

Aster Data Analytic Foundation

Accelerate Development of Advanced Analytics

Faced with exploding data volumes and sources as well as continued growth in demand for analytic insights to inform critical business decisions, organizations are demanding analytic applications that can rapidly deliver deep insights and scale to massive data volumes. However, updating existing applications and implementing new applications typically require significant time and expertise for development, testing, and deployment. Optimizing them to ensure that they perform and scale to meet current and future needs further increases the time and complexity involved. As a result, analytics teams find it increasingly difficult and costly to meet critical demand for these advanced analytic applications.

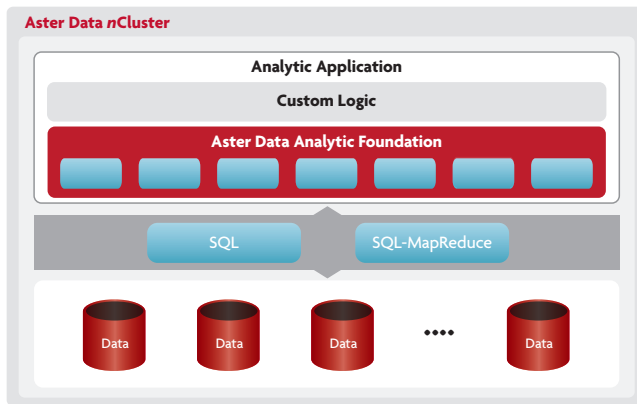


Figure 1: Aster Data Analytic Foundation provides building blocks for advanced analytic applications

The Aster Data Solution

Aster Data Analytic Foundation is a powerful suite of reusable MapReduce analytic functions that simplify and accelerate the development of advanced analytic applications. The functions in Analytic Foundation have been designed and optimized for high performance and scalability by leveraging the power of Aster Data's in-database MapReduce and SQL-MapReduce framework. These building blocks enable developers to simplify and accelerate deployment of analytic applications that deliver high performance and scalability with Aster Data nCluster - the first data-analytics server, a massively parallel (MPP) database with an integrated analytics engine.

Simplify and Accelerate Development

Aster Data Analytic Foundation provides a set of reusable building blocks that can be easily leveraged by custom analytic applications. Rather than spending time and effort on implementing, testing, and optimizing their own code for common analytic operations, developers can easily use the pre-built functions provided by Analytic Foundation within their applications through the SQL-MapReduce framework. As a result, they can focus on the higher-level logic that is unique to their applications and organization, speeding development and ensuring consistency across their applications.

Deliver Fast and Scalable Analytic Insights

The analytic functions included in Aster Data Analytic Foundation have been designed and optimized for high performance and scalability. They take full advantage of the Aster Data nCluster system, leveraging nCluster's in-database MapReduce implementation and the SQL-MapReduce framework to deliver high performance and scalability. These functions allow developers and analysts to create analytic applications that deliver rapid insights and scale to large data sizes with significantly less time and effort.

Highlights

- Powerful suite of reusable MapReduce analytics
- Easily incorporated into custom analytics through SQL-MapReduce
- Optimized for high performance and scalability
- Designed for parallel processing leveraging Aster Data nCluster's in-database MapReduce and SQL-MapReduce
- Easy to use with the Aster Data Developer Express visual development environment
- Includes functions for path analysis, statistical analysis, and relational analysis

“comScore sees significant value in the SQL-MapReduce functions that Aster Data provides. By using them to create analytics that run embedded within the database, we can enable our analysts to produce results that much faster. It reduces errors and makes the process a whole lot more efficient.”

Michael Brown, Chief Scientist
comScore



Enable Richer Analytics

Aster Data Analytic Foundation reduces the time and effort required to implement common analytic operations, enabling organizations to focus resources on developing analytic applications that provide deeper insight and more complex analysis of their data. By taking advantage of the performance and scalability of Aster Data *n*Cluster, Analytic Foundation also enables organizations to create applications that can analyze larger data sets and encompass more data sources than was previously practical. As a result, organizations can create analytic applications that deliver insights that were previously difficult or impractical to provide.

Aster Data Analytic Foundation Components

Aster Data Analytic Foundation provides a powerful set of MapReduce analytic functions that can be used as building blocks to create custom analytic applications. It includes analytic functions in the following areas:

- **Path Analysis:** Discover patterns in rows of sequential data for use in scenarios including time-series analysis, predictive analytics, and web analytics such as click-stream analysis. This group includes the following:
 - *n*Path: Perform complex sequential analysis for time series analysis and behavioral pattern analysis
 - Sessionization: Identify sessions from time series data in a single pass over the data
- **Statistical Analysis:** High-performance processing of common statistical calculations for use in a variety of applications including analysis of portfolios, market prices, consumer behavior, and security. This group includes the following:
 - Approximate percentiles and distinct counts: Calculate percentiles and counts within a specified variance
 - Correlation: Calculate the strength of the relationship between different columns of data
 - Regression: Perform linear or logistic regression between an output variable and a set of input variables
 - Averages: Calculate moving, weighted, exponential or volume-weighted averages over a window of data
- **Relational Analysis:** A collection of functions for discovering important relationships among data for use cases that include retail optimization, network analysis, and log file analysis. This group includes the following:
 - Basket analysis: Create configurable groupings of related items from transaction records in a single pass
 - Graph analysis: Find the shortest path from a distinct node to all other nodes in a graph
 - Tokenization: Split strings into words

About Aster Data

Aster Data is a proven leader in big data management and big data analysis for data-driven applications. Aster Data's *n*Cluster is the first MPP data warehouse architecture that allows applications to be fully embedded within the database engine to enable ultra-fast, deep analysis of massive data sets. Aster Data's unique "applications-within"™ approach allows application logic to exist and execute with the data itself. Termed a "Data-Analytics Server", Aster Data's solution effectively utilizes Aster Data's patent-pending SQL-MapReduce together with parallelized data processing and applications to address the big data challenge. Companies using Aster Data include Coremetrics, MySpace, comScore, Akamai, Full Tilt Poker, and ShareThis. Aster Data is headquartered in San Carlos, California and is backed by Sequoia Capital, JAFCO Ventures, IVP, and Cambrian Ventures, as well as industry visionaries including David Cheriton, Ron Conway, and Rajeev Motwani. For more information please visit www.asterdata.com, or call 1.888.Aster.Data.

Example Use Cases

- Analyze click streams to identify opportunities to optimize websites
- Perform statistical analysis on market prices in order to create and improve portfolio models
- Explore purchasing patterns to identify key cross-sell opportunities
- Process log files generated by software and systems
- Identify important connections between customers to find ways to improve customer retention