WHITE PAPER

SGI® Hadoop™ 10GigE Starter Kits for Big Data Analytics

October, 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:

- **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, 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 Starter Kit configurations are shown below.

Intel® Xeon® Processor E5-2600-based solutions using SGI ISS3012-RP6 half-depth servers include:

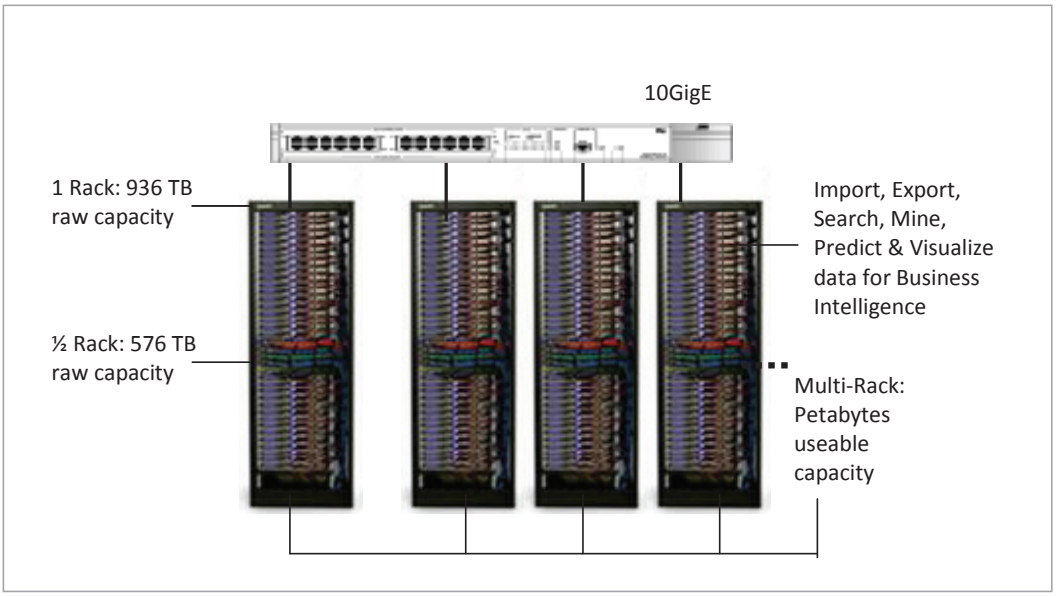

- C2005-RP1 servers as Hadoop NameNode / Secondary NameNode / JobTracker, each with 2x Intel® Xeon® Processor E5-2630 (2.3 GHz, six-core); 8x 8GB 1.35v 1333MHz DIMMs (64GB memory); 4x 2.5" 1TB 7200rpm SATA 6Gb/s drives in RAID 10 configuration; 1x Dual port 10GbE NIC; and a redundant power supply.
- ISS3012-RP6 servers as DataNodes / TaskTrackers, each with 2x Intel® Xeon® Processor E5-2630 (2.3 GHz, six-core); 8x 8GB 1.35v 1333MHz DIMMs (64GB memory); 12x 3.5" 1TB 7200rpm SATA 6Gb/s drives; and 1x Dual port 10GbE NIC.
- C2005-RP1 server as Application Node with 2x Intel® Xeon® Processor E5-2680 (2.7 GHz, eight-core); 16x 8GB 1.35v 1333MHz DIMMs (128GB memory); 4x 2.5" 1TB 7200rpm SAS 6Gb/s drives in RAID 10 configuration; 1x Dual port 10GbE NIC; and a redundant power supply.
- Network: 2x LG-Ericsson ES-5048XG 48-port 10GigE switches per rack; and 1x LG-Ericsson ES-4550G GigE switch for the SGI management network.

These servers are shown below.

DataNodes & TaskTracker: SGI ISS3012-RP6 servers with Intel Xeon Processor E5 Family

NameNode, Secondary NameNode, JobTracker: SGI Rackable C2005-RP1 servers with Intel Xeon Processor E5 Family

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

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 4 (hadoop-0.20.2-cdh3u4)
- SGI® Management Center 1.6.0
- An ecosystem of Business Intelligence applications from ISVs including Datameer, Kitenga, Pentaho and Quantum4D and a number that can run out-of-the-box on the Hadoop Starter Kit
- Data ingestion and integration capabilities from fast OLTP databases into HDFS

These are shown graphically below.

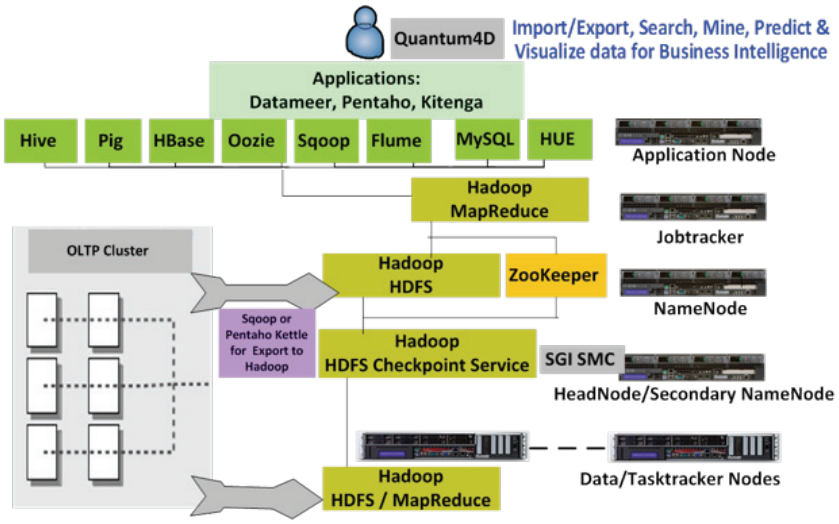

SGI Hadoop Software Stack

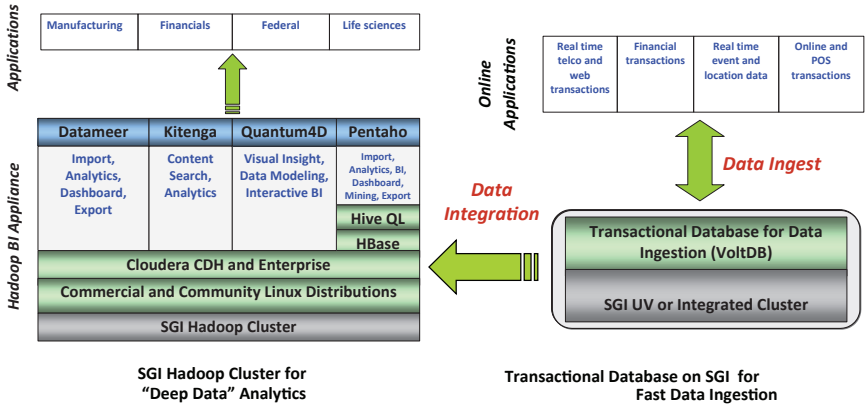

SGI Hadoop Business Intelligence Ecosystem with Data Ingestion followed by Integration to HDFS

4.0 Standard Benchmark Performance

To provide an indication of the performance of SGI Hadoop 10GigE Starter Kits in real world applications, the **Terasort** benchmark (see http://sortbenchmarkg.org) was run. The Terasort benchmark 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. Some of the results were compared to a bonded **GigE**-based, 16-node Hadoop cluster having the following configuration: 16x C2005-RP1 servers each with 2x Intel® Xeon® Processor E5-2667 (2.93 GHz, six-core) – clock set to 2.3 GHz; 8x 8GB 1.35v 1333MHz DIMMs (64GB memory); 10x 1TB 7200rpm SATA 6Gb/s drives Seagate HDD or 10x 300 GB 10K RPM HDD; GigE with 3-port bond; RHEL 6.2 OS; Cloudera's distribution including Apache Hadoop (CDH3u4).

4.1 TeraSort Benchmark

Results (from September 2012) show that an SGI ISS3012-RP6 10GigE Hadoop cluster is **10-20%** faster in performance for 8 and 16 nodes, compared to a C2005-RP1 GigE Hadoop cluster (3-port bond, both based on Intel® Xeon® Processor E5 series.

The 12-drive configuration in the ISS3012-RP6 is the primary contributing factor toward the small performance improvement over the 10-drive C2005-RP1 configuration. Both clusters were running Cloudera's distribution including Apache Hadoop (CDH3u4). A single node with 12 HDDs with no networking involvement is **25%** faster than one with 10 HDD, with the 1:1 ratio of cores to HDD being most important for the better performance.

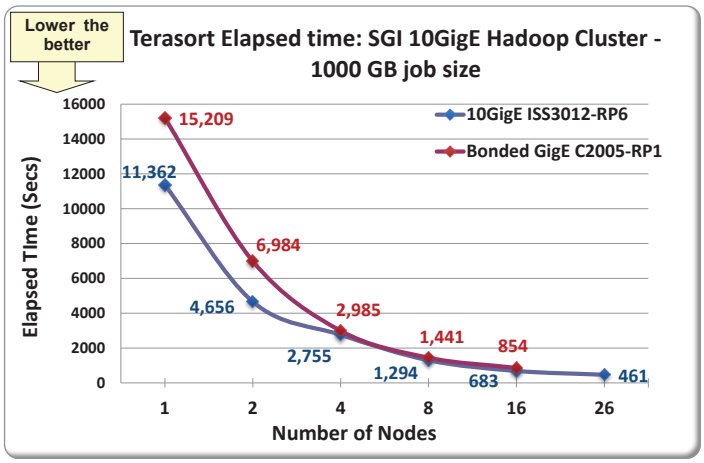

TeraSort Elapsed Time: SGI Hadoop Starter Kits on Intel® Xeon® Processor E5 Family

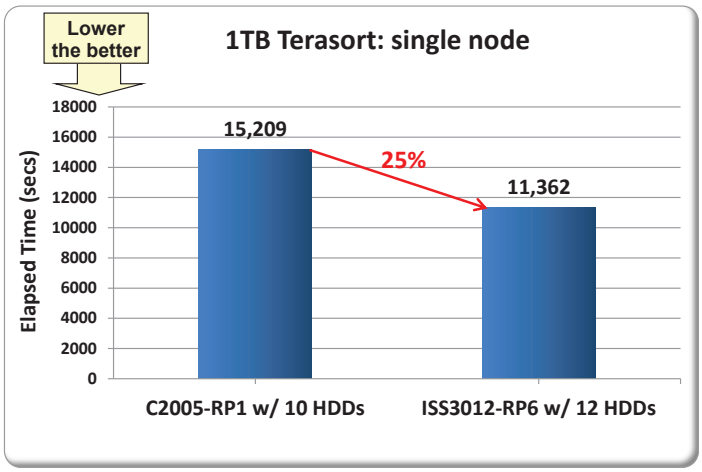

TeraSort Elapsed Times: A single SGI Hadoop data node, configured with 12 HDDs vs. 10 HDDs

Throughput for a Terasort@1TB on a 26-node SGI Starter Kit based upon ISS3012-RP6, 10GigE, and Intel® Xeon® E5 Family is **2.2 GB/s** and shows **perfect linear scaling**.

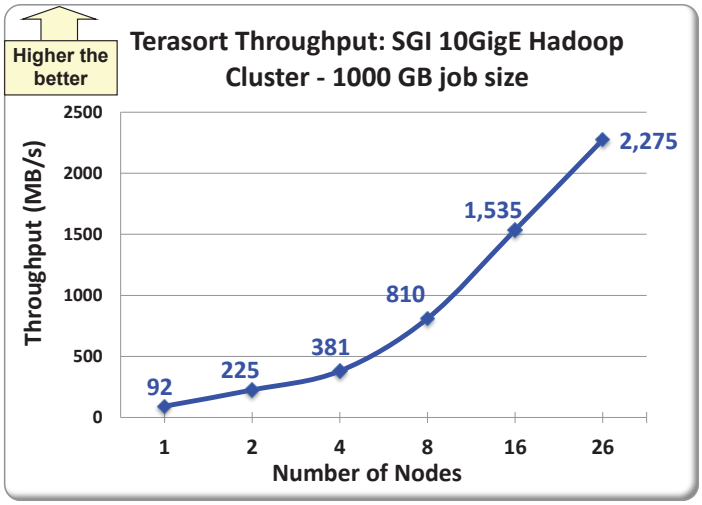

Terasort Throughput: SGI ISS3012-RP6 Hadoop Starter Kit on Intel® Xeon® Processor E5 Family

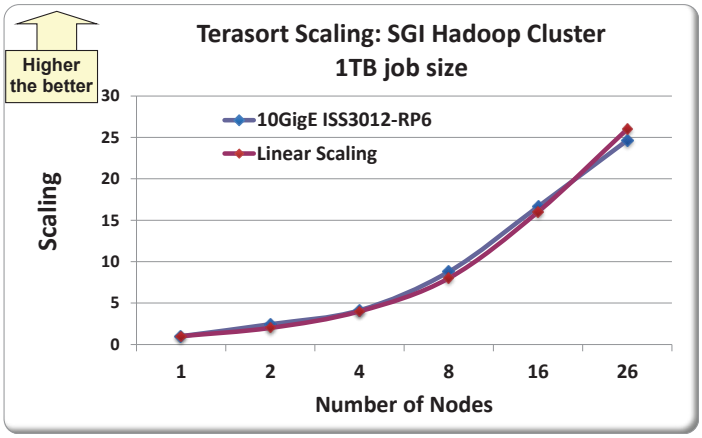

Terasort Scaling: SGI ISS3012-RP6 Hadoop Starter Kit on Intel® Xeon® Processor E5 Family

4.2 Price/Performance

The price/performance of a 16-node ISS3012-RP6 Hadoop 10GigE Starter Kit with 12 drives per node was compared to a C2005-RP1 Hadoop GigE Starter Kit with 10 drives per node. Price/performance is defined as $/TPM where TPM refers to Terasort-Throughput per minute @1TB jobs size on the Hadoop cluster.

The results show that the C2005-RP1 GigE Starter Kit with bonded GigE and 10 drives per node is 22% better in price/performance compared to the ISS3012 10GigE Starter Kit. The cost of the 10GigE infrastructure is the primary contributing factor to the higher cost of the ISS3012-RP6 Starter Kit. Both clusters were based on Intel® Xeon® C5-2630 processors and running Cloudera's distribution including Apache Hadoop (CDH3u4).

The conclusion is that 10GigE should be used when the workload needs the 10GigE network bandwidth.

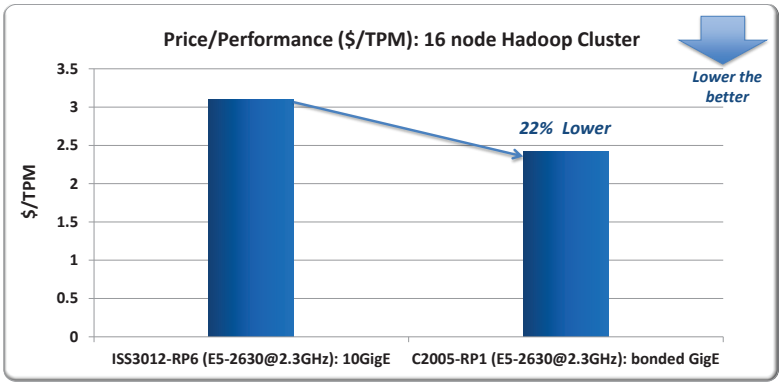

Price/Performance ($/TPM): SGI Hadoop Starter Kits on Intel Xeon Processor E5 Family Bonded GigE, 10 drives vs. 10GigE, 12 drives

5.0 Case Studies

5.1 Job Sizes

Two job sizes at 100GB and 1TB have been run on a 26-node ISS3012-RP6 10GigE Starter Kit to determine the variation of scalability of Terasort throughput. Results show that Terasort throughput with 1TB job size scales linearly compared to a 100GB job size.

A larger job size on a high capacity, high performance cluster allows for a denser distribution of the data set on each node with more work and lower communication performed across the nodes in reduce phase, thus resulting in better scaling.

A 100GB job size requires sparse distribution of data sets across multiple nodes resulting in more communication and sub-linear scaling.

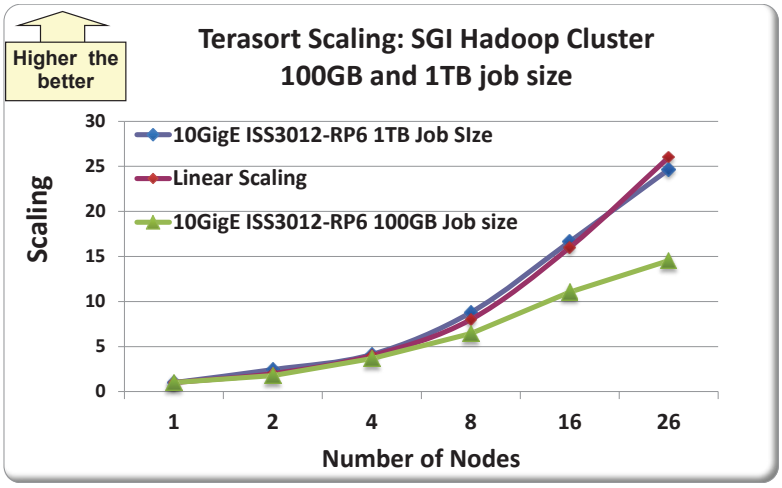

Terasort scaling: 100GB vs. 1TB job size

5.2 Network and CPU Usage

During the time interval of 1TB Terasort on the 26-node ISS3012-RP6 10GigE Hadoop Starter Kit:

Average Receive (kB/s): 172442.85 kB/s
Average Transfer (kB/s): 181429.98 kB/s
Max Bandwidth Used (kB/s): 537070.10 kB/s
Average CPU Usage (%): 66.193 %
Average IO Wait (%): 17.030 %
Average Disk Read (MB/s): 222.041 MB/s
Average Disk Write (MB/s): 307.237 MB/s
Max CPU Usage (%): 99.500%
Max IO Wait (%): 53.850 %
Max Disk Usage (MB/s): 1282.440 MB/s.

During the initial time interval of a Terasort run, for the Map phase, the operations are more CPU-intensive whereas during the latter part of the time interval, , for the Reduce phase, the operations are more network-intensive.

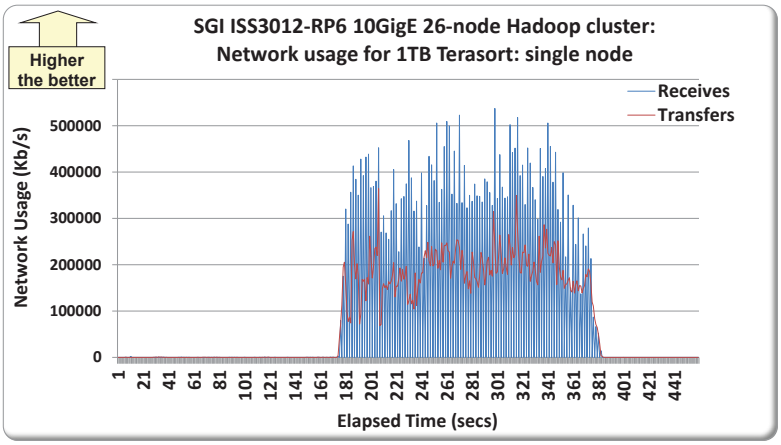

ISS3012-RP6 26-node 10GigE Starter Kit: Network usage by a node during 1TB sort time interval

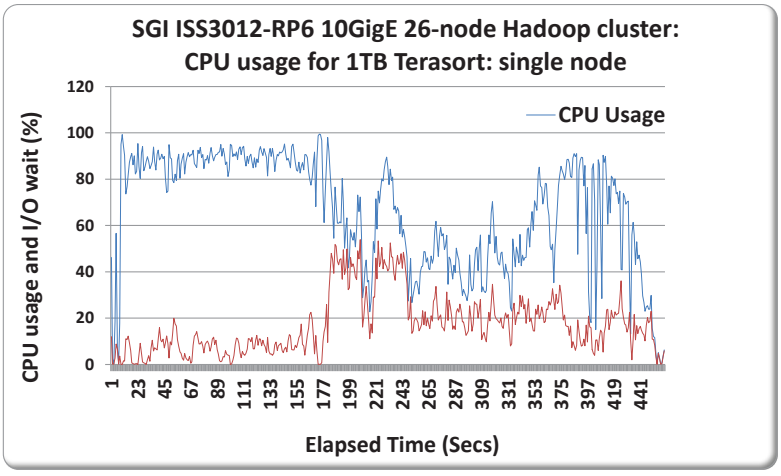

ISS3012-RP6 26-node 10GigE Starter Kit: CPU usage by a node during 1TB sort time interval

A comparative case study of network and CPU usage was done for ISS3012-RP6 and C2005-RP1 Starter Kits.

Results show that during reduce phase of 1TB Terasort on a 16-node ISS3012-RP6 10GigE Hadoop Starter Kit:

Average Receive: 237273.27 kB/s
Average Transfer: 227319.15 kB/s
Max Bandwidth Used: 854348.84 kB/s

and that on a 16-node C2005-RP1 bonded GigE Hadoop Starter Kit:

Average Receive: 290109.33 kB/s
Average Transfer: 319695.49 kB/s
Max Bandwidth Used: 609607.37 kB/s

The ISS3012-RP6 Starter Kit is leveraging the 10GigE network with a higher network bandwidth during the reduce phase, unlike the C2005-RP1 GigE Starter Kits, which is limited by the network.

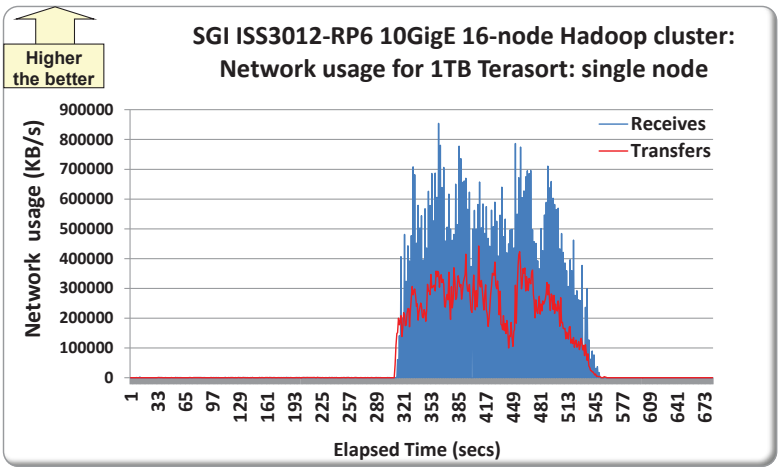

ISS3012-RP6 16-node 10GigE Starter Kit: Network usage by a node during 1TB sort

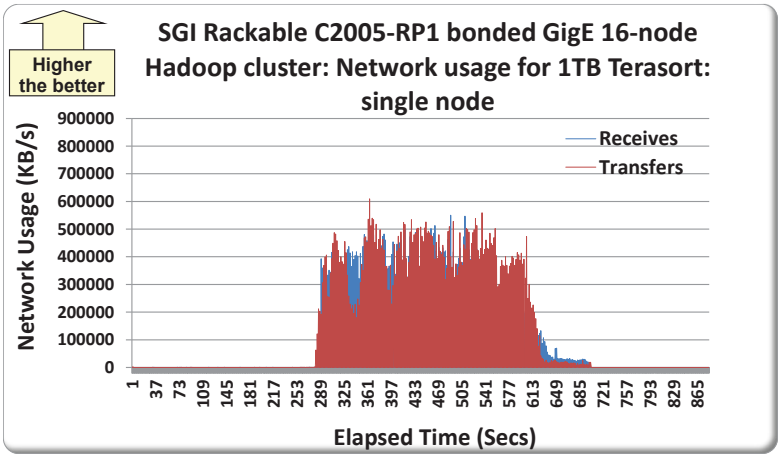

C2005-RP1 16-node GigE Starter Kit: Network usage by a node during 1TB sort

During a 1TB Terasort on a 16-node ISS3012-RP6 10GigE Hadoop Starter Kit, the CPU stats on a single node are as follows:

Average CPU Usage (%): 73.300 %
Average IO Wait (%): 23.060 %
Max CPU Usage (%): 99.420%
Max IO Wait (%): 72.890 %

During a 1TB Terasort on a 16-node C2005-RP1 bonded GigE Hadoop Starter Kit, the CPU stats on a single node are as follows:

Average CPU Usage (%): 67.649 %
Average IO Wait (%): 21.415 %
Max CPU Usage (%): 99.960%
Max IO Wait (%): 61.720 %

Both ISS3012-RP6 and C2005-RP1 Hadoop clusters show a sustained near 100% CPU usage without idling, ISS3012-RP6 slightly being better in CPU usage.

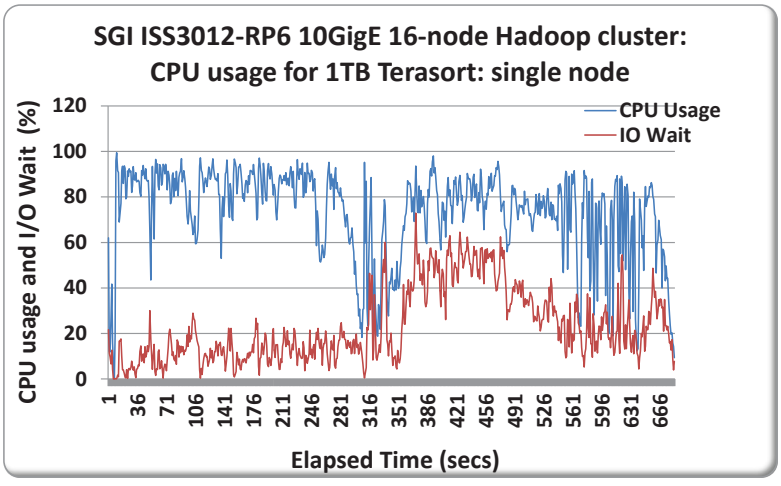

ISS3012-RP6 16-node 10GigE Starter Kit: CPU usage by a node during 1TB sort

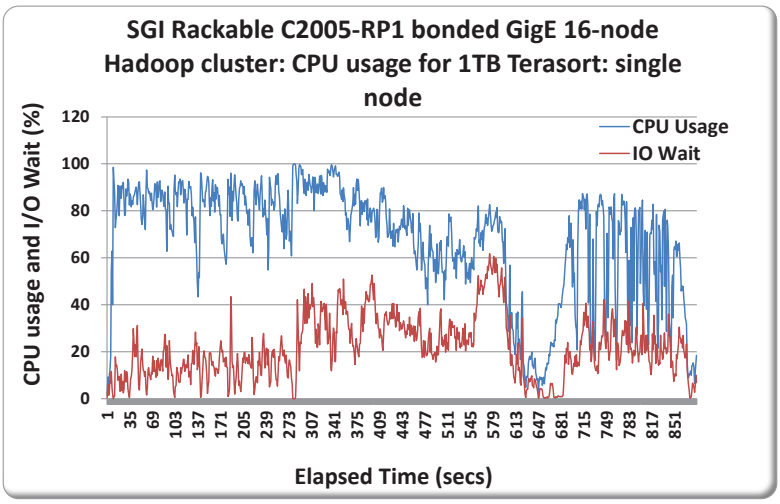

C2005-RP1 16-node GigE Starter Kit: CPU usage by a node during 1TB sort

5.3 Disk Usage

A comparative case study of disk usage has been done for ISS3012-RP6 and C2005-RP1 Starter Kits, configured with 12 drives and 10 drives per node, respectively.

During a 1TB Terasort on a 16-node ISS3012-RP6 10GigE Hadoop Starter Kit, the disk stats on a single node are as follows:

Average Disk Read(MB/s): 254.123 MB/s
Average Disk Write(MB/s): 321.377 MB/s
Max Disk Usage (MB/s): 1155.870 MB/s

During a 1TB Terasort on a 16-node C2005-RP1 bonded GigE Hadoop Starter Kit, the disk stats on a single node are as follows:

Average Disk Read(MB/s): 210.878 MB/s
Average Disk Write(MB/s): 266.104 MB/s
Max Disk Usage (MB/s): 898.660 MB/s

The ISS3012-RP6 Starter Kit with 12 drives per node is leveraging the 2 additional disks per node. Disk usage is **28% higher** than the C2005-RP1 GigE Starter Kit which has 10 drives per node, thus helping the ISS3012 cluster to complete the map/reduce tasks faster. The additional drives in ISS3012-RP6 Starter Kit are contributing to the small performance improvement over C2005-RP1 cluster (Figures 16 and 17).

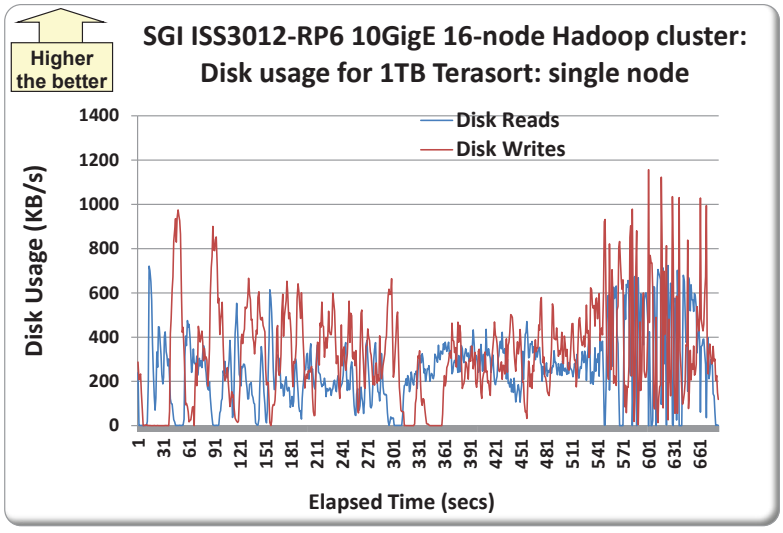

ISS3012 16-node 10GigE Starter Kit: Disk usage by a node during 1TB sort

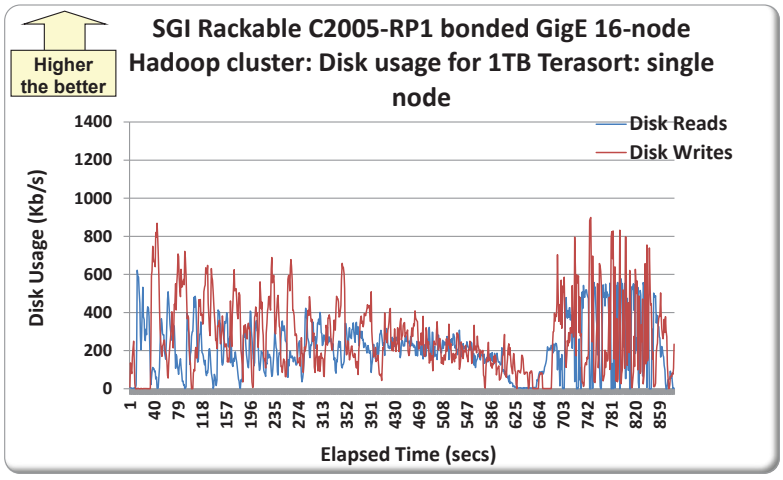

C2005-RP1 16-node GigE Starter Kit: Disk usage by a node during 1TB sort

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