



Flash Memory Summit

# Adding FPGA-based Acceleration to Flash Memory for Real-Time Analytics

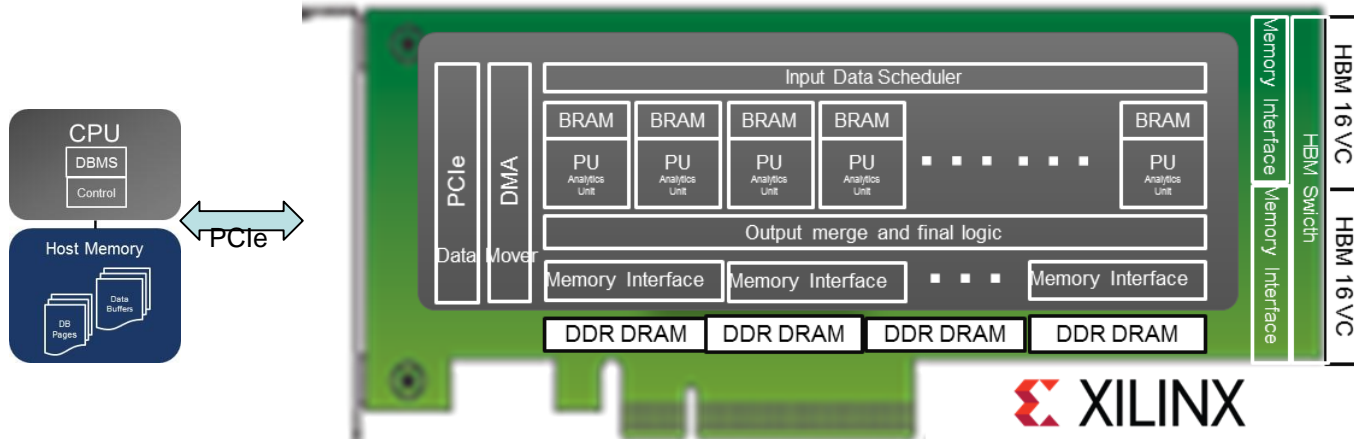
HK Verma,  
Distinguished Engineer, Xilinx Inc.

Increasing Performance by Moving Compute Closer to Data



# FPGA Platforms for Database Acceleration

## PCIe attached FPGA acceleration platform



- Target optimized hardware
- Enable massively parallel processing units
- DDR4 for high capacity, HBM for high bandwidth

VCU1525

Virtex Ultrascale+ XCVU9P  
4x DDR4 16GB, 2400MT/s, x64 with ECC

VCU1551

Virtex Ultrascale+ XCVU37P  
2HBM2 stack, 1024 @ 1.8 Gbps, 8GB  
4x DDR4 16GB, 2400MT/s, x64 with ECC



**FPGA Acceleration offers 10-50x compute efficiency improvement over CPUs**

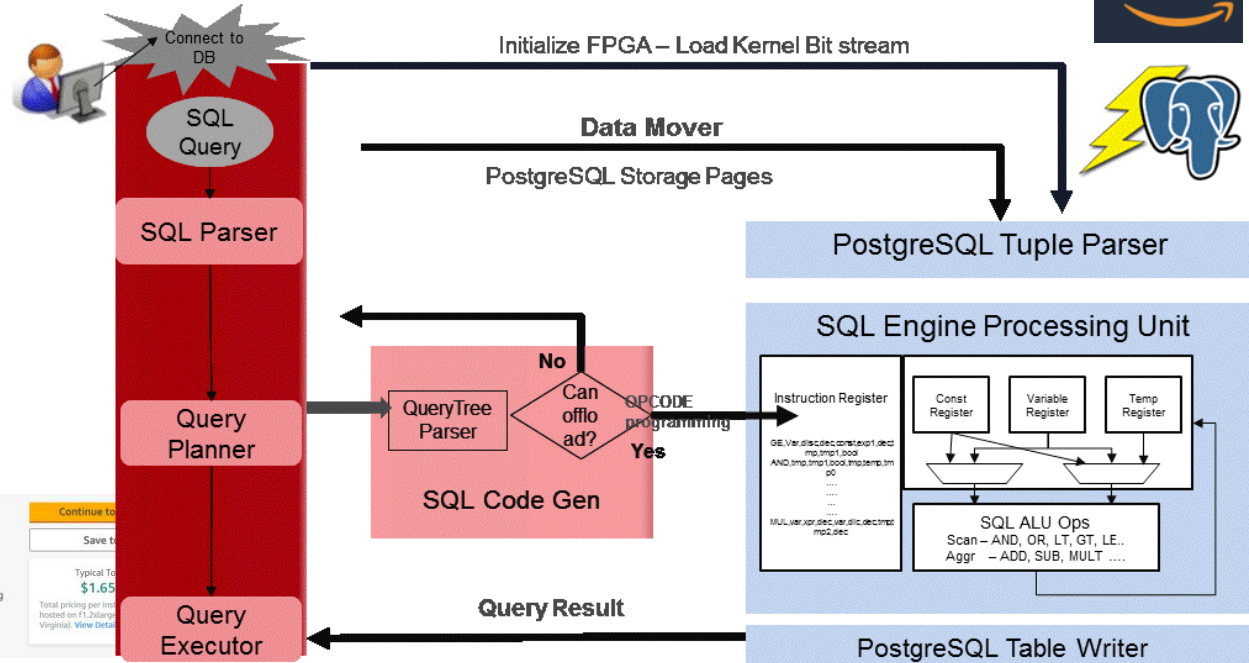


# Xilinx Accelerated Postgres on Amazon F1



## Flash Memory Summit

- Customers can run existing queries
- Uses 32 SQL PU in a parallel configuration
- Current PU implements scan & aggregate; extensible to hash, sort, or customer specific instructions



Data Analytics Acceleration Stack using Postgres, Preview Edition

Sold by: Xilinx Latest Version: 18.03.19

Xilinx's new data analytics acceleration stack enables users to easily evaluate FPGA-accelerated SQL query using Postgres. Users can execute their existing postgres queries using

Show more

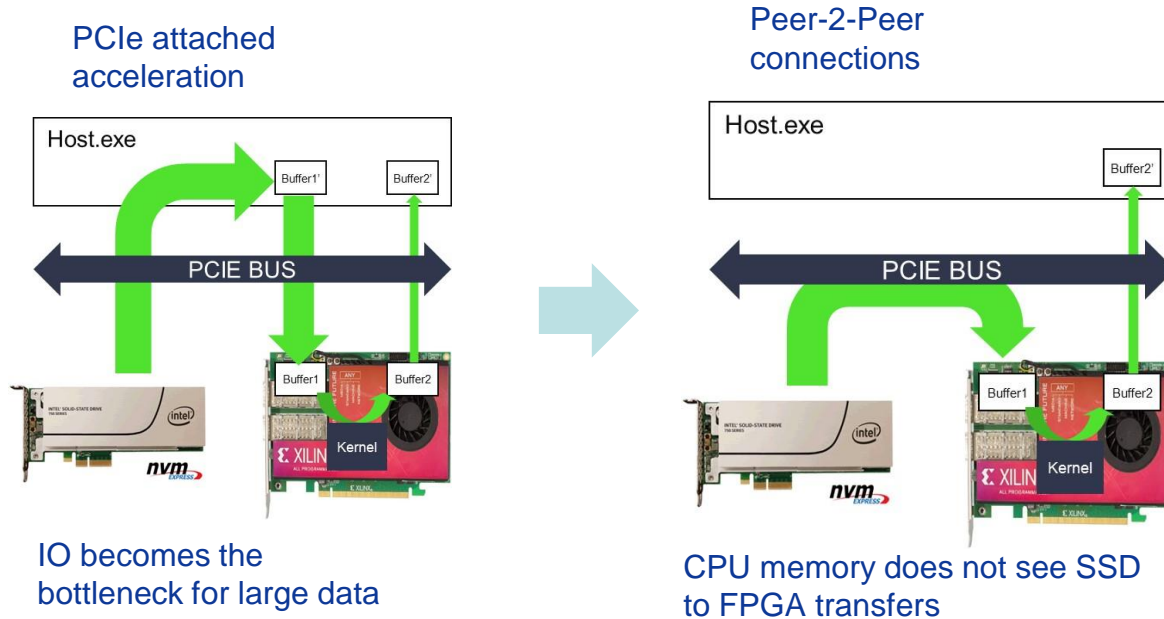
Linux/Unix ☆☆☆☆ (0)

<https://aws.amazon.com/marketplace/pp/B07BVSZL51>

**Offloads by hooking FPGA scan/aggregate into Postgres query plan**



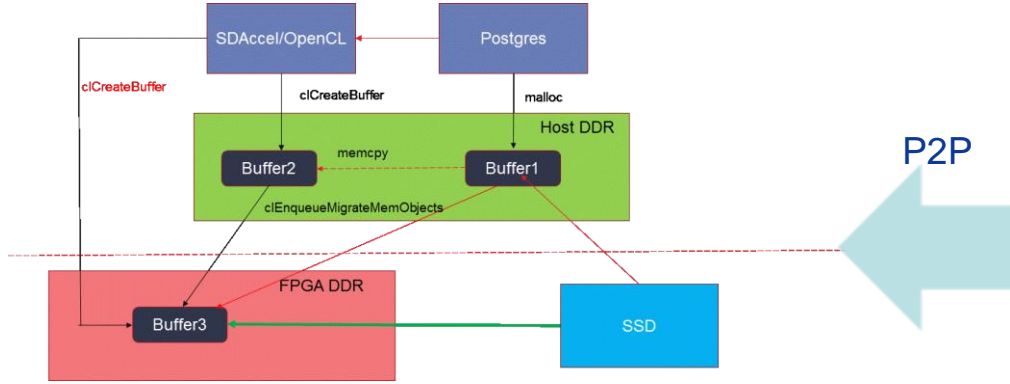
# Bringing compute closer to storage




**Integrated compute relieves IO bottlenecks, frees up CPU for higher performance**



# Enabling Peer-2-Peer Acceleration with SSD





Xilinx booth demo shows TPCH Query 6 in PostgreSQL using Peer2Peer connection to SSD



```

Aggregation
SELECT  * FROM L1_EXTRACT
WHERE  L1_DISCOUNT >= '1994-01-01' AND L1_SHIPDATE < '1994-01-01' AND L1_DISCOUNT BETWEEN .06 AND .09 AND L1_QUANTITY < 24
FROM  SHIPITEM
Data
Filter
  
```

Query 6 runs 30x faster on FPGA with 32 parallel processing units  
IO bandwidth limits the performance, CPU cycles are also inefficiently utilized



**Direct connection between FPGA and SSD relieves IO and CPU cycles**



# Xilinx Demo with Peer-2-Peer Acceleration with Storage Devices

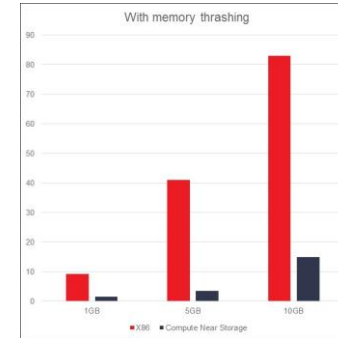
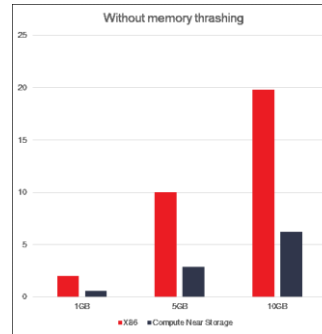
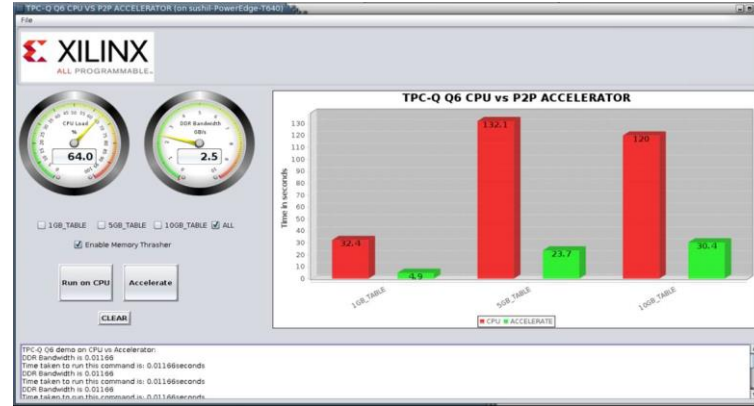
Flash Memory Summit

## Objective

- Showcase FPGA P2P capabilities for enabling efficient storage acceleration

## Application

- TPC-H Query 6 accelerated in Postgres using SDAccel stack and P2P implementation on Xilinx FPGA
- Database benefits by the direct P2P access from the storage to the acceleration kernel within FPGA



Please visit Xilinx Booth to see the demo !!!



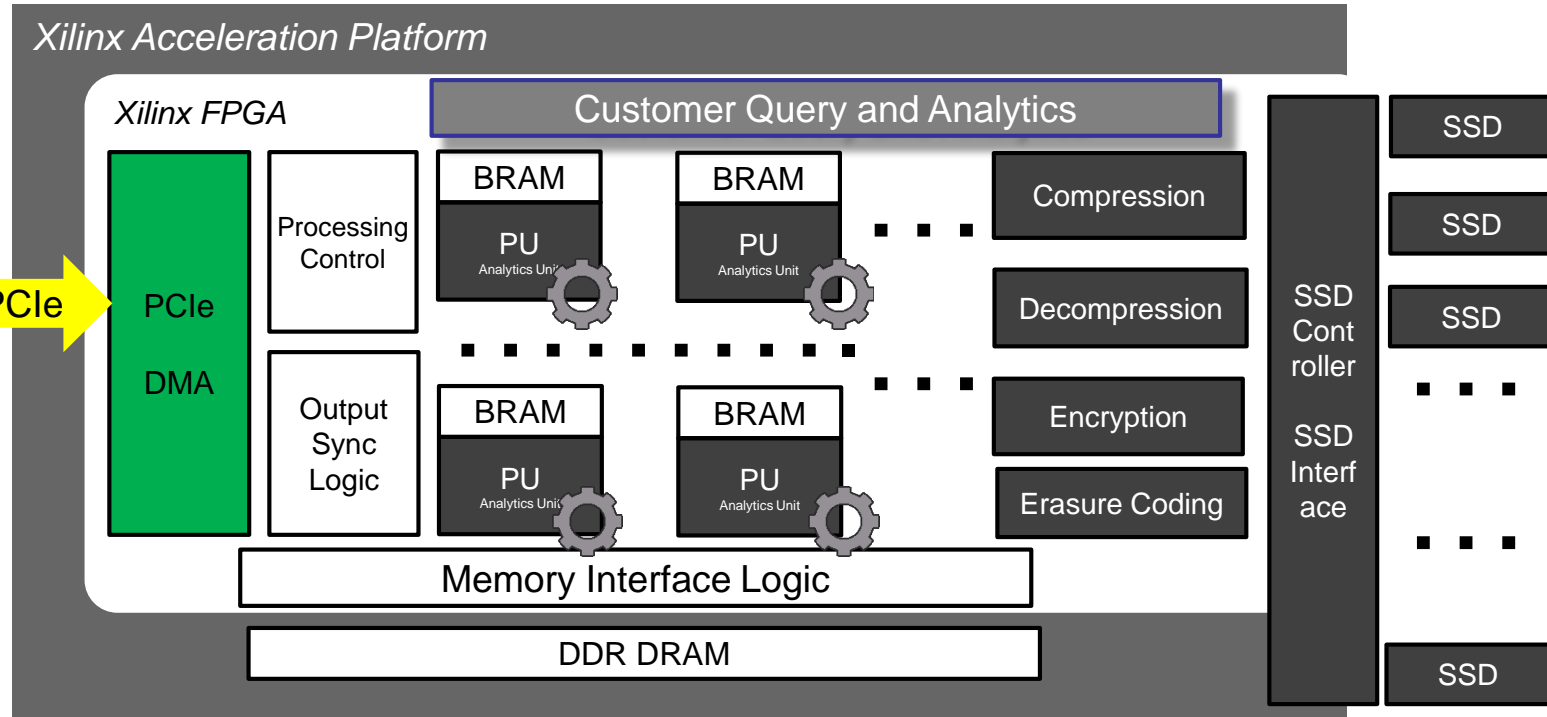
# Database in an integrated computational storage platform

Flash Memory Summit

Terabytes of Storage

Processed and PCIe filtered in FPGA

For gigabytes of IO transfer & CPU processing



**Improve database performance by bringing compute closer to storage!!**



# Summary

- Successful demo of a system architecture integrating FPGAs and flash storage with system software stacks
- Xilinx platforms available with SDAccel tools to move data from SSD to FPGA device memory without going to CPU memory space overcoming IO bottlenecks
- Many data analytics and processing blocks available to be implemented on FPGA