SNIA COMPUTE + MEMORY + STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023

Deploying SSDs with Computational Storage: The Promise & Reality

JB Baker
VP Product Management & Marketing
ScaleFlux

Abstract

- We've been hearing about Computational Storage for several years now. But, how is it playing out in the field?
  - What applications benefit from CS?
  - How are users deploying CS?
  - What challenges, changes, and integrations do they deal with to reap the rewards from CS?
- Join us to hear about a few examples of Computational Storage Drive deployments and use cases to improve efficiency, TCO, and uptime.

Computational Storage Expectations

- What is CS trying to achieve?

Computational Storage – Goals, Tactics, & Benefits

- GOALS
  - Alleviate Bottlenecks (CPU, Memory, Network)
  - Enable Parallel Computation
- TACTICS
  - Reduce data movement
  - Offload host processing
- BENEFITS
  - Increase Application Performance
  - Improve Infrastructure Efficiency (Power, TCO, Sustainability)

How are we doing against the goals?

- Example deployments

Hybrid Transactional Analytical Processing (HTAP)

User: Alibaba (cloud services)

Application: PolarDB (in-house)

Pain Points

- Network congestion
- Data processing hot spots
- Memory congestion

Solution

- Predicate filtering in the CSD
- Application integration

Benefits

- ✓ 50%+ reduction in PCIe traffic on TPC-H queries
  - 90%+ reduction for Q6 & Q14
- ✓ 70%+ reduction in Network traffic for TPC-H queries
- ✓ Up to 10x Performance gains

Online Transactional Database (OLTP)

User: Cloud Service Provider

Application: MySQL

Pain Points

- Delivering on their SLA
- Infrastructure costs
- Infrastructure reliability

Solution

- CSDs with Transparent Compression

Benefits

- ✓ Deliver the same SLA with 50% fewer servers and 75% fewer SSDs
- ✓ 50%+ Improved TCO
- ✓ 4x Reliability increase

Stream Data Capture

User: Financial Services

Application: Market data capture for High-Frequency Trading (HFT)

Pain Points

- Volume & rate of data capture
  - Throughput & latency
- Data storage bandwidth & capacity
- Flash Endurance

Solution

- CSDs with transparent compression

Benefits

- ✓ Store 2x data per GB of Flash
- ✓ Increase “max fill” threshold to 90%
- ✓ 2x+ increase in lifetime PBW per GB of Flash
- ✓ Reduced latency

Multiple apps…help us!!

User: Online Payments & Services

Application: Multiple databases involved in payment processing & fraud detection

Pain Points

- **HBase**: Shrinking window between archive/back-up offloads
  - Was months, now days
- **MySQL**: Cost of scaling capacity
- **Aerospike**: Hitting latency goals for rapidly growing transaction volume

Solution

- CSDs with transparent compression

Benefits

- ✓ HBase: Expand the window with 2x data storage per GB of SSD
- ✓ MySQL: 50% Lower Cost/GB of data
- ✓ Aerospike: Increase TPS per server

CS – How are we doing? Report Card

- GOALS
  - Alleviate Bottlenecks (CPU, Memory, Network) A-
  - Enable Parallel Computation A-
- TACTICS
  - Reduce data movement B
  - Offload host processing A-
- BENEFITS
  - Increase Application Performance A-
  - Improve Infrastructure Efficiency (Power, TCO, Sustainability) A

Challenges & Barriers to Adoption

- With all those “A” grades, why isn’t Computational Storage everywhere?

Challenges & Barriers to CS Ubiquity

- What's missing from the report card?
  - Real concerns & perceived issues with ***Complexity*** *(more in a minute)*
  - Standardization *(being addressed)*
  - Security *(being worked)*
  - Aligning the ecosystem
    - Who "does the work" vs who "gets the benefit"
    - Is there a big enough market for this function?

The old “cost-benefit” challenge

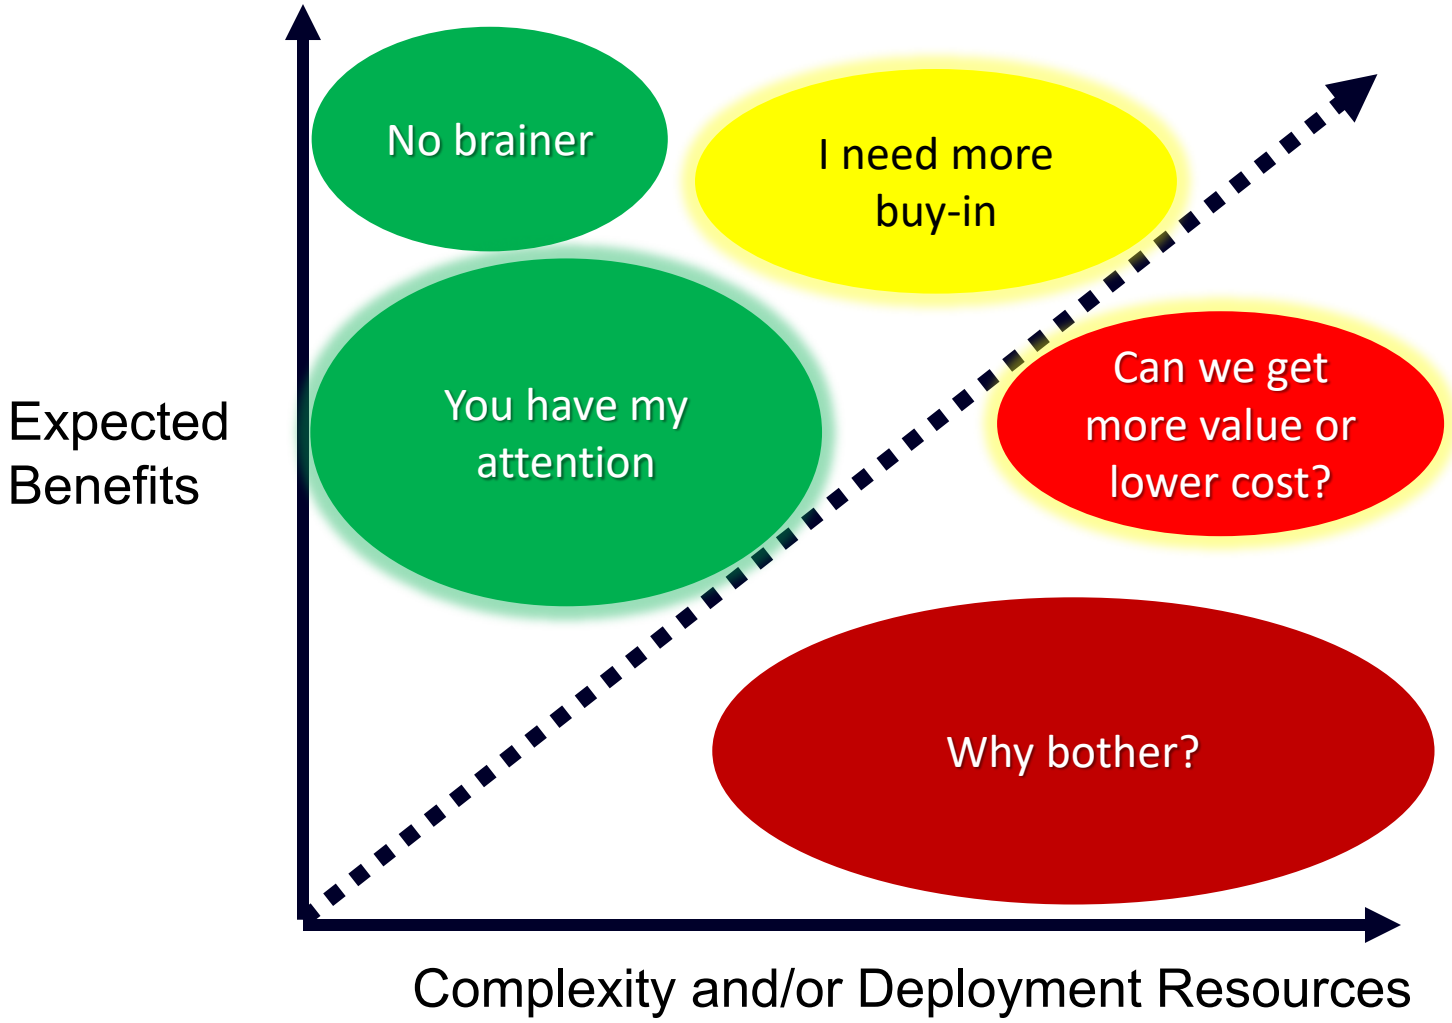

- Low complexity, good benefits
  - Transparent functions
  - No or minimal integration to applications or management infrastructure
  - Broad applicability
  - May require some change in procurement or user behavior
- High complexity, high benefits
  - Customized/co-developed
  - Lots of work to get approved
  - More factors in the benefit calc

Wrapping it up…

- Computational Storage has started delivering on its promises
  - Standardization efforts are helping
- But challenges are yet to be overcome…
  - Resources or skillsets to develop and deploy customized functions
  - “Peeling the onion” on TCO and Security

SNIA COMPUTE + MEMORY + STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023

Please take a moment to rate this session.

Your feedback is important to us.

