

STORAGE DEVELOPER CONFERENCE



Fremont, CA
September 12-15, 2022

BY Developers FOR Developers

A **SNIA** Event

Green Computing with Computational Storage Devices

Changho Choi, PhD

Senior Director of Engineering
Memory Solutions Lab
Samsung Semiconductor Inc.

September 14th 2022

Yangwook Kang, PhD

Senior Staff Engineer
Memory Solutions Lab
Samsung Semiconductor Inc.

September 14th 2022

Agenda

- **Industry trend**
 - Power hungry data center
 - Domain Specific Architecture
- **Computational Storage**
 - Samsung SmartSSD[®]
- **DB acceleration system architecture**
- **The 2nd Gen. SmartSSD[®] benefit**
- **Summary**

Industry Trend

- Power Hungry Datacenter

Energy Scale

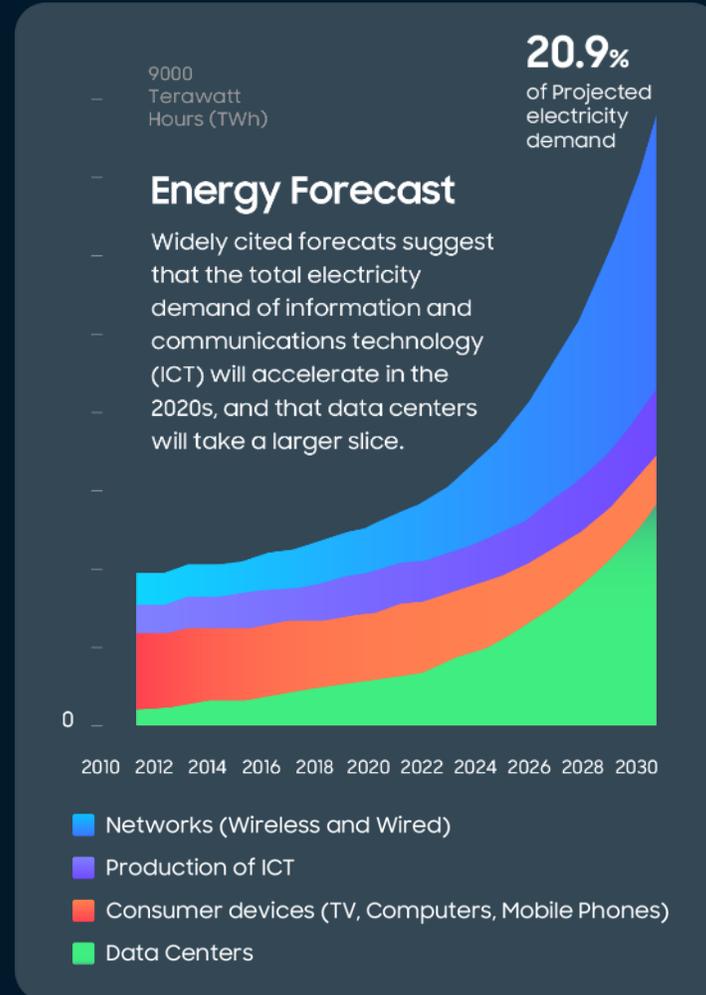
20,000 TWh
Global Electricity Demand

2,000 TWh
Electricity use by ICT

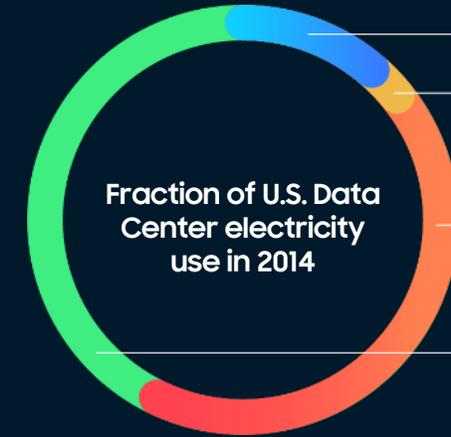
200 TWh
Data Center Electricity Demand

20 TWh
Bitcoin use by mid 2018

Source: nature, morganclaypool



Source: akcp



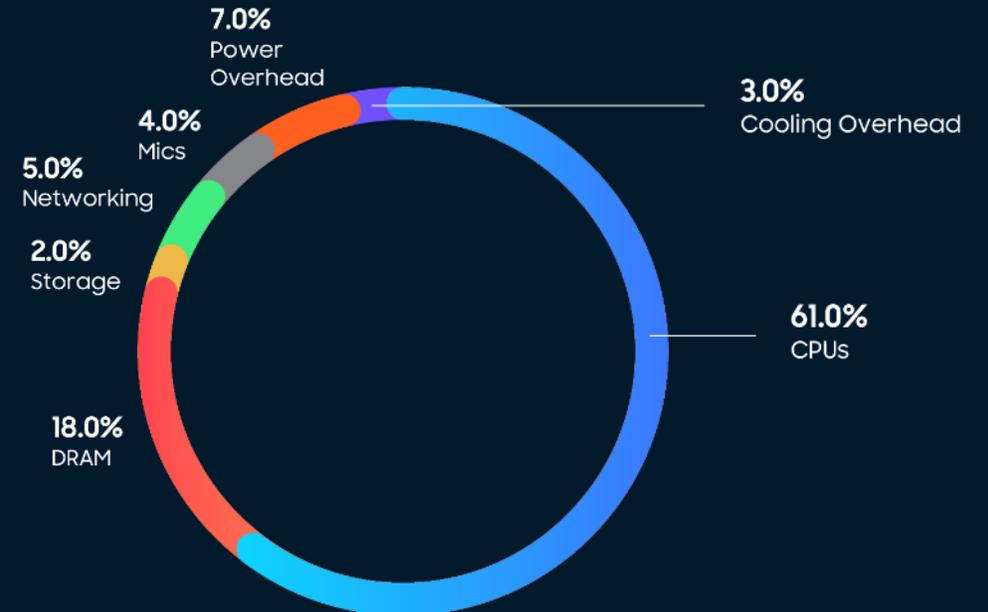
Source: energyinnovation

11%
Storage Drives

3%
Network

43%
Servers

43%
Cooling and Power provision systems

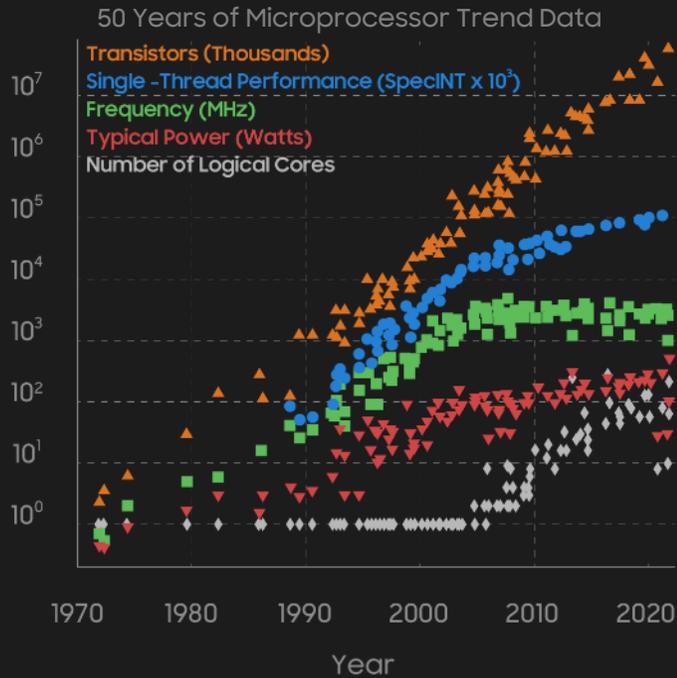


Source : morganclaypool

Industry Trend

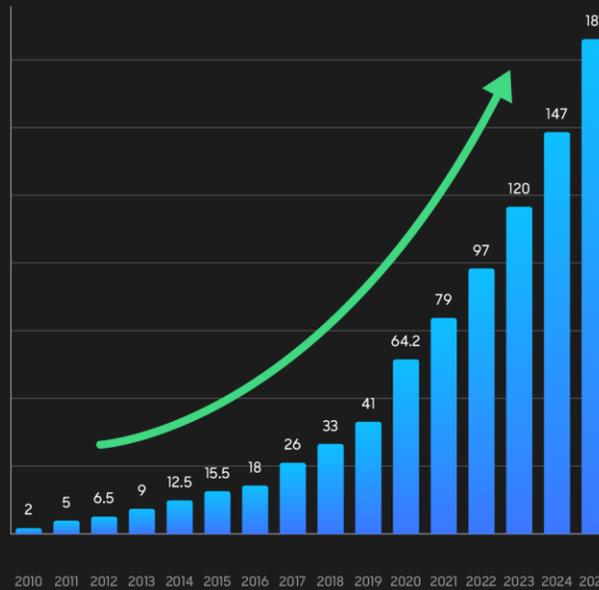
- Domain Specific Architecture

CPU Performance Growth is slowing down



Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
 New plot and data collected for 2010-2021 by K. Rupp

Data Growth

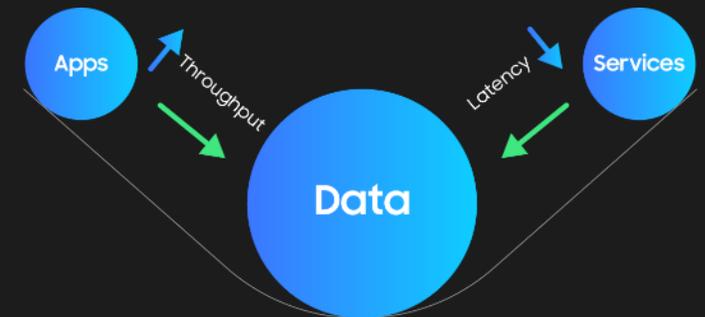


Volume of Data exponentially increases

Source: Statista

Data Gravity

Moving computation closer to data source can address the both problems



Source: Medium

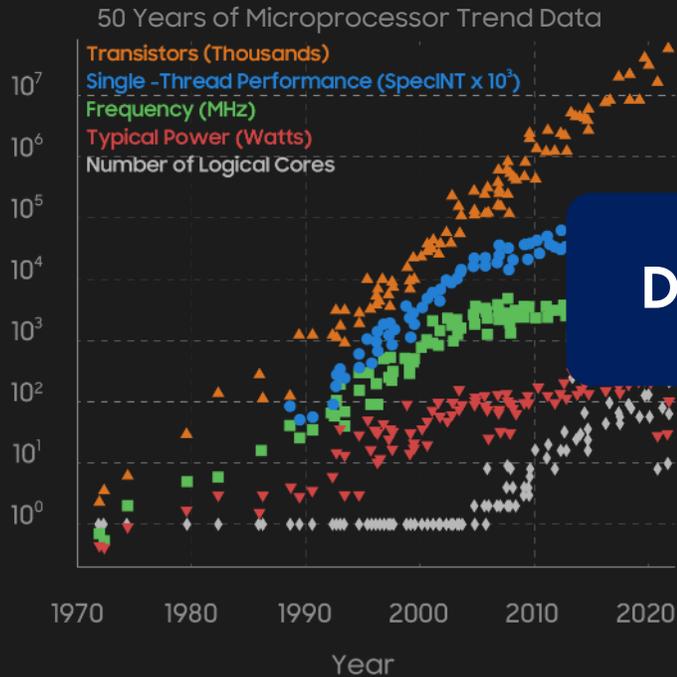
Why Compute Engines In Storage?

- Why do compute engines need to be in a storage device
 - Discrete GPU or FPGA boards on host
 - Each device occupies many PCI-e lanes
 - Limits the number of high-speed SSDs or allow a small number of compute devices
 - Requires complex control path and data path management in a host
 - Requires lots of power
 - In-storage computing
 - Internal IO is processed by a dedicated SSD controller
 - Internal HW resources and performance can be customized for the target domain including storage, big-data processing systems, and AI.
 - Various compute resources, including ARM, ASIC, FPGA, etc.
 - Sizes and types of internal DRAM
 - Do not require additional the PCI-e lanes allowing the compute engines to scale

Industry Trend

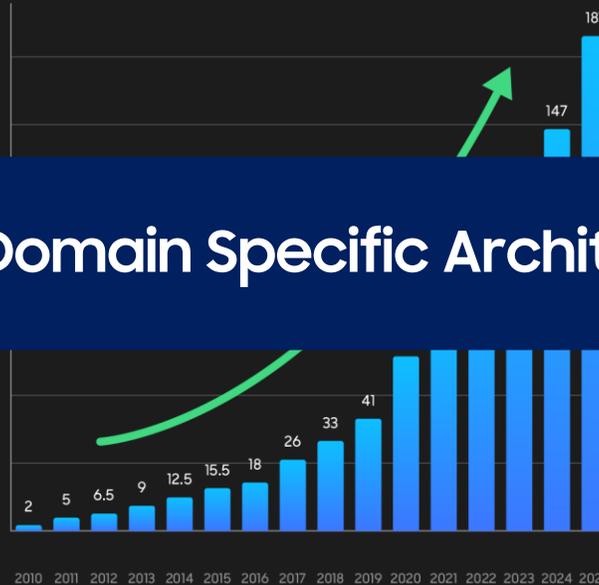
- Domain Specific Architecture

CPU Performance Growth is slowing down



Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2021 by K. Rupp

Data Growth

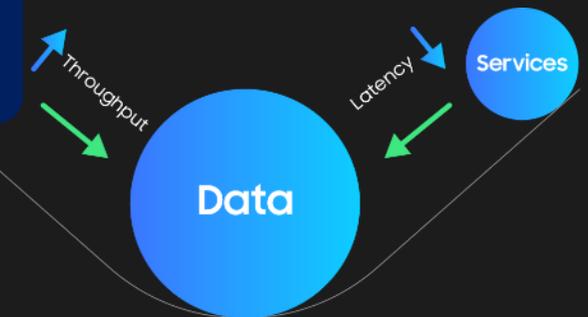


Volume of Data exponentially increases

Source: Statista

Data Gravity

Moving computation closer to data source can address the both problems



Source: Medium

Data-Domain Specific Architecture

Computational Storage

Data-Domain Specific Architecture

- Computational Storage
 - CSD, CSA, ...
- What is CSD (Computational Storage Drive)?
 - CSD = Persistent data storage + Computation
- Samsung SmartSSD®
 - SSD + HW acceleration engines
- Standard
 - NVMe computational storage (TP4091, TP4131)
 - SNIA
 - Computational storage architecture and programming model
 - Computational storage API

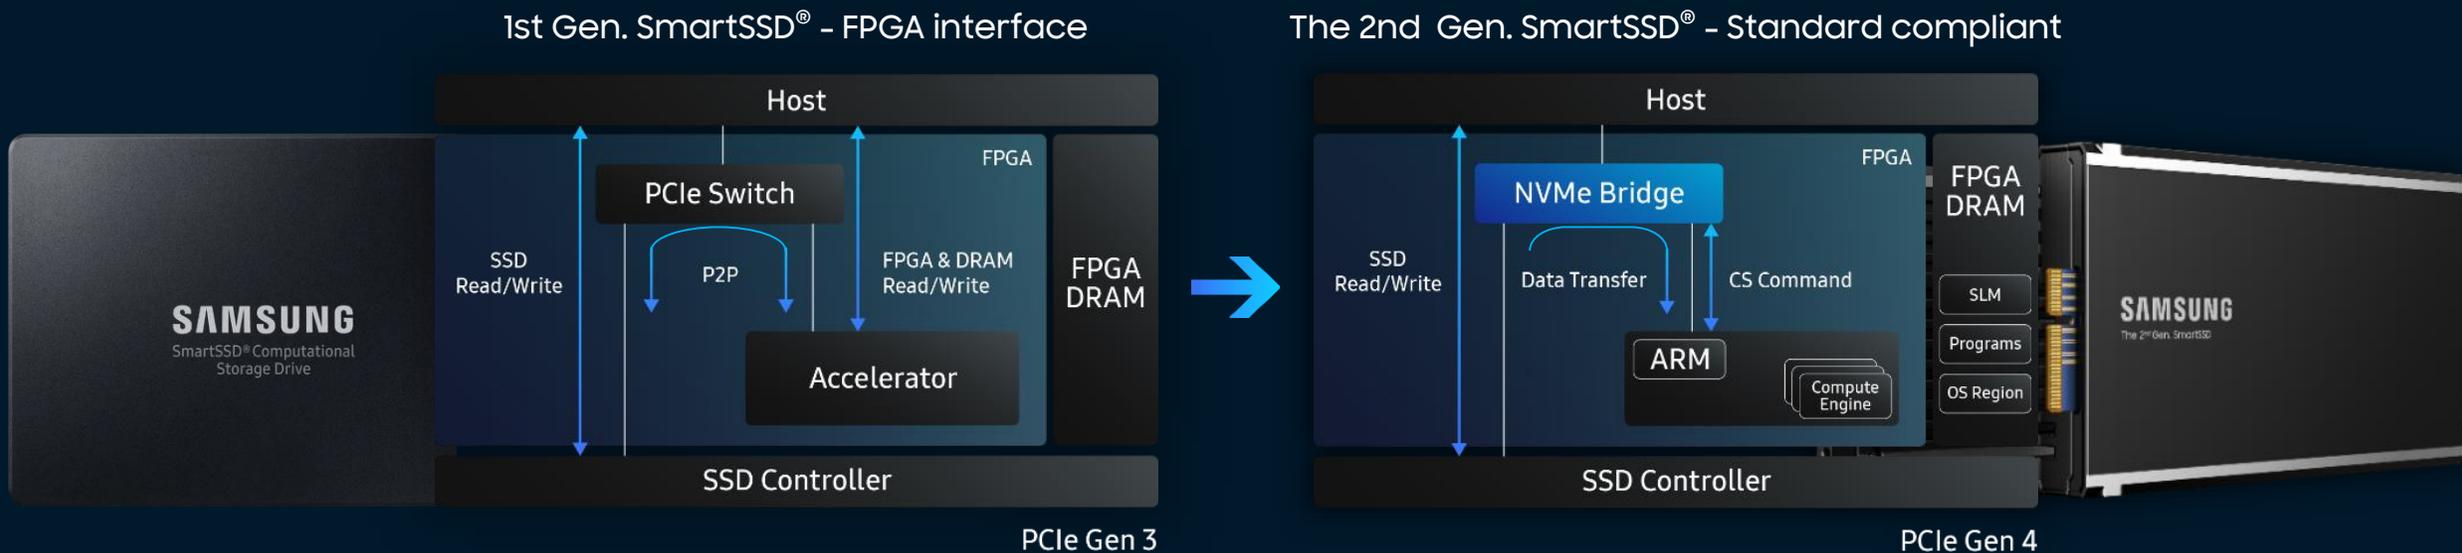


Computational Programs
Command Set Specifications

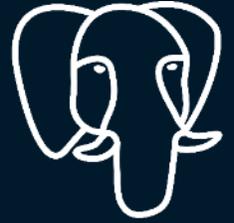


Samsung SmartSSD[®]

- SSD + HW acceleration engines
 - HW logic for data intensive operations (e.g., DB scan/filter, etc.)
 - At-Rest data processing
- The 1st Gen. SmartSSD[®] : FPGA interface based SmartSSD[®]
- The 2nd Gen. SmartSSD[®] : NVMe (TP4091) standard compliant SmartSSD[®]
 - Standard compliant eBPF for orchestration of offloaded SW + HW processing



PostgreSQL DB Engine



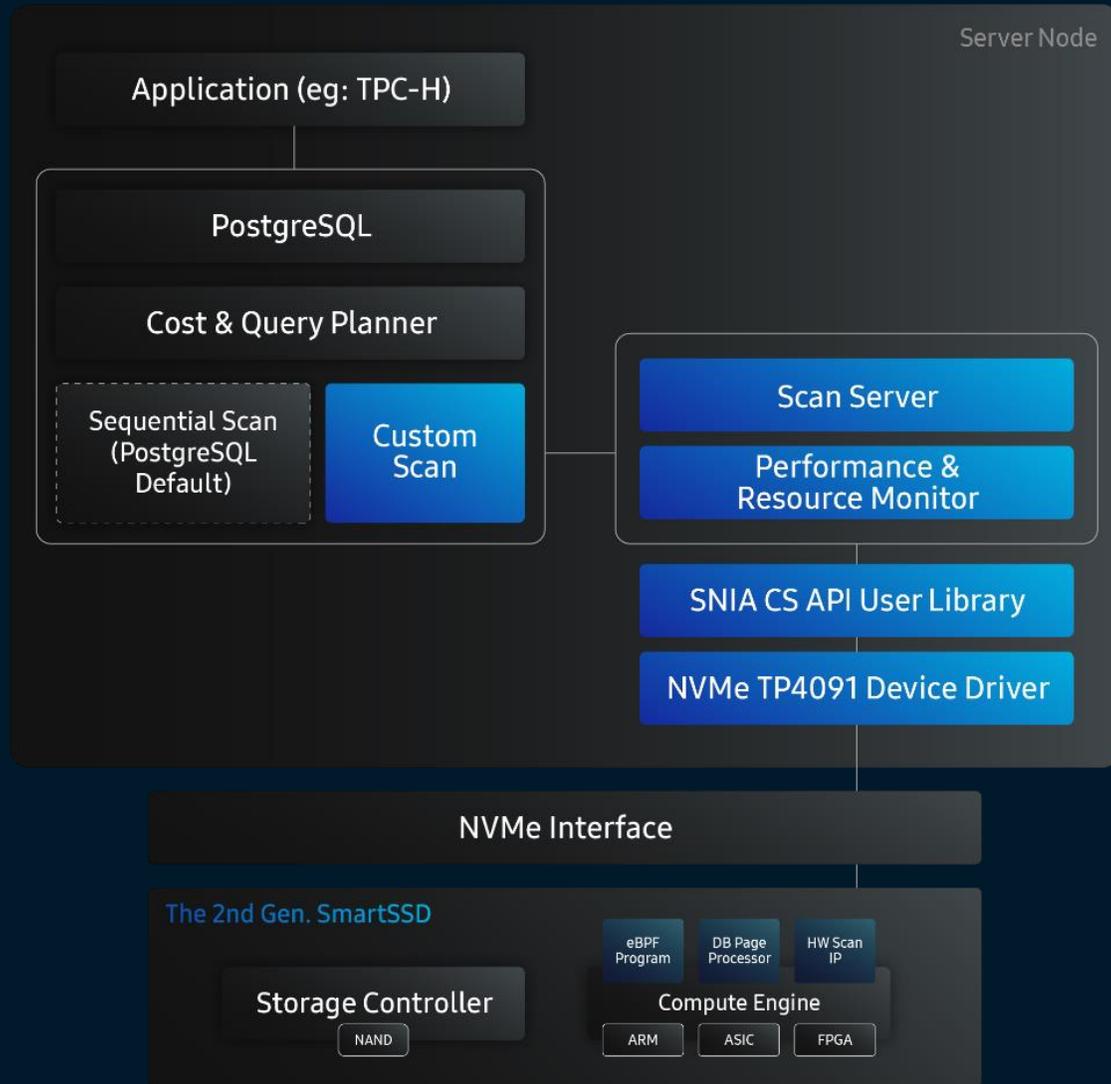
PostgreSQL

- Enterprise-class, full open-source
- Easily extensible plug-in module support for custom development
- The 4th most popular DB engine*
- PostgreSQL based DB engines deployed around 20% of data warehouse**

* https://db-engines.com/en/ranking_trend

** <https://www.datanyze.com/market-share/data-warehousing--240/ibm-netezza-100-market-share>

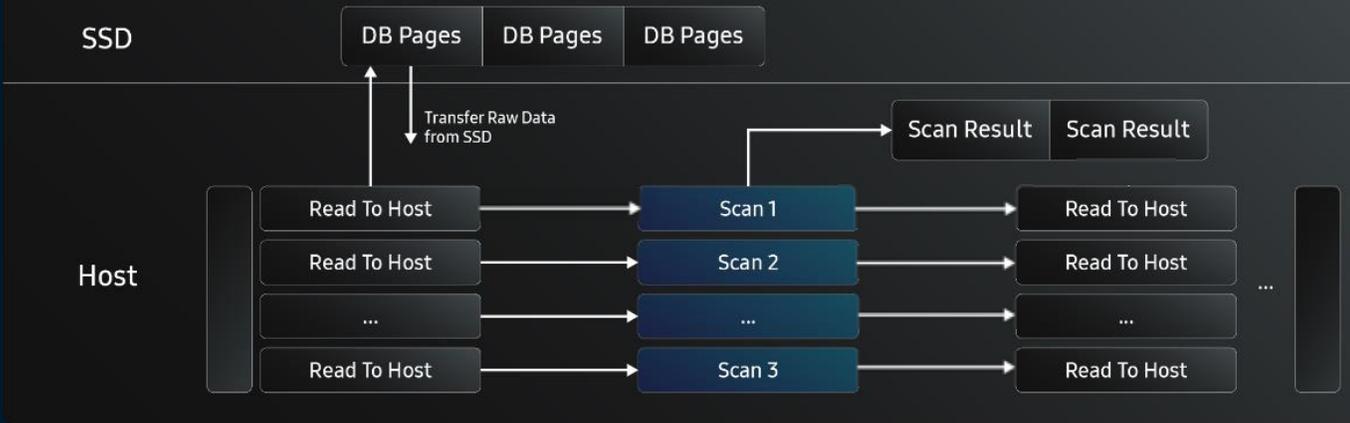
PostgreSQL DB Acceleration System Architecture



- Build Samsung software modules
 - Custom scan extension, scan server, CS API lib, etc.
 - Will be fully open sourced
- Enable PostgreSQL to use custom scan extension
 - Modify just one line of postgresql.conf
 - No PostgreSQL recompilation required
- Load Samsung NVMe device driver
- Run application

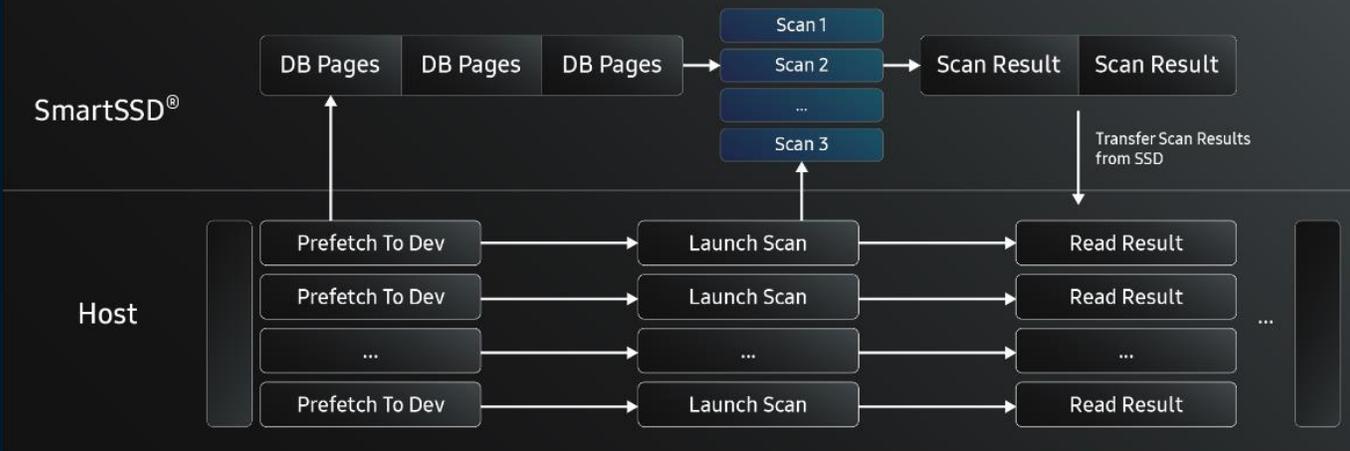
PostgreSQL DB Acceleration System Benefits

PostgreSQL + CPU - Sequential Scan Processing



- Large data movement (raw data)
- Under-utilized disk bandwidth
- Limited scalability (scale up)

PostgreSQL + SmartSSD® - Sequential Scan Processing



- Small data movement (scan result data only)
- High in-device IO bandwidth
- Lower host resource usage
 - Improved scalability (scale up)
- Low cost processing using in-device data processors

Performance Evaluation System

- System

- Supermicro SYS-420GP-TNR
- 2x 24-core Intel(R) Xeon(R) Gold 6342 CPU @ 2.80GHz
- DRAM: 128GB
- Up to 8x 3.84TB 2nd Gen. SmartSSD[®] PoC devices

- PostgreSQL

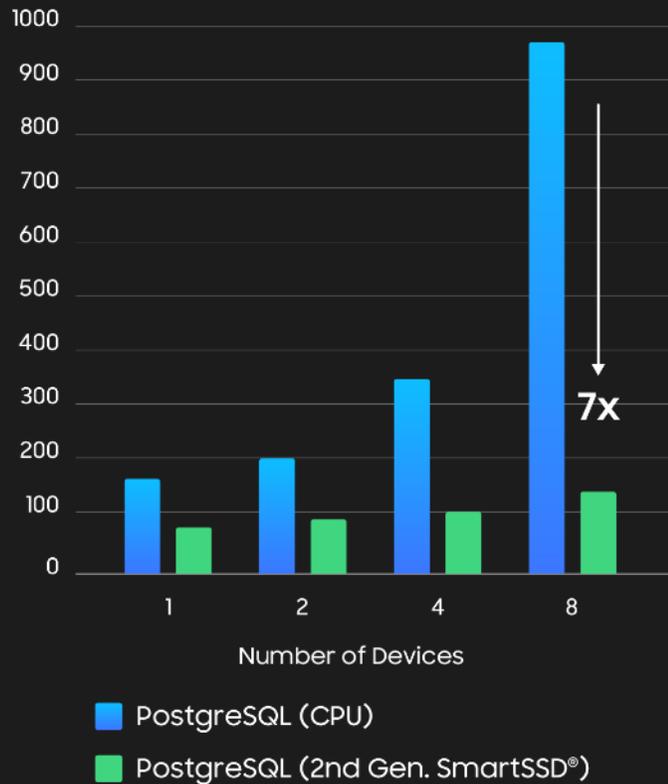
- PostgreSQL 13

- TPC-H Q6 SF1000 per device

The 2nd Gen. SmartSSD[®] Benefit

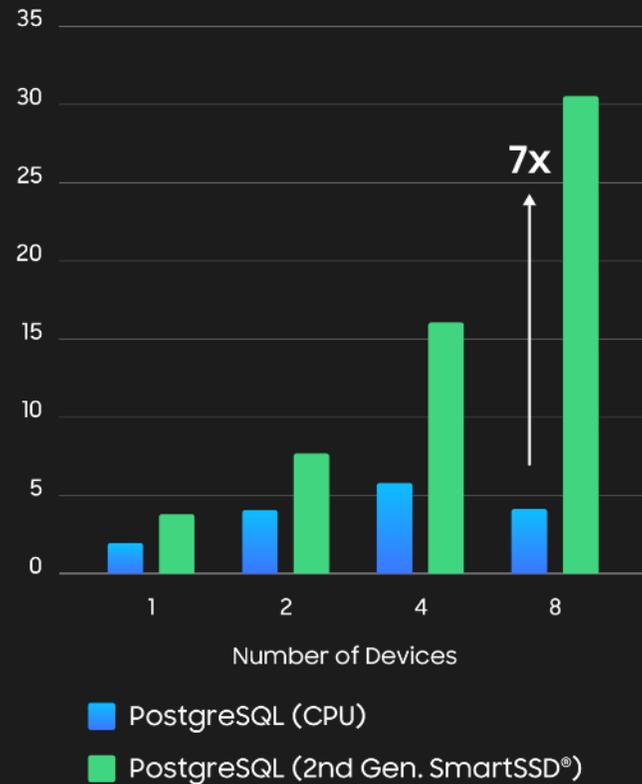
Energy Use (KJ)

Lower is Better



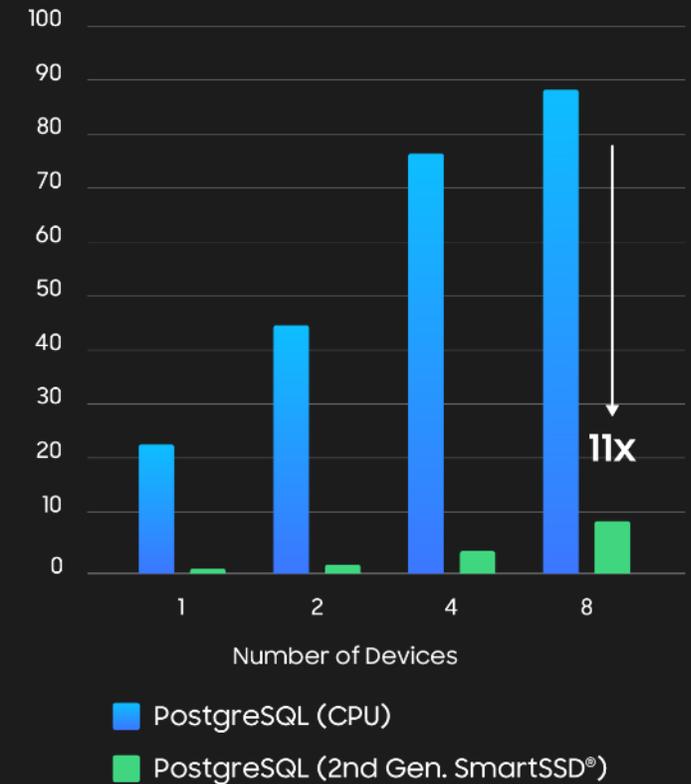
End-to-End Throughput (GB/s)

Higher is Better



CPU Utilization (%)

Lower is Better



Summary

- With the 2nd Gen. SmartSSD®
 - Green computing – more than 7x lower power usage
 - More than 7x End-to-End throughput enhancement
 - Better TCO with low host CPU utilization (more than 11x lesser utilization)
- Better security since data could stay in the at-rest storage
 - At-Rest data processing

Call for Action

- Visit Samsung demo
- Other sessions on Computational Storage
 - Samsung Keynote – Yang Seok Ki
 - Computational Storage APIs – Oscar Pinto
 - RETINA: Exploring Computational Storage (SmartSSD) Use case - Vishwanath Maram, Changwoo Min
 - Computational Storage: How Do NVMe CS and SNIA CS Work Together? – Bill Martin
 - Accelerating Near Real-time Analytics with High Performance Object Storage – Mayank Saxena
- Join the standardization efforts
 - SNIA, NVMe



Please take a moment to rate this session.

Your feedback is important to us.