

# **Bigstream: The Autonomous Data Platform for Accelerated AI**

**Bishwa Roop Ganguly**

**Chief Solution Architect**



# Performance Needs of Financial Applications



Risk Management



Algorithmic Trading



Fraud Analytics



Investor Intelligence



Backtesting




Security Analytics


# Hardware Accelerators Break Through the Processing Wall



FPGA



ASIC



GPU



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



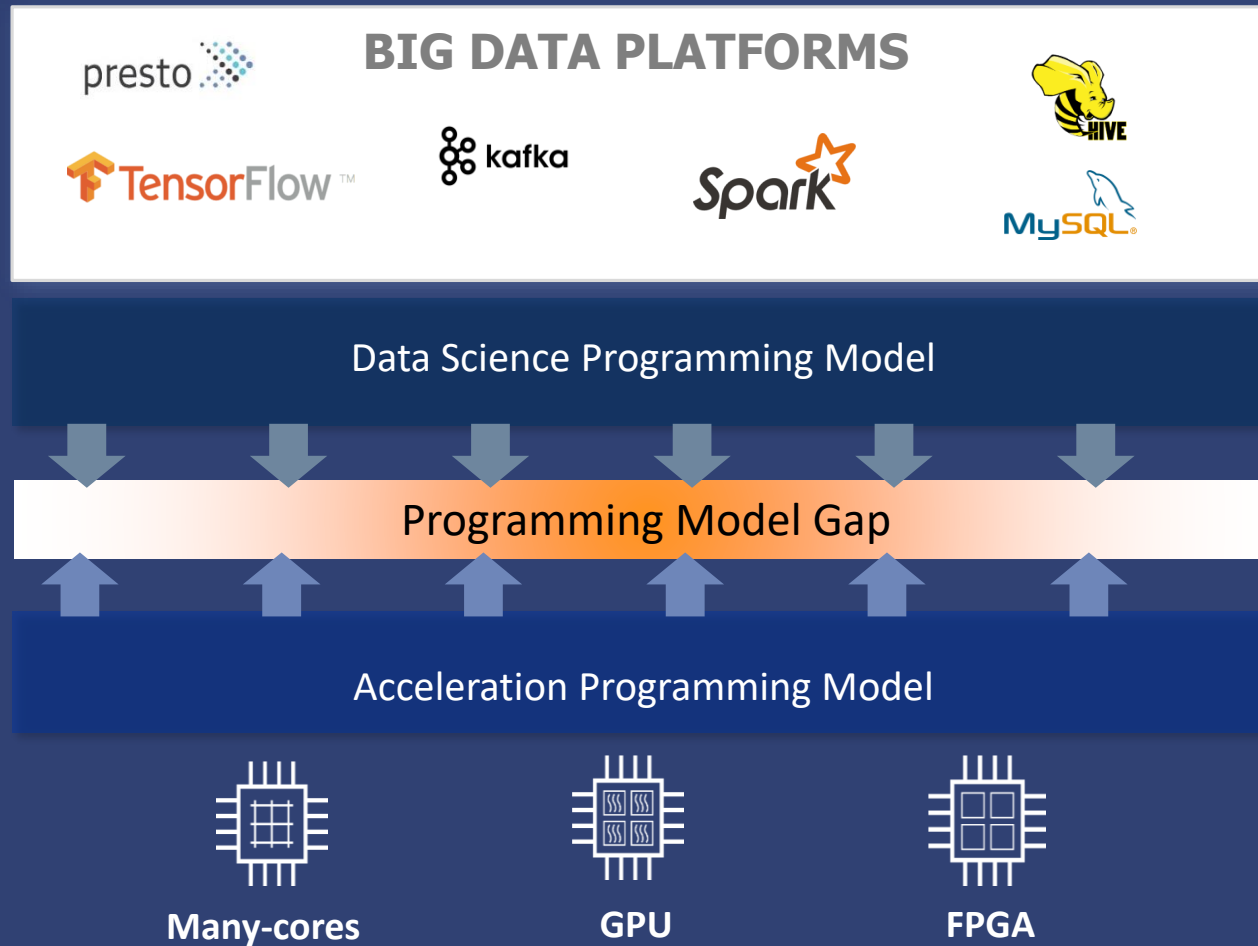
\$1B in datacenter revenues



193% growth in datacenter segment

# Inhibitor: Programming Model Gap for Hardware Accelerators

4



Data Scientists &  
Quants &  
Developers

Focus on  
Macroarchitecture

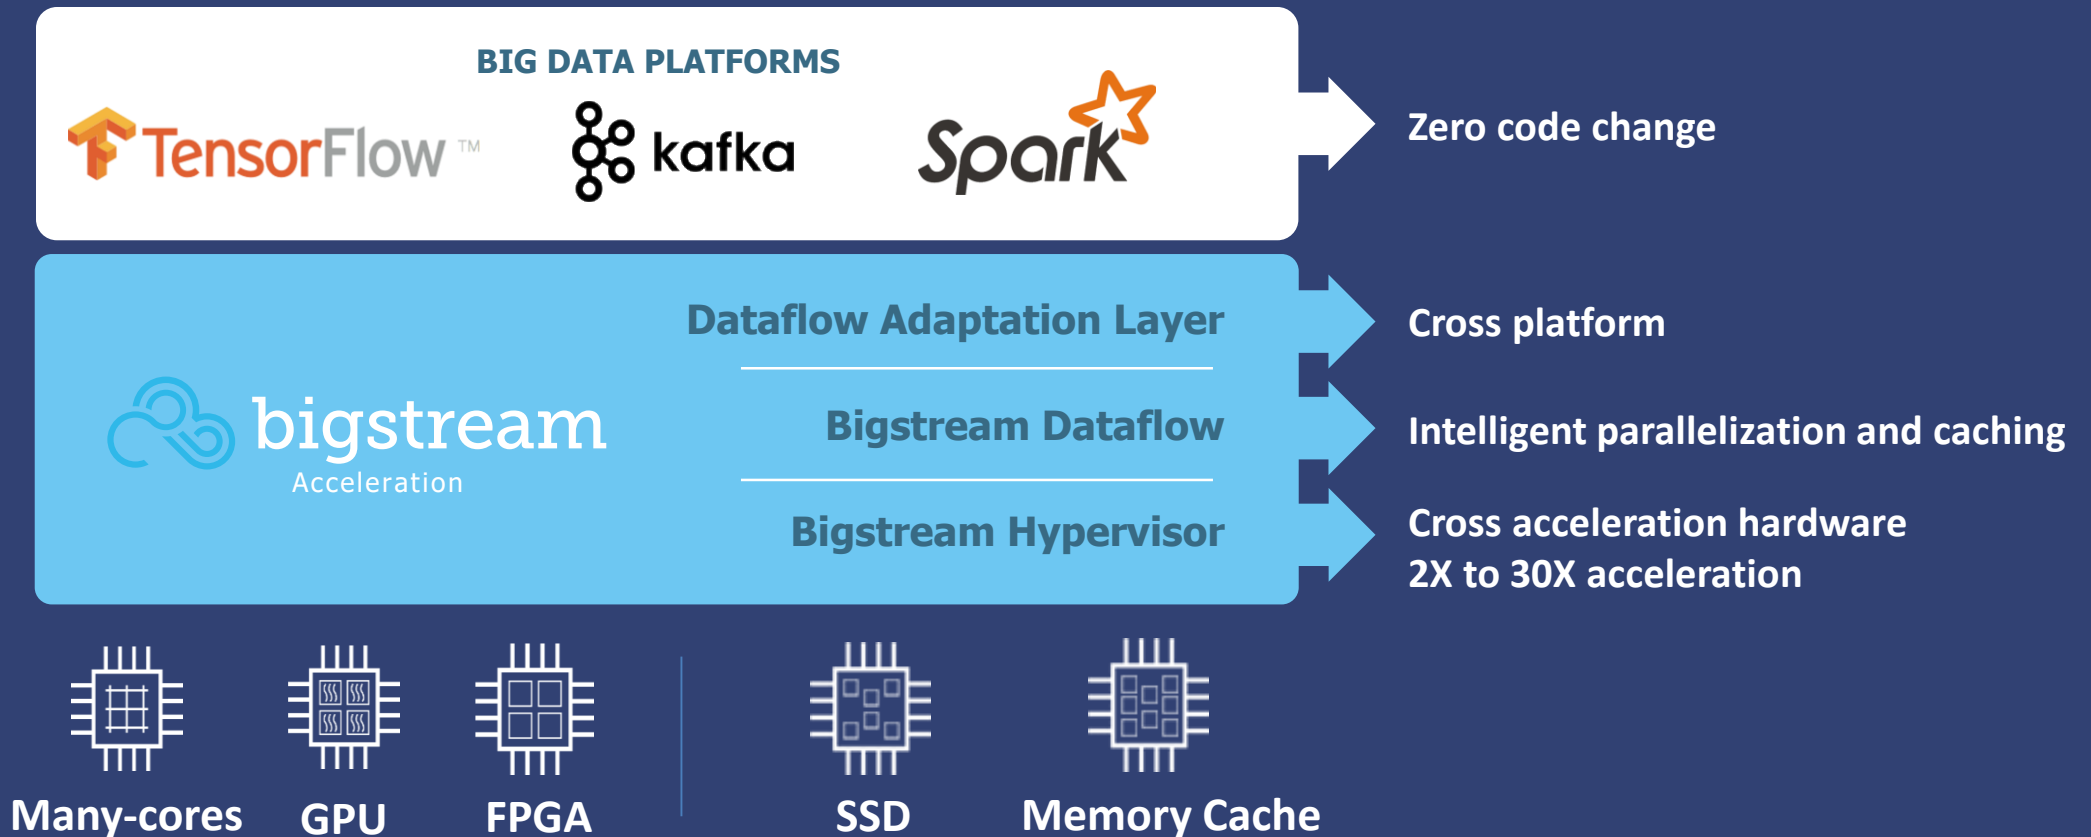
Skills Gap

Focus on  
Microarchitecture



Performance  
Engineers

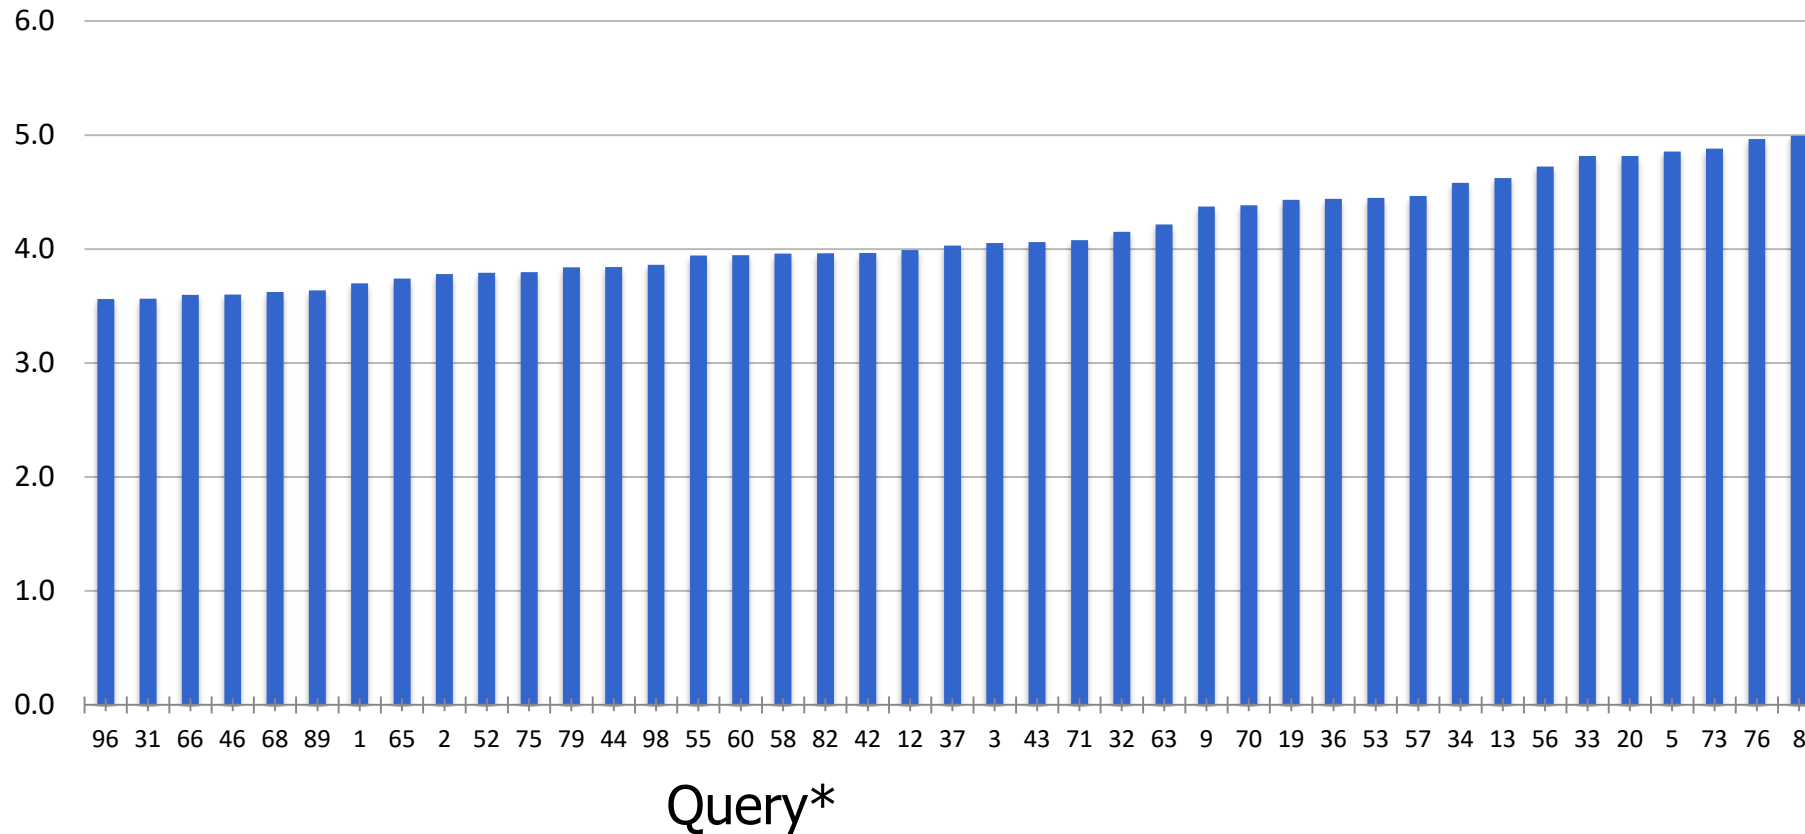
# Bigstream Hyper-acceleration Layer



# AWS F1 TPC-DS Speedup Results vs Spark

6

Spark vs. Bigstream Speedup For TPC-DS Queries on AWS F1



\*Not STAC Benchmarks

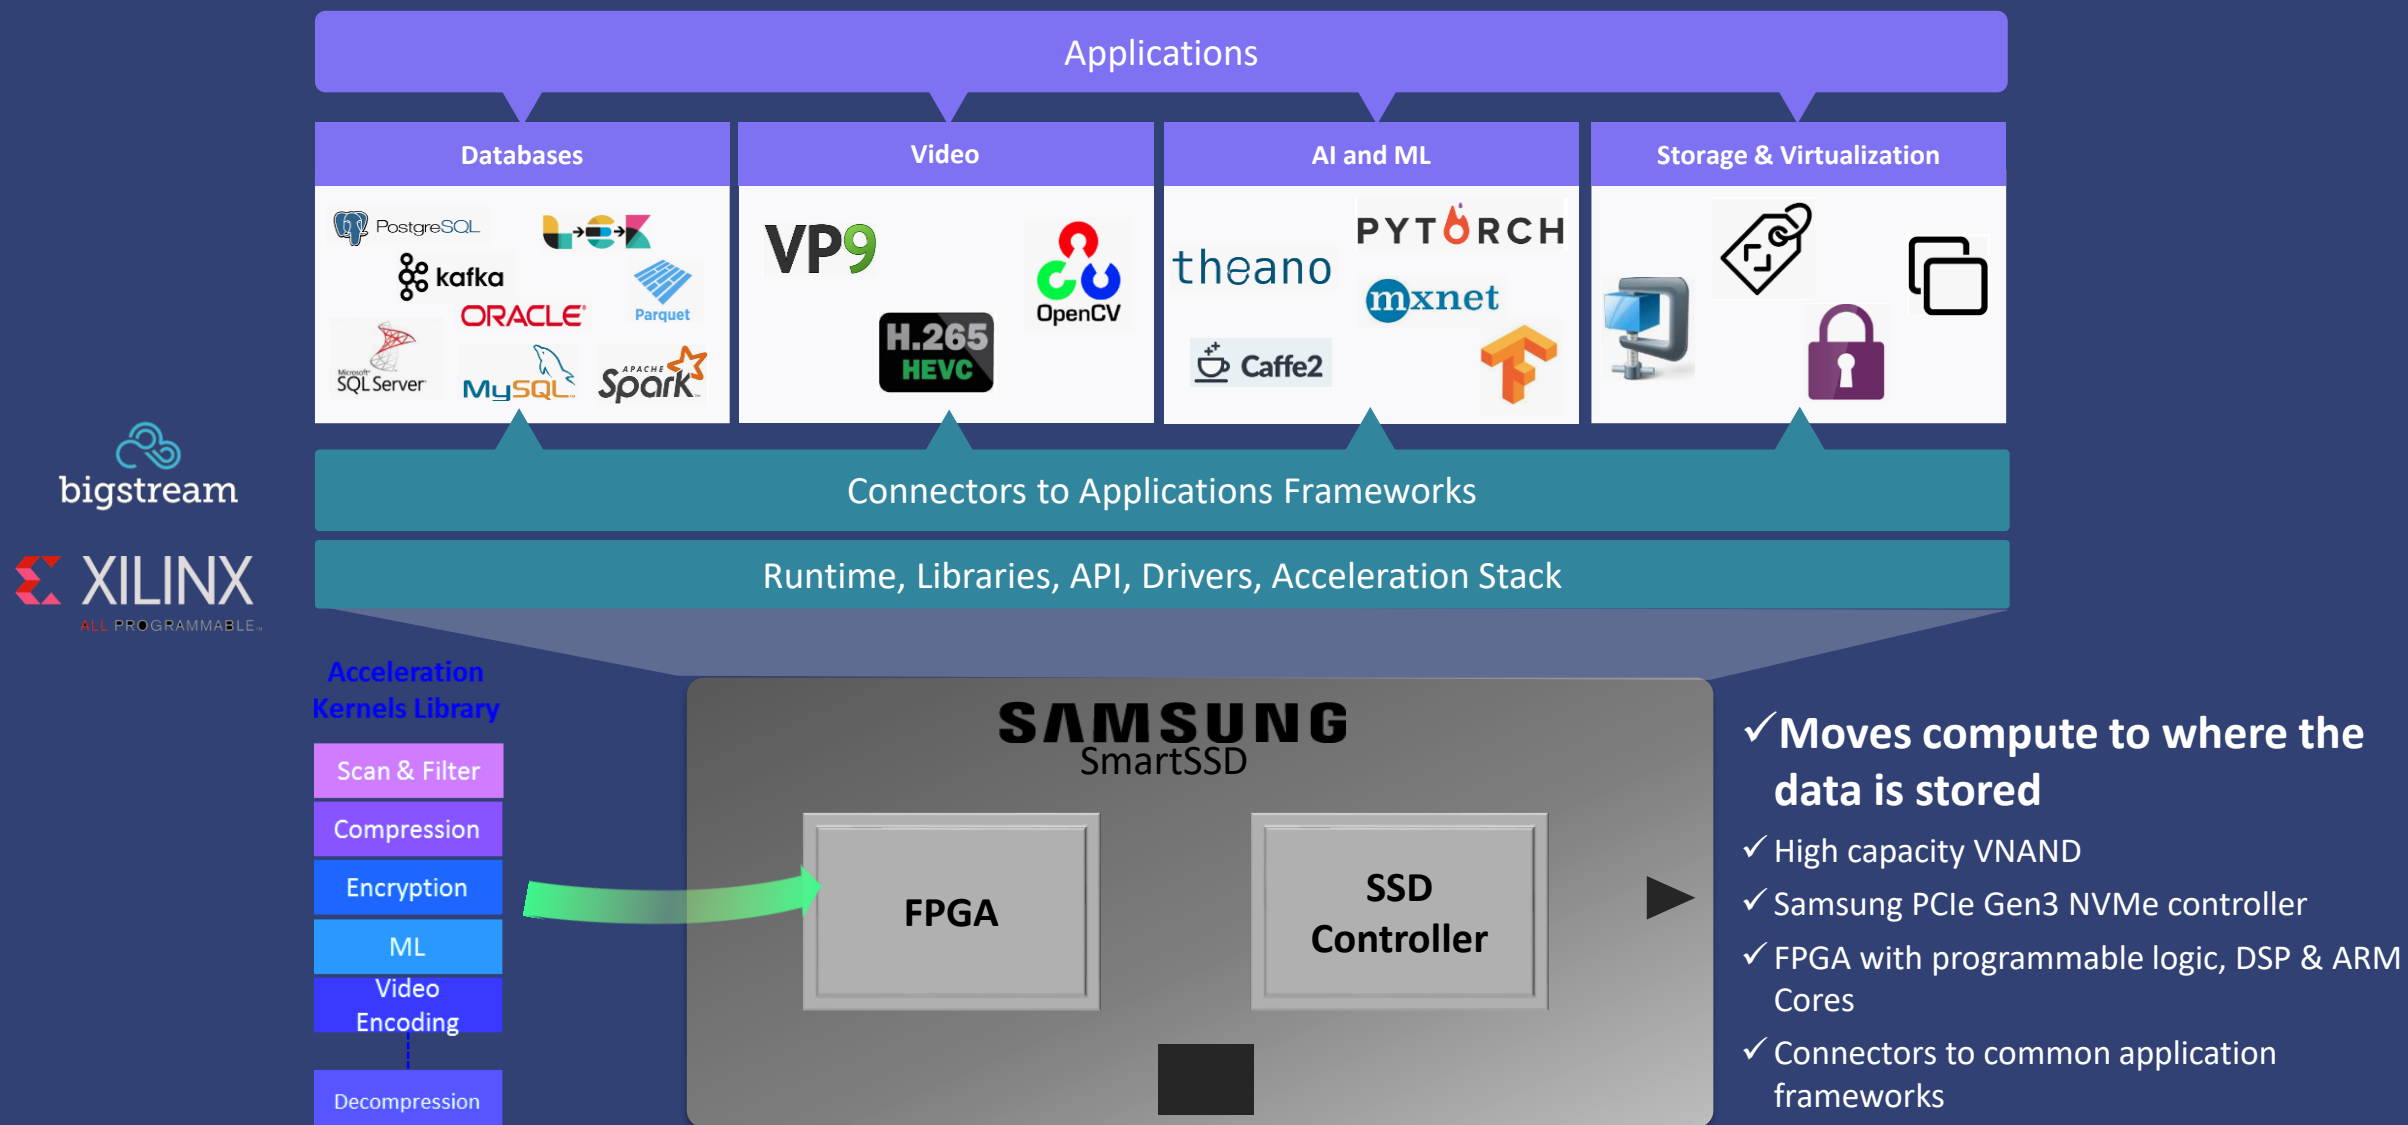
## • Summary

- Top 40 queries
- 40 sf JSON Data
- Server Config
  - Single f1.2xl
- Average
  - 4.3
- Maximum
  - 5.0
- Minimum
  - 3.2



# Samsung SmartSSD Platform

7




# Broad Range of SmartSSD Use Cases



Life Sciences and  
Genomics



Financial Services



Immediate Insights from  
Data Lakes



Instantaneous and Multi-Rules  
Fuzzy Search



Image Recognition and Object  
Detection



Real-Time Log Analytics



AI and ML



Speedup Complex Ad-hoc  
Queries



Enhance Business Intelligence &  
Data Warehousing



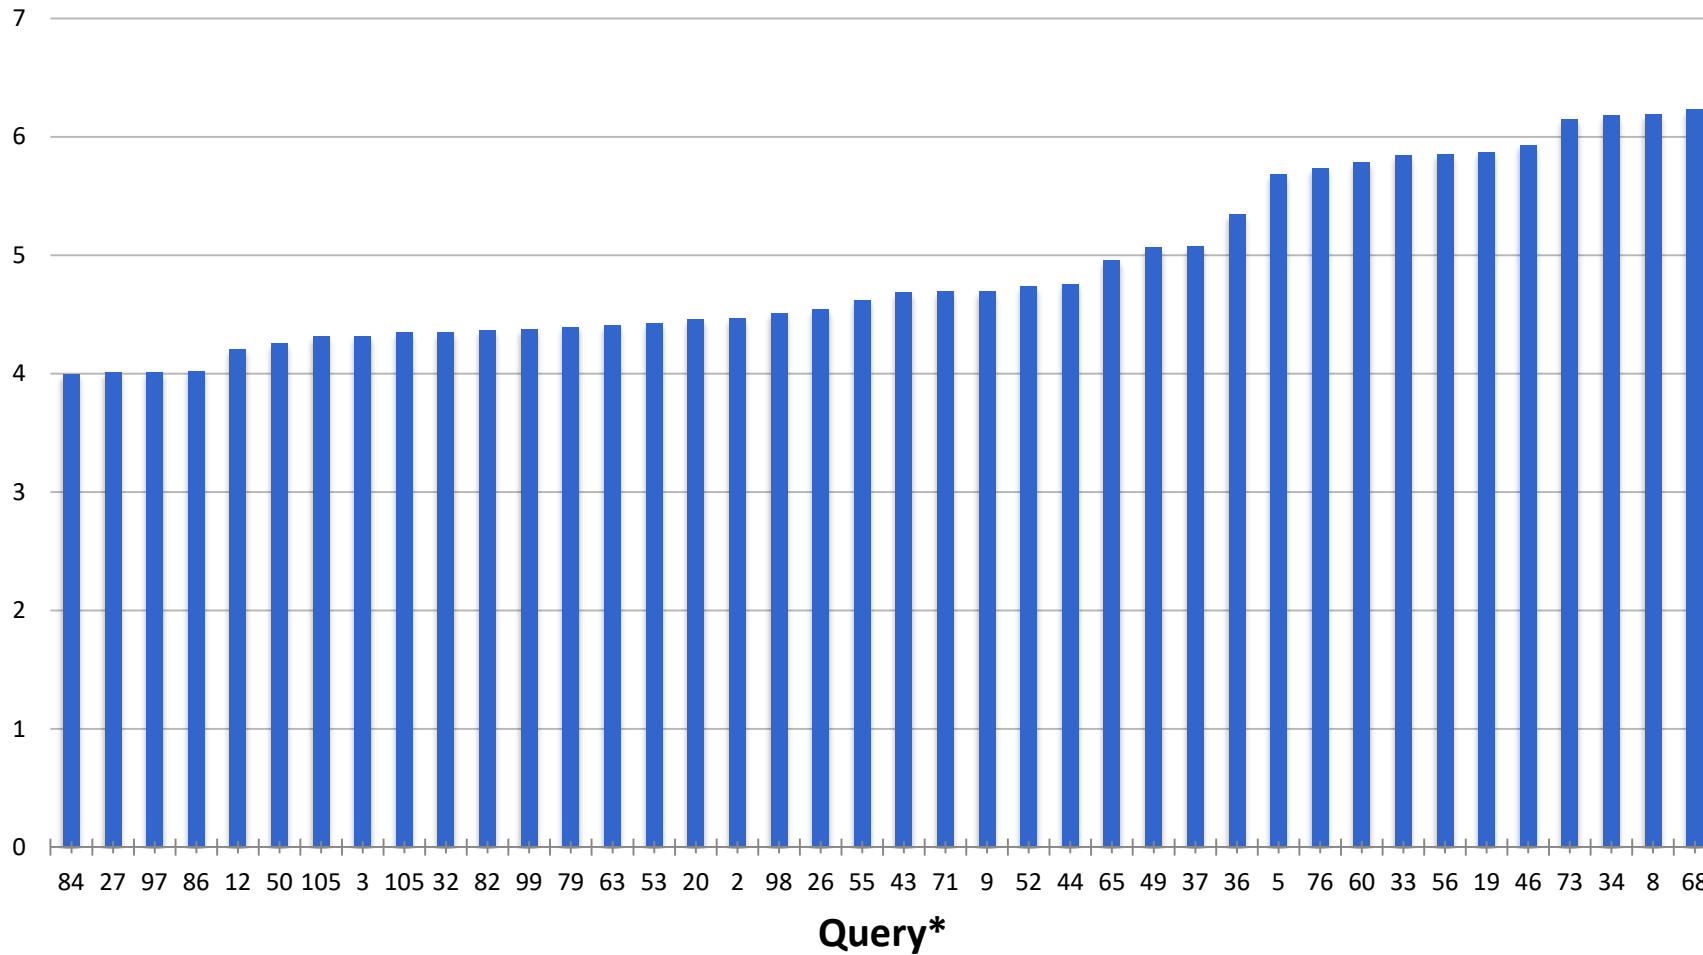
Video Transcoding



# Samsung SmartSSD Speedup Results vs SSD

9

Spark vs. Bigstream Speedup For TPC-DS Queries on SmartSSD



\*Not STAC Benchmarks

- Summary

- Top 40 Queries
- 200 sf JSON Data
- Server Config
  - Single SmartSSD
- Average
  - 4.9
- Maximum
  - 6.2
- Minimum
  - 4.0

**Thank You!**


**Check off [Bigstream](#) to start the conversation**  
**[roop@bigstream.co](mailto:roop@bigstream.co)**

# Samsung SpartSSD Results

1  
1

- **Summary**

Spark vs. Bigstream Speedup For TPC-DS Queries on SmartSSD (200sf Json)

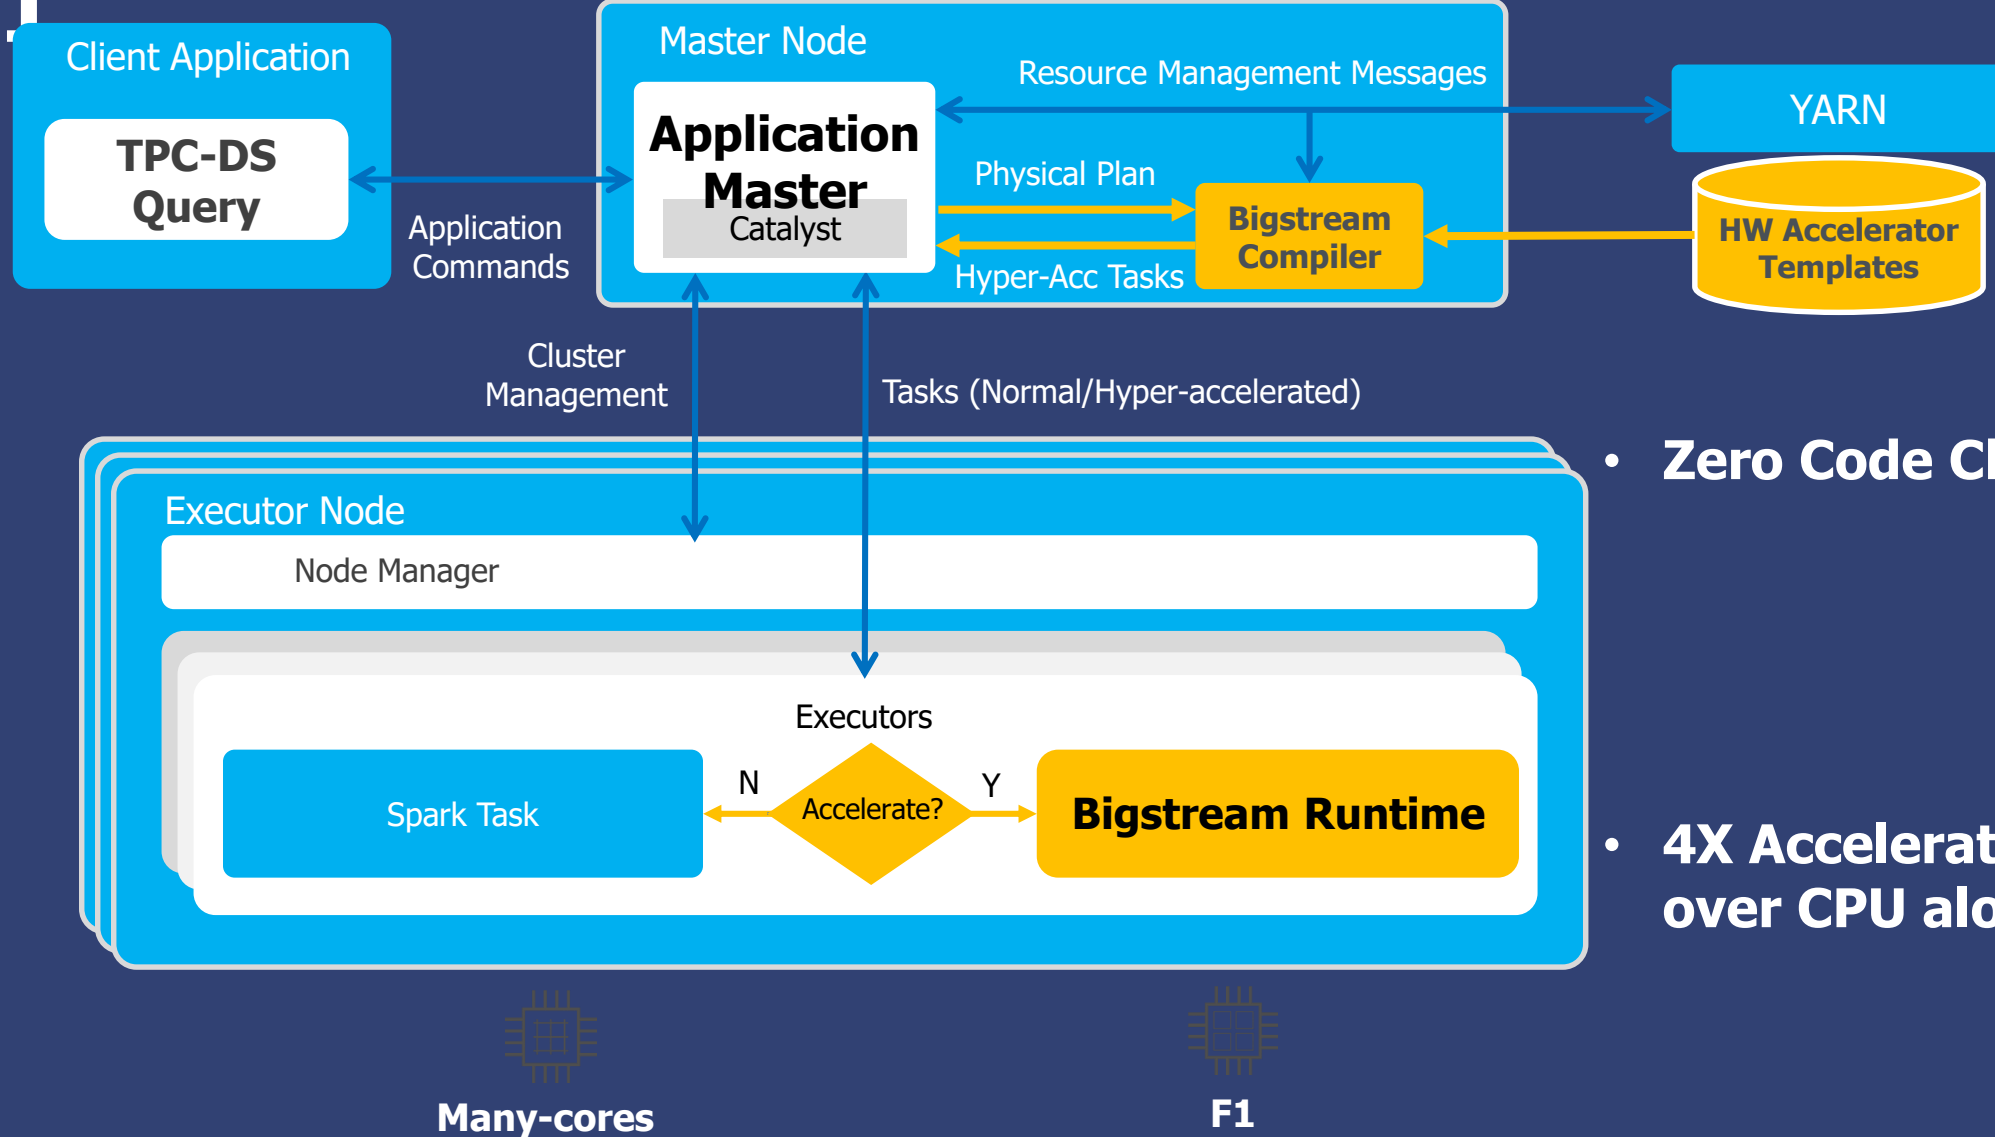


- **All 100 Queries**
- **200 sf JSON Data**
- **Server Config**
  - **Single SmartSSD**
- **Average**
  - **3.9**
- **Maximum**
  - **6.2**
- **Minimum**
  - **1.0**

# Bigstream Acceleration of TPC-DS Spark On AWS

12

F1



• **Zero Code Change**

• **4X Acceleration over CPU alone**