

Microsoft®

Microsoft SQL Server Project code-named “Madison”



Microsoft® SQL Server® Project code-named “Madison” is a highly scalable data warehouse appliance that delivers performance at low cost through a massively parallel processing (MPP) architecture.

Scale Your Data Warehouse from Terabytes to Petabytes

“Madison” 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.

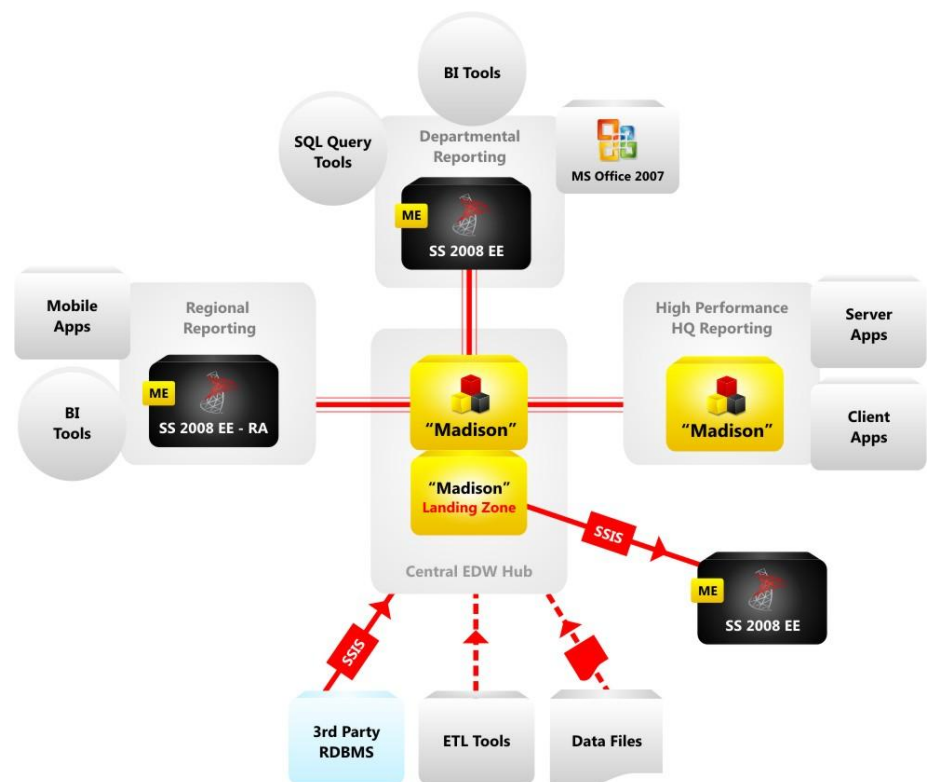
Madison uses massively parallel processing (MPP) to deliver the

high performance and scalability on SQL Server 2008, Windows Server® 2008 and industry-standard hardware.

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

KEY FEATURES

- Data warehouse scalability into the petabyte range
- Low cost of ownership through industry standard hardware
- Appliance model simplifies deployment and maintenance
- Integrates 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
- Balanced reference architectures deliver predictable performance
- Better agility and business alignment through hub-and-spoke architecture



A Madison 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 Madison 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 upon speed and scale.

A Madison MPP data warehouse 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, Madison 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 Madison and Microsoft BI tools. The Madison MPP data warehouse 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.”

Appliance Model

Simplify data warehouse deployment and maintenance. Madison appliances include carefully balanced, pre-assembled hardware and software from 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 Madison MPP data warehouse 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.

Better Agility and Business Alignment through Hub-and-Spoke

Using a true hub-and-spoke architecture, all enterprise data can be maintained on a Madison 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.

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 Madison MPP data warehouse appliance:

- As an economical, high-performance enterprise data warehouse or data mart
- As a platform for high scale Data Warehouses from terabytes to over 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

Additional Information

For more information about Project “Madison”, visit <http://www.microsoft.com/madison>

This data sheet is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS SUMMARY.

0308 Part No. 098-00000 (color)
Part No. 098-00000 (black and white)

Microsoft