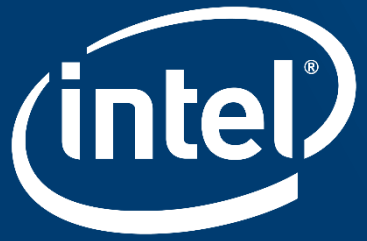

QATCodec: Past, Present and Future

Data Analytics Technology, Intel Corporation

LEGAL NOTICES

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

The products and services described may contain defects or errors known as errata which may cause deviations from published specifications. Current characterized errata are available on request.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting www.intel.com/design/literature.htm.

Intel, the Intel logo, Intel® are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others

FOR PERFORMANCE CLAIMS AND OPTIMIZATION NOTICE

Software and workloads used in performance tests may have been optimized for performance only on Intel® microprocessors. Performance tests, such as SYSmark* and MobileMark*, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. Notice Revision #20110804.

Performance estimates were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown." Implementation of these updates may make these results inapplicable to your device or system.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.
*Other names and brands may be claimed as the property of others
For more information go to http://www.intel.com/performance.

ABOUT US

Intel SSG OTC DAT

- Active open source contributions: About **30** committers from our team with significant contributions to a lot of big data key projects including **Spark**, **Hadoop**, **HBase**, **Hive**…
- A global engineering team (China, US, India) - Major team is based in Shanghai
- Mission: To fully optimize analytics solutions on IA and make Intel the **best** and **easiest** platform for big data and analytics customers

Edge

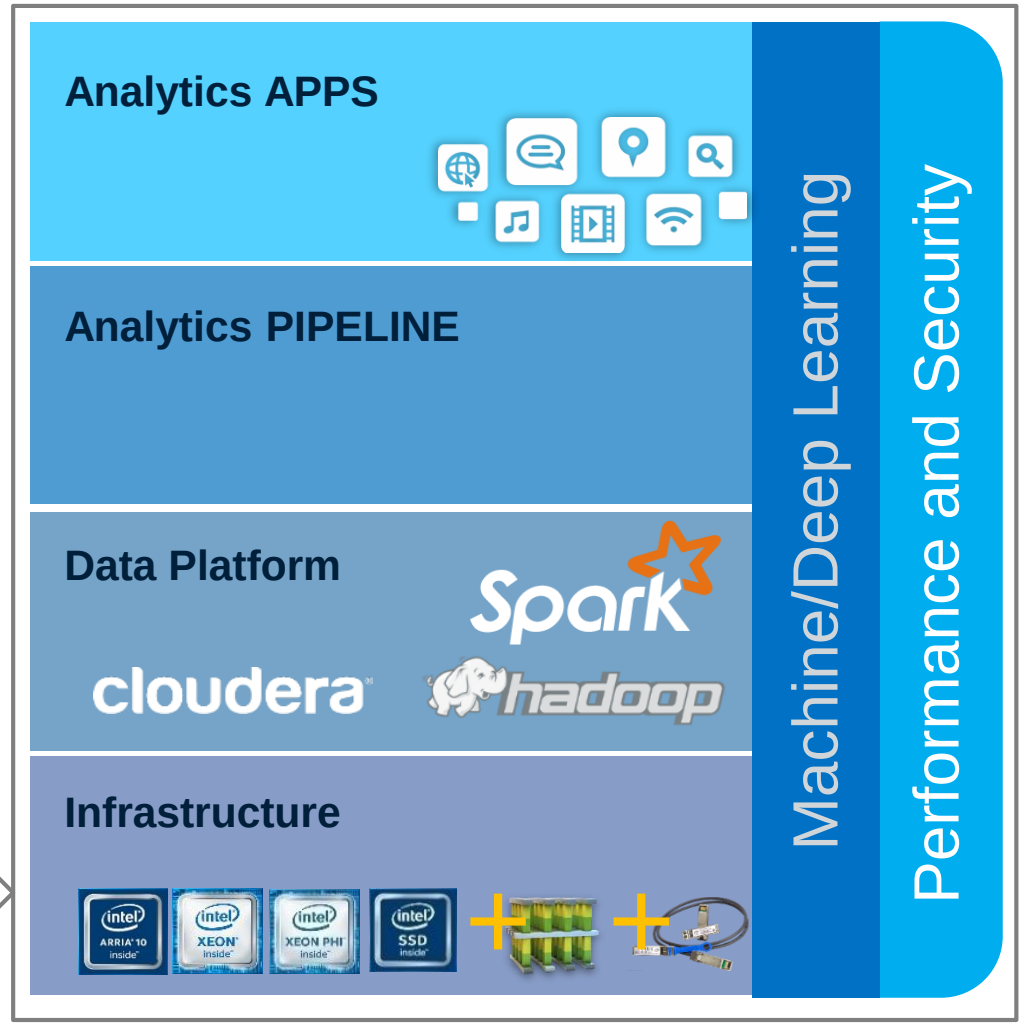

Data Center

AGENDA

- Introduction to Big Data
- QAT Acceleration Opportunity in Big Data
- High Level Architecture for QATCodec in Big Data
- Performance Evaluation
- Future Work for QATCodec

WHAT IS BIG DATA

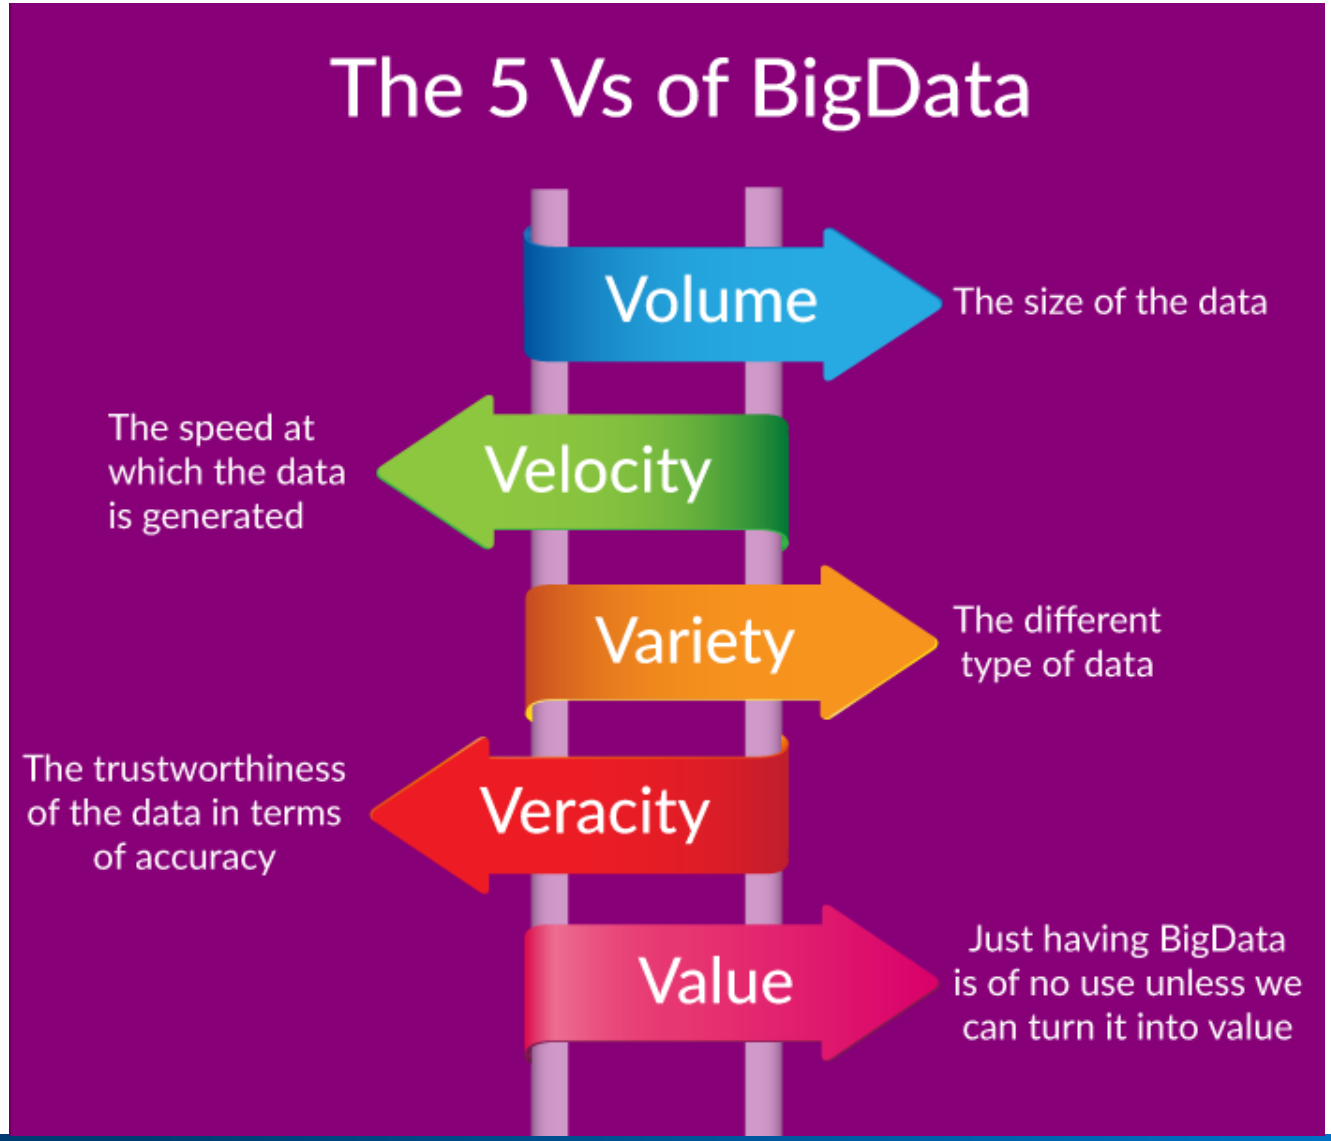

IMG Source: https://viblo.asia/p/why-we-need-modern-big-data-integration-platform-gAm5yx1Dldb

BIG DATA SOLUTION STACK AT INTEL

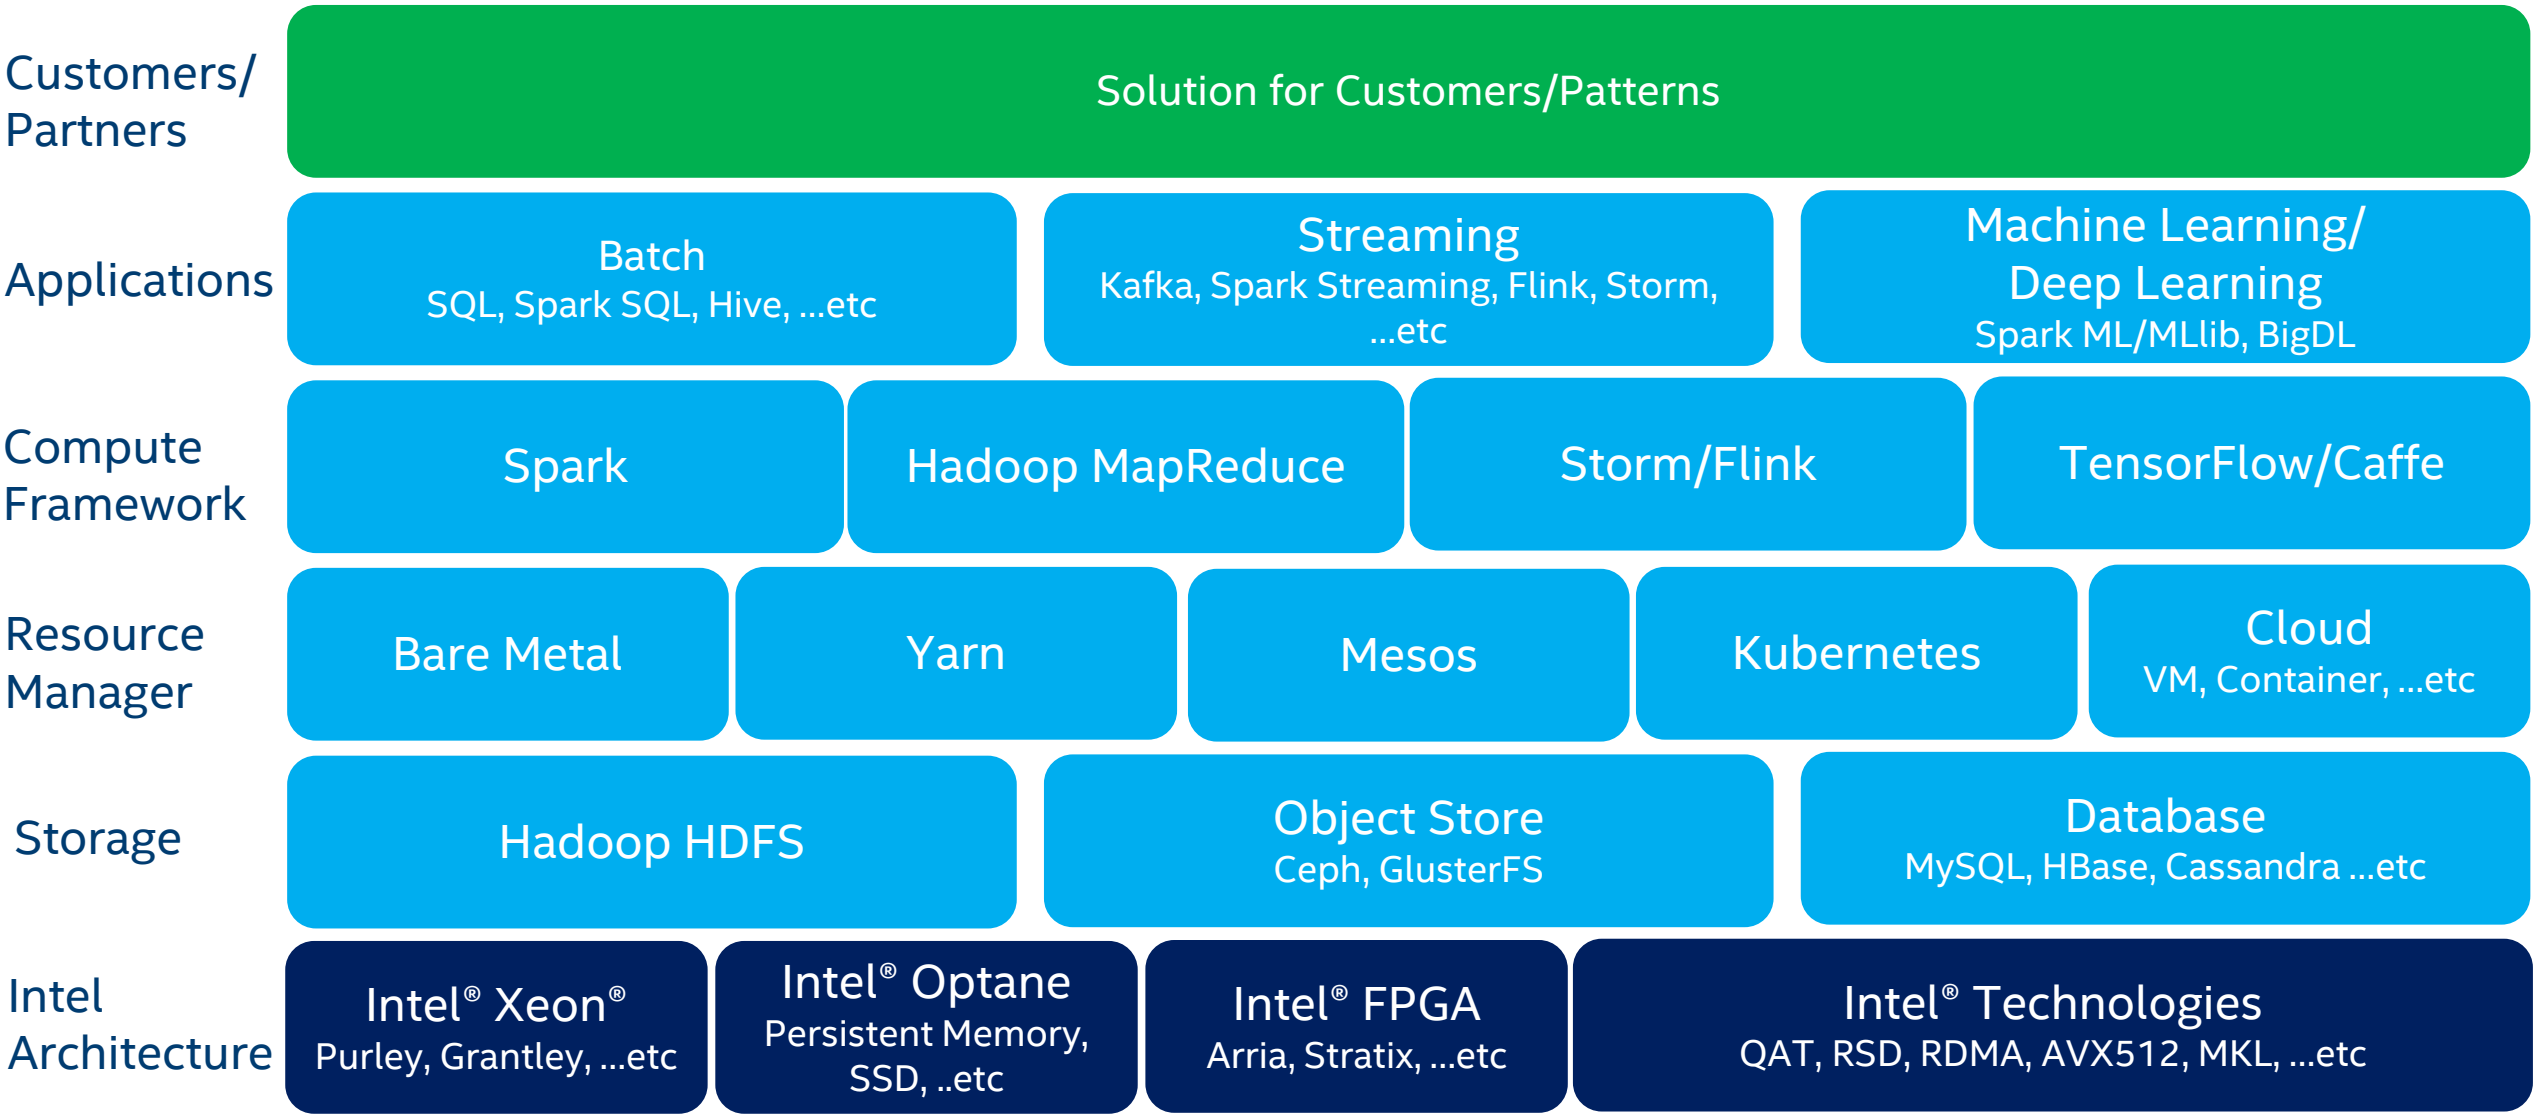

WHY QAT IS IMPORTANT TO BIG DATA

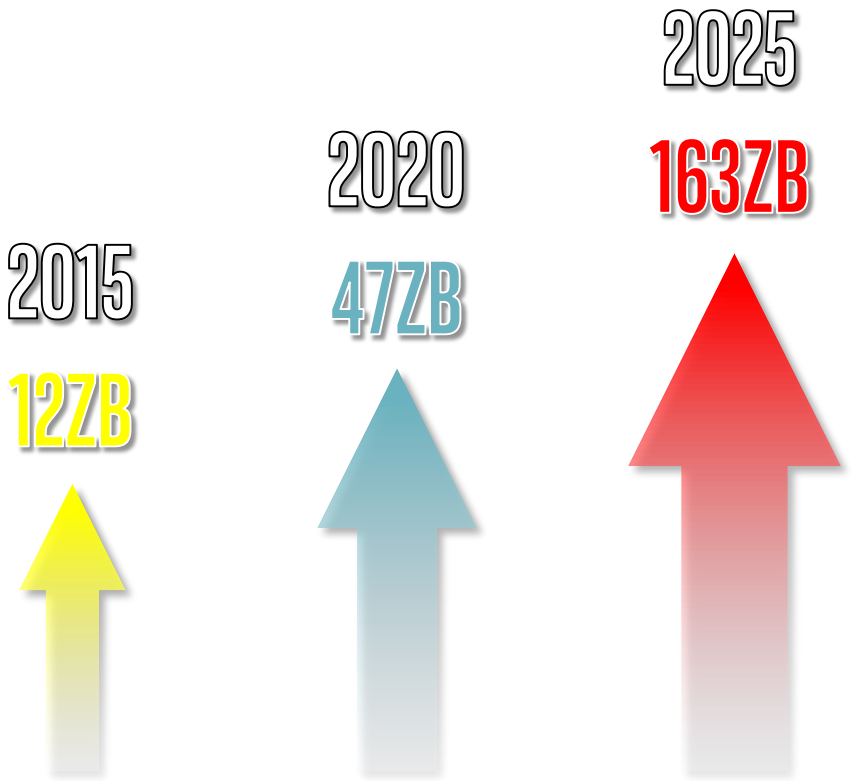

- **Compression Needs**
  - Reduce data volume and save storage space.
  - **Speed up the disk I/O operations** and **data transfer** across network, optimize workload performance.
- **Trade-off**
  - Computation overhead for high compression ratio codecs.

source[IDC and Seagate]: https://www.seagate.com/files/www-content/our-story/trends/files/Seagate-WP-DataAge2025-March-2017.pdf

DATA COMPRESSION PIPELINE IN BIG DATA FRAMEWOEK

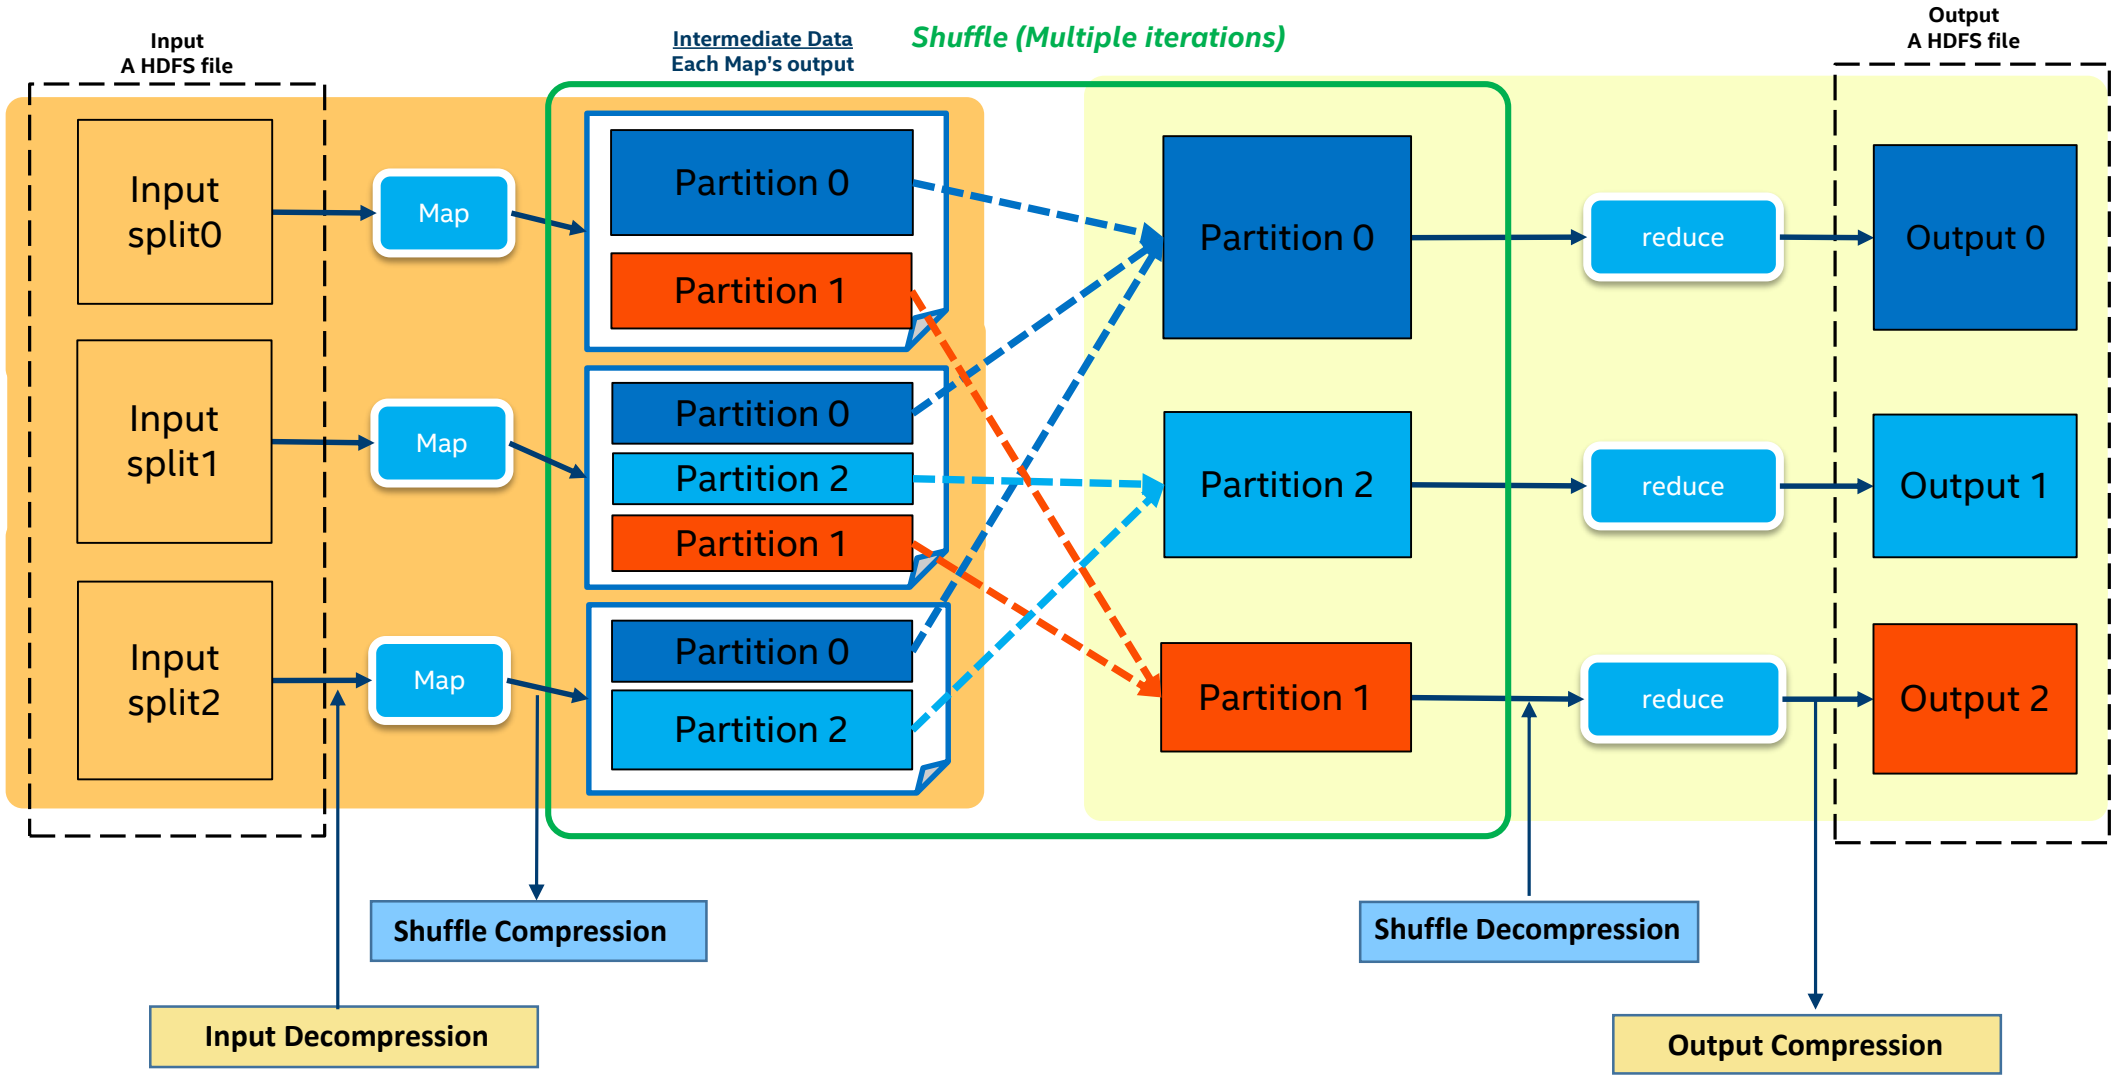

DATA COMPRESSION PIPELINE IN BIG DATA - I/O CHARACTERISTICS

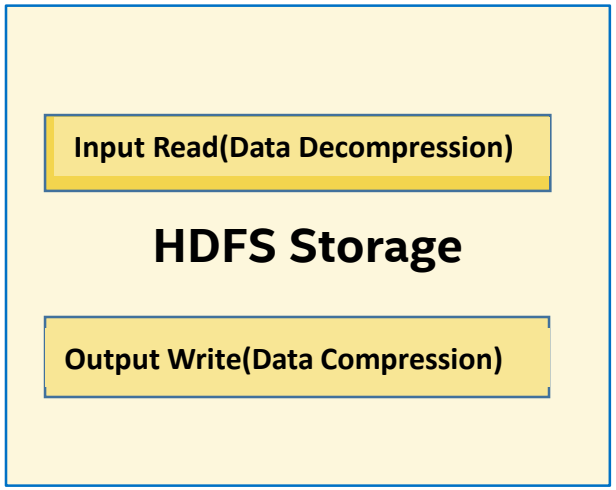

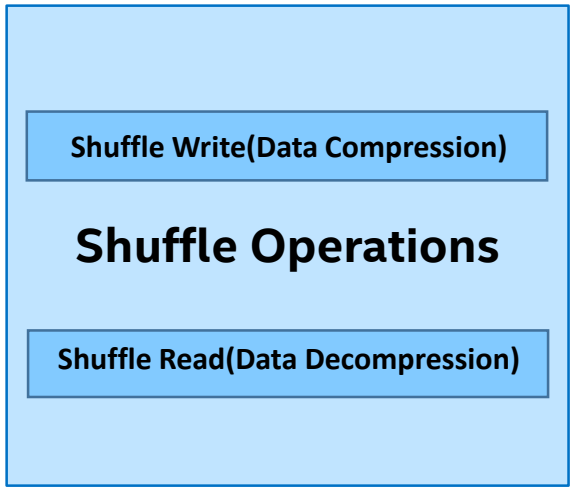

- HDFS Storage
  - Generally sequence read/write
  - Generally one time read/write

- Shuffle Operations
  - Random read/write
  - Multiple times read/write

QAT ACCELERATION OPPORTUNITY

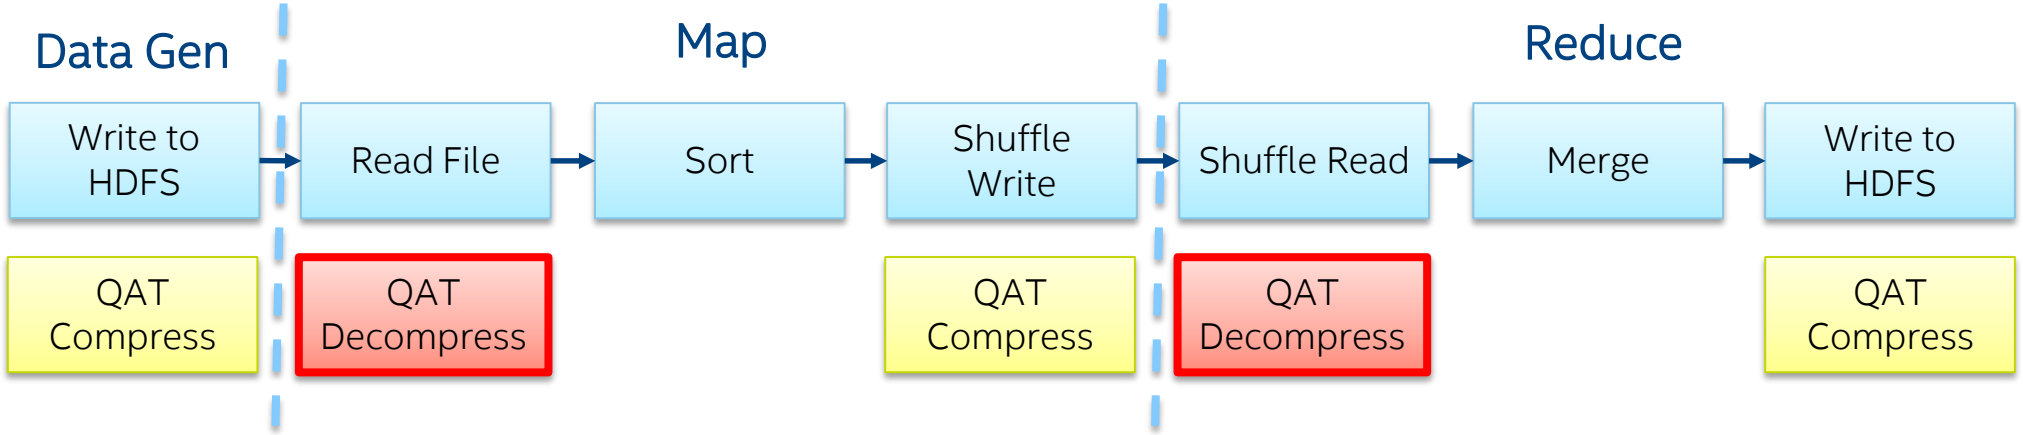

Compression/Decompression in Sort Workload

ABOUT QATCODEC

QAT Codec project provides compression and decompression library for Apache Hadoop/Apache Hive/Apache Spark using the Intel® QuickAssist Technology.

Its advantages:

- Higher performance with better compression ratio (against Snappy)

- Open
  - Great support in major big data components (Spark, Hadoop, Hive, Parquet)
  - It's an open source project and available at https://github.com/intel-hadoop/IntelQATCodec
  - Apply Intel QAT Open Lab https://www.sdnlab.com/intel_network_platform

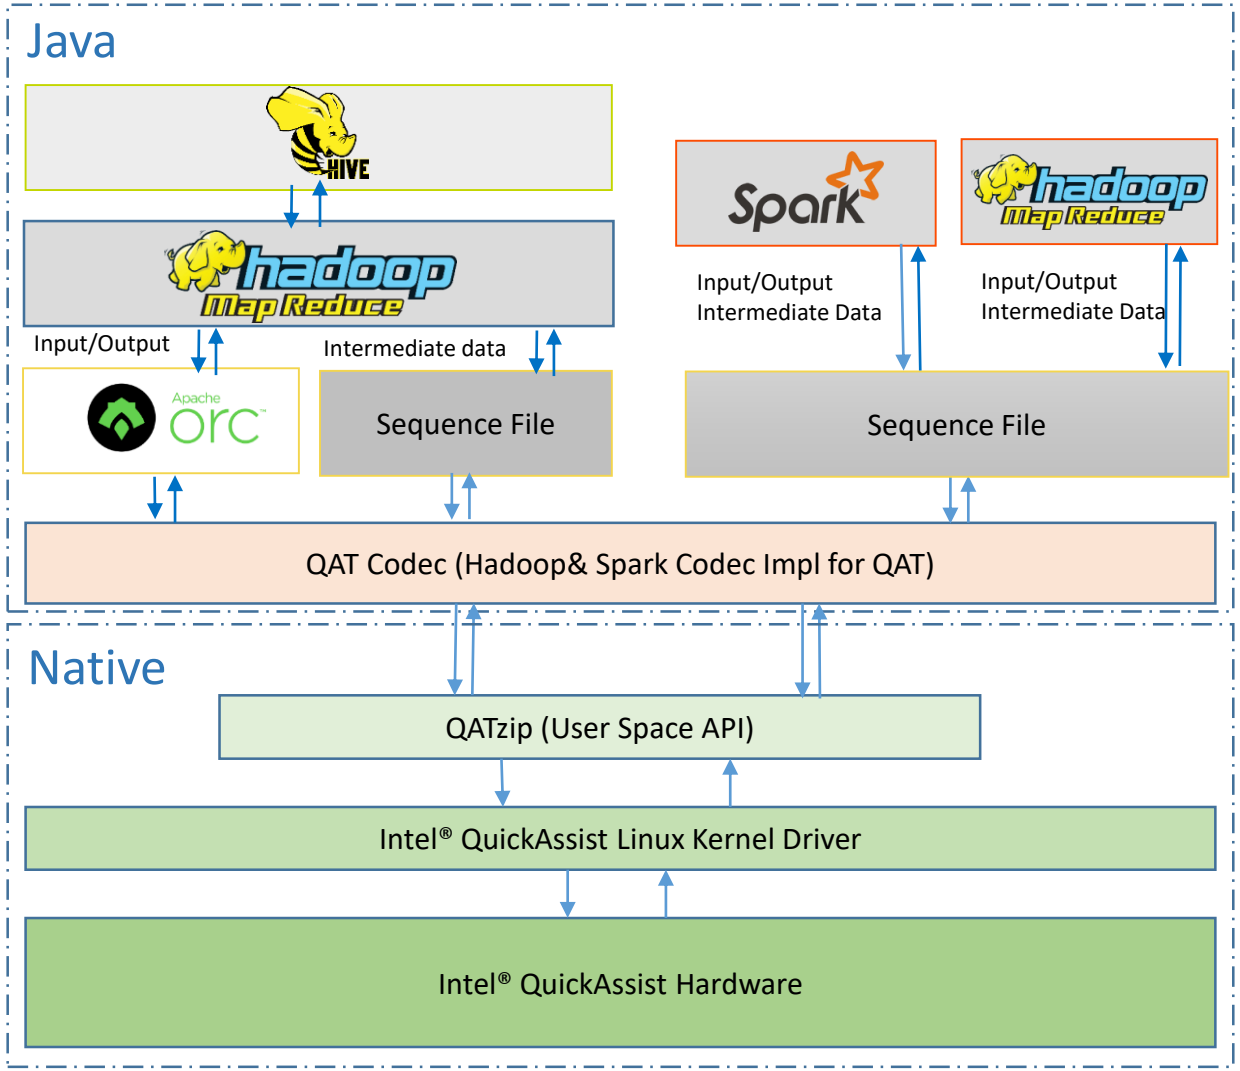

BENCHMARK CONFIGURATION

Hardware	
Number of worker nodes	3
CPU	Intel(R) Xeon(R) Platinum 8170 CPU @ 2.10GHz
Number of vcores per node	Limit to 64
Memory	384 GB
Disk	1* NVMe SSD 2 TB
QAT	3* Lewisburg
Network	10GbE
Software	
QAT	Version 1.1.0 Build 0010 (Early Access with CnV)
CDH	5.14.0

PERFORMANCE EVALUATION – SPARK SORT

- QAT VS. Snappy
  - Compression Ratio
    - Input size reduced: 7.5% =(431.8-399.3)/431.8
    - Output size reduced: 7.25% =(372.9-345.4)/372.9
  - Performance Gain: 14.3% =(697-597)/697

	Data Size (GB)
Input-QAT	399.3
Input-Snappy	431.8
Output-QAT	345.4
Output-Snappy	372.9

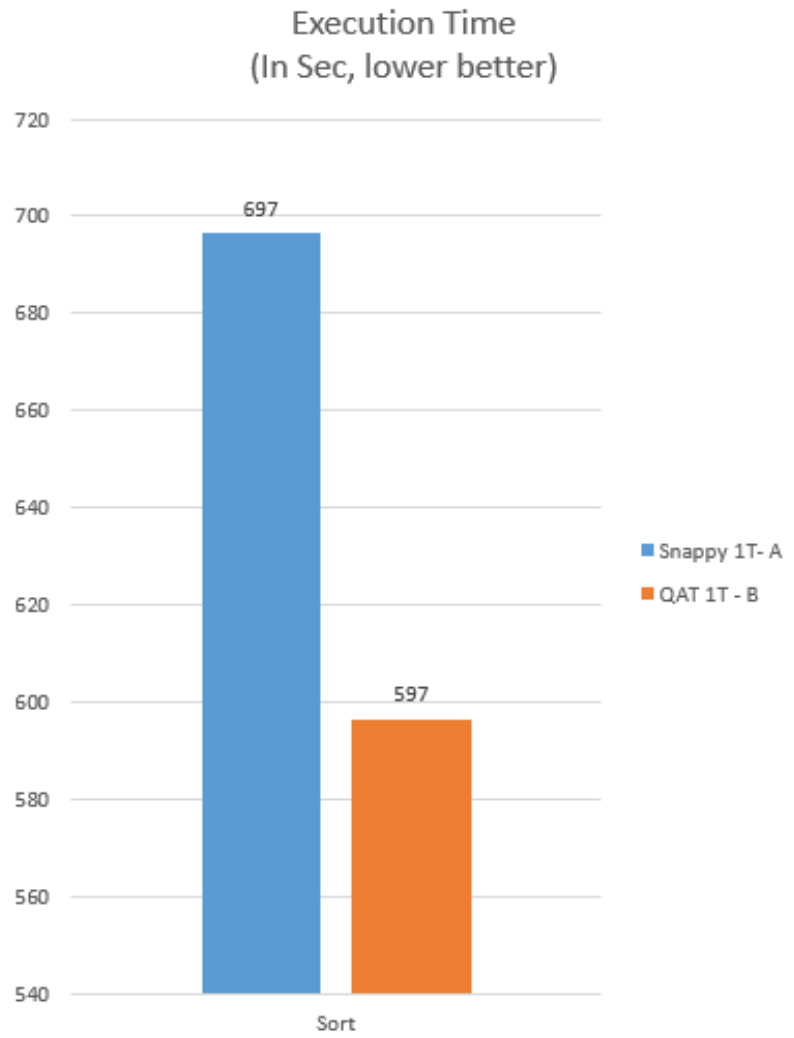

Performance estimates were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown." Implementation of these updates may make these results inapplicable to your device or system. Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to www.intel.com/benchmarks. Tests performed by Intel® company. Configurations: see slides 13

PERFORMANCE EVALUATION – MAP REDUCE

- QAT VS. Snappy
  - Compression ratio:
    - Input size reduced: 7.5% =(429.6-397.5)/429.6
    - Output size reduced: 8.04% =(424.2-390.1)/424.2
  - Performance gain:
    - Random TXT writer: 7.29% =(311-288.33)/311
    - Sort: 4.40% =(644.33-616)/644.33

	Data Size (GB)
Input-QAT	397.5
Input-Snappy	429.6
Output-QAT	390.1
Output-Snappy	424.2

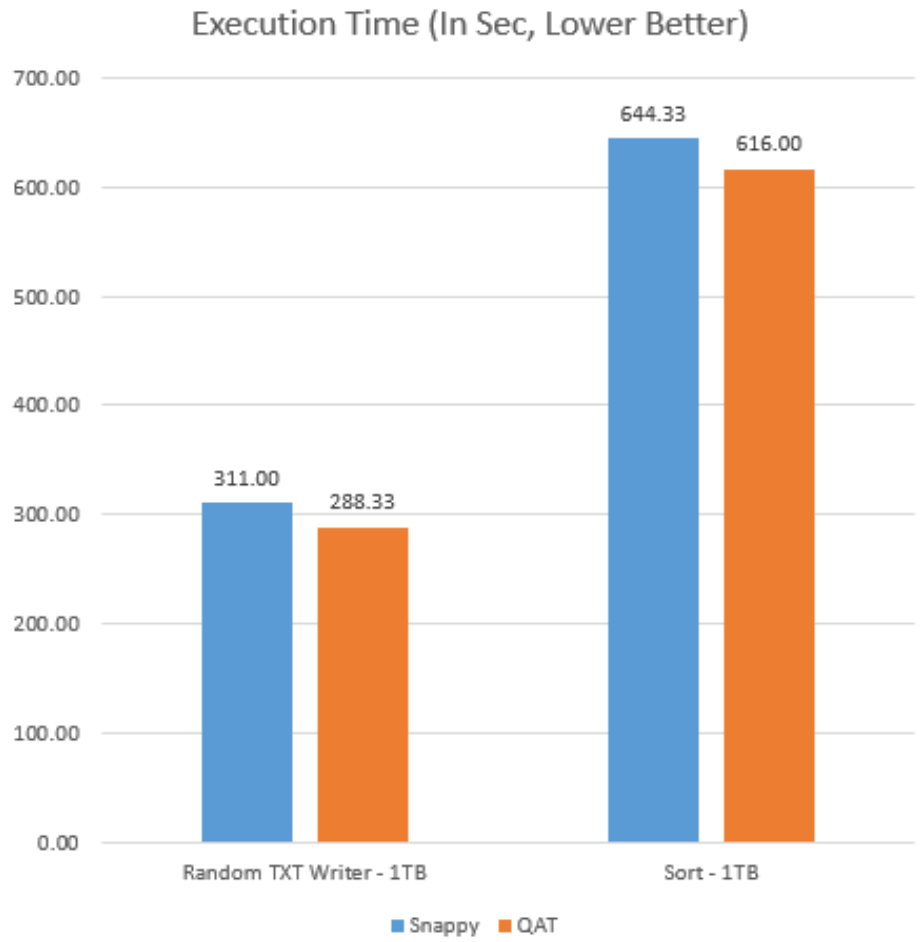

Performance estimates were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown." Implementation of these updates may make these results inapplicable to your device or system. Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to www.intel.com/benchmarks. Tests performed by Intel® company. Configurations: see slides 13

PERFORMANCE EVALUATION – HIVE ON MAP REDUCE

- QAT VS. Snappy
  - Compression ratio:
    - Data size reduced: 13.65% =(232.8-201)/232.8
  - Performance gain:
    - Query Execution Time: 12.98% =(11693-9974)/311

	Data Size (GB)
Raw data	871.2
Hive data / Snappy	232.8
Hive data / QAT dynamic	201.0

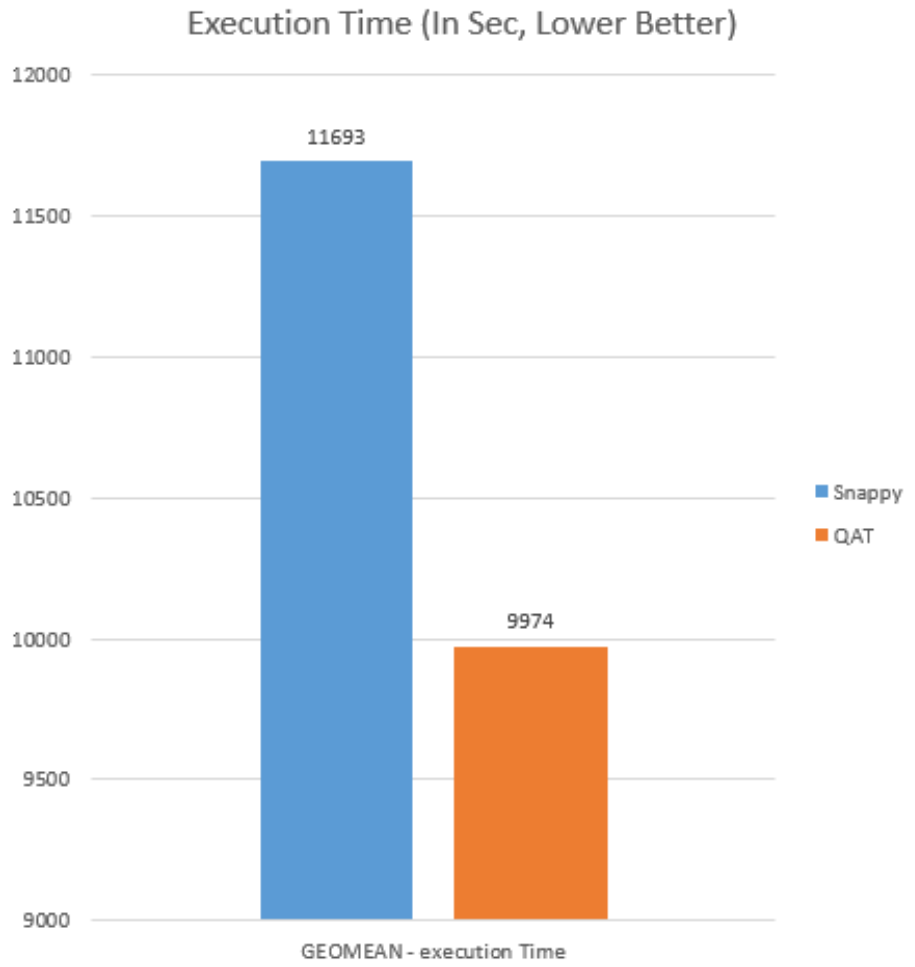

Performance estimates were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown." Implementation of these updates may make these results inapplicable to your device or system. Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to www.intel.com/benchmarks. Tests performed by Intel® company. Configurations: see slides 13

QAT VS. SNAPPY – COMPRESSION RATIO (1TB DATA SCALE)

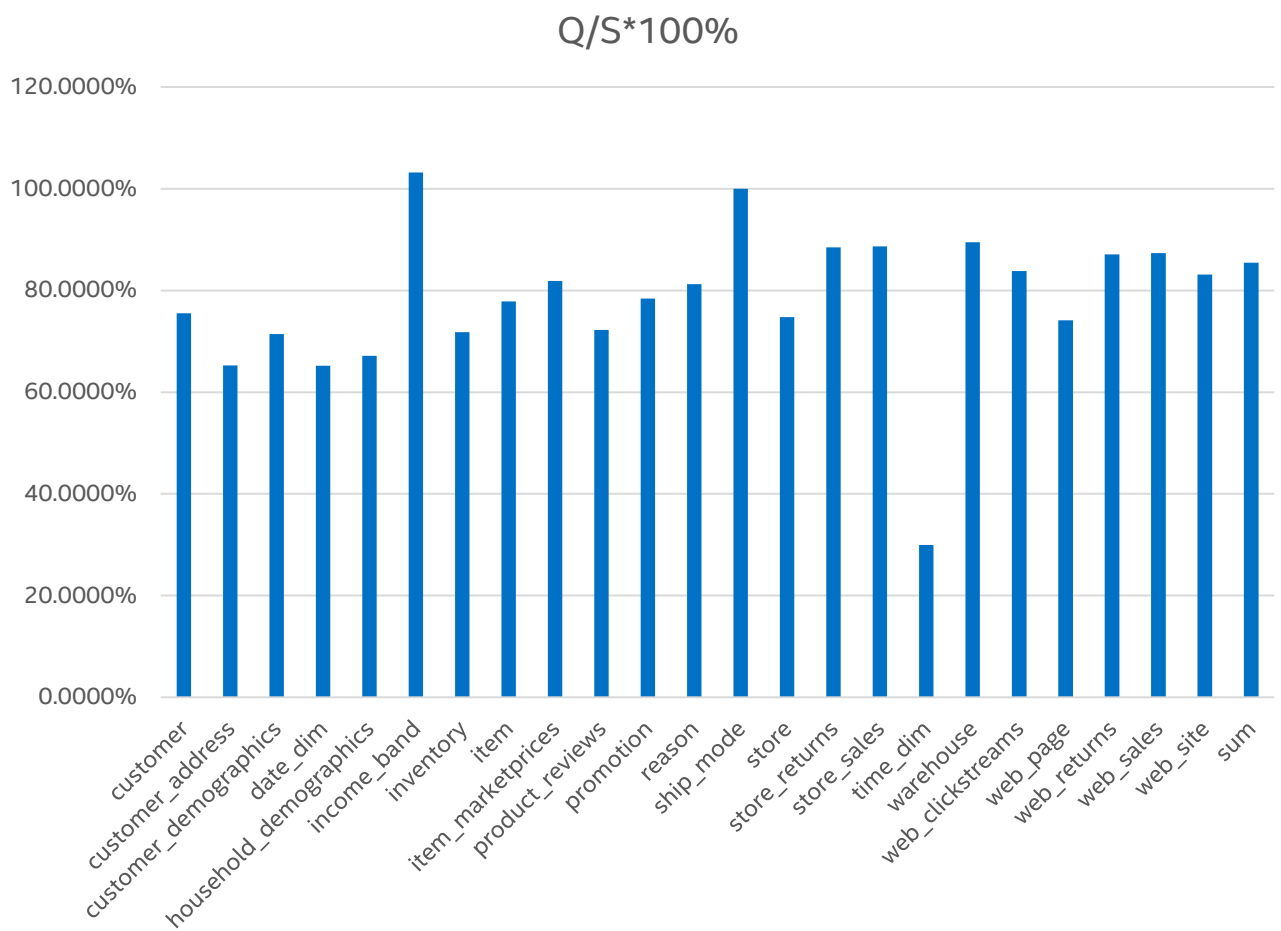

table	QAT	SNAPPY		Q/S*100%
customer	127.5	168.8	M	75.5332%
customer_address	25.7	39.4	M	65.2284%
customer_demographics	7	9.8	M	71.4286%
date_dim	1.5	2.3	M	65.2174%
household_demographics	14.5	21.6	K	67.1296%
income_band	420	407	B	103.1941%
inventory	5.6	7.8	G	71.7949%
item	77.2	99.2	M	77.8226%
item_marketprices	22.1	27	M	81.8519%
product_reviews	1.3	1.8	G	72.2222%
promotion	210.7	268.8	K	78.3854%
reason	21.2	26.1	K	81.2261%
ship_mode	2.1	2.1	K	100.0000%
store	35.9	48	K	74.7917%
store_returns	2.3	2.6	G	88.4615%
store_sales	43.9	49.5	G	88.6869%
time_dim	134.4	449	K	29.9332%
warehouse	3.4	3.8	K	89.4737%
web_clickstreams	80.8	96.4	G	83.8174%
web_page	59	79.6	K	74.1206%
web_returns	2.7	3.1	G	87.0968%
web_sales	62.2	71.2	G	87.3596%
web_site	7.4	8.9	K	83.1461%
sum	199	232.8	G	85.4811%

IMPACT ANALYSIS ON QUERIES - COMPRESSION RATIO

- Impacts
  - *Map Join, Parallelism & ORC records per strip*
- Map Join conversion
  - *hive.auto.convert.join.noconditionaltask.size*
  - *Joined table total size*
  - *GC overhead issue*
- Parallelism
  - *Input split size: mapreduce.input.fileinputformat.split.maxsize*
  - *File block size: smallest data unit for map task (e.g. orc.strip.size)*
  - *Determine Map task number*

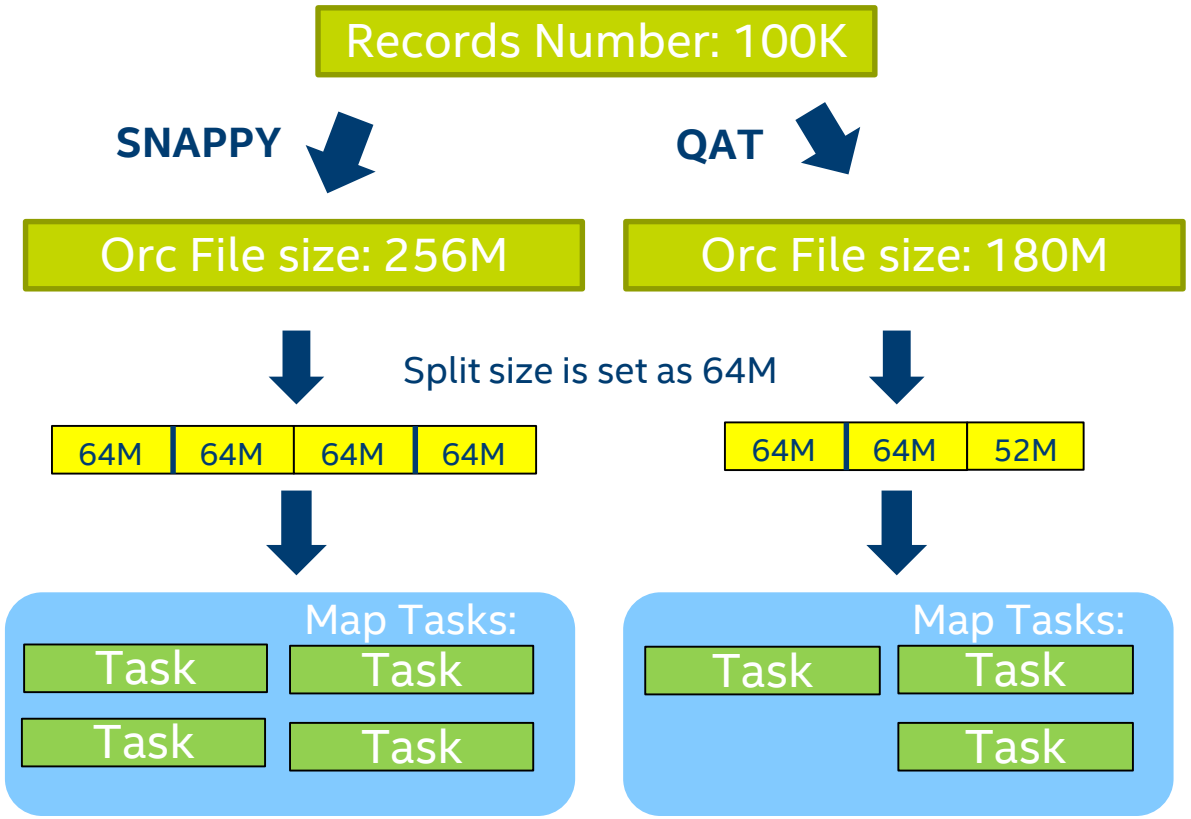

QUERY ANALYSIS

Map join conversion

Map join GC overhead issue

System Metrics

Q22 (MAP JOIN CONVERSION)

QAT

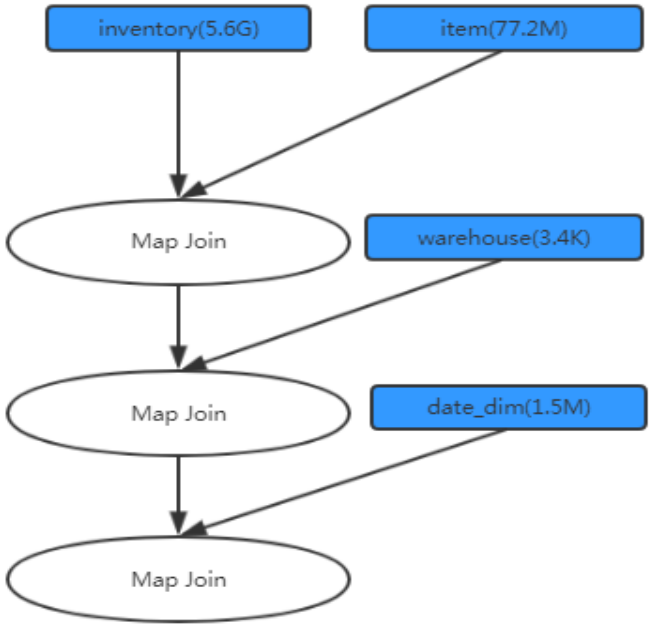

Snappy

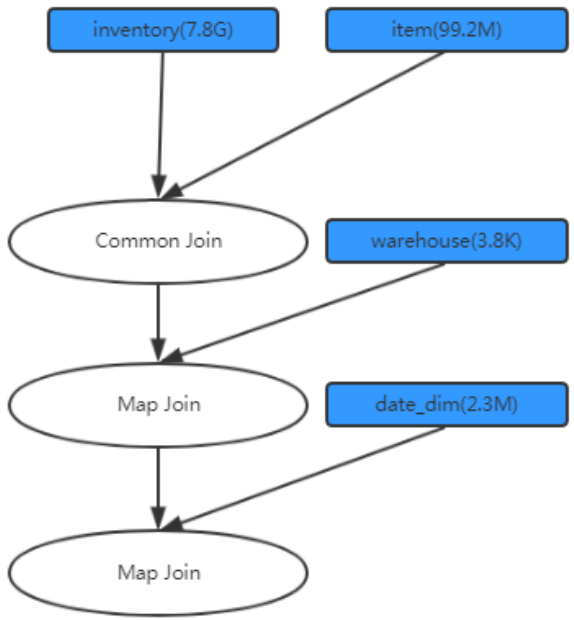

hive.auto.convert.join.noconditionaltask.size=100000000; (95MB)

Table size summary:

QAT	SNAPPY	table
1.5M	2.3M	date_dim
5.6G	7.8G	inventory
77.2M	99.2M	item
3.4K	3.8K	warehouse

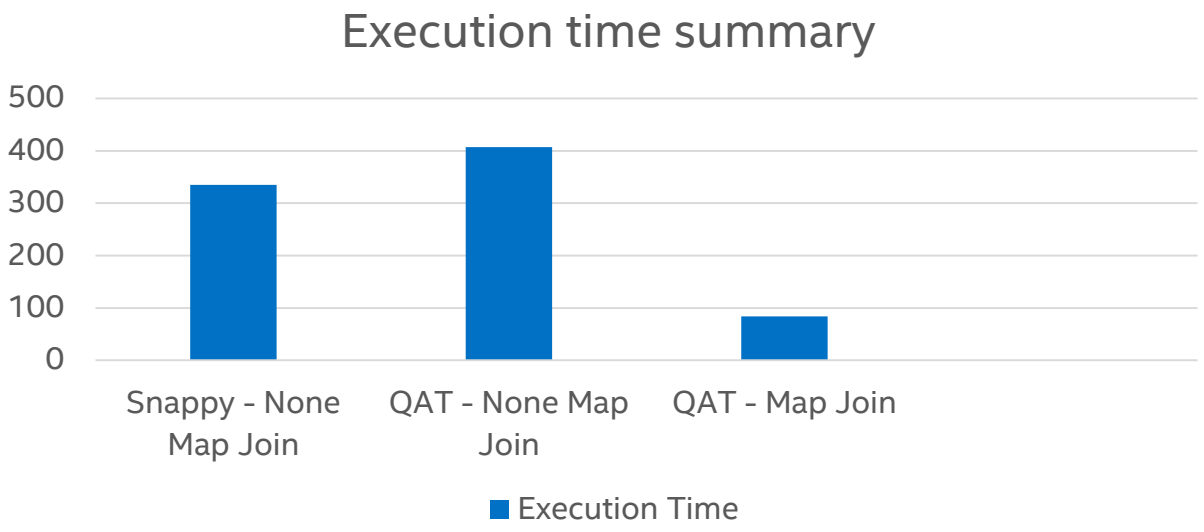

QUERY ANALYSIS

Map join conversion

Map join GC overhead issue

System Metrics

Q12 (GC ISSUE CAUSED BY MAP JOIN)

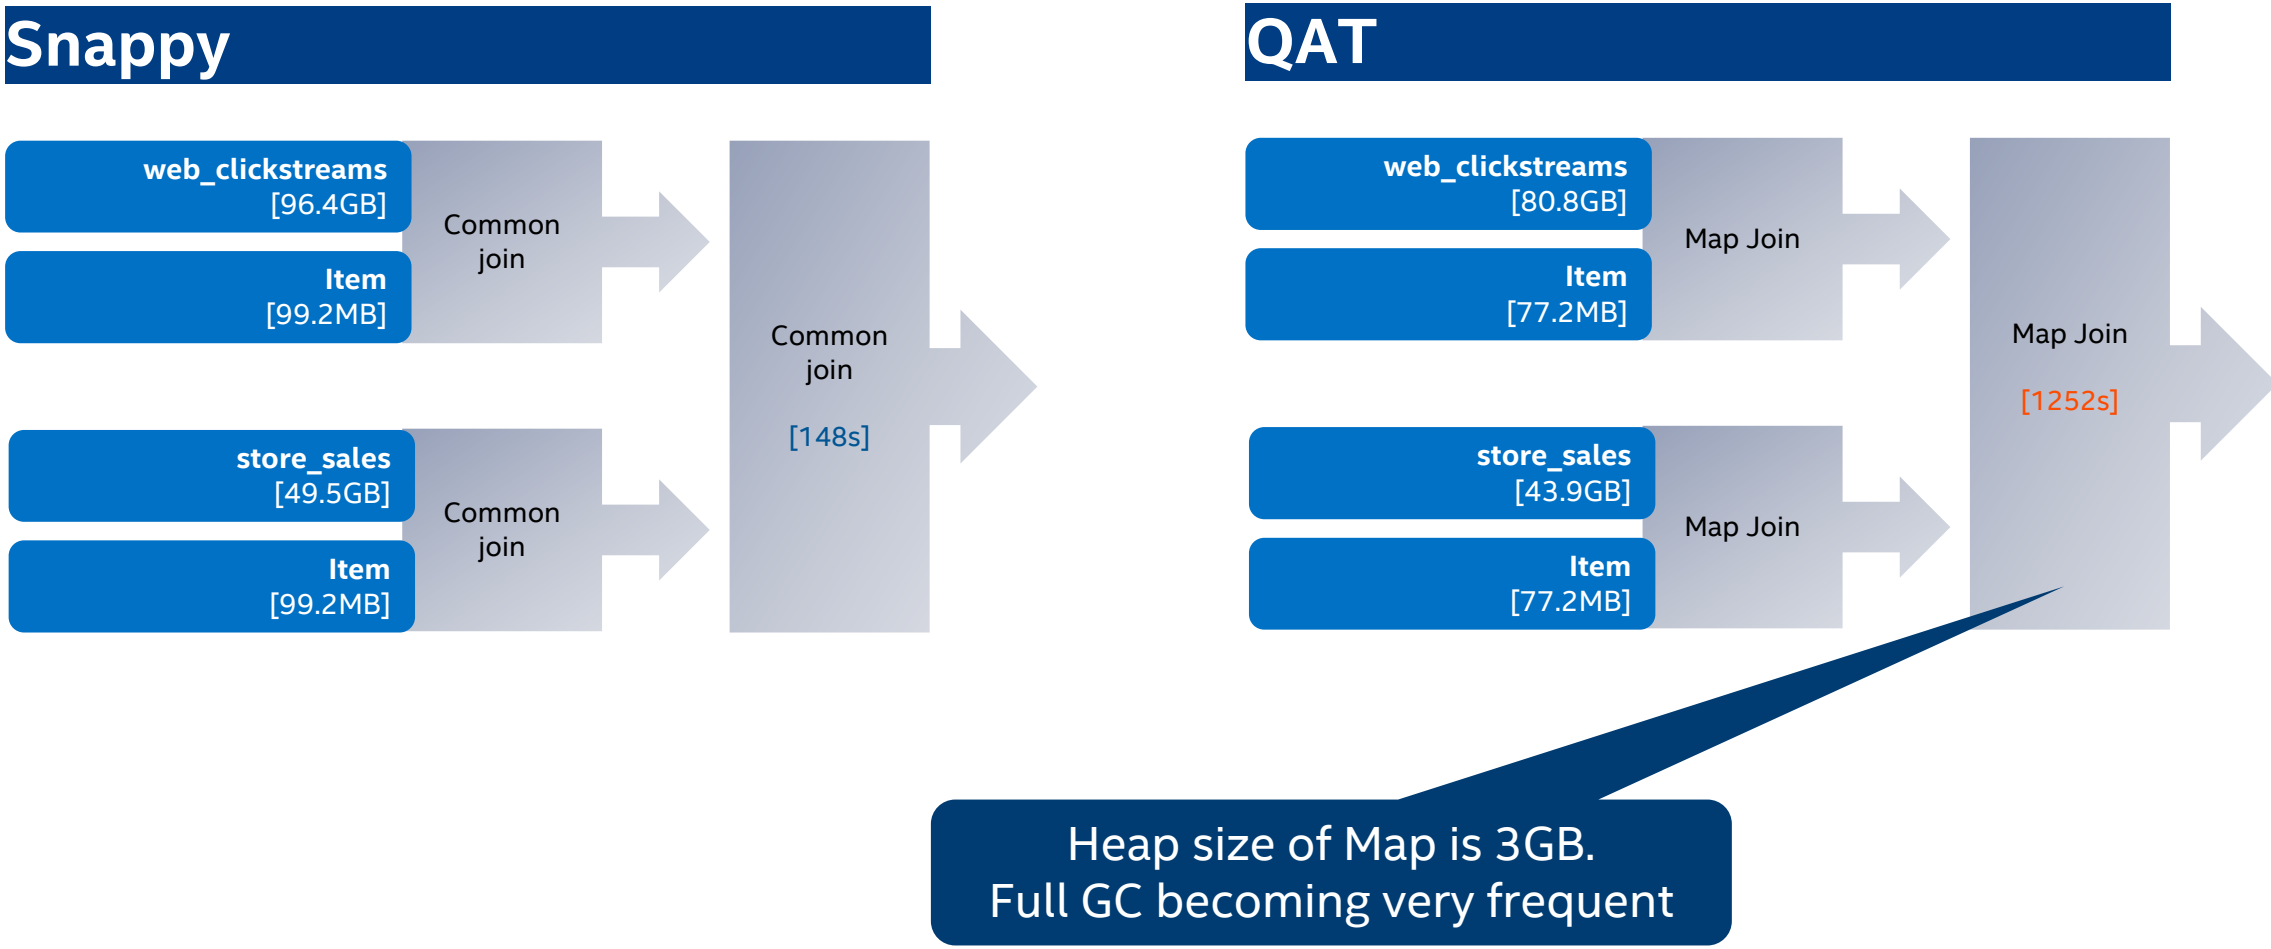

Q12 (GC ISSUE CAUSED BY MAP JOIN) – CONT'D

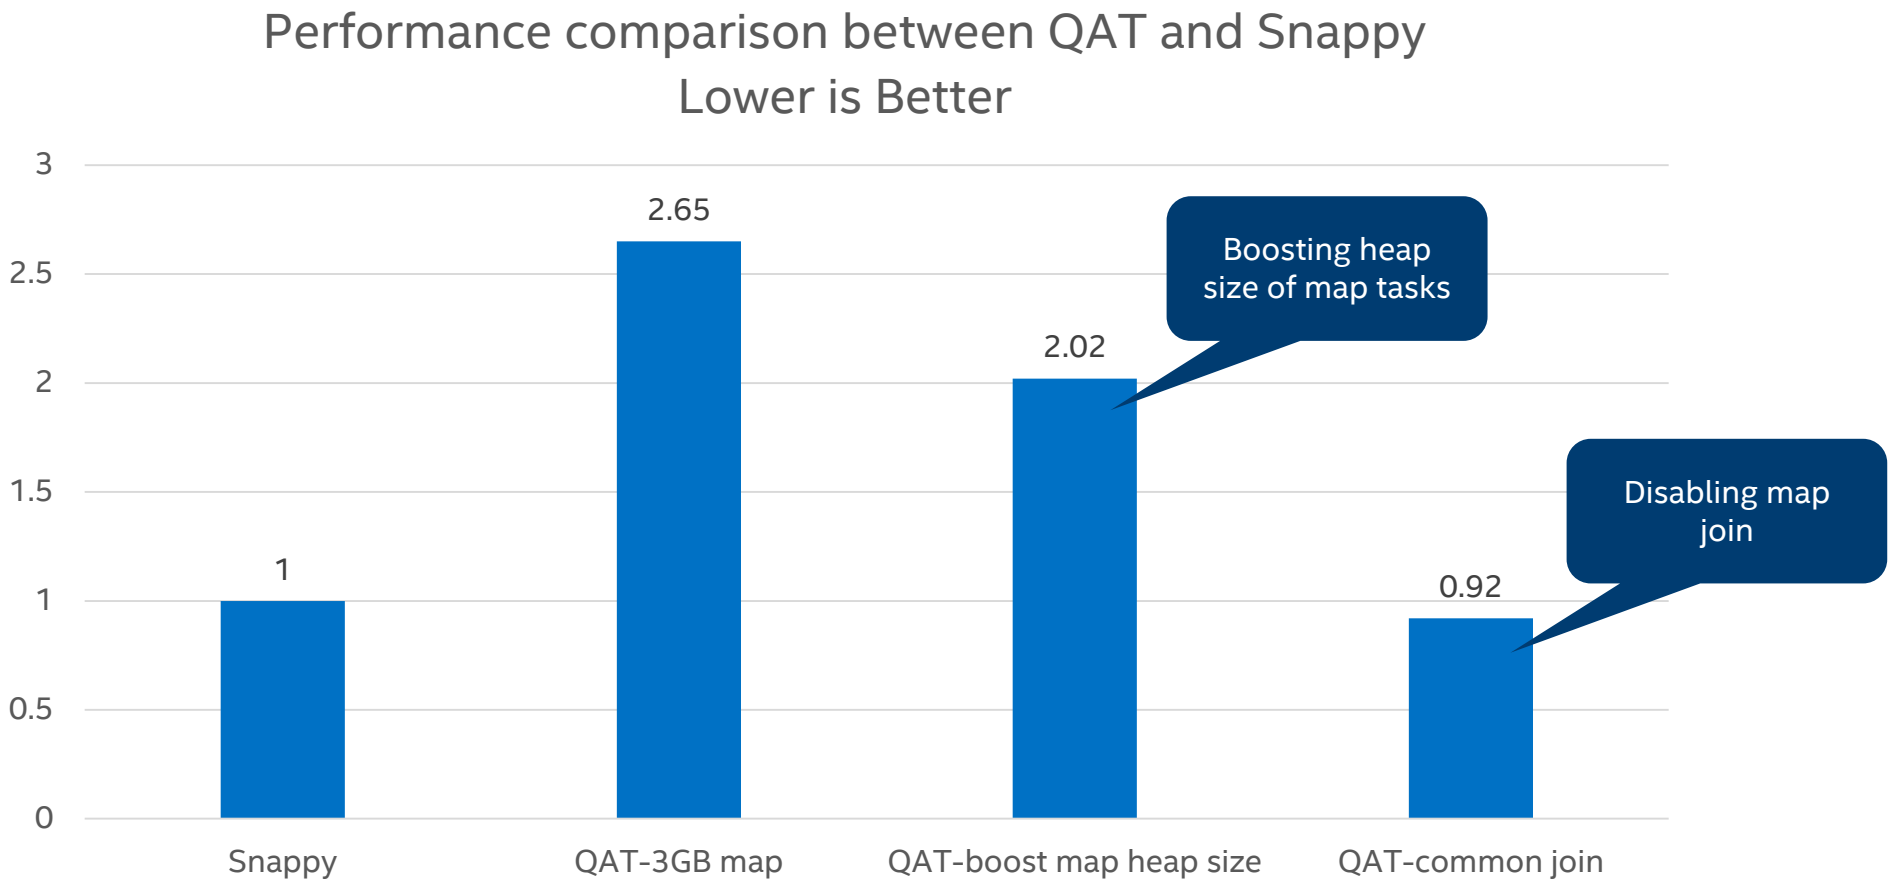

QUERY ANALYSIS

Map join conversion

Map join GC overhead issue

System Metrics

MICRO-VIEW QUERY COMPARISON - IO WAIT

CPU Utilization

IO Request

Snappy

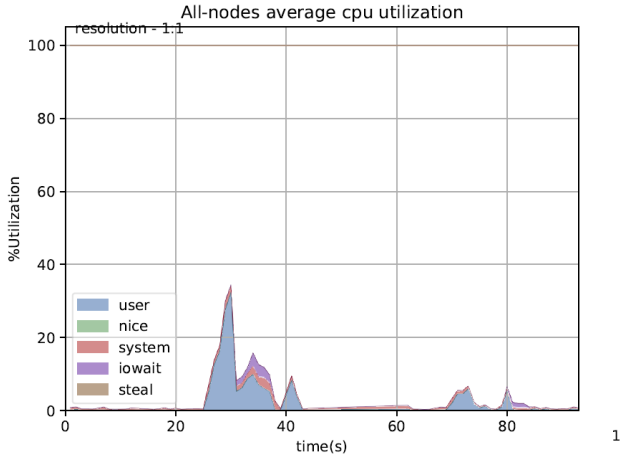

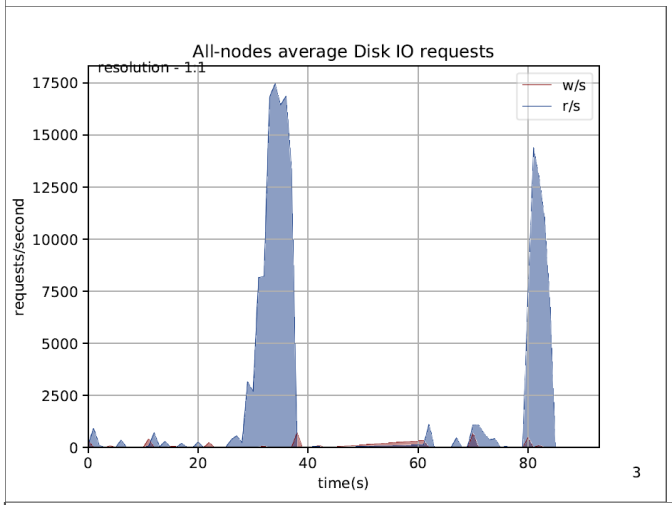

QAT

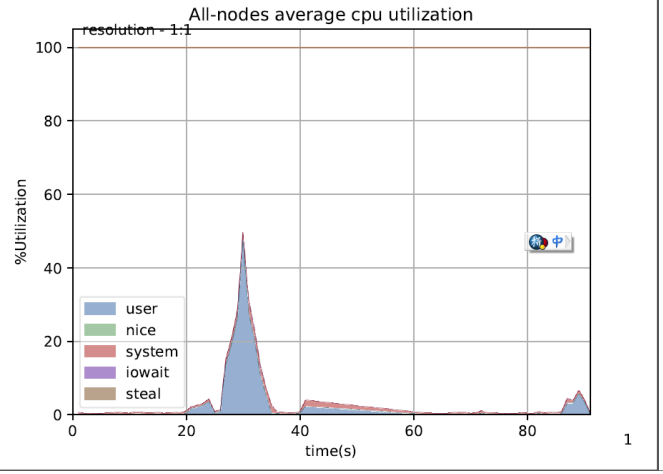

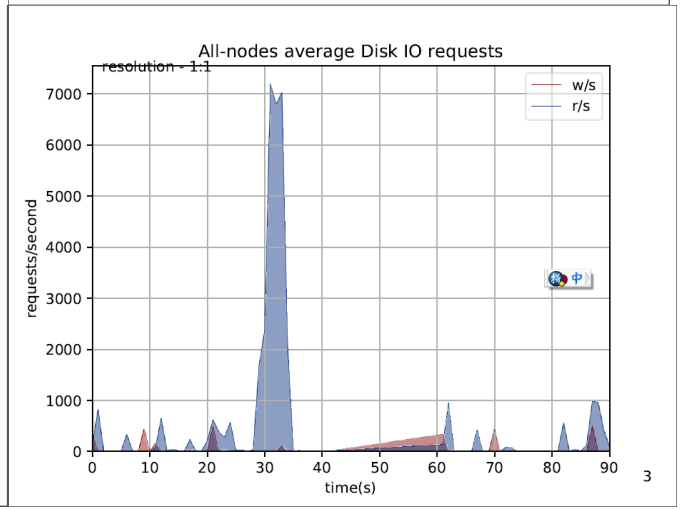

- Less IO request
- IO wait helps

MICRO-VIEW QUERY COMPARISON – NO IO WAIT

CPU Utilization

IO Request

Snappy

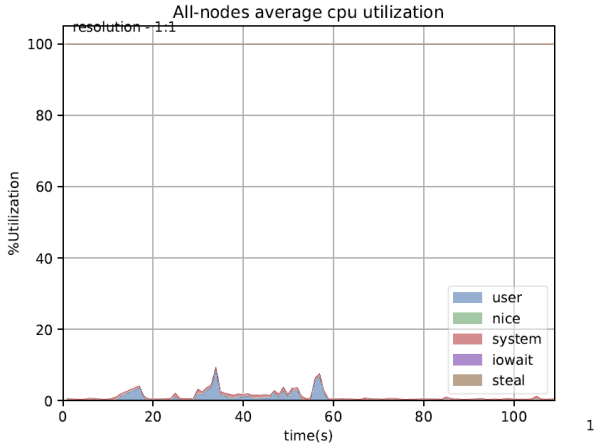

1

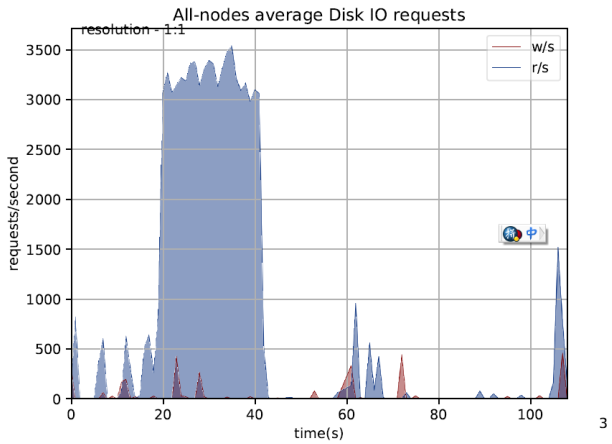

3

QAT

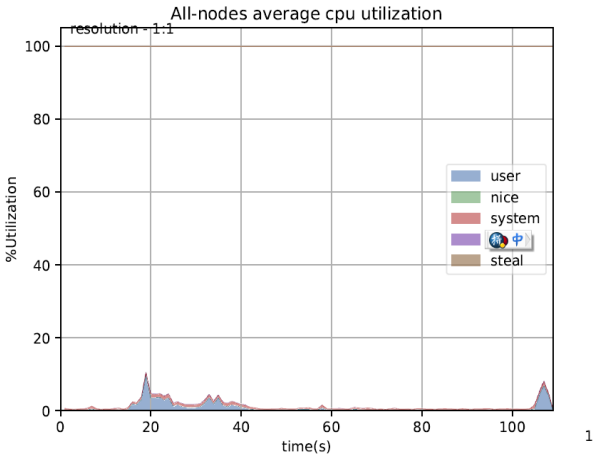

1

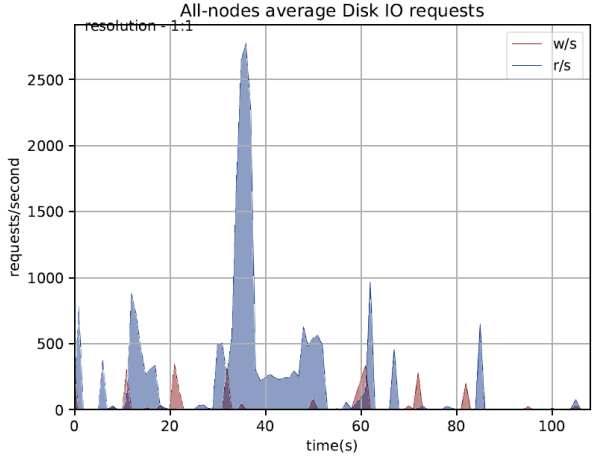

3

- Less IO request

ACHIEVEMENTS

Deal Win for QuickAssist Technology: We already have the first customer for Skylake with QAT and QATCodec, more customers have interesting in it.

Cloudera Certification of Intel® QuickAssist Technology Codec & QATZIP: Validation using MapReduce jobs with TeraGen and Terasort workloads

https://www.cloudera.com/partners/partners-listing.html?q=intel

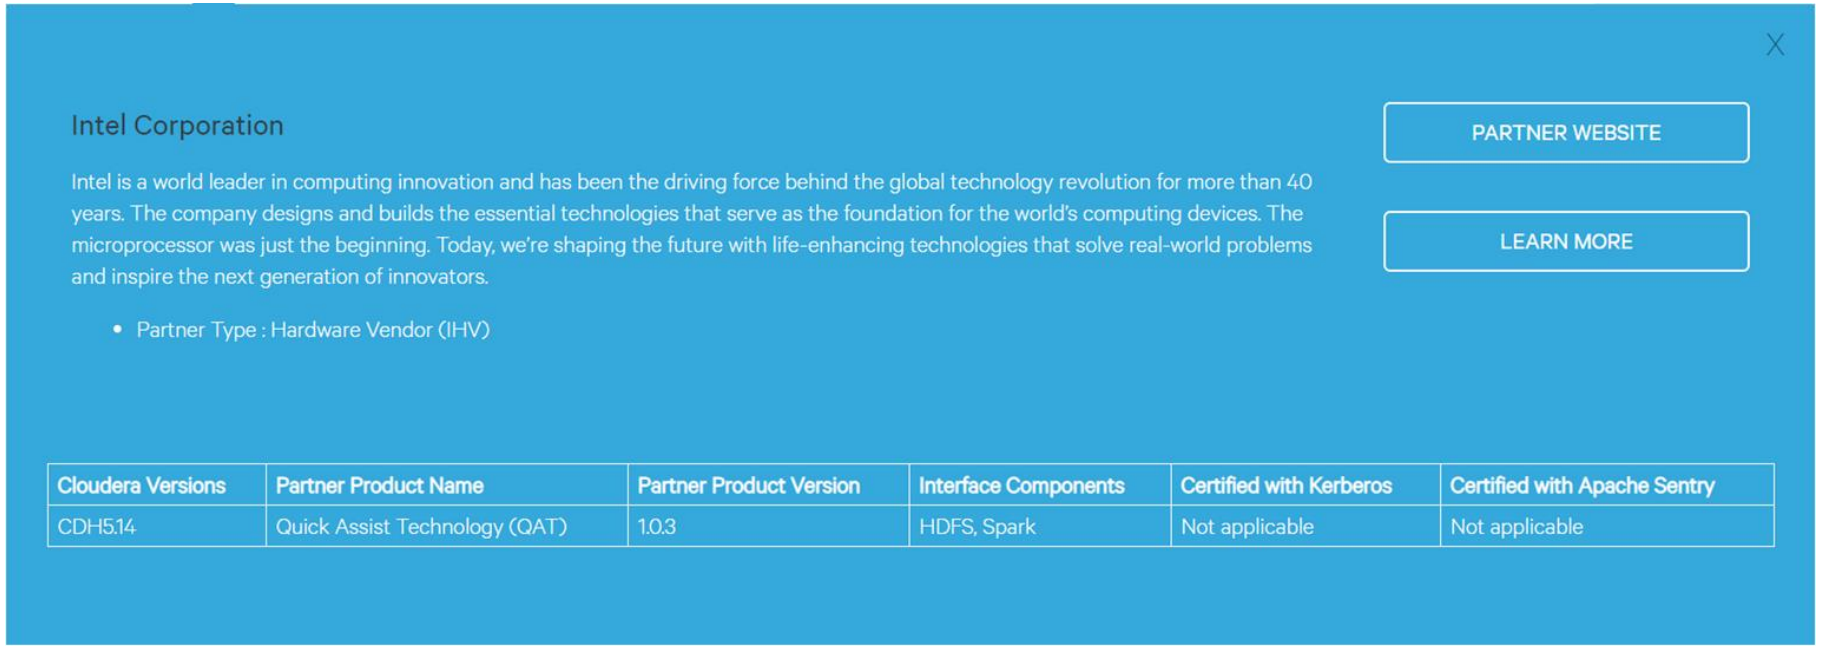

Intel Corporation

Intel is a world leader in computing innovation and has been the driving force behind the global technology revolution for more than 40 years. The company designs and builds the essential technologies that serve as the foundation for the world's computing devices. The microprocessor was just the beginning. Today, we're shaping the future with life-enhancing technologies that solve real-world problems and inspire the next generation of innovators.

- Partner Type : Hardware Vendor (IHV)

PARTNER WEBSITE

LEARN MORE

Cloudera Versions	Partner Product Name	Partner Product Version	Interface Components	Certified with Kerberos	Certified with Apache Sentry
CDH5.14	Quick Assist Technology (QAT)	1.0.3	HDFS, Spark	Not applicable	Not applicable

THE JOURNEY OF QATCODEC

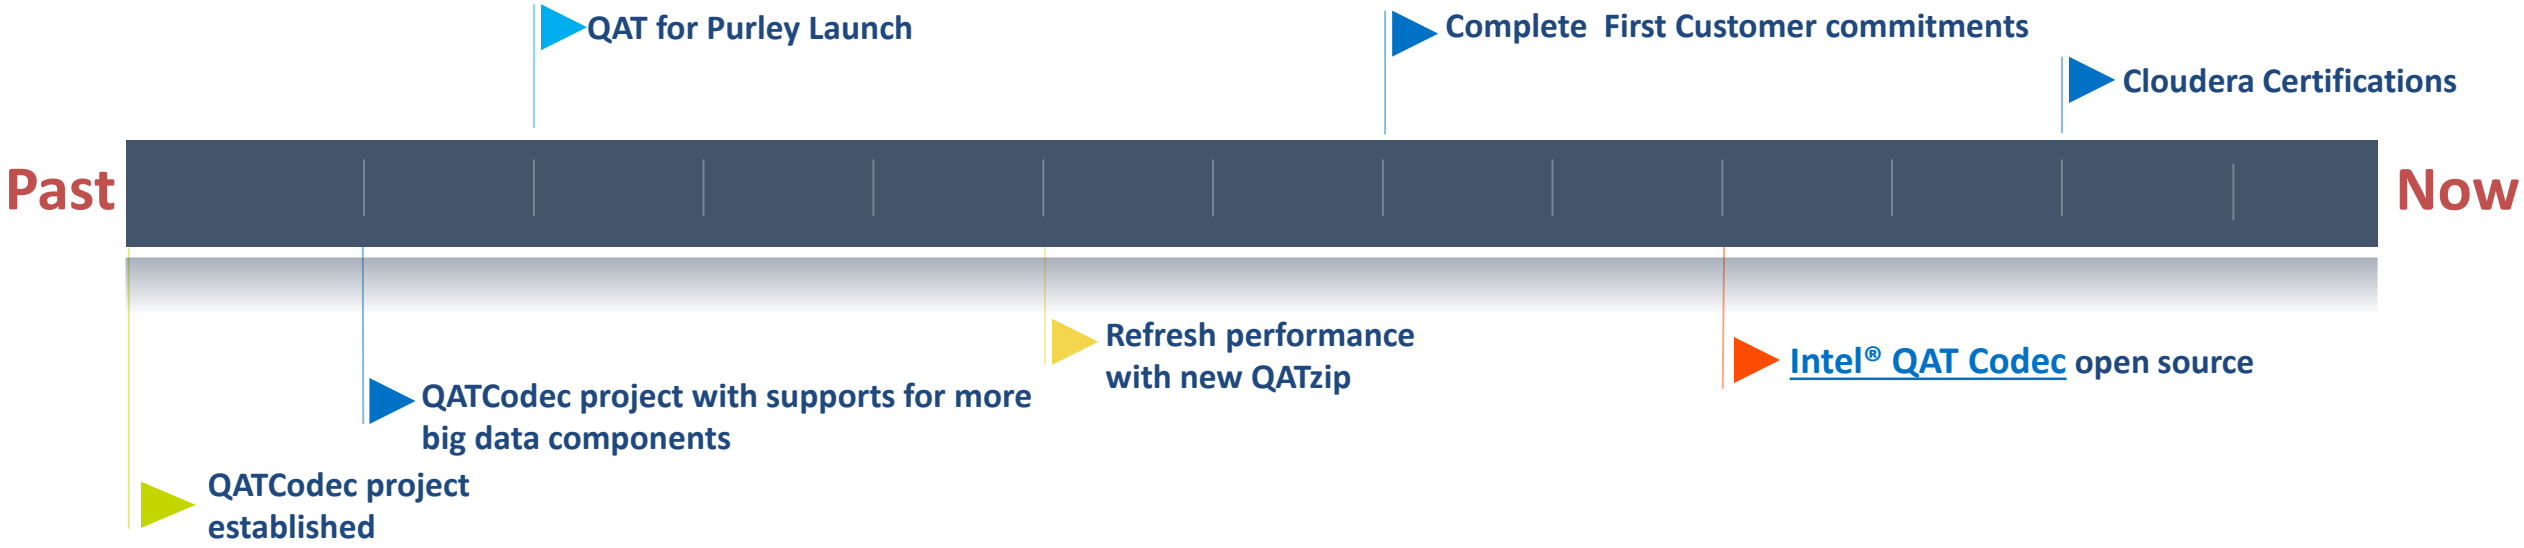

QATCodec project established:
- Begin with Hadoop QAT support
- Jan, 2015

QATCodec project with supports for more big data components
- With supports of Hive, Spark
- Jan, 2017

QAT for Purley Launch
- Report out, our final optimization brings benefits for Hive, Spark, Hadoop
- Aug 2017

Refresh performance with new QATzip
- Refresh performance for the new QATZip
- Nov 2017

Complete first customer commitments
- CDH 5.14.0 version support (rebase & package)
- Integration test/performance test with latest QAT 4.1.0 and QATZIP 0.2.3
- May 2018

Intel® QAT Codec open source
- 6 weeks to achieve Intel open source certification
- July 2018

FUTURE WORK

Continuing engineering collaborations

- More features support: hugepage optimization support
- More big data components supports: Apache CarbonData (Huawei)
- Continuing supports for coming CDH version: CDH 6.0 supports

More customers

- Seek for more deal wins like our first custormer case to bring the hardware accelerations to more big data users.

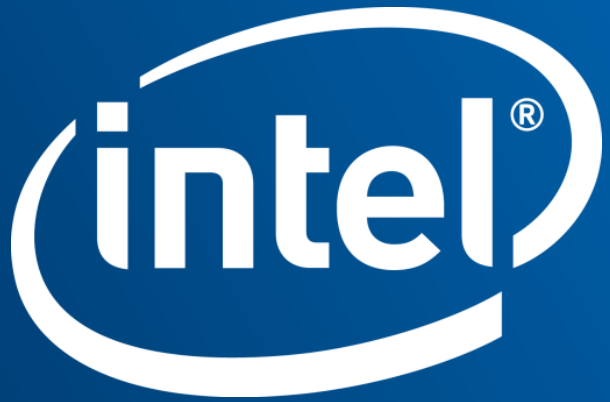
intel