



Microsoft® SQL Server® 2008 R2



MICROSOFT® SQL SERVER® 2008 R2 PARALLEL DATA WAREHOUSE

(PREVIOUSLY CODE NAMED PROJECT "MADISON")

Microsoft SQL Server 2008 R2 Parallel Data Warehouse is a highly scalable data warehouse appliance that delivers performance at low cost through a massively parallel processing (MPP) architecture.

Key Features

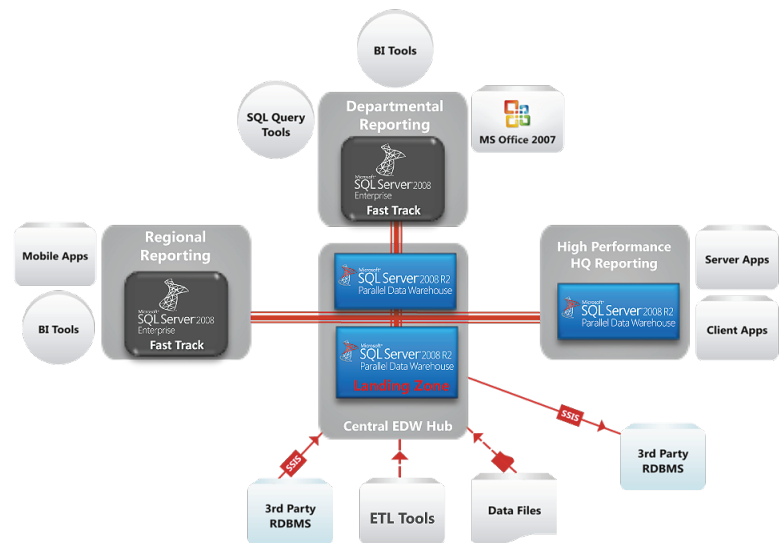
- Data warehouse scalability into the petabyte range.
- Low cost of ownership through industry-standard hardware.
- Simplified deployment and maintenance of appliance model.
- Integration with existing SQL Server 2008 data warehouses via hub-and-spoke architecture.
- Greater ROI from BI investments through integration with SQL Server 2008.
- Reduced risk through use of redundant, industry-standard hardware.
- Predictable performance delivery through balanced reference architectures.
- Better agility and business alignment through hub-and-spoke architecture.

Scale your data warehouse from terabytes to petabytes

Parallel Data Warehouse is a highly scalable appliance for enterprise data warehousing. It is the next step in the evolution of the data warehouse appliance created by DATAlegro.

Parallel Data Warehouse uses massively parallel processing (MPP) to deliver the high performance and scalability on SQL Server 2008, Windows Server® 2008, and industry-standard hardware.

The MPP architecture helps enables better scalability, better and more predictable performance, reduced risk, and a lower cost per terabyte than other data warehouse solutions.



A Parallel Data Warehouse MPP appliance acts as an enterprise hub that publishes data as needed to various business units or departments (spokes). For dedicated high-performance requirements, individual business units can redeploy SQL Server 2008 data marts, or deploy their own Parallel Data Warehouse appliances, as spokes.

Improve performance at a lower price per terabyte

In a traditional, symmetric multi-processing (SMP) architecture, query processing occurs entirely within one physical instance of a database. CPU, memory, and storage impose physical limits on speed and scale.

A Parallel Data Warehouse MPP appliance partitions large tables across multiple physical nodes, each node having dedicated CPU, memory, and storage, and each running its own instance of SQL Server, in a patented parallel design known as Ultra Shared Nothing. All components are balanced against each other to reduce performance bottlenecks, and all server and storage components are mirrored for enterprise-class redundancy.

A Control Node routes queries from applications to all Compute Nodes, then collects and returns the result. Because data is evenly distributed across multiple nodes

and processing occurs in parallel, queries can run many times faster than on single SMP database servers.

This architecture can cost less because, rather than relying on expensive proprietary processors or storage, Parallel Data Warehouse appliances use industry-standard hardware. As data volumes grow, scalability simply requires the addition of capacity to the appliance. There is no need for a “forklift upgrade,” where the entire appliance must be upgraded.

Improve ROI of existing BI investments

Extract more value from existing BI investments through tight integration of Parallel Data Warehouse and Microsoft BI tools. The Parallel Data Warehouse MPP appliance integrates with SQL Server technologies including Integration Services, Reporting Services, and Analysis Services. Existing SMP data marts can be redeployed and “plugged in” to the MPP cluster as spokes, receiving data from the MPP hub.

Benefit from an appliance model

Simplify data warehouse deployment and maintenance. Parallel Data Warehouse appliances include carefully balanced, pre-assembled hardware and software from Parallel Data Warehouse leading hardware vendors. This “appliance model” significantly accelerates your time to value and reduces deployment costs.

Deliver predictable performance

Through carefully balanced configuration, tight specifications, and testing, the Parallel Data Warehouse MPP appliance is designed to deliver predictable performance. Microsoft designs the reference configurations to suit different usage scenarios, including reporting and ad hoc queries.

Automated workload management and system resource balancing enable you to load data concurrently with queries, without sacrificing performance. As data grows and business needs change, Microsoft offers options to support more users performing different types of queries.

Gain better agility and business alignment through hub-and-spoke

Using a true hub-and-spoke architecture, all enterprise data can be maintained on a Parallel Data Warehouse hub while departments or business units keep their existing data marts to suit their needs. High-speed data transfer relieves traditional barriers to hub-and-spoke. Power users can even deploy a dedicated MPP appliance as a spoke so they can autonomously manage resources, while IT can enforce enterprise standards across all data.

Reap additional benefits

Automated storage management replaces complex space allocation. Available diagnostics detect and broadcast hardware issues. A management console provides a simple GUI interface and dashboard for tracking and resolving issues. Ultra-fast loading and high-speed backups reduce administration for even the largest data warehouses.

USE A PARALLEL DATA WAREHOUSE MPP APPLIANCE:

- As an economical, high-performance enterprise data warehouse or data mart.
- As a platform for high-scale data warehouses from terabytes to more than 1 petabyte.
- As the hub of a hub-and-spoke EDW.
- To offload applications from an EDW for high-speed aggregation and queries.
- As long-term data storage in multi-tier data warehouses.
- As a subject-oriented data mart or sandbox for new analytic applications.

GET ADDITIONAL INFORMATION

For more information about Parallel Data Warehouse, visit:

www.microsoft.com/madison
