Accelerating Near Real-Time Algorithms Using Disaggregated Computational Storage

(High Performance Object Storage - HPOS)

Mayank Saxena,
Memory Solution Labs, Samsung

HPOS S3 SELECT w/ Computational SSD Acceleration

High Performance Object Storage (HPOS)

Use Cases

Large scale real time analytics - Smart City, Smart Home, eHealth, IoT, Images, Video, Security

Benefits

- **Faster** queries
- **Lesser** network traffic
- **Lower TCO** due to reduced CPU and network traffic

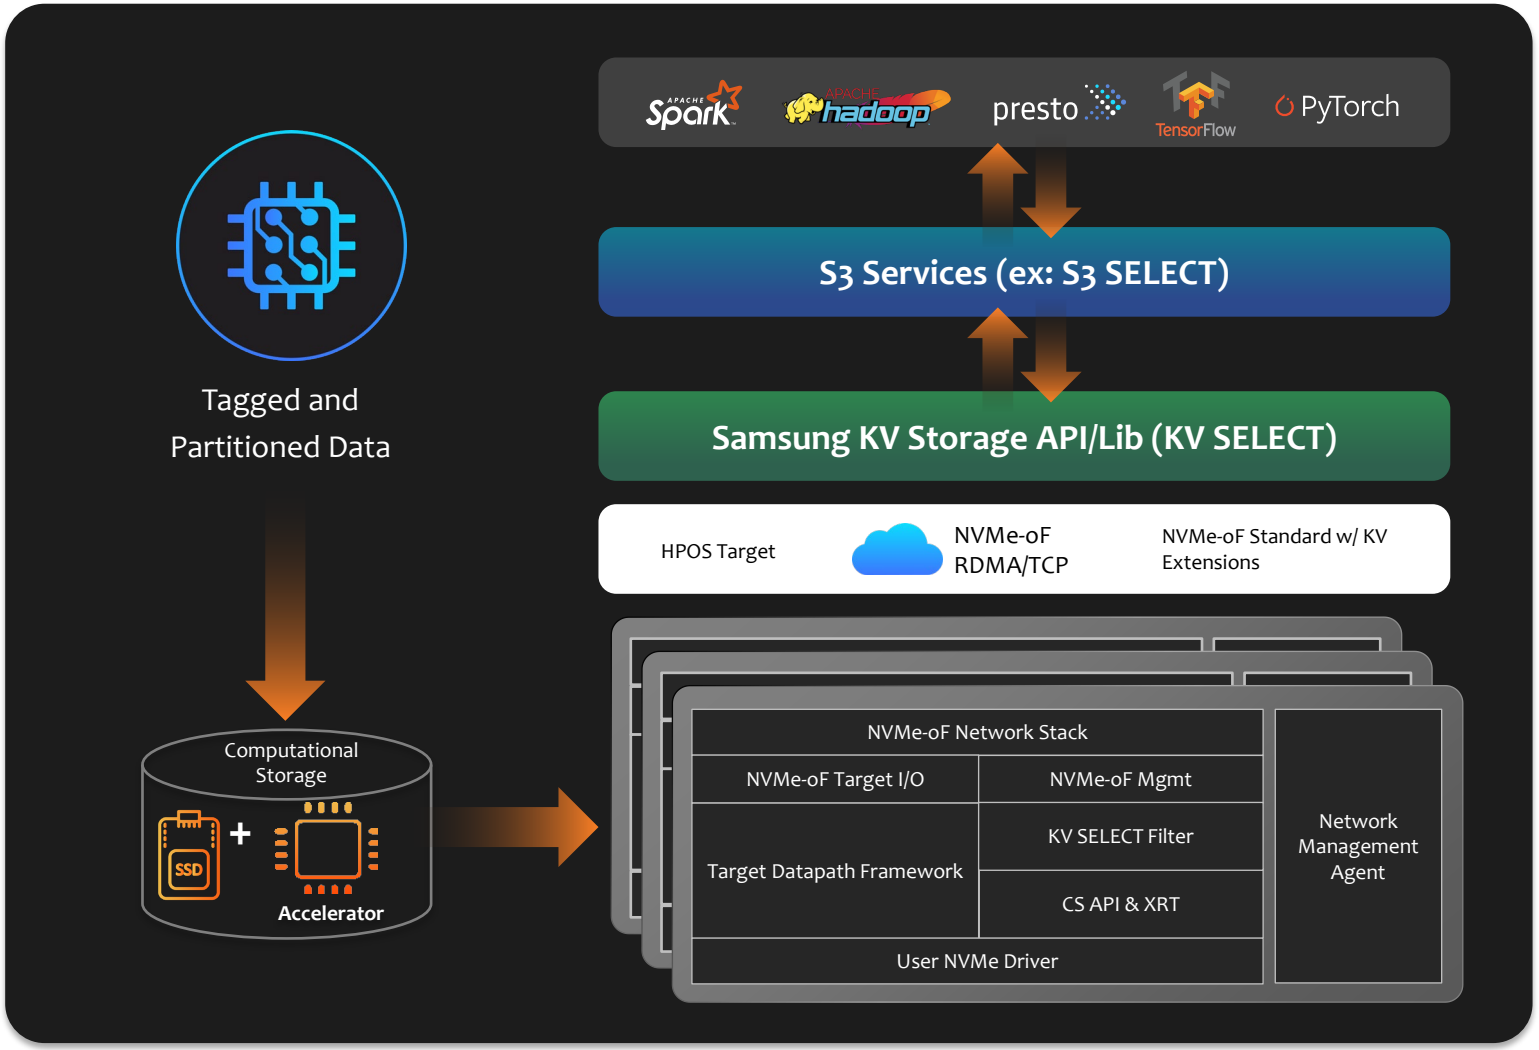

End-to-End S3 Select with HPOS

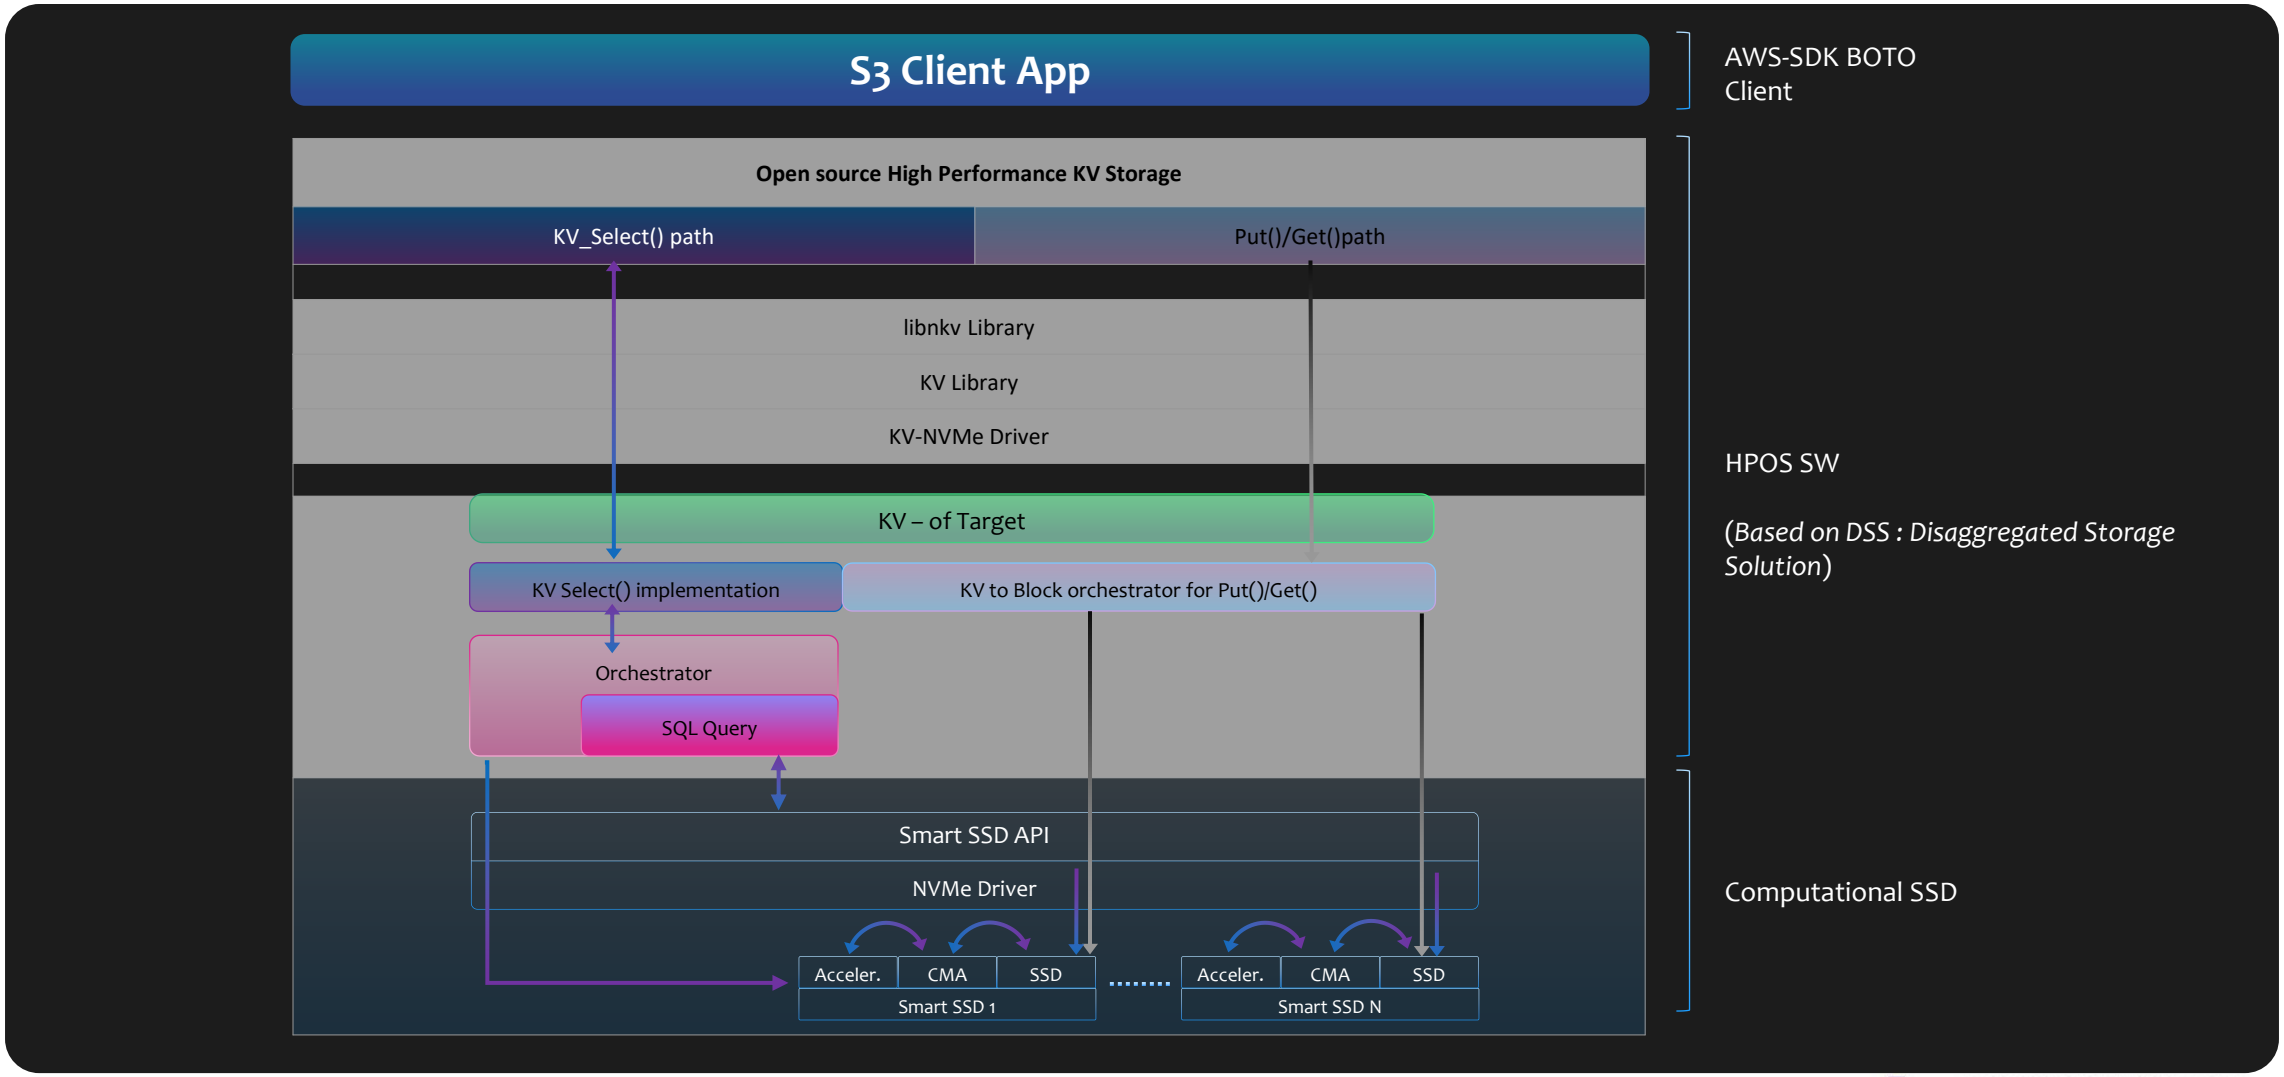

Normal data path Path for SmartSSD offload CMA Computational Memory Area

Near data processing for S3_Select

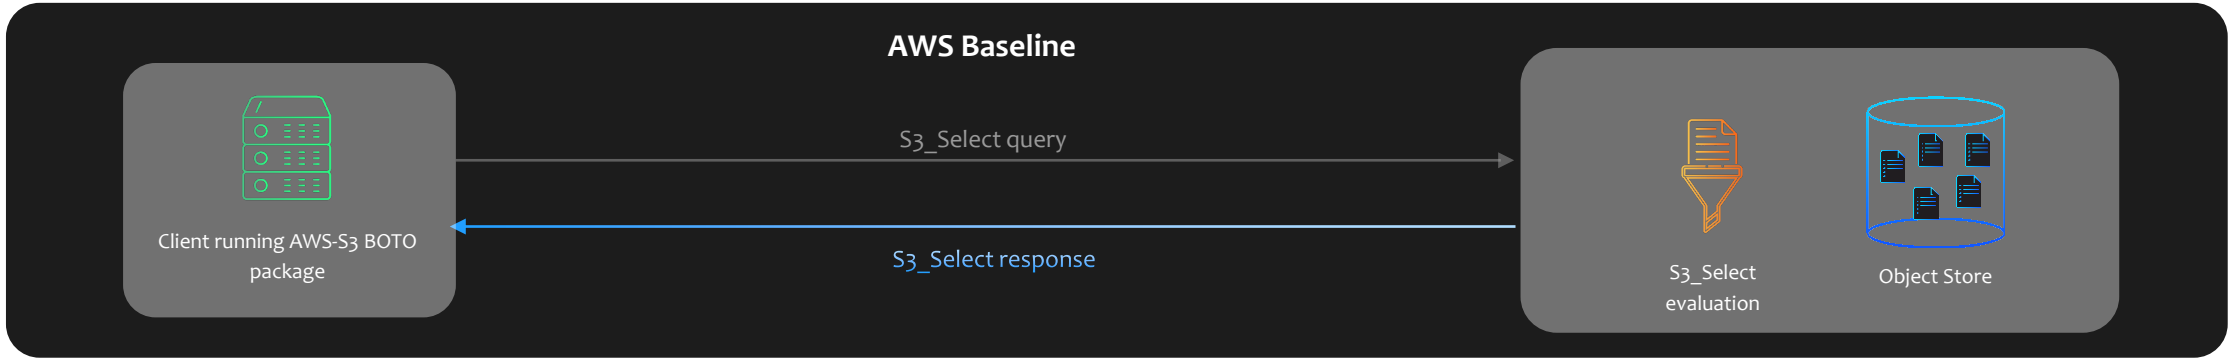

Customer's Sample Dataset

Index	Year	DayOfYear	Precipitation	ET0	tMax	tMin	RS
0	1	1	0	0	-26.6808	-29.7293	0
1	1	2	0	0	-27.0202	-30.9575	0
2	1	3	0	0	-28.0491	-31.0459	0
3	1	4	0	0	-26.2893	-29.3781	0
4	1	5	0	0	-25.1941	-27.9265	0
5	1	6	0	0	-23.5173	-27.2431	0
6	1	7	0	0	-23.8696	-29.4691	0
7	1	8	0	0	-24.4267	-29.0663	0

S3_Select Query ('Select COUNT(*) FROM S3Object s WHERE dayofyear < 55');

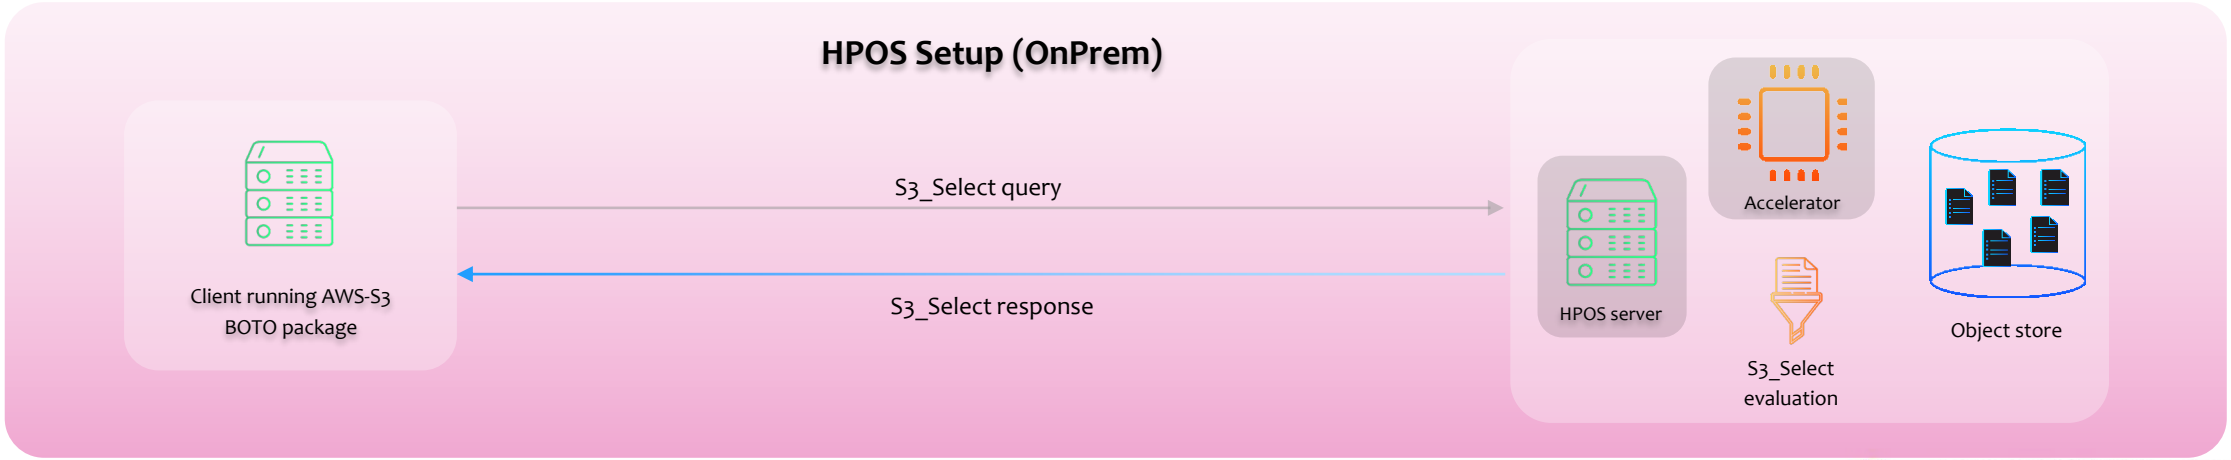

HPOS S3_Select Drag Race Demo

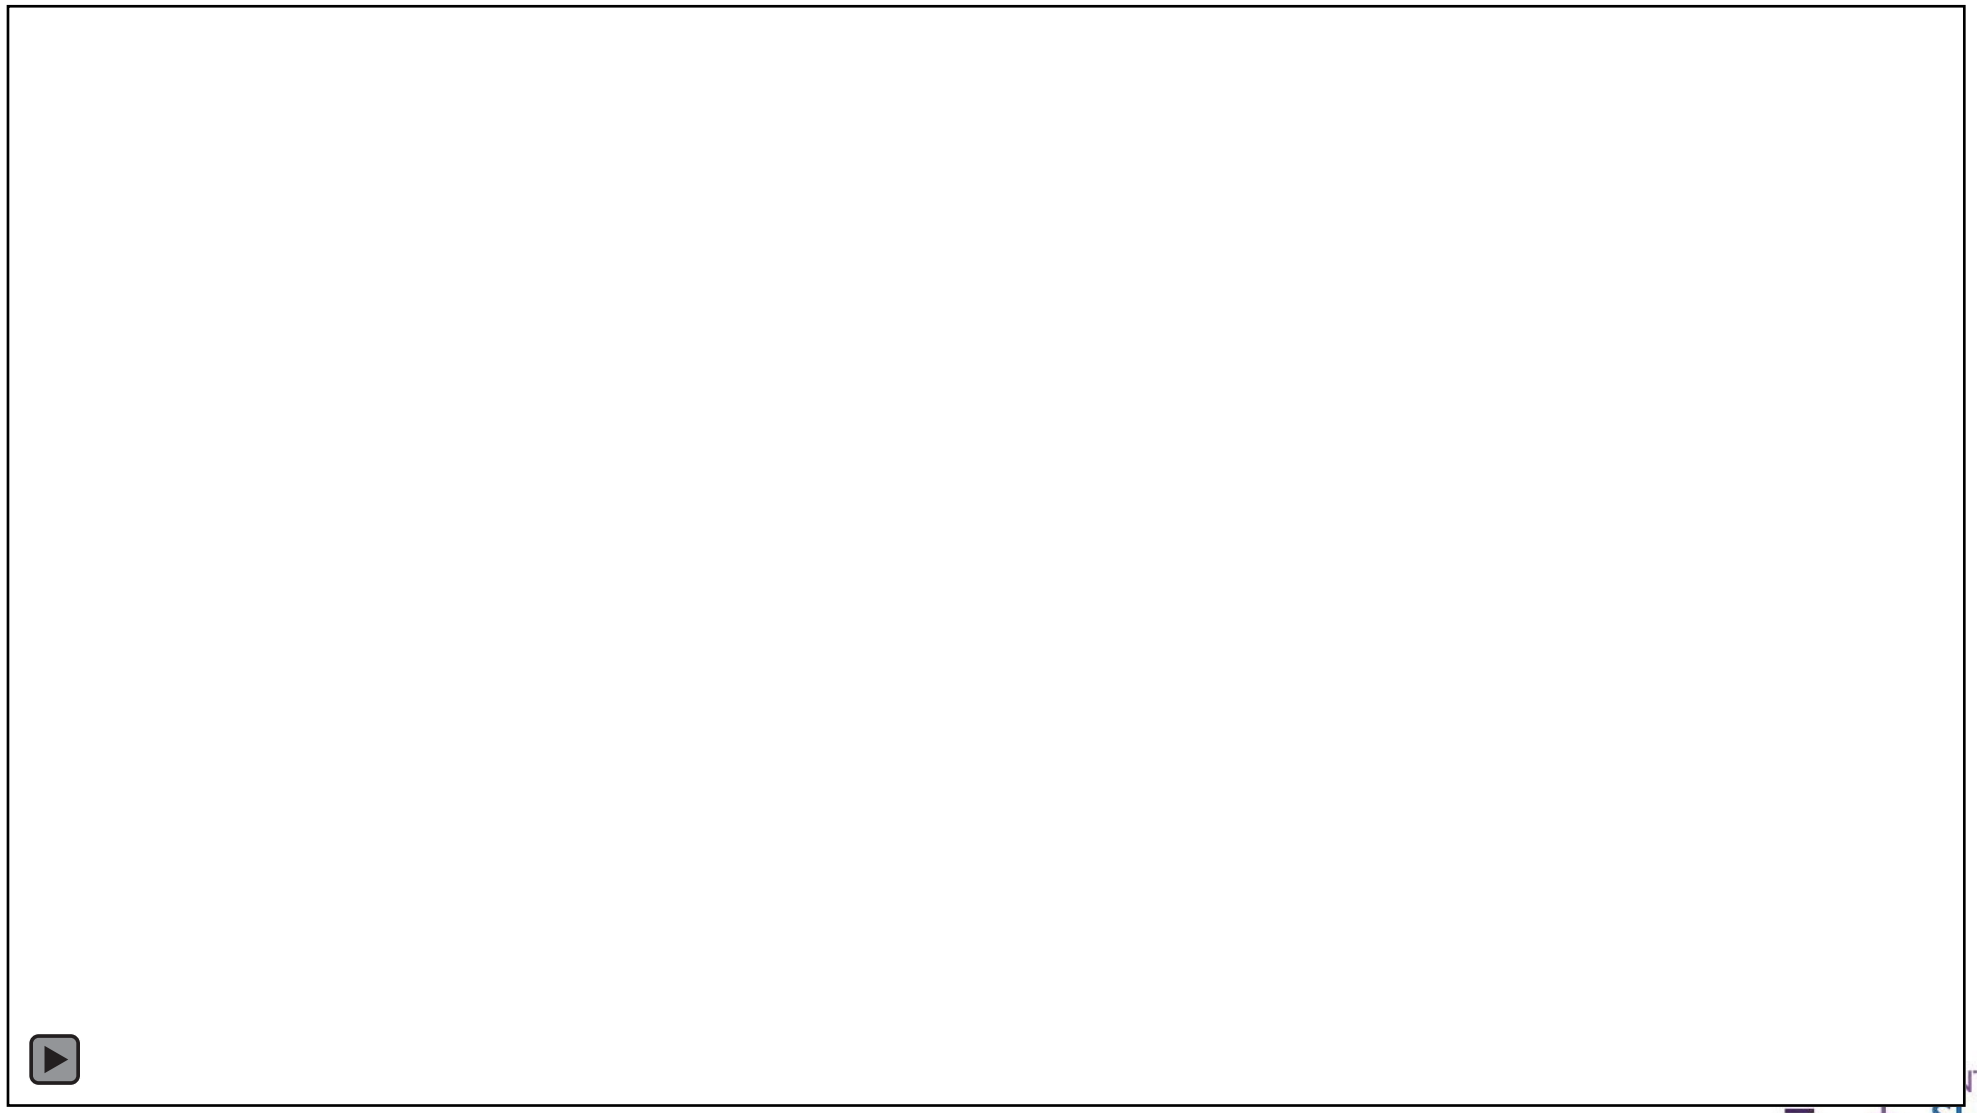

AWS S3 Select vs HPOS S3 Select

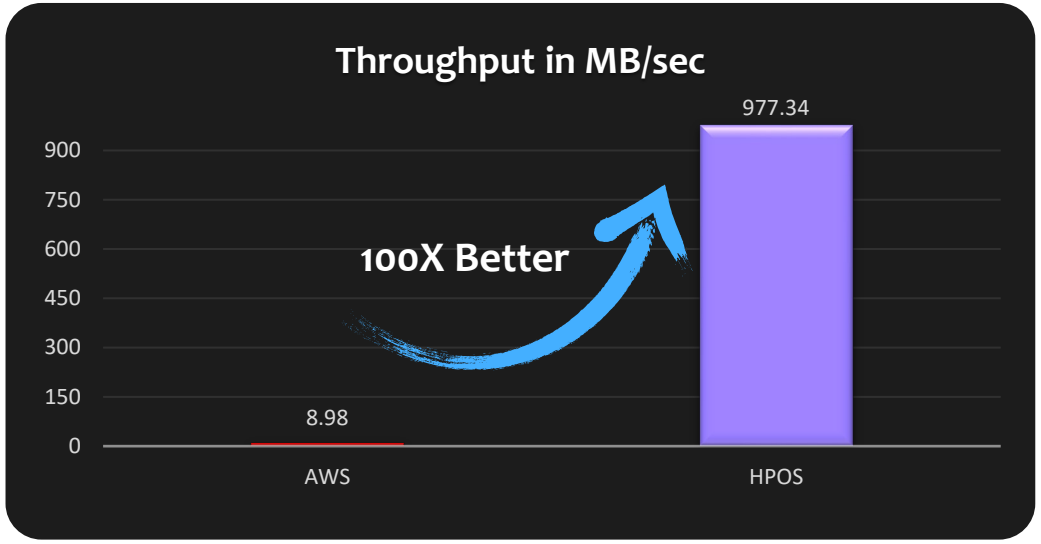

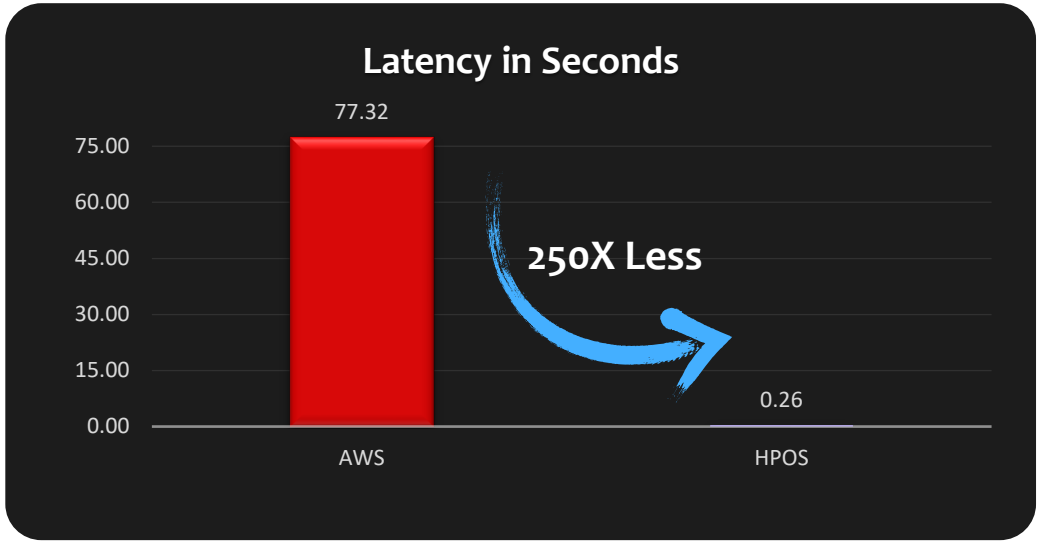

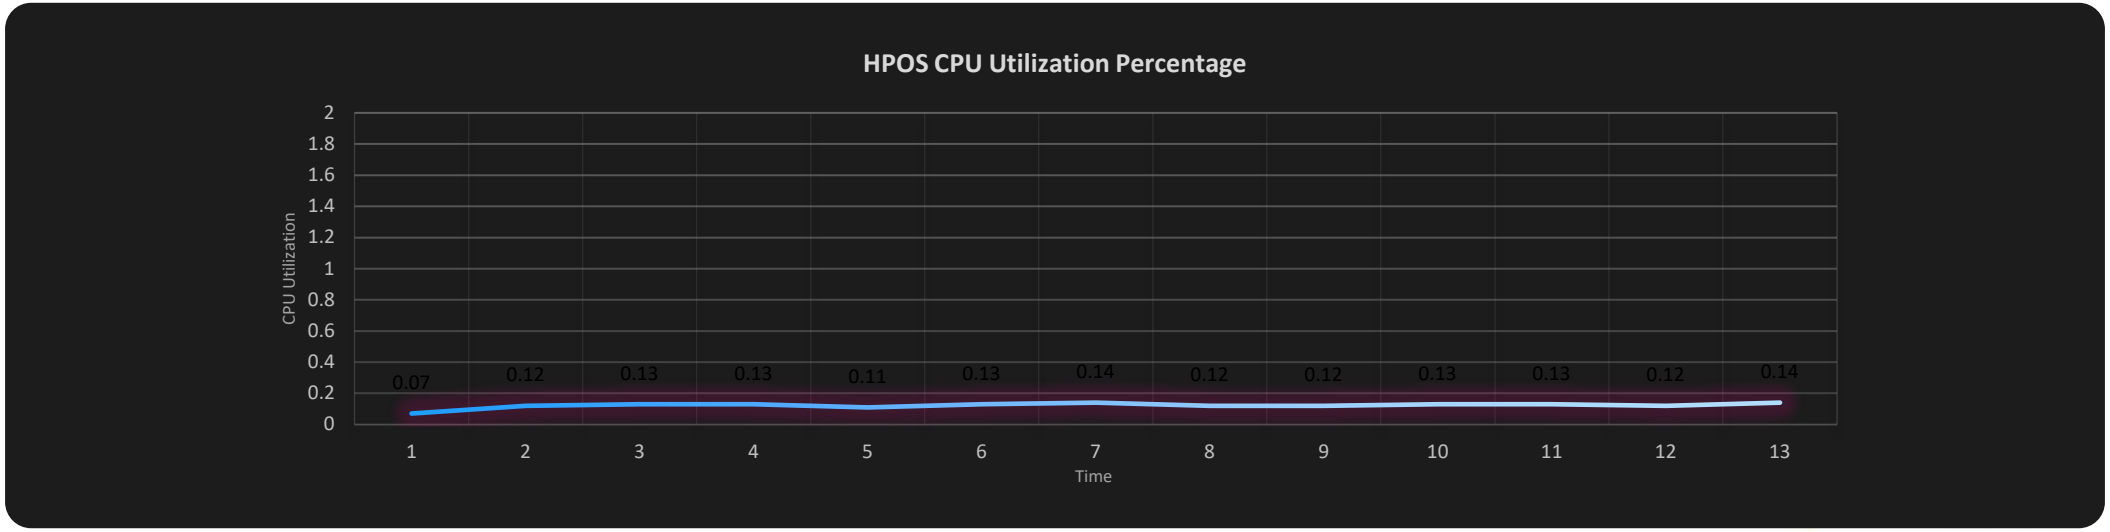

TCO - Power Savings Analysis

Observations

- On average there is ~6 Watts/hr of power saving using **single** Smart SSD
- Amounts to ~2 % of total server power usage

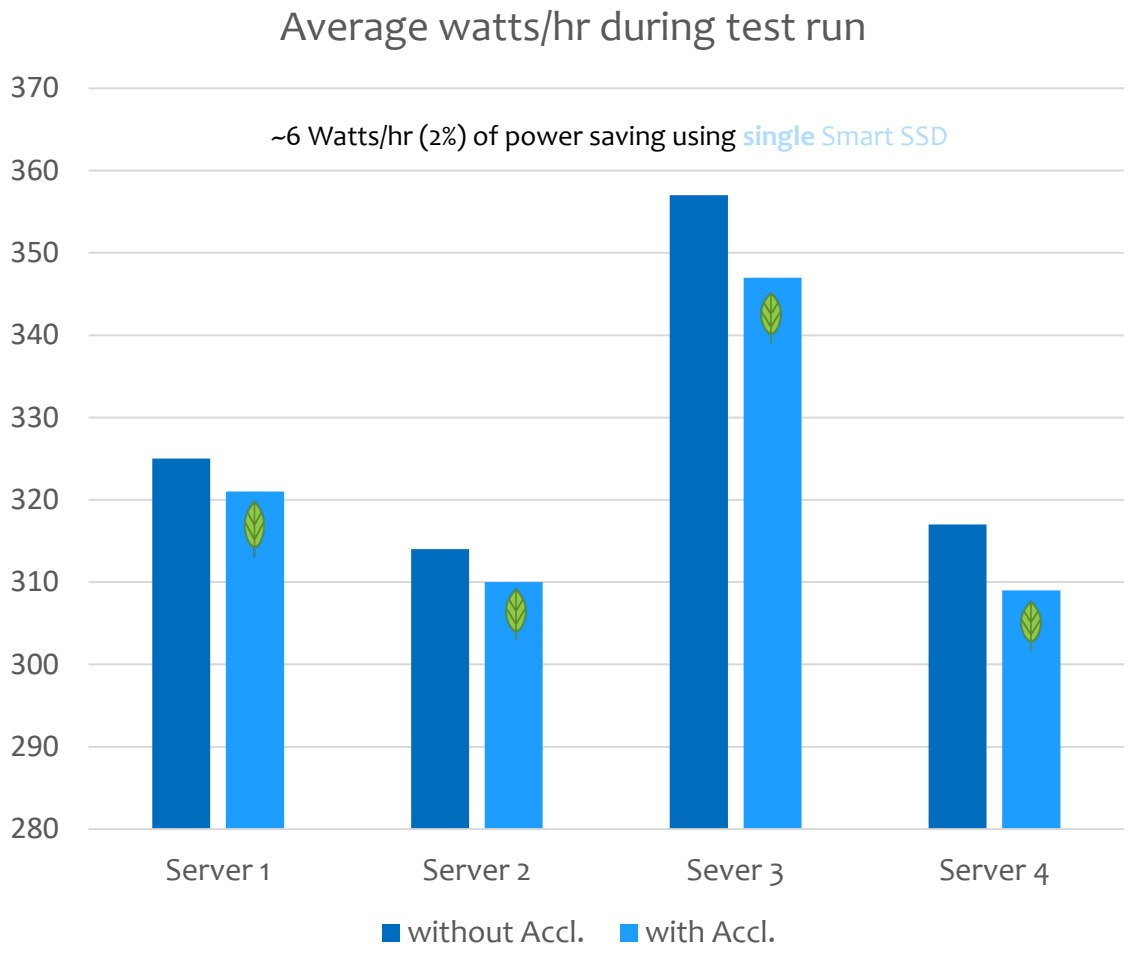

Open Source Availability

Disaggregated Storage Solution (DSS)

- Foundational Software for HPOS
- SNIA PM+CS Summit talk – "Storage for a new Generation of AI/ML" by Somnath Roy

Open Source'd

- https://github.com/OpenMPDK/DSS
- HPOS + CS specific modules in future

Complete Ecosystem

- AI Benchmarking Framework supporting user preferred training and models
- Client Wrappers supporting Pytorch and Tensorflow
- Host and Target Stack

Thank You

Mayank Saxena
(mayank.s4@samsung.com)

Please take a moment to rate this session.

Your feedback is important to us.

