WHITE PAPER

SGI® Hadoop™ Starter Kits for Big Data Analytics

July, 2012

TABLE OF CONTENTS

1.0 Why Hadoop?

IDC estimates that in the next five years, the amount of information created and replicated will grow by a factor of nine to more than 1.8 zettabytes (1.8 trillion gigabytes), while according to *The Economist,* only 5% of information created is structured - exacerbating the problem of how to analyze and derive quality business insights from the remaining 95% that is unstructured or semi-structured data.

Unstructured or 'Big Data' is data that comes in different shapes and sizes, is frequently changing, can be very short-lived, generally arrives in huge volumes, and originates from a broad variety of sources such as experimentation, web stream, point-of-sale, RFID tags, sensor arrays, etc. Unstructured data is also very difficult to analyze and manage using traditional means such as relational databases - which is why, by enabling this, the Hadoop™ middleware platform is proving so popular for organizations of all sizes in markets such as:

- **Massive In-Core Computation.** The SGI UV 2000 allows much larger and more detailed models and simulations of physical systems, or any large data set, to be entirely memory resident.
- **Biosciences** — for pharmacological trials
- **Federal & defense** — for fraud detection, predictive demographics, signal analysis, trend analysis and security analysis
- **Financial services** — for automated and algorithmic trading, risk analysis, and detecting fraud in credit transactions / insurance claims
- **Retail** — for analysis of customer buying behaviors and inventory management
- **Science and research** — for large-scale experiments (e.g. the Large Hadron Collider), continental-scale experiments and environmental monitoring, instruments and sensors (e.g. the Large Synoptic Survey Telescope)
- **Social media** — for click stream analytics, user search patterns and behavioral analysis
- **Telecom** — for customer trend analysis, network usage patterns and fraud detection

2.0 SGI and Hadoop

SGI has been supplying SGI® Rackable™ clusters for Hadoop since the technology's first commercial deployments - including the largest Hadoop clusters (up to 4,000 nodes) and installations (up to 40,000 nodes) in existence today. SGI Rackable clusters push the envelope in Hadoop performance (SGI holds the world record for the TeraSort data processing and analysis benchmark, set in October 2011), and customers need this unprecedented speed and scalability because the computational and storage capabilities in Hadoop clusters scale very quickly with the amount of data they are analyzing.

Design to Order enables SGI to deliver Hadoop clusters optimized for specific customer requirements in terms of data flows and analyses, data ingress, power, density, price/performance, or total cost of ownership (TCO). Through SGI® Management Center (including power optimization), SGI also excels in providing ease of use and manageability for very large configurations – traditionally a significant challenge for Hadoop; while through the SGI Management Center - Power Option, SGI clusters are the world's first green Hadoop clusters, offering a range of options allowing customers to run Hadoop within a data center power envelope, derive the maximum operations per watt for their applications, etc.

3.0 SGI Hadoop Starter Kits

For customers who cannot or do not want to build their own Hadoop clusters, or who need the shortest time to production, SGI offers solution-oriented clusters including pre-configured Starter Kits based on the latest Intel® Xeon® Processor E5 family and SGI Rackable half-depth or standard-depth servers, which

can be ordered 'off the shelf.' With the considerable growth in the Hadoop stack and applications, SGI has also entered a full partnership and reseller agreement with Cloudera - the premier provider of Hadoop software solutions - to distribute the company's entire product line, services and training, and pre-install and factory-integrate Cloudera software on SGI Hadoop clusters for easy, 'start up and go' deployment. All SGI Hadoop Starter Kits are therefore Cloudera certified.

This whitepaper summarizes the data integration capabilities of SGI Hadoop Starter Kits, together with the ecosystem of Business Intelligence (BI) applications supplied with them – including Datameer®, Kitenga®, Pentaho® and Quantum4D® solutions and their uses. To find out more about the detailed specifications, capabilities and performance metrics of SGI Hadoop Starter Kits in providing application developers, end customers or ISVs a best-in-class, differentiated BI solution built on Hadoop, please also see the accompanying whitepaper *'SGI Hadoop Starter Kits for Big Data Analytics'*.

3.1 Addressing Big Data's Major Pain Points

The major pain points associated with using Hadoop to analyze Big Data are the complexity, volume, velocity and variety of unstructured data, and how to optimize performance in finding value in this data. While Hadoop is an increasingly widely used technology for processing Big Data, challenges such as sizing, configuring and optimizing Hadoop to meet end-user requirements can still be major hurdles. So too can choosing the right applications for Hadoop, as these need to be written to scale, with streaming access to data, and write-once-read-many access to files.

Other challenges include how to apply knowledge of analytics in the structured space to the unstructured domain; and how to enable users to ingest and analyze high velocity data in real-time, and then collect that data into a historical data store like Hadoop for further analysis. Achieving this requires an integrated solution with the speed and scale to address the differing levels of throughput, processing rates and availability of information, and so deliver a seamless flow of information across the enterprise.

The key capabilities of SGI Hadoop Starter Kits in addressing these challenges are as follows:

- **Solving the complexity of Hadoop deployment**
  - With factory-installed, pre-integrated hardware, optimized Hadoop software and an ecosystem of analytical options, SGI Hadoop Starter Kits allow customers to run complex analytical applications on top of Hadoop, straight out of the box.
  - SGI Hadoop Starter Kits require little deployment effort, and ensure maximum rack density and low maintenance costs.
  - Built on top of Cloudera's standard Apache™ Hadoop™ distribution, SGI Hadoop Starter Kits provide a modular approach to vertical applications.
- **Solving the challenges of unstructured data volume, velocity and variety**
  - SGI Hadoop Starter Kits have the raw capacity to support hundreds of terabytes to petabytes of data in three densely-packed rack configurations, each with a choice of half-depth or standard-depth servers.
  - One of the special features of SGI Hadoop Starter Kits is also their characterization and performance of data integration to HDFS™.
- **Solving Hadoop optimization challenges with predictable performance**
  - SGI Hadoop Starter Kits are built for affordable price/performance, and provide linear TeraSort performance and optimal performance/watt - allowing organizations to focus on applications and application development rather than performance tuning.

- **Solving challenges in finding value in unstructured data**
  - SGI Hadoop Starter Kits contain pre-packaged analytical ISV software and demos – from vendors including Datameer, Kitenga, Pentaho and Quantum4D - which can be activated as the customer requires.
  - SGI Hadoop Starter Kits also support options for import/export, search, mine, predict, creating business models, and visualizing data for BI.

3.2 Specifications

SGI Hadoop Starter Kits are available in half-rack, single rack and multi-rack configurations based on SGI Rackable half-depth or standard-depth servers utilizing the latest Intel® Xeon® Processor E5 family. Starter kits based on servers using the Intel® Xeon® Processor E5-2600 family are available in both half-depth and standard-depth form factors; while those using the Intel® Xeon® Processor E5-2400 family are available in a half-depth form factor.

The basic capabilities of half-rack, single rack and multi-rack configurations are shown below.

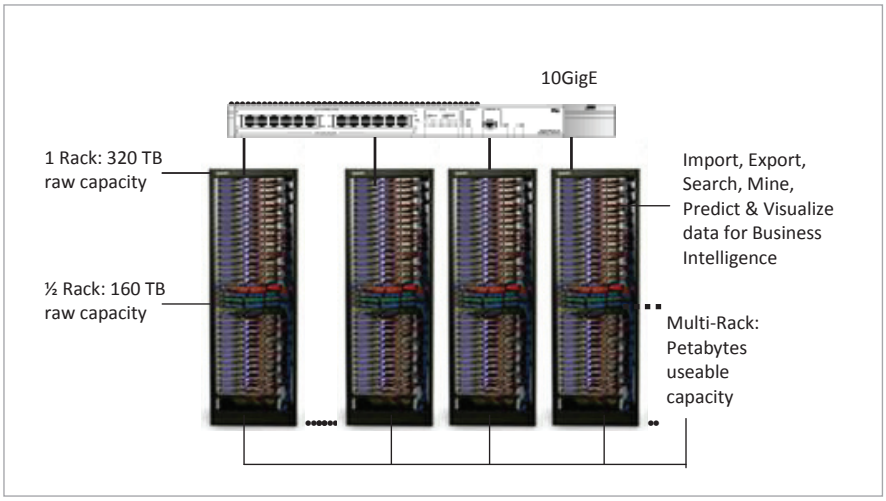

Intel® Xeon® Processor E5-2400/E5-2600-based solutions using SGI Rackable C2005-RN1 and C2005-RP1 half-depth servers include:

- C2005-RN1 servers as Hadoop NameNode / Secondary NameNode / JobTracker, each with 2x Intel® Xeon® Processor E5-2420 (1.9 GHz, six-core); 6x 8GB 1.35v 1333MHz DIMMs (64GB memory); 4x 3.5" 1TB 7200rpm SATA 6Gb/s drives in RAID configuration; 1x Dual port 10GbE NIC; and redundant power supply.
- C2005-RN1 servers as DataNode / TaskTracker, each with 2x Intel® Xeon® Processor E5-2420 (1.9 GHz, six-core); 6x 8GB 1.35v 1333MHz DIMMs (64GB memory); and 10x 2.5" 1TB 7200rpm SATA 6Gb/s drives.
- C2005-RP1 server as Application Node with 2x Intel® Xeon® Processor E5-2670 (2.6 GHz, eight-core); 16x 8GB 1.35v 1333MHz DIMMs (128GB memory); 4x 3.5" 1TB 7200rpm SAS 6Gb/s drives in RAID 5 configuration; 1x Dual port 10GbE NIC; and redundant power supply.
- Network: 2x LG-Ericsson ES-4550G 48-port GigE switches per rack; and 1x LG-Ericsson ES-5048XG 10GigE spine switch.

Intel® Xeon® Processor E5-2600-based solutions using SGI Rackable C2005-RP1 half-depth servers include:

- C2005-RP1 servers as Hadoop NameNode / Secondary NameNode / JobTracker, each with 2x Intel® Xeon® Processor E5-2630 (2.3 GHz, six-core); 8x 8GB 1.35v 1333MHz DIMMs (64GB memory); 4x 3.5" 1TB 7200rpm SATA 6Gb/s drives in RAID configuration; 1x Dual port 10GbE NIC; and redundant power supply.

- C2005-RP1 servers as DataNode / TaskTracker, each with 2x Intel® Xeon® Processor E5-2630 (2.3 GHz, six-core); 8x 8GB 1.35v 1333MHz DIMMs (64GB memory); and 10x 2.5” 1TB 7200rpm SATA 6Gb/s drives.
- C2005-RP1 server as Application Node with 2x Intel® Xeon® Processor E5-2670 (2.6 GHz, eight-core); 16x 8GB 1.35v 1333MHz DIMMs (128GB memory); 4x 3.5” 1TB 7200rpm SAS 6Gb/s drives in RAID 5 configuration; 1x Dual port 10GbE NIC; and redundant power supply.
- Network: 2x LG-Ericsson ES-4550G 48-port GigE switches per rack; and 1x LG-Ericsson ES-5048XG 10GigE spine switch.

These servers are shown below.

DataNode / TaskTracker: SGI Rackable C2005-RN1 or RP1 server with Intel® Xeon® Processor E5 Family

NameNode / Secondary NameNode / JobTracker: SGI Rackable C2005-RN1 or RP1 server with Intel® Xeon® Processor E5 Family

Application Node: SGI Rackable C2005-RN1 or RP1 server with Intel® Xeon® Processor E5 Family

Intel® Xeon® Processor E5-2600-based solutions using SGI Rackable C1110-RP6 full-depth servers (with a full-depth 1U form factor for maximum density) include:

- C1110-RP6 servers as Hadoop NameNode / Secondary NameNode / JobTracker, each with Intel® Xeon® Processor E5-2630 (2.3GHz six-core); 8x 8GB 1.35v 1333MHz DIMMs (64GB memory); 4x 3.5” 1TB 7200rpm SATA 6Gb/s drives in RAID 5 configuration; and redundant power supply.
- C1110-RP6 servers as DataNode / TaskTracker, each with Intel® Xeon® Processor E5-2630 (2.3GHz six-core); 8x 8GB 1.35v 1333MHz DIMMs (64GB memory); and 10x 2.5” 1TB 7200rpm SATA 6Gb/s drives.
- C1110-RP6 server as Application Node with Intel® Xeon® Processor E5-2670 (2.6GHz eight-core); 16x 8GB 1.35v 1333MHz DIMMs (128GB memory); 4x 3.5” 1TB 7200rpm SAS 6Gb/s drives in RAID 5 configuration; 1x Dual port 10GbE NIC; and redundant power supply.
- Network: 2x LG-Ericsson ES-4550G 48-port GigE switches per rack; and 1x LG-Ericsson ES-5048XG 10GigE spine switch.

These servers are shown below.

DataNode / TaskTracker: SGI Rackable C1110-RP6 servers with Intel® Xeon® Processor E5-2600 Series

NameNode / Secondary NameNode / JobTracker: SGI Rackable C1110-RP6 servers with Intel® Xeon® Processor E5-2600 Series

Application Node: SGI Rackable C1110-RP6 server with Intel® Xeon® Processor E5-2600 Series

3.3 Software Stack

The software stack provided with SGI Hadoop Starter Kits comprises:

- Operating system: Red Hat® Enterprise Linux® 6.2 (2.6.32-220.el6.x86_64)
- Cloudera® distribution including Apache™ Hadoop™ 3 update 2 (hadoop-0.20.2-cdh3u2)
- SGI® Management Center 1.5.0
- An ecosystem of Business Intelligence applications from ISVs including Datameer, Kitenga, Pentaho and Quantum4D
- Data ingestion and integration capabilities from fast OLTP databases into HDFS

These are shown graphically below.

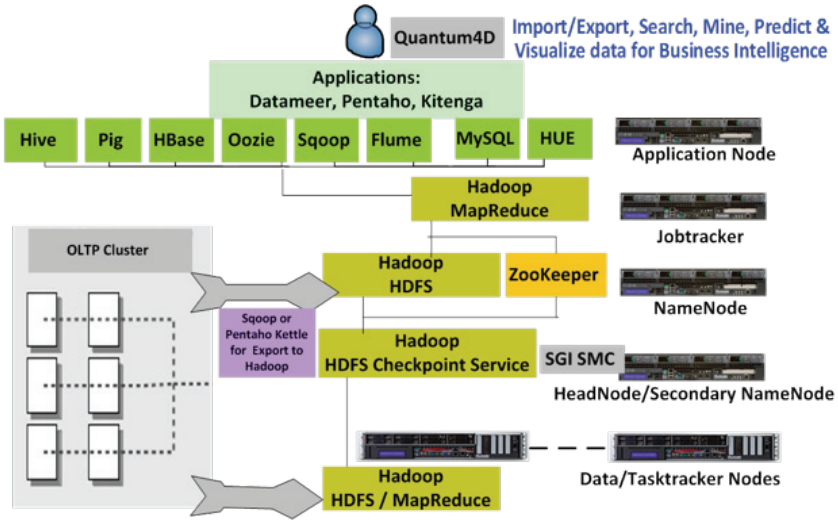

SGI Hadoop Software Stack

The key elements of the SGI Hadoop BI ecosystem and their use for BI applications are summarized in the diagram below, followed by detailed descriptions of their capabilities and benefits.

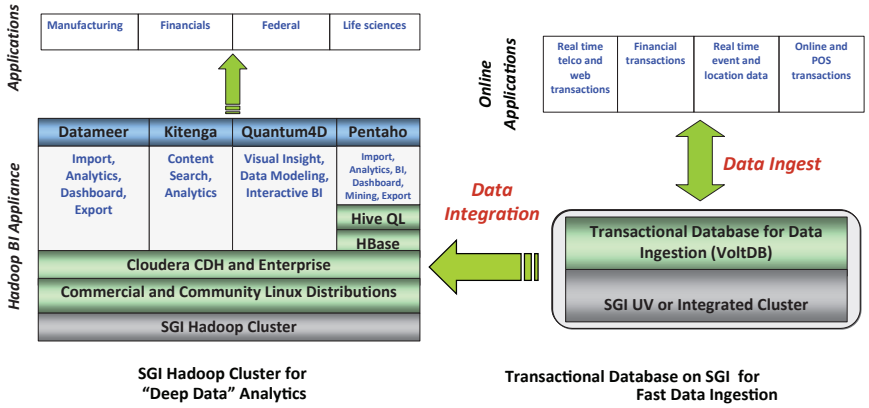

SGI Hadoop Business Intelligence Ecosystem with Data Ingestion followed by Integration to HDFS

4.0 Standard Benchmark Performance

To provide an indication of the performance of SGI Hadoop Starter Kits in real world applications, the following standard Hadoop benchmarks have been executed:

- **TeraSort** (for which SGI holds the world record) — which derives the sort time for 1TB or any other amount of data in a Hadoop cluster, and combines testing of the cluster's HDFS and MapReduce layers.
- **TestDFSIO** — a standard benchmark used to perform an I/O stress test for HDFS.
- **WordCount** — which reads text files and counts how often words occur. (The input and output are text files, each line of which contains a word and the count of how often it occurred, separated by a tab.)
- **Sort** — which uses the MapReduce framework to sort an input directory into an output directory. (The inputs and outputs must be Sequence files where the keys and values are BytesWritable.)

The benchmarks were executed on an SGI Hadoop Starter Kit based on the Intel® Xeon® Processor E5 family. References to measured performance on an SGI Rackable C2005-TY6 Hadoop cluster based on Intel® Xeon® Processor 5600 series have also been included where relevant and/or where SGI Rackable C2005-RP1 or C2005-RN1 results were not immediately available on a large 32-node configuration.

4.1 TeraSort

Results (from May 2012) show that an SGI Rackable C2005-RN1 Hadoop cluster based on Intel® Xeon® Processor E5-2400 is **similar in performance** to an SGI Rackable C2005-RP1 cluster based on Intel® Xeon® Processor E5-2600, and delivers around **6-11% higher** performance than an SGI Rackable C2005-TY6 cluster based on Intel® Xeon® Processor 5600 series. This is because the inherent distributed nature of Hadoop predominates over the compute performance of the CPUs in the cluster on which it is running.

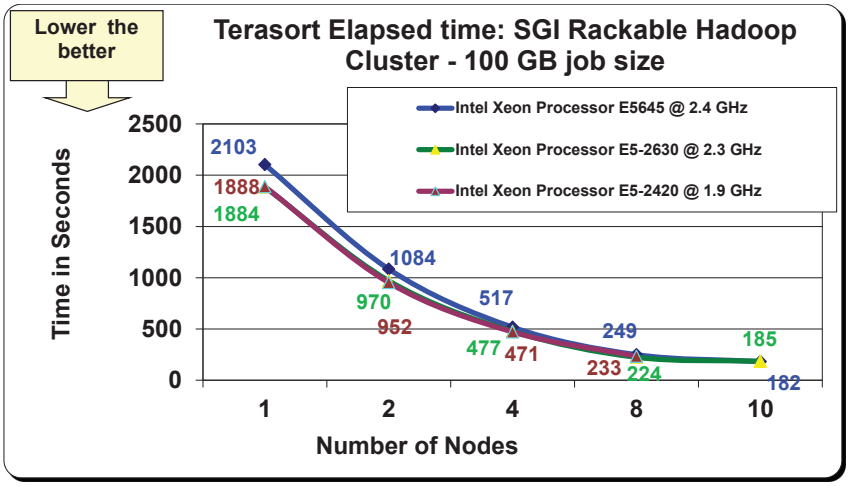

TeraSort Elapsed Time: SGI Rackable Hadoop clusters on Intel® Xeon® Processor E5 Family and 5600 series

With the clusters running Cloudera's distribution including Apache Hadoop (CDH3u2), TeraSort on SGI Rackable C2005-RN1 and C2005-RP1 Hadoop clusters scales almost **linearly**, and the scalability is similar to that based on the Intel® Xeon® Processor 5600 series.

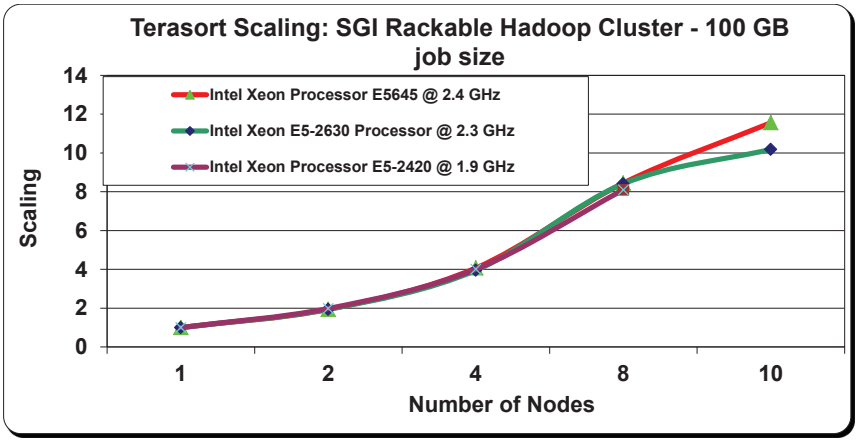

TeraSort Scaling: SGI Rackable Hadoop clusters on Intel® Xeon® Processor E5 Family and 5600 series

An SGI Rackable C2005-RN1 or C2005-RP1 Hadoop cluster is **~3x and ~2x faster** than an Oracle Sun X2270 cluster on a single node and 8 nodes respectively.

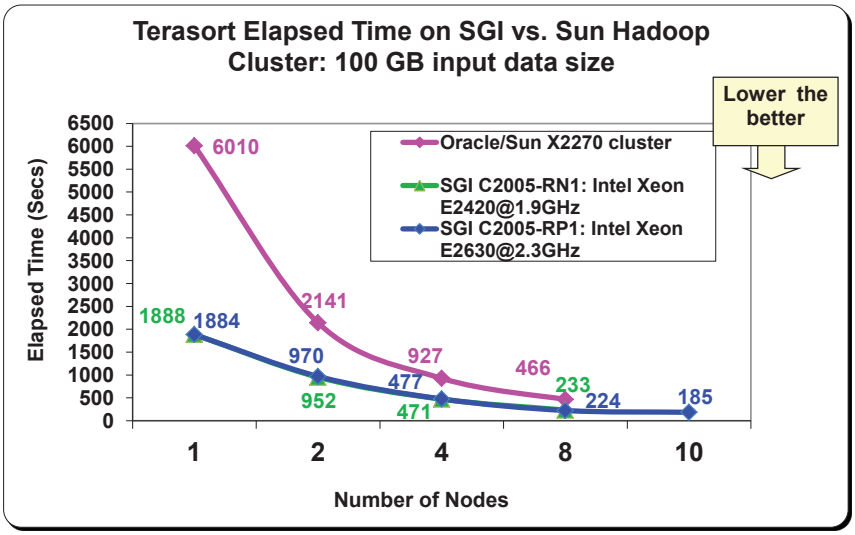

TeraSort Elapsed Times: SGI Rackable Hadoop cluster on Intel® Xeon® Processor E5 Family vs. Oracle Sun X2272 M2 cluster[1]

It is also important to note that a single SGI Rackable server acting as a Hadoop DataNode and configured with 10 HDDs is ~2x faster than with 4 HDDs, meaning that in general, a balance of cores to drives is important for Hadoop DataNodes.

[1] http://sun.systemnews.com/articles/152/1/server/23549

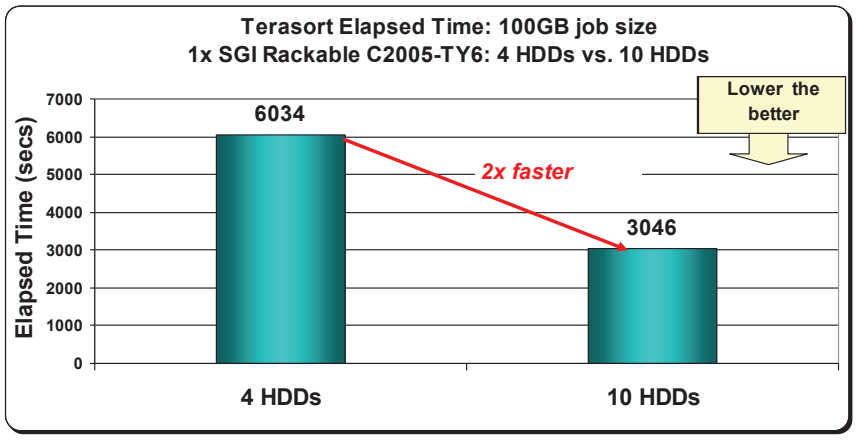

TeraSort Elapsed Time: A single SGI Rackable server as a Hadoop DataNode, configured with 4 HDDs vs. 10 HDDs

4.2 TestDFSIO

On a 10-node SGI Rackable Hadoop cluster based on the Intel® Xeon® Processor E5 family, in TestDFSIO **read tests**, uniform scaling occurs as all the read requests benefit from parallelism with more nodes in the cluster. In TestDFSIO **write tests**, the number of maps created is equal to the number of files specified as a parameter. For a run with 1,000 files, each 1GB in size, 1,000 map jobs (each of which will create a file 1GB in size) are created by the MapReduce engine.

- With two DataNodes, each map job or task will have to create 500MB locally and 500MB remotely. This means there is a network overhead where the map task will have to write half the file size over the network, that data will have to be written on the disks of the remote DataNode, and the map tasks will have to wait until the remote writes have successfully completed.
- As more nodes are added, the parallelism of map tasks gains over the amount of work per node, so the elapsed time reduces substantially.

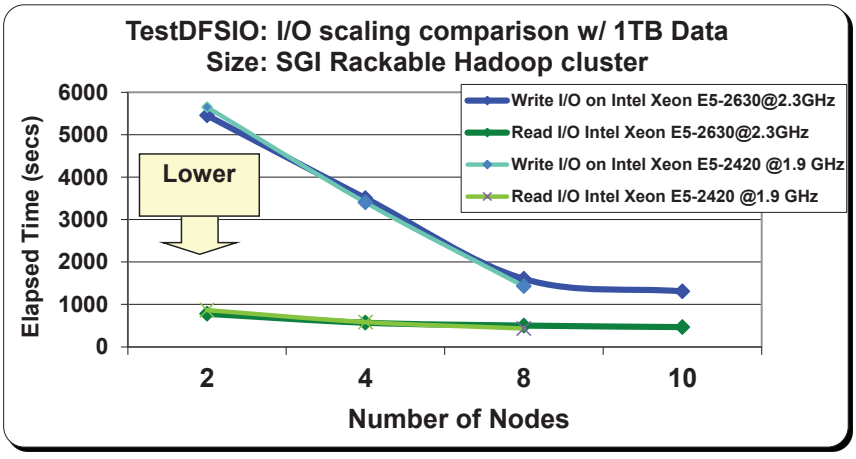

TestDFSIO Scaling on a 10-node SGI Rackable Hadoop cluster based on Intel® Xeon® Processor E5 Family

- In terms of the **best read performance with multiple files**, the sweet spot was with 1,000 files of 1GB each.
- For the **best write performance with multiple files**, the sweet spot was with 100 files of 10GB each.

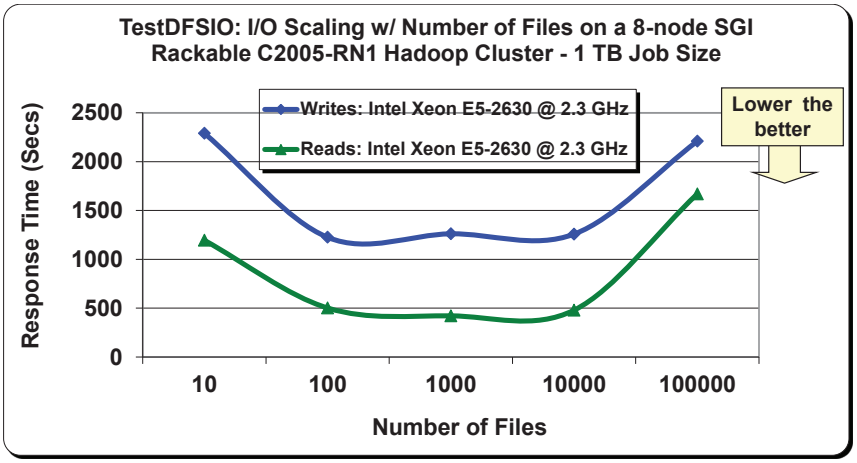

TestDFSIO Scaling with Number of Files

4.3 WordCount

WordCount with a 1TB job size scales **linearly** across a 10-node SGI Rackable C2005-RP1 or C2005-RN1 Hadoop cluster, with an enhanced scalability between 4 to 8 nodes.

WordCount has a **14x** capability to count occurrences of a word from text files, while using **4x** nodes in an SGI Rackable C2005-RN1 cluster. The 4 to 8-node enhanced scalability is a contributory factor.

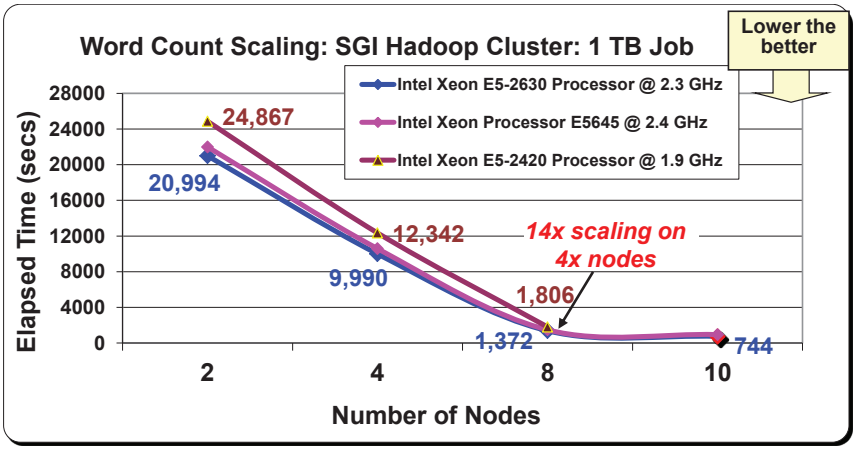

WordCount Scaling on SGI Rackable Hadoop clusters on Intel® Xeon® Processor E5 Family and 5600 series

4.4 Sort

On a 32-node SGI Rackable C2005-TY6 Hadoop cluster (the behavior of which is similar to a C2005-RP1 cluster), Sort time increases with an increase in sort size. The Sort benchmark creates input data proportional to the number of nodes in the cluster, so that for each additional node, 10GB of additional data is created to be sorted. The increased network traffic, and hence the traffic on the switch, become predominant with more nodes in the cluster.

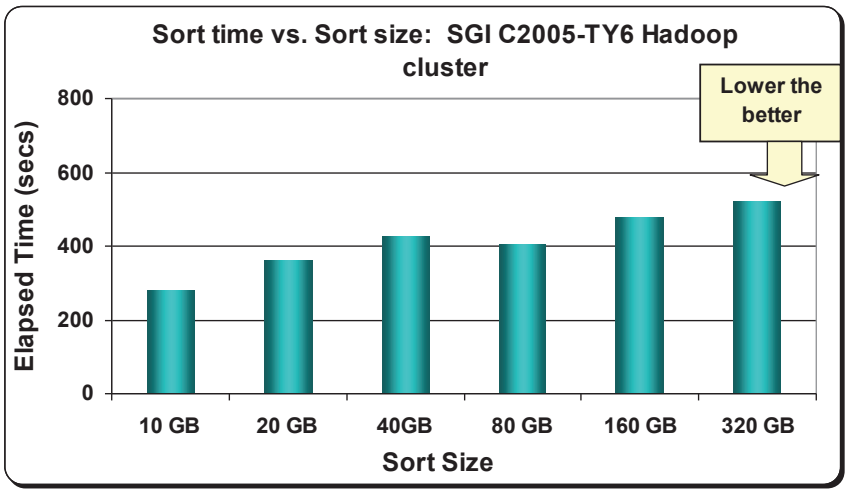

Sort Time vs. Sort Sizes on SGI Rackable C2005-TY6 Hadoop cluster

5.0 Price/Performance and Performance/Watt

The price/performance and performance/watt metrics of Hadoop clusters are defined in terms of **$/TPM** and **TPM/Watt**, where TPM refers to TeraSort-throughput per minute and is based on TeraSort throughput performance. This is then compared to the power consumption of the Hadoop cluster, measured in watts.

The price/performance (in terms of $/TPM) of an 8-node SGI Rackable C2005-RN1 Hadoop cluster based on Intel® Xeon® Processor E5-2400 series is **22% lower** (and therefore **better**) than an SGI Rackable C2005-TY6 Hadoop cluster based on Intel® Xeon® Processor 5600 series.

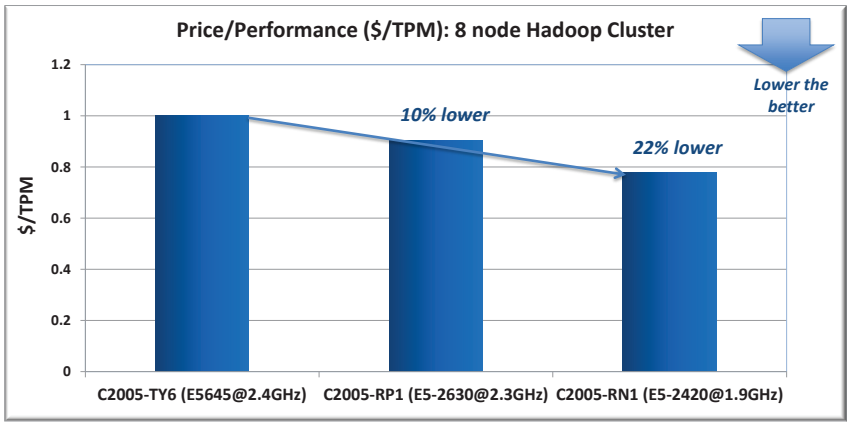

Price/Performance ($/TPM): SGI Rackable Hadoop clusters on Intel® Xeon® Processor E5 Family and 5600 series

Similarly, the performance/watt (in terms of TPM/Watt) of an SGI Rackable C2005-RN1 Hadoop cluster based on Intel® Xeon® Processor E5-2400 series is **27% higher** (and therefore **better**) than an SGI Rackable C2005-TY6 Hadoop cluster based on Intel® Xeon® Processor 5600 series.

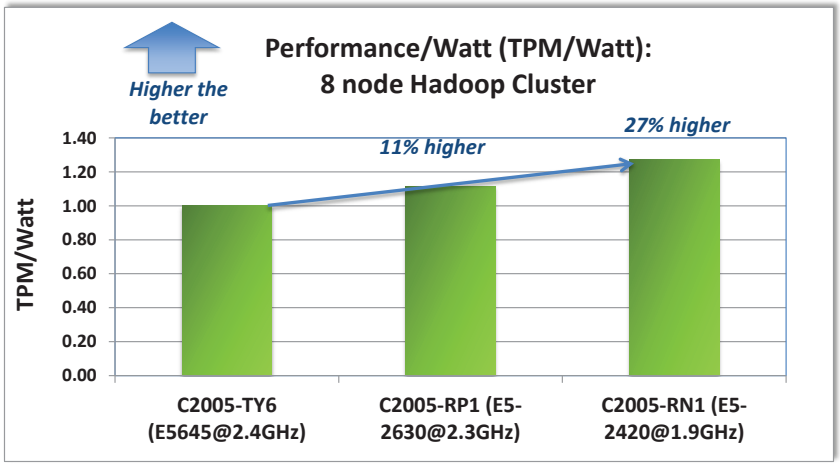

Performance/Watt (TPM/Watt): SGI Rackable Hadoop clusters on Intel® Xeon® Processor E5 Family and 5600 series

A Hadoop cluster based on the Intel® Xeon® Processor E5-2400 series is therefore a better choice in terms of both price/performance and performance/watt than one based on the Intel® Xeon® Processor 5600 series.

6.0 Why SGI Hadoop Starter Kits?

To summarize the benefits of SGI Hadoop Starter Kits in all of the above applications, they provide a pre-optimized, integrated platform for customers who have either already deployed Hadoop in their data center or are considering such a deployment. Delivered ready-to-run, and so alleviating any concerns around Hadoop adoption, SGI Hadoop Starter Kits:

- Combine linear TeraSort scalability with affordable price/performance and optimal performance/watt.
- Offer pre-integrated hardware, optimized software, data integration capabilities and a best-in-class BI solution, all out of the box.
- Leverage a factory-installed, optimized Apache Hadoop distribution, with an advanced ecosystem enabling customers to run complex analytical applications.
- Support fast processing and deep analytics against hundreds of terabytes to petabytes of data, enabling customers to import/export, search, mine, predict, create business models, visualize data for BI, and benefit from measured data ingestion and integration capabilities into HDFS.
- Meet the needs of customers looking for minimum deployment efforts, low maintenance costs, maximum rack density, predictable performance and optimized performance/watt – and so enable them to focus on application development instead of performance tuning.
- Provide an ideal solution for situations where extreme scalability, energy efficiency for server consolidation, and performance for Big Data analytics are the primary goals.

The table below summarizes the market-leading solution and value proposition provided by SGI Hadoop Starter Kits vs. alternative solutions.

Features/Advantages Comparisons	SGI Hadoop BI Reference Implementation	Integrated NoSQL/Hadoop Solutions					
		EMC Green Plum	HP Vertica	Oracle Big Data Appliance	Dell Hadoop Solution	Netapp Hadoop Solution	IBM Netezza + Hadoop
Appliance-based proprietary and closed architecture	✗	✓	✓	✓	✗	✓	✓
Out of the box factory installed and configured Hadoop stack	✓	✗	✗	✗	✗	✗	✗
Drag-and-drop analytics and indexing pipelines	✓	✓	✗	✓	✗	✗	✗
Open source core enabled customization and rapid innovation	✓	✓	✗	✗	✓	✓	✗
Hadoop enabled for scalability	✓	✓	✓	✓	✓	✓	✓
Hive enabled BI	✓	✓	✓	✓	✓	✓	✓
Spreadsheet-based analytics	✓	✓	✗	✗	✗	✗	✓
Native visualization of search results	✓	✗	✗	✓	✗	✗	✗
3D visualization of Big Data results	✓	✗	✗	✗	✗	✗	✗
Natural Language Processing for enhanced metadata	✓	✗	✗	✗	✗	✗	✗
Multiple Language/Character sets	✓	✗	✗	✗	✗	✗	✗
Document and content security	✓	✗	✗	✓	✗	✗	✓
Ingestion of multiple document formats	✓	✓	✗	✓	✓	✓	✓

SGI Hadoop Reference Implementation Feature Comparison

Sources	Url
Green plum Hadoop	http://www.greenplum.com/products/greenplum-hd
HP Vertica + Hadoop	http://www.vertica.com/the-analytics-platform/native-bi-etl-and-hadoop-mapreduce-integration/
Oracle Big Data Appliance	http://www.oracle.com/us/corporate/features/feature-obda-498724.html
Dell Hadoop Cluster	http://i.dell.com/sites/content/business/solutions/whitepapers/en/Documents/hadoop-enterprise.pdf
NetApp Hadoop Solution	http://www.netapp.com/us/solutions/infrastructure/hadoop.html; http://media.netapp.com/documents/ds-3237.pdf
IBM Netezza + BigInsights Hadoop	http://thinking.netezza.com/video/using-hadoop-ibm-netezza-appliance; http://thinking.netezza.com/blog/hadoop-netezza-synergy-data-analytics-part-2

In addition, because SGI Rackable half-depth servers combine affordable price/performance, super-linear scalability, disk capacity, rack density, and power and cooling efficiency, more of these servers can be added into a rack or scaled out to multiple racks to minimize energy usage and maintain optimal performance and capacity required for Hadoop deployments. SGI Hadoop clusters can therefore deal with the challenges related to the volume, velocity, variety and value of unstructured data for fast transactions and deep analytics; while their flexibility of choice in terms of memory, compute, capacity, I/O, network latency and low power helps them to:

- Optimize task assignment and expedite MapReduce tasks across distributed nodes with efficient parallelism
- Maintain file system metadata operations for HDFS
- Store large HDFS files and handle HDFS read/write requests
- Co-locate I/O with TaskTrackers for optimal data locality
- Achieve optimal performance-per-watt across all load levels

7.0 Why SGI for Hadoop?

Throughout the Big Data 'wave' of data ingest, Hadoop analytics, visualization, archiving and storage, SGI is uniquely positioned to offer integrated BI solutions based on proven, best-of-breed technologies delivering the highest speed and scale. As discussed above, SGI has been supplying SGI Rackable clusters for Hadoop since the technology's first commercial deployments - including the largest Hadoop clusters (up to 4,000 nodes) and installations (up to 40,000 nodes) in existence today – and the value proposition of the SGI Rackable product line on which SGI Hadoop clusters are based is truly compelling:

- Industry-leading configuration flexibility enables customers to meet their exact needs for any given application.
- With up to 2,688 cores per 42U cabinet, SGI Rackable servers deliver high density, allowing data centers to conserve precious space.
- With power increasingly at a premium, Eco-Logical™ technology ensures SGI Rackable servers typically consume significantly less power than competitive offerings, dropping Apex cost per server while allowing a larger number of servers to fit into a facility's power budget.
- Whether at the application level or in the server hardware, SGI Rackable configurations deliver high reliability, availability, serviceability and manageability.

To find out more, please visit www.sgi.com/products/hadoop.

Global Sales and Support: sgi.com/global