

# Big Data Acceleration

Francis Lam  
Huawei Technologies



# Huawei HPC Momentum



**170+**  
Countries

**\$75B**  
2016 Revenue


**#3** Q1 '17  
WW Server Shipment

**20**  
Top 500


**79,000**  
R&D Engineers

**36**  
Joint Innovation  
Centers

## World Class HPC Systems



## Chip Level Innovations




Management chip



NC interconnect



NIC chip



SSD controller chip



» Peta-scale Supercomputer



» Large Memory Analytics & Deep Learning



» Continuous Momentum with Automotive Manufacturing



» Advancing HPC on Cloud

# Huawei FusionInsight Big Data Platform



## Reshape Data Processing, Let the Data Speak Intelligently



Real-time  
risk control



Real-time  
marketing



Real-time credit  
investigation



One-click  
search



Detail query



Cloud services

### FusionInsight Big Data platform

Unified SQL

Real-time search

Real-time decision

Relationship engine

Multi-tenancy

Data integration



Data analysis



Data services



Offline/Near-line  
computing



In-memory  
computing



Real-time stream  
computing

MPPDB

Parallel database



Physical / Virtual  
Servers & Storage



Hybrid Cloud



Huawei or Partner  
Public Cloud



### Agile

- Open architecture delivering linear performance improvements
- Efficient O&M supported by various tools
- Powerful Structured Query Language (SQL) capabilities simplify service migration

### Smart

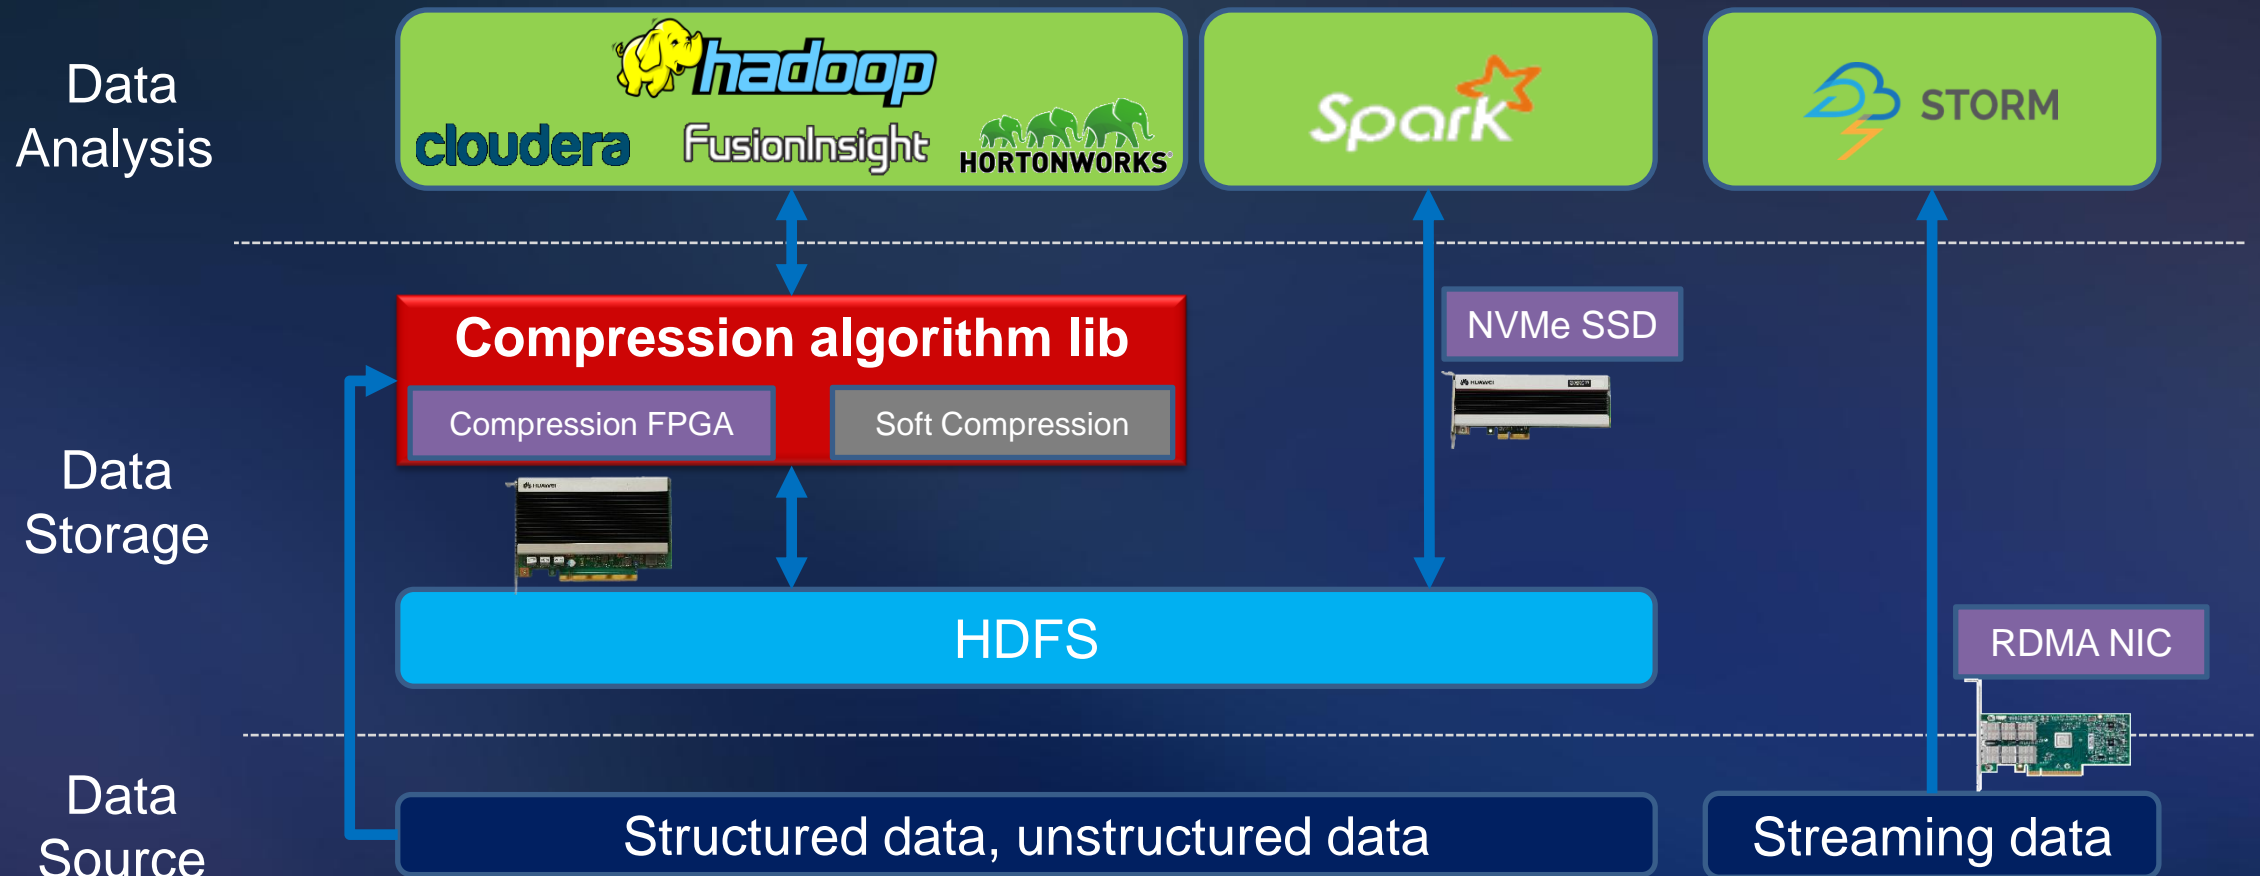
- Deep insight facilitated by full modeling
- Efficient and precise algorithms developed by Huawei

### Trustworthy

- High availability of all components, remote DR, and financial guarantees
- Trustworthy partners for win-win outcomes



# Big Data Acceleration






# HDFS Compression: Saved 30% of Data Nodes




Snappy soft compression

Vs



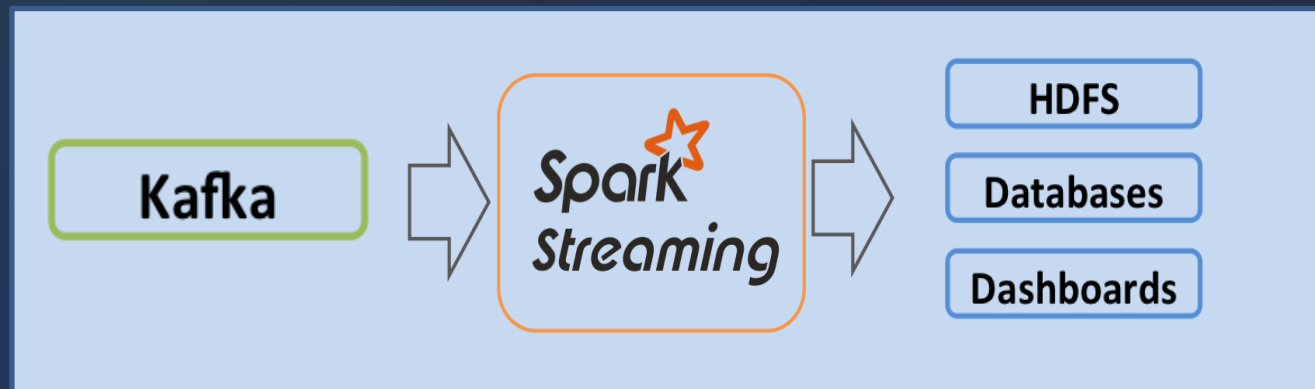

Data compression card

- 2.5X compression ratio
- Hardware Gzip compression 30% faster
- Release CPU resource – reduce CPU requirement
- Native software support
- 100% Transparent to application software



Data source: Huawei Server Solutions Lab

# Spark Acceleration: Real-time Analysis 40% Faster

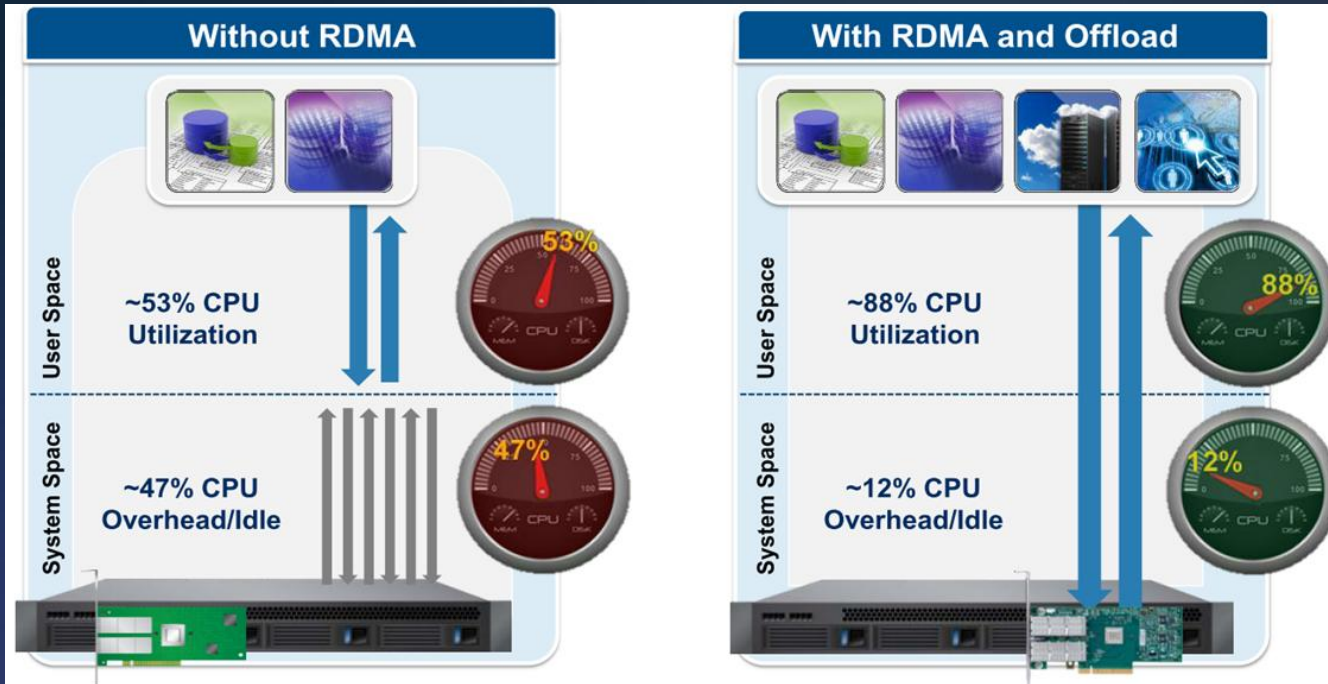


## Memory + SSD hybrid intelligent Cache

- During shuffle to disks, Shuffle Read requires intensive random read of small data block
- NVMe SSD delivers high BW, IOPS & latency

- Analysis time in case study **reduced by 40%**
- Performance increased by **75%**
- Boost performance / \$ by replacing HDD with SSD
- 100% transparent to application software
- Huawei SSD supports high reliability, built-in atomic write

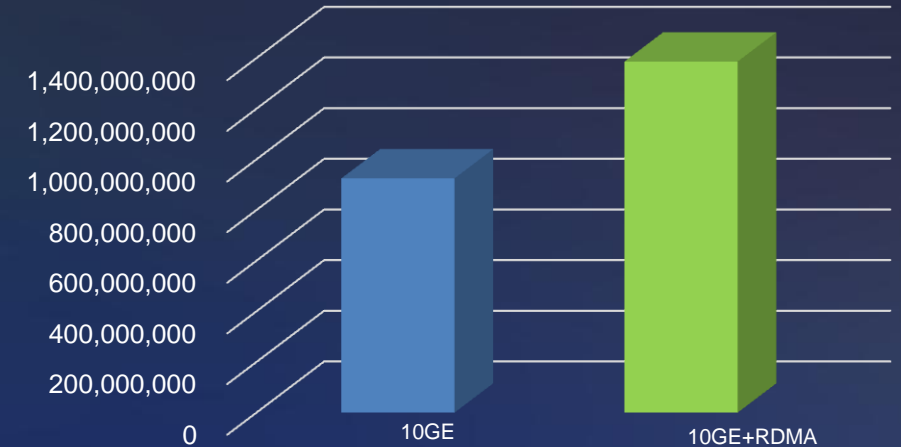
# Storm Acceleration: RDMA speeds stream processing




## RoCE Acceleration

- RoCE v2 based smart NIC
- Reduce CPU utilization, increase NIC throughput by 50%, reduce latency by 20%

## Storm wordcount benchmark



## latency



# Accelerating HPC



RH1288/2288



RH5288



RH5885



RH8100



KunLun  
32S System



## Fat Computing Nodes

FPGA Accelerator



ES3000 NVMe SSD



RDMA NIC



GPGPU



MIC



## Hardware Accelerators

GPU server



E9000



X6800



X6000



## High Performance Modular Systems





# Thank You