

Database Scan Acceleration With Computational Storage Drive

Changho Choi

PhD, Senior Director of Engineering
Memory Solutions Lab
Samsung Semiconductor Inc.

August 4th 2022

Agenda


1. Industry Trend – Domain Specific Architecture
2. Computational Storage
 - Samsung SmartSSD[®]
3. DB acceleration system architecture
4. The 2nd Gen. SmartSSD[®] benefit
5. Summary



Domain Specific Architecture Era


CPU Performance Growth is slowing down

50 Years of Microprocessor Trend Data



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


Volume of Data exponentially increases



Source: Statista

Data Gravity

Moving Compute closer to data source can address these issues




Source: Medium

Domain Specific Architecture Era


CPU Performance Growth is slowing down

50 Years of Microprocessor Trend Data



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


Volume of Data exponentially increases



Source: Statista

Data Gravity

Moving Compute closer to data source can address these issues



Source: Medium

Data-Domain Specific Architecture

Computational Storage

Data-Domain Specific Architecture



Flash Memory Summit

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](#)



NVM Express[™]

Computational Programs
Command Set Specifications



Samsung SmartSSD®

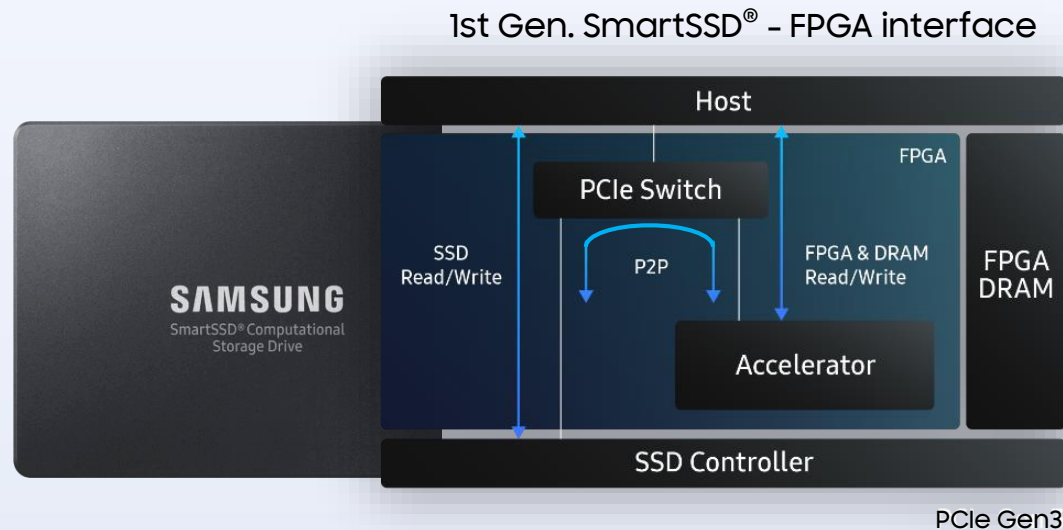
SSD + HW acceleration engines

- 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



The 2nd Gen. SmartSSD® - Standard compliant



PostgreSQL DB Engine



Enterprise-class, full open-source

Easily extensible plug-in module support for custom development

The 4th most popular DB engine*

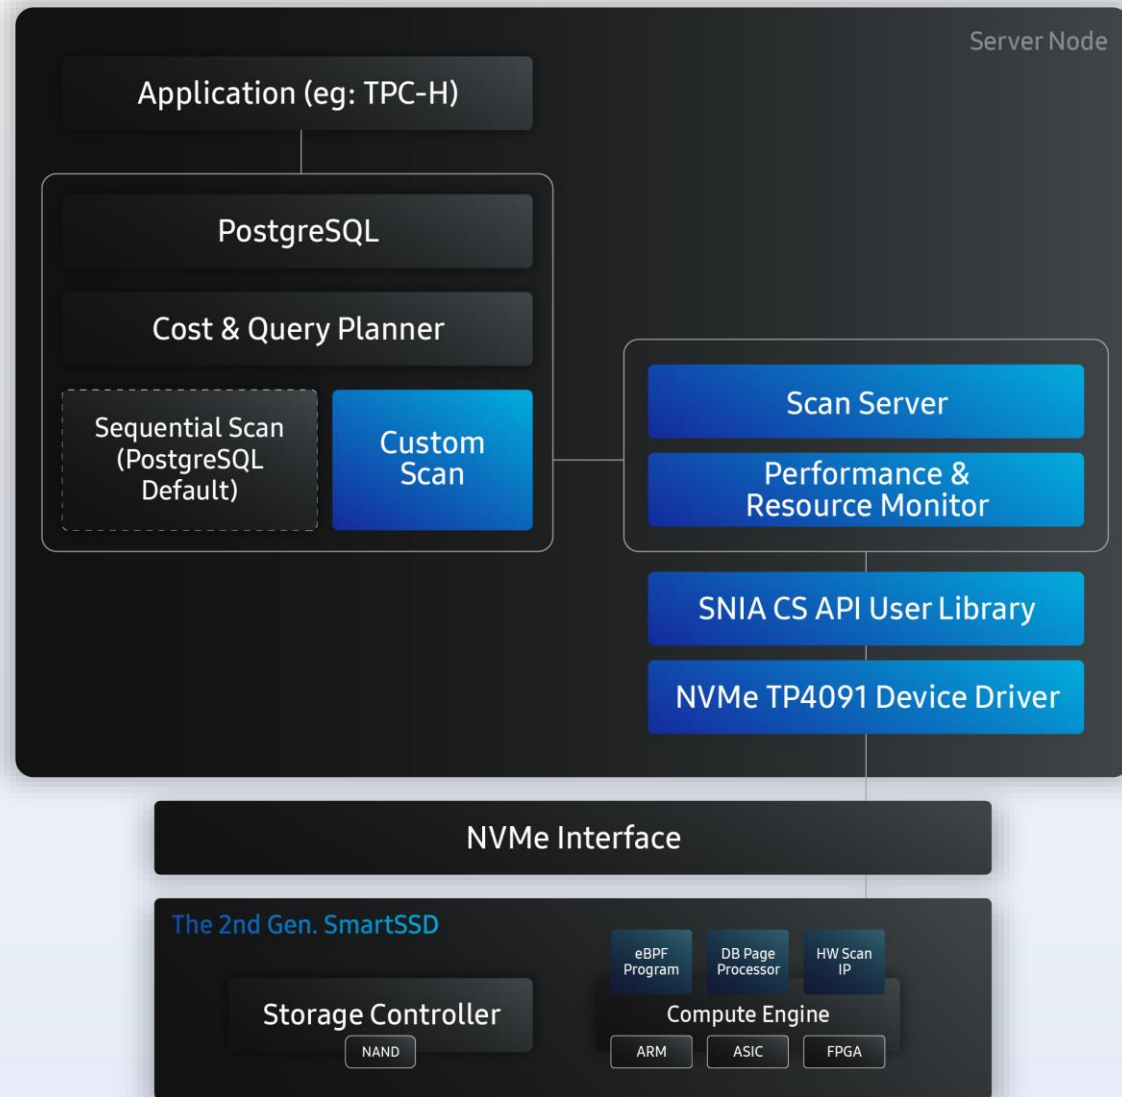
PostgreSQL based DB engines deployed around 20% of data warehouse**

* 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.

Enable PostgreSQL to use custom scan extension

- Modify just one line of postgresql.conf
- No PostgreSQL recompilation required

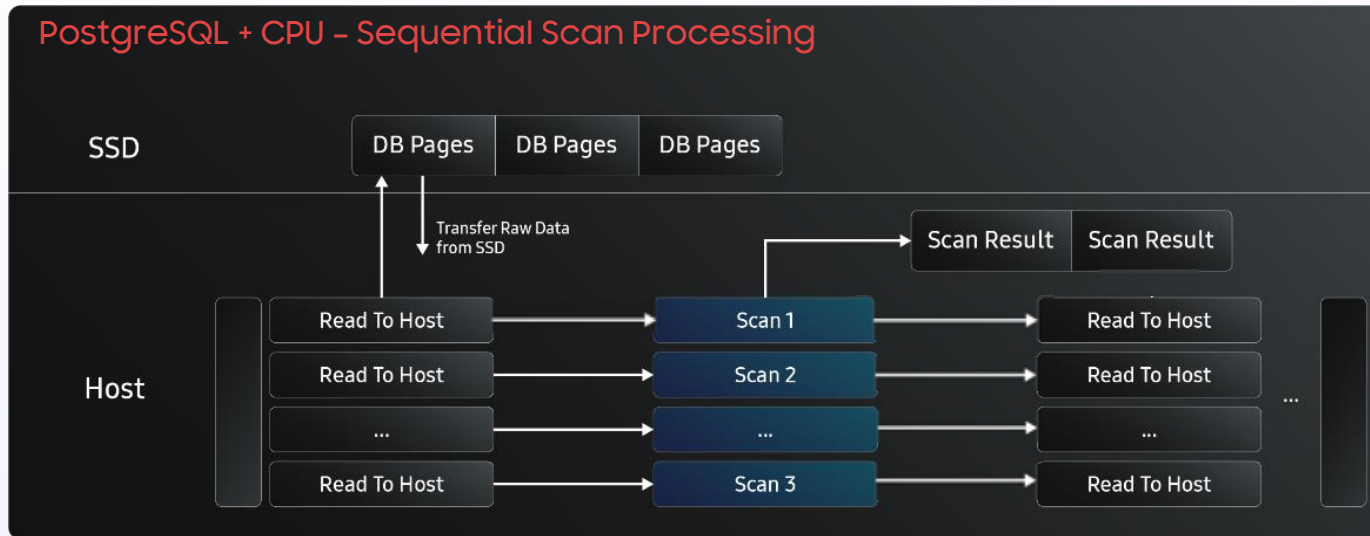
Load Samsung NVMe device driver

Run application

All SW modules will be fully open sourced

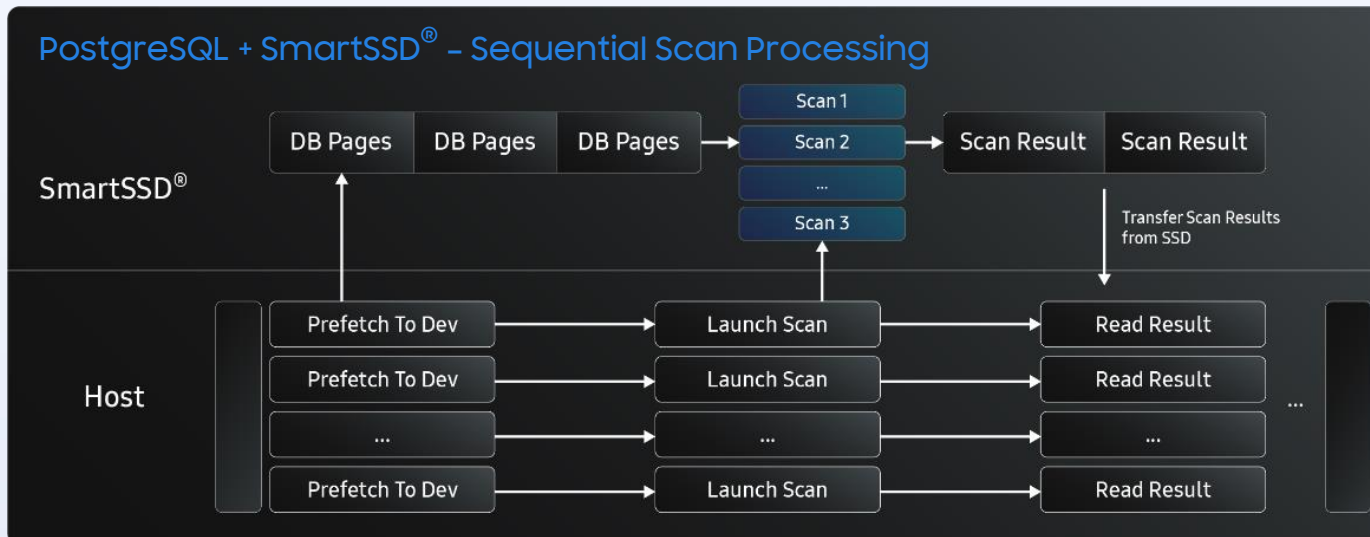
PostgreSQL DB Acceleration System Benefits

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

The 2nd Gen. SmartSSD[®] Benefit




Flash Memory Summit

End-to-End Throughput

End-to-End
Throughput
(GB/S)

Higher Better




■ PostgreSQL (CPU)
■ PostgreSQL (2nd Gen.SmartSSD[®])

CPU Utilization (%)

CPU
Utilization
(%)

Lower Better




■ PostgreSQL (CPU)
■ PostgreSQL (2nd Gen.SmartSSD[®])

Data Movement Device → Host (MB per sec)

MB

Lower Better



■ PostgreSQL (CPU)
■ PostgreSQL (2nd Gen.SmartSSD[®])

Summary

With the 2nd Gen. SmartSSD[®]

- More than 6.2x performance enhancement
- Better TCO with low host CPU utilization (33x lesser utilization)
- PCIe BW saving with less data movement (85x lower data movement rate)

Better security since data could stay in the at-rest storage

- At-Rest data processing

Samsung demo booth (#407) for more information



SAMSUNG

