

# Aster Database 5.0

## Deliver Powerful Big Data Analytics

The massive growth in data from new data types and sources is creating exploding volumes of multi-structured data that are not typically stored and processed in the enterprise data warehouse. At the same time, the need has emerged for analytic techniques to analyze this data quickly and efficiently, enabling iterative exploration at scale to find the high value data and deep insights locked inside. Taking full advantage of the value of this data requires a solution that simplifies and accelerates the development of rich analytic applications without requiring new skills and administrative resources.

### Aster Database Solution

Aster Database is a massively parallel software solution that embeds MapReduce analytic processing with data stores for big data analytics that incorporate new multi-structured data sources and types to deliver new analytic capabilities with breakthrough performance and scalability. Its unique Applications-Within<sup>®</sup> architecture runs analytic application logic inside the system, leveraging its massively parallel processing architecture and patented SQL-MapReduce<sup>®</sup> to fully parallelize processing for deep and ultra-fast analysis of massive data sets. Aster Database delivers easy management of data and analytic applications, continuous availability, and linear scalability. It can be deployed as a software-only solution on commodity hardware, as a cloud-based solution, or as part of Aster MapReduce Appliance on Teradata hardware. Combined with the visual development environment Aster Developer Express and a suite of optimized analytic modules — Aster MapReduce Analytics Portfolio — the analytic platform makes deep and rich analysis of big data that was previously difficult or impossible not only possible but also fast and easy.

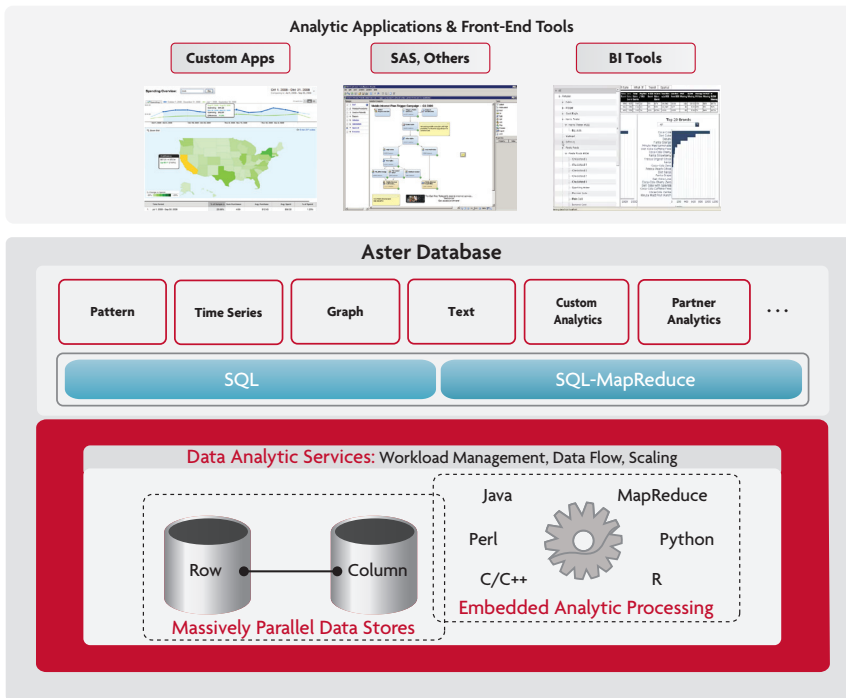


Figure 1: The massively parallel Teradata Aster MapReduce Platform enables analytics to be embedded with data for ultra-fast analysis.

### Breakthrough Performance and Scalability

Aster Database delivers leading performance and scalability for big data analytics. The results: 10x-1000x better performance than traditional systems and scalability to terabytes and petabytes of data.

- **“Always-Parallel” Pervasive Parallelism** – Aster Database’s massively parallel processing (MPP) architecture and embedded MapReduce analytics engine enable end-to-end parallelism of data and analytic processing. It executes loads, queries, exports, backups, recoveries, installs and upgrades in parallel to take full advantage of all resources, optimizing performance for analytic processing.
- **Embedded MapReduce** – Aster Database deeply embeds an implementation of the MapReduce processing framework, bringing the analytic power of MapReduce to developers of advanced analytics so they can easily

### Overview

Aster Database is a massively parallel software solution that embeds MapReduce analytic processing with data stores for big data analytics on new multi-structured data sources and types. It delivers new analytic capabilities for data scientists with breakthrough performance and scalability.

### Highlights

- Massively parallel processing (MPP) architecture for performance and scalability
- Patented SQL-MapReduce<sup>®</sup> framework embeds and parallelizes analytic applications to deliver ultra-fast, deep analytics on massive data volumes
- Suite of powerful building blocks for simplifying and accelerating development of analytics — Aster MapReduce Analytics Portfolio
- First visual integrated development environment for SQL-MapReduce analytic applications — Aster Developer Express
- Rich monitoring and easy management of data and analytic processing with the intuitive Aster Management Console
- Scalability to thousands of server cores and petabytes of active data
- Dynamic mixed workload management ensures scalable performance even with large numbers of concurrent users and workloads
- Fault-tolerant design and online management to minimize downtime
- Hybrid row and column stores with unified computation and management layers
- Next-generation loading architecture for leadership throughput to minimize time needed for loading data
- Deployed on certified commodity server hardware from Teradata partners, in the cloud, or on optimized Teradata hardware

deliver ultra-fast results on large data sets.

- **“Always-On” Fault Tolerance** – Massive-scale fault tolerance through replication, automatic failover, NIC bonding, failure heuristics, and clustered backup to prevent unplanned downtime due to hardware or software failures.
- **Hybrid Row and Column Stores** – Support for storing data in a row or column format for the highest levels of performance across diverse query workloads including ad hoc and interactive, investigative queries.
- **Dynamic Mixed Workload Management** – Advanced workload management allows for granular rule- and policy-based prioritization and dynamic resource allocation balance processing and compute resources for in-progress transactions, allowing administrators to adapt to changing priorities in real time.

## Next-Generation Data Science-based Analysis

Aster Database makes it easy to create advanced analytic applications that deliver ultra-fast analysis that scales to terabytes and petabytes of data. It allows analytic applications to be embedded within Aster Database with data for maximum performance, combines the analytic power of MapReduce with the familiarity of SQL, delivers the first visual development environment for SQL and MapReduce applications, and provides a suite of powerful analytic functions for accelerating development.

- **SQL-MapReduce** – Aster’s patented SQL-MapReduce framework makes it easy to leverage the power of MapReduce within the familiarity of SQL. SQL-MapReduce allows developers to write powerful and highly expressive functions in languages including Java, C, C++, C#, Python, and R and push them into the analytic platform where they can be called using standard SQL. SQL-MapReduce functions are simple to write and are seamlessly integrated with SQL statements.
- **Aster Developer Express** – Aster Developer Express makes it easy to develop, validate, and deploy advanced analytic applications with the first visual development environment for SQL-MapReduce applications. It enables visual development through integration with the Eclipse integrated development environment, provides a desktop testing environment for testing, and enables one-click push down of applications into Aster Database.
- **Aster MapReduce Analytics Portfolio** – This suite of powerful, reusable analytic functions accelerates the development of rich analytics with ready-to-use modules optimized to take advantage of the power of the SQL-MapReduce framework for fast analytics.

## Seamless Management and Connectivity to IT Ecosystems

Aster Database delivers rich visibility and control of not only data but also the analytic applications running inside the system. It provides intuitive tools for centralized management combined with powerful capabilities for simplifying and automating administration, streamlining and automating management of data and applications to minimize administrative work even as the system scales to 10s and 100s of servers.

- **Powerful Console for Monitoring and Managing Data and Applications** – Aster Management Console makes it easy 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. It also makes administration easy with single-click scaling and point-and-click access to workload management policies.
- **“Always-On” Online Maintenance** – Aster Database enables simultaneous load and export during queries, online backup and recovery, online restoration, and online scaling to avoid scheduled downtime.
- **Extensibility Framework** – Aster Management Console allows administrators to extend the capabilities accessible through the AMC with custom scripts and programs for common tasks.
- **Enterprise Adapters** – Aster-Teradata Adapter and Aster-Hadoop Adapter go beyond the capabilities of standard connectors for provide rapid and parallel data transfer between Aster Database and Teradata Database or Hadoop.

## About Aster Data

Teradata Aster MapReduce Platform is the market-leading big data analytics solution. Teradata Aster MapReduce Platform embeds MapReduce analytic processing for deeper insights on new data sources and multi-structured data types to deliver analytic capabilities with breakthrough performance and scalability. The solution utilizes Aster’s patented SQL-MapReduce® to parallelize the processing of data and applications and deliver rich analytic insights at scale.

## Technical Specifications

- **Server Hardware** – certified x86-based commodity server hardware from Teradata partners, or fully-supported Teradata hardware for appliance delivery
- **Operating Systems** – certified Linux-based operating systems
- **Drivers and APIs** – SQL, OLE DB, ADO.NET, ODBC, JDBC, Psycopg (Python)
- **SQL Standards** – ANSI SQL-92 compliant with SQL-99 and SQL-03 extensions
- **Business Intelligence and Data Integration** – compatible with leading business intelligence and data integration tools
- **Ecosystem Adapters** – Aster-Teradata Adapter and Aster-Hadoop Adapter for rapid and parallel data transfer between Aster Database and Teradata Database or Hadoop
- **Security** – compatible with Quest Authentication Services

*“Aster gives us the analytic capability to provide best-in-class digital marketing optimization for our clients, enabling more accurate marketing attribution. With Aster we can help our clients understand every marketing interaction with consumers over time and across their entire online market ecosystem, knowing the impact of every marketing dollar spent”*

*Sunil Kavi, Director of Technology,  
Business Intelligence  
Razorfish*

razorfish™