



Microsoft® SQL Server® 2008

SQL Server® Fast Track Data Warehouse for DELL



Microsoft® SQL Server® Fast Track Data Warehouse for Dell servers take the unknowns out of hardware selection and configuration for SQL Server data warehouse solutions. The new reference architectures provide server and storage guidance for various data warehouse workloads – designed to give you the most efficient hardware for your solution, saving you time and cost in choosing the right technology for your business needs.

KEY FEATURES

- Hardware components tuned for data warehouse solutions and sized for your identified workload
- Scale from 4 up to 24 terabytes using compression capabilities in SQL Server 2008 Enterprise
- Enjoy a low cost of ownership through great price performance, rapid deployment and industry-standard hardware
- Tested on proven Dell PowerEdge servers:
 - Dell™ PowerEdge™ 2950
 - Dell™ PowerEdge™ R900
- Storage configurations available with affordable and scalable EMC CX4-240 SAN technology.
- Balanced hardware approach ensures that all components are equally utilized from the CPU cores to the physical drives

Scale up your data warehouse using a Dell hardware platform configured with the new SQL Server Fast Track Data Warehouse offering. Help reduce your data warehousing cost while improving out-of-the-box performance.

The Dell reference configurations bring together the right mix of technology and software, integrating the powerful Dell PowerEdge server line, the robust EMC storage array and the data warehouse capabilities of SQL Server 2008 Enterprise.

Optimized for Data Warehouse Workloads

Many hardware configurations are ineffective at responding to data warehouse queries because they are tuned for transactional systems. The Microsoft® and Dell Fast Track architectures are designed to gain great throughput and scale by using the right approach:

- Targeting query workloads patterned for large sequential data sets rather than small random data transactions
- Optimizing rapid data reads and query aggregations



Balanced Hardware Approach

The Microsoft and Dell reference architectures begin by taking your query response time requirements and calculating the number of CPU cores necessary to achieve that throughput. The I/O channel and storage arrays are matched to help maximize the CPU throughput, taking advantage of all the hardware components in equilibrium, without under-utilizing or over-burdening any one component.

Core Consumption Rate and Estimated CPU Requirements

The Fast Track Data Warehouse configurations on the Dell platform are driven by the CPU core consumption rate, which is the data throughput that a single CPU core can handle based on the query load. The Dell PowerEdge servers with the 2.66 GHz Intel Xeon quad-core and six-core processor have a CPU core consumption rate of approximately 200 MB/s throughput based on data warehouse workloads. This allows competitive query response times on lower-cost commodity servers, a fraction of the cost of big iron servers.

Once the consumption rate is calculated, the number of cores can be estimated by factoring the concurrent query data volume divided by the total consumption rate divided by the target response time in seconds.

Storage Configuration

The EMC CX4-240 storage allows dual reads when drives are mirrored. For sequential data reads from data warehouse queries, this capability enables tremendous throughput per storage volume—up to 240 MB/s. For example, with only 8 pairs of drives, the storage array can achieve almost 2 GB/s when the solution is optimized for data warehouse workloads. When a storage array is optimized for transactional activity, but the system supports data warehouse workloads, you would need up to 2-3 times the number of drives to achieve the same throughput.

The Dell PowerEdge servers are configured with EMC storage arrays, multiple Fiber Channel controllers and a Fiber Channel switch to load balance the connections to the storage processors in the following storage system.

- EMC CX4-240
 - 300 GB FC 15K RPM
 - 16 to 64 drives dedicated to user data
 - 4 TB to 16 TB optimized
 - Up to 32 TB of storage

Powered by Dell PowerEdge Servers

The reference configurations are built on two Dell PowerEdge server platforms, each targeting a different tier of a SQL Server data warehouse solution:

- Dell PowerEdge 2950
 - (2) Intel Xeon, quad core, 2.66 GHz
 - 8 total CPU cores
 - 32GB 667MHz, 8X4GB, Quad Ranked DIMMS
 - 4 TB optimized storage (8 TB max)
 - 1600 MB/s tested throughput

The Dell PowerEdge 2950 is ideal for data marts with a smaller departmental footprint of query activity with a volume optimized at 4 TB of compressed data.

- Dell PowerEdge R900
 - (4) Intel Xeon, six core, 2.66 GHz
 - 24 total CPU cores
 - 96GB Memory, 667MHz, 24X4GB
 - 12 TB optimized storage (24 TB max)
 - 4800 MB/s tested throughput

The Dell PowerEdge R900, with four six-core processors, is the top scale-up model in the PowerEdge server line, capable of handling higher concurrency and query complexity and optimized at 12 TB of compressed storage.

Additional Information

For more information about the SQL Server Fast Track Data Warehouse, visit

<http://www.microsoft.com/fasttrack>.