



Intelligent Storage Solutions

Producing Efficient Scalable Systems



Dr. Vladimir Alves
CTO & Co-Founder

Scott Shadley
VP Marketing

Dr. Jae Young Do
Microsoft Research



Flash Memory Summit
BEST OF SHOW WINNER

Today's way of finding that needle...

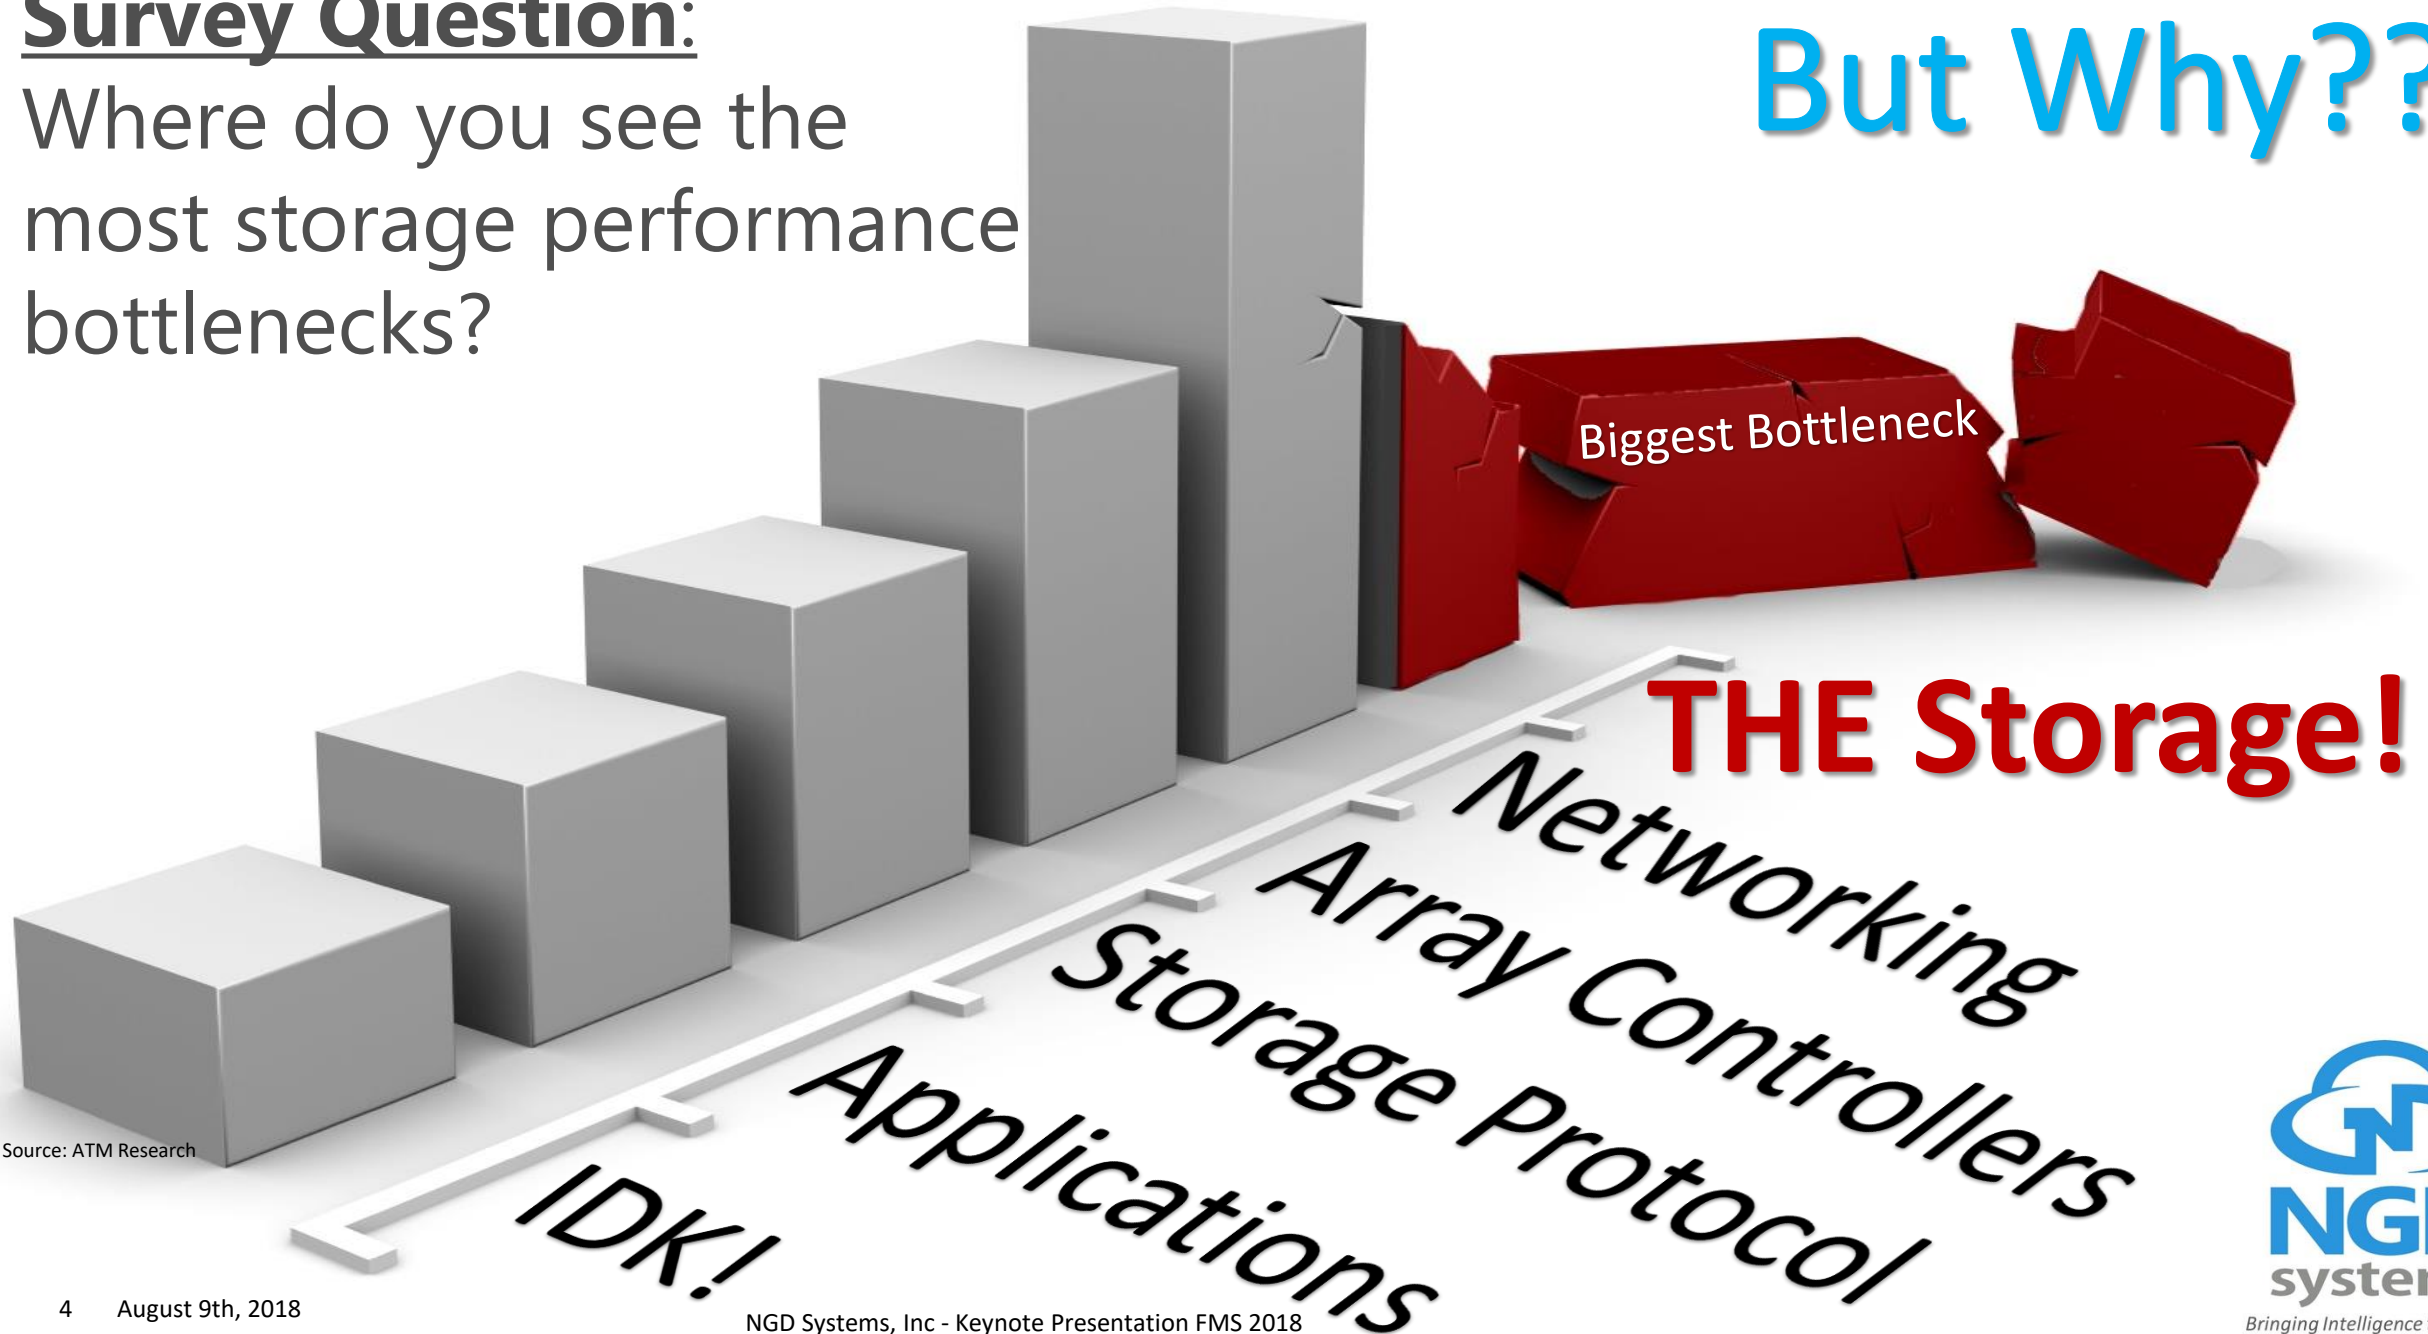
Facing the men is an immense display with 16 screens. It conveys live images from network of 180 CCTV cameras



Survey Question:

Where do you see the most storage performance bottlenecks?

But Why??



Source: ATM Research

3 Factors driving the Storage Problem... Tsunami



'Tsunami of data' could consume one fifth of global electricity by 2025

Keys To Harnessing The Data Tsunami



Jonathan Salem Baskin Contributor ⓘ
Jun 13, 2016, 10:00am • 1,486 views • #BigData

“We have a tsunami of data approaching.” —
Anders S.G. Andrae

the Analytical Scientist Defying the Data Tsunami

The Big Data Tsunami



Author: Matt Ferrari
Chief Technology Officer
ClearDATA

3 Factors driving the Storage Problem... Power

Gartner Says Data Center Power, Cooling and Space Issues Are Set to Increase Rapidly as a Result of New High-Density Infrastructure Deployments

STAMFORD, Conn., May 13, 2010



A heat-exchange process commonly used for cooling submarines to the underwater datacenter. The system pipes seawater directly through the radiators on the back of each of the 12 server racks and back out into the ocean

Free cooling

There are quite a few data centers that have embraced "free cooling" totally, i.e. using the cold air outside.


All you need is ... a mild climate



3 Factors driving the Storage Problem... Near-Data Compute

PUSHED TO THE EDGE

February 19, 2018 Timothy Prickett Morgan



AI Weekly: Computing power is shaping the future of AI

KHARI JOHNSON @KHARIJOHNSON MAY 18, 2018 7:14 PM

NEAR-DATA PROCESSING: INSIGHTS

Near-Data Computation: Looking Beyond Bandwidth

Published in: [IEEE Micro](#) (Volume: 34, [Issue: 4](#), July-Aug. 2014)

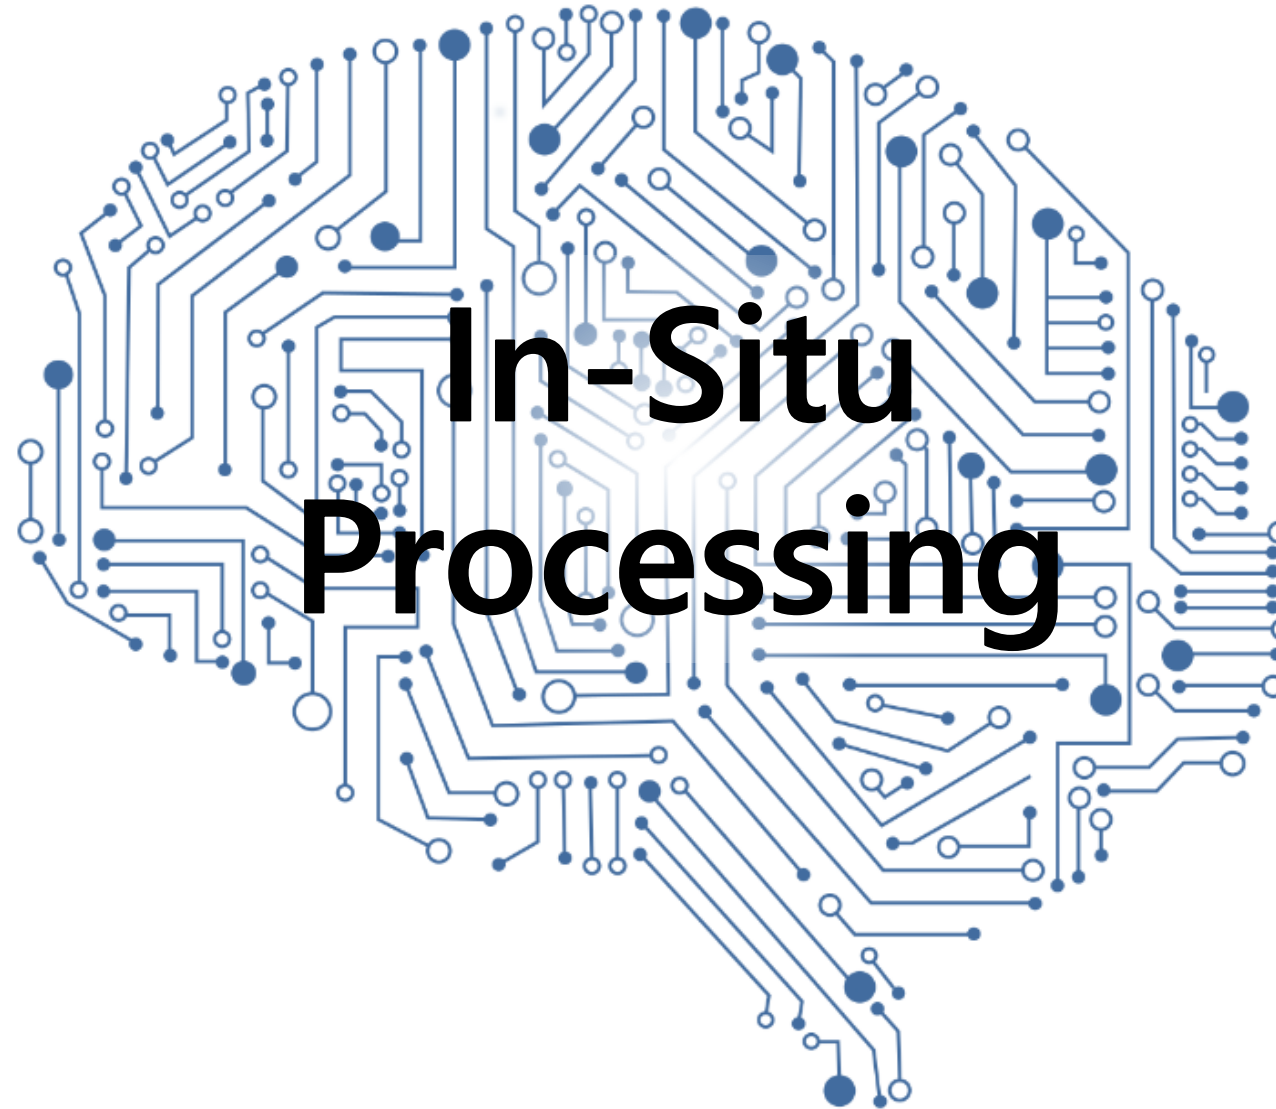
Three motivating factors for using Edge Computing

IBM


Internet of Things blog

1. Preserve privacy
2. Reduce latency
3. Be robust to connectivity issues

Solving the Data Growth and Compute Problem



Challenges with Moving (Big) Data



power density

Watts/Terabyte

volumetric density

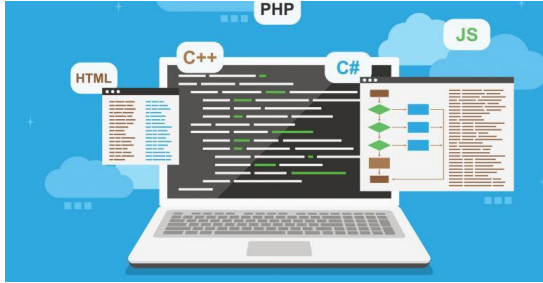
Terabytes/cm³

data bottleneck

SkinnyPipes

Using *Near Data Processing* to Tackle Data Bottlenecks


EFFECIENCY MATTERS



Seamless Programming Model




Scalability




Manage Capacity Growth

IN-SITU PROCESSING

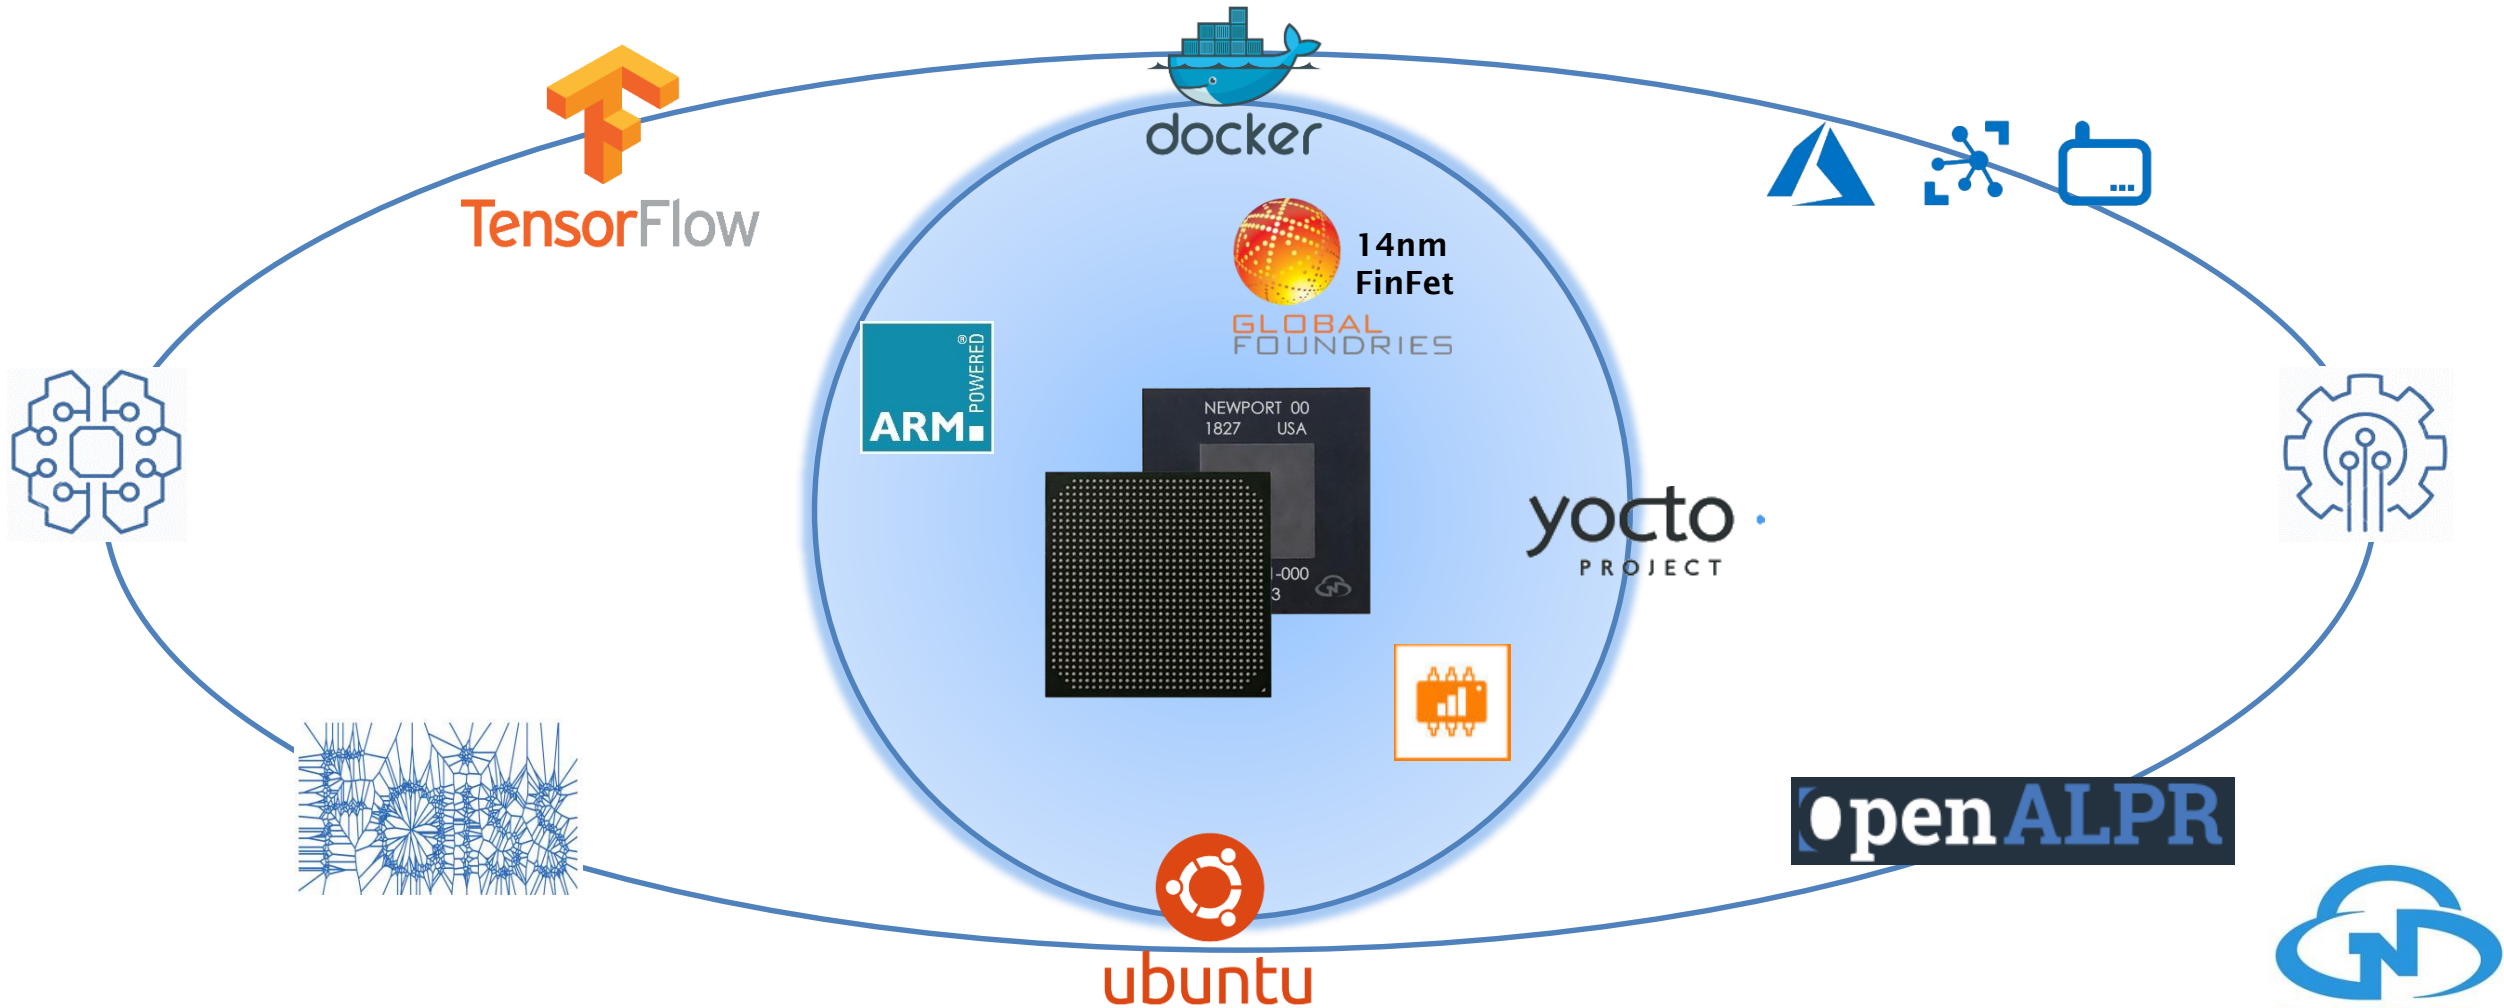


Brings harmony back to bandwidth needs

Dimensions that Enable Computational Storage



In-Situ Processing Ecosystem – The Newport Platform



Delivering the Solution – NGD Systems NVMe SSD Family


New Rack-Scale Form Factors



8TB / 8W



16TB / 12W



**Up to 64TB
16 flash channels**

Traditional Storage Form Factors



32TB / 12W



64TB / 13W

Use Cases


In-Situ Openalpr demo




Result:




```
plate0: 10 results
- V0DKAAE confidence: 92.0216
- V0DKAA confidence: 90.6883
- V0DKAAE confidence: 85.0866
- V0DKAAE confidence: 84.0379
- V0DKAA confidence: 83.7532
- V0DKAAE confidence: 83.5281
- V0DKAAE confidence: 82.9307
- V0DKAAE confidence: 82.7889
- V0DKAA confidence: 82.7046
- V0DKAAE confidence: 82.4195
```




Object Tracking – Azure IoT Hub



AI Image Classification



The Results of TensorFlow To Predict Across 1000 Labels:



Persian cat (score = 0.77295)

tiger cat (score = 0.03713)

tabby, tabby cat (score = 0.03207)

Egyptian cat (score = 0.02612)

lynx, catamount (score = 0.01572)

sea anemone, anemone (score = 0.00072)




ubuntu

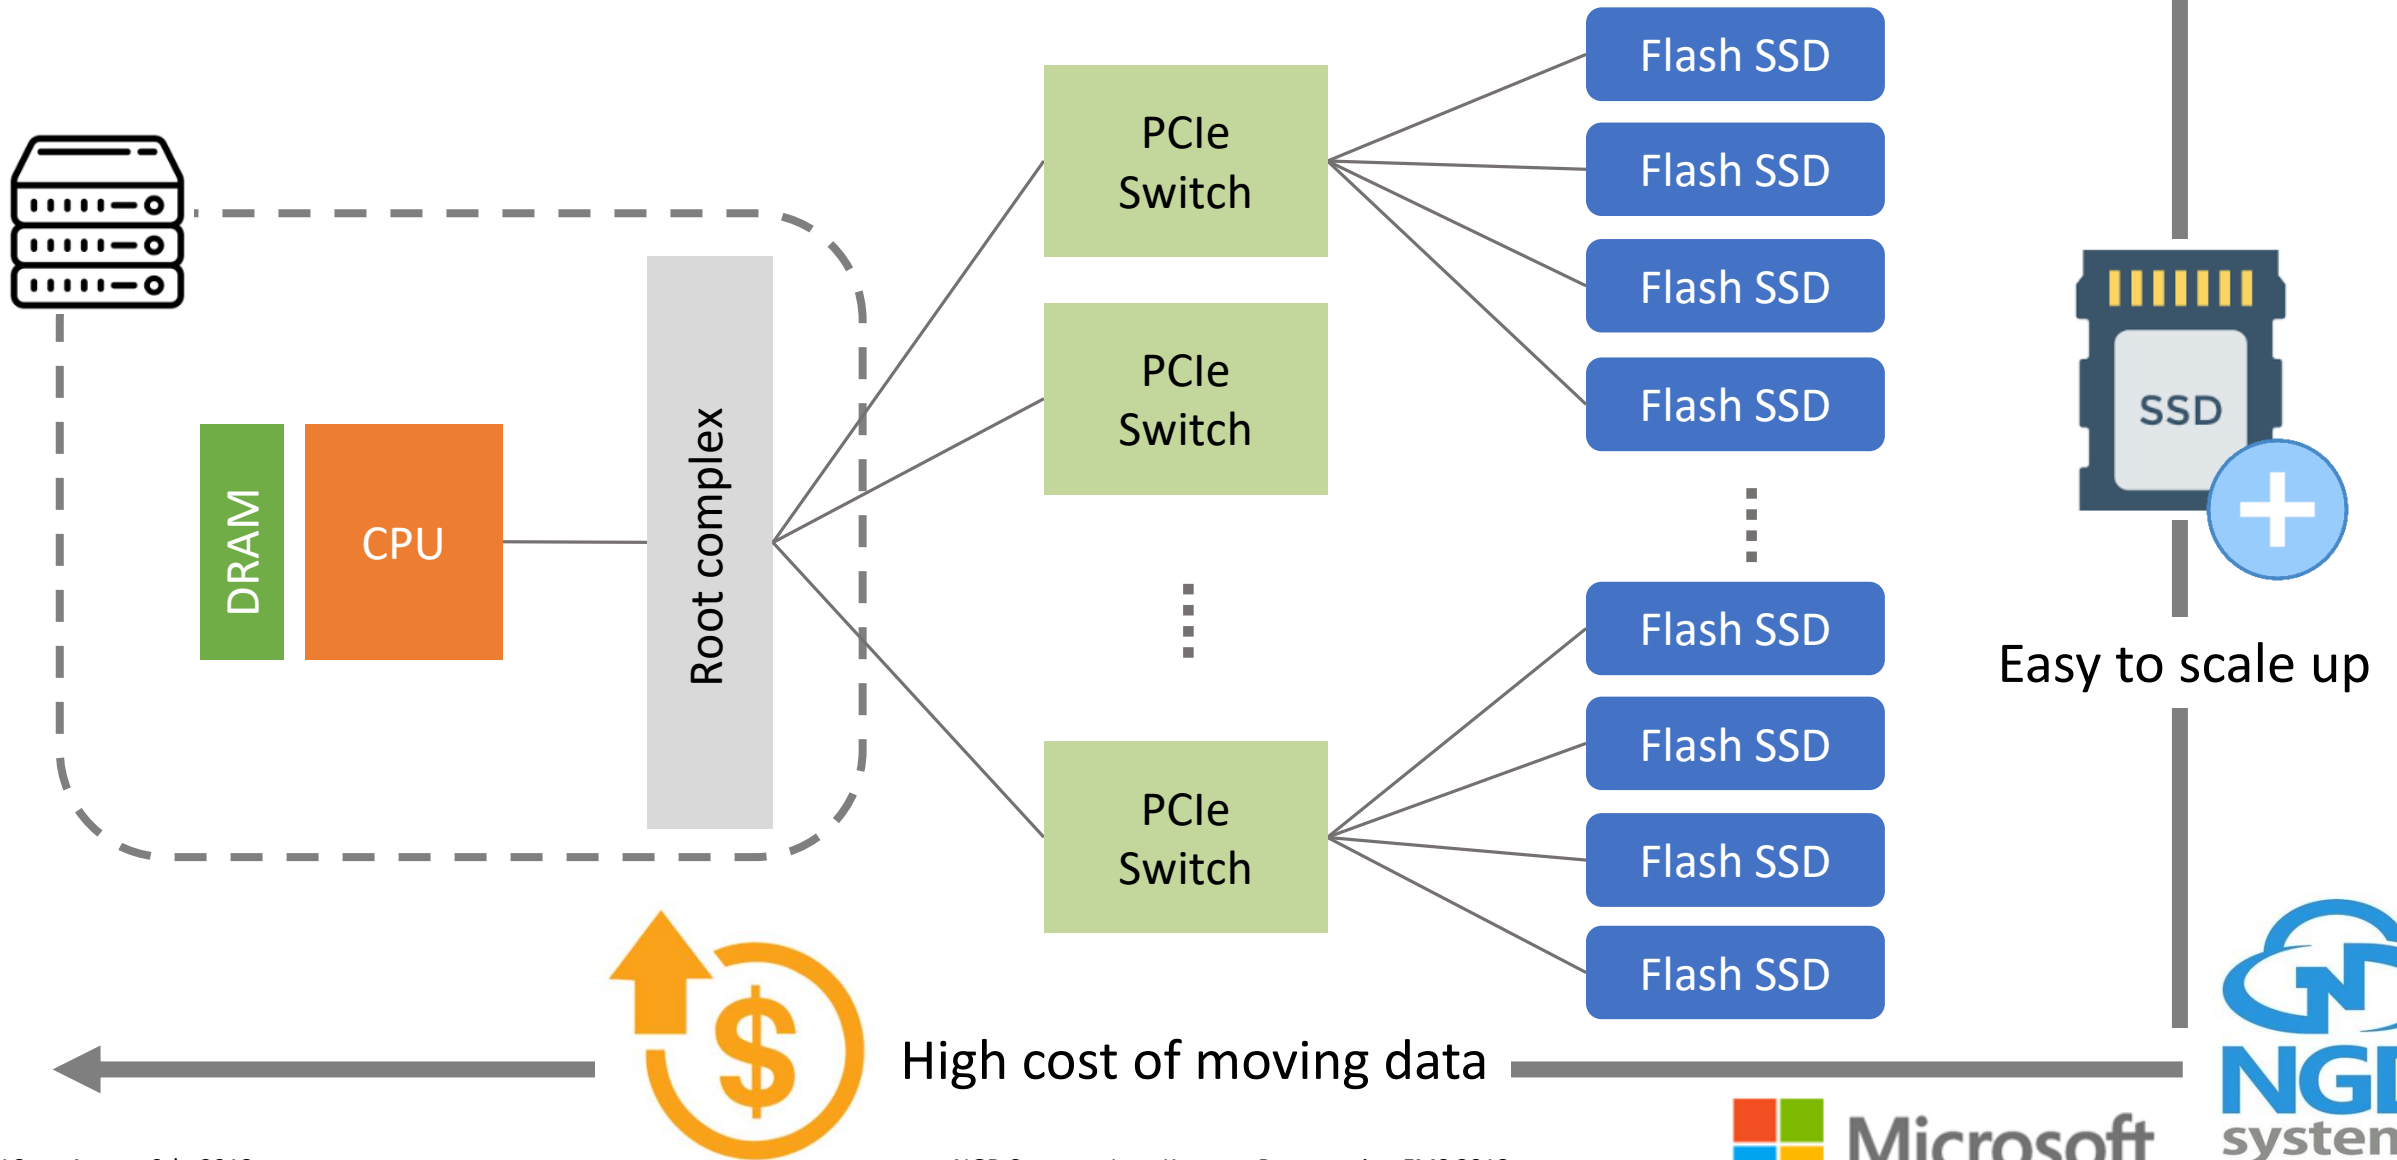


Bringing Intelligence to Storage

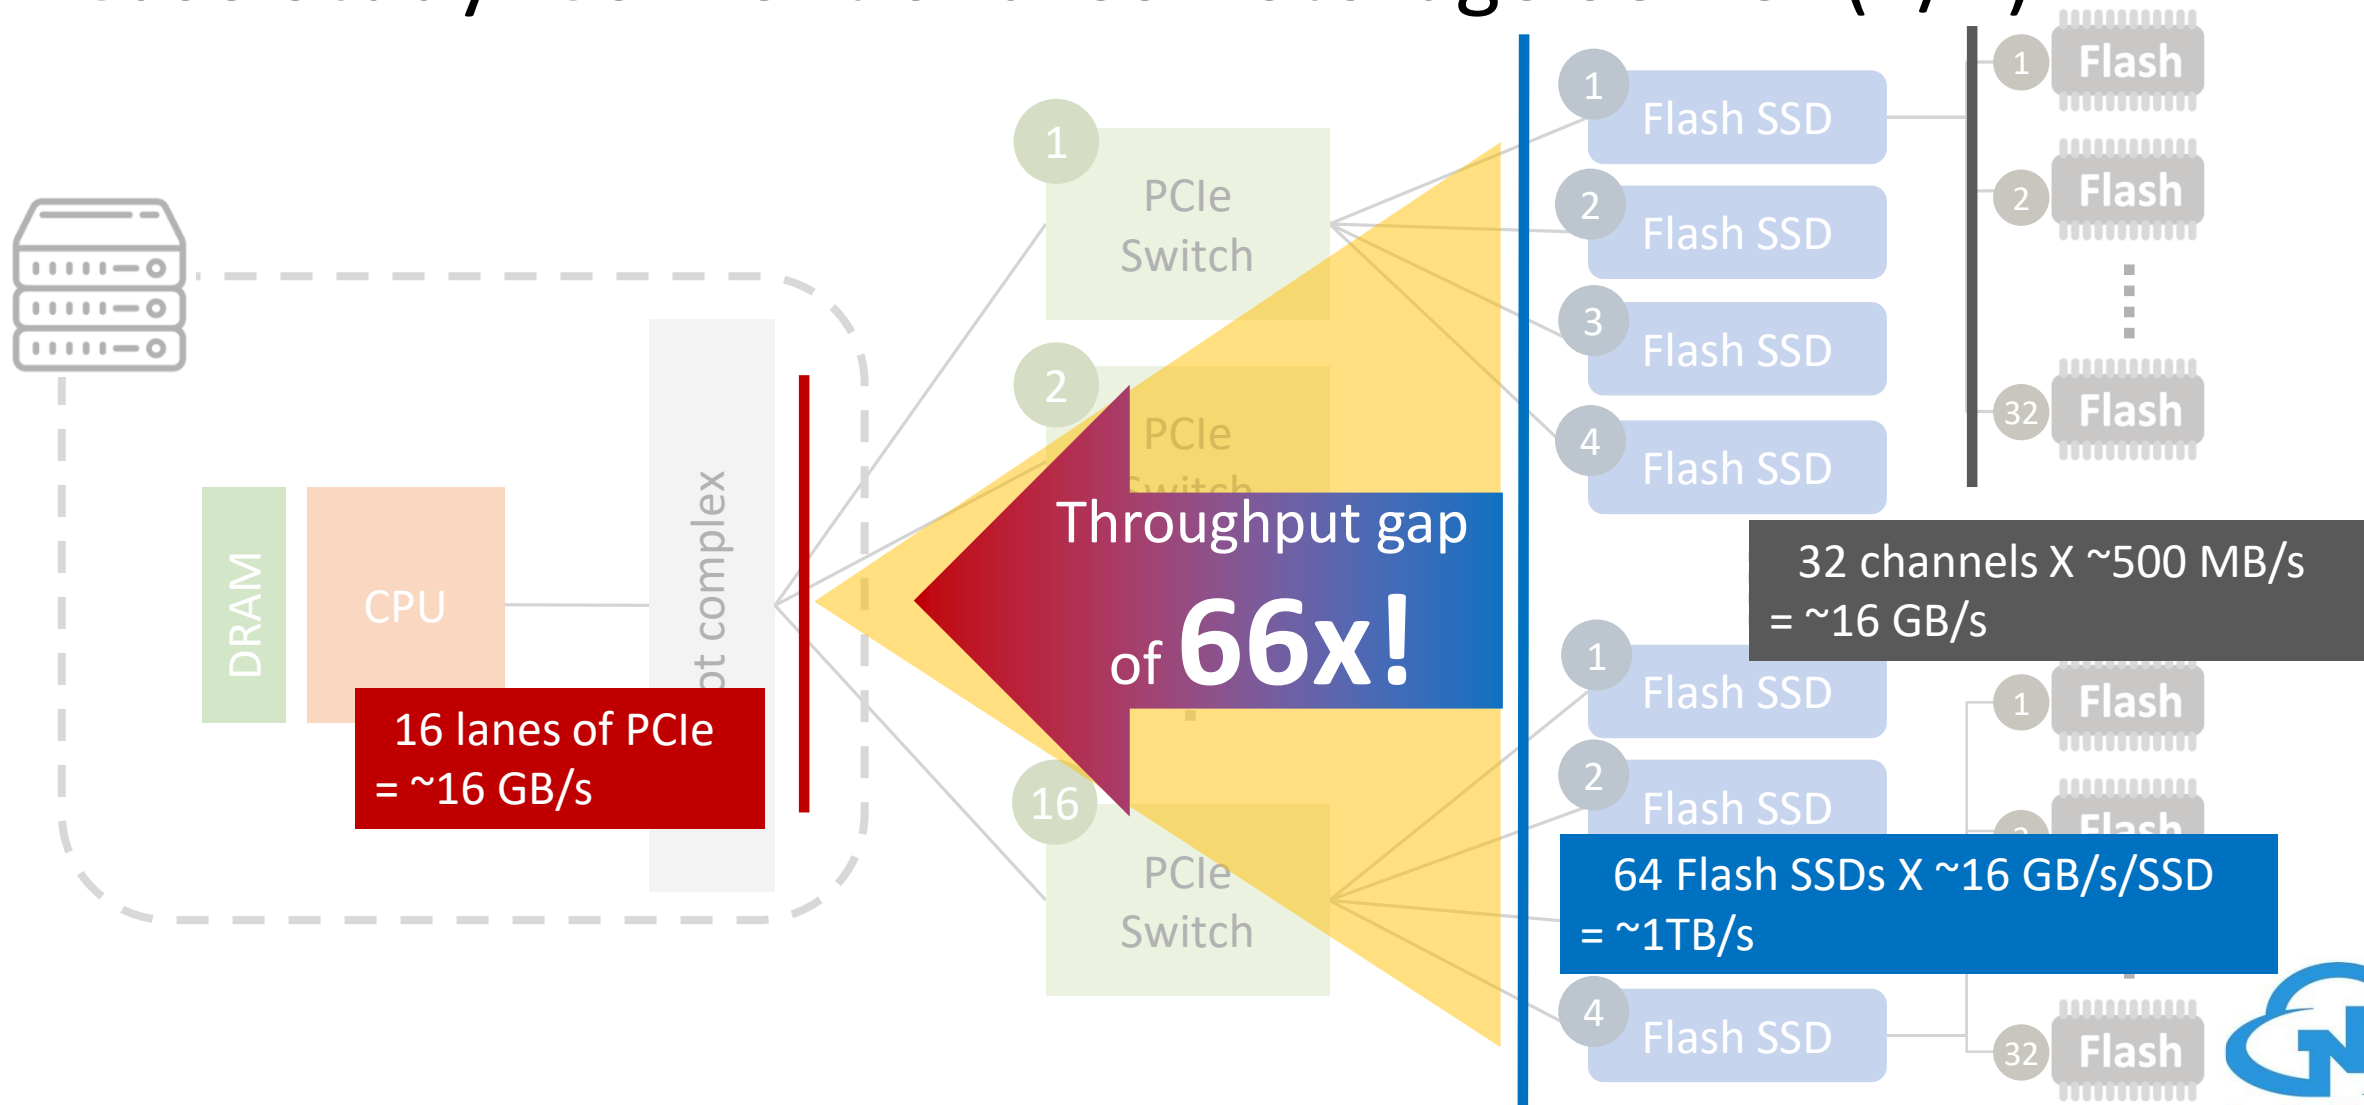
Introducing Jae Young Do – Finding the Right Path Forward



Case Study: Conventional SSD Storage Server (1/2)

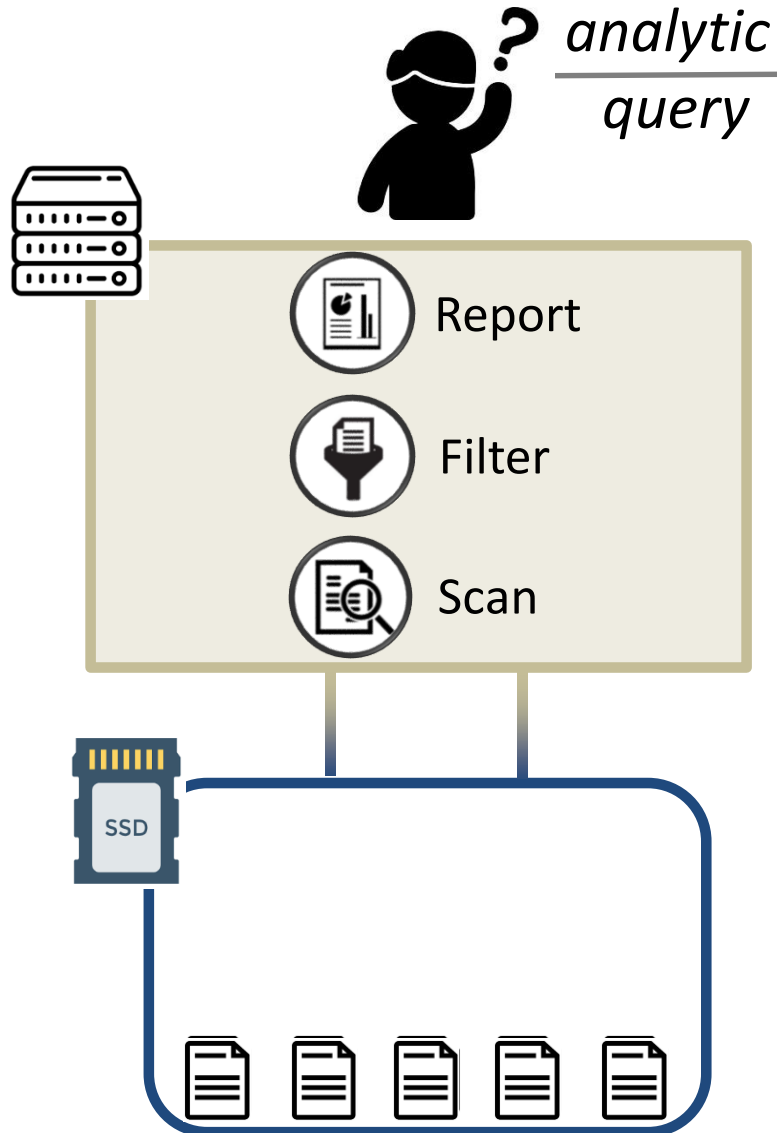


Case Study: Conventional SSD Storage Server (2/2)



Programming Attempts with standard SSDs

5 YEARS AGO



```
SELECT SUM (EXTENDEDPRICE*DISCOUNT)
FROM LINEITEM
WHERE SHIPDATE >= 1994-01-01 AND
SHIPDATE < 1995-01-01 AND
DISCOUNT > 0.05 AND
DISCOUNT < 0.07 AND
QUANTITY < 24
```

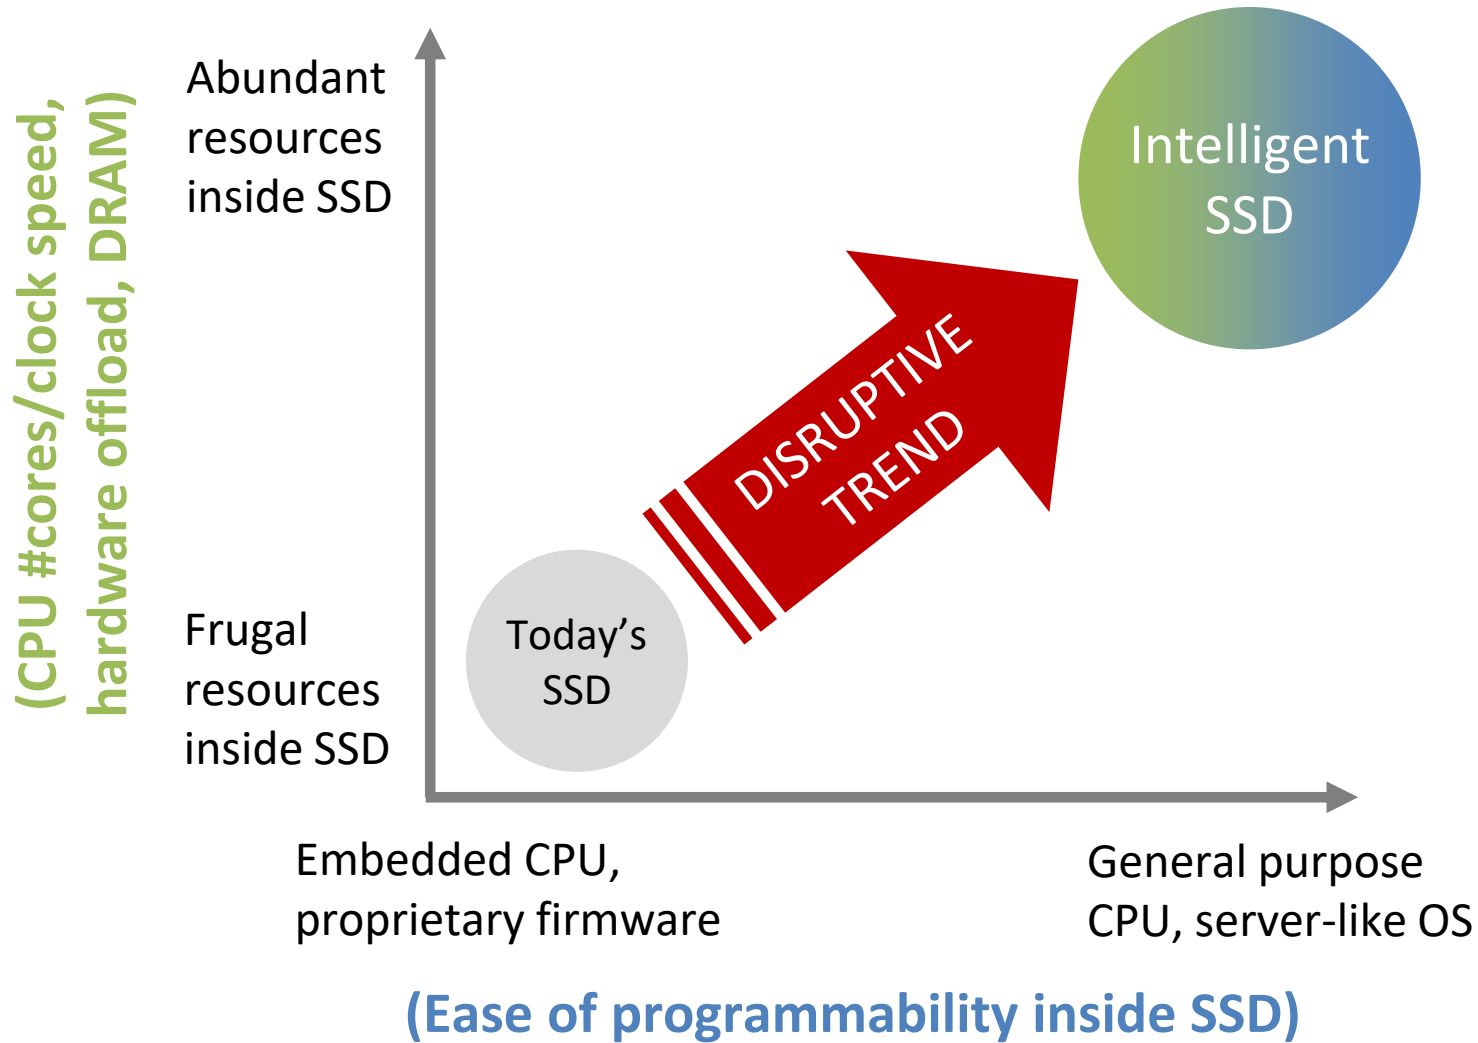
CHALLENGE

Not enough **spare** processing power

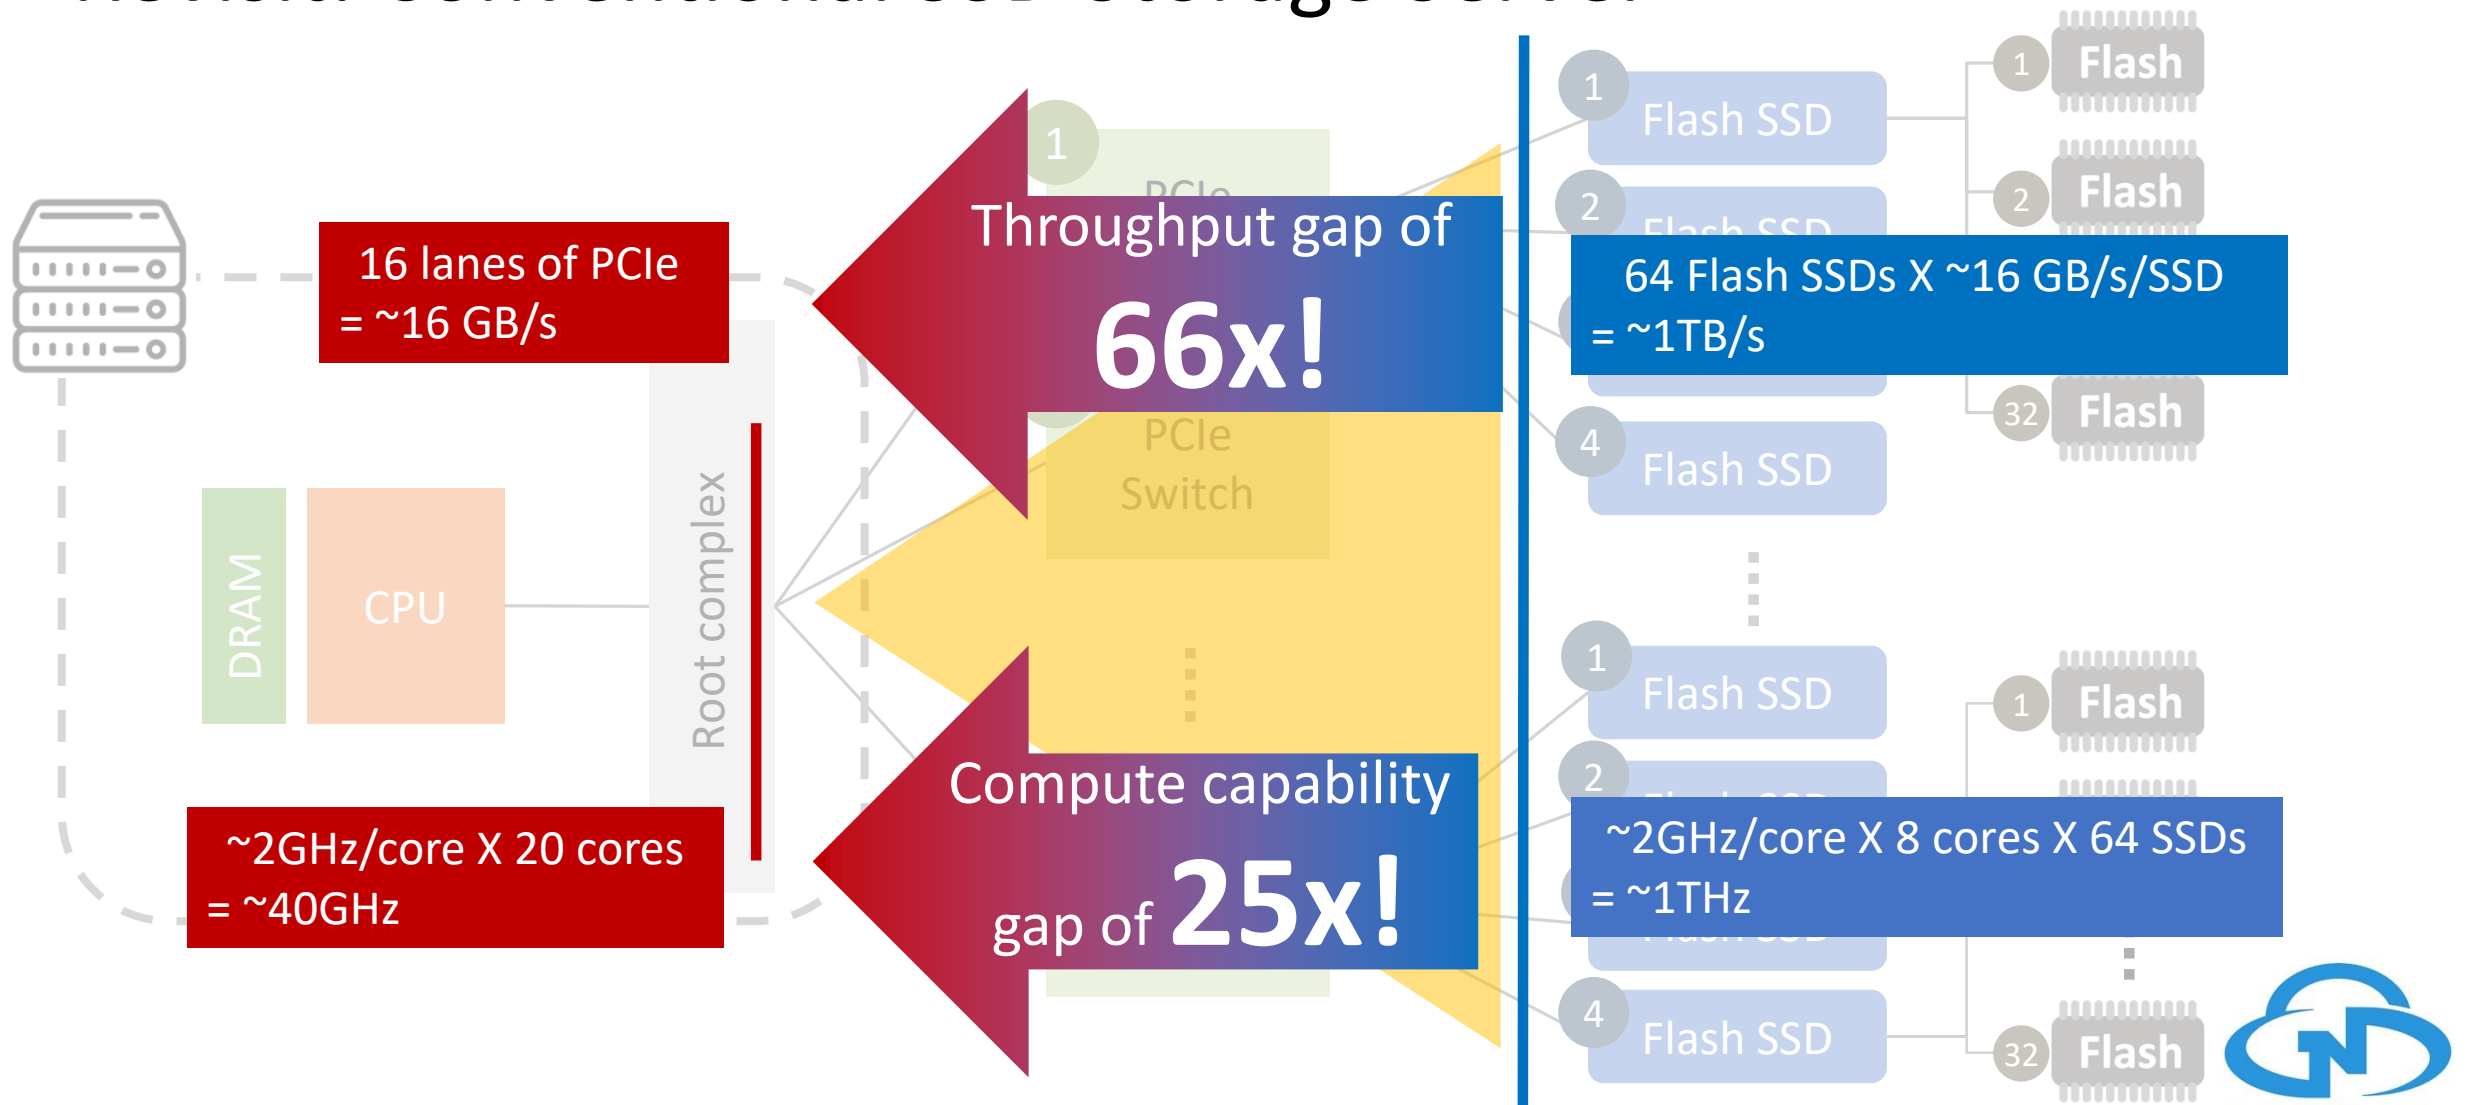
Not **dev-friendly** programming environment

Not **accessible** prototype devices

Disruptive Trends that Enable Intelligent SSDs



Revisit: Conventional SSD Storage Server



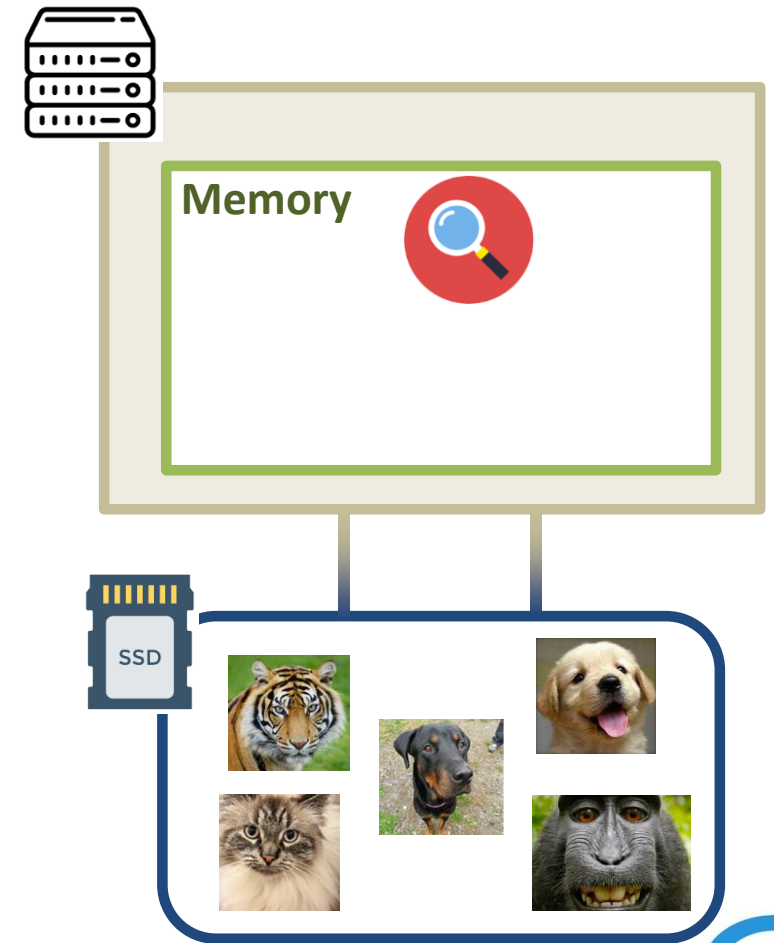
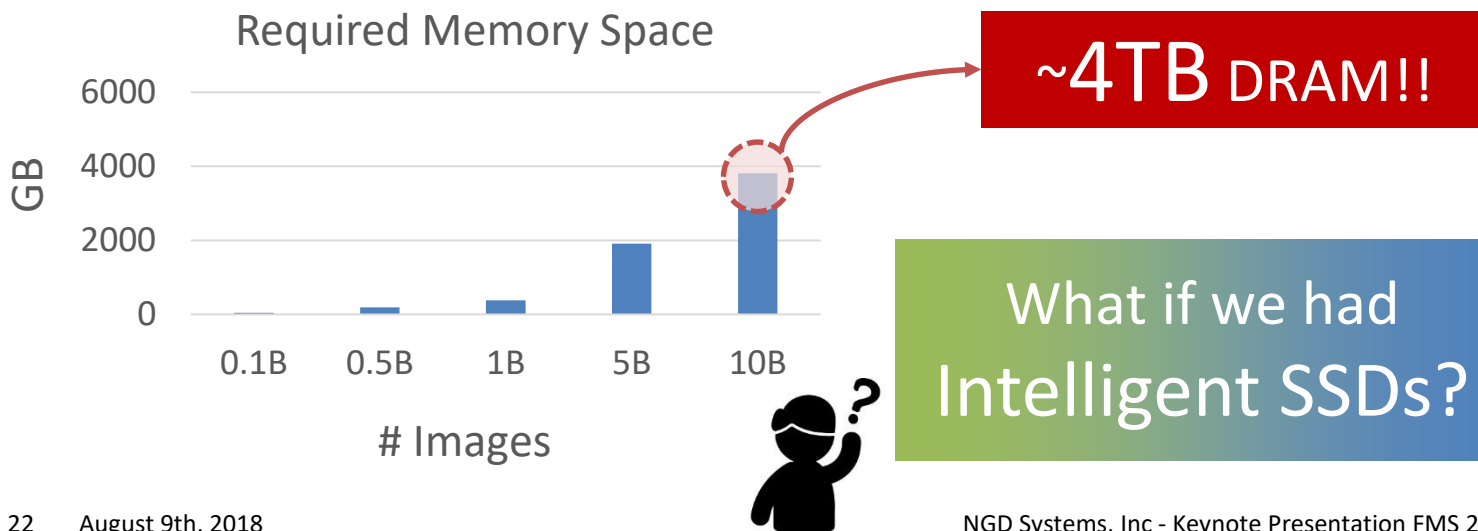
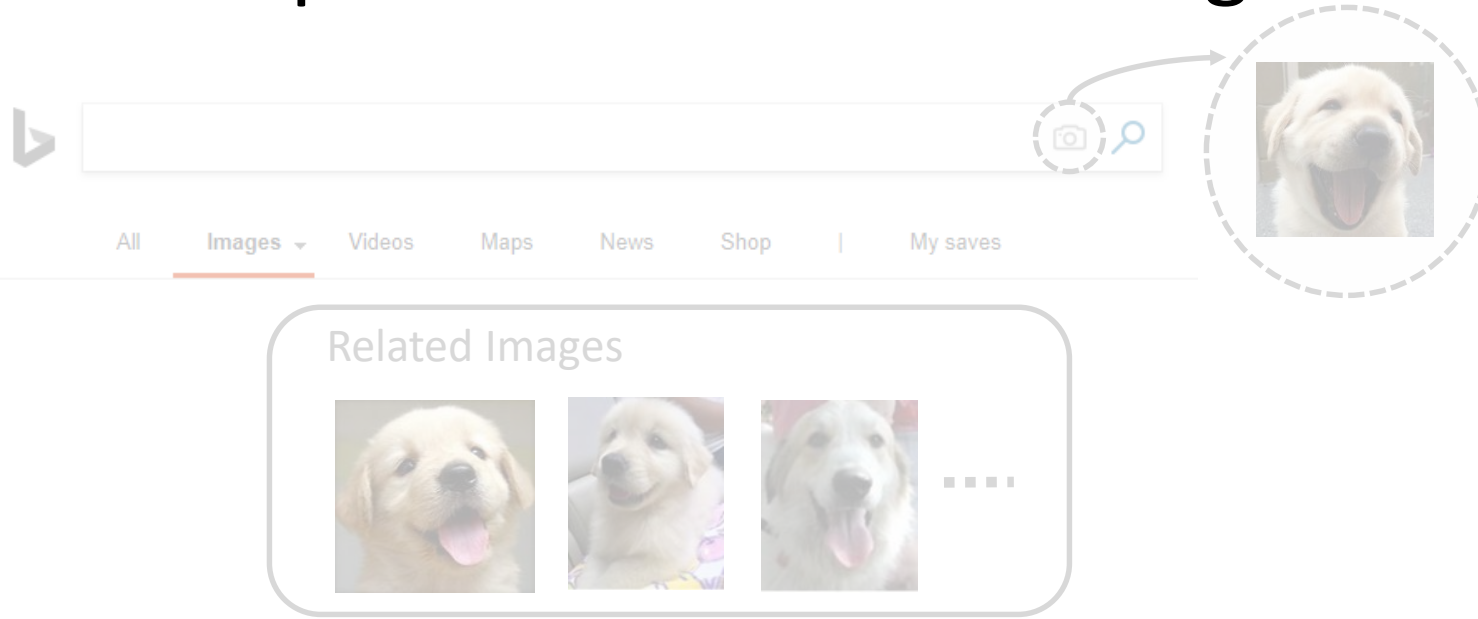
MSR SoftFlash Project




The **SoftFlash** project proposes to create a software-defined storage substrate of flash SSDs in the data center that is as programmable, agile, and flexible as the applications and operating systems accessing it from servers.

- Embrace flash SSDs as a first-class programmable platform in the cloud data center
- Add custom capabilities to storage over time
- Better bridge the gap between application needs and flash media capabilities/limitations
- Innovate in-house at cloud speed

Example Scenario with Intelligent SSDs – Image Search



Preliminary Results: Image Query Throughput



NOTE

The I/O cost of moving data is **NOT** considered!
More interesting results are ...



Key Takeaways

- Finding the Needle Faster

IN-SITU PROCESSING

- Bigger Pipes Feed Smaller Ones



- Smarter Storage Does Work



- Requires Intelligent Controllers



- Power is Factor - Always

Watts/Terabyte

NGD Systems Newport Platform Provides

On Drive Linux OS, Container Support
Dedicated Compute Cores

Mitigating Data Movement
Optimizing Application Execution

Partnerships for Success
Real World Implementation

Flash Agnostic – ONFI/Toggle, TLC/QLC
16 Channels - Capacities to 256TB

.35 W/TB @ 16TB





Flash Memory Summit
BEST OF SHOW WINNER



Thank You



Info@NGDSystems.com
www.NGDSystems.com

CIO 10 MOST PROMISING
STORAGE
Review SOLUTION PROVIDERS - 2018