

# Intel Data-Centric Technology Brings Optimizations to Retail Store Operations

**Intel® Optane™ Technology and Apache Spark\* bring cost-effective, fast data processing to retail stores, improving mobile operations and customer experience**



## Executive Summary

Intel, MTS and Huawei joined forces to use data-centric technologies to optimize retail store operations. The solutions, using Intel® Optane™ technology, provide efficient ways to mine large amounts of data and increase the efficiency and effectiveness of business decisions.

## Introduction

Big Data is ubiquitous in the world today. And the explosive growth of data is driven to a large extent by the mobile devices that have numerous applications in everyday life. From financial, social media and entertainment transactions to communications and healthcare applications, these are some of the usage models of mobile devices today that demand a reliable and efficient telecommunication service. Given the abundance of providers, products, and services in the telecom space, customer retention is a high priority for providers. There is a challenging ask for providers to constantly innovate and make their processes as efficient and streamlined as possible. Processes such as workforce scheduling and subscriber management contribute heavily to customer satisfaction — the key element of any successful business. It is not only the underlying mobile infrastructure, but also the operational processes that carry a heavy weight for the success of telecom providers.

Analytics, both traditional and advanced, are critical to operational processes. Where traditional analytics were limited to processing data as batch, advanced analytics can process large amounts of data in or near real-time. In-memory processing is a common technique employed to process data at fast, real-time speeds. Technologies such as Apache Spark\* leverage the fast processing speeds of DRAM to analyze vast amounts of data in a distributed fashion. While some applications may justify the high cost of large capacity DRAM for analytics, many applications are required to operate within budget limitations. Until now, these applications were required to find a trade off between high cost and performance, which created a chasm.

## Introducing Intel® Optane™ DC SSDs with Intel® Memory Drive Technology

The Intel® Optane™ SSD DC P4800X can be configured to use Intel® Memory Drive Technology to transparently expand system memory that provides closer to DRAM-like performance, but at a lower cost. It leverages the high responsiveness and high endurance of Intel® Optane™ SSDs along with a limited amount of DRAM to create a larger memory pool that memory capacity-driven applications can seamlessly leverage. Individual systems are no longer limited to 3TB memory capacity. Intel Memory Drive Technology can provide up to 24TB memory capacities on 2-socket servers. A wide spectrum of applications in Database, Cloud, HPC and Big Data/AI space are potential candidates for Intel Optane technology.

Intel Optane technology does not just increase addressable memory for data algorithms (i.e. quantitative changes), it enables a qualitative leap – transforming analytics/business tasks from traditional batch processing (offline) to close-to-real-time (online) processing without complex software re-design. The Intel Optane SSD with Intel Memory Drive Technology opens the way to move ~1TB datasets into a high core count Intel® Xeon® Scalable processor node and re-balance distributed workloads to enhance overall performance.

Intel Optane DC SSDs can bring transformative technology now — enabling the in-memory journey with existing software and hardware stack. Enterprise CTOs can see immediate results and greater capability with Intel Optane technology.

**Business Case**

MTS is the leading telecommunications group in Russia that provides fixed-line services including home telephony, broadband internet, and pay-tv services. MTS also provides financial services, such as MTS e-wallet to both individual and corporate clients, and cloud computing services that offer data analysis tools based on big data, cybersecurity systems, and intelligent IoT solutions for B2B clients etc.

MTS and Huawei, which specializes in telecommunications equipment, consumer electronics and technology-based services and products, have worked together in the IT domain for three years. Huawei provided MTS with Intel® Xeon® Scalable processor-based servers, which offer low failure rates and world-class performance. Huawei supplied a self-designed BMC chipset for unified management, and fast development to adopt new technologies, such as Intel Optane SSDs. Huawei also accelerated compatibility testing of Intel Optane technology to support MTS business requirements.

The goals for this project included:

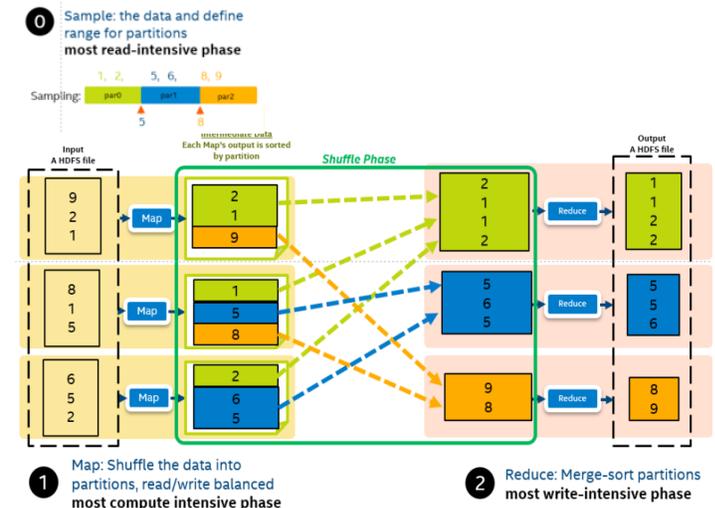
- Increase operational efficiency
- Provide mobile services for 1000+ retail stores
- Need for increasing the efficiency of workforce scheduling across stores
- Data mining and data processing of mobile subscribers' information to increase average monthly payments and improve customer experience
- Support searching of large amounts of data using parallel search queries
- Improve workday planning processes, including faster approvals

**Big Data Workload**

Data mining of unstructured datasets typically involves a sorting function, especially in the MapReduce programming world. Sorting is also a common step in search-related use cases, such as the one described above.

TeraSort, a popular micro benchmark for Spark, was used to compare the performance of sorting 1TB of data using Intel Optane DC SSD P4800X with Intel Memory Drive Technology with a baseline scenario that uses equivalent capacity of DRAM only for memory.

TeraSort is a read, compute, and write intensive workload over three stages (see Figure below).



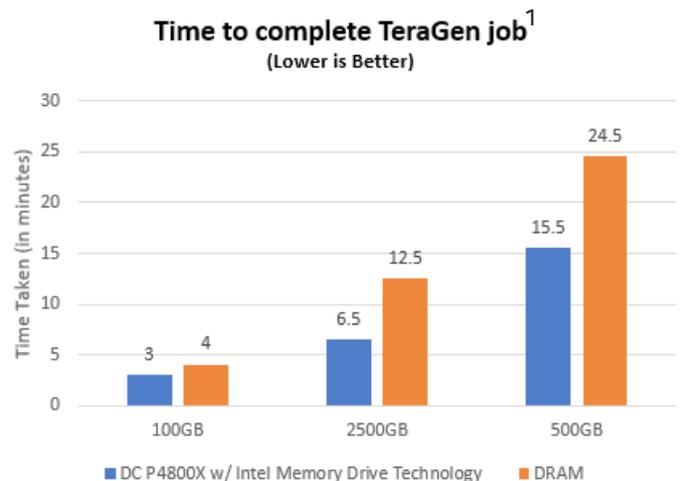
**Test Environment**

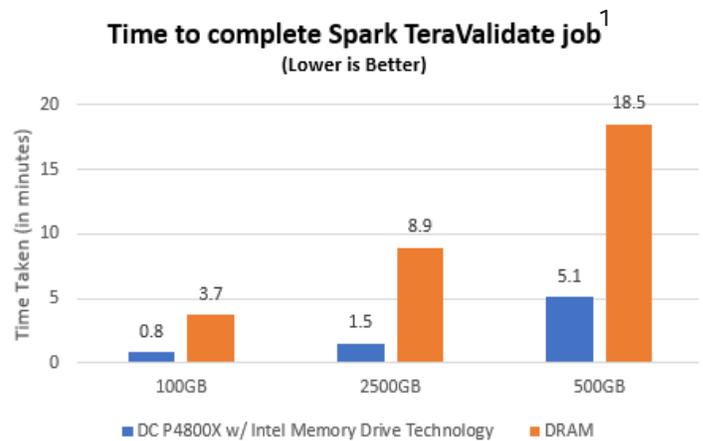
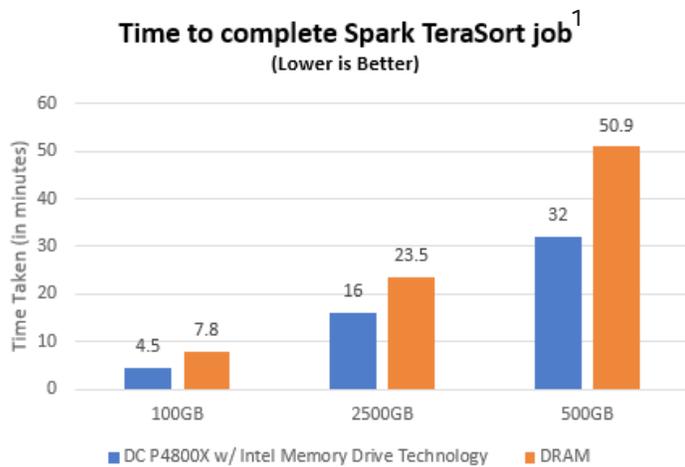
Tests were run in September, 2018 at the Intel Swindon HPC lab. The system used to test with DRAM included a single two socket server with Intel® Xeon® E5-2660 v3 CPU with 40 cores. The software stack used was HDP 2.6.1 with Spark2\* on Oracle Linux Server 7.3\* using kernel version 3.10.0-514.21.1 with a total of 384GB DRAM.

The system used to test with Intel Memory Drive Technology was a single two-socket server with Intel® Xeon® Gold 6148 CPU with 80 cores. The software stack included HDP 2.6.1 with Spark2\* on Red Hat Enterprise 7.3\* using kernel version 3.10.0-693.5.2. Intel Optane SSD DC P4800X configured with Intel Memory Drive Technology provided a total memory capacity of 1TB (DRAM + Intel® Optane™ SSDs).

**Results**

Below are the performance results comparing SSD DC P4800X with Intel Memory Drive Technology and baseline DRAM. With more memory made available at relatively lower cost using SSD DC P4800X with Intel Memory Drive Technology, it is possible to process bigger datasets more efficiently.





## Project Goals Achieved

The performance gains demonstrated by the TeraSort benchmark led to a deeper collaboration between MTS, Huawei and Intel. Based on the benchmark results, MTS decided to configure Intel Optane SSDs with Intel Memory Drive Technology for their data mining and analytics workload. As a result, MTS was able to achieve its goals of increasing operational efficiencies in workforce scheduling and faster processing of subscriber data for greater customer satisfaction.

- Built non-conflicting schedules in real-time (by optimizing the batch process that takes half a day) without significant budget increases
- Deep dependency extractions and rebuilt dependency chains
- Reached targeted state to submit the form at which point the new system rebuilt all the dependencies in real-time in one minute or less
- Moved data from relatively slow storage to a faster storage medium and achieved real-time processing speeds
- Built a platform for future expansion into deep learning and leverage AI technologies

Contributors: Grigorii Koval - Big Data Architect, MTS; Konstantin Yakovlev - Chief System Architect, MTS; Dmitry Lampsi - Lead Developer, MTS; Xu Hui - Director of Huawei EuroAsia, Intelligent Computing Business Unit; Ravikanth Durgavajhala - Enterprise Architect, Intel; Andrey Kudryavstev - Enterprise Architect, Intel; Michel Abou Diwan - Regional Alliances Director EMEA, Intel



Learn more now at [intel.com/optane](https://intel.com/optane)

Performance results are based on testing as of September 2018 and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at [intel.com](https://intel.com), or from the OEM or retailer.

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 complete information visit [www.intel.com/benchmarks](https://www.intel.com/benchmarks).

1. Source: Intel Tested at Intel Swindon HPC lab. System test with DRAM: single two socket server with Intel® Xeon® E5-2660 v3 CPU with 40 cores. Software stack: HDP 2.6.1 with Spark2\* on Oracle Linux Server 7.3\* using kernel version 3.10.0-514.21.1 with a total of 384GB DRAM. System test with Intel Memory Drive Technology: single two-socket server with Intel® Xeon® Gold 6148 CPU with 80 cores. Software stack: HDP 2.6.1 with Spark2\* on Red Hat Enterprise 7.3\* using kernel version 3.10.0-693.5.2. System BIOS: SE5C620.86B.00.01.0014.070920180847. Intel Optane SSD DC P4800X configured with Intel Memory Drive Technology provided a total memory capacity of 1TB (DRAM + Intel® Optane™ SSDs).

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase.

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.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document. The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

All rights reserved. Intel, Optane, Xeon 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.