**White Paper**
**Intel® Xeon® processor E7 v2**
Big Data Analysis

# Accelerate Big Data Analysis with Intel® Technologies

## Executive Summary

It's not very often that a disruptive technology changes the way enterprises operate. In the same way that the microprocessor, the personal computer, and virtualization changed the computing landscape, Apache Hadoop* and the MapReduce framework have forever changed the way that enterprises capture, store, and analyze information. But while yesterday's hardware technology helped lay the foundation for big data analysis, today's technologies let enterprises analyze more data faster than ever before.

Intel is leading the big data infrastructure charge with advancements in microprocessors, storage, and networking. These advancements can help increase the scalability and performance of large Apache Hadoop clusters. In internal tests that compared servers equipped with the previous generation Intel® Xeon® processor E5 family and servers equipped with the Intel® Xeon® processor E7 v2 family, the servers equipped with the Intel Xeon processor E7 v2 family demonstrated performance gains of up to 3.5 times across a spectrum of I/O- and CPU-intensive workloads.[1]

This paper highlights technologies available from Intel that enterprises can use to scale up Apache Hadoop clusters to handle the increasing volume, variety, and velocity of data. Enterprises can reduce the complexity and total cost of ownership (TCO) of their clusters by using fewer, more powerful servers, which can reduce operational costs up to 37 percent overall over a four-year period.[2]

## Apache Hadoop* Overview

Apache Hadoop is a distributed data storage and data processing platform that enterprises can use for storing and processing large amounts of semi-structured or unstructured data. Built on Java*, Apache Hadoop has open-source roots and enjoys the support of a large, active user and developer community. Apache Hadoop also benefits from the collaborative work of Java Virtual Machine (JVM) vendors and Intel engineers to increase Java performance on the latest Intel platforms. These traits help make Apache Hadoop a cost-effective, high-performance platform for enterprises to gather and analyze data from such varied sources as point-of-sale systems, credit card transactions, server log files, machine logs, and scientific sensors. Intel has worked with Java Virtual Machine vendors for more than 10 years to optimize Java performance on Intel hardware, as each new generation of Intel microarchitecture provides new features that can increase software performance. All of this capability enables advanced analytics for a range of tasks, from detecting credit card fraud to decoding the human genome.

An Apache Hadoop cluster can scale from a few servers to thousands. This flexibility makes Apache Hadoop an ideal platform across the data analysis spectrum, from

**Tim Allen**
Big Data Domain Expert
Intel Software & Services Group

**Eric Kaczmarek**
Big Data Performance Architect
Intel Software & Services Group

**Frank Jensen**
Performance Marketing Engineer
Intel Data Center Group Marketing

## Table of Contents



enterprises with small data analysis needs to enterprises with increasing data analysis requirements that need clusters that can scale with faster, more powerful hardware.

Three components make up the core of the Apache Hadoop version 1.x:

- **Apache Hadoop Distributed File System (HDFS*)**, which provides a high-performance file system that can span and replicate data across the nodes of an Apache Hadoop cluster. Important features of HDFS include fault tolerance and performance for large datasets.
- **MapReduce**, a processing framework that provides parallel processing across large, unstructured datasets. MapReduce includes two functions: map, which sorts and filters the data, and reduce, which further processes the output of map into a final result.
- **Apache Hadoop Common**, which ties HDFS and MapReduce together.

An Apache Hadoop cluster consists of master nodes and worker nodes. When a client sends a request to a master node, the node processes the request with two components:[3]

- ▪ **NameNode**, a component of HDFS that keeps track of data within the cluster nodes.
- ▪ **JobTracker**, which reduces an analysis request into smaller tasks based on where in the cluster the data resides, and then assigns those tasks to specific worker nodes.

After a master node processes a request, it communicates with three services on the worker nodes:

- **DataNode**, a component of HDFS that manages data on the worker nodes.
- **TaskTracker**, a service that receives and runs MapReduce tasks from a master node's JobTracker service.
- **MapReduce**, which performs the assigned tasks.

As MapReduce on each worker node finishes its assigned tasks, the worker nodes return the results to the master node. Since the tasks can run in parallel on multiple worker nodes, the master node waits for all of the tasks to complete on the worker nodes, compiles the results, and then returns the combined result to the client.

## Performance Bottlenecks

Apache Hadoop benefits from its distributed architecture, as worker nodes do not require high-availability configurations due to the HDFS ability to create multiple copies of data across the worker nodes. Any worker node within the cluster can fail without data loss or interruption to the rest of the cluster.

But as the number of worker nodes in an Apache Hadoop cluster increases, the strain on the master node—specifically the NameNode and JobTracker services—increases. As the volume and velocity of data increases, the master node services can become overwhelmed, reducing performance across the cluster.

Networking and storage I/O bottlenecks can also affect cluster performance. A master node must wait for all tasks on the worker nodes to complete before it can compile the results and return the results to the client. Therefore, slow worker nodes—whether they are hampered by CPU or I/O speeds—can hamper analytics and batch tasks. At the worker node, reading data from disk into memory to perform a task, and then sending the results across the network to the master node can introduce delays, especially where high-velocity data is concerned.

### Increase Apache Hadoop Cluster Performance with Intel® Technologies

Intel provides a number of technologies that can help dramatically improve Apache Hadoop performance across CPU- and I/O-intensive workloads. Combined, these technologies can help enterprises scale Apache Hadoop to address increasing performance requirements—such as analyzing high-velocity data—while relieving performance bottlenecks within the cluster.

### Intel® Xeon® Processor E7 v2 Product Family: Performance for CPU-Intensive Workloads

The Intel Xeon processor E7 v2 family, which is built on the Intel 22 nanometer process, reaches new levels of processing density with up to 15 cores and, with Intel® Hyper-Threading Technology (Intel® HT Technology)[4] enabled, 30 logical processors per socket. The family supports two-, four-, and eight-socket configurations natively, which provides a maximum of 120 cores and 240 logical processors per server. Enterprises can extend socket configurations even further with third-party controllers. With more cores and threads available, enterprises can deploy Apache Hadoop clusters with greater processing capabilities while using fewer servers.

A four-socket server configuration can support up to 6 terabytes of memory, while an eight-socket configuration can scale memory to 12 terabytes. Enterprises can use larger memory configurations to temporarily store frequently used or high-velocity data for analysis by Apache Hadoop services. For example, system engineers can configure servers to use a portion of a server's RAM for data storage space and then configure Apache Hadoop to use this space for temporary data storage. Since RAM is orders of magnitude faster than disk-based storage, Apache Hadoop can use the RAM-based storage to analyze larger amounts of high-velocity data faster than if the data were located on disk-based storage. After Apache Hadoop completes analysis, it can then write the data to slower disks for longer-term storage.

The Intel Xeon processor E7 v2 family also builds on a tradition of advanced reliability, availability, and serviceability (RAS) features that can give Apache Hadoop clusters better recovery from server hardware failures. Designed for systems with 99.999% uptime requirements, the Intel Xeon processor E7 v2 family provides continuous self-monitoring and self-healing capabilities that rival those of RISC-based systems.[5] Some of these features include:

- Machine Check Architecture (MCA) Recovery, which lets the CPU and operating system isolate errors that could normally crash a server, such as unrecoverable memory errors.
- MCA Recovery Execution Path, which handles uncorrectable data errors passed to the CPU. This feature enables operating systems and applications to assist in recovering from errors that cannot be corrected at the hardware level.
- MCA I/O, which provides information on uncorrectable I/O errors so that the operating system can take action. Operating systems or monitoring tools can use this information to determine the cause of system errors and enable preventive maintenance.
- Enhanced Machine Check Architecture (eMCA) Gen 1, which provides enhanced logging information to the operating system and applications that can better diagnose errors and proactively predict failures.
- PCIe Live Error Recovery (LER),[6] which provides recovery from and containment of PCIe errors.

Combined, the features of the Intel Xeon processor E7 v2 family complement those of Apache Hadoop to help enterprises increase the computing capabilities and reliability of their Apache Hadoop clusters with fewer servers, which can lower the overall TCO. Fewer servers means less complexity and lower power and cooling requirements.

### Intel® Solid-State Drives: High-Performance Storage for I/O-Intensive Workloads

Intel continues to be a leader in solid-state drive (SSD) technology, with drives that provide performance that is dramatically higher than that of mechanical hard disks, combined with greater reliability. Intel® SSDs eliminate the mechanical limitations of hard disks and provide higher I/O operations per second (IOPS) and increased mean time between failures (MTBF). Apache Hadoop clusters that require higher performance storage—such as those that perform real-time analysis on in-flight data—can benefit from Intel SSDs.

Intel SSDs can provide higher throughput than traditional mechanical hard drives, which can reduce the risk of storage bottlenecks on the cluster nodes. When tested against mechanical hard drives, SSDs have shown that they can deliver up to 2.7 times the throughput for I/O-intensive Apache Hadoop workloads, such as Sort.[7]

In addition to being more reliable, Intel SSDs also require less power and cooling than traditional mechanical hard drives, which can lead to greater node uptime and lower TCO.


**Figure 1:** Intel® SSDs can improve Apache Hadoop* throughput by up to 2.7 times over traditional mechanical hard disks.[7]

## Data Loading and Replication Performance

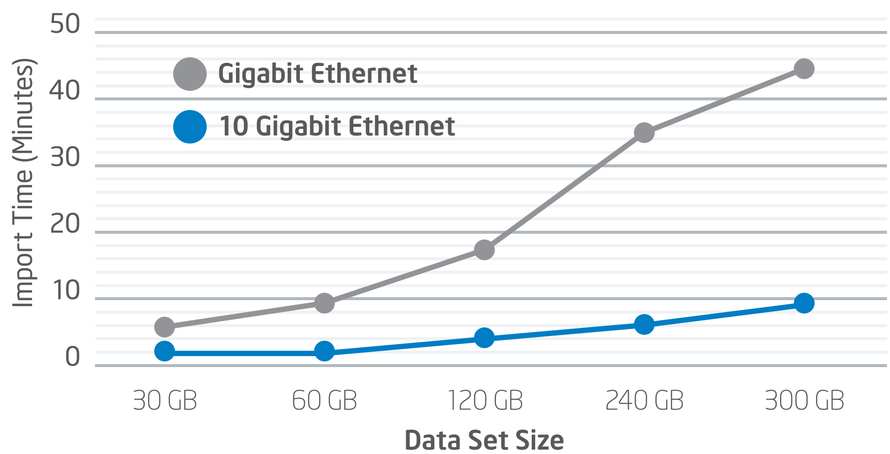


**Figure 2:** The 10 gigabit Intel® Ethernet Server Adapters demonstrate a five-fold increase in data loading and replication performance.[8]

### Intel® Ethernet Server Adapters: Higher Throughput for Distributed Clusters

The distributed architecture of Apache Hadoop depends heavily on fast and reliable network communication. Many enterprises use gigabit Ethernet (GbE) network fabrics to connect Apache Hadoop nodes, but as the frequency of workload requests and data velocity increases, combined with faster CPUs and storage, network speeds must also increase.

A common tool for increasing network throughput is Ethernet bonding, where multiple physical Ethernet ports are bonded together into a higher-bandwidth logical Ethernet port. This method can provide short-term performance gains, but it increases complexity and costs. The 10 gigabit Intel® Ethernet Server Adapters provide higher performance while decreasing port counts, cabling, and power consumption. More importantly, 10 GbE also provides scalability benefits for Apache Hadoop clusters.

In recent tests, Intel engineers tested the performance of 1 GbE and 10 GbE networks when importing data into an Apache Hadoop cluster and replicating it across worker nodes. The testing results demonstrated a five-fold increase in loading times using 10 GbE.[8]

## Putting It All Together: Benchmarking Apache Hadoop Clusters with Intel Technologies

In recent internal tests, Intel engineers combined high-performance Intel CPU, SSD, and networking technologies to determine the performance benefits across a range of CPU- and I/O-intensive Apache Hadoop workloads. These workloads included:

- **Sort**, an I/O-intensive workload that transforms data from one format to another. Sort is representative of a typical real-world MapReduce task.
- **TeraSort**, a popular industry-standard benchmark for large-size data sorting.
- **K-means**, a CPU-intensive workload that uses a well-known clustering algorithm for data mining and machine learning.
- **Apache* Hive Join**, a workload that is both CPU- and I/O-intensive. This workload provides performance benchmarks for more structured datasets.
- **Page Rank**, a MapReduce workload that uses a well-known search engine algorithm that ranks pages.

The tests consisted of running the workloads against two configurations:

- A six-node, fully optimized baseline configuration that used dual-socket servers based on the Intel Xeon processor E5-2680 with Intel SSDs and 10 gigabit Intel Ethernet Server Adapters.
- A three-node enhanced configuration that used four-socket servers based on the Intel Xeon processor E7-4890 v2 with Intel SSDs and 10 gigabit Intel Ethernet Server Adapters.

The results are normalized for the servers configured with the Intel Xeon processor E5 family.

The benchmarks demonstrate significant performance gains from the servers equipped with the Intel Xeon processor E7 v2 family over servers equipped with the previous generation Intel Xeon processor E5 family. The I/O-intensive workloads—Sort and Page Rank—showed a 2.6 and 2.7 times performance advantage, while CPU-intensive workloads—Apache Hive Join and K-means—showed the greatest performance advantage at 3.2 and 3.5 times the performance of the servers equipped with the previous generation Intel Xeon processor E5 family.[1] TeraSort, which is both I/O- and CPU-intensive, performed nearly 2.9 times faster.[1]

With memory configurations up to 12 terabytes of RAM, enterprises can increase the amount of in-memory buffer used by MapReduce tasks, which can significantly reduce data movement between memory and disk-based storage. In servers with smaller memory configurations, MapReduce tasks can quickly fill the buffer, which means that the contents of the buffer are automatically moved—or "spilled"—out of memory and onto disk. Moving data from memory to disk is time-consuming and also increases Java garbage collection tasks, which negatively impacts node performance. With a large memory server configuration, systems engineers can configure the buffer to use more available memory to avoid spilling to disk altogether, which reduces I/O and unnecessary Java garbage collection. The Apache Hive Join workload benchmark demonstrates the benefits of a larger memory configuration by yielding a 29 percent performance gain between the 512 GB and 1.5 TB memory configurations, with an estimated cost of only 20 percent for the additional memory.[1]

Enterprises can also realize lower power and cooling costs by using servers equipped with the Intel Xeon processor E7 v2 family. On the CPU-intensive K-means workload, the Intel Xeon processor E7-4890 v2 consumed approximately 1,162 watts of power, while the previous generation Intel Xeon processor E5 consumed approximately 466 watts of power. However, the previous generation Intel Xeon processor E5 took 3.5 times longer to complete the same workload, which resulted in the workload consuming a total of 1631 watts of power—29 percent more than the Intel Xeon processor E7 v2 when normalized to performance.[1]

10 to 3

Average Workload Performance Ratio

## Apache Hadoop* Workload Performance

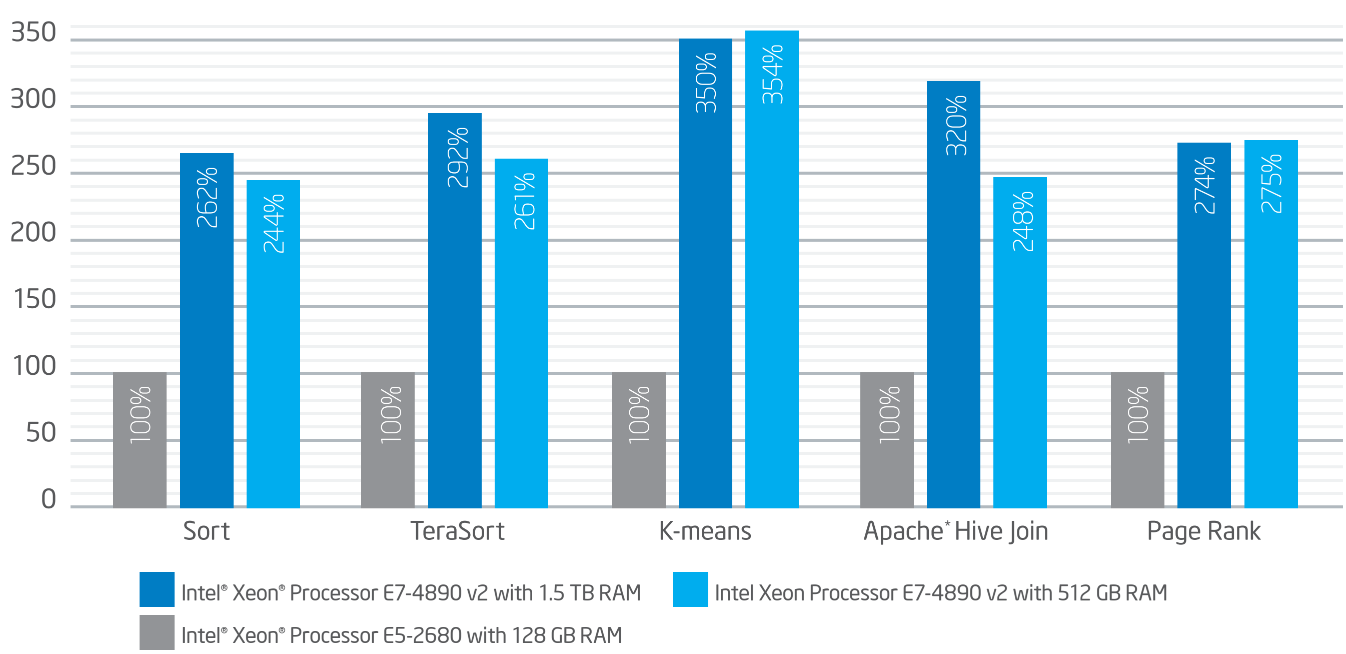


**Figure 3:** Servers equipped with the Intel Xeon E7-4890 v2 provide higher Apache Hadoop workload performance over servers equipped with the the previous generation Intel Xeon E5-2680.[1]

**TABLE 1: FOUR-YEAR OPERATIONAL COSTS COMPARISON BETWEEN A 1000-NODE INTEL XEON PROCESSOR E5-2680 CLUSTER AND A 334-NODE INTEL XEON PROCESSOR E7-4890 V2 CLUSTER.**

| ANNUAL COSTS PER SERVER[2] | MAINTENANCE | POWER AND COOLING YEARS 1 THROUGH 4 | RACK SPACE | NETWORK MAINTENANCE | TOTAL COSTS |
|---|---|---|---|---|---|
| 2-socket Intel® Xeon® processor E5-2680 (2.7 GHz, 8 cores) | $500 | $496 | $155 | $15 | $1166 |
| 4-socket Intel® Xeon® processor E7-4890 v2 (2.8 GHz, 15 cores) | $500 | $1443 | $310 | $30 | $2283 |
| **TOTAL 4-YEAR CLUSTER COST AT EQUIVALENT PERFORMANCE** | | | | | |
| 1000-node cluster with Intel Xeon processor E5-2680 | $2,000,000 | $1,984,185 | $620,000 | $60,000 | $4,664,185 |
| 334-node cluster with Intel Xeon processor E7-4890 v2 servers | $668,000 | $1,835,087 | $414,160 | $40,080 | $2,957,327 |
| **CLUSTER COST SAVINGS USING SERVERS BASED ON THE INTEL XEON PROCESSOR E7 V2 FAMILY** | 67 percent | 8 percent | 33 percent | 37 percent | 37 percent overall cost savings |

Table 1 provides a sample four-year operational cost comparison between a 1000-node Intel Xeon processor E5-2680 server cluster, and a 334-node cluster based on Intel Xeon processor E7-4890. Using the average of the Sort, TeraSort, K-means, Apache Hive Join, and Page Rank benchmarks on the 1.5 TB configuration as a baseline, a 334-node cluster can do the same amount of processing as the 1000-node cluster. The initial acquisition costs for both clusters vary based on the OEM and configuration, but the operational costs over a four-year period are considerably less for the 334-node cluster.[2]

By deploying higher performance servers with the Intel Xeon processor E7 v2 family, enterprises can scale up their Apache Hadoop clusters by using fewer, more powerful servers. A scale-up architecture reduces cluster complexity and can reduce operational costs, which can lead to lower TCO.

## Scale Up with Intel Technologies

Many enterprises have built their Apache Hadoop clusters on lower-performing commodity hardware. Yet advances in Intel CPU, networking, and storage technologies can greatly enhance the performance of big data analysis, especially in large Apache Hadoop clusters. These technologies include:

- The Intel Xeon processor E7 v2 family, which provides up to 15 cores per socket and up to 12 terabytes of RAM in an eight-socket configuration for faster access to data and greater processing speed. New RAS features can also increase the reliability of Apache Hadoop clusters.
- Intel SSDs, which provide greater reliability and performance that is significantly higher than traditional mechanical hard disks.
- 10 gigabit Intel Ethernet Server Adapters, which can increase the networking throughput across Apache Hadoop clusters while decreasing port count and cabling complexity.

For more information on how you can increase your enterprise's big data performance, go to www.intel.com/hadoop.

1 Based on Intel Corporation internal testing. The Intel® Xeon® E5-2680 Apache Hadoop cluster consisted of a master node and five worker nodes. All nodes were configured with dual Intel Xeon E5-2680 processors running at 2.7 GHz, 128 GB of RAM (8x16 GB DDR3 1600 Mhz), 3 Intel SSD DC S3700 Series SSDs (400GB, 2.5in SATA 6Gb/s, 25nm, MLC, Seq. Read: 500 MB/s, Seq. Write 460 MB/s), 3 Intel SSD 710 Series SSDs (200GB, 2.5in SATA 3Gb/s, 25nm, MLC, Seq. Read: 270MB/s, Seq. Write 210MB/s), 2 Intel SSD 710 Series SSDs (300GB, 2.5in SATA 3Gb/s, 25nm, MLC, Seq. Read: 270MB/s, Seq. Write 210MB/s, an Intel X25-E Extreme solid state drive for the operating system (SATA 2, 2.5 inch, SLC, High Performance), CentOS Release 6.2 with kernel 2.6.32-358.el6.x86_64, Apache Hadoop 1.0.4, Intel® 82599EB 10 gigabit Ethernet controller, and an LSI Logic SAS 2308 PCI-Express Fusion-MPT SAS-2 RAID controller. The estimated street price for each node is $15,000.

The Intel® Xeon® E7-4890 v2 Apache Hadoop cluster consisted of a master node and two worker nodes. All nodes were configured with four Intel Xeon E7-4890 v2 processors running at 2.8 GHz, 512 GB of RAM (32x16 GB DDR3 1600 MHz), 12 Intel SSD DC S3700 Series SSDs (400GB, 2.5in SATA 6Gb/s, 25nm, MLC, Seq. Read: 500 MB/s, Seq. Write 460 MB/s), an Intel X25-E Extreme 32 GB solid state drive for the operating system (SATA 2, 2.5 inch, SLC, High Performance), CentOS Release 6.2 with kernel 2.6.32-358.el6.x86_64, Apache Hadoop 1.0.4, Intel® 82599EB 10 gigabit Ethernet controller, and an Intel RAID Controller RS2BL040. The cluster with servers configured with 1.5 TB of RAM used the same configuration except for 96x16 GB DDR3 1600 MHz RAM modules. The estimated street price for each node configured with 512 GB of RAM is $65,000, and $80,000 for the nodes with 1.5 TB of RAM. Each node in both clusters used the following settings: BIOS (Intel® TurboBoost enabled, Intel® SpeedStep disabled), Apache Hadoop (DFS replication set to 1, JVM reuse set to enabled), operating system ('deadline disk scheduler', transparent large pages disabled, IPV6 disabled, iptables disabled, and time synchronized with NTP).

2 The results shown in Table 1 were generated using an Intel internal total cost of ownership tool (the Tool) as of December 2013, normalizing database transactional performance between the two options. These results and the scenarios described are estimates and should only be used as a guide to evaluate the cost/benefit or feasibility of a future purchase of systems with Intel Xeon Processors. Prior to purchasing any systems with Intel Xeon Processors you should consult with qualified IT professionals to ensure that any such purchases fit your particular circumstances. Intel does not and cannot guarantee the accuracy or reliability of the results generated by the Tool or that you will actually realize any cost savings as forecast by the Tool and shown in the table.

This analysis assumes an average of three (3) times performance increase in tested workloads between clusters equipped with two-socket servers based on Intel® Xeon® processor E5-2680 and 4-socket servers based on Intel Xeon processor E7-4890 v2 processors. Calculations include analysis based on performance, power, cooling, electricity rates, rack space and network maintenance. Assumptions include 42U racks, $0.10 per kWh, cooling costs 2x average server power consumption costs, Alinean* assumptions of $500 per server maintenance and $30 per server networking costs, average real estate cost per year from VMware* Cost-Per-Application Calculator at $310 per sq. foot * 10 sq. feet per rack divided by the number of servers per rack, 60% CPU utilization and PUE of 2.0. Cluster cost savings based on comparing the performance of deploying 1,000 two-socket servers based on Intel® Xeon® processor E5-2680 compared to equivalent performance of four-socket servers based on Intel Xeon processor E7-4890 v2. For more information about the VMware Cost-Per-Application Calculator, go to http://www.vmware.com/files/pdf/vmware-cost-per-application-calculator-methodology.pdf.

3 A master node can also host DataNode and TaskTracker.

4 Available on select Intel® Core™ processors. Requires an Intel® Hyper-Threading Technology-enabled system; consult with your PC manufacturer. Performance will vary depending on the specific hardware and software used. For more information, including details on which processors support Intel HT Technology, visit http://www.intel.com/info/hyperthreading.

5 No computer system can provide absolute reliability, availability or serviceability. Requires an Intel® Xeon® processor E7-8800/4800/2800 v2 product families or Intel® Itanium® 9500 series-based system (or follow-on generations of either.) Built-in reliability features available on select Intel® processors may require additional software, hardware, services and/or an internet connection. Results may vary depending upon configuration. Consult your system manufacturer for more details.

6 PCIe Live Error Recovery implementations are specific to an OEM's implementation. Consult your OEM for specific recovery capabilities.

7 Based on Intel Corporation internal testing. The Apache Hadoop cluster consisted of a master node and two worker nodes. The master node was configured with dual Intel® Xeon® X5570 processors running at 2.93 GHz and an Intel 82574L gigabit Ethernet adapter. The worker nodes were configured with dual Intel Xeon E5-2680 processors running at 2.7 GHz, 128 gigabytes of DDR3 RAM, 16 hardware threads (Intel Hyper-Threading disabled), and an Intel 82598EB 10 gigabit Ethernet adapter. The software configuration consisted of CentOS 6.3 64-bit with kernel 2.6.32-279.19.1.el6.x86_64, Java Hotspot™ 1.7.0_13 64-bit server VM (build 23.7-b01, Apache Hadoop 1.0.4 with performance tuning applied, and compression provided by Snappy 1.1. The tests focused on I/O-intensive symmetric workloads. The tests consisted of a Sort test using 250GB of seed data, and a Terasort test with highly compressible data (1TB of data compressed to 160GB with zlib). Note that Apache Hadoop requires advanced tuning in order to take full advantage of SSD throughput. See http://software.intel.com/sites/default/files/Optimizing%20Hadoop%20Deployments.pdf for more details.

8 Based on Intel Corporation internal testing. Performance comparison using geometric mean of SPECint*_rate_base2006, STREAM*_MP Triad, and Linpack* benchmark results. Baseline geometric means score of 166.75 on prior generation 2S Intel® Xeon® processor X5690 platform based on best published SPECrate* scores to www.spec.org and best Intel internal measurements on STREAM_MP Triad and Linpack as of 5 December 2011. New geometric mean score of 306.74 based on Intel internal measured estimates using an Intel® Rose City platform with two Intel® Xeon® processor E5-2690, Turbo and EIST Enabled, with Hyper-Threading, 128 GB RAM, Red Hat Enterprise Linux* Server 6.1 beta for x86_6, Intel® Compiler 12.1, THP disabled for SPECfp_rate_base2006 and enabled for SPECint_rate_base2006.

Results have been estimated based on internal Intel analysis and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance.

Intel does not control or audit the design or implementation of third party benchmark data or Web sites referenced in this document. Intel encourages all of its customers to visit the referenced Web sites or others where similar performance benchmark data are reported and confirm whether the referenced benchmark data are accurate and reflect performance of systems available for purchase.

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 SSE3 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

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.



Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See www.intel.com/products/processor_number for details.



Printed in USA 0114/TA/PRW/PDF Please Recycle 330050-001US