

Aster Database: Performance and Scalability

Massively parallel data-application-servers are designed to manage and process big data being generated in and around today's enterprises. Whether fed by data from hundreds of applications, thousands of ATMs, or millions of mobile phones, the massively parallel data-application server provides fast access to a persistent, comprehensive, and up-to-the-minute record of the enterprise's mission-critical data while providing a parallelization platform for ultra-fast analytic application execution.

Traditional MPP database and data warehouse systems do not scale to meet the requirements of today's analytics-intensive applications. Using traditional systems, it's often difficult to even load the data at its rate of generation, let alone analyze it fast enough to provide timely and meaningful insights.

Aster Database delivers the first analytic platform, a massively parallel (MPP) row and column database with an integrated analytics engine. Aster Database's analytics engine delivers a unified SQLMapReduce® computation layer on top of an architecture distinguished by a hybrid row and column DBMS for the highest level of performance. It 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. Teradata Aster's solution effectively uses Teradata Aster's patent-pending SQL-MapReduce together with parallelized data processing and applications to address the big data challenge.

Online Precision Scaling – Tiers of Discrete-Purpose Servers

The Aster Database architecture is unique in that it provides dedicated server groups for specific purposes. Separating workloads according to the nature of the task enhances performance by eliminating resource contention and enables independent scalability for different server groups:

- Queen servers are responsible for accepting user queries, globally-optimized query planning, coordinating query execution among Worker servers, and sending the final response back to users.
- Worker servers are responsible for storing data partitions, doing locally-optimized query planning, and executing queries according to the global plan generated by the Queen servers.
- Loader servers are responsible for processing incoming data and distributing it across Worker servers. They can also be used for data exports.
- Backup servers receive data from Worker servers, store backups, and send data to the Workers when recovery is required..




Figure 1: Teradata Aster's Online Precision Scaling enables on-demand incremental scaling of each functional server group.

Overview

Aster Database leverages its "Always Parallel" architecture to deliver high performance for all database operations such as queries, backup and recovery, data loading, etc. Aster Database has field-proven scalability to hundreds of servers running in a single cluster and can efficiently leverage the vast collective resources of such servers to deliver a very high level of performance.

Highlights

- Massive parallelization of all database operations across all servers and their processor cores provides high performance
- Dynamic workload manager for predictable performance
- Tiers of discrete-purpose servers provide independent scaling for specific tasks
- Network optimizations address the interserver network bottleneck issues to enhance performance and scalability
- High performance data storage techniques such as providing a hybrid row/column architecture and direct-attached storages for fast reads
- Large number of servers in the cluster provide large system memory for high performance

"Teradata Aster delivers superior scalability on commodity hardware – giving us a tremendous cost and performance advantage."

Shawn Farshchi, Chief Operating Officer
Coremetrics



While the separation of workloads provides performance benefits, it also enables independent scaling according to the workload requirements. For example, if high data loading rates are required, more Loader servers can be provisioned in the cluster. Even more, Online Partition Splitting enables granular splitting and load balancing of virtual worker (v-Worker) partitions maintains maximum parallelism across CPU-cores and servers for massive “no limits” scalability.

Massive Parallelization of All Database Operations

Aster Database is built on an “Always Parallel” data warehouse architecture that leverages the power of massive parallelism to provide high performance and scalability across all MPP database operations. All operations in Aster Database are automatically parallelized across servers of the cluster and across all processor cores of each server. Massive parallelization on tens to hundreds of servers and thousands of their processor cores results in tremendous performance gains and enables the system to scale efficiently. In particular:

- Queries are planned at a global level by the Queen servers and executed on Worker servers in a massively parallel manner.
- Data loads and exports are massively parallelized on Loader/Exporter servers.
- Backups are massively parallelized on Backup servers. Worker servers establish direct links with Backup servers and send massively parallel streams of data in a many-to-many fashion, reducing backup time.
- Similarly, restores from the Backup servers are also massively parallelized. Backup servers send massively parallel data streams directly to the Worker servers, reducing recovery time.
- Even the software installation and upgrade processes in Aster Database are massively parallel. During installation and upgrade, the Queen server automatically parallelizes installation/ upgrade of Aster Database software on other servers. This significantly reduces the time required for deployment and ongoing maintenance.

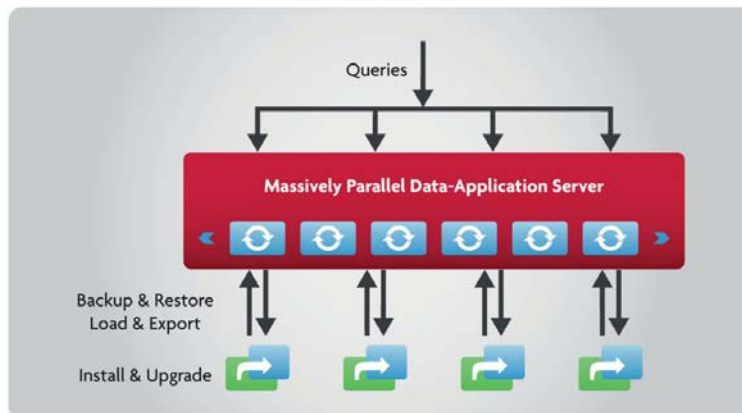


Figure 2: Teradata Aster’s “Always Parallel” architecture enables massive parallelization across all system operations.

Optimizations for Debottlenecking the Inter-Server Network

MPP (Massively Parallel Processing) databases use horizontal scaling to address the scalability limitations of the traditional SMP-based RDBMS architecture. But the full power of parallelization can only be realized when the most scarce resource – the inter-server network – is optimally used. In MPP databases, data is partitioned across servers at load time but a lot of inter-server network traffic is generated by queries involving joins, aggregations, etc. Such traffic limits the performance and scalability of other MPP databases. Aster Database includes patent-pending network optimizations to minimize such traffic and make efficient use of the network when such traffic is necessary. Aster Database database leverages these network optimizations to deliver high performance and scale to a very large number of servers without being slowed by the network bottleneck:

- **Network-Optimized Query Planning** – Aster Database's network-optimized global query planner uses innovative techniques to create execution plans that significantly reduce inter-server data shuffling required by queries involving multi-table joins, aggregations, etc.
- **Efficient Data Transport** – When data shuffling is necessary, Aster Database uses its Optimized Transport feature for network data compression and parallel streaming for manifold improvement in network performance.
- **Network-Optimized Data Loading and Export** – Separation of loads and exports to dedicated Loader and Exporter servers allows them to be placed strategically in the cluster such that the network traffic is minimized.
- **Network Aggregation** – Aster Database has the capability to leverage multiple network cables/ports in parallel for manifold increase in the network bandwidth for each server. Using the Network Aggregation feature, all links on a server can be seamlessly used by all partitions, providing true sharing of network bandwidth for all queries.

Dynamic Workload Manager for Predictable Performance

When hundreds or thousands of mixed workloads are executing simultaneously, it becomes increasingly difficult to prioritize and intelligently allocate the right amount of system resources to the right workloads at the right time. Traditional manual database tuning cannot possibly keep pace with rapidly changing workload demands. Using the Aster Dynamic Workload Manager, administrators can define rules that reallocate resources on the fly across hundreds of distributed nodes to adapt to new workloads and changing priorities in real time. The result is highly predictable performance and guaranteed service levels for the complex mixed workloads of an enterprise data warehouse and analytic-intensive applications.

- Fine-grained policy controls allow administrators to define pre-admission control and manage diverse workloads to meet the organization's business priorities.
- Workload management rules are easily created and managed using the Aster Management Console. Rules are written as easy-to-read SQL predicates, eliminating complex tuning.
- Dynamic reallocation of CPU and storage resources based on in-progress transactions ensures that time-critical queries can be processed immediately.

High-Performance Data Storage Strategies

Delivering high performance predictably and consistently so that analysts can focus on the content of their queries rather than response times, requires system optimization across a number of dimensions including the physical storage of data. Aster Database is a hybrid row and column DBMS, supporting equally both row and column physical storage with unified data access across both stores provided by SQL-MapReduce. Deployed on off-the-shelf commodity servers, Teradata Aster also leverages direct-attached storage on commodity servers for fast data access.

- Administrators can select row, column, or a combination of both storage techniques to optimize performance for their business' particular workloads.
- Aster Data Model Express is a recommendation tool which suggests the optimal data model, row, column, or a combination of both, based on actual query workloads.
- Use of direct-attached storage significantly improves I/O throughput – such servers in the cluster can collectively read data at very high rates. For example, a 100-server cluster can read data at $100 \times 400 \text{ MBps} = 40 \text{ gigabytes per second}$ (assuming 400MBps disk controller throughput). This is much higher than a costly SAN used by many other database systems, which can serve data only at a few gigabits per second.

Use of Direct-Attached Storage for Fast Data Access

Aster Database leverages off-the-shelf commodity servers that include direct-attached storage. Use of direct-attached storage significantly improves I/O throughput – such servers in the cluster can collectively read data at very high rates. For example, a 100-server cluster can read data at $100 \times 400\text{MBps} = 40$ gigabytes per second (assuming 400MBps disk controller throughput). This is much higher than a costly SAN used by many other database systems, which can serve data only at a few gigabits per second. The throughput increases with the number of servers and hence, Aster Database's market-leading scalability to hundreds of servers can provide extremely high data throughput.

Large System Memory

Similar to a large number of processor cores and high collective disk I/O, Aster Database servers can collectively provide a very large amount of system memory (RAM). The ability to hold large data volumes in memory provides a significant performance advantage for Aster Databases. As Aster Database servers are commodity servers, this memory footprint comes at a very affordable price. For example, a high-end SMP system may have 128 GB RAM with a very high price tag. On the other hand, an Aster Database analytic platform running on 100 servers would provide $100 \times 32\text{ GB} = 3.2\text{ TB}$ of RAM at a fraction of the SMP price. Because Aster Database can scale to a large number of servers, the MPP database can have access to a very large system memory and deliver the high performance required by today's analytically-intensive applications.

About Teradata Aster

The Teradata Aster MapReduce Platform is the market-leading big data analytics solution. This analytic 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. Teradata Aster's solution utilizes Aster's patented SQL-MapReduce® to parallelize the processing of data and applications and deliver rich analytic insights at scale.