



Business Intelligence Solutions on the IBM System x3850 X5



Redguides
for Business Leaders

Michael Lawson
Vinay Kulkarni



- Third paper in the series on Business Intelligence solutions from IBM System x and IBM System Storage
- Based on the IBM System x3850 X5 with Microsoft SQL Server 2008 R2 BI for top-notch decision making
- Presents IBM Business Intelligence reference configurations for small-sized to large-sized businesses



Executive overview

Business Intelligence (BI) is defined as the technology used to extract useful information from the vast amount of data that a business has been collecting over time. The information extracted can then be used by decision makers to guide their business to success.

IBM® builds, tests, and publishes reference configurations to provide clients with guidelines when sizing their BI configurations, and with the release of the new eX5 servers with more processors and memory than ever before in an x86 platform, this IBM Redguide™ shows you how you can get the most out of one of these high-performance servers in a BI configuration. The eX5 servers allow you to scale to new highs in performance, process your cubes fast, and get query results even faster.

IBM has tested and is releasing new configurations built using these key components:

- ▶ IBM System x3850 X5
- ▶ IBM System Storage™ DS5000
- ▶ Microsoft® Windows® Server 2008 R2
- ▶ Microsoft SQL Server 2008 R2

This Redguide is the third in a series that describes Business Intelligence reference configurations on 4-socket IBM x3850 X5 running Intel® Xeon 7500 series processors. The first two Redguides™ covered reference configurations on IBM x3850 M2 and Microsoft Windows Server 2008 and SQL Server 2008.

This document provides sizing information for small, medium, and large businesses that plan to use the Microsoft Business Intelligence solution utilizing IBM System x3850 X5 server, IBM System Storage DS5000 storage, and Microsoft Windows Server 2008 R2 and SQL Server 2008 R2. Three reference configurations were investigated, as listed in Table 1.

Table 1 Business Intelligence reference configurations with results

	Small	Medium	Large
DW database size	1.8 TB	1.8 TB	3.7 TB
Server and storage	32 cores 64 GB memory 1 DS5300 56 disks	64 cores 128 GB memory 1 DS5300 56 disks	64 cores 256 GB memory 1 DS5300 101 disks
Cube building time	6 hours, 36 minutes	3 hours, 42 minutes	6 hours, 24 minutes
% of 1X	100%	56%	97%

Table 1 provides a high-level summary of the reference configurations discussed in this document and the processing results. These results show excellent scaling when processing the same sized data warehouse (DW) with more powerful hardware (going from the small to medium configuration) or when processing double the DW size (going from small to large). Essentially, doubling the hardware cuts the time in half and processing double the DW size can be done in the same amount of time as the original DW using more powerful hardware.

IBM offers the IBM System x3850 X5 server, a unique and compelling high-end modular server based on the Intel Xeon® 7500 series processor and running Windows Server 2008 R2 Enterprise. This server is a modular *building-block* server. The x3850 X5 can be configured with one or two chassis with an additional one or two 1U MAX5 memory drawer to build a scale-up server with up to 128 logical processors and 3 TB of memory.

This document reports on three scenarios in the small, medium, and large range, based on data warehouse size. Comparable runs to the small and large configurations in this Redguide with a previous Redguide took approximately 11 hours using the x3850 M2, so the X5 model provides an improvement of 60% on this real-world workload.

In conclusion, we show that the new eX5 servers allow processing your data and deriving intelligent information from it faster than you could ever do it before.

Introduction

The purpose of Business Intelligence is to take a company's historical and operational data, process it with analytics software, and present the data in an easy-to-read and familiar format to enterprise users. The resulting data empowers users by aiding them in making better business decisions to help improve their company's bottom line. Examples of BI are:

- Using the retail sales trends by geographical areas over the past few years to decide which products go on the shelf
- Using the results of various sales promotions offered historically to decide on future promotions

IBM builds reference configurations and tests them to provide clients with guidelines when sizing their BI configurations. IBM uses a workload developed by Microsoft using their internal proprietary data. The workload used for testing the configurations listed in this document includes BI cube building. Cube building is a processing step in the BI workflow that

pre-aggregates the data warehouse data in a form suitable for fast retrieval when users are performing business analytics. Building a cube can take many hours and is dependent on the hardware resources provided and the specifications of the cubes themselves (for example, the amount and complexity of aggregation).

IT professionals must know approximately how long their cubes will take to build in order to plan effectively and have the cubes ready when their customers need them. To assist with this task, this document provides three examples of data warehouses and hardware configurations of increasing size. See “Data warehouse profile” on page 9 when attempting to compare our workload with yours. The reference configurations have been tested by measuring how long it takes to build the cube.

The advantages of the reference configurations that we describe in this document are:

- ▶ Configurations are tested using the Microsoft BI software stack, providing assurance that the combined solution of software, server hardware, and data storage will provide the best return on investment in hardware and people resources.
- ▶ There are three sizes of hardware configurations—small, medium, large—to cover a relevant range of possible workloads.

The remainder of this Redguide is made up of three sections:

- ▶ “Solution architecture” on page 3
- ▶ “Workload description” on page 8
- ▶ “Reference configurations” on page 10

Solution architecture

Solution components include the high-end IBM System x3850 X5, IBM System Storage DS5300, and Microsoft Windows Server 2008 R2 and Microsoft SQL Server 2008 R2. This section introduces each of these components.

IBM System x3850 X5 server

The IBM System x3850 X5 server is the fifth generation of the Enterprise X- Architecture, delivering innovation with enhanced reliability and availability features to enable optimal performance for databases, enterprise applications, and virtualized environments. Figure 1 shows the IBM System x3850 X5.



Figure 1 IBM System x3850 X5

Features of the x3850 X5 include:

- ▶ Increased performance with more processors and memory than ever before on x86
- ▶ Intel Xeon X7500 series processors with up to eight cores and HyperThreading
- ▶ Up to eight sockets and 128 DIMMs with QPI scaling for larger databases, enterprise, and mission-critical workloads
- ▶ Up to 64 DIMM slots per chassis, delivering up to 1024 GB of high speed PC3-10600 double data rate (DDR3) memory
- ▶ Memory reliability and availability with memory ProteXion with Chipkill, memory mirroring, memory sparing, Intel SMI Lane Failover, SMI packet retry, and SMI Clock failover
- ▶ Low-power, cost-effective memory with Advanced Buffer eXecution chip
- ▶ High-performing databases and fast time to value for database workloads
- ▶ Advanced networking capabilities with Emulex 10Gb Ethernet Adapter, including support for Fibre Channel over Ethernet (FCoE) as a future feature entitlement upgrade
- ▶ Integrated Management Module (IMM) for enhanced systems management capabilities
- ▶ Serial Attach SCSI (SAS) plus RAID to maximize throughput and ease installation of a RAID card

Intel Xeon 7500 series processors

The x3850 X5 is powered by the highly intelligent and massively scalable Intel Xeon processor 7500 series. Built to handle your most demanding applications, the Intel Xeon processor 7500 series delivers a major increase in enterprise-computing performance.

The Intel Xeon processor 7500 series features Intel Advanced Reliability Technology that provides automatic detection and correction of errors, dynamic reassignment of workloads

across processors, interconnect error detection/recovery, and individual virtual machine recovery in virtualized environments.

IBM System Storage DS5000

The IBM System Storage DS5000 family, which includes the DS5100 and the DS5300, sets a new standard for performance, scalability, reliability, availability, and flexibility for midrange storage systems. The most powerful mid-range storage system available from IBM, the DS5000, is the ideal platform for a database environment that can keep pace with an organization's business growth. Organizations can buy only the capacity needed initially, and then dynamically upgrade and reconfigure additional capacity and features later to meet changing business requirements, all without any system downtime. Figure 2 shows the IBM System Storage DS5300.



Figure 2 IBM System Storage DS5300

The DS5000 delivers class-leading performance and is equally adept at supporting:

- ▶ Transactional applications, such as databases and OLTP
- ▶ Throughput-intensive applications, such as HPC and rich media
- ▶ Concurrent workloads, well-suited for consolidation and database

The DS5000 includes the following key features:

- ▶ Flexible and unique host interface options, for example, 4 or 8 Gbps Fibre Channel (FC) and 10 Gbps iSCSI-ready interface cards that are designed to provide investment protection and life-cycle longevity to improve efficiency and lower costs.
- ▶ Field-replaceable host interface cards (HICs), two per controller. The current release supports 4 and 8 Gbps FC HICs (16 total host ports).
- ▶ Two performance and capacity models with the ability to field upgrade. The base model DS5100 is field upgradeable to high-end model DS5300.
- ▶ *Pay-as-you-grow* scalability up to 256 drives for DS5100 and 448 drives for the DS5300 for the most capacity-demanding requirements. Support for RAID levels 1, 3, 5, 6, and 10.
- ▶ Options include 4, 8, or 16 GB of cache per controller with persistent cache backup. Drive-level encryption with no performance degradation. Innovative self-encrypting disk solutions for mid-market clients takes the worry out of exposing sensitive data on drives that are returned for repair, retired, or repurposed with near zero performance impact.
- ▶ Architecture efficiently handles compute-intensive parity calculations, enabling exceptional disk-based performance that is ideally suited for RAID-5 and RAID-6 configurations.
- ▶ Remote Volume Mirroring and FlashCopy® premium features for Volume Shadow Copy (VSS) supported backups and flexible DR scenarios.

- ▶ Designed to support high availability with dual-active, hot-swappable controllers, power supplies, and many nondisruptive firmware upgrades.
- ▶ Full Microsoft MPIO support via a custom DSM
- ▶ Management support via Microsoft System Center Operations Manager (3Q10)

The IBM System Storage DS5000 systems provide excellent performance, availability, and scalability for your most demanding SQL Server applications. It has the performance capabilities to handle the highly demanding I/O bandwidth needs of a Business Intelligence solution.

Furthermore, the ability to intermix high-capacity SATA II drives (currently up to 2 TB) with high performance SSDs and Fibre Channel drives allows for tiered storage implementation within a single array. Advanced copy services, such as FlashCopy and Remote Volume Mirroring, allow for nondisruptive backups and advanced disaster recovery scenarios. When combined with products such as Tivoli® Storage Manager or FlashCopy Manager, your SQL Server Business Intelligence environment is completely and easily protected.

For additional guidance on deploying SQL Server on IBM System Storage products see:

<http://ibm.com/systems/storage/solutions/isv/#microsoft>

Microsoft Windows Server 2008 R2 and SQL Server 2008 R2

Microsoft SQL Server 2008 R2 running on Microsoft Windows Server 2008 R2 enables organizations to build comprehensive, enterprise-scale analytic solutions that deliver intelligence where customers want it.

Key features include:

- ▶ Datacenter Edition, which supports up to 256 logical processors, a 4x increase from the previous product.
- ▶ Parallel Data Warehouse, an appliance product based on the SQL Server and designed to provide high-end data warehousing for petabyte-sized volumes of data.
- ▶ PowerPivot, which offers self-service Business Intelligence, allowing business users to create BI solutions in Excel. Users can publish their solutions to SharePoint to collaborate with other users, whereas database professionals can continue to administer the data in the solutions.
- ▶ Master Data Services, which provides the capability to manage an organization's master data to become, through integration, the single source of master data across the organization.
- ▶ Data-tier application (DAC), which provides a new way to develop, deploy, and manage the database objects used by an application.
- ▶ 64-bit native DBMS.
- ▶ Dynamic large memory support.
- ▶ High-availability through failover clustering, database mirroring, log shipping, and replication.
- ▶ Database snapshots, a point-in-time, instantaneous read-only copy of a database, which can be used for reporting and reverting back.
- ▶ Data compression for reduced on-disk space utilization, faster data retrieval, and faster backups and restores.

- Policy-based management for managing SQL Server instances throughout the enterprise.
- All other components necessary to provide a comprehensive, low-cost business intelligence solution.

SQL Server 2008 BI components

SQL Server 2008 includes the SQL relational database engine, analysis services, integration services, and reporting services. Together these components form a complete data warehousing and business intelligence solution.

Key features include:

- The database engine provides the services for the relational databases, which include the operational databases and the data warehouses.
- Integration services provide the means for extracting, transforming, and loading (ETL) the data from the operational databases into the data warehouses.
- Analysis services provides the means for creating, processing, and serving online analytical processing (OLAP) cubes.
- Reporting services provides a comprehensive reporting solution for SQL Server data stored in relational databases and cubes.

Figure 3 shows the IBM BI reference configuration architecture.

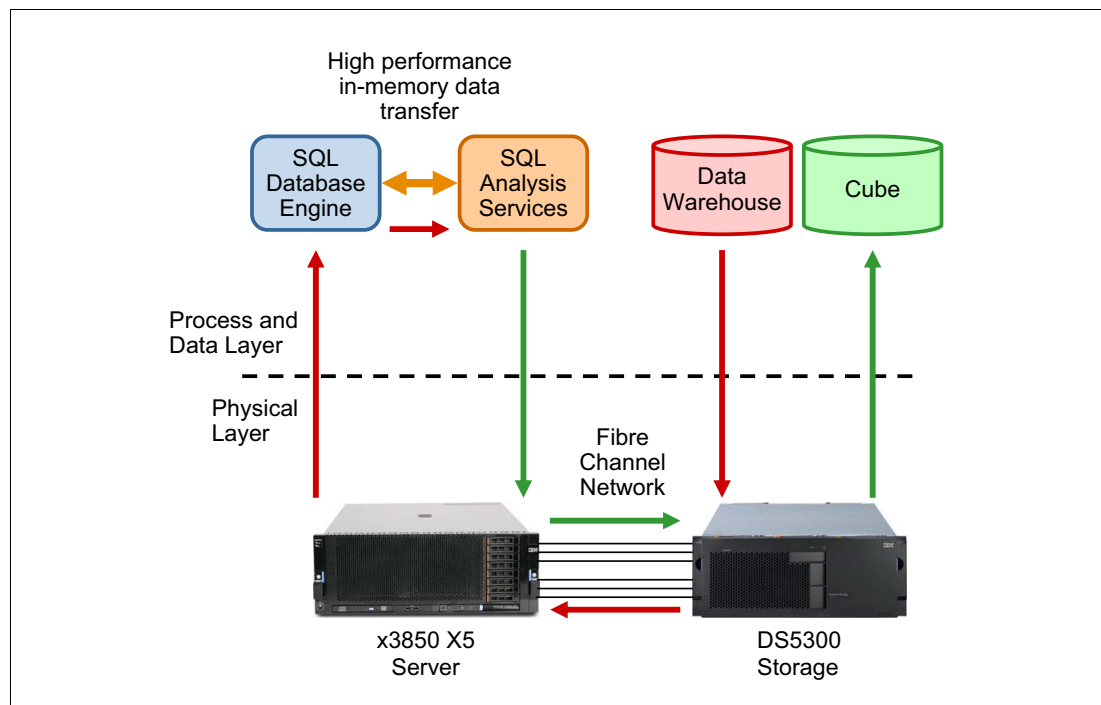


Figure 3 IBM BI reference architecture

Figure 3 shows the various components in the BI Reference configuration and how the data flows between them. During the cube-building process, analysis services (AS) requests data from the DW, aggregates it, and writes it into the cube. The data flows (shown by the arrows) from the DW to the DS5300, across the Fibre Channel network, and into the x3850 X5 server where database engine (DE) is running. The DE passes the data to AS via a high-performance in-memory data transfer. AS, also running on the x3850 X5, sends the data across the Fibre Channel network into the DS5300 and from there into the cube.

Workload description

The cube building (also called cube processing) workload used to test these reference configurations comes from a data warehouse of Web tracking information. The workload and the data warehouse were received from Microsoft.

To aid describing the cube-building workload, we first provide a brief explanation of several key BI concepts:

- ▶ Dimensions
- ▶ Fact tables
- ▶ Measure groups
- ▶ Cubes

Then we provide a profile of the DW and cube design and processing used in these tests. Finally, we describe how the 2X (doubled) and 4X (doubled again) data warehouses were created.

BI terminology

Key terms used in Business Intelligence include:

- ▶ Fact tables

Fact tables contain the individual transactions or events that are relevant to the business. For example, a retail business would have a fact table of all its product sales over time.

- ▶ Dimensions

Dimensions are hierarchical entities that categorize and group fact table rows. For example, time is nearly always a dimension with a hierarchy of year, month, day, hour, minute, and second. Another dimension example is geographic location with a hierarchy of country, state, city, and zip code. A third dimension example is a product with a hierarchy of product category, product subcategory, product code, and product model.

- ▶ Measure and measure groups

A measure is a numeric attribute of a fact table that can be aggregated (for example, the sale amount in dollars of a transaction). The sale amount can be summed for all the transactions in a particular day, giving the gross total sales for that day. A measure group is a specific combination of measures and dimensions.

- ▶ Cubes

A cube is composed of measure groups, which, in turn, are composed of fact tables and dimensions. In order to quickly return queries from the cube, the cube also contains pre-aggregated measures from the fact tables along the related dimensions. For example, the cube might contain the total sales for each product in each city for each day. Then a query, asking for the total sales of all bicycles sold in Seattle, WA, on January 15, 2009, can be quickly returned without having to scan through all the data and sum up the result. A similar query requesting the same data for the entire month of January can be returned by summing up each of the (pre-aggregated) 31 days of January. This illustrates that not every possible pre-aggregation must be stored. Higher-level aggregations can be created *on the fly* from lower-level ones.

A cube can be queried efficiently by a decision maker to explore and plan new business actions. Cubes can be queried using the industry-standard Multidimensional Expressions (MDX) query language. Several products support MDX query, including reporting services.

Reference database design

The Web tracking database and application provided by Microsoft fits in the general category of mining the information captured during user Web sessions. In this case, events that occur during the Web sessions are recorded in terms of success or failure and the BI analytics are used to reduce the failures. This section describes the profile of the data warehouse used in our testing. This data is provided here so that customers can compare their database profile to the one used in testing the reference configurations listed in this paper.

Data warehouse profile

The data warehouse, referred to below as the small (1X) DW used in this test, is 1.77 TB. It contains 26 dimension tables. The largest dimension has 77 million rows. There are five dimensions in the 100,000–300,000 row range, three dimensions in the 10,000–100,000 row range, and three dimensions in the 1,000–4,000-row range. The remaining dimensions have fewer than 1,000 rows.

There are three main fact tables. On average, about 12 dimension tables are referenced by the three fact tables. Each fact table appears in a standard star-schema configuration with the dimensions referenced. The largest fact table (row length 115) has 5 billion rows, the next largest fact table (row length 125) has 3.6 billion rows, and the third fact table (row length 95) has 2.4 billion rows, giving roughly 11 billion rows total. The fact tables are partitioned by date with a clustered index using an integer surrogate key ($\text{key} = \text{YYYY} * 10,000 + \text{MM} * 100 + \text{DD}$) into 14 major partitions by month.

Between the three fact tables there are five non-unique, non-clustered indexes that are aligned on the same partitioning scheme as the clustered indexes. The arrangement is such that it can support the *sliding window* partition update scenario. There are 14 files, one for each of these months, so three partitions and their indexes reside in each file, one for each of the fact tables for that file's month. The partitions (months) contain an uneven number of rows. The largest file is 190 GB and the smallest is 49 GB.

The number and size of the dimensions along with the row lengths and number of rows given for each of the three fact tables provides a point of reference for the volume and complexity of data processed in these reference configurations. Depending on how similar your own DW and cube definition are to this description, you might be able to extrapolate how your workload will perform on the reference configurations provided.

Cube design and processing

There are four measure groups. One measure group is created for each of the two largest fact tables. Two measure groups are created for the smallest fact table. The fact tables represent events. There are computed columns that define binary (0 or 1) attributes about these events (for example, success or failure). These binary attributes are counted and comprise the measures computed in the measure groups. Thus, the cube provides visibility into the frequency of various outcomes related to the events. For example, the cube can support a query showing the percentage of successful events across a slice (dimension) of the fact data.

In these tests, when the cube is processed, each measure group is processed one at a time in sequence. Some improvement in elapsed time can be realized by processing all measure groups at once. Processing the small (1X) DW produced a cube approximately .58 TB in size.

Doubling the data warehouse

In order to demonstrate the scaling of the hardware, we needed a data warehouse double the size of the original. We created this larger DW by doubling (copying) the fact table rows and

the largest dimension table. The new fact table rows were each modified to reference a corresponding new row in the dimension table. This allowed us to compare building the cube using the small (1X) DW with the medium (2X) DW. By creating a medium hardware configuration, which is twice the size and power of the small hardware configuration, we can evaluate the scalability of the solution. Good scaling results in the medium (2X) configuration running on the medium (2X) DW and taking about the same amount of time as the small (1X) configuration running on the small (1X) DW.

We used the doubling procedure again on the medium (2X) DW to produce the large (4X) DW mentioned in the query performance section.

Reference configurations

In this section we provide the reference configurations tested, including the parameters settings used (see appendix) and the overall elapsed time taken to build the cube in each configuration. These reference configurations show how customers can process larger and larger cubes in roughly the same amount of time by adding processor and memory resources to the x3850 X5 server, highlighting its scale-up features. Table 2 describes the reference configurations tested.

Table 2 Reference configurations

	Small (1X/1S)	Medium (1X/2S)	Large (2X/2S+)
Data warehouse size	1.8 TB 93% space used 11 billion fact rows	1.8 TB 93% space used 11 billion fact rows	3.7 TB 68% space used 22 billion fact rows
Cube size	.58 TB	.58 TB	.99 TB
Cube building time	6 hours, 36 minutes	3 hours, 42 minutes	6 hours, 24 minutes
Time % of 1X	100%	56%	97%
Server			
Server	One x3850 X5 server Two 8-core 2.26 GHz CPUs 64 GB memory	One x3850 X5 server Four 8-core 2.26 GHz CPUs 128 GB memory	One x3850 X5 server Four 8-core 2.26 GHz CPUs 256 GB memory
Storage			
Total Drives	56	56	101
DW – Data	4 RAID-5 arrays 7 drives each	4 RAID-5 arrays 7 drives each	7 RAID-5 arrays 7 drives each
SQL Server TempDB data	1 RAID-1 array 6 drives	1 RAID-1 array 6 drives	3 RAID-1 arrays 6 drives each
SQL Server			
TempDB log	1 RAID-1 array 2 drives	1 RAID-1 array 2 drives	1 RAID-1 array 4 drives
OLAP cube	1 RAID-5 array 15 drives	1 RAID-5 array 15 drives	1 RAID-5 array 15 drives
OLAP temp	1 RAID-5 array 5 drives	1 RAID-5 array 5 drives	1 RAID-5 array 15 drives

Scaling results

In this section we discuss two scaling results, 1X/1S to 1X/2S and 1X/1S to 2X/2S+, and an additional comparison where we held the DW size fixed at 4X and changed the server configuration by varying the amount of memory. The X above refers to the size of the data warehouse. The 1X DW is 1.8 TB in size, 2X DW is 3.7 TB, and 4X DW is 6.7 TB. The S above refers to the size of the server. The 1S server is 32 cores/64 GB. The 2S server is 64 cores/128 GB. The 2S+ server is 64 cores/256 GB.

1X/1S to 1X/2S scaling

The scaling results are excellent going from small (1X/1S) to medium (1X/2S), as evidenced by the elapsed time being nearly half (56%). This shows a nearly linear benefit (reduction in half) from doubling the processor and memory resources while keeping the database size constant. The number of disks was unchanged.

1X/1S to 2X/2S+ scaling

The scaling results are also excellent going from small (1X/1S) to large (2X/2S+), as evidenced by the elapsed time being nearly the same (6.5 hours). In going to the 2X/2S configuration we doubled the number of cores, quadrupled the amount of memory, and doubled the number of disks. We also used the same set of tuning parameters with proportional settings. For example, analysis services was given 50% of the memory and the SQL Server was given 30%, leaving the remaining 20% for the OS and the system file cache. The SQL Server was given 16 out of 32 processors (50%) in the 1X/1S configuration and 32 out of 64 processors (also 50%) in the 2X/2S+ configuration.

Performance metrics

In this section we provide basic performance metrics for the three reference configurations. The basic metrics provided are:

- ▶ Disk bandwidth usage
- ▶ Rows read from the DW
- ▶ Rows written to the cube

Disk bandwidth usage

Figure 4 shows the total MB per second (read plus write) of disk traffic when building the cube in each of the three test configurations.

The chart on the left compares the small configuration with the medium configuration. The medium 1X/2S configuration reads and writes at a higher rate than the small 1X/1S configuration and processes the same amount of data in half the time.




Figure 4 Disk bandwidth usage

The chart on the right compares the small configuration with the large configuration. The large 2X/2S+ configuration reads and writes at an accelerated rate and processes twice the amount of data in the same amount the time. In these cases, the MB per second does not have to be twice the rate of the small 1X/1S configuration because the additional memory in the medium and large configurations reduces the amount of IO required (for example, because temporary work files can be kept in memory).

Rows read/second from the DW

Figure 5 shows the number of rows read from the DW while building the cube. Higher is better. The number of rows read from the DW increases in proportion to the increased workload.




Figure 5 Rows read/second from the DW

Rows written/second to the cube

Figure 6 shows the number of rows written to the cube while building the cube. Higher is better. As the DW size increases, so does the number of rows written to the cube.




Figure 6 Rows written/second to the cube

Memory effect on cube processing

In addition to the three reference configurations that we list in this document, we also carried out a slightly different scaling study to show the benefit of doubling the memory resources while keeping the database size constant. We ran three cube processing tests on a 4X DW 6.7 TB database. The cube produced was 1.84 TB.

The disk configuration was the same for these tests with 202 total disks. The database size was constant at 6.7 TB. The goal of the exercise was to understand the scaling that can be achieved by increasing the amount of memory.

We found that at the smallest memory configuration, 64 GB, the cube processing took over 20 hours on the 4X DW. In the largest memory configuration, 256 GB, the cube processing took 13 hours and 48 minutes.

Figure 7 shows the beneficial decrease in cube processing time as memory increased.




Figure 7 Memory effect on cube processing

Query performance

After building the cube we ran 25 concurrent users, each running six queries against the cube, and recorded the response time in seconds. We did this for each of the cube sizes:

- ▶ 1X
- ▶ 2X
- ▶ 4X

Figure 8 shows the ninetieth percentile response time in seconds for 25 concurrent users, each running the same six queries. This chart shows that even with 25 concurrent users running queries against the cube, the server provides sub-second ninetieth percentile response times.




Figure 8 Query performance

Summary

In this document we described three reference configurations for the IBM Business Intelligence solutions with Microsoft SQL Server suitable for small, medium, and large businesses. Based on our analysis, these are the conclusions reached:

- ▶ These configurations show excellent 100% scaling when going from a 1.8 TB DW on a 2-socket 32-core x3850 X5 to a 3.7 TB DW on a 4-socket 64-core x3850 X5.
- ▶ These configurations show a 56% reduction in elapsed time when applying a more powerful hardware configuration to the same workload (in this case, cube processing on a 1.8 TB DW).
- ▶ Customers who are constrained to a time window for cube processing might benefit by seamlessly adding additional memory. In our tests, this resulted in an overall reduction of approximately 100% of the elapsed time.
- ▶ Complex cube processing and querying can be performed in a timely manner using the x3850 X5 server, which supports the extraction of actionable insights from a company's transactional data, which is the goal of business intelligence solutions.

Appendix

Hardware specifications

The specifications were the same for all configurations.

Table 3 Hardware specifications

Component	Description
Server	IBM System x3850 X5, 8-core 2.26 GHz processors
Server local drives: OS and SQL	1 RAID-1 array, 2 drives, 146 GB, 10K rpm, SFF
Memory	DDR3
SAN disks	300 GB, 15K rpm

Software specifications

Table 4 Software specifications

Component	Description
SQL Server	Microsoft SQL Server 2008 R2 (CTP): 10.50.1352.12 (X64) Enterprise Edition (64-bit)
Windows	Windows Server 2008 R2 Enterprise Edition x64 Build 7600

Tuning parameters

Table 5 Tuning parameters

	Small (1X/1S)	Medium (1X/2S)	Large (2X/2S)
SQL Server			
Tempdb files	64	64	64
sp_configure			
Max server memory	30%	30%	30%
Affinity mask	0xFFFF	0xFFFFFFFF	0xFFFFFFFF
Affinity64 mask	0x0	0xFFFFFFFF	0xFFFFFFFF
Max degree of parallelism	1	1	1
Analysis Services			
Msmdsrv.ini			
OLAP\Process\AggregationMemoryLimitMin	1%	1%	1%
OLAP\Process\AggregationMemoryLimitMax	5%	5%	5%
Memory\TotalMemoryLimit	50%	50%	50%
Memory\LowMemoryLimit	30%	30%	30%
Memory\PreAllocate	35%	35%	35%
ExternalCommandTimeout	3,600	3,600	3,600

	Small (1X/1S)	Medium (1X/2S)	Large (2X/2S)
Create cube xmla			
MaxActiveConnections	32	64	64
Process cube xmla			
Parallel MaxParallel	32	64	64

References

Refer to the following resources for more information:

- ▶ *IBM System x Business Intelligence Solutions for Small-Medium Businesses with Microsoft SQL Server 2008*, REDP-4552
<http://www.redbooks.ibm.com/abstracts/redp4552.html>
- ▶ *IBM System x Business Intelligence Solutions for Small to Large Businesses with Microsoft SQL Server 2008*, REDP-4587
<http://www.redbooks.ibm.com/abstracts/redp4587.html>
- ▶ *Microsoft AS Performance Tuning Guide*
<http://sqlcat.com/whitepapers/archive/2009/02/15/the-analysis-services-2008-performance-guide.aspx>
- ▶ SQL Server books online
<http://www.microsoft.com/downloads/details.aspx?FamilyId=765433F7-0983-4D7A-B628-0A98145BCB97&displaylang=en>
- ▶ SQL Customer Advisory Team (CAT) Analysis Services resources
<http://sqlcat.com/tags/Analysis+Services/default.aspx?PageIndex=1>

The team who wrote this guide

Michael Lawson is a Database and Virtualization Specialist at the IBM Center for Microsoft Technologies in Kirkland, WA. He has a Bachelor of Arts degree in Information Sciences and Mathematics from the University of California at Santa Cruz. Michael has 30 years of experience as a database administrator, the last 14 years with Microsoft SQL Server. Since joining IBM in 1999, he has supported IBM customers running SQL Server on System x® servers, especially very large NUMA configurations. He was the principal IBM SQL Server DBA for the Sydney 2000 Olympics. He is a regular speaker at the IBM Executive Briefing Center in Kirkland, WA. He hosted a Windows Solution Lab technical engagement resulting in the largest Microsoft ERP solution in the United States. He holds the Microsoft Certified Database Administrator certification. Michael is also an author of other IBM Redbooks® publications, including *Running Microsoft SQL Server 2008 on the IBM System x3950 M2*, REDP-4476, and the two other IBM Redguides in this series.

Vinay Kulkarni has a Bachelor of Computer Engineering degree from Shivaji University in India. He has his master's degree in computer science from Minnesota State University. He started his career on the AS/400® File system test team in Rochester, MN. Then he worked on the IBM AS/400 performance benchmark team, publishing many industry-leading benchmark results. In his current role, he is an IBM System x Performance Engineer working on-site at the Microsoft Redmond campus. He has been working with Microsoft over the past eight years to optimize performance of System x servers running Microsoft Windows and SQL Server software. He works closely with the Microsoft Windows and SQL Server performance teams to ensure good performance of IBM System x servers running software from Microsoft. He also works with IBM clients to tune the performance of System x and storage environments. He is the co-author of the two other IBM Redguides in this series.




Figure 9 The authors (l-r): Michael and Vinay

The authors would like to acknowledge the Microsoft SQL team for providing the Web Tracking DW and the workload to go with it, and especially Tim Shea, in the Microsoft SQL Performance group, for his consultation on this project.

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author - all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our papers to be as helpful as possible. Send us your comments about this paper or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

ibm.com/redbooks

- Send your comments in an e-mail to:

redbooks@us.ibm.com

- Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on Facebook:

<http://www.facebook.com/pages/IBM-Redbooks/178023492563?ref=ts>

- Follow us on twitter:

<http://twitter.com/ibmredbooks>

- Look for us on LinkedIn:

<http://www.linkedin.com/groups?home=&gid=2130806>

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>

- Stay current on recent Redbooks publications with RSS Feeds:

<http://www.redbooks.ibm.com/rss.html>

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

This document, REDP-4662-00, was created or updated on March 29, 2010.




Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>



The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AS/400®	Redguide™	System x®
FlashCopy®	Redguides™	Tivoli®
IBM®	Redbooks (logo)  ®	
Redbooks®	System Storage™	

The following terms are trademarks of other companies:

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel Xeon, Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.