WHITE PAPER

SGI® DataRaptor™ with MarkLogic® Database

October, 2012

TABLE OF CONTENTS

1.0 Eat Big Data for Lunch

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. Many market segments can benefit from optimal solutions for managing and analyzing Big Data 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 Analysis in Real Time

Big Data can bring terrific insights to enterprises and government agencies. Many popular methods for analysis, though, work in a batch mode, whereas often the insights are more valuable when they are revealed in real time. Some of the most critical Big Data applications in operation today use the MarkLogic enterprise NoSQL database and SGI and MarkLogic have created a combined solution to expand the world's access to Big Data capabilities. The combined solution is specially tuned by SGI and MarkLogic engineers working together and arrives at your site, ready to be plugged in and providing answers right away.

Many Big Data needs can be addressed in a batch mode, but the solution becomes better when the failing manufacturing process is improved as soon as possible, or the fraud is stopped before the loss occurs. Also, using current methods, many organizations need months to develop analysis tools to harness the Big Data in their midst. The following are examples of applications where Big Data analysis in real time has big pay-offs:

- Manufacturers need to optimize process yields based on sensor data
- The intelligence community needs to coordinate diverse sources of intelligence data to predict events
- In genomics, sequence data needs to be mapped to reference genomes
- Media firms want to bring the right content to users at the right time
- Financial services firms need to conduct market analysis, as well as fraud analysis, risk analysis and compliance processing
- Retailers need to perform customer buying analysis and inventory optimization

3.0 Traditional RDBMS vs. the MarkLogic Database

While one maintains a traditional RDBMS for processing and analyzing structured data, dealing with the volume, velocity and variety of unstructured and semi-structured information becomes a challenge. The rigidity of the relational model, in which all data is held in tables with a fixed structure of rows and columns, has increasingly been seen as a limitation when handling information that is richer or more varied in structure.

One solution is NoSQL technology comprised of XML and object databases, for example. An XML database is a data persistence software system that allows data to be stored in XML format. This data can then be queried, exported and serialized into the desired format. XML databases are usually associated with document-oriented databases designed for storing, retrieving, editing and managing documents, semi-structured data and information. MarkLogic is one such XML database. Unlike many other NoSQL databases, MarkLogic is an ACID-compliant, W3C-standardized XML database. It is especially useful for loading, searching, unifying diverse content formats for real time access with substantial efficiency in the delivery of content. MarkLogic also combines a database with a search engine and an applications platform in one package, versus other products that include only the database.

4.0 SGI DataRaptor: SGI and MarkLogic Together

Even though NoSQL technology is now becoming popular for addressing such Big Data and information needs, challenges like sizing, configuring and optimizing the NoSQL databases to meet end-user requirements is a big conceptual hurdle. Next, applications need to be written to scale with this content processing framework. The other challenge lies in applying the knowledge of search and analytics in the structured space to the unstructured domain. Lastly, users need to ingest, search and analyze high velocity content in real time. Thus, an integrated solution for speed and scale is required to address the different phases like ingestion, search, retrieval and insight for a seamless information flow across an enterprise.

To address these needs and more, SGI and MarkLogic have partnered to deliver solutions that reduce time to insight and remove the risk of harnessing Big Data. SGI brings its heritage in high-performance computing, producing hardware and software optimized for managing large data ingestion, split-second computation and real-time reaction to a myriad of inputs. The SGI and MarkLogic solution is based upon the Intel® Xeon® E5-2600 processor and includes:

- Hardware and software in a complete, factory-integrated package
- A choice of either performance-optimized or capacity-optimized solutions
- Support offered by SGI for the complete solution including MarkLogic software, worldwide

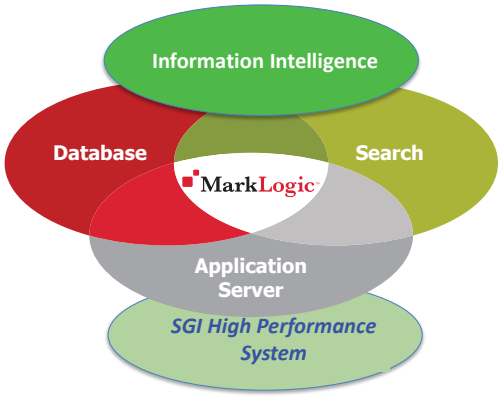

5.0 SGI DataRaptor

The core to the SGI DataRaptor is the Rackable™ ISS3112 or ISS3124 storage server with similar capabilities but different disk sub-chassis for different size disks.

- "RP2" two-socket motherboard with the latest Intel® Xeon® E5-2600 microprocessor
- 16 cores of computing power per server
- 384 GB of main memory, up to 1600 MHz in speed per server
- Six PCIe gen 3 x 8 slots, full-height, or two PCIe gen 3 x 8 slots, full-height and two PCIe gen3 x 16 slots, double-width
- Up to twelve 3.5 inch disks or twenty-four 2.5 inch disks, SAS, SATA or SSD

Database Nodes (with high performance drives): SGI ISS3124-RP2 servers with Intel Xeon Processor E5 Family

Database Nodes (High Capacity drives): SGI ISS3112-RP2 servers with Intel Xeon Processor E5 Family

The configuration of the database nodes is the following:

Feature	Performance Configuration	Capacity Configuration
Server Type	ISS3124-RP6	ISS3112-RP6
Processor	Intel® Xeon® E5 Series	Intel® Xeon® E5 Series
Processing cores	16	16
Hard-disk Drives	20x 2.5 in. 900GB 10k RPM SAS	10x 3.5 in. 3 TB 7.2k RPM SAS
SSD	4x 100 GB SLC	2x 200GB SLC
Boot Drive	1x 80 GB SLC	1x 80 GB SLC
Main Memory	8x 8GB DIMMs	8x 8GB DIMMs

The SGI DataRaptor has a number of configurations in both a performance and a capacity flavor: quarter-rack, half-rack, full-rack and multi-racks (out to nine, currently). The configurations are listed in the following tables:

Feature	1/4-Rack Performance Configuration	1/4-Rack Capacity Configuration
Storage Server	ISS3124-RP2	ISS3112-RP2
Number of servers	5	5
Disk Drives per Server	20	10
Type of Disk Drive	SAS, 6 GB/s, 2.5 in., 10K RPM, 900 GB capacity	SAS, 6 GB/s, 3.5 in, 7.2K RPM, 3 TB capacity
SSDs per Server	4	2
Type of SSD	SATA, 3 Gb, 2.5 in. SLC 100 GB capacity	SAS, 6 Gb, 2.5 in, SLC, 200 GB capacity
Memory per Node	16x 8GB DIMMs	16x 8GB DIMMs
Raw Storage/Rack	92 TB	152 TB

Feature	1/2-Rack Performance Configuration	1/2-Rack Capacity Configuration
Storage Server	ISS3124-RP2	ISS3112-RP2
Number of servers	10	10
Disk Drives per Server	20	10
Type of Disk Drive	SAS, 6 GB/s, 2.5 in., 10K RPM, 900 GB capacity	SAS, 6 GB/s, 3.5 in, 7.2K RPM, 3 TB capacity
SSDs per Server	4	2
Type of SSD	SATA, 3 Gb, 2.5 in. SLC 100 GB capacity	SAS, 6 Gb, 2.5 in, SLC, 200 GB capacity
Memory per Node	16x 8GB DIMMs	16x 8GB DIMMs
Raw Storage/Rack	184 TB	304 TB

Feature	Full-Rack Performance Configuration	Full-Rack Capacity Configuration
Storage Server	ISS3124-RP2	ISS3112-RP2
Number of servers	21	21
Disk Drives per Server	20	10
Type of Disk Drive	SAS, 6 GB/s, 2.5 in., 10K RPM, 900 GB capacity	SAS, 6 GB/s, 3.5 in, 7.2K RPM, 3 TB capacity
SSDs per Server	4	2
Type of SSD	SATA, 3 Gb, 2.5 in. SLC 100 GB capacity	SAS, 6 Gb, 2.5 in, SLC, 200 GB capacity
Memory per Node	16x 8GB DIMMs	16x 8GB DIMMs
Raw Storage/Rack	386 TB	638 TB

6.0 Software, Management and Networking Infrastructure

SGI DataRaptor is built around the following key software components:

Operating System	Redhat RHEL 6.2 (2.6.32-220.el6.x86_64)
Database	MarkLogic® 6 v.6.0-1.1
Management	SGI Management Center (SMC) 1.6.0

SGI Management Center is used to manage and monitor the ISS servers in the solution, using a Rackable™ C1110-RP6 server as the host server, which is in turn connected via a GigE fabric and Ericsson GigE switch to the ISS servers.

The ISS servers are connected to each other via a dual 10GigE network, providing performance and high-availability to the solution.

MarkLogic Rack Networking Layout

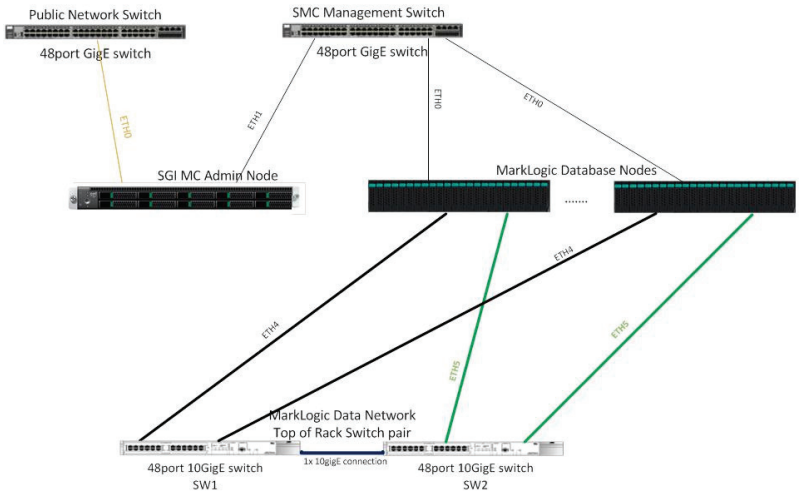

In a multi-rack configuration, Rack 1 would have the following configuration:

- 2x 48 port 10GigE switches for Data
- 1x 48 port GigE switch for Management
- 2x 48 port 10GigE spine switches

Rack 2 and beyond would have the following config:

- 2x 48 port 10GigE switches for Data
- 1x 48 port GigE switch for Management

7.0 SGI DataRaptor Speeds and Feeds

Full Rack Features	Performance Configuration	Capacity Configuration
Uncompressed Disk Bandwidth	Up to 47 GB/s	Up to 26 GB/s
Uncompressed Flash Data Bandwidth	Up to 20 GB/s	Up to 10 GB/s
Database Disk IOPS	Up to 144,000 8K reads	Up to 32,500 8K reads
Database Flash IOPS	Up to 2,100,000 8K reads	Up to 1,500,000 8K reads
Formatted Disk Capacity (after RAID 6)	300 TB	504 TB
Uncompressed Usable Capacity	~135 TB	~ 224 TB
Rack	42U/standard rack (24 in. width)	42U/standard rack (24 in. width)
Nodes	21	21
CPU	42 Intel E5	42 Intel E5
Cores	336	336
Memory	2688 GB	2688 GB
Disk / Flash	420x 900GB 10k RPM + 84x 100 GB SLC SSD	210x 3TB 7.2k RPM + 42x 200 GB SLC SSD
Networking	2x 10 GigE Data, 1x GigE for Management	2x 10 GigE Data, 1x GigE for Management

8.0 MarkLogic 6 Database Key Features

Features	Description	Advantages
Universal Index	Stores information in the database w/o the need for modeling or a schema	Allows for any type of heterogeneous, unstructured information to be captured and stored on the fly
XML Data Model	Stores information using a W3C-standard model with support for XQuery, XSLT	Embeded markup in XML allows granular search using context; makes it extensible to new data elements; makes it flexible to capture any unpredictable information
ACID-compliant transaction controller	Inserts and indexes information in real-time (due to Multi-version concurrency control) simultaneously with queries	Processes millions to billions of queries against a fast incoming feed of new information
Search, Analytics and Visualization	Uses context search for relevance and range indexes for structuring values for precise information; supports in-database analytics with Xquery; supports REST API for building fully functional MarkLogic applications in any programming language; supports java API for paginated search with facets and snippets; supports visualization widgets with application builder	Makes it an efficient search and analytical engine for unstructure data, followed by information intelligence
MarkLogic Content Pump (mlcp)	Command-line tool for fast, parallel load and content migrations	Allows speed loading and exporting data; takes advantage of hadoop cluster for load parallelism
Shared Nothing Architecture	Scale-out, flexible, high performance architecture to address unanticipated growth of unstructure information	Allows for linear scaling with user load within minimal response time

9.0 Benchmark

The benchmark used to test SGI DataRaptor is the CPoX benchmark, **C**ontent **P**rocessing framework **O**ver **X**ML. The benchmarks consists of a Wikipedia workload and simulates web users performing concurrent searches, updates and navigation—basically a search application, with faceted navigation and document updates.

The CPoX benchmark has two phases:

- Document loading. Wikipedia articles are loaded in using multiple threads. Each thread simulates a user repeating the same request without sleeping. The metric is documents per sec (docs/s)
- Runtime requests. The runtime phase simulates users performing a search, viewing an article, or performing an edit/confirm cycle for an article. The metric is requests per second (requests/s) where a search, a view and an edit/confirm are each viewed as a request

In the CPoX benchmark, each thread behaves as a normal Wikipedia user most of the time – performing searches and viewing articles.

- About 10% of the time, the thread would behave as a Wikipedia editor. An article would be retrieved and the thread would add its unique signature at the end of the article
- Immediately after this edit, the same thread would perform a search for the number of articles that contain this signature. This number should increase by 1 every time. If not, CPoX would report an error

The threads used for both loading and runtime do not simulate user sleep time. Each thread would perform loads or requests as soon as the previous load or request is complete. Because of this, average response time can be calculated as:

1000ms/ (overall_throughput/number_of_threads), or
*(number_of_threads*1000ms)/overall_throughput.*

The data set consists of 61.8 million documents (4KB/doc) in 11 languages. After fully loaded into MarkLogic server with the indexing option specified by the benchmark, it would consume about 1TB of disk space.

The benchmarks have been run to validate performance of the appliance with high performance and high capacity drives.

Results as of the date this paper was written (September, 2012) show that CPoX@1TB throughput on a 10-node (1/2 rack) appliance cluster with high performance drives, scales almost linearly from 50% to 100% of the system size delivering 358 requests/sec with a total of 160 physical cores or concurrent user threads. Hyperthreading provides only a 5% improvement in throughput over 160 physical threads, with a steep increase in response time beyond a threshold of ~500 ms (Figures 4 and 5).

With no sleep time between thread initiations, the CPU usage peaks to 90% at 160 threads, the total system size. So in such scenarios, it is recommended to use only up to a total number of physical cores, while running user transactions (160 user threads in this case) in order to maintain a response time below the threshold of ~500ms.

The CPoX@1TB throughput scales nearly linearly from 1 to 5 and 10 nodes in the appliance with high performance drives, delivering up to **358 requests/sec** on a ½ Rack.

The ingestion rate using the xqsync load utility scales learly linearly from 1 to 5 and 10 nodes in the cluster, delivering up to **15,470 docs/sec** on a ½ Rack.

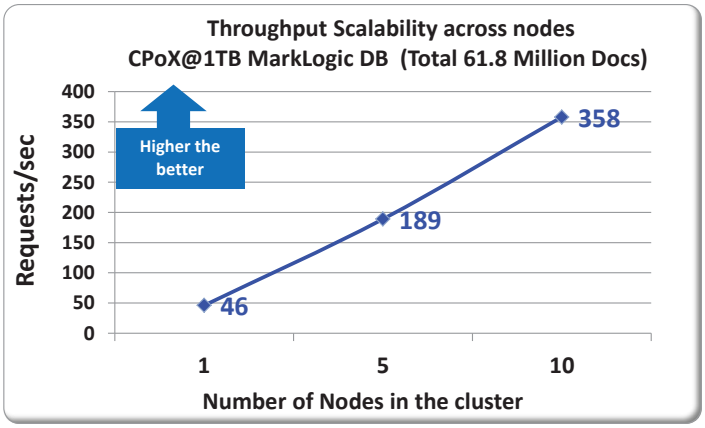

CPoX@1TB Throughput: Multi-node scalability

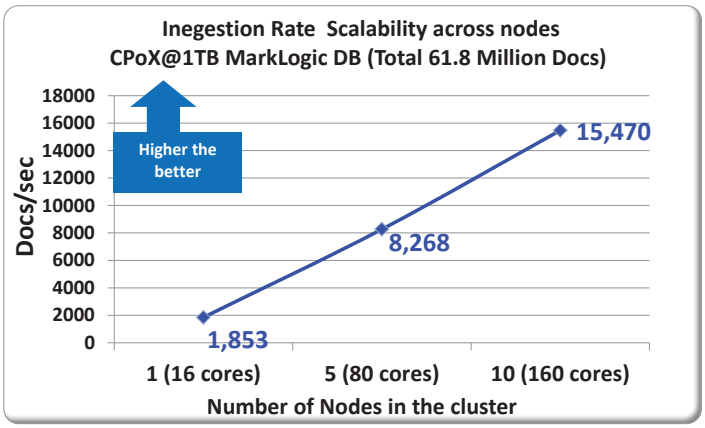

CPoX@1TB Ingestion Rate: Multi-node scalability

10.0 SGI DataRaptor with MarkLogic Database Summary

SGI DataRaptor with MarkLogic Database provides:

- All you need for Big Data built into a single system, factory-integrated and with plug-and-play capability that allows you to be up and analyzing Big Data in hours, not days or weeks
- Rock solid Big Data capability, containing the MarkLogic database
- Real ROI for your business or enterprise, with rapid prototyping to develop applications quickly so that you can deploy new applications in record time, monetize your Big Data and crowd source your decisions quickly
- Big Data. Ready. Set. Go. Starting at 100s of TBs and expanding quickly, at ease and under load, the SGI DataRaptor offers one price and one support line to call

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

Global Sales and Support: sgi.com/global