



SGI

Industry

High Performance Computing

Website

www.sgi.com

Company Overview

The trusted leader in high performance computing (HPC), SGI (NASDAQ: SGI) is focused on helping customers solve their most demanding business and technology challenges by delivering technical computing, Big Data analytics, cloud computing, and petascale storage solutions that accelerate time to discovery, innovation, and profitability.

Product Overview

SGI[®] InfiniteData™ Cluster enables business and government agencies to derive value from Big Data at greater speed, scale, and efficiency. Factory-integrated with the Cloudera Hadoop distribution running on Red Hat Linux and SGI Management Center, SGI's Cloudera-certified solution is ready to power up and go.

Achieve Faster, Greater Insight at Lower Cost with Cloudera-Certified SGI[®] InfiniteData™ Cluster

Organizations increasingly recognize the potential of Big Data to transform their business—improving customer retention and acquisition, increasing operational efficiencies, enabling better products and service delivery, and generating new business insights. An analytics-ready data management platform is critical to take full advantage of Big Data.

Historically utilized for complex, computationally intensive problems, High Performance Computing (HPC) is now being applied to Big Data analytics. SGI has been a driving force in HPC for nearly two decades and drawing upon vast expertise in clustered computing, was at the forefront of Apache Hadoop with installations currently totaling tens of thousands nodes.

The potential gains surrounding Big Data are far reaching and also present new challenges: how to derive value at greater speed, scale, and efficiency. SGI InfiniteData Cluster with Cloudera is designed to meet this challenge and enables business and government agencies perform Big Data analytics with faster, greater insight and at lower cost.

Applying HPC Leadership to Big Data Analytics

Performance + Density

InfiniteData Cluster delivers optimized Hadoop performance using a unique SGI node architecture that provides over 2X the compute storage density per rack of any enterprise-class HPC cluster. SGI's 3rd generation tray design features dual, 6-core Intel[®] Xeon[®] E5-2600 v2 processors and up to twelve 3.5" hard drives per server/storage tray, delivering 1:1 core to storage spindle ratios to speed analytics and up to forty 2U cluster nodes and 1.9PB in a single InfiniteData Cluster rack.

Accelerate Time to Solution

InfiniteData Cluster solutions for Hadoop are factory-integrated with Cloudera running on Red Hat Linux and with SGI Management Center, and arrive with all hardware and software pre-racked, pre-configured, and pre-tested. Ready to "power up and go", InfiniteData Cluster provides business with a world-class analytics platform in days vs. months and saves IT hundreds to thousands of man hours in set up time.

Solution Highlights

- Optimized Performance in ½ the Footprint – 1:1 core to storage spindle ratio delivering over 2X the compute/storage density over other HPC solution
- Save Countless Man Hours – Solutions arrive pre-racked, pre-configured, and pre-tested
- Unparalleled Scalability – Flexible deployment size with linear scalability
- Superior Serviceability and Energy Efficiency – Cold-aisle accessible with < 10W per tray for cooling
- Expert Services and Cloudera Certified – Cloudera certified with 24x7 support

Highly Scalable Architecture

InfiniteData Cluster is designed to expand as your data and compute needs grow. Solutions can start with a small number of cluster nodes in a single rack, or encompass multiple racks with thousands of nodes and linear scalability. The unique form factor of the half-width 2U server/storage tray can accommodate 3.5" SATA III drives with capacities ranging from 500GB to 4TBs, providing up to 48TBs per data node. Highly resilient network switching includes GbE for cluster management, 10GbE for node-to-node connectivity, and 40GbE for multi-rack connectivity.

Serviceability and Energy Efficiency

All hardware components—trays, interconnects, and I/O—are cold-aisle accessible to support energy-efficient hot-aisle/ cold-aisle floor plans. The server/storage trays pull out easily for service and disk drives can be added or replaced in minutes. InfiniteData Cluster also helps lower energy costs as trays utilize high efficiency 450W power supplies and require under 10W per tray for cooling.

Cloudera-Certified with World Class Services

InfiniteData Cluster Solutions for Hadoop are fully certified by Cloudera, with pre-configured options for Cloudera software, and interoperability with Cloudera's broad analytics partner community. SGI and Cloudera services personnel bring expertise and experience to help architect your Hadoop environment and turn powerful technology into business solutions. SGI products are fully backed by warranty with technical support available around the globe up to 24 by 7.

SGI InfiniteData Cluster Specifications

Core Features	Specifications		
Cluster Nodes	Name, Standby Node, Job Tracker Nodes - SGI C1110-RP6 <ul style="list-style-type: none"> • 1U full width • (2) Intel Xeon Processor E5-2630 v2 series, 6-core • 8 x 8GB memory • 4 x 3.5" SATA III drives, RAID 10 • 1 x dual port 10GbE NIC • Redundant power supply 	Application Node – SGI C1110-RP6 <ul style="list-style-type: none"> • 1U full width • (2) Intel Xeon Processor E5-2680 v2 series, 10-core • 16 x 8GB memory • 4 x 3.5" SATA III drives, RAID 10 • 1 x dual port 10GbE NIC • Redundant power supply 	Data Nodes– SGI IDC3212 –RP4 <ul style="list-style-type: none"> • 2U ½ width • (2) Intel Xeon Processor E5-2630 v2 series, 6-core • 8 x 8GB memory • 12 x 3.5" SATA III drives, JBOD • 1 x dual port 10GbE NIC • Redundant power supply
Cluster Network	10GbE Node-to-Node <ul style="list-style-type: none"> • (2) Extreme Networks Summit X670V, active dual-path 	40GbE Rack-to-Rack (included with multi-rack clusters) <ul style="list-style-type: none"> • (2) Mellanox SX1012, active dual-path 	GbE Cluster Management <ul style="list-style-type: none"> • (1) Edge-Core ECS4610-50T
Operating System	<ul style="list-style-type: none"> • Red Hat Enterprise Linux 6.4 		
Cluster Management	<ul style="list-style-type: none"> • SGI Management Center • Supports Intel Remote Management Module 4 		
Scalability	<ul style="list-style-type: none"> • Solutions begin at eight data nodes and scale to thousands of data nodes utilizing multiple racks with linear performance 		
Cloudera	<ul style="list-style-type: none"> • Cloudera Standard Hadoop Distribution with Cloudera Manager • CDH5 Enterprise Edition (optional) 		
Hadoop Environment	<ul style="list-style-type: none"> • CDH includes HDFS, MapReduce, Flume, Hive, Pig, etc. 		
Enclosure	<ul style="list-style-type: none"> • 42U Cabinet with 4U extension; 85.6" x 24" x 47.4" • Up to 40 data nodes (1,920TB; 1,920 cores) per rack 		
Power and Cooling	<ul style="list-style-type: none"> • AC PSU at the tray level • Front-to-back airflow; (2) tray level fans 		
Support	<ul style="list-style-type: none"> • Up to 24 x7 SGI global support; optional Cloudera support (Cloudera Enterprise) 		

About SGI

Building on 25 years of innovation with over 600 patents, SGI (NASDAQ: SGI) has over 1,400 employees worldwide, serves over 6,500 customers, and is distributed in 50 countries through our direct and indirect sales force. For more information visit www.sgi.com/solutions/bigdata/hadoop.

About Cloudera

Cloudera is revolutionizing enterprise data management by offering the first unified Platform for Big Data: The Enterprise Data Hub. Cloudera offers enterprises one place to store, process and analyze all their data, empowering them to extend the value of existing investments while enabling fundamental new ways to derive value from their data. Founded in 2008, Cloudera was the first and is still today the leading provider and supporter of Hadoop for the enterprise. Cloudera also offers software for business critical data challenges including storage, access, management, analysis, security and search. With over 15,000 individuals trained, Cloudera is a leading educator of data professionals, offering the industry's broadest array of Hadoop training and certification programs. Cloudera works with over 700 hardware, software and services partners to meet customers' big data goals. Leading organizations in every industry run Cloudera in production, including finance, telecommunications, retail, internet, utilities, oil and gas, healthcare, biopharmaceuticals, networking and media, plus top public sector organizations globally. www.cloudera.com.

cloudera.com

1-888-789-1488 or 1-650-362-0488

Cloudera, Inc. 1001 Page Mill Road, Palo Alto, CA 94304, USA

© 2015 Cloudera, Inc. All rights reserved. Cloudera and the Cloudera logo are trademarks or registered trademarks of Cloudera Inc. in the USA and other countries. All other trademarks are the property of their respective companies. Information is subject to change without notice.

cloudera-sgi-solutionbrief-104