

TERADATA ASTER BIG ANALYTICS APPLIANCE: AN INDUSTRY FIRST

New generation big data technologies can often be quite innovative. Yet, if they are not well thought out, they can be disruptive. Moreover, to get the most from a big data solution, a lot of capital and resources may have to be put at risk. 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 analytics-optimized environment for rapid, on-the-fly data exploration, this is the first appliance in the industry to combine big analytics technologies, such as Hadoop, MapReduce, graph, pattern, and path analysis, with ecosystem compatibility with BI and ETL tools, business-friendly ANSI-standard SQL, out-of-the-box MapReduce analytics and mature, enterprise-grade system management solutions. It is a tightly coupled hardware and software solution that contains Aster Database and Apache Hadoop to process structured, unstructured, and semi-structured data. The appliance features key capabilities including:

- ~ A truly hybrid architecture that includes Aster Database, Aster SQL-MapReduce® and Apache Hadoop for a complete big analytics solution.
- ~ Industry's deepest stack integration, across the Aster and Apache Hadoop platforms. Enables business analysts to issue ANSI-standard SQL, via SQL-H™ and SQL-MapReduce®, on both Aster and Hadoop data.
- ~ Best-in-class enterprise big analytics via industry's only SQL-MapReduce® interface and the most comprehensive library of MapReduce analytic functions. The analytic library features more than 50 pre-built MapReduce functions, providing out-of-the-box analytics for graph, text, behavior, and marketing analytics, among others.
- ~ Industry-leading performance and scalability via optimized hardware and software configuration for Aster and Hadoop and 40 Gb/s InfiniBand networking. Proven 35x faster for interactive SQL and MapReduce analytics than alternative Apache Hadoop solutions.
- ~ Truly 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 Aster Big Analytics Appliance is a powerful, ready-to-run platform that is pre-configured and optimized specifically for big data storage and analysis. A purpose-built, integrated hardware and software solution for analytics at big data scale, the appliance runs the Aster patented SQL-MapReduce® and SQL-H™ technology on a time-tested, fully supported Teradata hardware platform. Depending on workload needs, it can be configured with Aster nodes exclusively, Hortonworks Data Platform (HDP) Hadoop nodes exclusively, or a mixture of Aster and Hadoop nodes. Additionally, integrated backup nodes are available for data protection.

By minimizing the number of moving parts required for deployment, the appliance offers easy, integrated management of an enterprise-ready information discovery solution with the benefits of optimized performance, continuous availability, and linear scalability. The result? You can bring information discovery 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 BIG INSIGHTS

The appliance comes with Aster Database version 5, which features more than 50 pre-packaged SQL-MapReduce functions to enable faster insights. The SQL-MapReduce framework, created by Teradata Aster, allows developers to write powerful and highly expressive SQL-MapReduce functions in languages such as Java, C#, Python, C++, and R, and push them into the discovery platform for advanced in-database analytics. Business analysts can then invoke SQL-MapReduce functions using standard SQL 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. Also included is the new SQL-H™ functionality, which makes Hadoop enterprise-ready. Optional Hadoop nodes can be used for staging, data transformation, and long-term data archival. Using SQL-H™, business analysts can easily access Hadoop data from within the Aster Database via standard SQL commands and BI tools. The result, through technologies like SQL-MapReduce® and SQL-H™ and tight metadata integration between Aster and Hadoop, business analysts for the first time, can seamlessly analyze structured and unstructured data from integrated and unified SQL-based query and analytics interfaces and their choice of BI tools.

In addition, the discovery platform extends and enhances the value delivered by your data warehouse by enabling it to process new data types and sources that are not typically stored in a data warehouse, and by leveraging that data already stored to enhance its analysis. Plus, the appliance lets you quickly find the high-value data and deep insights you need for strategic and operational intelligence. With the integrated, high-speed Aster-Teradata Adaptor, those insights can be easily integrated into your integrated data warehouse.

BUILT FOR EASE OF START-UP

Delivered as a fully-integrated system that's ready to plug-in, the appliance is purpose built for big data

analytics on multi-structured data. It features a complete Aster Database, including the patented Aster SQL-MapReduce framework, on a proven Teradata hardware platform with dual 8-core Intel® Sandy Bridge processors, SUSE® Linux operating system, and enterprise-class storage—all preinstalled into a power-efficient unit. That means you can have the system up and running live in just a few hours for rapid time to value.

DELIVERS RELIABLE, SCALABLE AND POWERFUL PROCESSING

Massively parallel processing architecture, embedded SQL, 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. Independent benchmarks have shown up to 400x and an average of 35x performance improvements on Aster for interactive SQL and MapReduce analytics compared to commodity Apache Hadoop solutions.

ENTERPRISE READY SOLUTION

The solution is easily accessible, delivering big data analysis, without requiring specialized software or hardware management skills in your data center or enterprise architecture. Existing Teradata customers will find the management consoles and capabilities familiar and complementary to their existing database and data warehouse deployments.

A suite of adaptors 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 included to enhance your enterprise solution.

Fast and Reliable Network Interconnect

A 40 Gb/s InfiniBand network connects the Aster and Hadoop nodes. This fast, dedicated and reliable network interconnect can also be used to connect multiple Aster Big Analytics Appliances and you can integrate with the Teradata Analytical Ecosystem via 10 GbE.

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 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 included:

~ 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 Aster and Hadoop.

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

The diagnostic information collected by TVI assists support personnel in identifying and quickly resolving problems; conversely, omission of this proactive software may prolong the time it takes to resolve incidents. In fact, internal Teradata studies indicate that it uncovers 62% to 70% of all system incidents.

~ **Integration with Teradata Viewpoint** offers a common management console for the Aster Database, Teradata Integrated Data Warehouse, and Apache Hadoop.

Teradata Viewpoint offers simpler, faster, and more comprehensive system management by providing a browser-based portal that delivers management intelligence to DBAs and users alike. Viewpoint 14.01 enables Aster portlets such as the Aster Node Monitor for Aster system monitoring. Apache Hadoop software monitoring will be available in Viewpoint 14.10.

- ~ **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.

In addition, Teradata offers a full array of support advantages, including Teradata Vital Infrastructure services which provide:

- ~ Industry certified regional and global **Customer Care Centers** with experienced service representatives who are available 24x7.
- ~ **End-to-end support** for the entire solution, from the hardware/OS to the Aster and Hadoop software.
- ~ **Secure remote connectivity options** to pinpoint problems and react rapidly.

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.

ASTER BIG ANALYTICS APPLIANCE DESCRIPTION

THE ASTER BIG ANALYTICS APPLIANCE FEATURES

- Aster Database 5 or higher version
 - SQL-H™: Allows data access in Hadoop through a standard SQL interface
 - Over 50 SQL-MapReduce® analytical functions
- Hortonworks HDP 1.1 or higher
- SUSE Linux 11 64-bit Operating System
- Integrated Cabinet
 - Dual Intel Eight Core Sandy Bridge Xeon® processors @ 2.6GHz per node (Aster Queen/Aster Worker/Hortonworks Hadoop Master)
 - Dual Intel Six Core Sandy Bridge Xeon® processors @ 2.0GHz per node (Aster Backup/Aster Loader/Hadoop Data)
 - 40 Gb/s InfiniBand Node Interconnect
- System Management Infrastructure, Teradata Administration, and Teradata Viewpoint
- Aster Teradata Adapter: high-speed data connector to the Teradata data warehouse
- RAID High Availability
- Disk Capacities – uncompressed (3x compression assumed typical)
 - Aster Worker Nodes with 900GB drives: 5.5TB
 - Data Nodes for Hadoop with 3TB drives: 9.5TB
 - Aster Backup Nodes with 3TB drives: 25.1TB
 - Full Aster Cabinet (2 Queens/16 Workers) – 88TB user space
 - Full Hadoop Cabinet (2 Masters/16 Data) – 152TB user space
- Scalability to petabytes with expansion cabinets and larger network switch

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 (41.9F), High end moisture: 60% RH and 15C DP (59F DP)

Electrical

- North America Voltage Range: 200 - 240V, 3~+PE
 - Current:
 - 30A, 3-p, 4-wire, 4-cord
 - 60A, 3-p, 4-wire, 2-cord
- International Voltage Range: 220 - 240 / 381 - 415, 3~ +N +PE with three phase
 - Current:
 - 32A/30A, 3-p, 5-wire, 2-cord, All continents (including North America with EU style power)
- Frequency: 50Hz – 60Hz
- Maximum Power: 12.25kW
- 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
- Software Implementation



10000 Innovation Drive Dayton, OH 45342

teradata.com


TERADATA.
**THE BEST
DECISION
POSSIBLE™**