Microsoft® SQL Server® 2012

SQL Server 2012 Performance White Paper

Published: April 2012

Applies to: SQL Server 2012

Copyright

The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This document is provided "as-is." Information and views expressed in this document, including URL and other Internet Web site references, may change without notice. You bear the risk of using it.

This white paper is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED, OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT.

This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2011-2012 Microsoft Corporation. All rights reserved.

Microsoft, SQL Server, Windows Server, SharePoint, Excel are trademarks of the Microsoft group of companies.

All other trademarks are property of their respective owners.

Table of Contents

Introduction

Organizations today are facing an explosion in information volumes and data types. Gartner estimates that the total worldwide volume of data is growing at a rate of 59% per year—and that 70-85% of data is unstructured[1]. This represents a dramatic change in the amount and type of data that is consumed. To be productive today, organizations expect to access, analyze and share a wide range of data types at near instantaneous speeds. This is especially true as end users are increasingly accustomed to accessing and sharing unstructured data in real-time from nearly any source, including the Web, social media, and mobile devices. To meet this demand, a new level of performance and scalability is required of database systems, which must process very large data sets quickly.

In this paper, we will cover the performance and scalability capabilities of SQL Server 2012. SQL Server 2008 delivered major enhancements to performance and scale and the new technologies in SQL Server 2012 take performance a generational leap forward, delivering a whole new level of performance, alongside a range of tools that help make that performance more consistent and predictable. We will closely examine how SQL Server 2012 helps customers manage the massive growth in data volume and complexity through major enhancements for predictable and next-generation performance.

Generational Performance Leap

In SQL Server 2012, Microsoft introduces xVelocity in-memory technologies that deliver a generational leap in performance for data warehousing and business intelligence workloads. xVelocity leverages major advances in performance technology to deliver in-memory column store index, vector-based query execution, and aggressive compression, that fundamentally change the level of performance that can be achieved for data warehousing and business intelligence scenarios.

Figure 1: The graph compares elapsed time of nine complex DW queries measured before and after adding column store indexes to the fact tables.
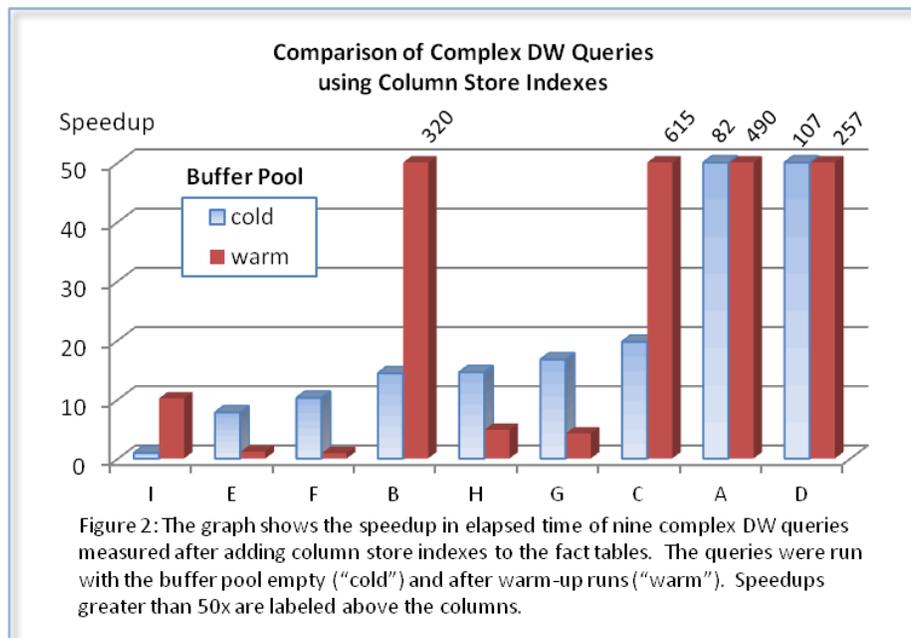
xVelocity for Data Warehousing

xVelocity in-memory column store index is a new technology in SQL Server 2012 that enables massive

---
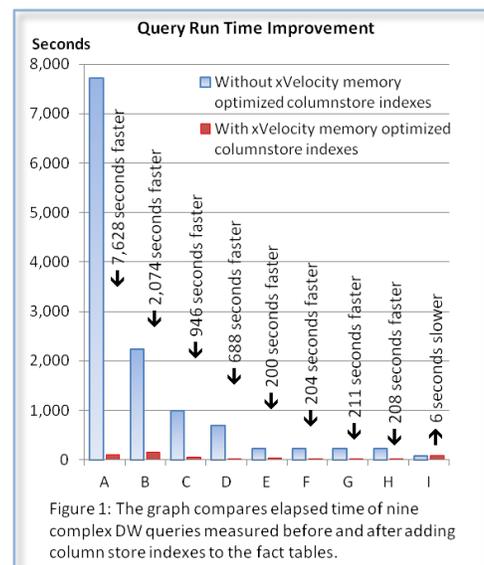
[1] Source: Gartner Symposium Presentation 'Information Management Goes 'Extreme': The Biggest Challenges for 21st-Century CIOs, Mark Beyer, October 2011

gains in performance for data warehousing workloads. Column store indexes group and store data for each column and then join all the columns to complete the whole index. It builds a new view on top of existing tables and fetches only those columns needed, which requires less I/O, reduces CPU usage, and improves memory utilization. xVelocity in-memory column store index has proven to increase query performance by 10-100x[2] for star join and similar queries and can routinely give a tenfold speedup for a broad range of decision support queries.

SQL Server is the first and only major database with this next-generation performance optimization. SQL Server column store index technology is especially appropriate for typical data warehousing data sets—enabling significantly improved performance for common data warehousing queries such as filtering, aggregating, grouping, and star-join queries. The column store index can also help to reduce development costs and ETL times because it limits or eliminates the need to rely on pre-built aggregates, including user-defined summary tables, and indexed (materialized) views.

Figure 2: The graph shows the speedup in elapsed time of nine complex DW queries measured after adding column store indexes to the fact tables. The queries were run with the buffer pool empty ("cold") and after warm-up runs ("warm"). Speedups greater than 50x are labeled above the columns.

The new vector-based query execution in SQL Server 2012 addresses a range of modern-day performance realities by reducing overhead, taking advantage of larger cache sizes, and limiting the number of CPU instructions. Vector-based queries are part of batch mode processing, which uses a new iterator model for processing data a-batch-at-a-time instead of a-row-at-a-time. Each column within a batch is stored as a vector in a separate area of memory. Batch mode processing uses algorithms that are optimized for the multicore CPUs and increased memory throughput that are found on modern hardware. Batch mode reduces overhead from metadata

---

[2] http://www.microsoft.com/sqlserver/en/us/product-info/benchmarks.aspx

SQL Server 2012 Performance White Paper

and other sources and operates on compressed data when possible. The result is better parallelism and faster performance for data warehousing workloads.

xVelocity for Business Intelligence

In SQL Server 2012, Microsoft is introducing xVelocity in-memory analytics engine on our standalone Analysis Services models, which offers major performance gains for business intelligence applications. This technology is based on the same engine behind PowerPivot, which was introduced by Microsoft in SQL Server 2008 R2. In Power Pivot, this in-memory analytics engine powered Excel to process hundreds of millions of rows with sub-second response time on nothing more than an end user's desktop PC. It leveraged aggressive compression and column-oriented storage to achieve high performance on personal desktops and enabled a fluid model development experience.

In SQL Server 2012, xVelocity in-memory analytics engine takes the powerful technology that was developed for the desktop and applies it to the server, delivering a generational leap in performance for BI projects in the enterprise. The technology enables scan rates of 10s of billions of rows per second on typical industry hardware, and it delivers compression rates of up to 300x in Microsoft tests.

The xVelocity in-memory analytics engine also benefits the IT department because it does not require manual hand query optimization, and it provides fine-grained security and data management capabilities, reducing the time to manage and administer the solution.

Predictable Performance

SQL Server 2012 offers capabilities built-in that help make performance more consistent and predictable. With enhancements to Resource Governor, you can better monitor and control the way different workloads utilize CPU and memory resources on SQL Server instances in a hosting or private cloud environment. SQL Server Management Studio provides familiar management tools for proactive troubleshooting and maintenance. In addition, new performance and scale enhancements are built into the SQL Server 2012 database engine, enabling improved support for the largest workloads.

Consistent Performance, Even for Multitenant Environments

Resource Governor Enhancements

With the increasing trend toward server consolidation and multi-tenant environments found in private and public cloud environments, organizations need a way to ensure consistent performance across workloads and applications that share the same server. Managing resource

allocation in these scenarios can be complex. Dynamic Management Views in Resource Governor enable you to monitor and manage CPU and memory utilization in multi-tenant environments—and even provide information and controls that can be helpful in chargeback scenarios. Resource limits can also be reconfigured in real time with minimal impact on workloads that are executing.

Consistent Performance for Multiple Workloads

In an environment where multiple distinct workloads are present on the same server, Resource Governor enables you to differentiate these workloads and allocate shared resources as they are requested based on the limits that you specify. This helps to ensure consistent performance for concurrent and mixed workloads.

Increased Resource Pools

A resource pool in the Resource Governor represents the physical resources of the server, and each resource pool can contain one or more workload groups. In SQL Server 2012, Resource Governor has increased the number of resource pools to 64, benefitting multi-tenant environments where resources need to be managed across a large number of distinct workload groups.

Enabling a Hard Cap on Resource Usage

In SQL Server 2012, Resource Governor enables you to create a hard cap on CPU resource usage. This is an important capability because it enables you to not only assign and monitor resources, but also to manage and enforce restrictions on resource usage.

Affinitized Resource Pools

Resource Governor enables you to affinitize your resource pools for CPU schedulers, groups of schedulers, and NUMA nodes.

Proactive Management Tools in a Single Console

SQL Server 2012 delivers new capabilities that help you perform proactive troubleshooting and administration with tools built into SQL Server Management Studio, such as Performance Data Collector, Policy-Based Management, and SQL Server Profiler. SQL Server Management Studio provides an integrated environment for accessing, configuring, managing, administering, and developing all components of SQL Server. It provides a single console for efficient management using a range of capabilities and familiar tools.

Data Collector

The Data Collector in SQL Server 2012 provides a central point for data collection across your database servers and applications. The Data Collector is a system that can collect data from any T-SQL query, like DMVs, Windows Performance Counters, and the SQL Trace, and it bundles any selection of these items into a Collection Set, which can be configured to meet your requirements.

Policy-Based Management

Policy-Based Management enables you to define and enforce common, enterprise-wide policies for the SQL Server instances you manage throughout the organization. With this capability, you can create policies for the management of entities on a server, such as the instance of SQL Server, databases, or other SQL Server objects. Policy-based management can help to enable a number of useful scenarios, such as imposing the surface area configuration settings of one database instance to another, creating and enforcing naming convention policy, and supporting scalability by exporting and importing policies or by applying policies to a server group.

SQL Server Profiler

SQL Server Profiler gives you the ability to monitor and capture server events for real-time diagnosis. It does this by showing how SQL Server resolves queries internally, which enables you to see exactly what Transact-SQL statements or Multi-Dimensional Expressions are submitted to the server, and then how the server accesses the database or cube to return result sets. You can configure SQL Server Profiler to monitor only the events of interest to you, and you can filter traces based on the information you want, so that only a subset of the event data is collected. This is helpful when traces become too large. Monitoring too many events adds overhead to the server and the monitoring process, and can cause the trace file or trace table to grow very large, especially when the monitoring process takes place over a long period of time.

Database Engine Tuning Advisor

SQL Server Database Engine Tuning Advisor (DTA) helps to improve query processing by selecting and creating an optimal set of indexes, indexed views, and partitions without requiring an expert understanding of the structure of the database. Database Engine Tuning Advisor analyzes a workload and the physical implementation of one or more databases, and it can recommend adding, removing, or modifying physical design structures in your databases. The advisor can also recommend what statistics should be collected to back up physical design structures, such as clustered indexes, non-clustered indexes, indexed views, and partitioning.

Multi-Server Management

SQL Server 2012 includes multi-server management capabilities that improve insight into resource utilization to help identify potential issues and make proactive changes before

problems occur. SQL Server Management Studio provides dashboard and drilldown utilization views of instances and applications, and it enables quick set up and instance enrollment into a central control point. Policy evaluation and trending analysis of instances and applications enable you to gain the insight needed to ensure that resources are fully optimized.

Speed and Scale Enhancements, Built into the Engine

SQL Server 2012 provides new enhancements to performance and scale that are built into the engine. Support for up to 15,000 partitions enables large sliding window scenarios that make it easier to perform large scale data management. Enhancements to Full-Text Search help make it faster to query information across unstructured data. A range of additional enhancements offer a major boost to performance and scale, including statistical semantic search, enterprise-grade compression capabilities, and support for up to 640 cores and 4TB of memory.

Support for up to 15,000 Partitions

SQL Server 2012 offers expanded partition support from 1k to 15k, which enables large sliding windows scenarios. For applications like SAP, that take tens of thousands of snapshots of data in short periods of time, the length of time where data is held before it's pushed out to allow for new data to enter can be significantly extended. This makes managing large amounts of data easier, and it helps to maintain large sets within filegroups that need data switched in and out to meet the needs of the data warehouse.

Enhancements to Full-Text Search

Full-Text Search (FTS) provides powerful search capability within and across unstructured documents. Full-text queries perform linguistic searches of words and phrases based on the rules of a particular language. FTS underwent major performance improvements for 2012 with improved query execution and concurrent index update mechanisms. Additionally, layered on top of FTS is statistical semantic search, which uses algorithms in SQL Server to deliver results with high statistical relevance, based on search inquiries. SQL Server delivers this capability built right into the engine for the highest performance across large volumes of complex data types.

Enterprise-Grade Compression Capabilities

SQL Server 2012 continues to deliver enterprise-grade compression capabilities that can cut growing data volumes approximately 50-60%[3]. This can help to dramatically speed performance for I/O intensive workloads because the data is stored in fewer pages, which means that queries need to read fewer pages from disk.

---

[3] http://msdn.microsoft.com/en-us/library/bb964719(v=SQL.100).aspx

Increased Processing and Memory Support

Another major boost in performance in SQL Server 2012 can be gained with the ability to scale up to 640 cores and 4TB of memory. This helps provide the scalability and high-performance you need for very large workloads or consolidation scenarios using some of the industry's most powerful hardware.

Conclusion

SQL Server 2012 delivers built-in capabilities for predictable and next-generation performance to address today's most demanding applications and business requirements. xVelocity in-memory technologies enables a generational leap in performance with sub-second query responses, rapid column-oriented processing, and faster time to business insight. Meanwhile, enhancements to Resource Governor help enable improved visibility and control across workloads for multi-tenant environments. Powerful management tools, delivered in a familiar consolidated interface, help to proactively ensure consistent and predictable performance. A range of additional performance and scalability enhancements are built into the SQL Server database engine, including increased partition support, Full-Text Search improvements, and greater processing and memory capacity. Together, these capabilities help enable SQL Server 2012 to deliver a generational leap in performance, helping organizations to better manage massive volumes of data and information, and empowering information workers to gain new insights at the speed of thought.