

BIG DATA APPLIANCE

ORACLE BIG DATA APPLIANCE

BIG DATA FOR THE ENTERPRISE

KEY FEATURES

- Massively scalable infrastructure to store and manage big data
- Big Data Connectors delivers unprecedented load rates between Big Data Appliance and Oracle Exadata
- Cloudera's Distribution including Apache Hadoop delivers managed and proven Hadoop to the enterprise
- Cloudera Manager simplifies management of Hadoop
- Advanced analytics with Oracle R on Hadoop data
- Oracle NoSQL Database Community Edition pre-installed and configured
- 648 TB of raw storage
- InfiniBand Connectivity between nodes and across racks

KEY BENEFITS

- Optimized and Complete Big Data Solution
- Integrated with Oracle Exadata to analyze all your data
- Risk free installation and quick time to value
- Single vendor support for your entire big data solution

Oracle Big Data Appliance brings Big Data solutions to mainstream enterprises. Built using industry-standard hardware from Sun and Cloudera's Distribution including Apache Hadoop, Big Data Appliance is designed and optimized for big data workloads. By integrating the key components of a big data platform into a single product, Oracle Big Data Appliance delivers an affordable, scalable and fully supported big data infrastructure without the risks of a custom built solution. Oracle Big Data Appliance integrates tightly with Oracle Exadata and Oracle Database using Oracle Big Data Connectors, and seamlessly enables analysis of all data in the enterprise – structured and unstructured.

Bringing Big Data to the Enterprise

Oracle Big Data Appliance is pre-installed and pre-configured for large-scale big data management. Using Cloudera's Distribution including Apache Hadoop (CDH) and Oracle NoSQL Database as data management capabilities, the Big Data Appliance runs on Oracle Linux and Oracle HotSpot JVM. It includes Cloudera Manager for cluster-wide administration and monitoring of CDH. For deep analysis of big data, it also includes an open-source distribution of the statistical environment R. As compared to building a big data system from scratch, the Big Data Appliance eliminates the time-consuming efforts of choosing and configuring hardware, determining the proper open-source components and versions, and integrating and tuning the overall configuration. The entire solution is installed and configured out of the box for high performance and high availability.

Big Data Appliance
Integrated Software
Oracle Enterprise Linux 5.6
Oracle Hotspot Java Virtual Machine
Cloudera's Distribution including Apache Hadoop
Cloudera Manager
Open Source Distribution of R
Oracle NoSQL Database Community Edition

The Oracle Big Data Appliance is designed for enterprise production environments. Oracle provides enterprise support for the entire solution, both hardware and software. For high availability, the Big Data Appliance has complete redundancy built in. The Big Data Appliance has redundant InfiniBand connectivity and redundant Power Distribution Units (PDU), and the servers all have hot-swappable power supplies for high availability. In addition, data is automatically mirrored across servers with multiple replicas to protect against data loss and ensure data accessibility.

Designed and Optimized for Big Data

The Big Data Appliance has been designed and optimized for Big Data from the ground up. With 648 TB of raw storage and 288 CPU processing cores in a single rack, the Big Data Appliance delivers large-scale data processing. The Big Data Appliance scales by adding additional racks; up to 8 racks can be connected by connecting InfiniBand cables, and larger

RELATED PRODUCTS AND SERVICES

Oracle Big Data Appliance brings a low risk, highly scalable big data platform to the enterprise.

RELATED PRODUCTS

The following are related products available from Oracle:

- Oracle Exadata
- Oracle Exalytics
- Oracle Exalogic

RELATED SERVICES

The following services are available from Oracle Support Services:

- Advanced Customer Services
- Product Support Services
- Consulting Services
- Oracle University Courses

configurations are supported with additional external InfiniBand switches.

The Big Data Appliance has been specifically optimized for CDH to ensure scalable performance. The Big Data Appliance is configured with pre-tuned operating system parameters, file system settings and Java VM configuration properties to allow optimized performance. Moreover, the Big Data Appliance's use of InfiniBand is well-suited for MapReduce processing, since data redistribution during the "map" processing phase can be distributed much faster between nodes.

Big Data Appliance has been optimized and pre-configured with Oracle NoSQL Database. Oracle NoSQL Database provides fast low-latency access to massive amounts of data, supporting both data retrieval and data updates using a flexible key-value pair programming paradigm. Enterprises may choose the Big Data Appliance as a system to run only CDH or to run both NoSQL Database and CDH simultaneously on the same cluster (for example, to capture data using NoSQL Database, and then integrate and process the data for historical analysis using CDH).

Big Data: The Big Picture

The true benefits of big data are realized when new insights from big data are combined with insights from an enterprise's existing information. A big data solution is not simply a system that executes MapReduce programs against a new set of data; instead, a big data solution requires an architecture that manages both structured and unstructured data, provides deep analytical capabilities against that data, and delivers information directly to business users to help drive business decisions. Oracle Big Data Appliance was designed to be part of an enterprise's overall big data solution. When combined with Oracle Exadata Database Machine for managing structured data, executing deeper analytics, and supporting large numbers of users, and Oracle Exalytics In-Memory Machine for providing speed-of-thought analytical environment to business users, Oracle Big Data Appliance is a core component of a complete analytic infrastructure.

Big Data Appliance, built with Infiniband, allows scalable interconnectivity with all of Oracle's other engineered systems: not only Oracle Exadata and Oracle Exalytics, but also Oracle Exalogic and Supercluster. At 40Gb/sec, Infiniband provides more bandwidth as well as lower latency than ethernet-connected systems.

Oracle Big Data Connectors, a separately licensed software-only product, provides a high-performance Hadoop to Oracle Database integration solution. Big Data Connectors allows data to be loaded from CDH into Oracle Database or can be accessed in-place on the Big Data Appliance from an Oracle Database. Such high-performance data movement and access is necessary infrastructure for enterprises who want to realize the full value of big data.

Big Data Appliance	
Hardware Specification and Details	
18 Compute and Storage nodes	Per Node: <ul style="list-style-type: none"> • 2 x Eight-Core Intel® Xeon® E5-2260 Processors (2.2 GHz) • 64 GB Memory (expandable to 256GB) • Disk Controller HBA with 512MB Battery backed write cache • 12 x 3TB 7,200 RPM High Capacity SAS Disks • 2 x QDR (40Gb/s) Ports • 4 x 10 Gb Ethernet Ports • 1 x ILOM Ethernet Port
2 x 32 Port QDR InfiniBand Switch	<ul style="list-style-type: none"> • 32 x InfiniBand ports • 4 x 10Gb Ethernet ports
1 x 36 Port QDR InfiniBand Switch	<ul style="list-style-type: none"> • 36 InfiniBand Ports

Additional Hardware included:	<ul style="list-style-type: none"> • Ethernet switch for administration of the Appliance • 2 x Redundant Power Distributions Units (PDUs) • 42U rack packaging
Spares Kit Included:	<ul style="list-style-type: none"> • 2 x 3 TB High Capacity SAS disk • InfiniBand cables

Big Data Appliance – Environmental Specifications

Physical Dimensions	
Height	42U, 78.66" - 1998 mm
Width	23.62" - 600mm
Depth	47.24" - 1200 mm
Weight	
Installed	2005Lbs
Power	
Maximum	10.2 kVA 10.0 KW
Typical	7.15 kVA 7.0 KW
Cooling	
Maximum	34,142 BTU/hour 34,838 kJ/Hour
Typical	23,940 BTU/hour 24,420 kJ/Hour
Airflow Maximum	1,573 CFM
Airflow Typical	1,103 CFM
Further Environmental Specifications	
Operating temperature/humidity: 5 °C to 32 °C (41 °F to 89.6 °F), 10% to 90% relative humidity, non-condensing	
Altitude Operating: Up to 3,048 m, max. ambient temperature is de-rated by 1° C per 300 m above 900 m	
Regulations ¹	
Safety: UL 60950-1 2nd Ed, EN60950-1:2006 2nd Ed, CB Scheme with all country differences	
RFI/EMI: FCC CFR 47 Part 15 Subpart B Class A, EN 55022:2006+A1:2007 Class A, EN 61000-3-11:2000, EN 61000-3-12:2005, ETSI EN 300 386 V1.4.1 (2008)	
Immunity: EN 55024:1998+A1:2001:+A2:2003	
Certifications ¹	
Safety: UL/cUL, CE, BSMI, GOST R, S-Mark, CSA C22.2 No. 60950-1-07 2nd Ed, CCC	
EMC: CE, FCC, VCCI, ICES, KCC, GOST R, BSMI Class A, AS/NZ 3548, CCC	
Other: Complies with WEEE Directive (2002/96/EC) and RoHS Directive (2002/95/EC)	

¹ In some cases, as applicable, regulatory and certification compliance were obtained at the component level.

Big Data Appliance Support Services

- Hardware Warranty: 1 year with a 4 hour web/phone response during normal business hours (Mon-Fri 8AM-5PM), with 2 business day on-site response/Parts Exchange
- Oracle Premier Support for Systems: Oracle Linux and integrated software support and 24x7 with 2 hour on-site hardware service response (subject to proximity to service center)
- Oracle Premier Support for Operating Systems
- Oracle Customer Data and Device Retention
- System Installation Services
- Software Configuration Services
- Oracle Auto Service Request (ASR)

Big Data Appliance Upgrade Services

Multi-Rack Connection

Connect any combination of Big Data Appliance or Exadata Database Machine Full Rack or Half Rack via included InfiniBand fabric

Other configuration considerations:

- Up to 8 racks can be connected without requiring additional InfiniBand switches
- InfiniBand cables to connect 3 racks are included in the rack Spares Kit
- Additional optical InfiniBand cables required when connecting 4 or more racks

Memory Upgrades

- Upgrade any number of servers within Big Data Appliance from 64GB up to 256GB

Contact Us

For more information about Oracle Big Data Appliance, visit oracle.com or call +1.800.ORACLE1 to speak to an Oracle representative.



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2012, Oracle and/or its affiliates. All rights reserved.

This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners. Cloudera, Cloudera CDH, and Cloudera Manager are registered and unregistered trademarks of Cloudera, Inc.,

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0611

Hardware and Software, Engineered to Work Together