

Aster Data® and Dell

Proven Large-Scale Advanced Analytic Solutions

Aster Data and Dell have partnered to deliver performance-optimized solutions for large-scale data warehousing and advanced analytics. The solutions provide a compute-rich hardware platform of Dell PowerEdge™-C Series servers, ideally suited for the Aster Data nCluster® massively parallel (MPP) data warehouse platform which has been optimized for advanced analytics on very large data volumes. To enable advanced analytics on large data sets, Aster Data's solution allows analytic applications to be fully embedded within the database engine, allowing ultra-fast and deep analysis. Additionally, unlike other MPP platforms, Aster Data excels in making it easy and practical to build and run advanced analytics in a cloud environment.

Aster Data and Dell have been proven in some of the most demanding analytic environments. Aster Data's unique "applications-within®" approach allows analytic application logic to exist and execute with the data itself, providing breakthrough performance for data-driven companies including Intuit, MySpace, LinkedIn, Coremetrics, Akamai, ShareThis, and Full Tilt Poker.

The powerful Aster Data and Dell solutions for large-scale data warehousing and advanced analytics provide:

- **On-demand extreme scaling with no downtime**
- **Always-on data cloud availability**
- **High performance, next-generation analytics for big data**

On-Demand Extreme Scaling with No Downtime

Aster Data and Dell solutions allow organizations to scale quickly on-demand to meet ever-increasing data and analytic capacity and processing needs. Plug-and-play x86 hardware from Dell together with Aster Data's Online Precision Scaling™ provide easy incremental scale-out. This gives you fast access to the analytic power you need for mission-critical business operations. And, Aster Data nCluster enables incremental scaling with Dell PowerEdge-C servers while the system is live—without degradation of query performance—to provide:

- **Rapid data cloud expansion (and contraction) with Online Precision Scaling** – Provides linear scalability across loads, queries, and backups independently or in unison to meet requirements with no downtime. Granular splitting and load balancing of virtual partitions ensures maximum parallelism for massive "no limits" scalability. As capacity and processing requirements trail off, you can easily scale down to conserve resources and power.
- **On-demand incremental scaling by function** – Allows you to scale by any function (loading, processing, backups) based on Aster Data's multi-tier scaling architecture.

"Always-On" Data Cloud Availability

Aster Data's unique "always-on" architecture enables all key operations to run concurrently, providing uninterrupted query and analytics processing. Data and analytics scaling, backups, application processing, and all critical operations can be run in parallel providing Dell PowerEdge-C customers with 24x7 data and application availability.

- **"Always-On" Fault Tolerance** – Aster Data nCluster leads the industry in massive-scale fault tolerance with replication, automatic failover, NIC bonding, failure heuristics, and clustered backup to prevent unplanned downtime due to hardware or software failures.

Overview

Dell PE-C2100 was designed with big data in mind. It is the ideal platform for MPP architectures that combines a compact 2U form factor with powerful Intel Xeon processors, high performance memory system and expansive disk storage. This maximizes space, power and cost efficiency in data centers where memory and storage density are critical: MapReduce, Web analytics, database, and cloud computing.

Highlights

- Massively parallel processing hybrid row/column architecture for breakthrough performance and scalability
- Application processing services that allow applications to be fully parallelized inside the database
- Unlimited scalability to thousands of server cores and terabytes to petabytes of active data
- In-database analytics to enable ultra-fast, deep analytics on massive data using SQL and SQL-MapReduce
- Patent-pending SQL-MapReduce framework for powerful analytics
- Dynamic mixed workload management for thousands of concurrent users
- Powerful management tools for efficient administration
- "Always-on" fault-tolerant design and online management to minimize downtime
- Up-to-the-second loading of fresh data
- Faster analytic insights on cost-effective Dell PowerEdge-C servers

- **“Always-On” Online Administration** – *nCluster* enables simultaneous load and export during queries, online backup and recovery, online restoration, and online scaling to avoid scheduled downtime—saving you time and money.

High Performance, Next-Generation Analytics for Big Data

Aster Data’s platform uniquely enables high performance, advanced analytics on terabytes to petabytes of data. Key to this is Aster Data’s patent-pending SQL-MapReduce parallel processing framework and Aster Data’s Analytic Foundation—a suite of ready-to-use analytic functions. This powerful combination eases development of advanced analytic applications and delivers deep business insights in minutes to seconds for highly complex, diverse analytics. Dell customers can now easily migrate or deploy new advanced analytic applications on Aster Data’s massively parallel data warehouse platform powered by Dell PowerEdge C-Series servers. Key features include:

- **SQL-MapReduce** – Aster Data’s patent-pending SQL-MapReduce (SQL-MR) framework makes it easy to leverage the power of MapReduce within the familiarity of SQL. Developers write powerful and highly expressive functions in languages such as Java, C#, Python, C++, and R and push them into the database where they can be called using standard SQL. SQL-MR functions are simple to write and are seamlessly integrated within SQL statements.
- **Embedded Applications** – *nCluster*’s Applications-Within approach makes it simple to embed analytic applications within the *nCluster* database to eliminate the overhead of moving large data sets from the database to analytic applications. Enterprises can easily push existing applications into *nCluster* to see significant performance benefits.
- **Aster Data Developer Express** – Developer Express makes it easy to develop, validate, and deploy advanced analytic applications with the first visual development environment for SQL and MapReduce applications. It enables visual development through integration with the Eclipse integrated development environment, includes wizards to automate integration of applications with *nCluster*, provides a desktop testing environment for easy testing, and enables one-click push down of applications into *nCluster*.
- **Aster Data Analytic Foundation** – 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 Aster Data’s SQL-MR framework for fast analytics.

About Aster Data

Aster Data is a market leader in data management and advanced analytics for diverse and big data, enabling the powerful combination of cost-effective storage and ultra-fast analysis of relational and non-relational data. Aster Data *nCluster* is an analytic platform that incorporates a massively parallel processing (MPP) hybrid row and column database with an integrated analytics engine, allowing application logic to execute with data to deliver breakthrough performance and scalability. Aster Data’s solution utilizes Aster Data’s patent-pending SQL-MapReduce to parallelize processing of data and applications and deliver rich analytic insights at scale. Companies including Barnes & Noble, Intuit, LinkedIn, Akamai, Full Tilt Poker, and MySpace use Aster Data to deliver applications such as deep clickstream analysis, recommendation and personalization analytics, real-time fraud detection, and churn analysis.

Technical Specifications

- **Hardware** – Dell PowerEdge C2100 rack servers
- **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
- **Compatible with leading business intelligence tools**
- **Hadoop Connector** – Rapid and parallel data transfer between *nCluster* and Hadoop
- **Data Integration** – Leading tools including Informatica, GoldenGate, Microsoft SQL Server Integration Services (SSIS), Pentaho, Talend
- **Security** – Question Authentication Services