

Bigstream Hyperacceleration

The Analytics Performance Advantage




Spark Use Cases Across Financial Services Sectors

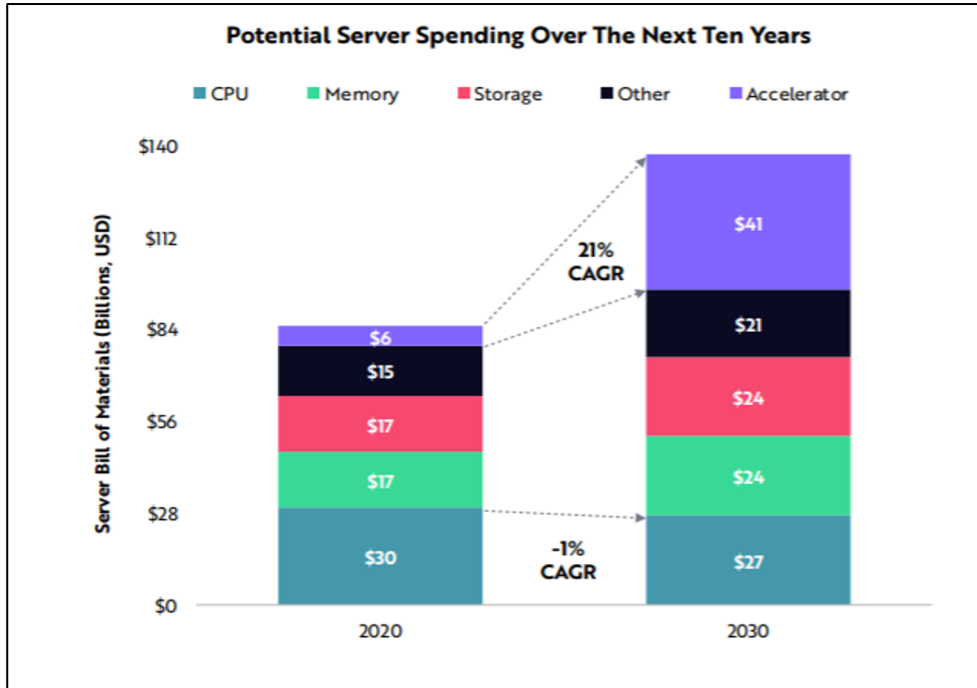
<u>Financial Services Sector</u>	Customer Insights & Retention	Fraud Detection	Valuation & Market Models	Risk & Compliance	Telematics
Retail Banks	X	X		X	
Payments	X	X			
Investment banks			X	X	
Investment Management	X		X	X	
Insurance	X	X			X
Tech / FinTech	X	X			
Exchanges / Govt		X	X	X	

Hardware Accelerators

- **Problem: CPU cluster scaling is running into limitations**
- **Programmable hardware is an attractive alternative: FPGA, GPU**
- **Designed for efficient execution of specialized code. A departure from general-purpose CPUs**
- **FPGA, GPU: provide flexibility, but don't natively connect to big data platforms**
 - **Middleware needed to automate this**
- **Both can provide power and cost advantages w.r.t CPU scale-up, scale-out**
- **Designed to physically attach to existing servers simply (i.e. PCIe slot)**



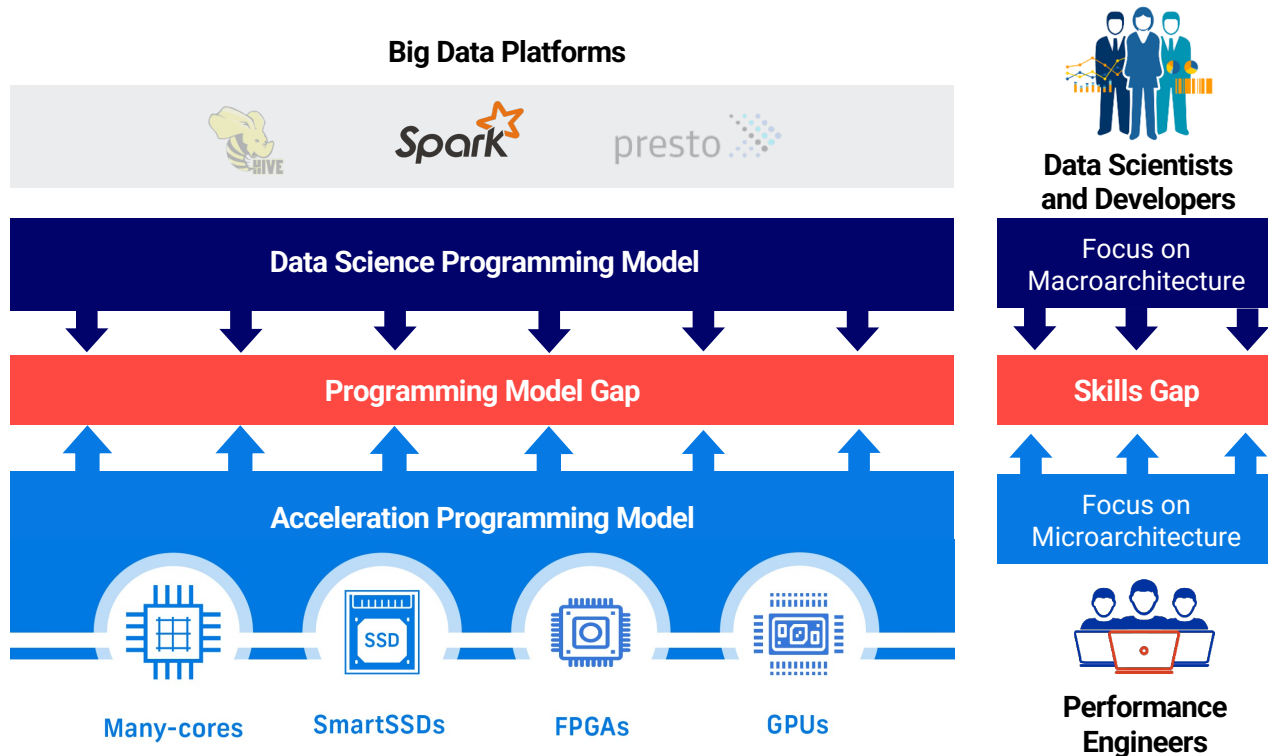
HW Accelerator Market is Trending



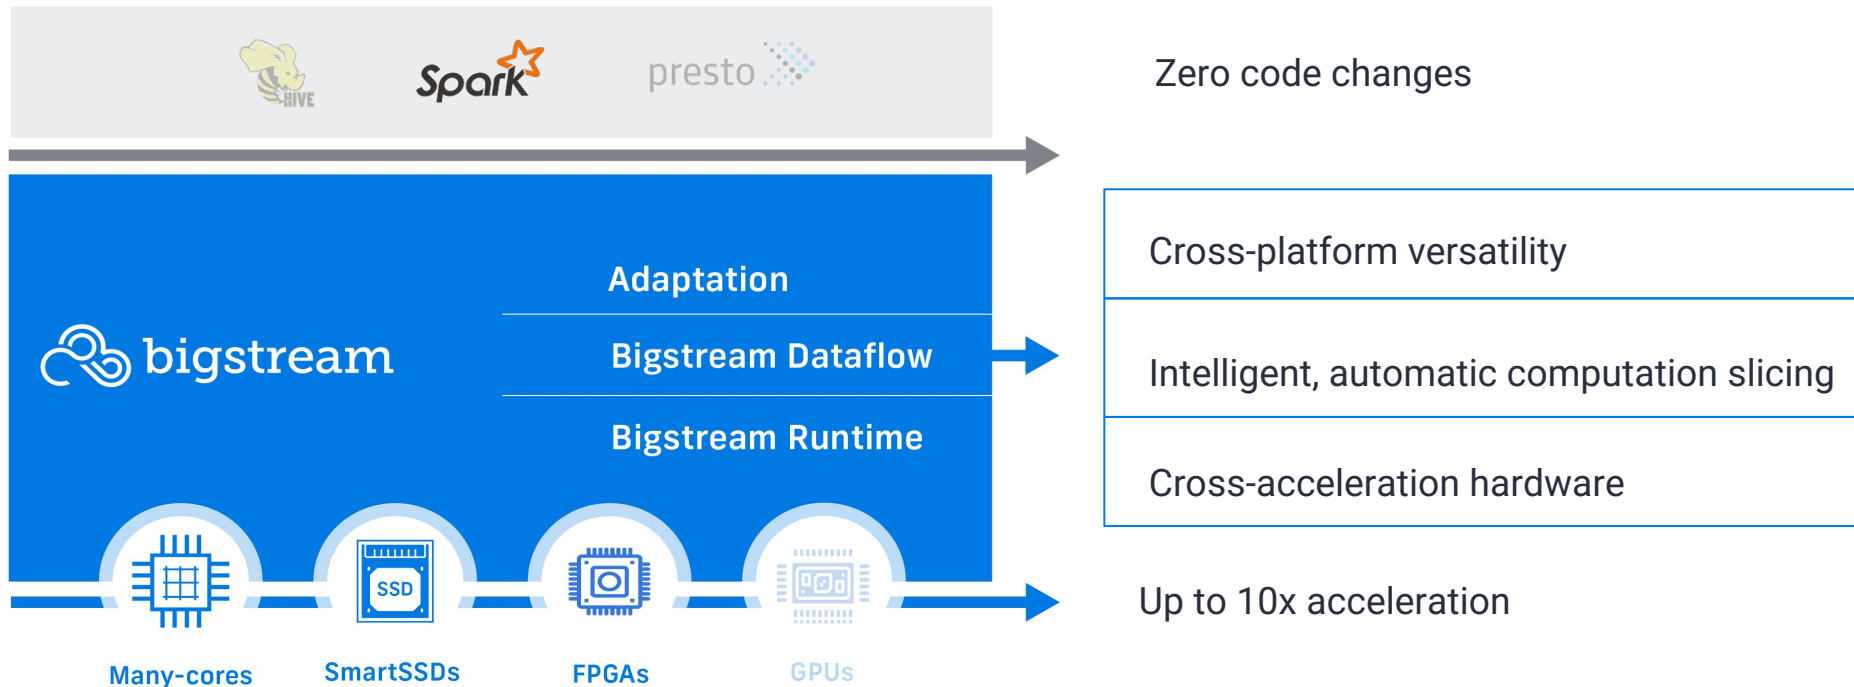
ARK Invest "Big Ideas 2021"

- Accelerators = GPUs, ASICs, and FPGAs
- \$41 Billion industry in next 10 years, surpassing CPUs.
- Driven by big data analytics and AI

Programming Model Gap



Bigstream Hyperacceleration Layer

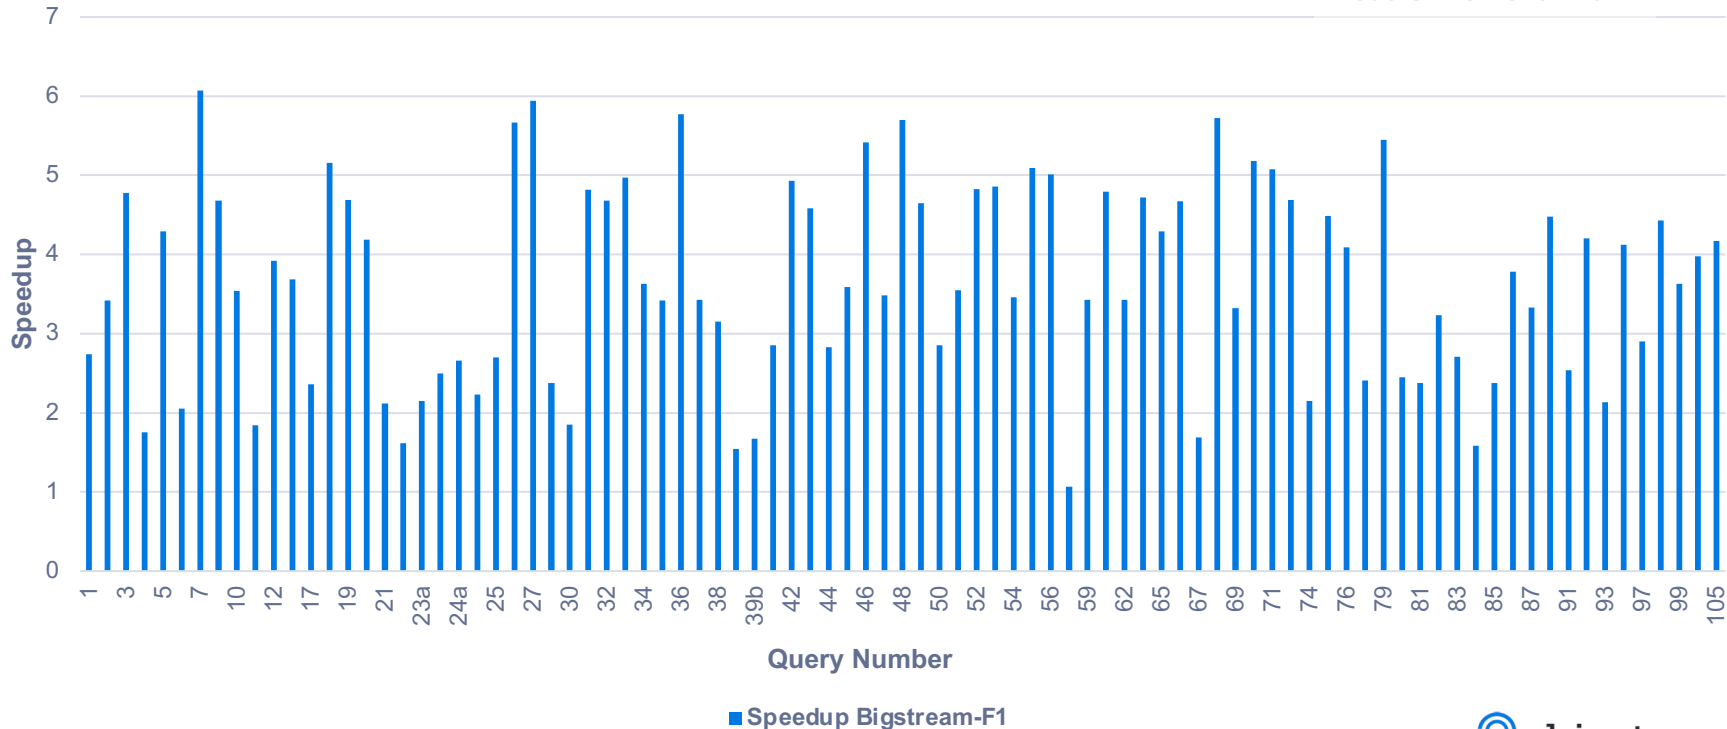


Bigstream-F1 FPGA Speedup Results

- All runs on identical md5.2xl EMR clusters - Baseline: Spark 3.0
- 4 workers/cluster, S3 storage
- 250SF CSV gzip compressed standard data (approx. 72GB)

Bigstream Average Speedup: 3.6x

Not a STAC Benchmark




Bigstream-SmartSSD Speedup Results


- All runs an identical Intel-based server - Baseline: Spark 2.4.6
- Single server, 3 SmartSSD
- 250SF JSON gzip compressed standard data (approx. 110GB)

Bigstream Average Speedup: 4.3x

Not a STAC Benchmark



Bigstream Spark Acceleration: Options and Roadmap

 Deployment Option	Software-based Acceleration	Hardware-based Acceleration	
		FPGA	SmartSSD
Cloud (AWS)	✓ AWS EMR	✓ AWS EC2 + F1	
Cloud(Azure)	1H '21	1H '21	
On-Prem	✓	✓	✓

Primary Use Cases (current customers)

- Ingest / ETL / ELT / ML Prep
- Batch Analytics

Secondary Use Cases (road mapped)

- Ad Hoc Analytics
- ML Training

Summary/Next Steps

- Bigstream: most mature and comprehensive Spark acceleration portfolio
 - From complex to easy
 - From fast to fastest
- Best customer fit - deployments
 - Managed cloud: Bigstream + EMR = plug & play, start today at Marketplace
 - FPGA on cloud – single script deploy
 - On prem – single script deploy
- Guaranteed TCO reduction vs existing infrastructure
- Fastest Spark with zero code change

Thank You!

Please select **Bigstream** to start the conversation
roop@bigstream.co

