


IBM System x Business Intelligence Solutions for Small to Large Businesses with Microsoft SQL Server 2008



Redguides
for Business Leaders

Vinay Kulkarni
Michael Lawson



- Second paper in the series on Business Intelligence solutions from IBM System x
- Shows how IBM System x3950 M2 configurations are an ideal match for small, medium and large BI customer requirements
- Shows the improvements possible with right hardware and software configuration



Executive summary

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

Figuring out the hardware and software needed to build BI configurations is a complicated and time-consuming process. IBM® has built some reference configurations and tested them to provide clients with guidelines when sizing their BI configurations. These configurations are being built using these key components:

- ▶ IBM System x® 3850 M2 and x3950 M2 servers
- ▶ IBM System Storage™ DS5000
- ▶ Microsoft® Windows® Server 2008
- ▶ Microsoft SQL Server 2008

This Redguide™ is the second in a series that describes BI reference configurations that scale with database size. The first Redguide, *IBM System x Business Intelligence Solutions for Small-Medium Businesses with Microsoft SQL Server 2008*, REDP-4552, covered reference configurations at a constant database size. It is available from the following Web page:

<http://www.redbooks.ibm.com/abstracts/redp4552.html>

This document provides sizing information for small, medium, and large businesses that are planning to use the Microsoft BI solution using IBM System x servers, IBM System Storage DS5000 storage, and Microsoft Windows and SQL Servers. Three progressively larger data warehouses are investigated.

The three reference configurations tested are listed in Table 1.

Table 1 Reference configurations

	Small (1X)	Medium (2X)	Large (4X)
DW Database Size	1.8 TB	3.7 TB	6.7 TB
Server and Storage	12 cores, 64 GB memory, 1 DS5300, 56 HDDs	24 cores, 128 GB memory, 1 DS5300, 103 HDDs	48 cores, 192 GB memory, 1 DS5300, 206 HDDs

IBM offers a unique and compelling high-end, modular, Intel® processor-based x3950 M2 server running Windows Server 2008 Enterprise. This server is a modular “building-block” server. The x3950 M2 can be configured with one, two, three, or four chassis. This paper reports on three scenarios in the small, medium, and large range. The first and second scenarios use one chassis, and the third scenario uses two chassis. Additionally, we compare the large 4X configuration case using two different system configurations to show the scale-up advantages of the IBM x3950 M2. In the next Redguide in this series, we report on a very large configuration that uses four x3950 M2 chassis.

In conclusion, we show that 100% scaling was achieved going from the small to medium configuration. The results show that customers can go to larger and larger IBM hardware configurations and reduce the cube building and processing times. A 50% reduction in cube building time can be achieved by doubling the number of cores and memory on the IBM x3950 M2. The results also show that IBM x3950 M2 and DS5300 are capable of handling the BI needs of large enterprise customers with database sizes over 6 TB.

Introduction

The idea behind BI 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. Some 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 has built some reference configurations and tested them to provide clients with guidelines when sizing their BI configurations. IBM used a workload developed by Microsoft using their internal proprietary data. The workload used for testing the configurations listed in this paper 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. IT professionals need to know approximately how long their cubes will take to build. To assist with this task, this paper provides three examples of data warehouses of increasing size. The reference configurations have been tested by measuring how long it takes to build the cube.

We describe the following advantages of the reference configurations in this paper:

- ▶ Configurations are tested using the Microsoft BI software stack, providing confidence in the combined solution (software + server + storage).
- ▶ There are three sizes of hardware configurations (small, medium, large) to suit your workload.
- ▶ Seamless upgrade is possible for more powerful configurations as the database and workload grow.

The remainder of this paper contains the following sections:

- ▶ “Solution architecture” on page 3
- ▶ “Workload description” on page 6
- ▶ “Reference Configurations” on page 8
- ▶ “Conclusion” on page 12

Solution architecture

Solution components include the high-end System x 3850 M2 and 3950 M2 servers, the high-performance mid-range DS5300 storage system from IBM and Windows Server 2008 and SQL Server 2008 from Microsoft. These components are described in the following sections.

IBM System x3950 M2 server

The IBM System x3950 M2 server is based on eX4, the fourth generation chipset of the proven IBM Enterprise X-Architecture®. The x3950 M2 has a feature called the ScaleXpander, which lets you connect four x3950 M2 chassis together to form one large single-image system with up to 16 sockets and 96 cores.

Also available is the IBM System x3850 M2, which is an x3950 M2 without the ScaleXpander feature. This means that it can have at most four sockets and 24 cores. The ScaleXpander feature can, however, be added to the x3850 M2 at a later date fully upgrading it to an x3950 M2.

The x3950 M2 and x3850 M2 include the following key features:

- ▶ Ability to scale-up on processor and memory based on a modular “pay-as-you-grow” design
- ▶ High performance, as demonstrated by leadership TPC benchmark results
- ▶ Mainframe-like reliability in an x86 environment
- ▶ True 2-to-16-socket scalability up to 96 cores
- ▶ Superior performance with Intel Xeon® multi-core 7400 Series Processors and 16 MB L3 cache
- ▶ Up to 1 TB of registered DIMM memory for better workload density and up to 30% less power consumption than fully buffered DIMM technology
- ▶ IBM Memory ProteXion with redundant bit-steering, which offers twice the memory resilience of other systems
- ▶ Fourth-generation snoop filter to ensure maximum efficiency across all processors
- ▶ IBM Predictive Failure Analysis, on hard drives and memory as well as on processors, power supplies, fans, and voltage regulator modules
- ▶ Local memory latency that is 40% lower than the nearest competition
- ▶ Up to 32 registered DDR2 4 GB DIMMs that offer higher system performance at lower power consumption enabling faster speeds and higher capacities needed for high-end enterprise servers like the x3950 M2

IBM System Storage DS5000

The IBM System Storage DS5000 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.

The DS5000 delivers class-leading performance and is equally adept at supporting transactional-applications such as OnLine Transaction Processing (OLTP) and throughput-intensive applications such as HPC/DW and rich media. The DS5000 is also well-suited for concurrent workloads such as database consolidation.

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 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.
- ▶ Up to 700,000 IOPs and 6,400 MBps performance capability.
- ▶ Options include 4, 8, or future 16 GB of cache per controller.
- ▶ Persistent cache backup in the event of a power outage.
- ▶ Drive-level encryption with no performance degradation. Innovative self-encrypting disk solution 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.
- ▶ Support for RAID levels 6, 5, 3, 10, 1, and 0.
- ▶ 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 nondisruptive firmware upgrades.

Microsoft SQL Server 2008

Microsoft SQL Server 2008 running on Microsoft Windows Server 2008 enables organizations to build comprehensive, enterprise-scale analytic solutions.

Key features include:

- ▶ 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 use, faster data retrieval, and faster backups and restores
- ▶ Policy-based Management for managing SQL Server instances throughout the enterprise
- ▶ All other components needed 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 includes the operational databases and the data warehouses.
- ▶ Integration Services provides 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 1 shows the IBM BI reference configuration architecture.

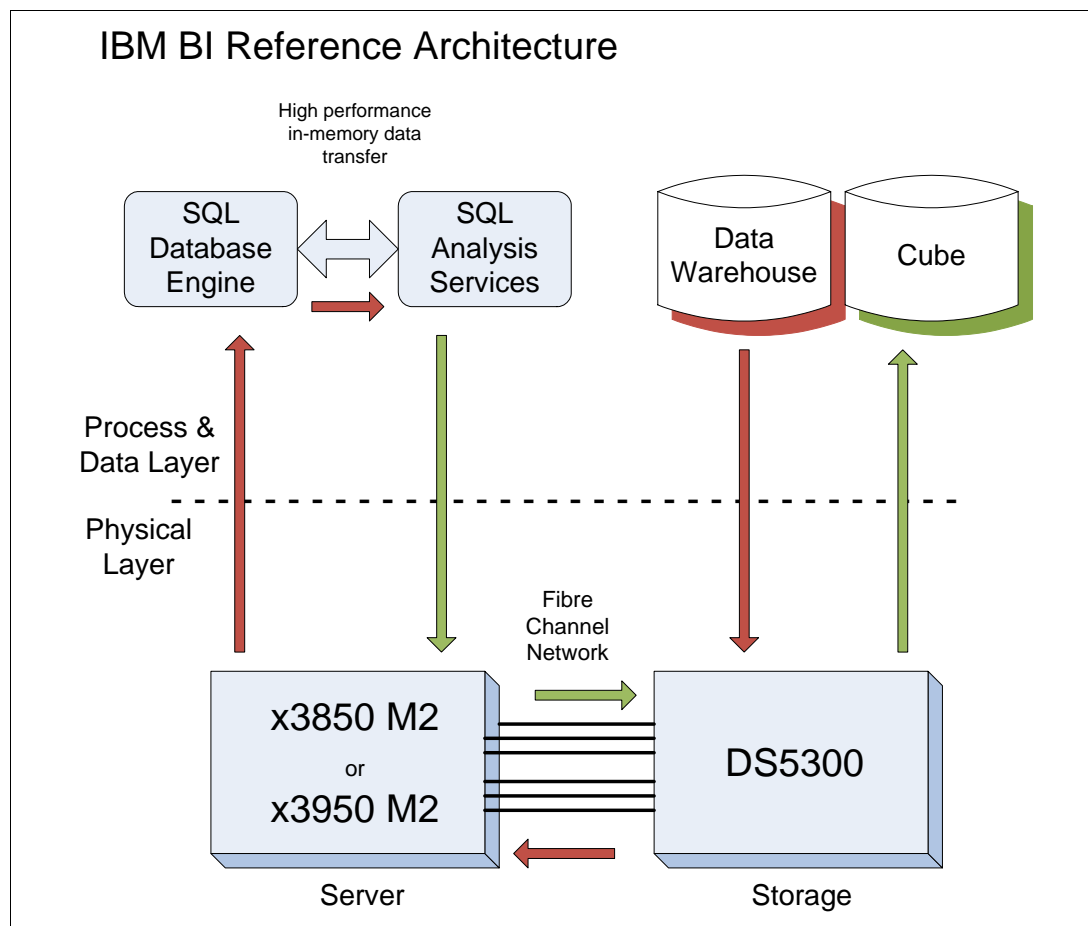


Figure 1 IBM BI reference configuration architecture

Figure 1 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 Data Warehouse (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, into the x3950 M2 (or x3850 M2) server where the Database Engine (DE) is running. The DE

passes the data to AS through a high performance in-memory data transfer. AS, also running on the x3950 M2 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. In order to describe the cube-building workload, in the sections below, we give a brief explanation of several key BI concepts:

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

Then we give a profile of the Data Warehouse (DW) and Cube design and processing used in these tests. And we describe how the 2X (doubled) and 4X (doubled again) data warehouses were created.

BI terminology

There are a few key terms used in BI that are worth introducing at this point:

- ▶ 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 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 postal 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 could 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. To return queries from the cube quickly, 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, could 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 could be returned by summing up each of the (pre-aggregated) 31 days of January.

This illustrates that not every possible pre-aggregation needs to 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 MDX (Multidimensional Expressions) query language. Several products support MDX query, including Reporting Services.

Web tracking

The Web Tracking database and application provided by Microsoft fits in the general category of mining the information captured during a user's 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. The following paragraphs describe the profile of the data warehouse used in our testing. This data is provided here so customers can compare their database profile to the one used in testing the reference configurations in this paper.

Data Warehouse Profile

The data warehouse (DW), 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 to 300,000-row range, three dimensions in the 10,000 to 100,000-row range, and three dimensions in the 1,000 to 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 five 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.

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 could support a query showing the percentage of successful events across some 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 .44 TB in size.

Doubling the Data Warehouse

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 could evaluate the scalability of the solution. Good scaling would result 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.

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 x3950 M2 server, highlighting its scale-up features.

Table 2 describes the reference configurations tested.

Table 2 Reference configurations

	Small (1X)	Medium (2X)	Large (4X)
Data Warehouse			
Data Warehouse size	1.8 TB, 93% space used, 11 billion fact rows	3.7 TB, 68% space used, 22 billion fact rows	6.7 TB, 75% space used, 44 billion fact rows
Cube size	0.44 TB	0.78 TB	1.5 TB
Cube Building Time	10 hours 45 minutes	10 hours 57 minutes	14 hours 11 minutes
% of 1X	100%	102%	132%
Server			
Server	One x3950 M2 server: <ul style="list-style-type: none"> Two 6-core 2.6 GHz processors 64 GB memory Two Emulex LPe12002 dual-port 8 Gb FC HBAs 	One x3950 M2 server: <ul style="list-style-type: none"> Four 6-core 2.6 GHz processors 128 GB memory Four Emulex LPe12002 dual-port 8 Gb FC HBAs 	Two x3950 M2 servers <ul style="list-style-type: none"> Eight 6-core 2.6 GHz processors 192 GB memory Eight Emulex LPe12002 dual-port 8 Gb FC HBAs
Storage			
Total Drives	56	103	206
DW – Data	Four RAID-5 arrays, seven drives each	Seven RAID-5 arrays, seven drives each	14 RAID-5 arrays, seven drives each
SQL Server TempDB Data	One RAID-1 array, six drives	One RAID-1 array, 10 drives	One RAID-1 array, 20 drives

	Small (1X)	Medium (2X)	Large (4X)
SQL Server			
TempDB Log	One RAID-1 array, two drives	One RAID-1 array, four drives	One RAID-1 array, eight drives
OLAP Cube	One RAID-5 array, 15 drives	One RAID-5 array, 30 drives	Two RAID-5 arrays, 30 drives each
OLAP Temp	One RAID-5 array, five drives	One RAID-5 array, 10 drives	One RAID-5 array, 20 drives

Scaling results

In this section, we discuss three different scaling results: 1X to 2X, 2X to 4X and an additional comparison where we held the DW size fixed at 4X and changed the server configuration from 1- node x3850 M2 to 2-node x3950 M2. The X above refers to the size of the data warehouse. The 1X DW is 1.8 TB in size, 2X DW is 3.7TB and 4X DW is 6.7TB in size.

1X to 2X scaling

The scaling results are excellent going from small (1X) to medium (2X) as evidenced by the elapsed time being nearly the same (just under 11 hours). What this says is that you can double the workload and double the hardware, and it will take about the same amount of time to build the cube. This is an ideal result.

In going to the 2X configuration we doubled the number of cores and the amount of memory. The number of disks was roughly doubled also. We also used the same set of tuning parameters with proportional settings. For example, Analysis Services was given 50% of the memory and SQL Server was given 30%, leaving the remaining 20% for the OS and the system file cache. SQL Server was given four out of 12 processors (33%) in the 1X-configuration and eight out of 24 processors (also 33%) in the 2X-configuration.

2X to 4X scaling

The scaling results going from medium to large at 60% are not as good as the 1X to 2X at 100%. The 4X-configuration ran about one-third longer than the 2X-configuration. We get some scaling, but not linear scaling. Some of this we expect because of contention between the large numbers of components. We were able to double the number of disks and cores, but not the amount of memory. Doubled memory would have been 256 GB, but we had only 192 GB memory. This also reduced our scaling expectation. The reduced memory put pressure on the disks, and we saw evidence of disk bottlenecks on the TempDB and OLAP Temp disks.

In the 4X configuration, to take advantage of the large number of cores, we used different tuning parameters than in the 1X and 2X configurations. We limited the system file cache to 20% (through the LimitSystemFileCacheSizeMB option) of memory and reduced SQL Server from 30% to 20% of memory. We also increased other parameters to give Analysis Services more threads. Still, we were not able to use all the cores effectively and suspect that there are some scalability issues with Analysis Services in this largest configuration.

4X DW 1-Node versus 2-Node Scaling

Apart from the three reference configurations we have listed in this Redguide, we carried out one other slightly different scaling test to show the benefit of doubling the processor and memory resources while keeping the database size constant. We ran two tests on a 4X DW 6.7 TB database:

- ▶ For the first test, the x3850 M2 server was configured with 24 cores and 96 GB memory.
- ▶ For the second test, the server was upgraded to a 2-node x3950 M2 server with 48 cores and 192 GB memory. This is the same configuration as the third reference configuration listed in the paper.

The disk configuration was the same for these tests with the total number of disk at 206. The database size was constant at 6.7 TB. The goal of the exercise was to understand the scaling that can be achieved by just attaching an additional node to an existing single-node x 3850 M2 configuration. This is the sort of issue a customer faces when their cube processing time begins to take longer and longer and eventually does not fit in the allotted time window. In this scenario we are doubling the cores and memory, but leaving the data warehouse size fixed at 4X and the number of disks fixed at 206.

We found that on a single node the cube processing took almost 24 hours on the 4X DW. Adding the second node brings the elapsed time down to just over 14 hours, an overall reduction of approximately 60% less time. This shows that a customer can achieve 1.6x scaling just hooking up an additional x3950 M2 to a current single-node x3850 M2 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 Data Warehouse
- ▶ Rows written to the cube

Disk bandwidth usage

As the number of spindles increased in each configuration, the MBps increased, showing good scaling in the storage subsystem. See Figure 2 on page 11.




Figure 2 Disk bandwidth usage

Figure 2 shows the total MB per second (read plus write) of disk traffic when building the cube in each of the three test configurations. It is expected and desirable that the disk traffic increases in proportion to the increased workload.

Rows read per second from the data warehouse

Figure 3 shows the number of Microsoft Analysis Services (MSAS) rows read from the DW while building the cube. It is expected and desirable that the number of rows read from the DW increases in proportion to the increased workload.




Figure 3 Rows read per second from the Data Warehouse

Rows written per second to the cube

Figure 4 shows the number of rows written to the cube while building the cube. It is expected and desirable that the number of rows written to the cube increases in proportion to the increased workload.




Figure 4 Rows written per second to the cube

Conclusion

In this paper we have described three reference configurations for the IBM BI 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 2-socket 12-core x3850 M2 to a 3.7 TB DW on a 4-socket 24-core x3850 M2.
- ▶ The scaling from a 3.8 TB DW on a 4-socket 24-core x3850 M2 to 6.7 TB DW on an 8-socket 48-core x3950 M2 was measured at 1.6x.
- ▶ Customers who are constrained to a time window for cube processing may benefit by seamlessly adding a second x3950 M2 node to a single node configuration. In our tests, this resulted in an overall reduction of approximately 60% of the elapsed time.

Appendix

Hardware specifications

The specifications were the same for all configurations. See Table 3.

Table 3 Hardware specifications

Component	Description
Server	IBM System x3950 M2, 6-core 2.66 GHz processors
Server local drives: operating system and SQL Server	1 RAID-1 array, 2 drives, 72 GB, 15K RPM, SFF
Memory	DDR2
SAN disks	Each 300 GB, 15K RPM

Software specifications

Table 4 Software Specifications

Component	Descriptions
SQL Server	Microsoft SQL Server 2008 (SP1) – 10.0.2714.0 (x64) Enterprise Edition (64-bit)
Windows	Windows Server 2008 Datacenter Edition x64 Build 6001: Service Pack 1

Tuning parameters

Table 5 Tuning parameters

	Small (1X)	Medium (2X)	Large (4X)
SQL Server			
Tempdb files	12	24	10
sp_configure: max server memory	30%	30%	20%
sp_configure: affinity mask	0xF	0xFF	0xFFFF
sp_configure: max degree of parallelism	1	1	0
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	36,000
Analysis Services: Create cube xmla			
MaxActiveConnections	10	20	40
Analysis Services: Process cube xmla			
Parallel MaxParallel	12	24	48

Additional settings were used for the Large (4X) configuration.

Table 6 Addition tuning parameters for the Large (4X) configuration

Analysis Services	Large (4X)	Default
LimitSystemFileCacheSizeMB	39322	
ThreadPool\Query\MaxThreads	128	10
ThreadPool\Process\MaxThreads	256	64
CoordinatorExecutionMode	-8	-4
OLAP\LazyProcessing\MaxCPUUsage	1.0	0.5
DatabaseConnectionPoolMax	64	50
DataStorePageSize	65536	8192

Acknowledgements

We would like to acknowledge the Microsoft SQL team for providing the Web Tracking DW and the workload to go with it.

The team who wrote this guide

Vinay Kulkarni started his career on the AS/400® File system test team, in Rochester, MN. He then 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 first IBM Redguide in this series, *IBM System x Business Intelligence Solutions for Small-Medium Businesses with Microsoft SQL Server 2008*, REDP-4552. He has a Bachelor of Computer Engineering degree from Shivaji University in India. He studied his masters in computer science at Minnesota State University.

Michael Lawson is a database and virtualization specialist at the IBM Center for Microsoft Technologies in Kirkland, WA. 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 2000 Olympics in Sydney, Australia. He is a regular speaker at IBM Executive Briefing Center in Kirkland, WA. He hosted a Windows Solution Lab technical engagement resulting in the largest Microsoft ERP solution in the US. 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 first IBM Redguide in this series, *IBM System x Business Intelligence Solutions for Small-Medium Businesses with Microsoft SQL Server 2008*, REDP-4552. He holds the Microsoft Certified Database Administrator certification. He has a B.A. in Information Sciences and Mathematics from the University of California at Santa Cruz.




Figure 5 The authors (left to right): Michael and Vinay

References

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>

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 CAT (Customer Advisory Team) Analysis Services resources

<http://sqlcat.com/tags/Analysis+Services/default.aspx?PageIndex=1>

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.




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®	Redbooks®	System Storage™
FlashCopy®	Redguide™	System x®
IBM®	Redbooks (logo)  ®	X-Architecture®

The following terms are trademarks of other companies:

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.

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