IBM System x Business Intelligence Solutions for Small-Medium Businesses with Microsoft SQL Server 2008

Redguides
for Business Leaders

Vinay Kulkarni
Michael Lawson

- Presents scalable Microsoft BI reference configurations designed for small-medium sized businesses
- Shows that Microsoft SQL Server 2008 offers improved performance over SQL Server 2005
- Describes how the IBM System Storage DS5000 and the IBM System x3950 M2 provide an ideal platform for BI applications

Executive summary

Business intelligence (BI) can get a lot of attention as companies realize they have a vast amount of data that needs to be analyzed and that can be used to intelligently guide their business decisions going forward.

The issues that companies face today include:

- Data is spread out in different locations and in different formats.
- Data needs to be extracted, transformed, and integrated into one common database to be analyzed.
- When analyzed, reports need to be generated that can help executives make better business decisions.

Overall, business intelligence market opportunity is in the order of $40 to $50 billion, according to Gartner. Gartner reports also show that BI applications remain the top technology investment area in a survey of CIOs for the last 4 years.

IBM® has tested BI software components that ship with the Microsoft® SQL Server® product on System x® scale-up servers. In this IBM Redguide publication, we discuss three such tested reference configurations, as shown in Table 1. These configurations highlight the cost-effective standard Intel® Xeon processor-based, server-based scale-up BI solutions by IBM.

Table 1 Reference configurations using IBM System x3950 M2 servers with Intel Xeon® processors

	Small	**Small-Medium**	**Medium**
Configuration	16 processor cores 64 GB memory 50 disks	24 processor cores 96 GB memory 63 disks	32 processor cores 128 GB memory 87 disks
Scaling result	1x (baseline)	1.17x	1.32x scaling 24 to 32 cores 1.54x scaling 16 to 32 cores

These reference configurations have the following characteristics:

- Are tested on standard servers using Intel Xeon processors
- Are tested and tuned along with Microsoft SQL Server team
- Are configured using System x3950 M2 modular scale-up servers that are built on the IBM X4 chipset, which incorporates mainframe-like RAS features
- Can be easily scaled up, because adding processor, memory, and I/O resources when needed is as easy as connecting additional servers to the first server with scalability cables and then rebooting the single-image complex
- Show how customers can reduce cube processing time
- Describe the servers and storage that you need to invest in based on database size and the needed processing speed
- Are targeted at the small and medium business (SMB) market with database sizes up to 2 TB.

Introduction

Business intelligence is the accurate and timely interactive access to the most important information in the company. It has become a valued tool for employees in nearly every job in the modern enterprise. Integrated and accurate information about customers, products, prices, inventories, and hundreds of other subjects has grown into an essential tool for workers in offices, stores, and factories; for mobile employees in sales and service; for partners; for suppliers; and for customers.

In general, executives find that, as long as access is managed properly to safeguard confidential information, putting the correct information into the hands of each person in the extended enterprise is a powerful way to enhance service, customer satisfaction, and business performance. The idea behind BI is to take company 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 workers by aiding them in making better business decisions to help improve the company's bottom line.

Examples of BI include using the retail sales trends by geographical areas over the past few years to decide which products go on the shelf, and using the results of various sales promotions offered historically to decide on future promotions.

After a company decides to implement BI, the IT department must choose from the myriad of Business intelligence tools that are available from many different vendors in the market. This guide addresses the BI needs of the market with a cost-effective scale-up solution that consists of IBM system x3950 M2 server with Intel Xeon Processor 7400 Series processors, IBM System Storage™ DS5000 and Microsoft SQL Server BI components, specifically Analysis Services.

IBM and Intel work closely with Microsoft to ensure that our products are optimized for SQL Server 2008 deployments. This guide is the result of a joint collaboration between the companies.

In this guide, we discuss the components used in the solution. We also include specific server and storage configurations that are used to show improvements in SQL Server 2008 over SQL Server 2005. We also describe the three reference configurations tested by IBM.

This guide includes the following topics:

- IBM business intelligence solution architecture
- SQL Server 2008 performance improvements
- IBM BI reference configurations

IBM business intelligence solution architecture

Figure 1 shows a high-level overview of the IBM BI solution and how the components are related. On the left side of the diagram, SQL Database Engine and Analysis Services are the software components that execute on the x3950 M2 server hardware. On the right side of the diagram, the Data Warehouse and Cube are the data repositories that reside on the DS5000 storage hardware.

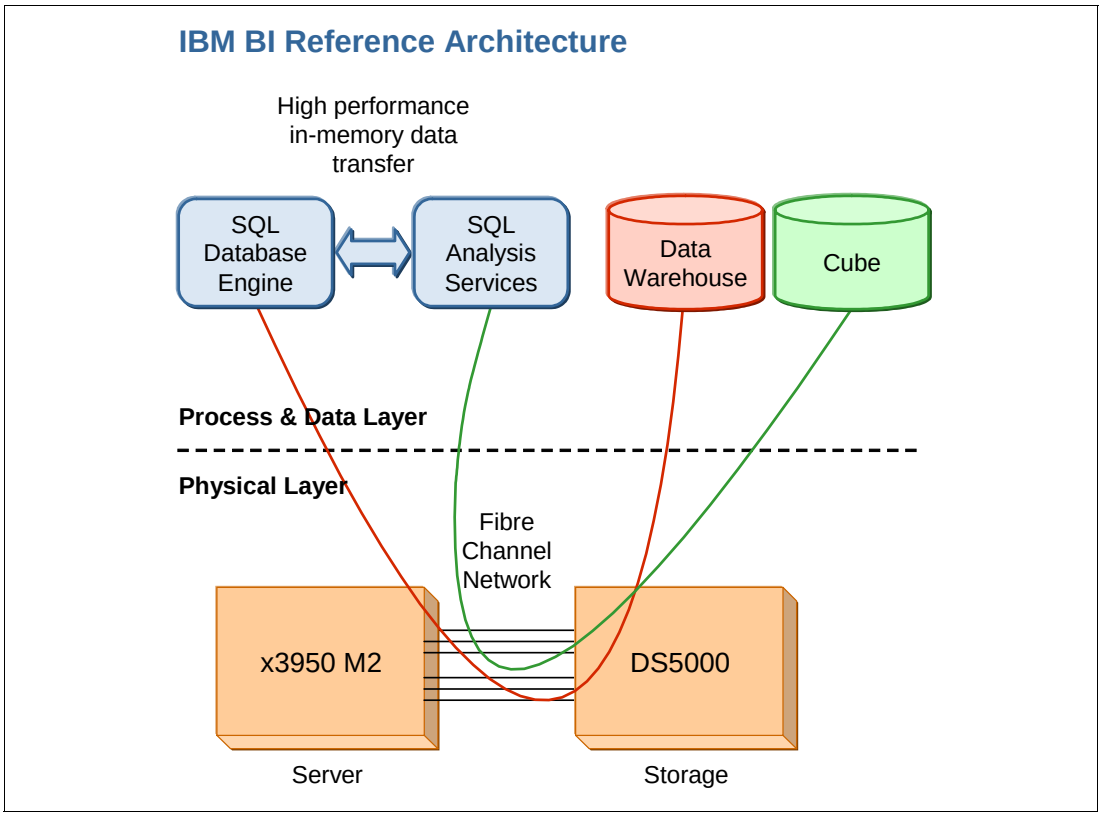

Figure 1 IBM BI solution architecture

The following basic steps are involved in the cube building process using this reference architecture:

1. Analysis Services requests data from the data warehouse using the database engine.
2. The database engine queries the data warehouse, which causes the data on the DS5000 storage to travel through the Fibre Channel network into the memory of the x3950 M2 server with Intel Xeon Processor 7400 series processors.
3. The database engine then passes that data directly to Analysis Services as an in-memory data transfer.
4. Analysis Services aggregates and processes that data and stores the result in the Cube, causing the data to travel back through the Fibre Channel Network to be stored on the DS5000 storage.

We describe the key components in the architecture in the following sections.

IBM System x3950 M2 server

The IBM System x3950 M2 server is based on eX4, the fourth generation of the proven IBM Enterprise X-Architecture® chipset. The eX4 chipset, which includes the memory controller, provides the x3950 M2 server with the following 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

The IBM System x3950 M2 includes the following key features:

- True 2–to–16-socket scalability up to 96 cores
- Superior performance with multi-core processors, Intel Xeon Processors 7400 Series 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
- Memory latency that is 40% lower than the nearest competition

Up to 32 registered DDR2 4 GB DIMMs are used in these configurations. These DIMMs 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 new standards 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 databases and OLTP; throughput-intensive applications, such as HPC and rich media; and 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 to lower costs.
- Field-replaceable host interface cards (HIC), 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 demanding capacity requirements.
- Up to 700,000 IOPs and 6,400 MBps.
- Designed for a future upgrade up to 32 GB of dedicated data cache (16 GB per controller)
  - Options include 4, 8, or future 16 GB of cache per controller
  - Dedicated cache mirroring channels
  - Persistent cache backup in the event of a power outage
  - Field-upgradeable
- 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.
- End-to-end media encryption, with support for enterprise and mid-range disk and tape.
- 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 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 many 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 that deliver intelligence where customers want it.

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 utilization, 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.

Emulex 8Gb Fibre Channel host bus adapter

The Emulex® 8Gb Fibre Channel Dual-port HBA for IBM System x 3950 M2 host bus adapters (HBAs) provide outstanding scalability, streamlined installation and management, and industry-leading database support well-suited for small-to-large enterprises and Microsoft Windows Server® 2008 and SQL Server 2008 FC storage area network (SAN) environments.

With powerful management tools and broad System x support, the LightPulse family of IBM-branded 4 Gbps and 8 Gbps HBAs (IBM Server Proven validation) delivers high performance for a broad range of applications and environments. The Emulex LPE 12002 HBAs and the accompanying drivers have built-in support for Message Signaled Interrupts (MSI-X) and NUMA I/O and take advantage of NUMA I/O improvements in Windows server 2008.

Key features include:

- Exceptional performance and full-duplex data throughput
- Comprehensive database capabilities with support for N-Port ID Database (NPIV)
- Simplified installation and configuration using AutoPilot Installer®
- Administration using HBAnyware® integrated with IBM Systems Director

SQL Server 2008 performance improvements

In this section, we show the server and storage configurations used to show improvements in SQL Server 2008 over SQL Server 2005. Improvements were made in Analysis Services in SQL Server 2008 to remove some of the single-threaded code and replace it with multi-threaded code. This change allowed us to run more threads in parallel thus reducing cube processing time.

Server configuration

Figure 2 shows the single-node, 16-core server configuration. The single-node IBM System x3950 M2 server configured with Intel Xeon Processor 7400 Series used for these tests is connected to the DS5000 storage controller with six direct connections from the three Emulex HBAs to six ports on the controller.

There are seven PCIe slots on the x3950 M2 server. Slots 1, 3, and 5 are populated with the Emulex dual-port adapters. The Microsoft multi-path bus driver is used to manage the multiple paths. The DS3000/DS4000® DSM is used with the *least queue depth* policy. The least queue depth policy compensates for uneven loads by distributing proportionately more I/O requests to lightly loaded processing paths. Other policies available are *failover only*, *round robin*, and *weighted paths*.

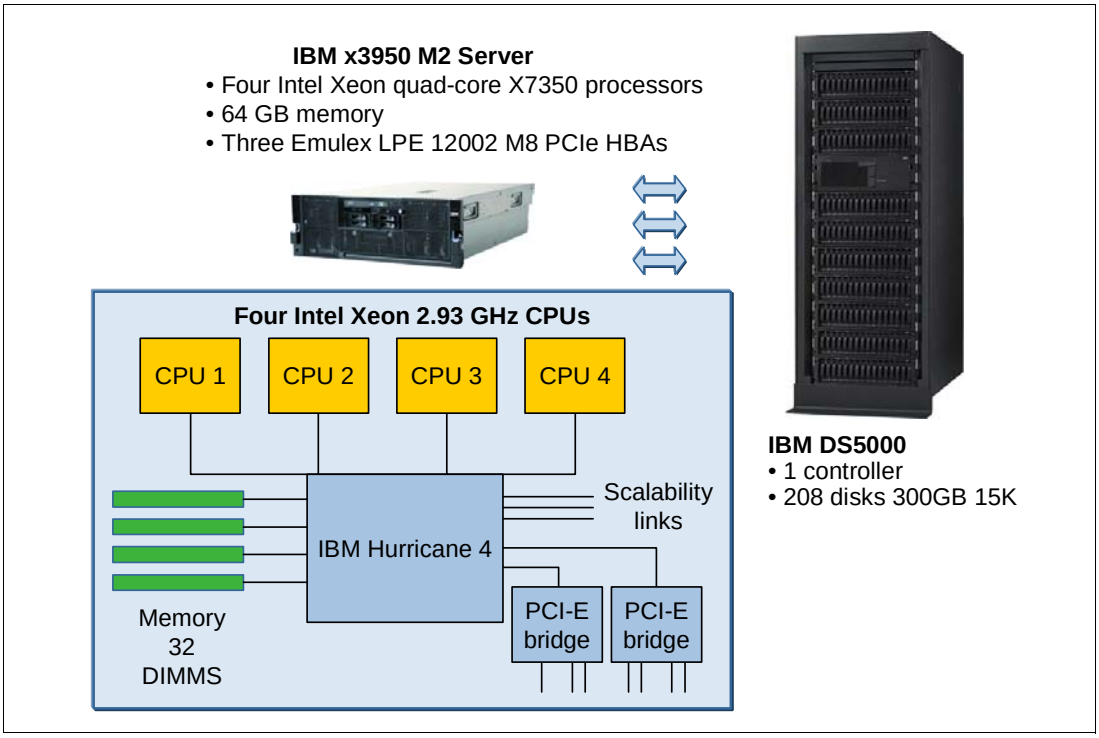

Figure 2 Server and storage configuration

Storage configuration

We attached 15 DS4000 EXP810 enclosures to the DS5300 dual-controller. Each EXP810 enclosure is populated with 16 300 GB 15,000 RPM drives for a total of 208 drives. The data warehouse consisted of 14 months of data. The drives are used as follows:

- Fourteen RAID-5 LUNs each of seven disks are designated to hold the 14 months of data.
- Two miscellaneous LUNs are designated to hold the remaining data in the Data Warehouse.
- Six disks in each of 14 enclosures, for a total of 84 disks, are set up as one large RAID-10 LUN for the OLAP cube.
- Four log LUNs are set up with four disks each.
- One of the enclosures with 16 drives is set up as a RAID-5 LUN for archiving the backup of the database.

Figure 3 shows the cabling between the DS5300, the EXP810 enclosures, and the x3950 M2 server. The cabling achieves the twin objectives of redundancy and balancing the traffic across all the components.

Any single component (that is HBA, DS5300 controller, or cable) can fail and the configuration will continue to operate correctly. In addition the traffic from the HBAs are divided equally between the two DS5300 controllers and each controller can access all of the disk enclosures with a maximum of two hops.

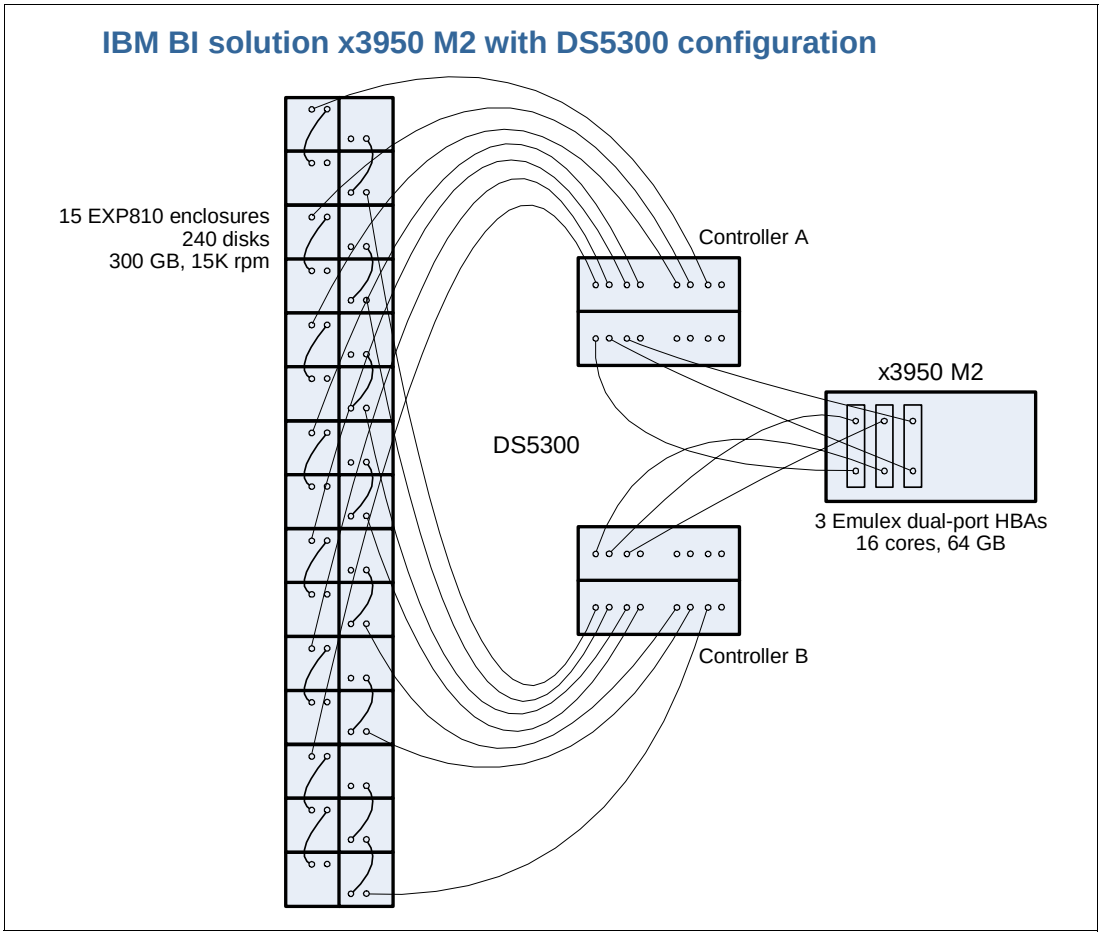

Figure 3 Cabling diagram for the x3950 M2 server and DS5300 storage

Figure 4 shows how the LUNs are configured. In Figure 4, the disks with the same color are part of the same LUN. The data warehouse LUNs run vertically, the cube LUN is the large mass of blue disks and the log and backup LUNs run horizontally.

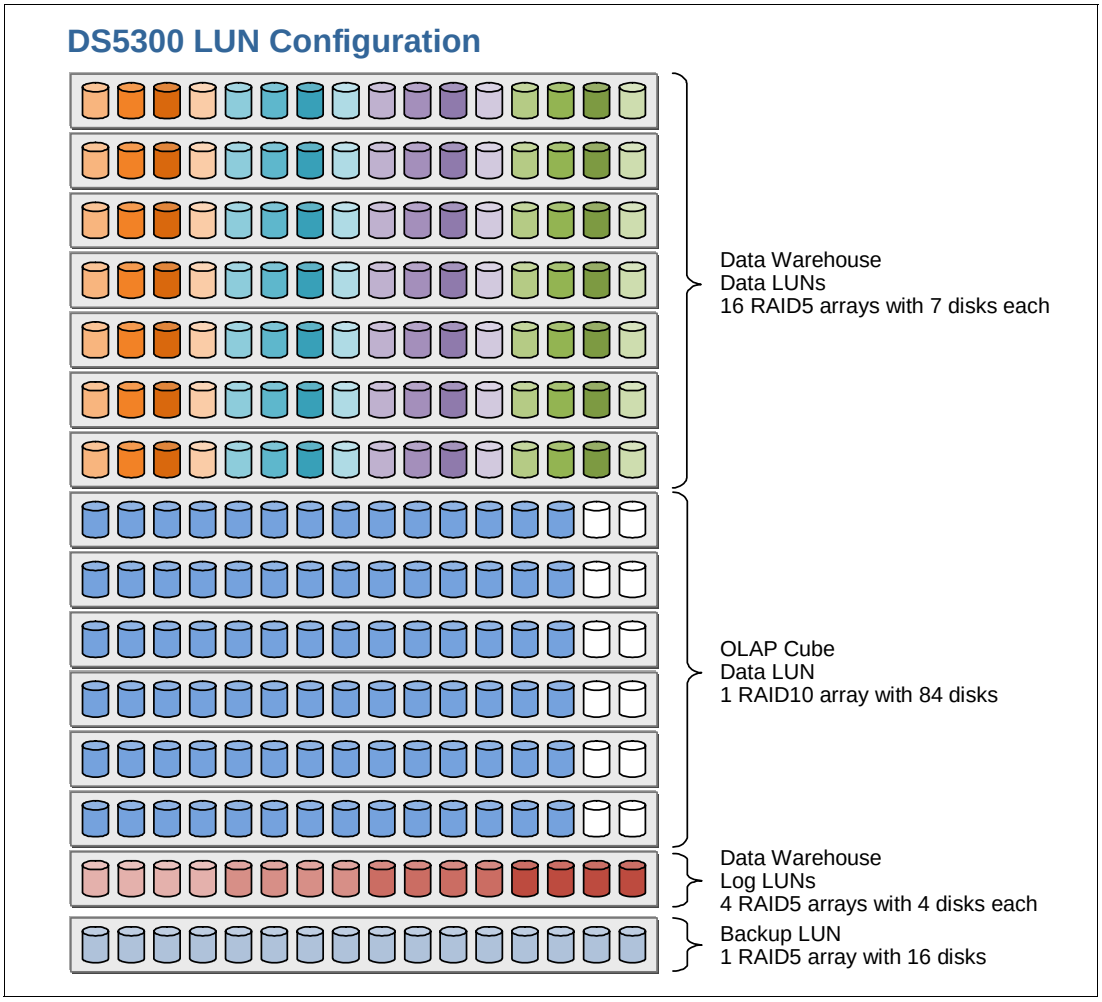

Figure 4 Configuration for the DS5300 LUN

Workload description

The tests in this configuration use a data warehouse of Web tracking information as a data source to create a multidimensional cube. A simple cube object is composed of basic information, dimensions, and measure groups. Basic information includes the name of the cube, the default measure of the cube, the data source, and the storage mode.

Dimensions are the actual set of dimensions used in the cube. All dimensions have to be defined in the dimensions collection of the database before being referenced in the cube. Measure groups are sets of measures in the cube. A *measure group* is a collection of measures that have a common data source view and a common set of dimensions. A measure group is the unit of process for measures. Measure groups can be processed individually and then browsed. The generated OLAP cube is then processed with default measure groups and partition sizes.

Data warehouse profile

The data warehouse used in this test is around 1.77 TB. It contains 26 dimension tables. The largest dimension has 77 million rows. There are 5 dimensions in the 100,000-row range, 3

dimensions in the 10,000-row range, 3 dimensions in the 1,000-row range. The remainder of the dimensions have fewer than 1,000 rows.

There are three main fact tables. The largest fact table has 5 billion rows, the next largest has 3.6 billion rows, and the third fact table has 2.4 billion rows. The fact tables are partitioned by date with a clustered index using an integer surrogate key (key = YYYY * 10,000 + MM * 100 + 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.

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. One measure group is created for each of the three fact tables. 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.

BI workload size

The data warehouse database size used in the test is around 1.77 TB, and the cube size ended up being around 412 GB. The cube consisted of three measure groups. The three measure groups ranged from 123 GB to 146 GB in size.

SQL Server 2008 configuration

SQL Server 2008 database engine and Analysis Services were installed on one SMP server and run in a shared configuration. Processor and memory for the database engine were limited by using the `sp_configure` settings for affinity mask and max server memory. Of the 16 processors and 64 GB memory on the server, the database engine was limited to 8 processors and 28,000 MB (28 GB) memory. We change the `sp_configure` settings shown in Example 1.

Example 1 Changes to sp_configure settings

```
sp_configure ‘affinity mask’, 0xFF
go
sp_configure ‘max server memory’, 28000
go
```

Cube processing

An Analysis Services Cube is a multidimensional structure that enables clients to access high volumes of pre-aggregated data and to extract useful data very quickly. A cube is a set of related measures and dimensions that is used to analyze data. When a cube is processed, the aggregations designed for the cube are calculated, and the cube is loaded with these calculated aggregations and data.

An example is product sales information aggregated by Time (Year, Month, and Day) and Product (Bicycles, Vendor, and Model). 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 OLAP query, including Reporting Services.

Results

A significant 41% reduction in cube processing times can be achieved by upgrading from SQL Server 2005 to SQL Server 2008 as depicted in Figure 5. A lower value is better.

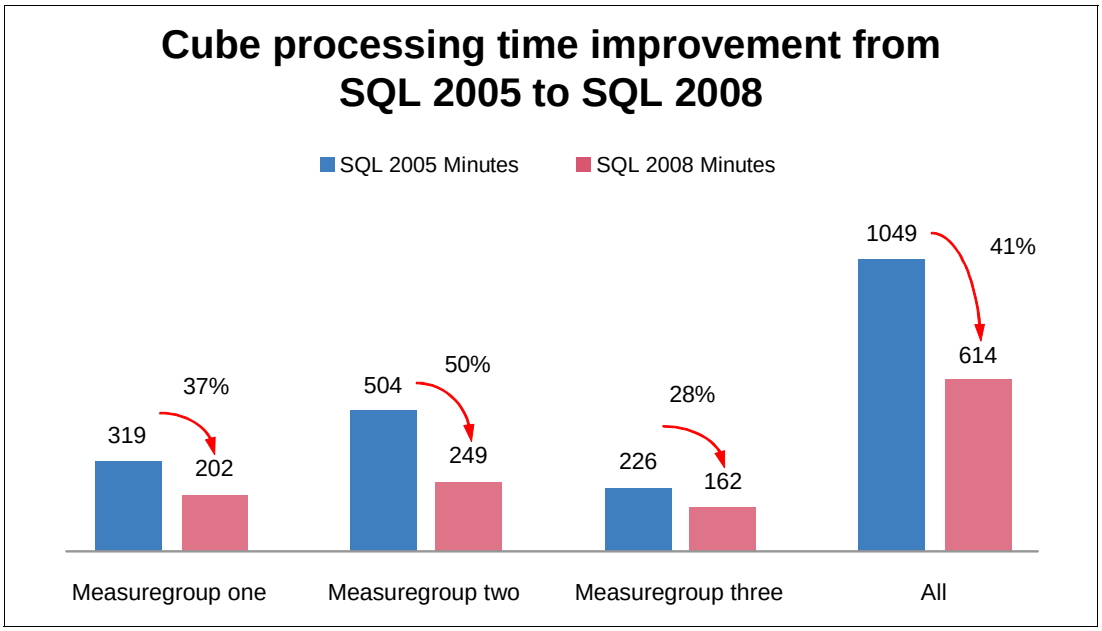

Figure 5 Cube processing improvement going from SQL Server 2005 to 2008

The data associated with Figure 5 is shown in Table 2.

Table 2 Cube processing times showing SQL Server 2005 to SQL Server 2008 improvements

Phase	Size in GB	SQL Server 2005 minutes	SQL Server 2008 minutes	Time reduction	Percentage improvement
Measure group one	123	319	202	117	37%
Measure group two	121	504	249	254	50%
Measure group three	146	226	162	64	28%
Total	412	1049	614	435	41%

The results in Table 2 show the performance effects of using x3950 M2 server with the improvements in SQL Server 2008. Improvements made in Analysis Services in SQL Server 2008 to run better on multi-core servers helped the workload to scale better on x3950 M2. These changes allowed us to run more threads in parallel thus reducing cube processing time.

IBM BI reference configurations

In this section, we describe the three different reference configurations tested by IBM. These reference configurations show how cusomters can reduce cube processing times by adding processor and memory resources to x3950 M2 servers highlighting it scale-up features.

In these configurations, the Relational Data Warehouse (RDW) size is 1.8 TB and the OLAP Cube size is 412 GB.

Table 3 Tested reference configurations

	Configuration 1	**Configuration 2**	**Configuration 3**
Server configuration	► Single-node x3950 M2 server ► 4 Intel Xeon X7350 Processors (QC, 2.93 GHz) ► 64 GB DDR2 memory ► 3 Emulex 8 Gb HBAs	► Single-node x3950 M2 server ► 4 Intel Xeon X7460 Processor (6-core, 2.66 GHz) ► 96 GB DDR2 memory ► 3 Emulex 8 Gb HBAs	► 2-node x3950 M2 server ► 8 Intel Xeon X7350 Processors (QC, 2.93 GHz) ► 128 GB DDR2 memory ► 6 Emulex 8 Gb HBAs
Storage configuration – RDW	► 1 DS5300 controller ► 2 x 72 GB RAID-1 15,000 SFF drives (OS and SQL drive) ► 4 RAID-5 7-disk arrays for fact data ► 1 RAID-10 6-disk array for tempdb, database log, and OLAP temp directory	► 1 DS5300 controller ► 2 x 72 GB RAID-1 15,000 SFF (OS and SQL drive) ► 5 RAID-5 7-disk arrays for fact data ► 1 RAID-10 6-disk array for tempdb, database log, and OLAP temp directory	► 1 DS5300 controller ► 2 x 72 GB RAID-1 15,000 SFF (OS and SQL drive) ► 7 RAID-5 7-disk arrays for fact data ► 1 RAID-10 6-disk array for log ► 1 RAID-10 6-disk array for tempdb and OLAP temp directory
Storage configuration – OLAP Cube	1 RAID-10 14-disk array	1 RAID-10 20-disk array	1 RAID-10 30-disk array
Cube building and processing time	7 hours and 5 minutes	6 hours and 4 minutes	4 hours and 36 minutes
Scaling	1x	1.17x	1.32x over 24 cores 1.54x over 16 cores

Figure 6 shows a chart of the elapsed time for each of the three configurations when processing the cube.

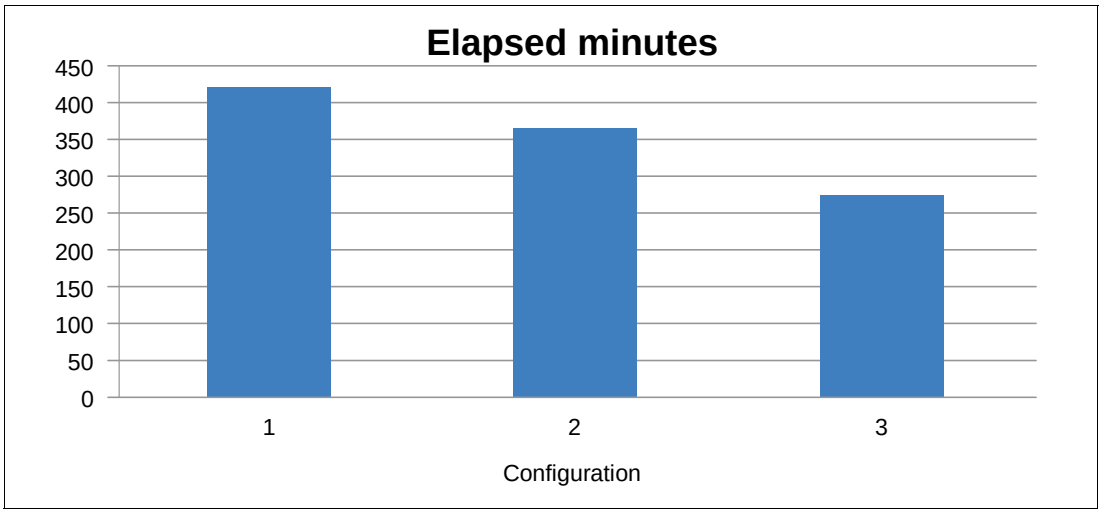

Figure 6 Elapsed time in minutes for each configuration (lower is better)

Figure 7 shows the rate at which rows are read and written during cube processing for each configuration. Observe that rates increase as we add more processors and memory to the configuration. This results in reduced cube processing times demonstrating the scalability of the platform. With the modular scale-up design of the x3950 M2 and the scalable Intel Xeon Processor 7400 Series, adding processors and memory to your configuration is simple and just requires cabling the additional chassis and rebooting the server for the operating system and SQL server to start using the added resources.

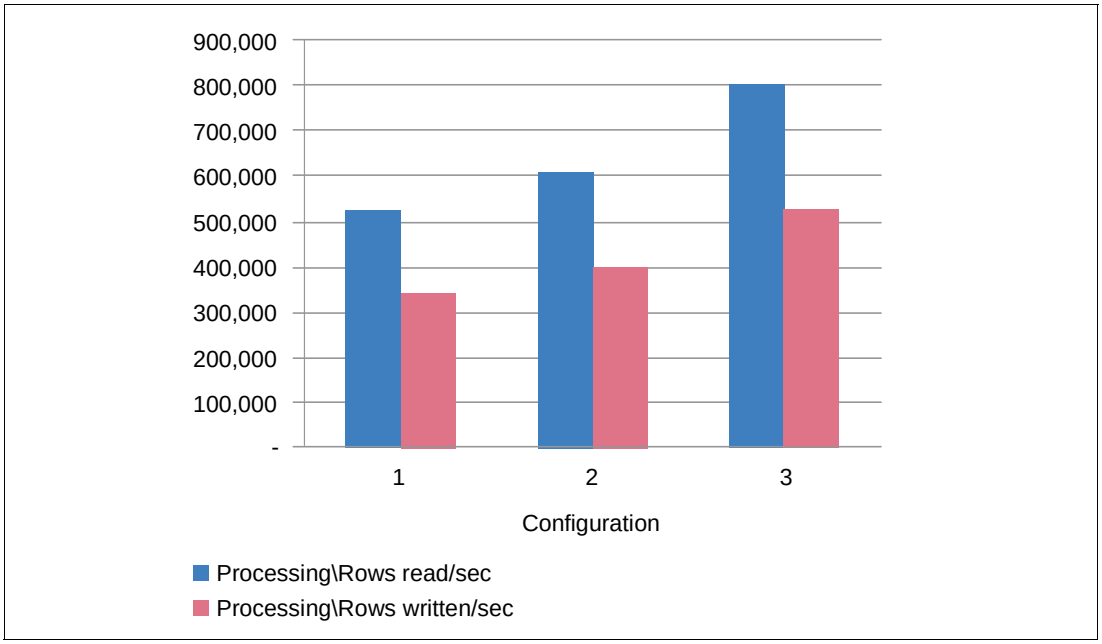

Figure 7 Rate of rows read and written during cube processing (higher is better)

Table 4 shows the data that corresponds to Figure 7.

Table 4 Data for Figure 7

Microsoft SQL Server 2008 Analysis Services counter in Windows PerfMon	Configuration 1	Configuration 1	Configuration 3
MSAS 2008 : Processing : Rows read/sec	523,268	610,705	803,715
MSAS 2008 : Processing : Rows written/sec	343,036	400,378	526,185

Table 5 shows that the elapsed time multiplied by the rows read/sec rate gives approximately the same number of total rows (13.3 billion rows read for each configuration). Likewise, for the rows written/sec rate, the total is approximately 8.7 billion rows written for each configuration. The data in the table shows that the same amount of data was processed in each configuration what with the database size being constant.

Table 5 Row processing rate x elapsed time gives approximately equal number of total rows

Row rate * Elapsed time	Configuration 1	Configuration 2	Configuration 3
Total rows read	13,343,327,013	13,337,798,401	13,309,522,503
Total rows written	8,747,418,842	8,744,260,237	8,713,622,524

Tuning SQL Analysis Services for DW and BI workloads

This section describes the changes made to the Analysis Services properties to get the best performance for the reference configurations. Table 6 lists the properties that we changed.

Table 6 Performance improvements

Property	Optimized value	Original value
CoordinatorExecutionMode	-8	-4
Memory\LowMemoryLimit	70	75
OLAP\Process\AggregationMemoryLimitMin	1	10
OLAP\Process\AggregationMemoryLimitMax	5	80
OLAP\Process\DatabaseConnectionPoolMax	64	50
ThreadPool\Process\MaxThreads	128	64

The Analysis Services properties are as follows:

- CoordinatorExecutionMode

  This property determines the maximum number of jobs that can execute in parallel for the current processing and query operations. A negative value specifies the maximum number of parallel jobs that can start per core per operation. A value of 0 indicates no limit. A positive value specifies an absolute number of parallel jobs that can start per server.

  The default value for the CoordinatorExecutionMode is -4, which indicates that four jobs will be started in parallel per core. This value was changed to -8 to indicate that 8 jobs will be started in parallel per core. This was necessary to exploit the large number of processor threads available on the servers and translate that into better performance.

- Memory\LowMemoryLimit

  This property defines when the server is deemed to be low on physical memory. Values greater than 100 are absolute values. Values less than 100 are calculated as a percentage of memory. The reference value of memory (that corresponds to 100) is calculated as the

minimum between the physical memory and the virtual memory available for a process. An 8 GB RAM 64-bit Windows has an 8 GB limit.

The default is 75 (75%), which is appropriate most of the time, but if other processes are running on the same machine, it might be preferable to lower this setting. If you want to know what the actual and running absolute limit is, you can get the exact running value reading the MSAS 2008 Memory/Memory Limit High KB performance counter. We measured the best performance when this value was changed to 70.

- OLAP\Process\AggregationMemoryLimitMin

  This property is a signed 64-bit double-precision floating-point number property that defines the minimum amount of memory that can be devoted to aggregation processing, expressed as a percentage of physical memory. A larger value can speed up aggregation processing at the cost of memory usage. Alternatively, a smaller value might permit more partitions to be processed in parallel because the memory is spread out among more processes.

  The default value for this property is 10, indicating that a minimum of 10% of physical memory will be devoted to aggregation processing. The optimal value was determined to be 1 in our tests. Lowering this limit made allowed more partitions to be processed giving us an improvement in overall processing time.

- OLAP\Process\AggregationMemoryLimitMax

  This property defines the maximum amount of memory that can be devoted to aggregation processing, expressed as a percentage of physical memory.

  The default value for this property is 80, indicating that 80% of physical memory can be devoted to aggregation processing. The optimal value was determined to be 5 in our configuration. Again, lowering this limit made allowed more partitions to be processed giving us an improvement in overall processing time.

- OLAP\Process\DatabaseConnectionPoolMax

  This property specifies the maximum number of pooled database connections.

  The default value for this property is 50. We increased the value to 64 to allow more database connections for better performance.

- ThreadPool\Process\MaxThreads

  This value determines the maximum number of available threads to Analysis Services during processing. On large, multiple-processor machines, the default value of this setting might be too low to take advantage of all processor cores. However, as you increase this counter, bear in mind that increased parallelism of processing also has an effect on queries running at the system level. As you dedicate more processors and threads to processing, less processor will be used for query responses. Of course, if you are processing the cubes during a batch window, this might not be an issue. The optimal value we determined for our configuration was 128 which increased the processor utilization giving better performance. The default value is 64.

Conclusions

The major components used in these reference configurations are:

- IBM System x3950 M2 server with Intel Xeon processors
- IBM System Storage DS5000 storage running
- Microsoft Windows Server 2008 and
- Microsoft SQL Server 2008

Based on the results and analysis that we present in this paper, these reference configurations provide an excellent platform for clients to design a cost-effective, scalable BI solution in an x86 environment. Microsoft has made many improvements in Analysis Services, which is seen in improved scaling as demonstrated in this paper.

The modular design and industry-standard x86 architecture of IBM System x servers and storage enable the client to start small and grow as their databases grow and their computing needs increase.

This is demonstrated in the reference architectures where:

- Adding processors and memory seamlessly to the BI configurations reduces the cube processing time significantly. The results show the benefit that the additional memory and cores in each configuration provides.
- We see increases in the rate at which rows are processed in each successive configuration, which in turn reduces the total elapsed time required to build the cube. Since the size of the cube was the same for each configuration, the elapsed time is shorter for each more powerful hardware configuration.

When a cube is processed faster, the data is available to analysts sooner and it is possible to process larger amounts of data. With analysts operating on a more current and comprehensive view of the company's business data, the business intelligence derived from the data is more relevant and valuable for making decisions strategic to the business.

The team who wrote this guide

This guide was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO).

Vinay Kulkarni has a Bachelor of Computer Engineering degree from Shivaji University in India. He studied his masters in computer science at 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 8 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.

Michael Lawson is a database and virtualization specialist at the IBM Center for Microsoft Technologies in Kirkland, WA. He has a B.A. 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 Sydney 2000 Olympics. 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. He holds the Microsoft Certified Database Administrator certification. Michael is also an author of two other IBM Redbooks® publications, including *Running Microsoft SQL Server 2008 on the IBM System x3950 M2*, REDP-4476.

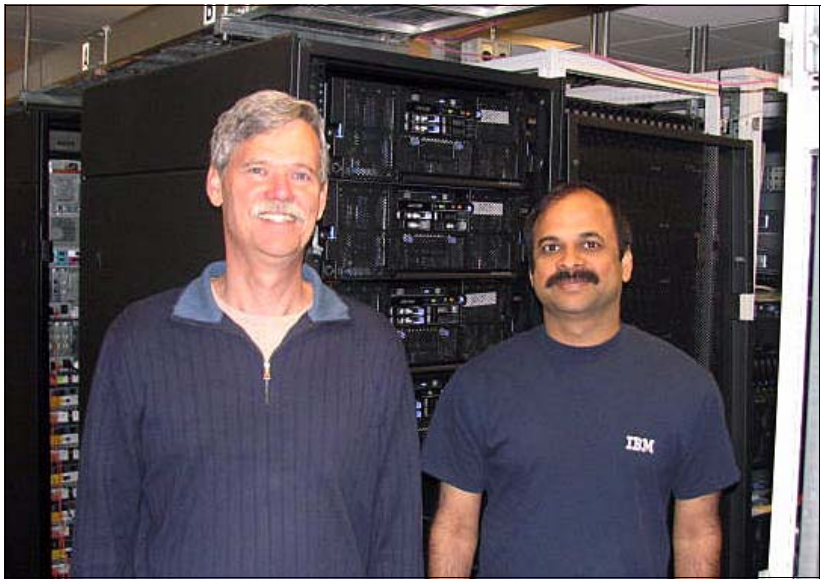

Figure 8 The authors (left to right): Michael and Vinay

References

You might also find the following resources helpful in the discussion of this topic:

- 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.

This document, REDP-4552-00, was created or updated on July 24, 2009.

Trademarks