HITACHI
Inspire the Next

Accelerating Business Analytics with Flash Storage and FPGAs

Satoru Watanabe

Center for Technology Innovation
- Information and Telecommunications
Hitachi, Ltd., Research and Development Group

Aug.10 2016

© Hitachi, Ltd. 2016. All rights reserved.

Disclaimer

HITACHI
Inspire the Next

- The contents of this presentation are based on early research results.
- This presentation does not reflect any product plan or business plan.
- All information is provided *as is*, with no warranties or guarantees.

© Hitachi, Ltd. 2016. All rights reserved.

Overview of Pentaho

- Pentaho has two components: data integration and business analytics.
- This presentation focus on business analytics on analytical database.

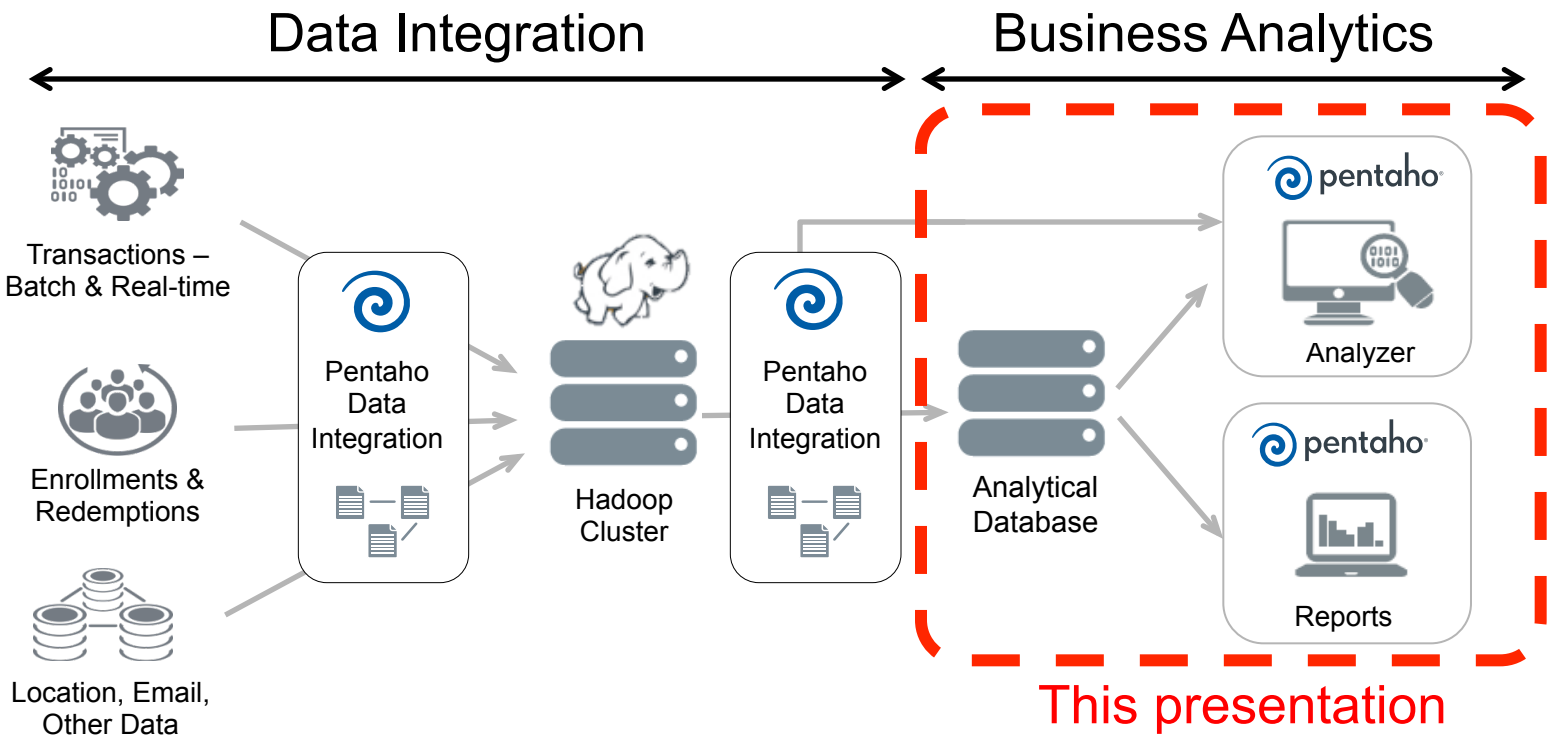

© Hitachi, Ltd. 2016. All rights reserved.

Use Case: Self-Service Business Analytics (BA)

HITACHI
Inspire the Next

- Spreading for improving daily business operations.
- Needs more powerful analytical database(DB) than traditional BA to support daily or interactive analytics and hundreds to thousands of users.

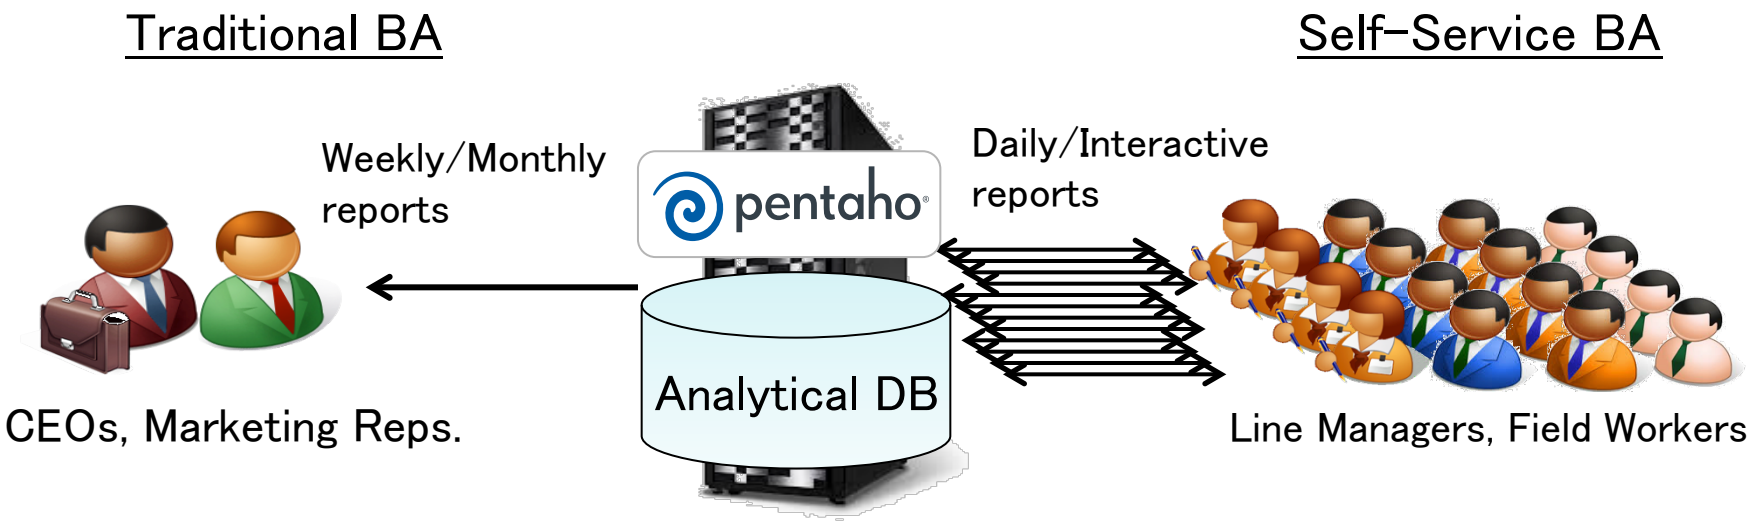

© Hitachi, Ltd. 2016. All rights reserved.

Bottleneck of Data Analytical Systems

HITACHI
Inspire the Next

- New technologies have changed system bottleneck from storage to CPU.

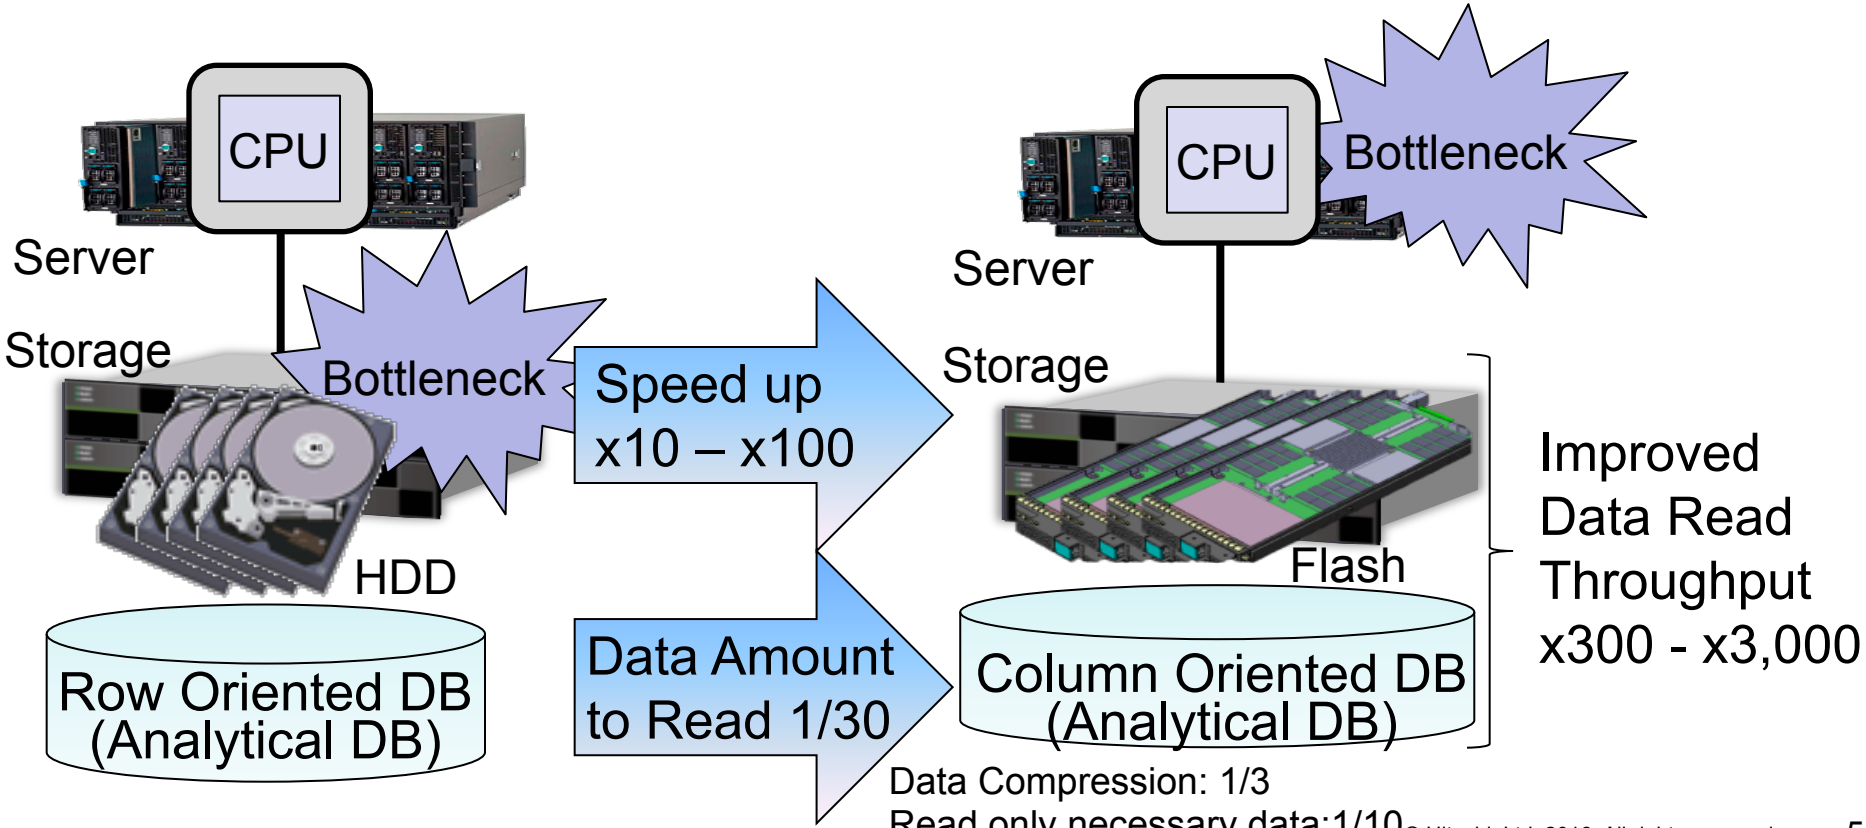

© Hitachi, Ltd. 2016. All rights reserved.

Combining CPUs and FPGAs to Pioneer New Computer Architecture

HITACHI
Inspire the Next

- FPGA accelerators overcome the CPU bottleneck.
- Widely used in-memory DB does not ease the CPU bottleneck.

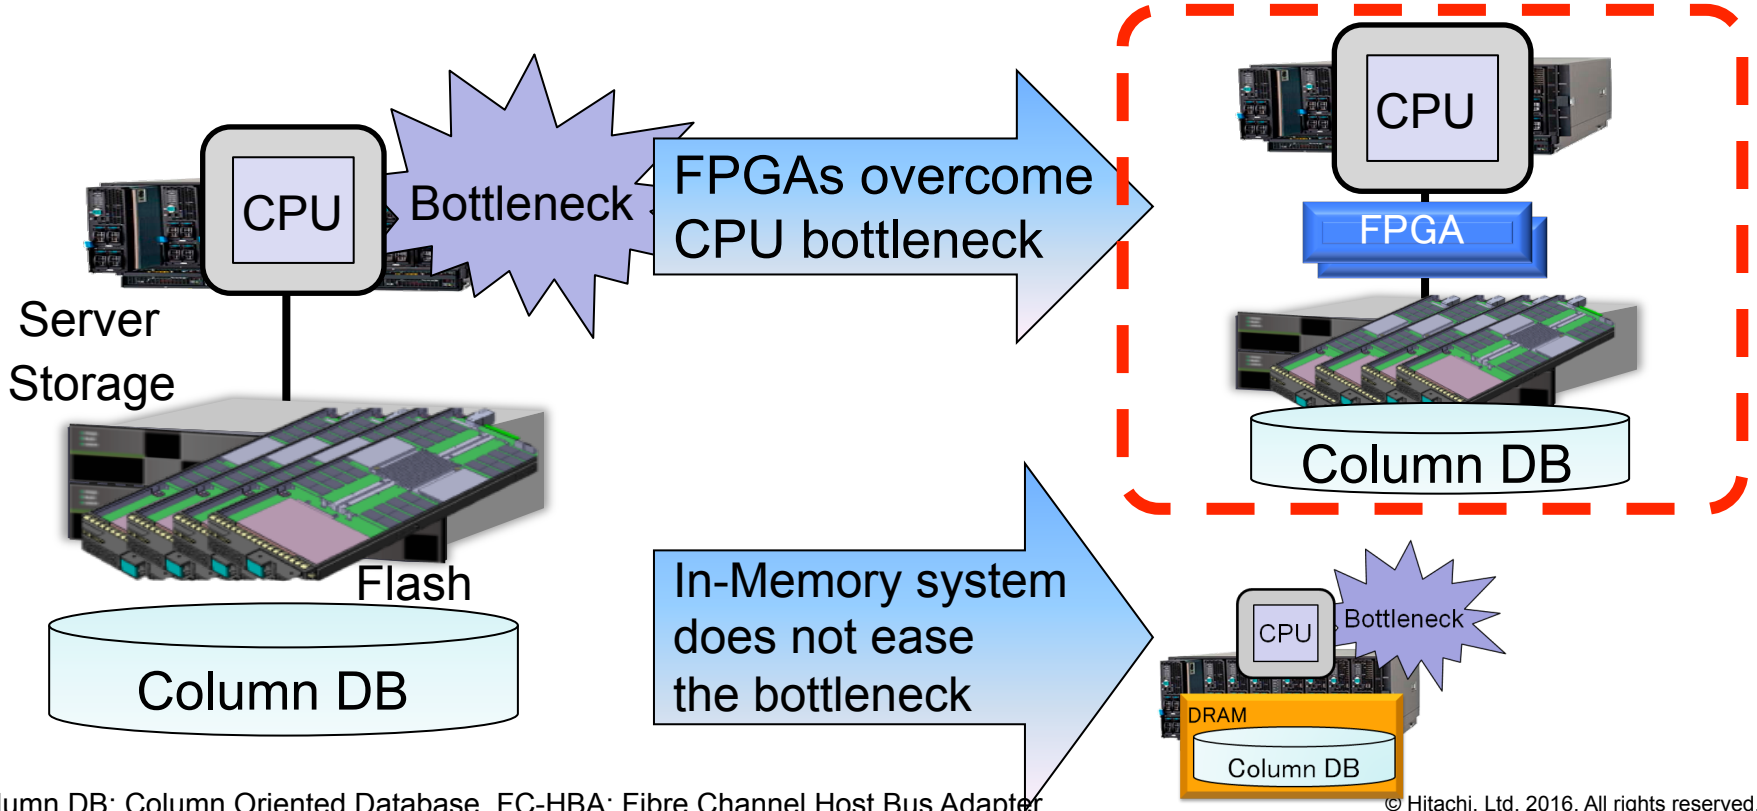

Column DB: Column Oriented Database, FC-HBA: Fibre Channel Host Bus Adapter

© Hitachi, Ltd. 2016. All rights reserved.

Demonstration

- BioMe is a sample application of Pentaho business analytics.
- Healthcare information is explored on a web with GUI.

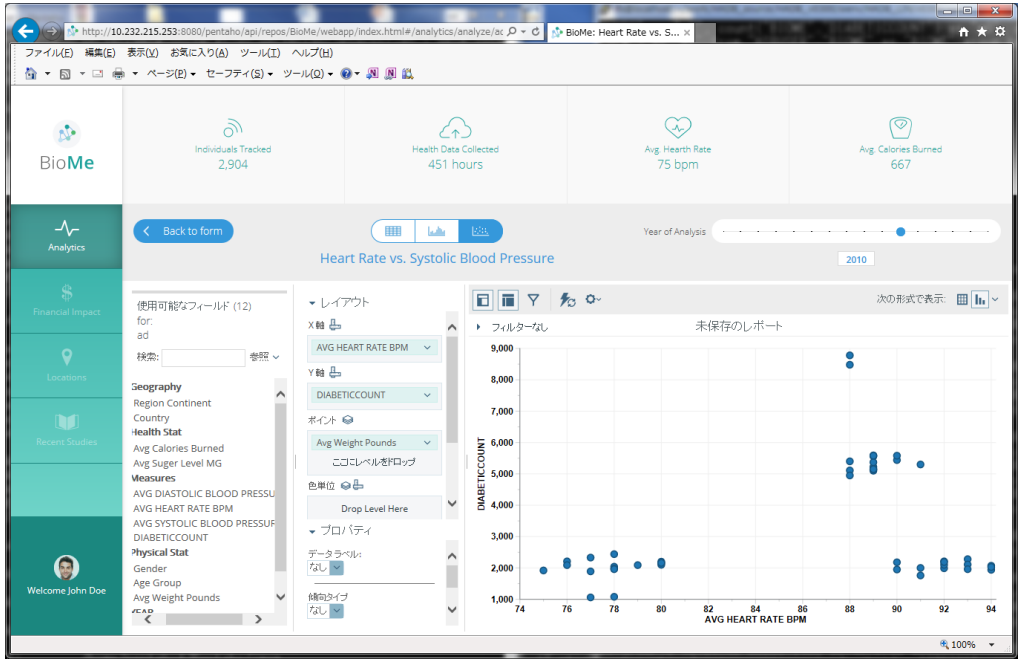

GUI: Graphical User Interface

© Hitachi, Ltd. 2016. All rights reserved.

Performance Comparison

HITACHI
Inspire the Next

- CPU bottleneck limits performance of In-memory Database.
- FPGA accelerator overcomes CPU bottleneck

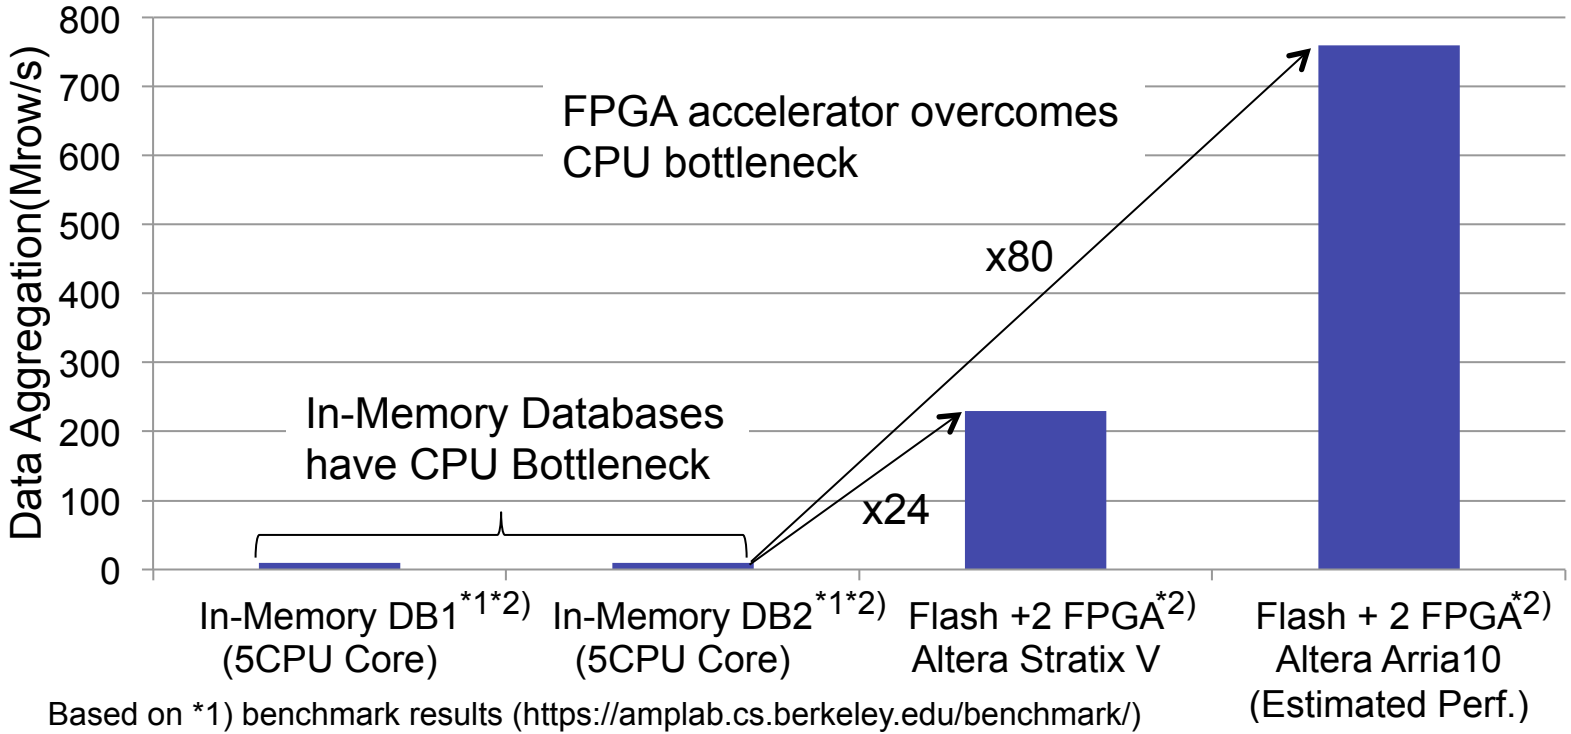

Based on *1) benchmark results (https://amplab.cs.berkeley.edu/benchmark/) and *2) performance measurements by CTI, Hitachi, Ltd. R&D Group

© Hitachi, Ltd. 2016. All rights reserved.

FPGA Accelerator between Server and Storage

HITACHI
Inspire the Next

- Filter & Aggregation are major operations in Analytics and have bottleneck in CPU.
- FPGA(DB accelerator) accelerates them by parallel and pipeline processing.
- Our FPGA-Column-DB runs behind Pentaho BA and is transparent to users.

13 patent applications have been filed

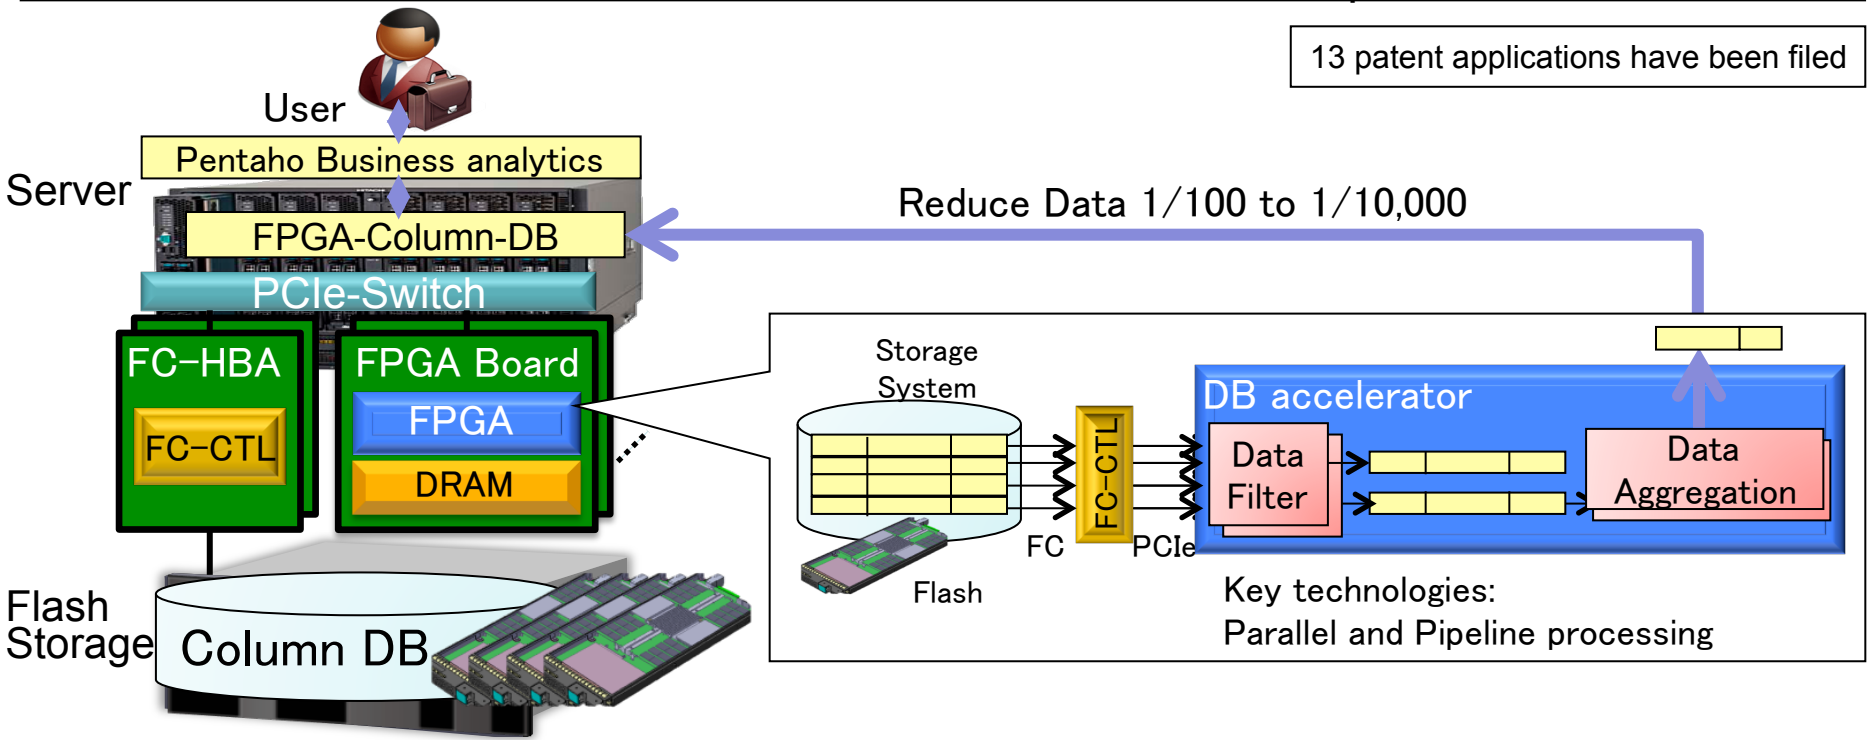

Column DB: Column Oriented Database, BA: Business Analytics, FC-HBA: Fibre Channel Host Bus Adapter

© Hitachi, Ltd. 2016. All rights reserved.

Implementation

HITACHI
Inspire the Next

- Pentaho business analytics issues SQLs to FPGA-Column-DB.
- FPGA-Column-DB interprets SQL to control FPGA with custom command on NVMe.

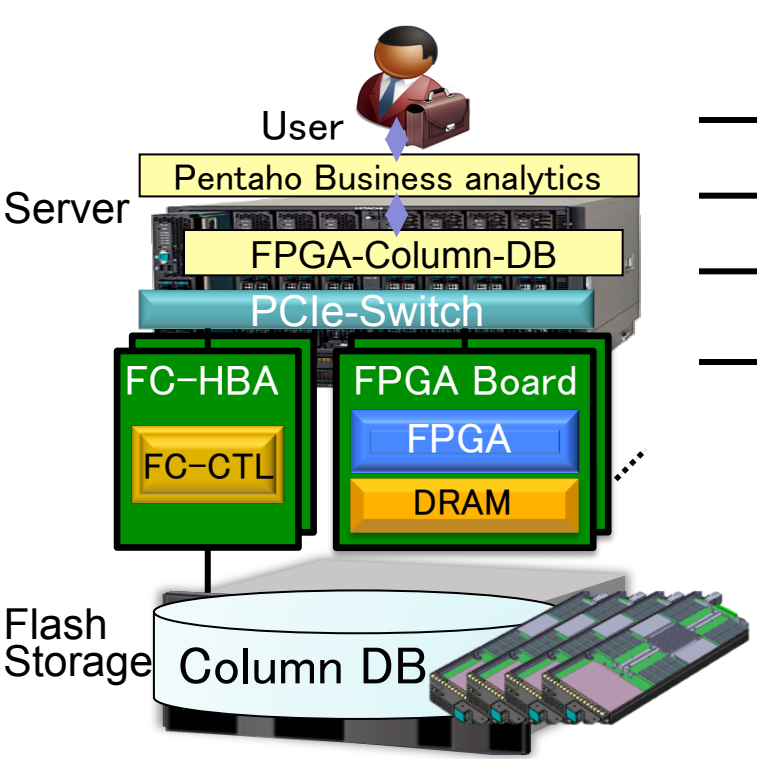

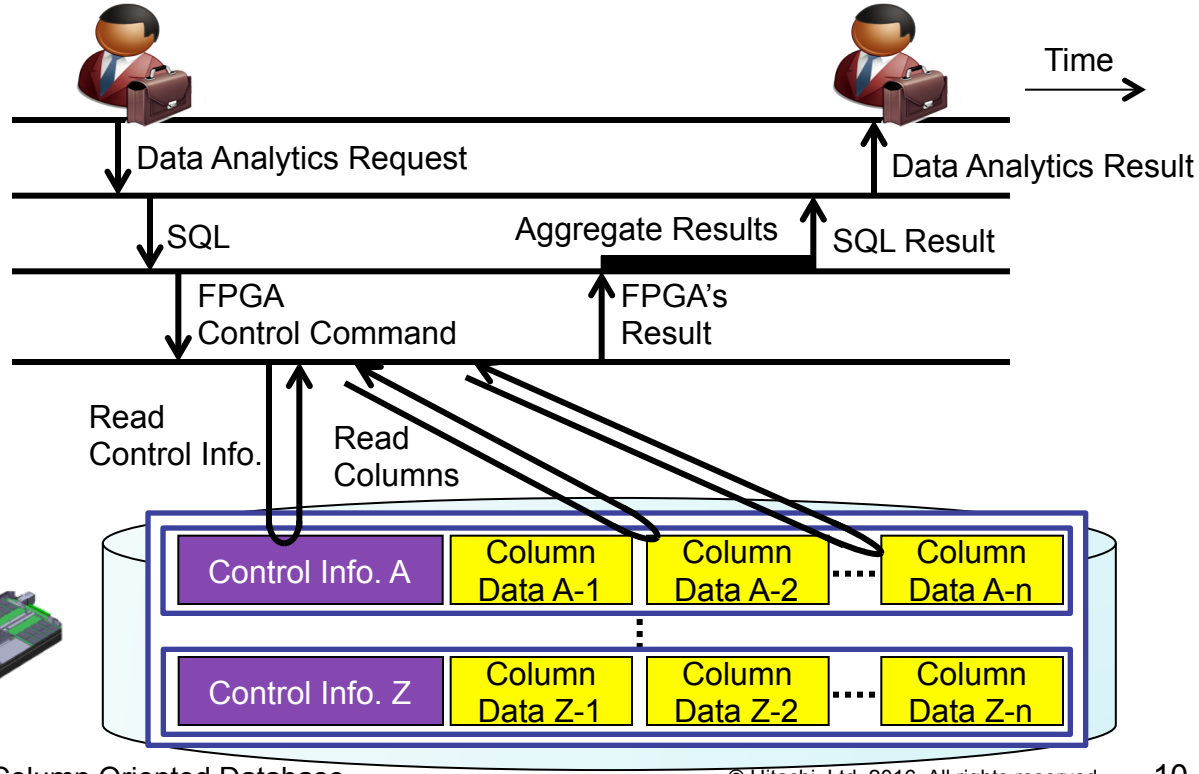

Column DB: Column Oriented Database

© Hitachi, Ltd. 2016. All rights reserved.

Core Technology

HITACHI
Inspire the Next

- Frequent accesses to DRAM connected to FPGA cause performance deterioration.
- We developed column oriented database optimized for FPGA processing.
- Control information fitted in SRAM size minimizes DRAM accesses to maximize performance.

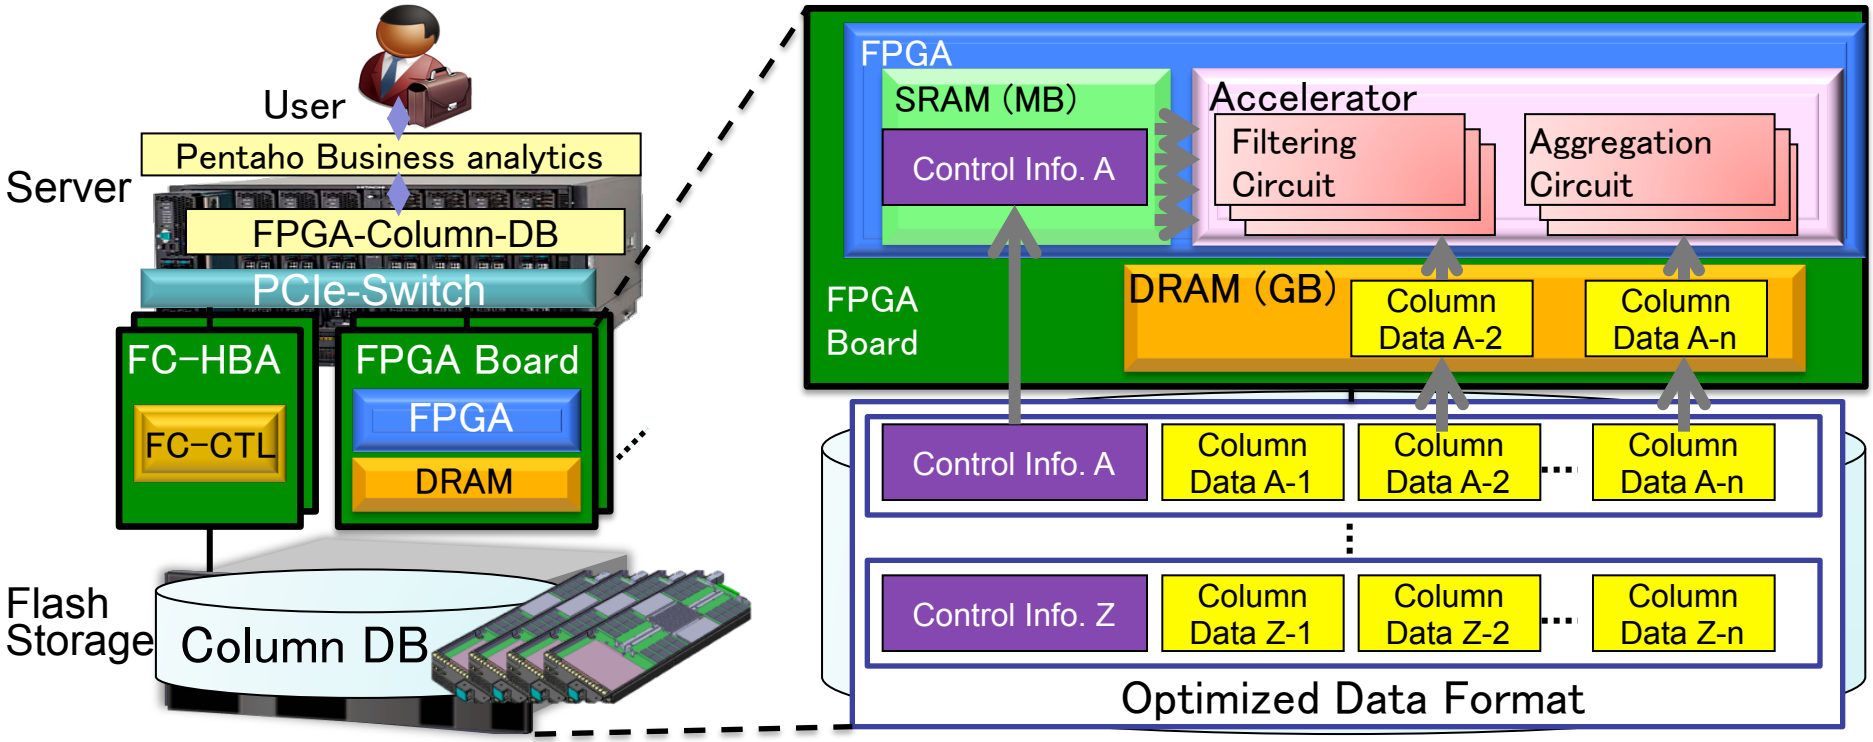

Column DB: Column Oriented Database

© Hitachi, Ltd. 2016. All rights reserved.

Storage Database vs. In-memory Database

HITACHI
Inspire the Next

- In-memory DB has higher data throughput than Flash, but bottleneck still remains in CPU.
- Moreover, cost of DRAM is over x10 times higher than Flash Storage.
- In-memory database needn't be used when there is a CPU bottleneck.

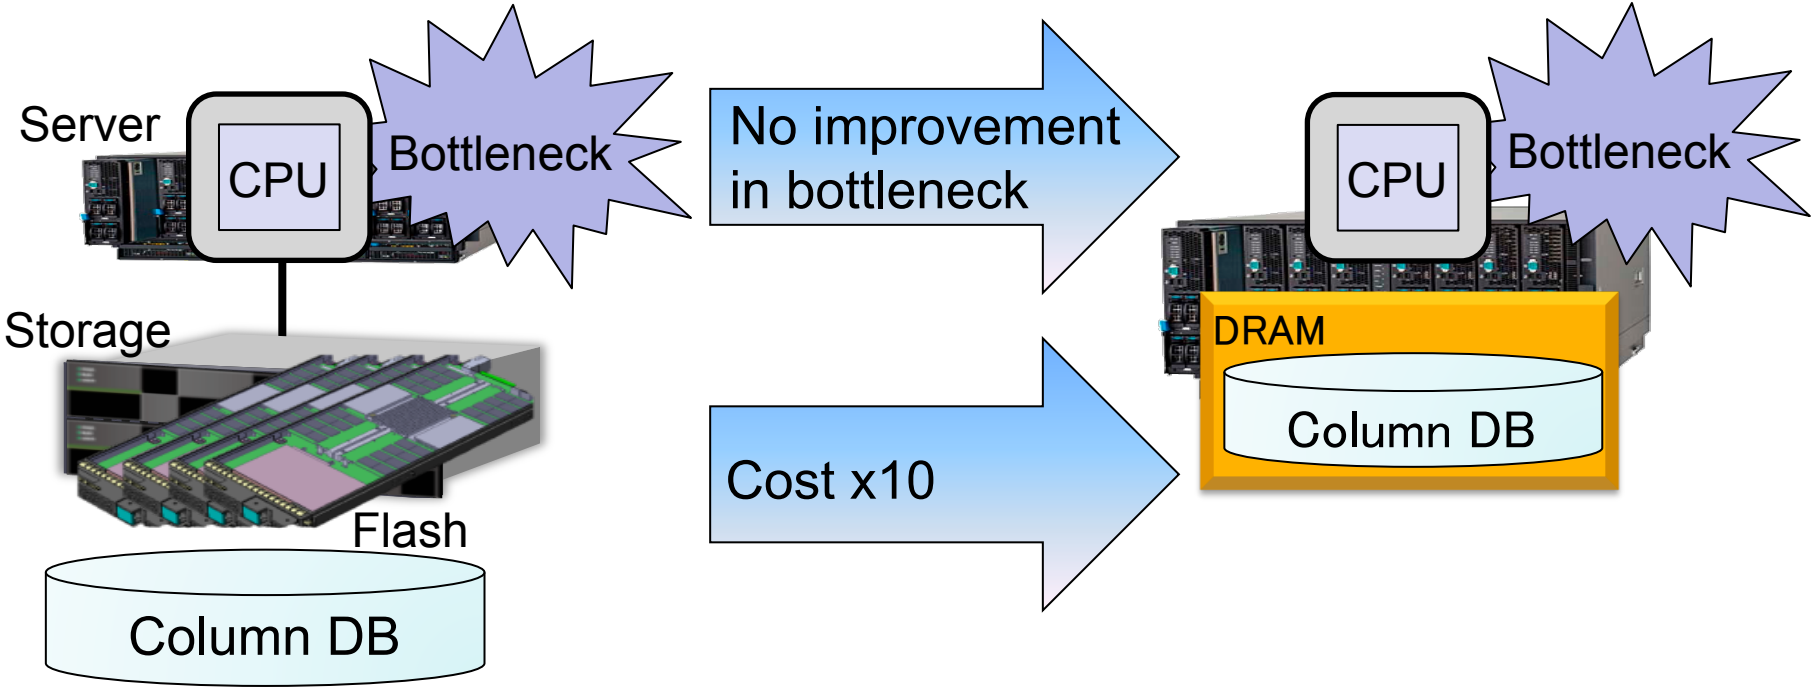

Column DB: Column Oriented Database

© Hitachi, Ltd. 2016. All rights reserved.

Bottleneck Detail of In-memory Database

HITACHI
Inspire the Next

- 40-CPU-core-server requires 800MB/s – 8GB/s.
- About x10 higher performance of DRAM (76.8GB/s) is not required in such a data analytics.

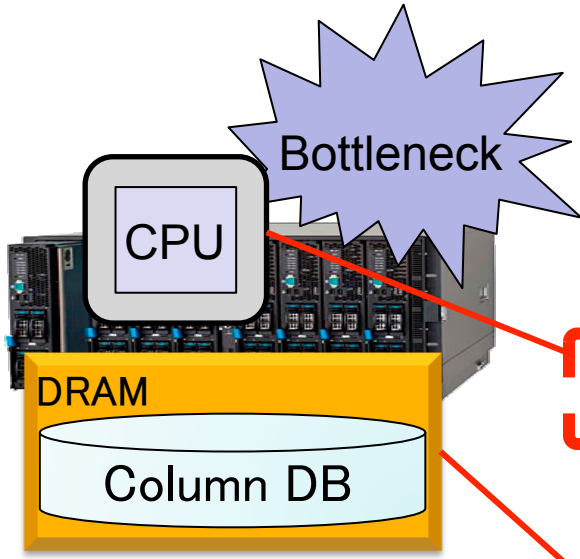

CPU Performance

	Performance	Basis and Assumption
1 CPU core	1M row/s/core – 10M row/s/core	Benchmark Results (https://amplab.cs.berkeley.edu/benchmark/) Measurement Results at Hitachi Lab.
1 CPU core	20 MB/s/core – 200 MB/s/core	Assume 1 row = 20 Bytes, because of data compression and reading data necessary for analytics
40 CPU Cores	800MB/s – 8GB/s	2-Socket-Server (two 20-core-CPUs, E.g. : Xeon E5-2698 v4)

DRAM Performance

	Performance	Basis and Assumption
4ch-DDR4	76.8GB/s	E.g.: Xeon E5-2698 v4, DDR4 2400MHz

Column DB: Column Oriented Database

© Hitachi, Ltd. 2016. All rights reserved.

FPGA Accelerators Exploit Flash Storage's Ability

HITACHI
Inspire the Next

- FPGA accelerators fill the gap between flash storage and server-CPU.
- Our 8 FPGA accelerators will balance with high performance flash storage.

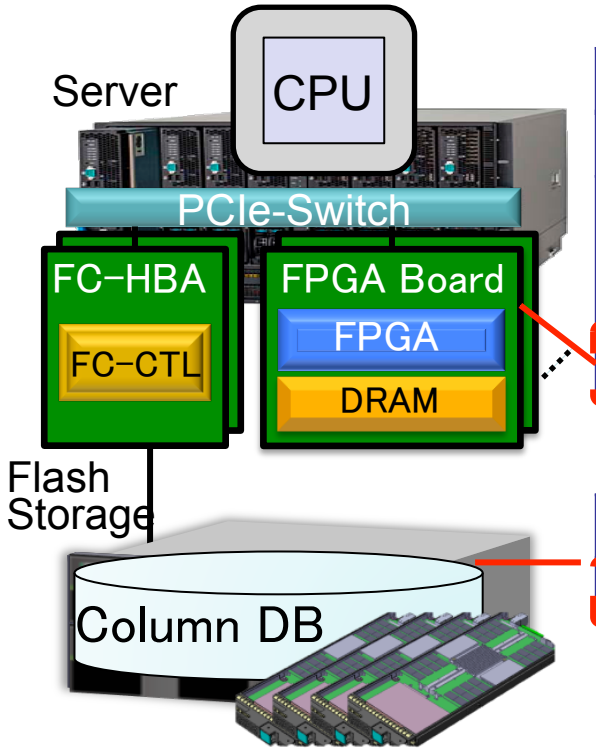

FPGA Performance

	Performance	Basis and Assumption
1 FPGA	380Mrow/s	Our Estimated Performance of Arria10
1 FPGA	7.6 GB/s	Assume 1 row = 20 Bytes, because of data compression and reading data necessary for analytics
8 FPGAs	48 GB/s	= 7.6 x 8, but limited by PCIe BW (6 x 8)

Flash Storage Performance

	Performance	Basis and Assumption
RD throughput	48 GB/s	16Gbps Fibre Channel x 32

Column DB: Column Oriented Database

© Hitachi, Ltd. 2016. All rights reserved.

Transaction System and Analytical Database

HITACHI
Inspire the Next

- In-Memory DB suits transaction systems because random read and write are dominant.
- Flash Storage + FPGA accelerators suits analytical database because sequential read is dominant.

	Transaction System	Data Analytical System
Data access pattern	Random Read / Write	Sequential Read
In-memory database	✓	
Flash Storage + FPGA Accelerators		✓

© Hitachi, Ltd. 2016. All rights reserved.

FPGA and Other Accelerators

HITACHI
Inspire the Next

- FPGAs have superior energy efficiency and performance.
- FPGA acceleration can be applied to database processing in data centers to save energy.

	Advantages
FPGA	✓High energy efficiency ✓High performance
Processor (ARM, Phi)	✓Easy to change data processing ✓Easy to develop programs for data processing
Graphic Processor Unit	✓Multiple parallel calculations ✓Good development environment

© Hitachi, Ltd. 2016. All rights reserved.

Summary

HITACHI
Inspire the Next

- We proposed a system architecture integrating flash storage and FPGAs for accelerating business analytics.
- Flash storage provides sufficient read throughput for business analytics, and FPGAs overcome CPU bottleneck.
- Our integrated system is 10 – 100 times as fast as In-memory systems in business analytics.

© Hitachi, Ltd. 2016. All rights reserved.

Visit us at booth #811

© Hitachi, Ltd. 2016. All rights reserved.

HITACHI

Inspire the Next