency through the integration of compute resources (outside of the traditional compute & memory architecture) either directly with storage or between the host and the storage. The goal of these architectures is to *enable parallel computation and/or to alleviate constraints on existing compute, memory, storage, and I/O*

The Disaggregation of Computation

- Computation is spreading beyond the CPU into a variety of new devices
  - Computational Storage Drives (CSDs)
  - Computational Storage Arrays (CSAs)
  - Computational Storage Processor (CSP)
- All share the same concept of offloading certain aspects of processing from the CPU to devices that exist in the storage domain
- Kind of like Apache Hadoop at a more granular level – move the computation to where the data is, instead of moving very large amounts of data to the CPU

Challenges to Computational Storage

- **Application Modification**: Whenever applications need to be modified to accommodate a technology, adoption times are significantly lengthened
- **Go-To-Market Channel Adoption**: Are computational storage devices available from standard IT sources (SIs, resellers, OEMs)
- **"Crossing The Chasm"**: Many IT buyers won't adopt a new technology until it has been adopted by "mainstream" companies
- **Standardization**: Do standards for the technology exist to help avoid vendor lock-in

Panelists

Achronix
Data Acceleration

Tom Spencer
Sr. Product Marketing Manager
www.achronix.com

PLIOPS

Tony Afshary
Sr. Director, Product Line Mgmt
www.pliops.com

ScaleFlux™

JB Baker
Vice President, Marketing
www.scaleflux.com

G2M RESEARCH

Mike Heumann
Principal Analyst
www.g2minc.com

Achronix®
Data Acceleration

Tom Spencer

Sr. Manager, Product Marketing

www.Achronix.com

G2M
Computational Storage

Tom Spencer – Senior Manager, Product Marketing – July 13, 2021

Achronix®
Data Acceleration

WW Data Explosion

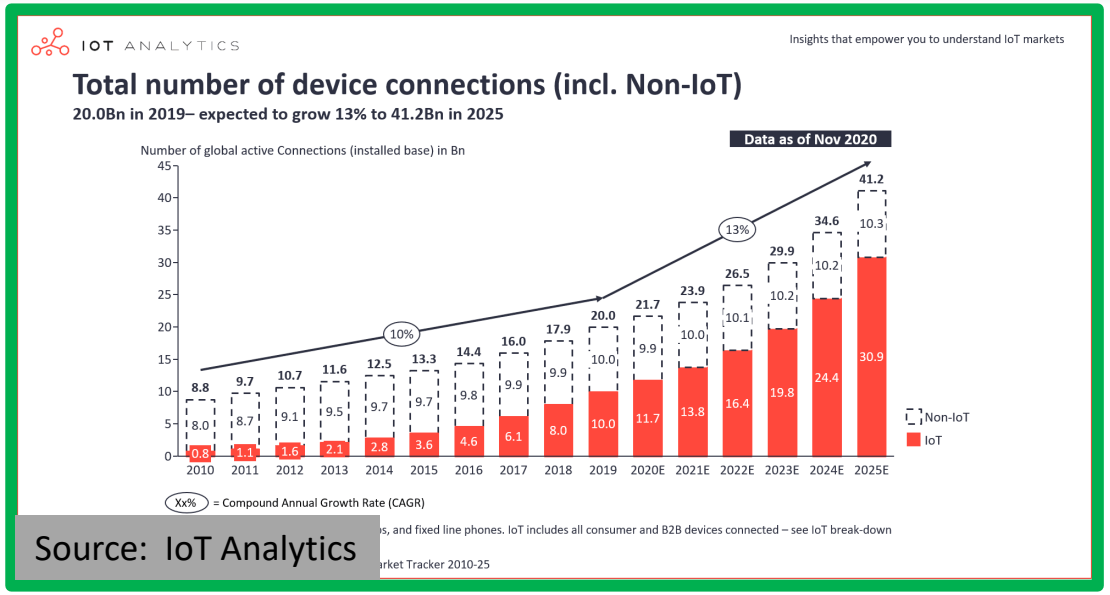

Source: IoT Analytics

30+ Billion IoT Connections by 2025

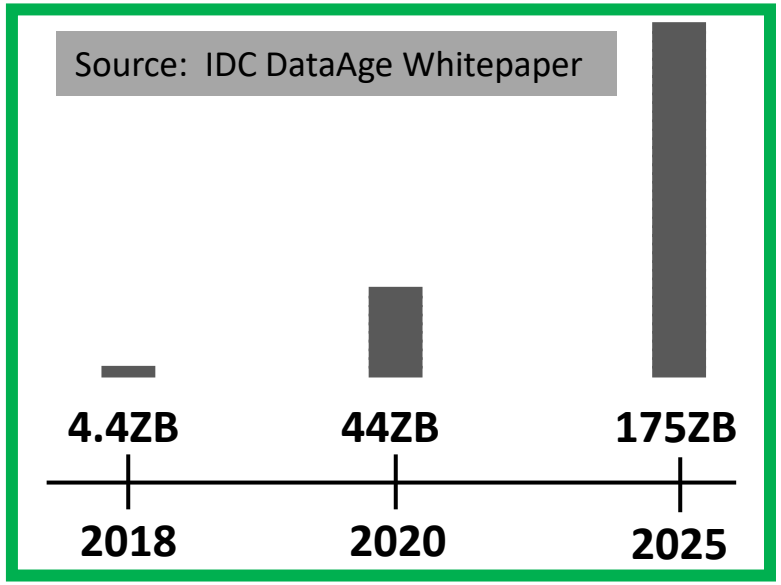

WW data 175ZB by 2025

5G vs 4G:

- 100X faster
- 250X higher density
- 100X lower latency

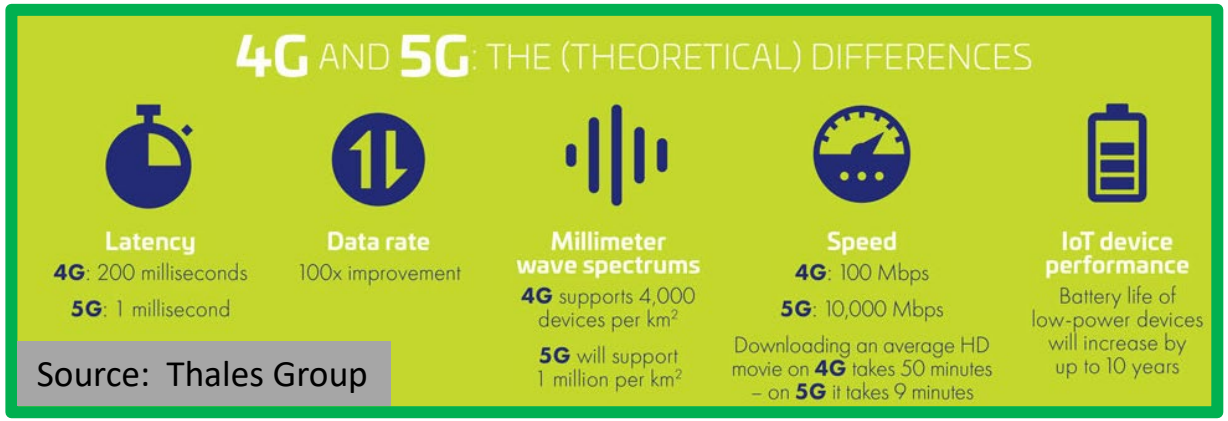

Source: Thales Group

Achronix
Data Acceleration

Data Accelerators Reduce Server Count

Server Only

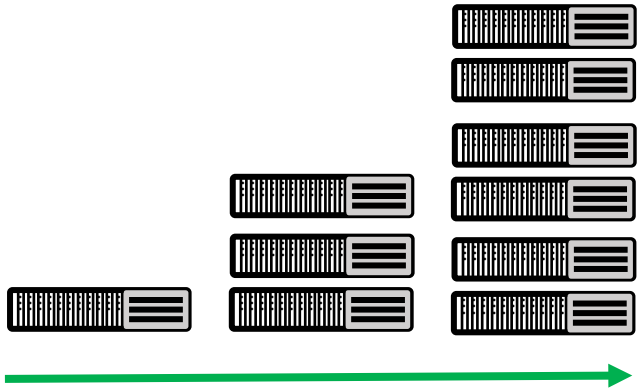

Increasing Data

Server + Data Accelerator

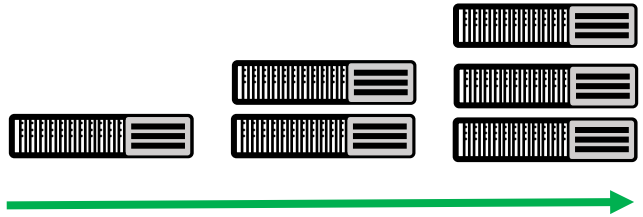

Increasing Data

Problem Statement

- As data increases, more servers used to process data
- A typical servers can use 300-400W and cost $5k-$10k.
- Up to 10X higher latency
- Much less deterministic

Solution → Data Accelerators

Reduction in…..

- Latency
- Power
- Cost
- Real estate

Increase in…..

- Performance

"The data center accelerator market is expected to register a CAGR of 41% over the forecast period (2020 - 2025)." ResearchAndMarkets.com

Achronix
Data Acceleration

Computational Storage Use Cases

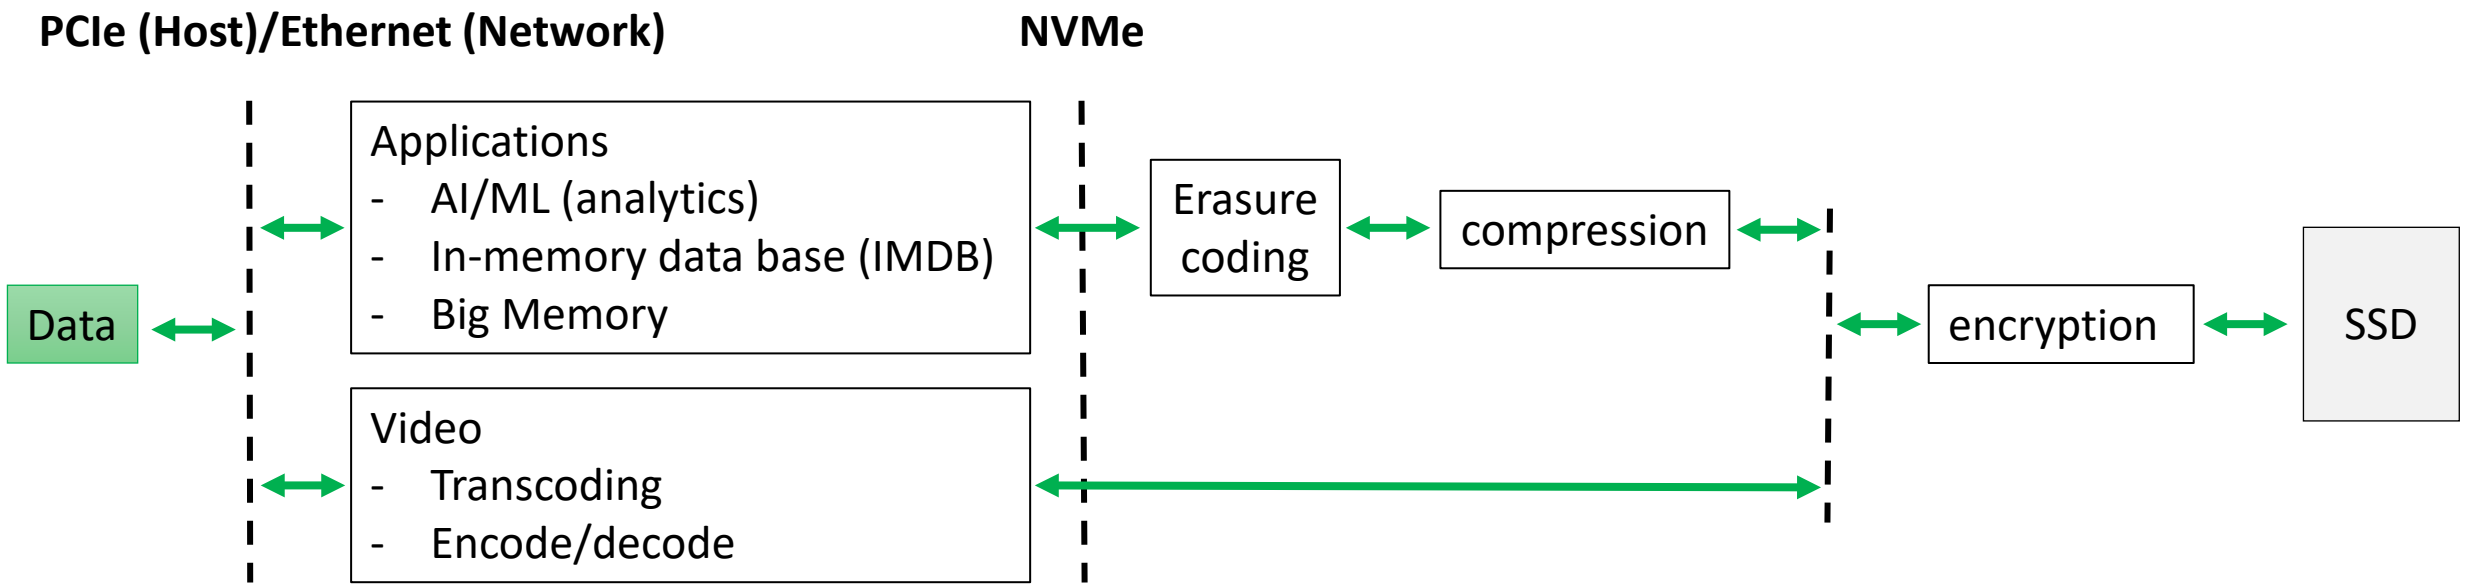

FPGA Data Accelerators

- provide the needed flexibility for changing workloads and algorithms
- programmable like a CPU/GPU/IPU but run at ASIC-like speeds

Achronix
Data Acceleration

Addressing Challenges to Computational Storage

Application Modification:

- Need for flexibility
- Deploy once, repurpose based on workload/use case

Go-To-Market Channel Adoption:

- Provide a flexible, data acceleration platform that easily integrates into standard OEM server and storage hardware

"Crossing The Chasm":

- By providing the flexibility, risk of requirement migration is mitigated

Standardization:

- As standard evolve, a flexible, data accelerator will accommodate changes along the way

The FPGA Data Accelerator

- ASIC-like performance
- Flexibility of a CPU/GPU/IPU

Achronix FPGA Acceleration Solutions

Speedster7t
FPGA Chip

Highest Bandwidth FPGA

- 2D NoC (>20 Tbps)
- 112Gbps SERDES
- 400 GbE
- PCIe Gen 5
- GDDR6 and DDR4/5
- Machine learning processors (MLP)

Speedcore
eFPGA IP

Embedded FPGA IP

Build your own ASIC or SoC using Achronix IP

- TSMC-based libraries
- 16FFC
- 12FFC
- N7
- N5 (coming soon)

VectorPath
Accelerator Card

Bittware Platform

- Speedster®7t AC7t1500 FPGA
- PCIe Gen4 x16
  - Upgrade path to Gen5
- Up to 400GbE
- 16 x GDDR6
- Onboard BMC
- Linux and Windows
- Developer toolkit

Achronix
Data Acceleration

Summary

- Data growing exponentially → expect 175ZB worldwide by 2025
- Data Accelerators reduce infrastructure cost and dramatically increase performance
- FPGA data accelerators provide the programmability of a CPU/GPU/IPU but run at ASIC-like speeds
- FPGAs data accelerators will help bridge many of the market adoption challenges: acceptance of new technology, changing specifications, varying workloads.
- Achronix offers unique portfolio of **discrete FPGAs**, **embedded FPGA IP and Accelerator Cards**

Achronix
Data Acceleration

How to Get Started with Achronix?

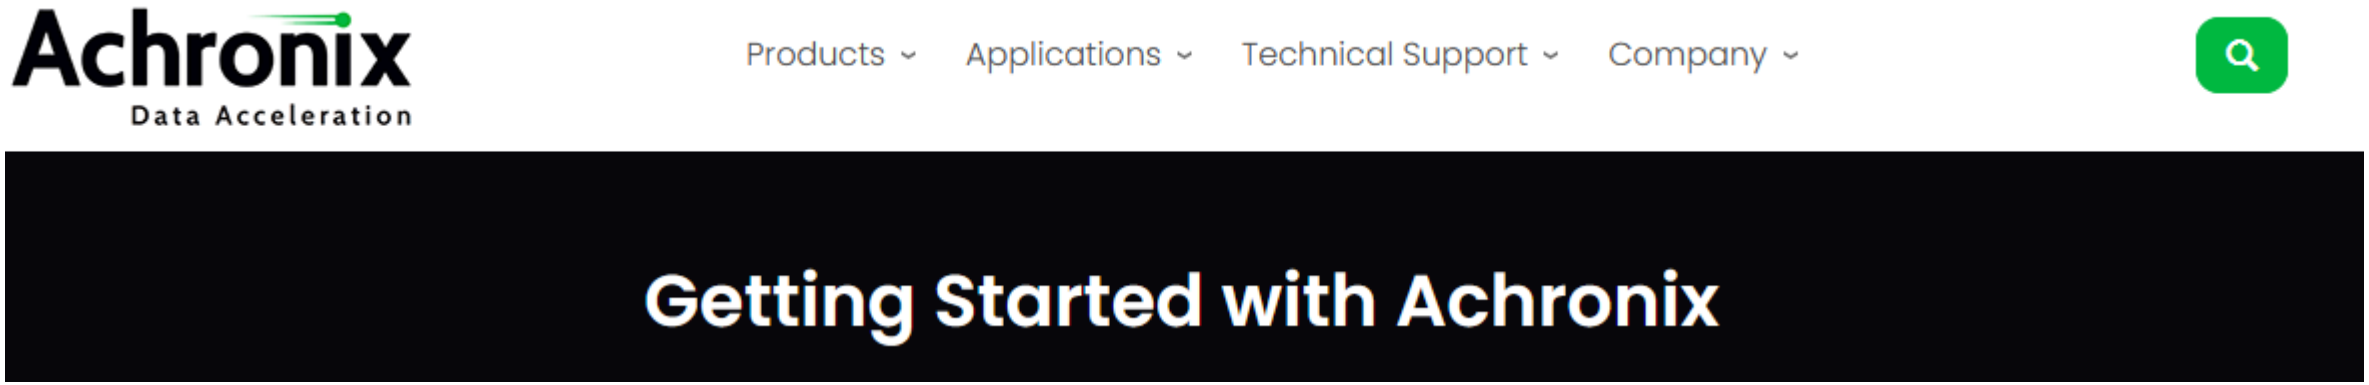

- Register online
- Free evaluation licenses for Achronix design tools
- Rapid prototyping and production-worthy accelerator cards

https://www.achronix.com/getting-started-achronix

Achronix
Data Acceleration

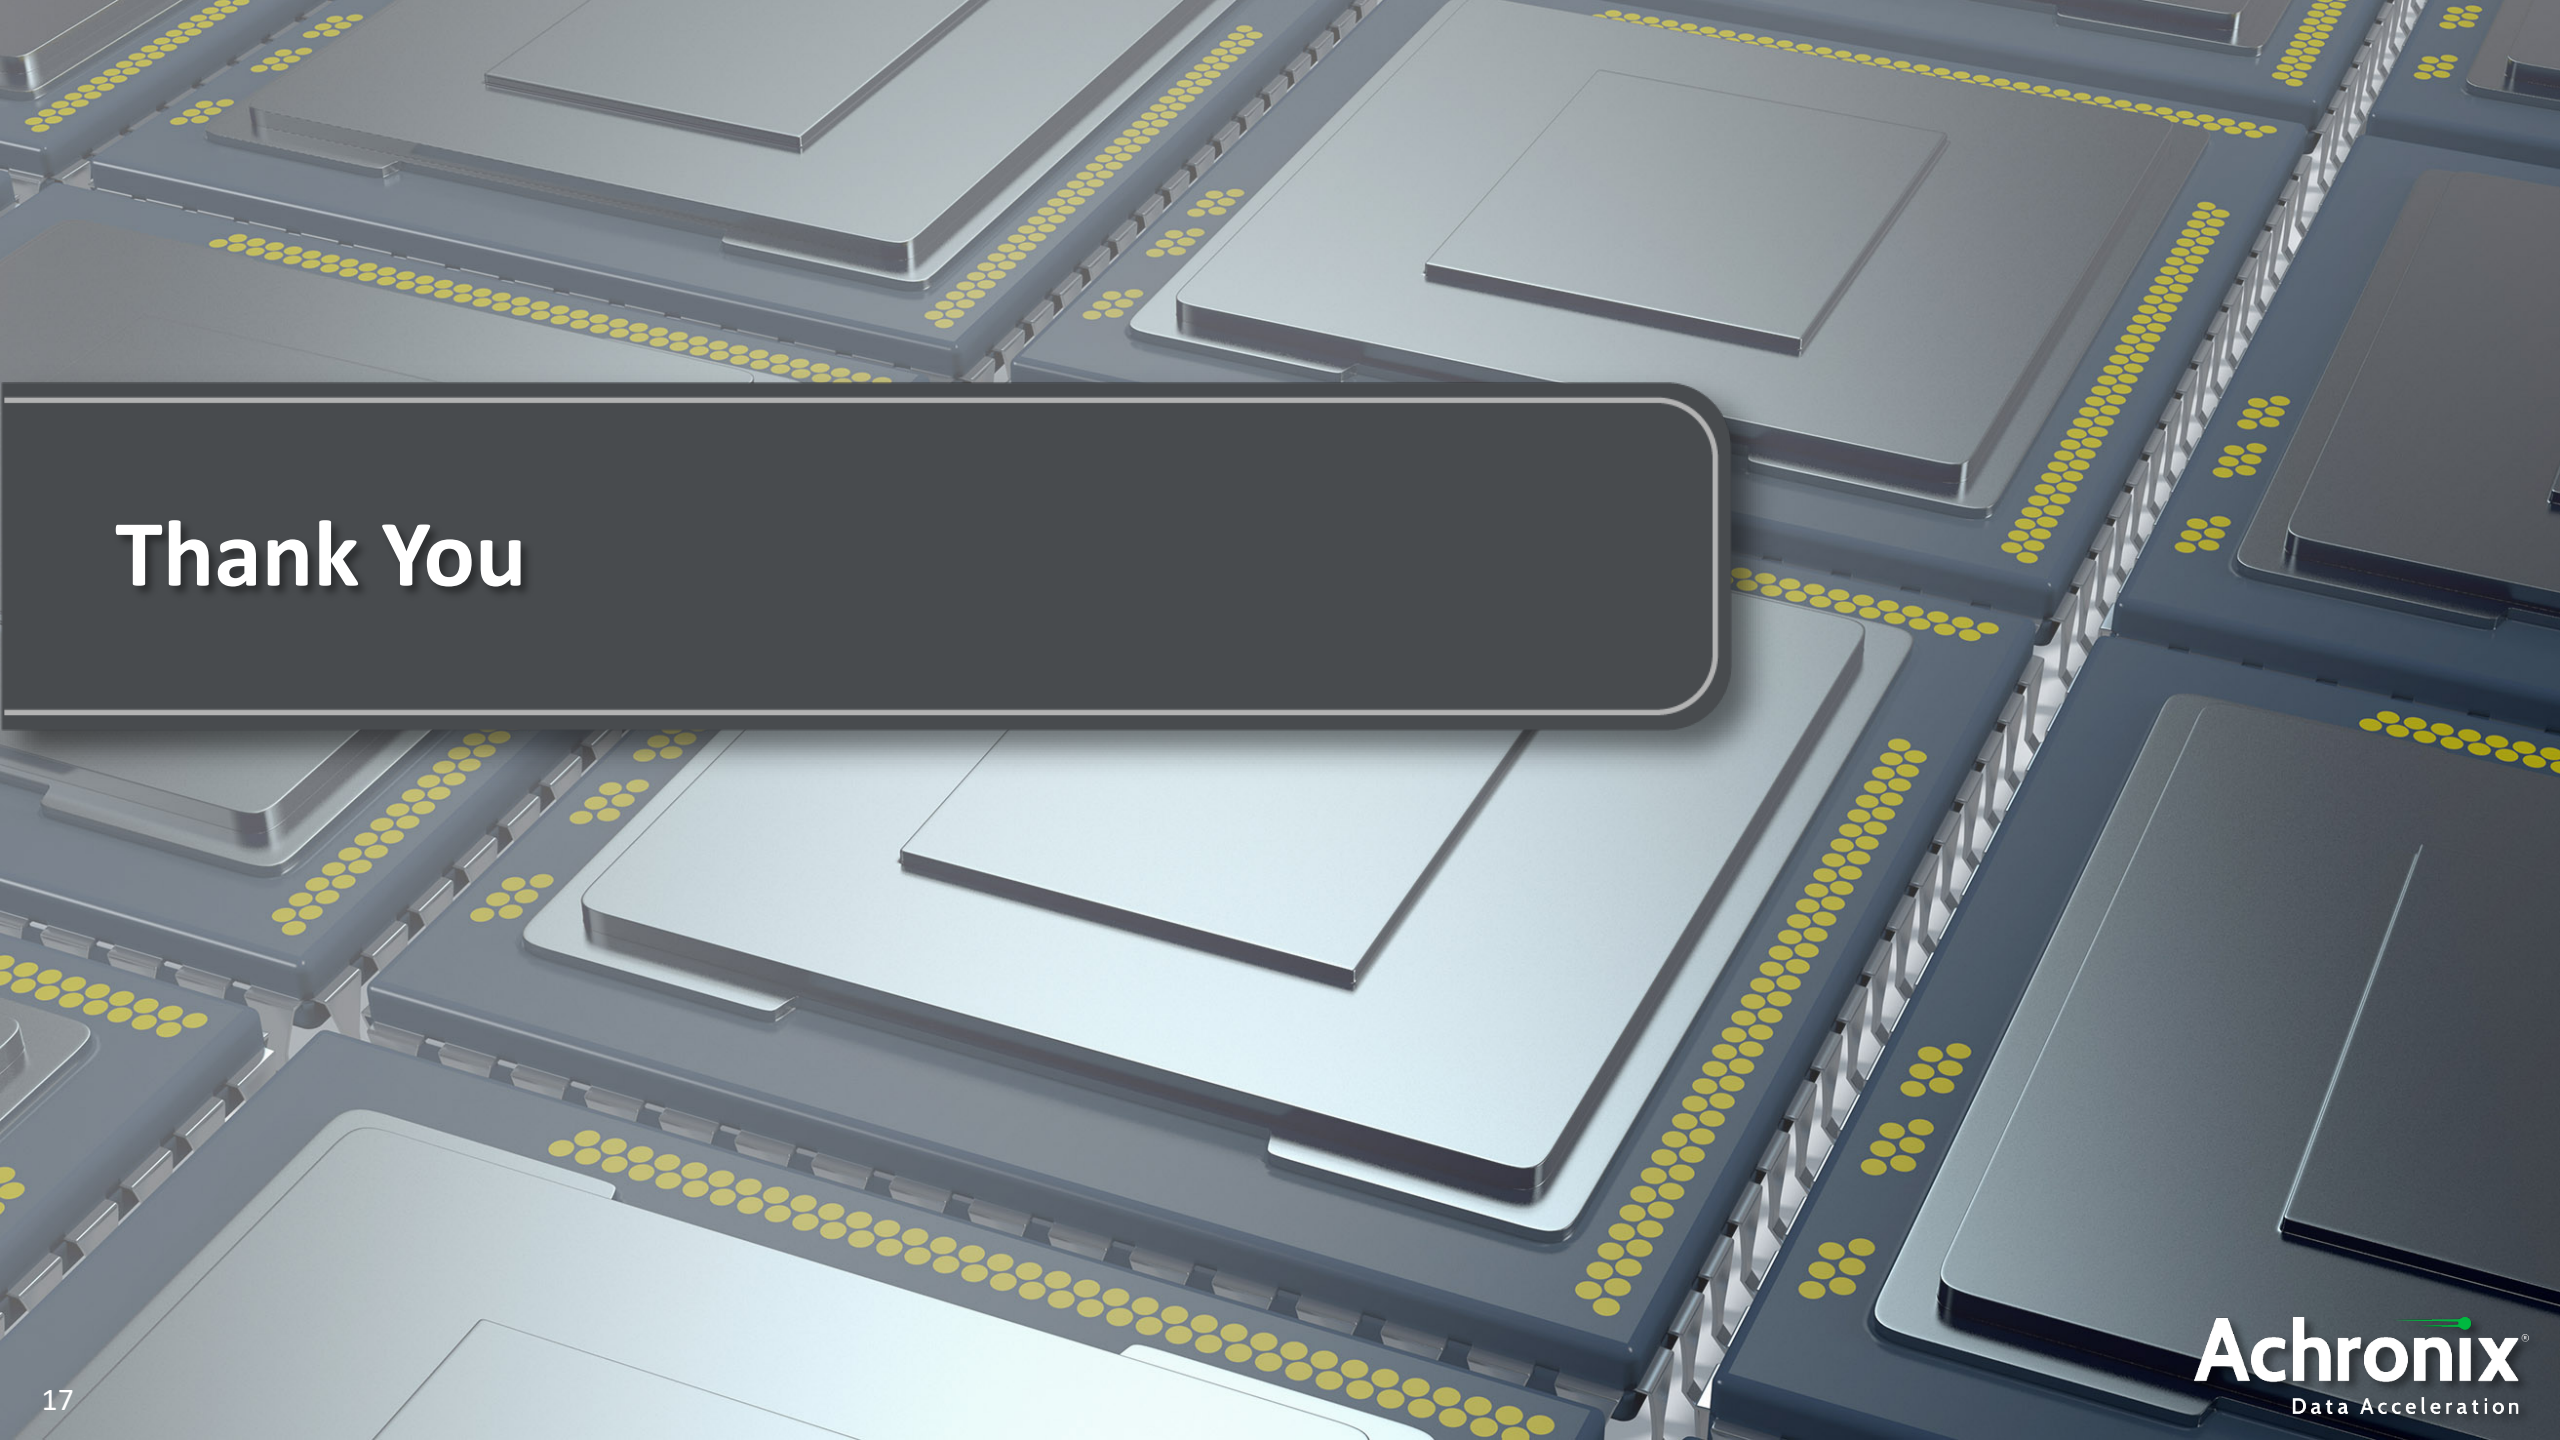

Achronix
Data Acceleration

Tony Afshary
Sr. Director, Product Line Management
www.pliops.com

Pliops Profile

Mission

To massively accelerate performance and dramatically lower infrastructure costs for flash-based data-intensive applications including Databases, Analytics, AI/ML, 5G, IoT, and more

Team

Experts in database, flash storage, and semiconductors from industry leaders including Samsung, Intel, Kioxia, Amazon, Microsoft, Yahoo, VMware, Dell/EMC, Western Digital, Fusion-io, HPE, Apple, Nvidia, Cisco and Lenovo

Customers

More than 20 Fortune 500 cloud and enterprise companies

Strategic Investors

SoftBank Ventures Asia

intel Capital

NVIDIA

Western Digital

XILINX

PLIOPS

Industry Recognition

2021: CRN Top 10 Cool Tech Companies that Raised Funding in February

2021: Enterprise Storage Forum Top Computational Storage Companies

2020: Most Innovative Flash Memory Enterprise Business Application Product

2020: CRN Top 10 Hottest Semiconductor Startup

2019: Most Innovative Flash Memory Startup

The Data-Compute Bottleneck

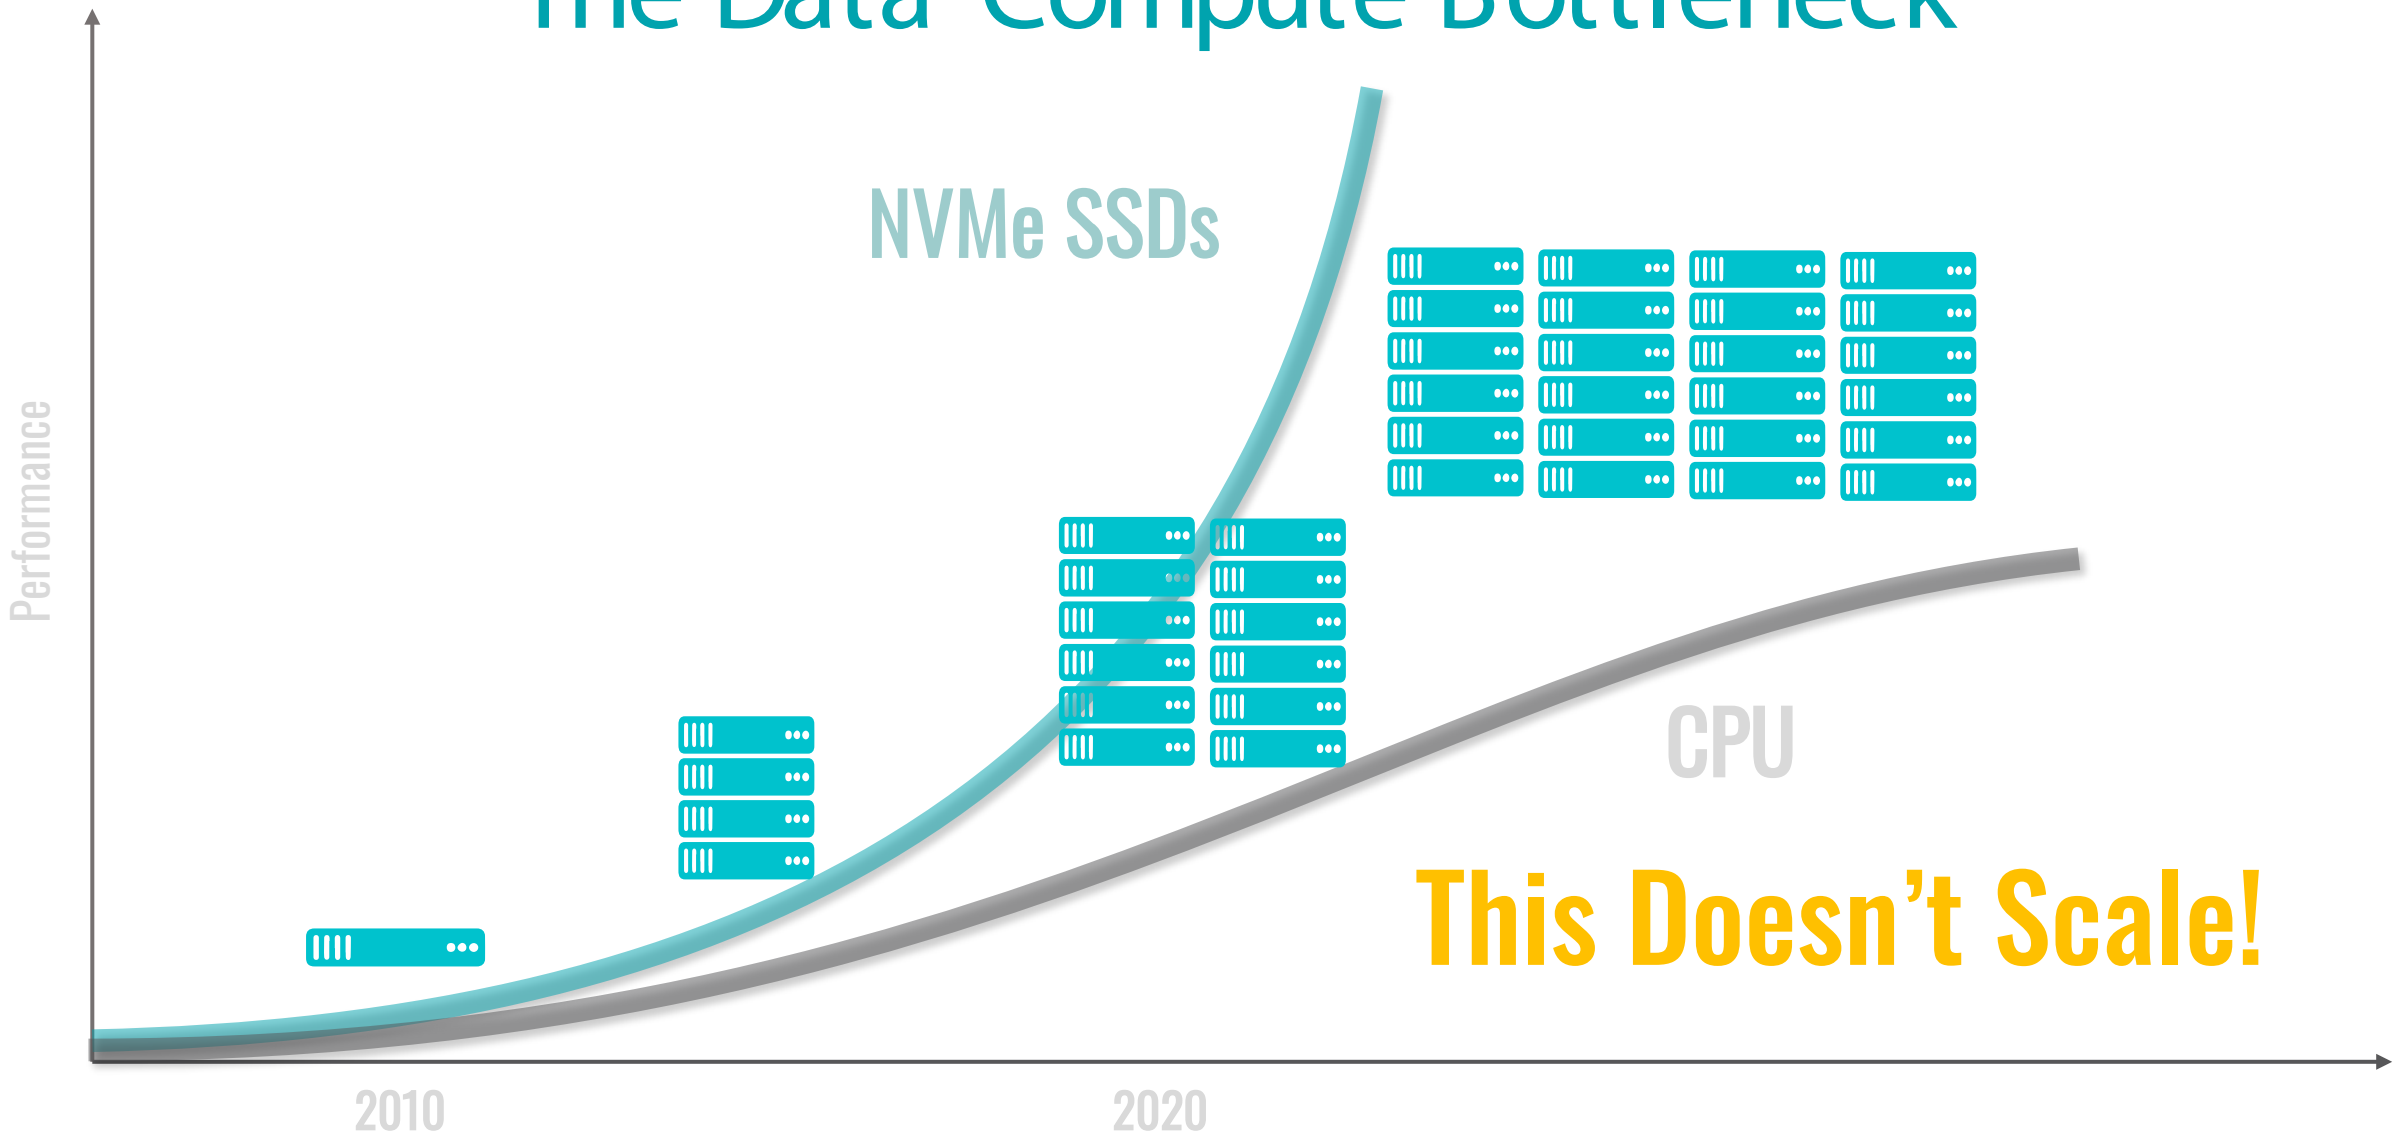

PLIOPS

Yesterday's
COMPUTE CENTRIC

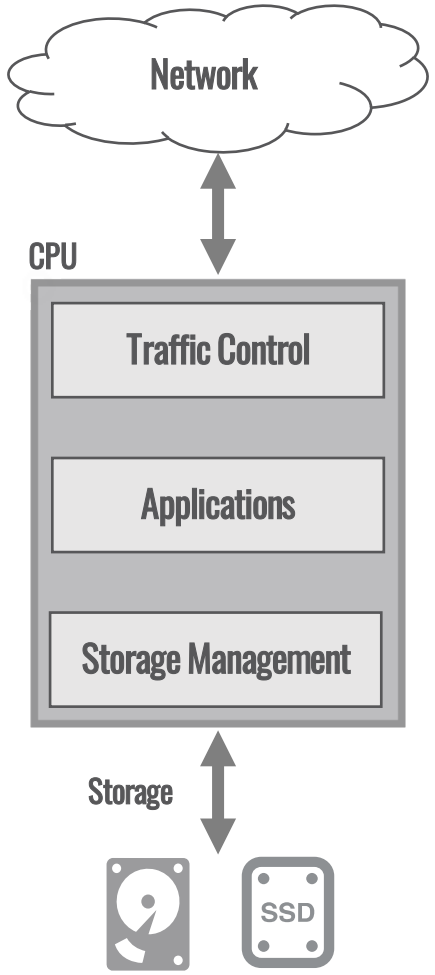

Today's
DATA CENTRIC

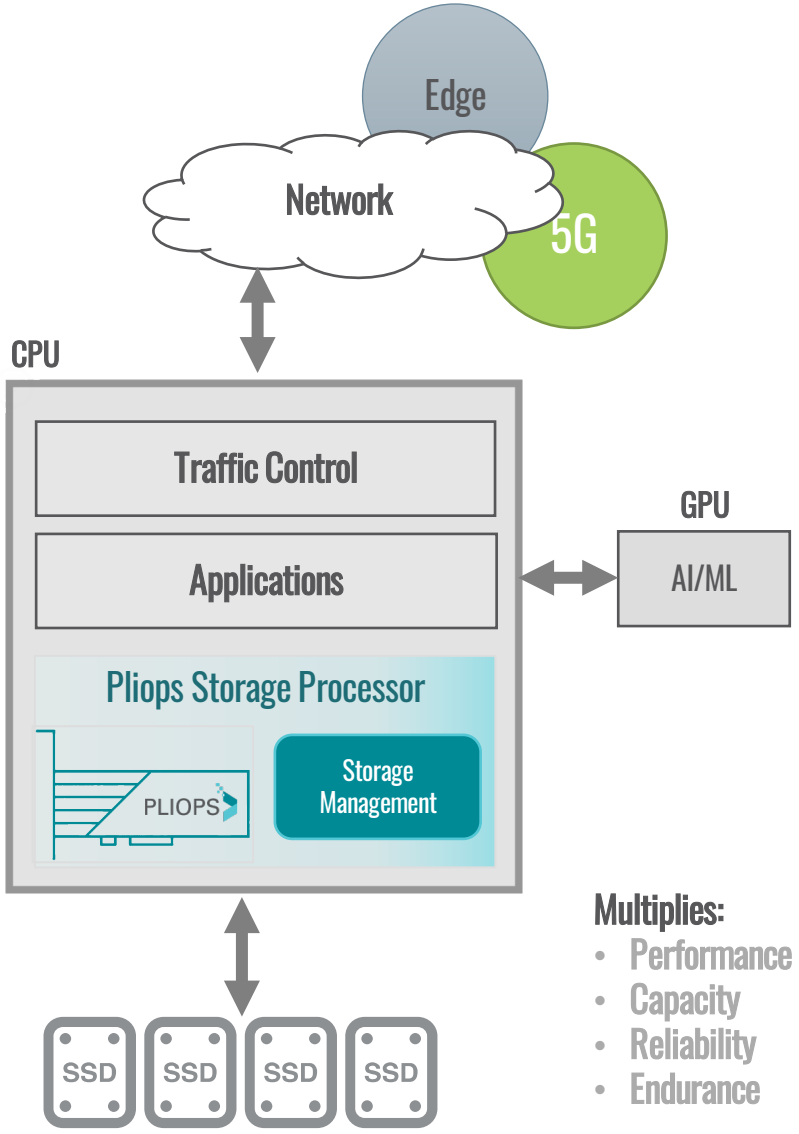

Multiplies:

- Performance
- Capacity
- Reliability
- Endurance

PLIOPS

Pliops Storage Processor

Reliability Drive Fail Protection 2x > RAID 0

Performance 2-15x

Capacity Up to 6x

Efficiency TLC, QLC for any workload

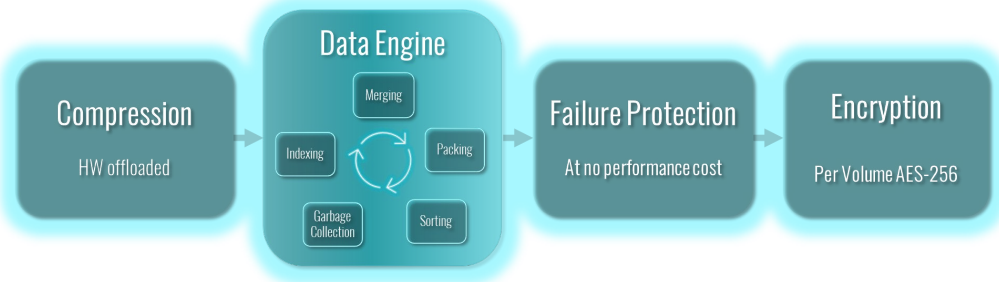

PLIOPS

System Integration Overview

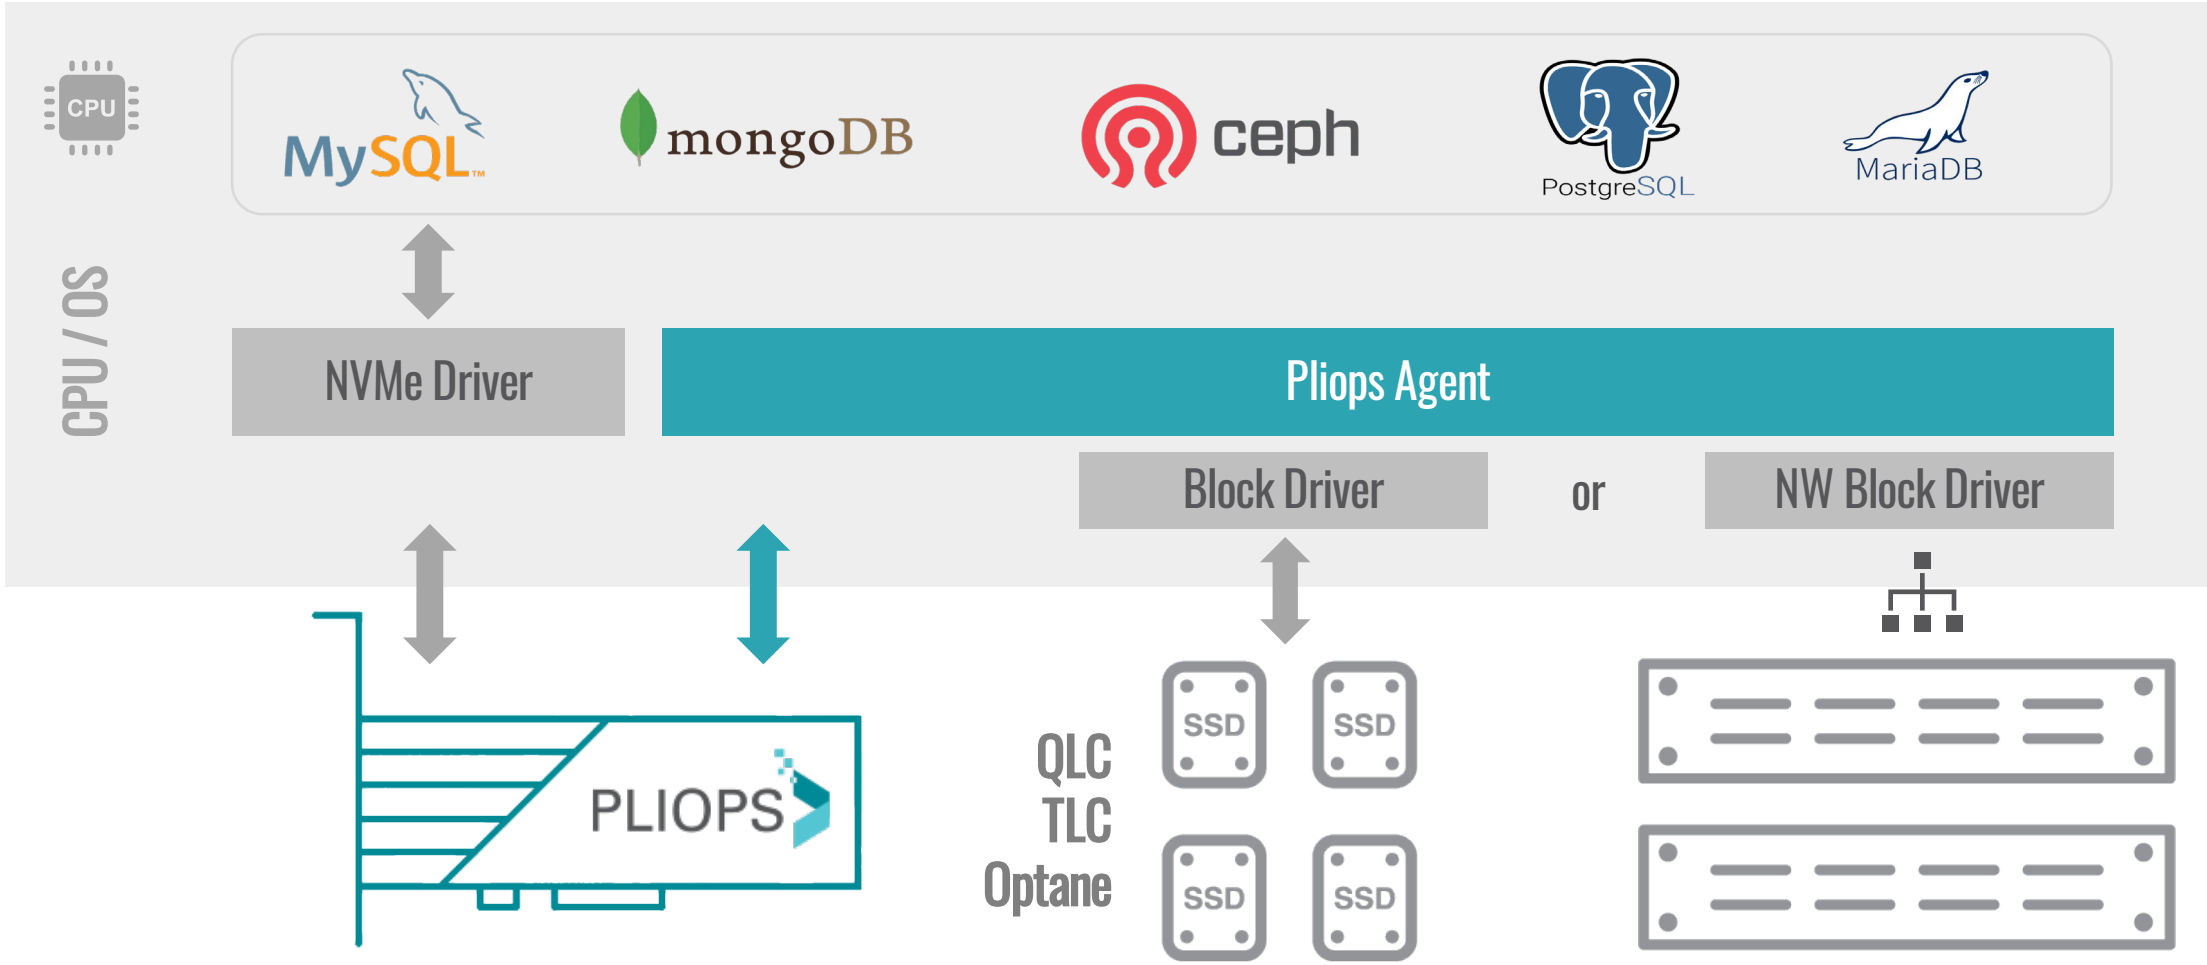

Any Application —Any Standard Server — Any SSD — Direct or Disaggregated

PLIOPS

Cloud Deployments

Disaggregating Infrastructure

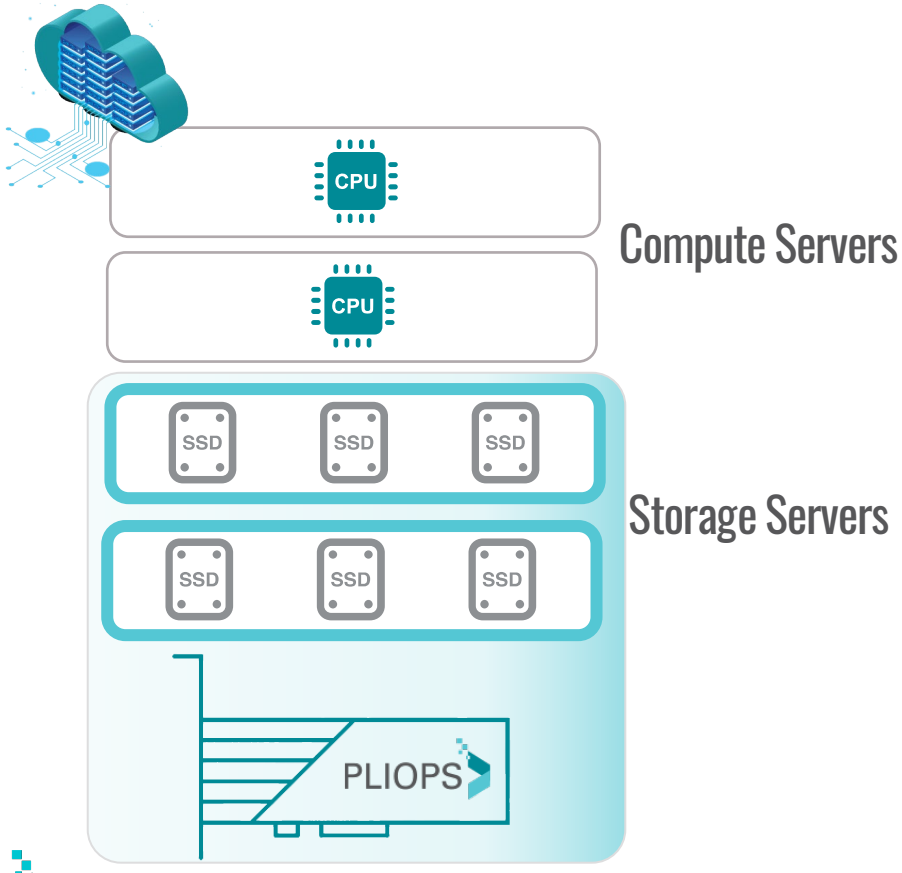

Compute Servers

Storage Servers

Empowering the Application Server

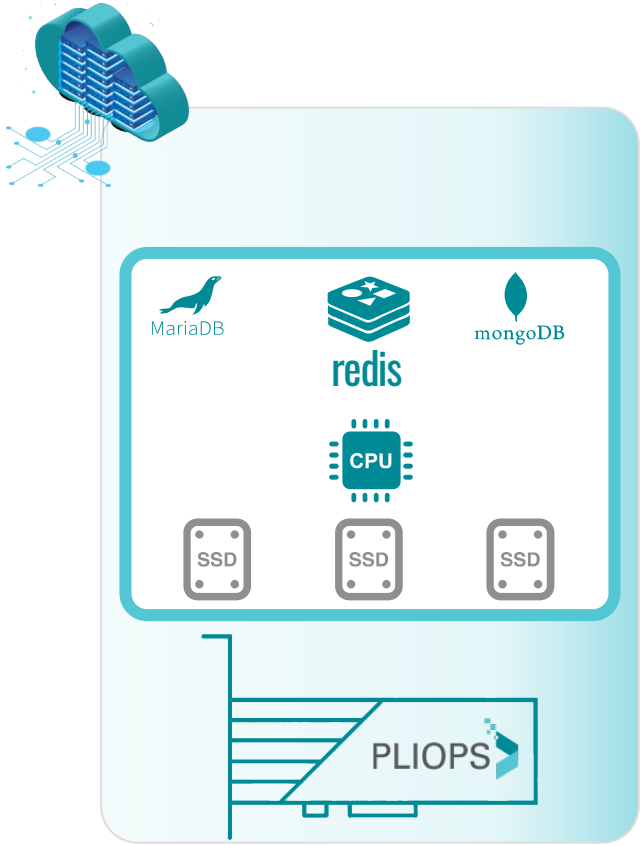

Application Servers

PLIOPS

Use Case: Public Cloud Provider Pliops Impact

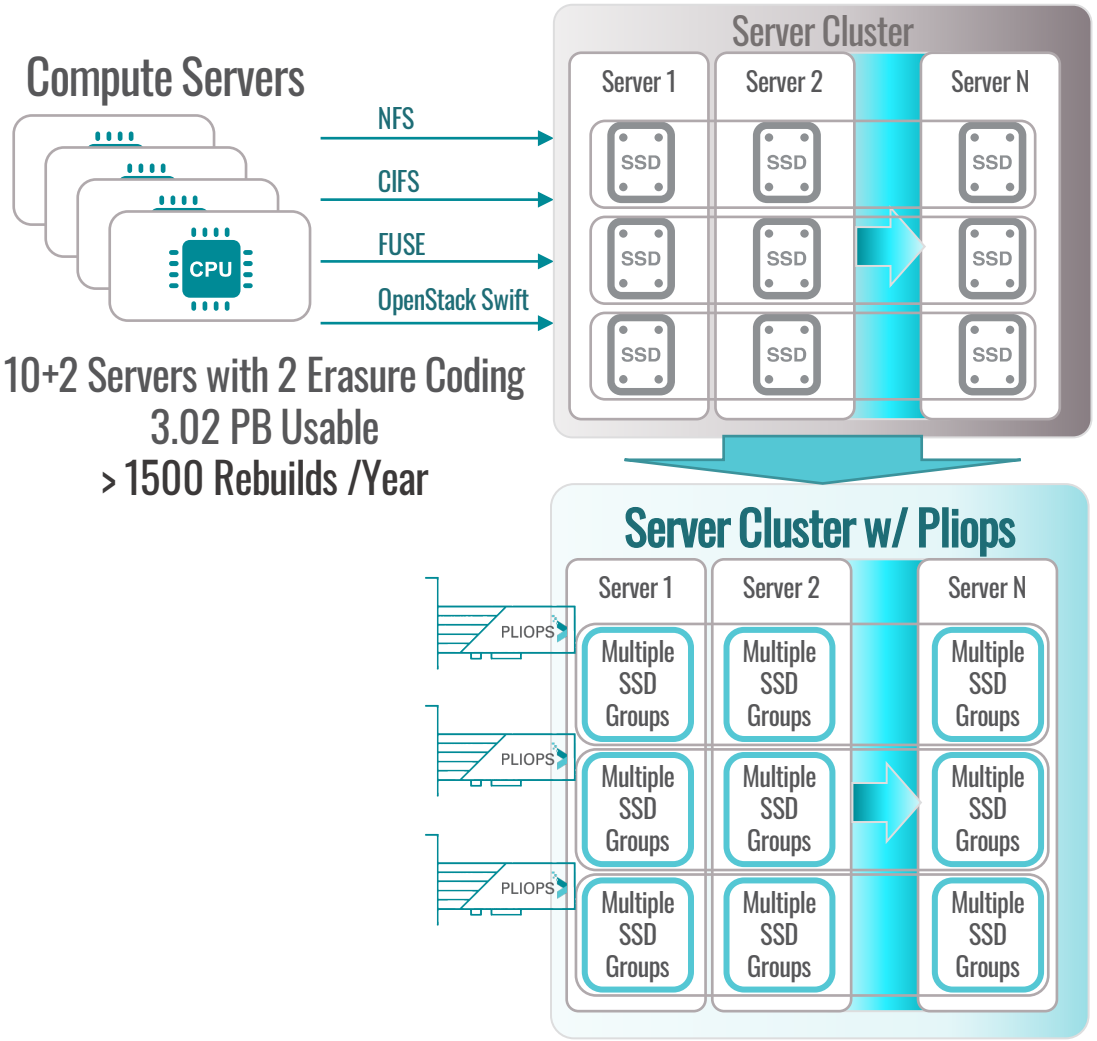

Capex Benefit

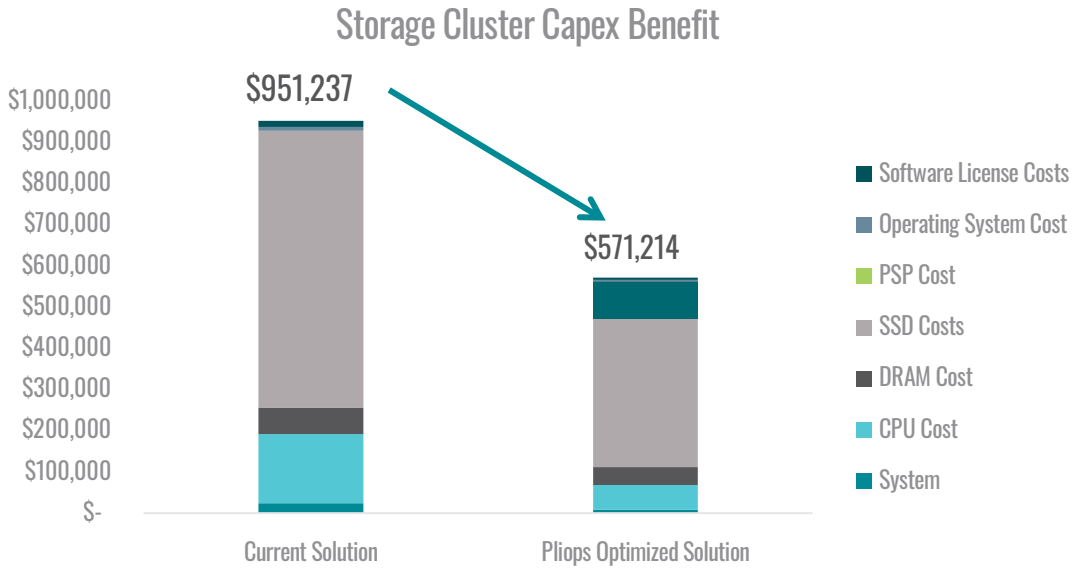

Cost	40%↓
Reliability	100% for free
Endurance	2X vs TLC Setup

PLIOPS

Use Case: Top SaaS Provider TCO

Current Software Based Solution

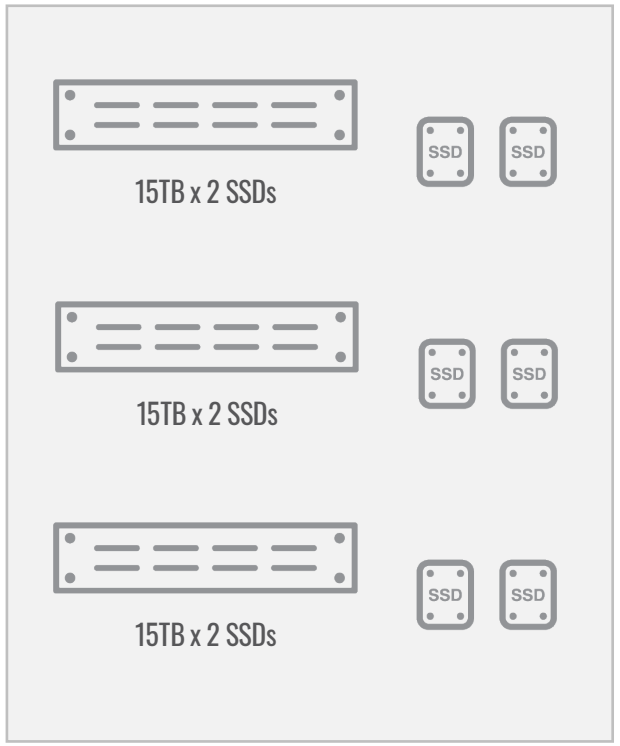

15 DB Instances
41 TB Usable, RAID 0
600 Failures/Year

Pliops Accelerated Solution

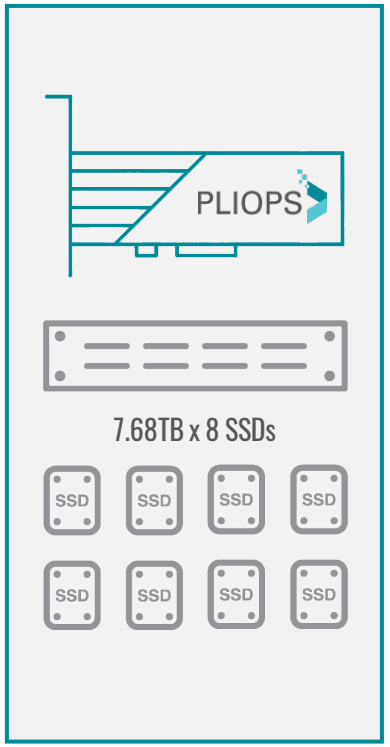

20 DB Instances
66TB Usable, PSP Drive Fail Protection
0 Failures/Year

Capex Benefit

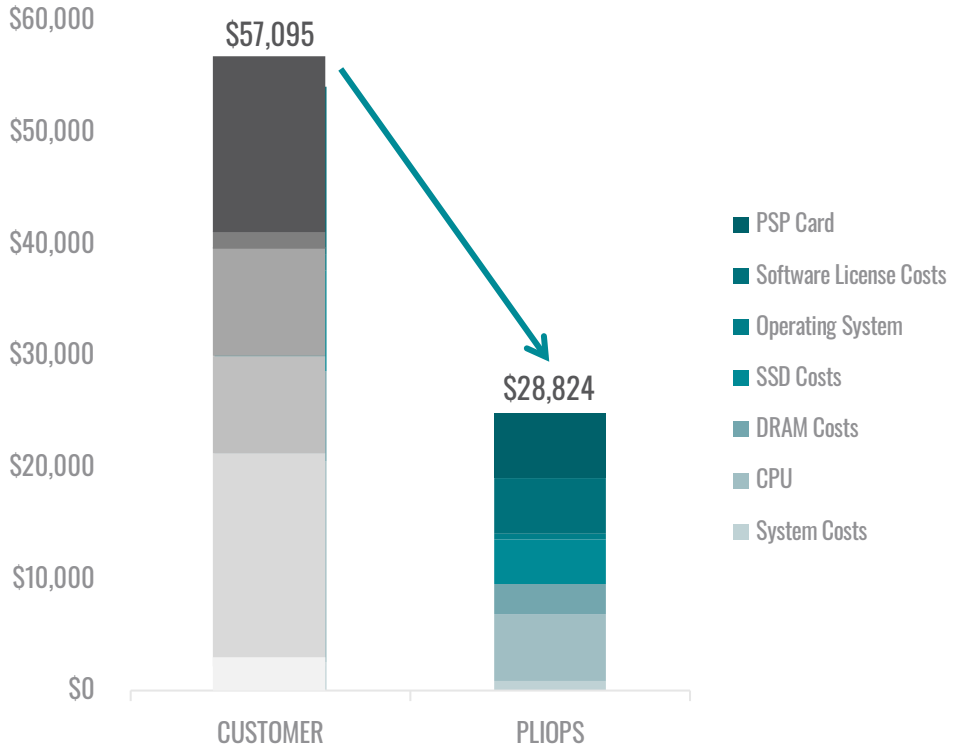

Cost	50%↓
Scaling (DB Instances)	33%↑
Capacity	66%↑
Reliability	100% for free

PLIOPS

Pliops Deployments in the Cloud

Reliability – Increased Data protection and database uptime

Performance –Storage and Database Application scaling

Capacity – Use low-cost QLC or TLC , store more data with no performance cost

Efficiency – Database/Storage consolidation, reduce infrastructure footprint

PLIOPS

ScaleFlux™

JB Baker

VP of Marketing

www.ScaleFlux.com

ScaleFlux®

Computational Storage Drives:
Evolving SSDs for Data-Centric Architectures

July 2021

ScaleFlux: Computational Storage Leader

Low Latency Flash Storage

Compute Engines

Industry Standard Infrastructure

Data Driven Applications

Database | Big Data | AI/ML | IOT/Edge | CDN

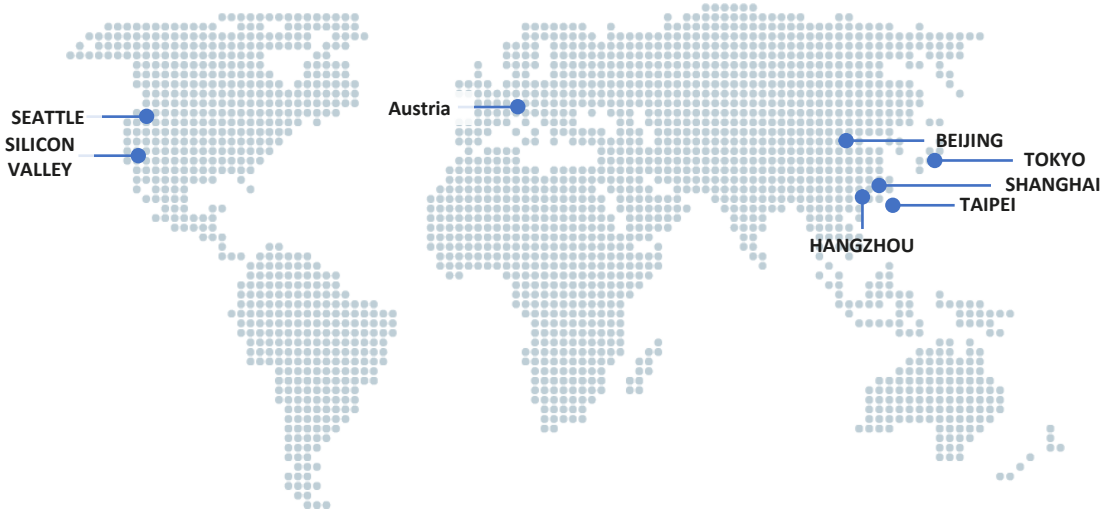

- Founded in 2014
- **1st Production Computational Storage**
- **2nd Generation of CSD shipping**
- US & China manufacturing sites
- Strong Patent Portfolio
- Well Funded
- Hyperscale, webscale & enterprise customers

2018
FMS - Most Innovative Startup
ITBrand Pulse - Market Leader

2019
FMS - Most Innovative Customer Implementation
IDC Innovator – Computational Storage

2021
Data Breakthrough Awards – Flash Solution of the Year
Gartner – Cool Vendors in Storage & Hybrid Infrastructure

What is a Computational Storage Drive (CSD)?

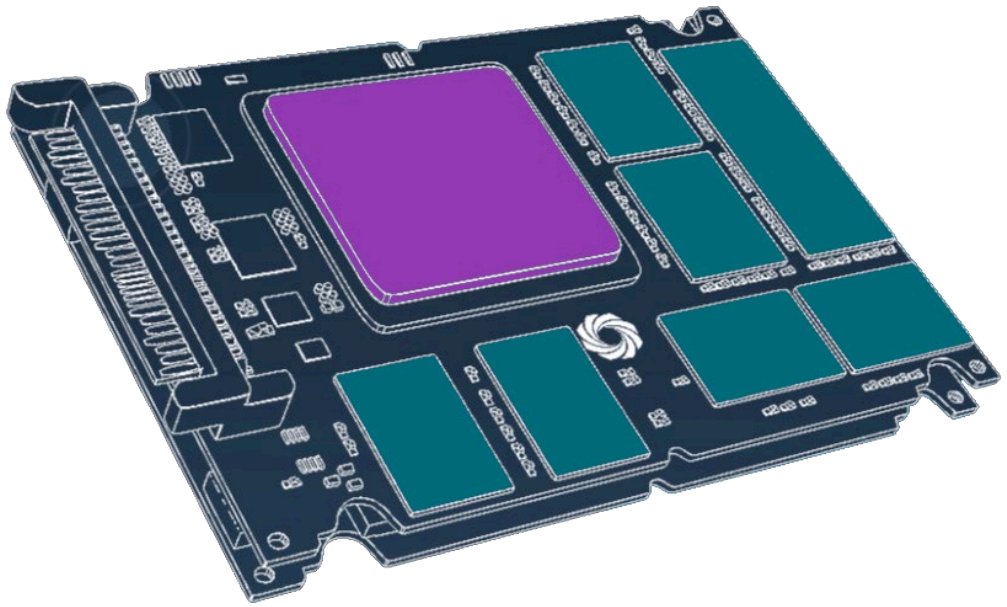

Enterprise PCIe SSD + Compute Offload Engines

ScaleFlux CSD 2000:

Enterprise/Datacenter PCIe SSD

- ✓ Industry standard form factors: U.2 & HHHL Add-in Card
- ✓ 4TB, 8TB & 16TB Raw Capacities
- ✓ TLC & QLC Options
- ✓ Typical enterprise features

Plus…

- ✓ Transparent Compression/Decompression
- ✓ Extendable Capacity
- ✓ Configurable Overprovisioning
- ✓ Atomic Writes

Why use Computational Storage?

Processor-Driven Architecture

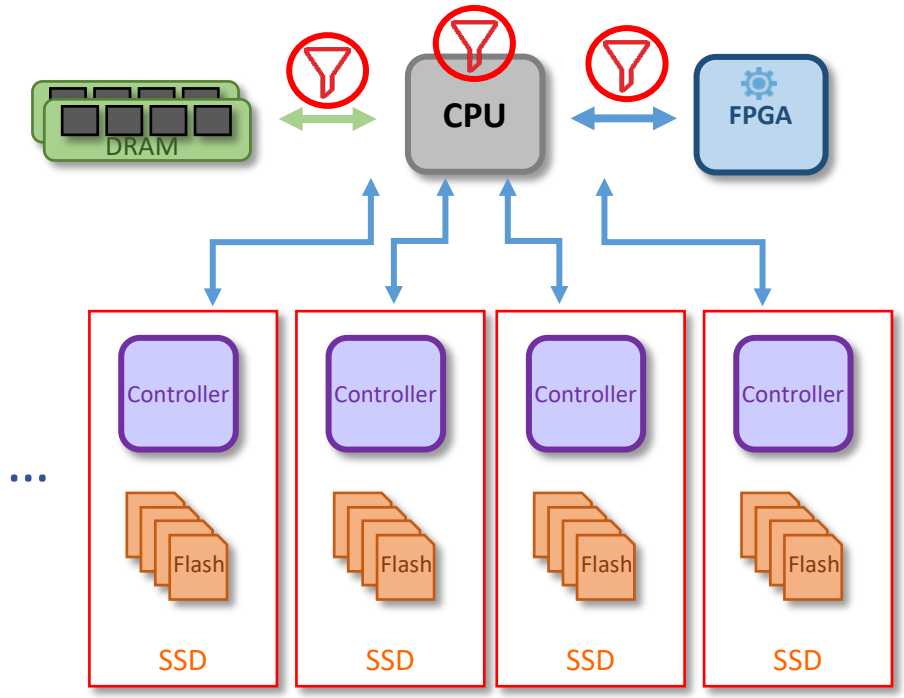

- CPU & Memory I/O bottlenecks
- Limited FPGAs, specific sockets required
- Massive data movement
- Scaling challenges
- **No compute parallelism**

Data-Driven Architecture

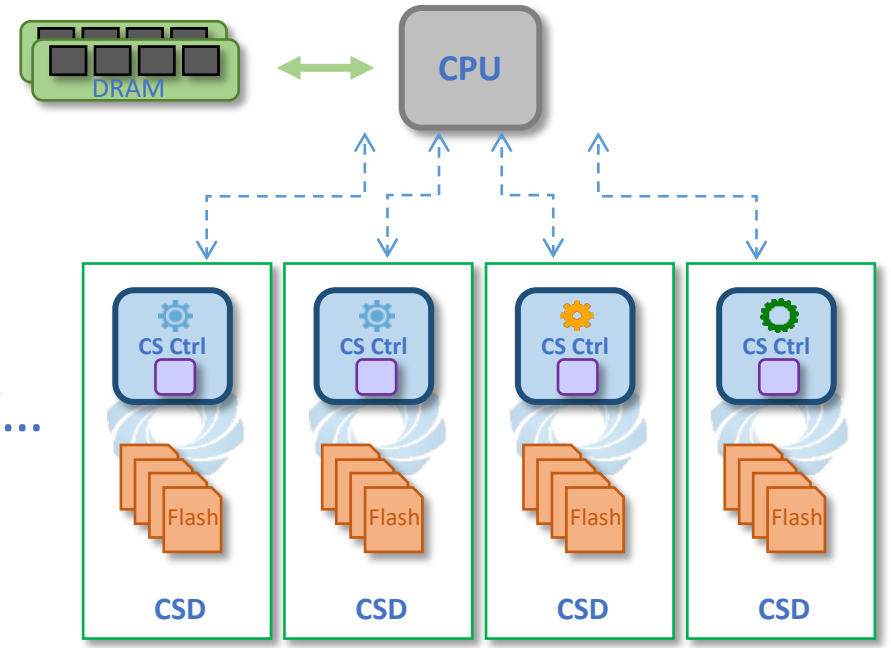

CSD: Computational Storage Drive
CS Ctrl: Computational Storage Controller

- Balanced compute & storage I/O
- Multiple CS Engines, integrated with storage
- Minimize data movement
- Scales compute and storage
- **Maximum compute parallelism**

Example Application Benefits

	Performance*	Cost & Capacity*
AEROSPIKE	• ***98% Better Latency*** • 2x TPS	• 3-4x Storage Capacity • 60%+ Cost Savings
MySQL MariaDB	• Latency consistently lower • Up to 3x QPS	• 3-4x Storage Capacity • ***60%+ Cost Savings***
PostgreSQL	• Latency consistently lower • Up to 1.4x QPS	• 3-4x Storage Capacity • ***60%+ Cost Savings***
redis	• Latency consistently lower	• 4x Storage Capacity • ***Up to 8x Endurance***

*With CSD 2000 vs Ordinary NVMe SSDs

Thank You

97 East Brokaw Road, Suite 260
San Jose, CA 95112
www.scaleflux.com #compute2data

Panel Questions and Audience Surveys

Panel Question #1

Flash storage devices seem to be going in (at least) two different directions: hyperscalers are asking for devices that they can control at the “micro level” (Open Channel SSDs), while many datacenter devices are adding more intelligence and (in some cases) network interfaces. How does this impact the positioning of computational storage in the datacenter?

- Tom Spencer - Achronix
- Tony Afshary – PLIOPS
- JB Baker - ScaleFlux

Audience Survey Question #1

How does your organization view the positioning of computational storage in the datacenter versus conventional storage/computation virtualization architectures? (check one):

- Computational storage is critical – uniquely solves a number of very important application problems: 24%
- Computational storage is a relevant approach to solve some important applications problems: 35%
- There may be a few “niche” applications that computational storage has significant value for, but it is not a “wide” application accelerator: 24%
- Computational storage may not be viable for applications, but can provide “background” functions (indexing, encryption, compression, etc.): 0%
- Computational storage doesn’t have meaningful place in the datacenter: 0%
- Don’t know/no opinion: 18%

Panel Question #2

One set of use cases where computational storage provides ‘value add‘ is extremely large dataset (‘petabyte-scale’) problems, where it eliminates the need to move that data across the PCIe bus. How common are these type of computational problems, and how does computational storage compare to other alternatives?

- Tony Afshary – PLIOPS
- JB Baker – ScaleFlux
- Tom Spencer - Achronix

Audience Survey Question #2

What experience does your organization have with computational storage? (check all that apply):

- Explored information on its use (conferences, articles, etc.): 31%
- Talked to computational storage vendors: 25%
- Defined potential computational storage projects: 13%
- Started one or more proof-of-concept evaluations: 19%
- Budgeted for production computational storage deployments: 6%
- Deployed computational storage in production: 6%
- No experience/don't know 38%

Panel Question #3

One of the stumbling blocks in getting datacenter traction for computational storage has been the perceived need to modify, port, or adapt applications. How big of an issue is this, and are there examples of successfully tackling this?

- JB Baker – ScaleFlux
- Tom Spencer – Achronix
- Tony Afshary – PLIOPS

Audience Q&A


Thank You For Attending!