


Teradata Aster Big Analytics Appliance

09.15 EB6434 INTEGRATED BIG DATA ANALYTICS AND DISCOVERY SOLUTION

New generation big data technologies can often be quite innovative. Yet, if they are not well thought out, they can be negatively disruptive accompanied by large capital and resource outlays. But that doesn't have to be the case. The Teradata Aster Big Analytics Appliance meets the needs and requirements of all types of organizations to harness the power of 21st century big data.

An optimized environment for rapid, iterative data exploration and advanced analytics, this is the first appliance to combine big data and analytics technologies, such as Hadoop, open source R, MapReduce and graph analysis. This appliance can be extended to include the industry leading Teradata integrated data warehouse (IDW) to meet your organizations data and analytic requirements. This environment is compatible with BI and ETL tools and is a tightly coupled hardware and software solution that contains Teradata Aster Database, Teradata IDW and a choice of Hadoop distribution software from Hortonworks or Cloudera to process structured, unstructured and semi-structured data. Key capabilities of the appliance include:

- Best-in-class enterprise big data analytics via industry's only SQL-MapReduce® and SQL-GR™ interface and the comprehensive analytic portfolio of data acquisition, data preparation, advanced analytics and visualization functions. The portfolio features more than 100 pre-built functions, providing out-of-the-box analytics for text and sentiment parsing, time series behavior, network centrality measurement, data cleansing and preparation, machine learning, location analysis and natural language processing.
- A truly hybrid architecture that includes Aster Database, Teradata Database, and Hadoop for a complete big data analytics solution.
- Industry's deepest stack integration, across the Aster, Hadoop, and Teradata IDW systems either within a single cabinet or across multiple appliances. Enables business analysts to issue ANSI-standard SQL or R commands on Aster, Hadoop, or Teradata IDW data via Teradata QueryGrid™.
- Industry-leading performance and scalability via optimized hardware and software configuration for Aster, Teradata IDW, and Hadoop.
- High speed 40Gb/s InfiniBand network connectivity across the Teradata Unified Data Architecture™ enables expansion to multiple Aster Big Analytics Appliances and communication with other Teradata purpose-built appliances.
- Enterprise-ready software and hardware management, via Teradata Server Management, Teradata Viewpoint, and Teradata Vital Infrastructure services.
- Fully supported by the most trusted name in enterprise data management and analytics, Teradata.



An Out-of-the-Box Big Analytics Solution

The Teradata Aster Big Analytics Appliance is a powerful, ready-to-run platform that is configured and optimized specifically for big data analytics and management. A purpose-built, integrated hardware and software solution for analytics at big data scale, the appliance runs the Teradata patented SQL-MapReduce®, SQL-GR™, R,

Teradata Aster AppCenter, and Teradata QueryGrid™ technologies on a time-tested, fully supported Teradata hardware platform. Depending on workload needs, it can be configured with Aster nodes exclusively, Hadoop nodes with Hortonworks Data Platform (HDP) or Cloudera Enterprise exclusively, or a mixture of Aster, Hadoop and Teradata IDW nodes.

By minimizing the number of moving parts required for deployment, the appliance offers easy, integrated management of an enterprise-ready big data analytics and discovery solution with the benefits of optimized performance, continuous availability, and linear scalability. Now you can bring advanced analytics to the business by just plugging the appliance into existing infrastructure, thereby leveraging your current investments in technology and resources.

Takes Business Analysts from Big Data to Powerful Insights

The appliance comes with Aster analytics, which features more than 100 pre-packaged SQL-MapReduce and Graph functions to enable faster insights. The SQL-MapReduce® and SQL-GR™ frameworks, created by Teradata, allows developers to write powerful and highly expressive MapReduce and graph functions in languages such as Java, C#, Python, C++, and R, and push them into the discovery platform for high speed, scalable analytic processing. Business analysts can then invoke SQL-MapReduce and SQL-GR™ functions using standard SQL or R through Aster Database, the first discovery platform that allows applications to be fully embedded within the database engine to enable ultra-fast, deep analysis of massive data sets. Using Teradata QueryGrid™, business analysts can easily access Hadoop and Teradata data from within the Aster Database via standard SQL commands and BI tools. Optional Hadoop nodes can be used for staging, data transformation, and long-term data archival. Optional Teradata nodes can be used to manage production data and analytic results while off-loading the actual analytic processing on to Aster nodes. Through technologies like SQL-MapReduce®, SQL-GR™ and QueryGrid™, business analysts can now seamlessly analyze structured and unstructured data from integrated and unified SQL-based query and analytics interfaces and their choice of BI tools.

Built for Ease of Start-up

The Teradata Aster Big Analytics Appliance is purpose built for big data analytics on multi-structured data. It features a complete Aster Database, including the patented SQL-MapReduce and SQL-GR™ frameworks, on a proven Teradata hardware platform with dual 12-core Intel®

processors, SUSE® Linux operating system, and enterprise-class storage—all preinstalled into a power-efficient unit. That means you can have the Aster system up and running live in just a few hours for rapid time to value.

Delivers Reliable, Scalable and Powerful Processing

Massively parallel processing (MPP) architecture with embedded SQL, R, graph and MapReduce engines enable end-to-end parallelism of data and analytic processing. Aster Database executes loads, queries, exports, backups, recoveries, installs, and upgrades in parallel to take full advantage of all resources and thus optimizing processing performance. The appliance delivers high availability through a disk RAID and a software-redundant design, and Teradata's advanced Server Management features proactively monitor the system with diagnostic tools to quickly detect and address any issues to maximize availability.

Enterprise Ready Solution

The solution is optimized, certified and fully supported by Teradata, eliminating the need for specialized software or hardware management skills required to cobble together a big data analytic solution. Existing Teradata customers will find the management consoles and capabilities familiar and complementary to their existing database and data warehouse deployments.

A suite of adapters support out-of-the-box enterprise integration for a complete ecosystem of enterprise data management systems. Certified ODBC and JDBC support for major business intelligence, visualization, and ETL tools; Teradata integrated data warehouse high-speed data transfer infrastructure; and native Hadoop connectivity are all available to enhance your enterprise solution.

The Teradata Aster Big Analytics Appliance offers the flexibility of AppCenter server and generic Teradata Managed Servers which can be used to run Aster or third party applications such as RStudio. These servers communicate with Aster nodes across the high speed InfiniBand network infrastructure.

Fast and Reliable Network Interconnect

A dual 40 Gb/s InfiniBand network connects the Aster, Hadoop, and Teradata nodes within a cabinet. This fast, dedicated and reliable network interconnect can also be used to connect multiple Aster Big Analytics Appliances and enables tight integration with the Teradata Unified Data Architecture.

Aster Big Analytics Appliance Description

Features:

- Aster Database 6.0 or higher
 - Over 100 SQL-based functions for data acquisition, data preparation, analysis and visualization
 - Open Source R engine within the Aster database
- Optional Hortonworks HDP 2.3 or Cloudera Enterprise 5.4 distribution
- Optional Teradata Database 14.0 or higher
- SUSE® Linux Enterprise Server (SLES) 11 64-bit Operating System
- Integrated Cabinet
 - Aster nodes
 - Dual Intel® Twelve-Core Xeon® processors @ 2.5GHz per node (Queen and Worker)
 - Dual Intel Eight-Core Xeon processors @ 2.6GHz per node (Backup and Loader)
 - Optional Hadoop nodes
 - Dual Intel Twelve-Core Xeon processors @ 2.5GHz per node (Master, Edge, Balance Data and Performance Data nodes)
- Dual Intel Eight-Core Xeon processors @ 2.6GHz per node (Capacity Data node)
- Optional Teradata Data Warehouse Appliance 2800 (2 or 4 nodes in first cabinet only)
- Optional Teradata Data Mart Appliance 680 SMP (1 node in first cabinet only)
 - Dual Intel Fourteen-Core Xeon processors @ 2.6GHz per Teradata node
- High throughput BYNET v5 on a dual 40Gb/s InfiniBand interconnect for performance and redundancy
- Optional 10GB switch for connecting nodes to customer LAN
- Optional Servers for AppCenter, ViewPoint and 3rd party applications
- RAID High Availability
- Disk Capacities – uncompressed (3x compression assumed typical)
 - Aster Worker Nodes with 1.2TB drives: 7.3TB
 - Balanced data nodes and Capacity nodes for Hadoop with 12 - 4TB drives: 16TB user space
- Performance data nodes for Hadoop with 24 - 1.2 TB drives: 9.6TB user space
- Teradata nodes with 300/600/900/ 1200GB drives: 3.8/7.5/11.4/26.7TB
- Aster Backup Nodes with 4TB drives: 36TB
- Full Aster Cabinet (2 Queens/16 Workers) - 117TB user space
- Teradata Data Warehouse Appliance 2800 - up to 107TB user space
- Teradata Data Mart Appliance 680 - up to 12TB user space
- Scalability to petabytes with expansion cabinets and larger network switch
- Teradata software developed and/or optimized for Teradata Aster and Hadoop
 - System Management Infrastructure, Teradata Administration, and Teradata Viewpoint 14.10 or higher
 - Teradata Query Grid: Allows data access in Aster, Hadoop, and Teradata through a standard SQL interface

Complete Centralized Management, Service, and Support

The platform features simplified administration, control, and monitoring through the single operational view on the Teradata Server Management and web-based Administration Workstation portal. The appliance delivers rich visibility and control not only of data, but also the SQL, R, graph and MapReduce analytic processes running inside the system.

It combines intuitive tools for centralized management with powerful capabilities for facilitating and automating administration, and managing data and processes to minimize administrative work, even as the system scales.

All the tools and utilities you need to build and maintain your analytic applications are available:

- **Integration with Teradata Server Management** enables proactive monitoring of hardware and software events, such as disk or node failures. With the Teradata proactive support software available on each Teradata Aster

Big Analytics Appliance, support information gathered is automatically routed to Teradata Customer Services' back-end support infrastructure, so issues can be addressed quickly thereby minimizing system downtime.

Teradata Server Management is a set of components that detect and report hardware and software exceptions (faults), hardware and software asset data, and supporting diagnostic data. Teradata Server Management runs in virtual machines on a series of Virtual Management Server (VMS) nodes/servers. The VMS supports Teradata hardware and Aster Database and Apache Hadoop software. It also provides Teradata Vital Infrastructure support for Teradata, Teradata Aster, and Hadoop.

Teradata Vital Infrastructure is the end-to-end solution for delivering Server Management data to Teradata Customer Service's back-end support infrastructure. When enabled and warranted, it escalates (forwards) the alerts and other types of data to the Teradata Customer Services back-end infrastructure. The diagnostic information collected by Teradata Vital Infrastructure

Specifications

Cabinet

- Height: 80.5 in. (204.5 cm)
- Width: 24 in. (61 cm)
- Depth: 49 in. (124.5 cm) 47 in. (119.4 cm) w/o the front and rear doors
- Weight: 2,228 lbs. (1011 kg) fully loaded with crate
- Installed Weight: 1,828 lbs. (830 kg) fully loaded w/o crate

Operating

- Operating Temperature: Allowable: 59°F to 90°F (15°C to 32°C);
 - Recommended: 64.8°F to 80.6°F (18°C to 27°C)
- Relative Humidity: Allowable: 20% to 80% (non-condensing);
 - Recommended: Low end moisture: 5.5C DP (41.9°F), High end moisture: 60% RH and 15C DP (59°F DP)

- Electrical
- Worldwide: 200 – 240V (Phase to Phase or Single Phase)
 - 30A/32A, 4-cord
 - North America: 200 - 240V, 3~+PE
 - 30A, 3-p delta, 4-cord
 - 60A, 3-p delta, 2-cord
 - International: 220 - 240 / 381 - 415, 3- +N +PE
 - 30A/32A, 3 phase wye, 2-cord, (including North America with EU style power)
- All plugs are IEC 60309 CEE17
- Frequency: 50Hz / 60Hz
- Maximum Power: 13,300 Watts
- Dual AC: Standard
- Compliant with U.S. and International Safety and Emissions Standards

Support Services

Maintenance and Support

- Integrated hardware and software maintenance and support
- Secure remote connectivity
- Fast response times
- Flexible coverage hours
- Robust diagnostic capabilities with Teradata Vital Infrastructure
- Easy access to software updates via Teradata @ Your Service
- Proactive system monitoring
- Implementation Services
- System Installation

assists support personnel in identifying and quickly resolving problems; conversely, omission of this software may prolong the time it takes to resolve incidents.

- **Integration with Teradata Viewpoint** offers a common management console for the Aster Database, Teradata IDW, and Hadoop. Teradata Viewpoint offers simpler, faster, and more comprehensive system management and monitoring capabilities by providing a browser-based portal that delivers management intelligence to DBAs and users alike. Viewpoint provides two dedicated portlets of Aster completed processes and Node Monitor for Aster system monitoring.
- **The Aster Management Console (AMC)** is used to configure, manage, and monitor data, applications, and infrastructure. An intuitive graphical interface enables easy monitoring with summary dashboards, graphical views of query and process execution, and easy drill-down. When using Viewpoint, the AMC is only needed for Aster environment configuration.
- **Always-On Online Maintenance** allows the Aster Database to support simultaneous load and export

during queries, online backup and recovery, online restoration, and online scaling to avoid scheduled downtime.

Why Teradata?

Teradata is the world's largest company focused on analytic data solutions through integrated data warehousing, big data analytics, and business applications. Only Teradata gives organizations the advantage to transform data across the organization into actionable insights empowering leaders to think boldly and act decisively for the best decisions possible.

For More Information

Find out how the Aster Big Analytics Appliance from Teradata can make your entry into big data enterprise analytics fast, efficient, and cost effective while you improve your decision-making capabilities and grow a stronger, more productive business. Contact your local Teradata representative or visit Teradata.com.

10000 Innovation Drive, Dayton, OH 45342 Teradata.com

SQL-GR, SQL-H is a trademark, and QueryGrid, SQL-MapReduce, Teradata and the Teradata logo are registered trademarks of Teradata Corporation and/or its affiliates in the U.S. and worldwide. Intel, the Intel logo, and Xeon are registered trademarks of Intel Corporation. SUSE is a registered trademark of Novell, Inc. Teradata continually improves products as new technologies and components become available. Teradata, therefore, reserves the right to change specifications without prior notice. All features, functions, and operations described herein may not be marketed in all parts of the world. Consult your Teradata representative or Teradata.com for more information.

Copyright © 2015 by Teradata Corporation All Rights Reserved. Produced in U.S.A.

09.15 EB6434



TERADATA.