

SGI DataRaptor with MarkLogic Database

Hardware. Software. Complete.

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.



Deriving Key Value from Big Data Today, Not Tomorrow

Big Data can bring tremendous value to customers in multiple industries and those firms that use Big Data successfully can have a strategic advantage in the marketplace and political landscape of today:

- 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

All of the above 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 MarkLogic and SGI solution

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

MarkLogic, the Big Data Database

NoSQL databases are the answer for transactional Big Data. And MarkLogic brings enterprise-class to NoSQL. By avoiding the Extract-Transform-Load (ETL) cycle inherent with relational databases and with APIs and other tools for fast application development, MarkLogic takes months off the process of getting value from your Big Data. Included among the features of the MarkLogic database are:

- Operational, NoSQL, enterprise-class database
- ACID transactions for reliable, consistent data operations
- High-availability, disaster recovery and security to safeguard your data
- Built-in search capabilities, including full-text, geospatial and real time alerting
- Rapid prototyping of new applications



A Complete, Expandable, and Flexible Solution

SGI® DataRaptor™ with MarkLogic® Database was built via a collaboration between SGI and MarkLogic core and applications engineering teams. Leveraging the experience of hundreds of MarkLogic customers including BBC, the U.S. Army and JPMorgan Chase, MarkLogic engineers specified the requirements that SGI engineers turned into a complete solution.

- All you need for Big Data is built into a single system, factory-integrated and with plug-and-play capability that allows a customer to be up and analyzing Big Data in hours, not days or weeks. A validation database is included with the solution so that you can run some standard queries to test functionality
- Rock solid Big Data capability, the solution contains the MarkLogic database, the best of breed NoSQL solution with ACID transactions, solving the problems that can't be solved with relational databases
- Real ROI for your business or enterprise, with rapid prototyping allowing you 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 hundreds of TBs, and expanding quickly, at ease and under load the SGI® DataRaptor™ with MarkLogic® Database offers one price and one support line to call

Configured for Either Performance or Capacity

The core to the SGI DataRaptor with MarkLogic Database is the SGI® 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 x16 slots, double-width
- Up to twelve 3.5 inch disks or twenty-four 2.5 inch disks, SAS, SATA or SSD

Feature	Full-rack Performance Configuration	Full-rack Capacity Configuration
Storage Server	ISS3124-RP2	ISS3112-RP2
Number of servers	21	21
Total Processing Cores	336	336
Total RAM Memory (GB)	2688	2688
Total Formatted Storage (TB)	300	504
Uncompressed Disk Bandwidth	Up to 47 GB/s	Up to 26 GB/s
Database Disk IOPS (8K Reads)	Up to 144,000	Up to 32,500
Networking	10GigE Data, GigE Admin	10GigE Data, GigE Admin

SGI Big Data Leadership

SGI has been providing Big Data solutions since long before Big Data become popular. Whether these are 24x7 installations ingesting large amounts of data and producing various analysis by-products, or 1000+ node Hadoop installations, SGI has been there with the fully-integrated and supported solutions to bring rapid time to insight. SGI has provided:

- Unprecedented speed and scale with our heritage in the areas of high-performance and technical computing
- End-to-end solutions, hardware, software, storage and compute, that are tested by engineering as a complete solution offering before they are fully integrated in the factory for shipment to customers.
- Power-optimized solutions to save OPEX dollars while you compute and store your data.
- Plug In and Go! SGI understands that organizations need solutions that deploy quickly and bring about ROI out of Big Data



Global Sales and Support: sgi.com/global

©2012 Silicon Graphics International Corp. All rights reserved. SGI and the SGI logo are registered trademarks or trademarks of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries. All other trademarks are property of their respective holders. 10102012 4389

