

Accelerating Intelligence

Bishwa Roop Ganguly

Chief Solution Architect



Performance Needs of Financial Applications



Hardware Accelerators Break Through the Processing Wall




On track to millions of servers w/FPGAs




Introduces FPGA powered server instances




TPUs help avoid cost of 12-15 data centers





\$16.7B acquisition of Altera





\$1B in datacenter revenues





193% growth in datacenter segment



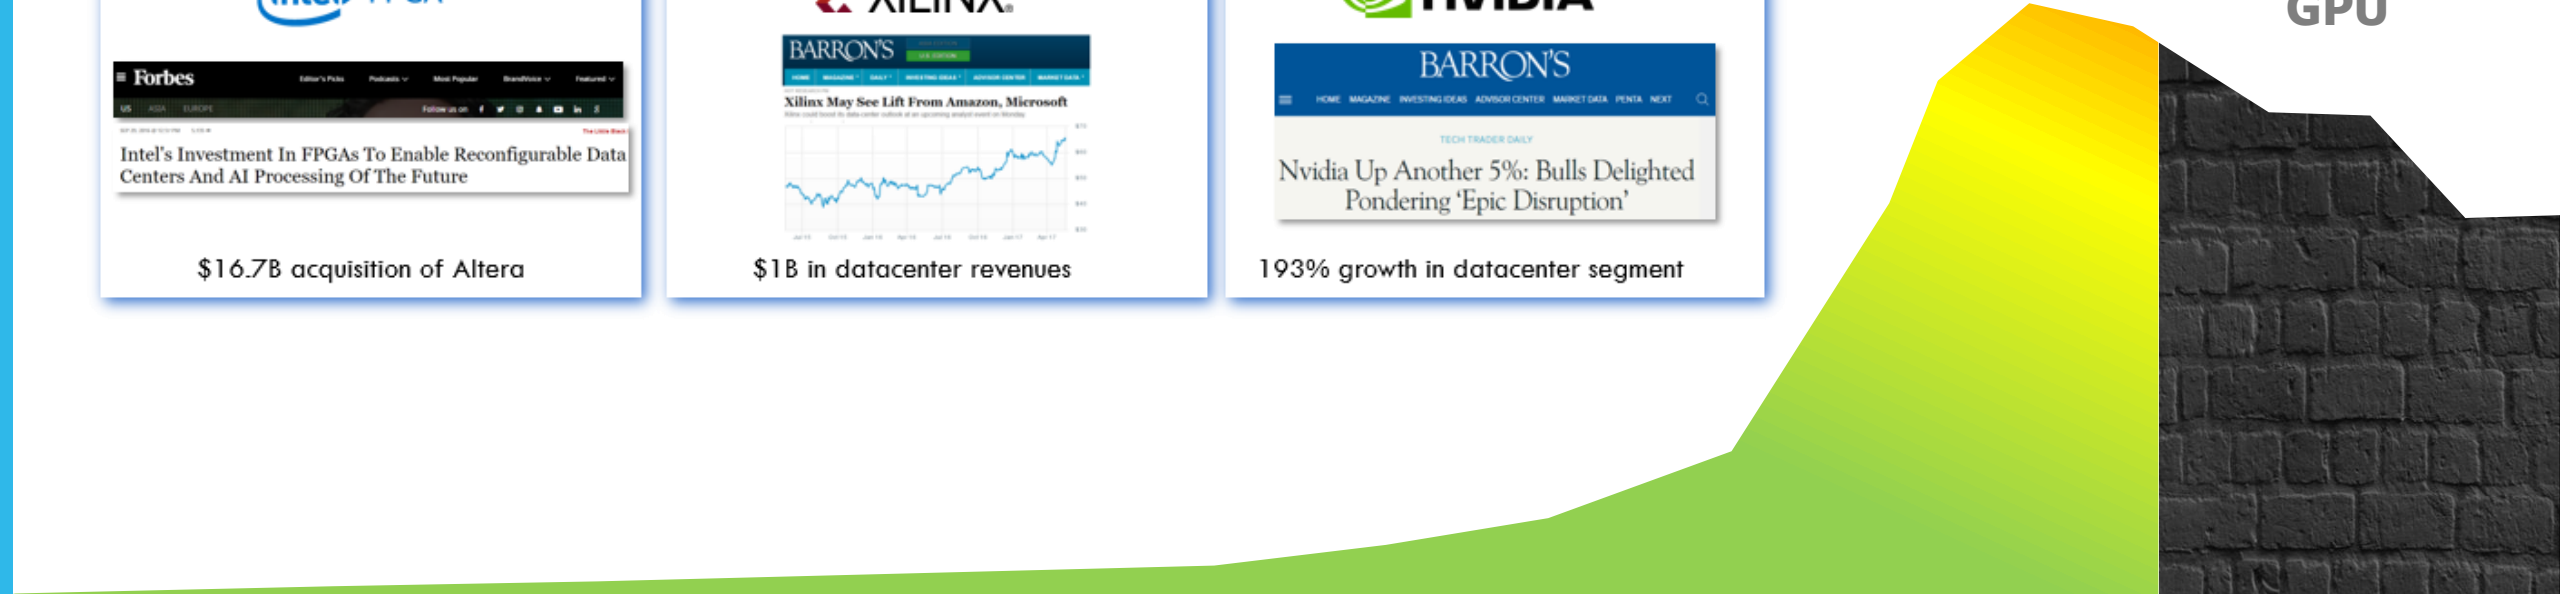

FPGA



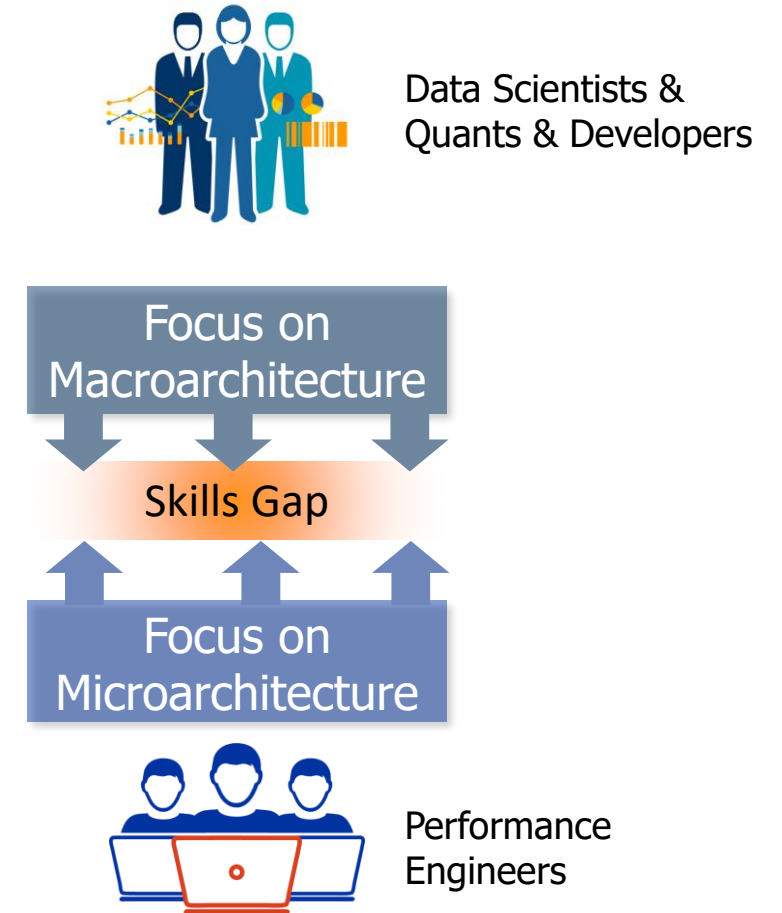
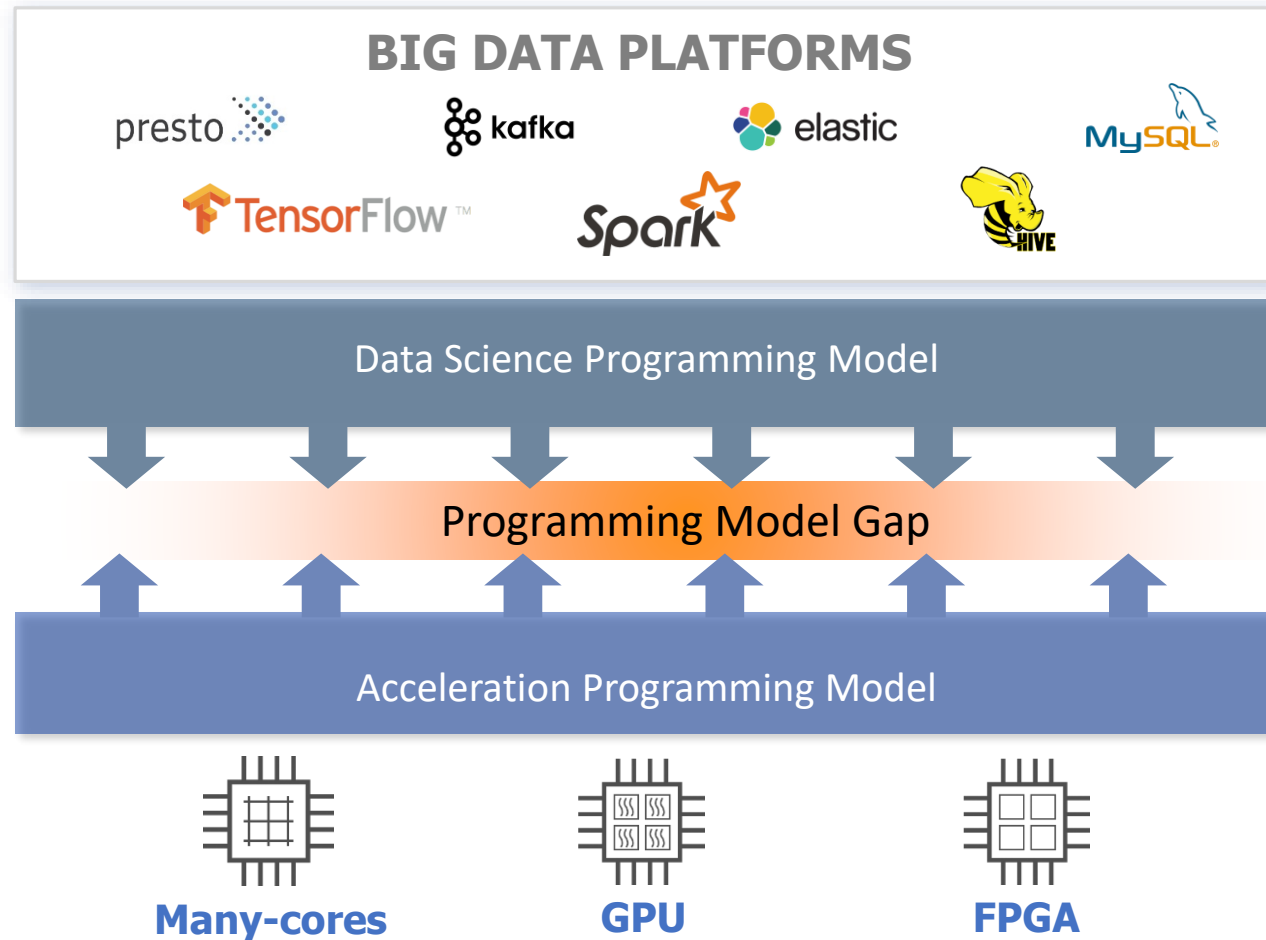
ASIC



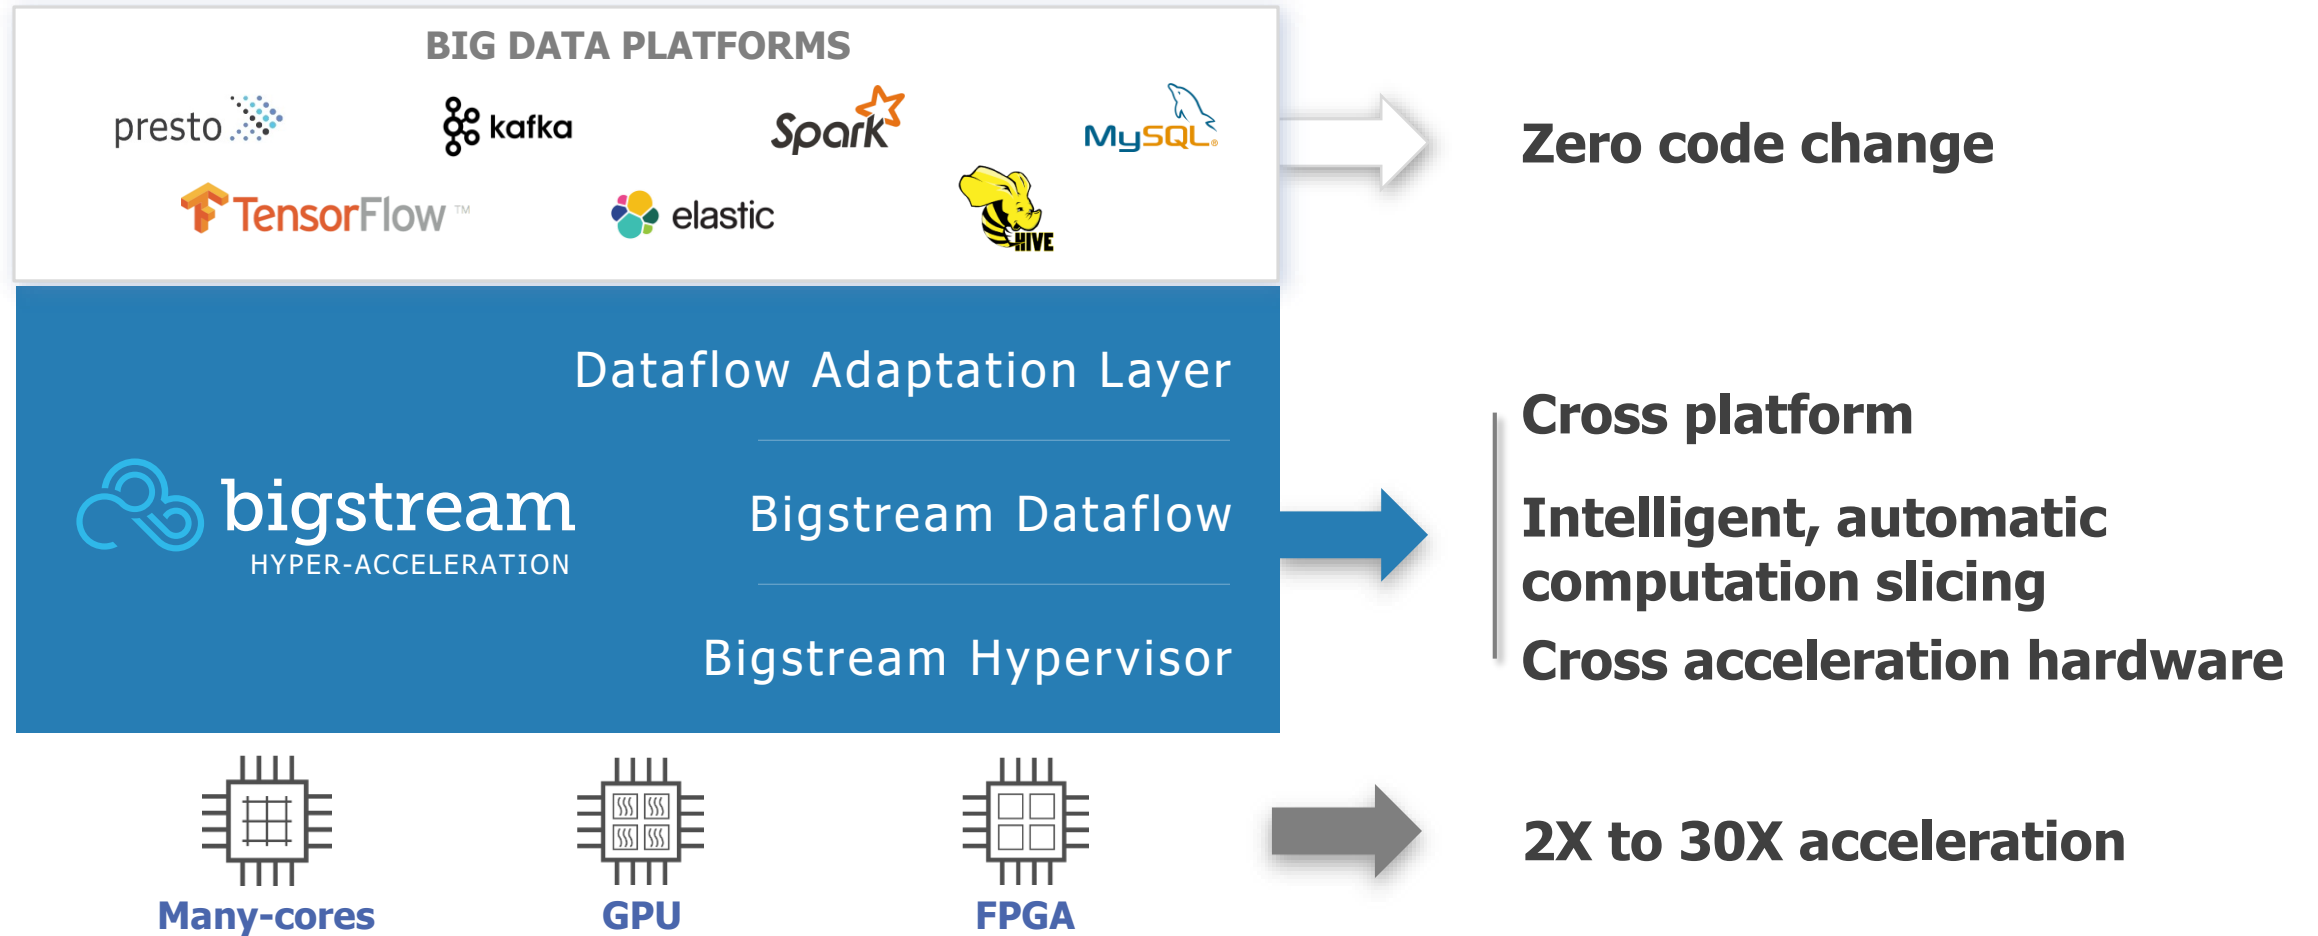
GPU



Inhibitor: Programming Model Gap for Hardware Accelerators



Introducing: Bigstream Hyper-acceleration Layer



Announcing: Acceleration of Spark on AWS F1 Instances



What: Hyper-accelerated Spark on F1 instance



How: Standard AWS EC2 GUI with Bigstream software

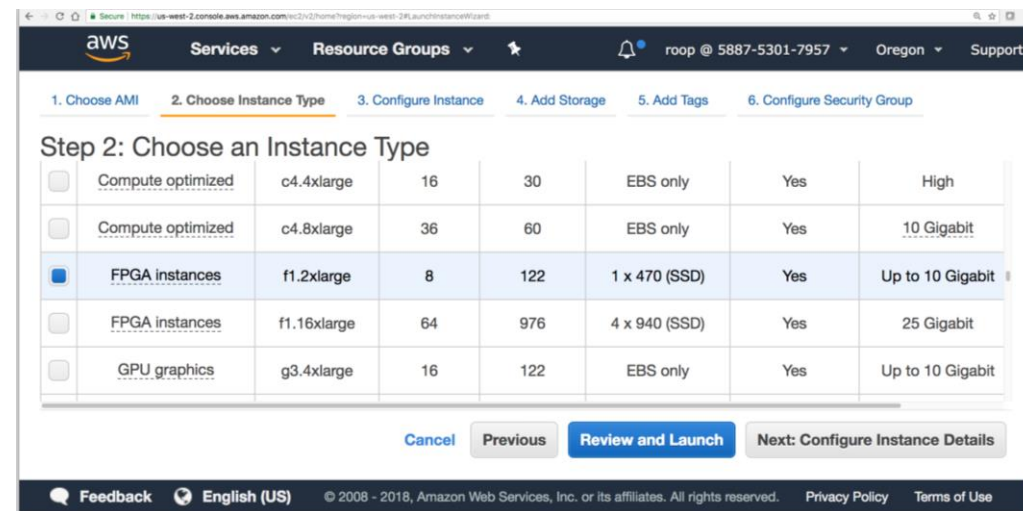


Result: Up to 5x end-to-end speedup

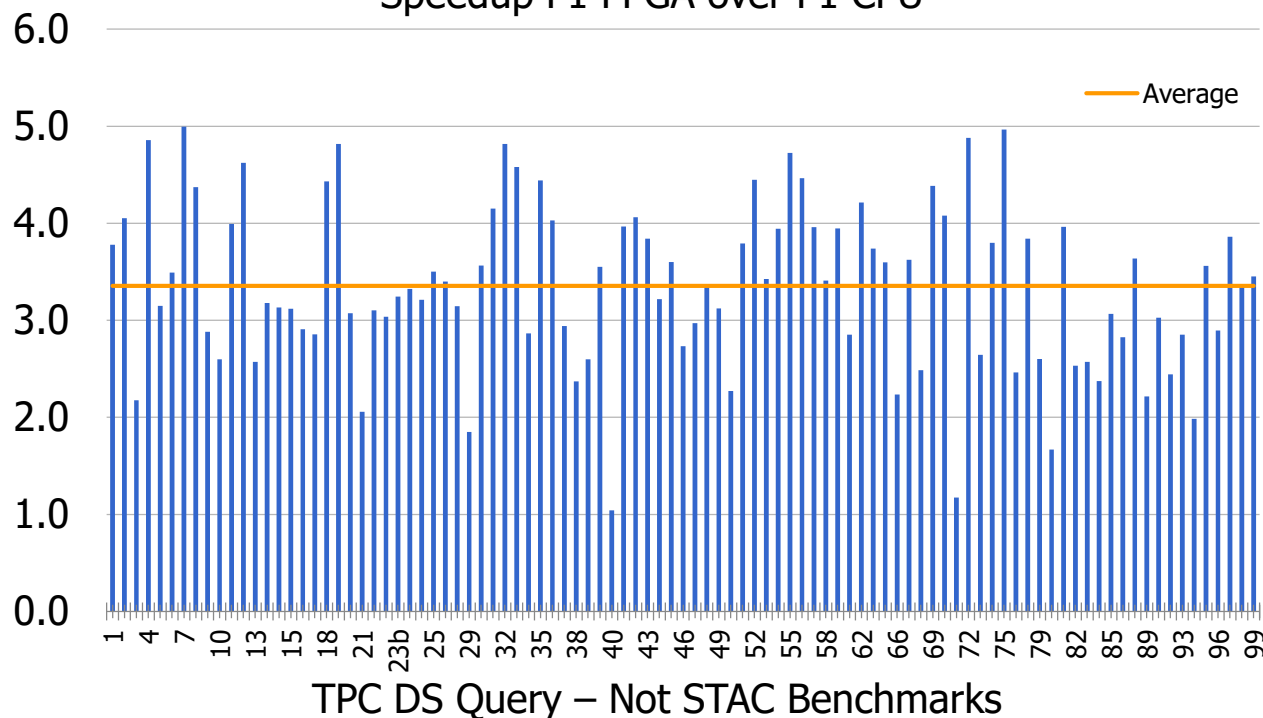
- CPU-only time vs. CPU/FPGA+Bigstream time
- **3.3x** average end-to-end speedup
- 45% FPGA utilization



Results applicable to on-premise customer deployments

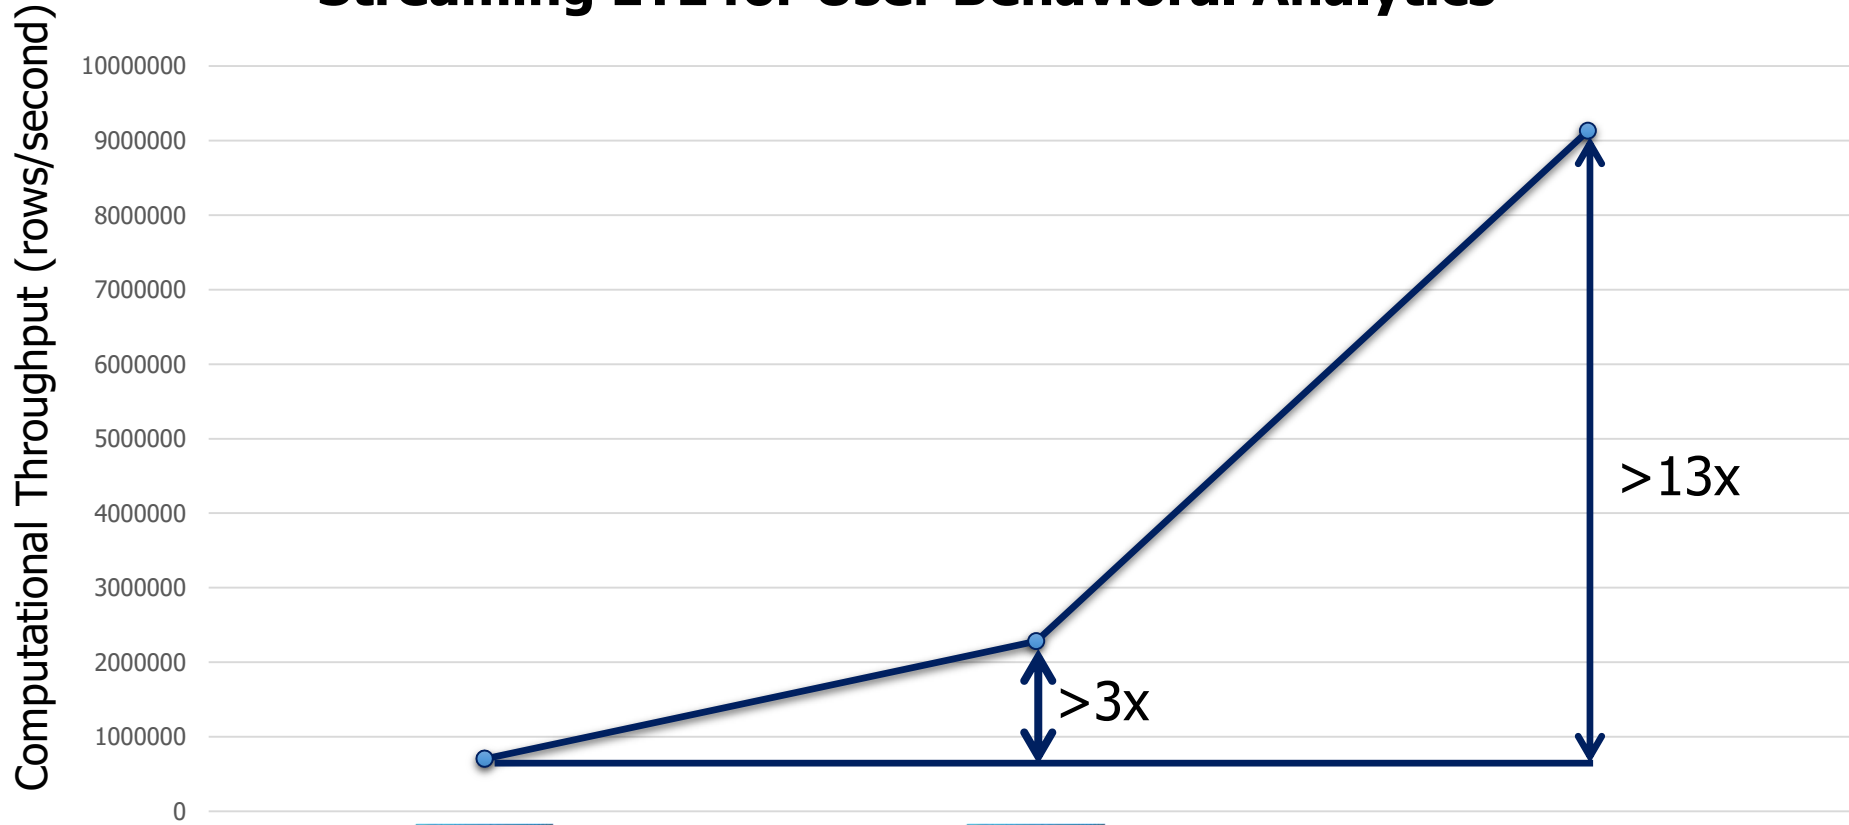


Speedup F1 FPGA over F1 CPU



On-Premise User Sentiment Analysis Pipeline

Streaming ETL for User Behavioral Analytics



Spark
700,000
rows/second



Bigstream NS
(SW Only Acceleration)
2,824,62 rows/second



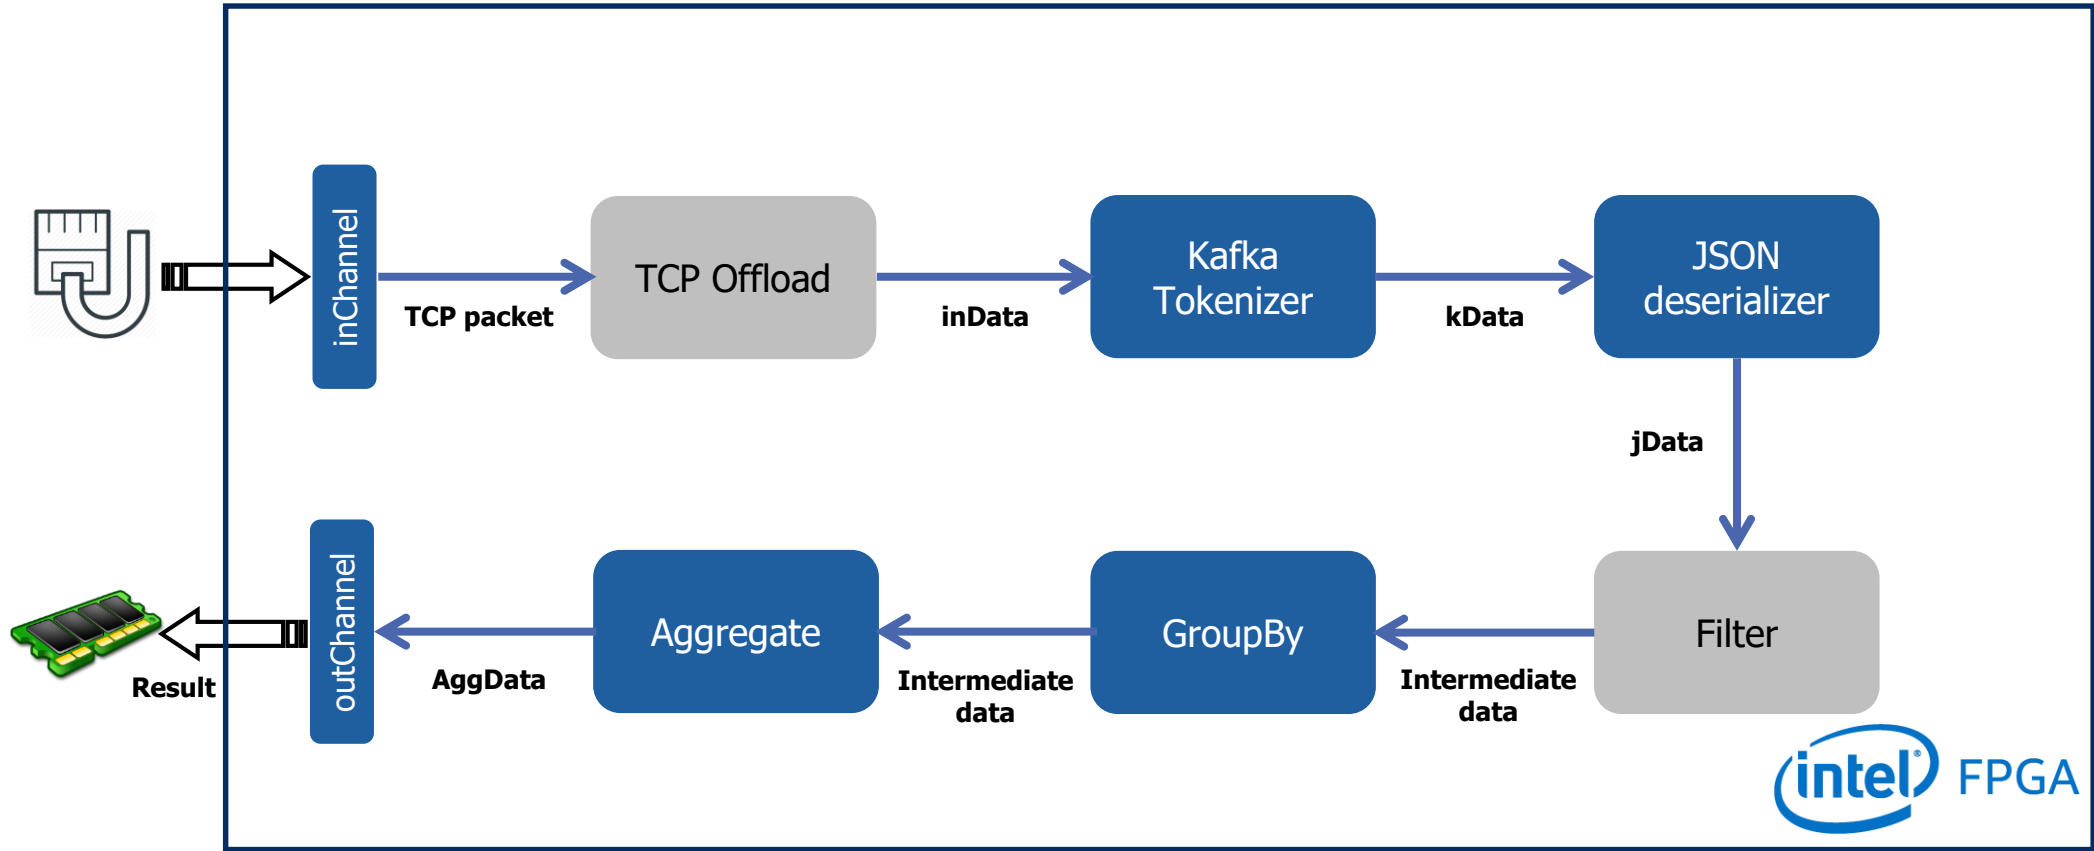
Bigstream XS
(FPGA Acceleration)
9,130,000 rows/second

Not STAC Benchmarks



HW Architecture: User Behavior Pipeline In-line Accelerator

Intel Stratix 5 FPGA



- 3rd Party IP
- Bigstream IP



Thank You

Check off Bigstream for More Info
roop@bigstream.co

Bigstream Acceleration of TPC-DS Spark On AWS F1

