



Business Intelligence Solutions on the IBM System x3850 X5 Using XIV Storage



Redguides
for Business Leaders

Vinay Kulkarni
Michael Lawson

- Fourth paper in the series on Business Intelligence solutions from IBM System x and IBM System Storage
- Tested reference configurations using IBM System x3850 X5 with Microsoft SQL Server 2008 R2 BI
- IBM XIV System Storage benefits for Business Intelligence: Ease of use, performance, and data protection



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 IBM eX5 servers with more processors and memory than ever before in an x86 platform, this IBM Redguide™ publication shows you how you can get the most out of one of these high-performance servers in a BI configuration. The IBM eX5 servers allow you to scale to new performance highs, process your cubes quickly, and get query results even faster than on previous-generation M2 servers.

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

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

This Redguide is the fourth in a series that describes Business Intelligence reference configurations on the 4-socket IBM x3850 X5 platform based on Intel® 7500 series processors. The first three Redguides™ covered reference configurations on IBM x3850 M2, IBM System Storage® DS5000, and Microsoft Windows Server 2008 and SQL Server 2008.

This document provides sizing information for small, medium, and large businesses that are planning to use the Microsoft Business Intelligence solution utilizing IBM System x3850 X5 server, IBM XIV Storage System, Microsoft Windows Server 2008 R2, and SQL Server 2008 R2 components. Three reference configurations were tested, which are listed in Table 1.

Table 1 BI reference configurations with results

| | Small (1X) | Medium (2X) | Large (4X) |
|--------------------|--|---|---|
| DW database size | 1.7 TB | 3.6 TB | 6.5 TB |
| Server and storage | <ul style="list-style-type: none"> ▶ 32 logical processors ▶ 64 GB memory ▶ 1 XIV, 79 TB usable | <ul style="list-style-type: none"> ▶ 64 logical processors ▶ 128 GB memory ▶ 1 XIV, 79 TB usable | <ul style="list-style-type: none"> ▶ 64 logical processors ▶ 256 GB memory ▶ 1 XIV, 79 TB usable |
| Cube building time | 6 hours 19 minutes | 6 hours 50 minutes | 12 hours 2 minutes |
| % of 1X | 100% | 108% | 190% |

Table 1 provides a high-level summary of the reference configurations discussed in this paper and the processing results. These results show excellent scaling (going from the small to medium configuration) when processing a data warehouse (DW). Doubling the size using double the hardware takes approximately the same amount of time (108%) to build a cube. When scaling from the medium to large configuration, we doubled the server memory, but because SQL Analysis Services cannot use more than 64 logical processors, we did not increase the number of logical processors. In this case, we expected approximately double the processing time, and we actually get a little better result than that (190%), which can be attributed to doubling the memory. This is another data point confirming good scaling.

The IBM System x3850 X5 server is a unique and compelling high-end modular server with up to four Intel Xeon® 7500 series processors 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 to build a scale-up server with up to 128 logical processors and 2 TB memory. The x3850 X5 can be configured with a 1U MAX5 memory drawer to build a 4-socket, 1.5 TB memory configuration.

This Redguide reports on three scenarios in the small, medium, and large range based on DW size. Comparable runs to the small and medium configurations in this paper with a previous Redguide takes approximately the same amount of time when using the IBM System Storage DS5000, confirming that the XIV is delivering comparable performance.

In conclusion, we show that the new eX5 servers and XIV storage process your data and derive intelligent information from it faster than you could ever do it before.

Introduction

The purpose of 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. 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 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 and is very dependent on the hardware resources provided and the specifications of the cubes themselves (for example, the amount and complexity of aggregation).

IT professionals need to know approximately how long their cubes will take to build so that they can plan effectively and have the cubes ready when their customers need them. To assist with this task, this guide provides three examples of data warehouses and hardware configurations with 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 guide 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 guide is made up of three sections:

- ▶ Solution architecture
- ▶ Workload description
- ▶ Reference configurations

Solution architecture

Solution components include the high-end IBM System x3850 X5, IBM XIV Storage System, 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 7500 series processors with up to 8 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, 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 scalable Intel Xeon processor 7500 series. Built to handle the most demanding applications, the Intel Xeon processor 7500 series delivers a quantum leap 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 XIV Storage System

The IBM XIV Storage System is a revolutionary, high-end, disk storage system with enterprise-level, consistent performance through massive parallelism, disk utilization, and unique caching mechanisms, while at the same time eliminating the complexity of storage

administration and making storage management easy. Figure 2 shows the XIV Storage System.



Figure 2 IBM XIV Storage System

The IBM XIV Storage System is characterized by the following set of features:

- ▶ iSCSI and Fibre Channel (FC) interfaces
- ▶ Multiple host access
- ▶ Management software, including graphic user interface (GUI) and command-line interface (CLI)
- ▶ Volume cloning (snapshots)
- ▶ Thin provisioning
- ▶ Asynchronous and synchronous replication of a volume to a remote system
- ▶ Easy assignment and reassignment of storage capacity

Note: The term *storage capacity* here refers to the total storage capacity, and does not take into account the amount of storage used for data-redundancy or mirroring and other data-related tasks.

- ▶ Remote configuration management
- ▶ Notifications of events through email, SNMP, or SMS messages
- ▶ Truly no single-point-of-failure
- ▶ Fault tolerance, failure analysis, and self-healing algorithms
- ▶ Non-intrusive maintenance and upgrades

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:

- ▶ Windows and SQL Server 2008 R2 Datacenter Editions support up to 256 logical processors, a 4x increase from the previous product.
- ▶ PowerPivot 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, while database professionals can continue to administer the data in the solutions.
- ▶ Master Data Services 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) 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 needed to provide a comprehensive, low-cost business intelligence solution.

SQL Server 2008 R2 BI components

SQL Server 2008 R2 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 On Line 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.

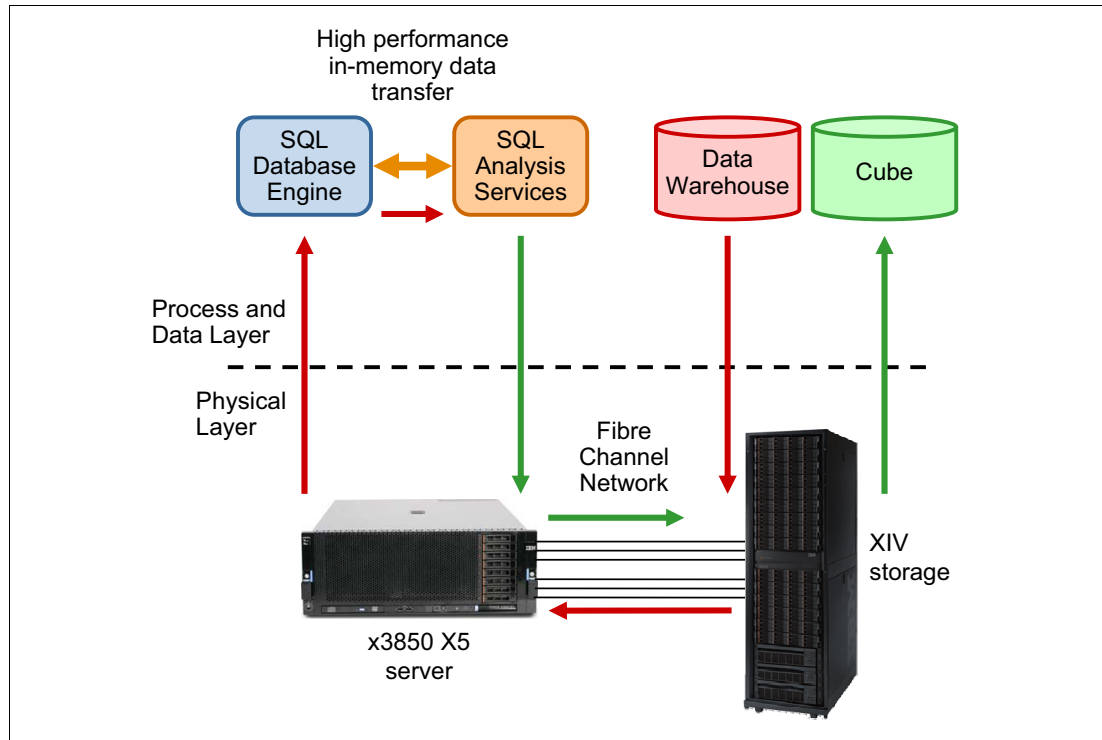


Figure 3 IBM BI reference architecture

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 XIV, across the Fibre Channel network, and into the x3850 X5 server where the 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 XIV, 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 DW come from Microsoft.

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

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

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

BI terminology

This section lists a few key terms used in BI.

- ▶ Fact tables

Fact tables contain the individual transactions or events that are relevant to the business. For example, a retail business has 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 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 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, for example, 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 Multidimensional Expressions (MDX) query language. Several products support MDX query, including Microsoft Reporting Services.

Reference database design

The Web Tracking database and application provided by Microsoft performs database mining of 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 help 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 guide.

Data warehouse profile

The data warehouse, referred to below as the small (1X) DW used in this test, is 1.7 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 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 data warehouse profile, 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 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 .58 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 mentioned in the query performance section.

Reference configurations

In this section we provide the tested reference configurations, including the parameters settings used (see “Appendix” on page 16) 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) | Medium (2X) | Large (4X) |
|------------------------|--|--|--|
| Data warehouse | | | |
| Size | 1.7 TB 93% space used 11 billion fact rows | 3.6 TB 93% space used 22 billion fact rows | 6.5 TB 68% space used 44 billion fact rows |
| Cube size | .58 TB | 1.01 TB | 1.88 TB |
| Cube building time | 6 hours 19 minutes | 6 hours 50 minutes | 12 hours 2 minutes |
| Time % of 1X | 100% | 108% | 190% |
| Server | | | |
| Model | One x3850 X5 server | One x3850 X5 server | One x3850 X5 server |
| Processors | Two Intel Xeon X7560 processors 8C, 2.26GHz | Four Intel Xeon X7560 processors 8C, 2.26 GHz | Four Intel Xeon X7560 processors 8C, 2.26GHz |
| Memory | 64 GB memory | 128 GB memory | 256 GB memory |
| HBAs | Four Emulex FC HBA ports | Four Emulex FC HBA ports | Four Emulex FC HBA ports |
| IBM XIV Storage | | | |
| Total drives | 180 | 180 | 180 |
| DW – data | 4 TB, 1.7 TB used | 7 TB, 3.6 TB used | 14 TB, 6.5 TB used |
| SQL Server TempDB data | 1 TB, 0.6 TB used | 1 TB, 0.6 TB used | 1 TB, 0.6 TB used |
| SQL Server | | | |
| logs | 500 GB, 116 GB used | 500 GB, 116 GB used | 500 GB, 116 GB used |
| OLAP cube | 1 TB, .6 TB used | 2 TB, 1 TB used | 4 TB, 2 TB used |
| OLAP temp | 500 GB, 115 GB used | 500 GB, 115 GB used | 500 GB, 115 GB used |

Storage configuration

The XIV was extremely easy to use, performed flawlessly, and provided excellent data protection. When compared to conventional and previous-generation storage systems, the

XIV provides better value because ease of use translates into lower administrative costs, while providing equal or better performance and data protection.

XIV architecture

Benefits come from the revolutionary XIV architecture. Each logical volume in the XIV is divided into multiple stripes of one megabyte. These stripes are distributed evenly across all disks in the system using a sophisticated pseudo-random distribution mechanism. Every data write is written simultaneously to two disks for redundancy. This takes the hard work out of storage configuration. All you need to specify is the size of the logical volumes. The XIV architecture also makes it simple to resize volumes as workload and storage needs change.

Comparing conventional and previous-generation storage systems

With conventional and previous-generation storage systems, it is necessary to make many detailed decisions when configuring the storage system. Configuration choices include selecting the individual disk drives or the number of disk drives that make up a logical volume, whether to share disk drives between logical volumes, choosing the RAID level, choosing a preferred controller, and selecting hot spare disks and the policy for their use. These decisions are made on the basis of capacity, data protection, and performance needs, which are competing objectives and thus result in making tradeoffs.

Once the configuration is made, it is difficult to change. This means that significant time and effort must be expended up front when initially configuring the storage, and again any time that the workload changes. The time and effort required to change the configuration is even greater when large amounts of existing data must be moved. A degree of risk to the data is involved, so it is prudent to back up the data prior to reorganizing the configuration, adding to the expense of any change operation.

XIV ease of use

Creating volumes on XIV only requires specifying the size of the volume. No placement information or RAID setting is needed. Therefore, it is simple to create the volumes. Another benefit is that the volumes are available for use instantly, without any performance impact. There is no waiting time while the volumes are initialized. On conventional storage systems, it can sometimes take many hours to initialize a large volume. During this time the volume might not be available or the initialization process might affect performance. It is also simple to grow and shrink volumes, and volume size does not affect disk performance. Again, the XIV compares favorably with conventional storage. Certain conventional storage systems do not allow resizing volumes, or only allow growth, or impose a performance penalty on the volume once it is resized.

XIV performance

The XIV performed well on this cube-building workload because of the XIV architecture. All 180 1TB disk drives were used by all data volumes, providing sufficient I/Os per second and sufficient bandwidth.

XIV data protection

Because the XIV always writes data to two locations, the risk of data loss is much lower than RAID-5, the most common data protection choice with conventional storage. Availability is better in failure scenarios, as rebuild time is much longer with RAID5 than with the XIV method, which is very similar to RAID10 (mirroring).

Scaling results

In this section, we discuss two scaling results:

- ▶ Small (1X) to medium (2X)
- ▶ Medium (2X) to large (4X)

The X above refers to the size of the data warehouse. The 1X DW is 1.7 TB in size, 2X DW is 3.6 TB, and 4X DW is 6.5 TB in size.

Small to medium scaling

The scaling results are excellent going from small (1X) to medium (2X), as evidenced by holding the run duration to nearly the same while doubling the size of the DW. This shows a nearly linear benefit from doubling the processor and memory resources while doubling the database size.

Medium to large scaling

The scaling results are also excellent going from medium to large, as evidenced by the elapsed time being approximately double for double the work (because the DW is twice the size). In going to the large configuration, we keep the same number of cores, double the amount of memory, and keep the same number of disks. We also use the same set of tuning parameters with proportional settings. For example, Analysis Services uses 50% of the memory and SQL Server uses 30%, leaving the remaining 20% for the operating system and the system file cache. SQL Server uses half the logical processors, 32 out of 64 processors in both configurations.

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 disk traffic in total MBps (read plus write) when building the cube in each of the three test configurations. The medium configuration reads and writes at a higher rate (31%) than the small configuration, and processes twice the amount of data in the same amount of time. The large configuration reads and writes at a slightly more accelerated rate (4%) than the medium configuration, and processes twice the amount of data in twice the amount of time.




Figure 4 Disk bandwidth usage

Rows read per second from the DW

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

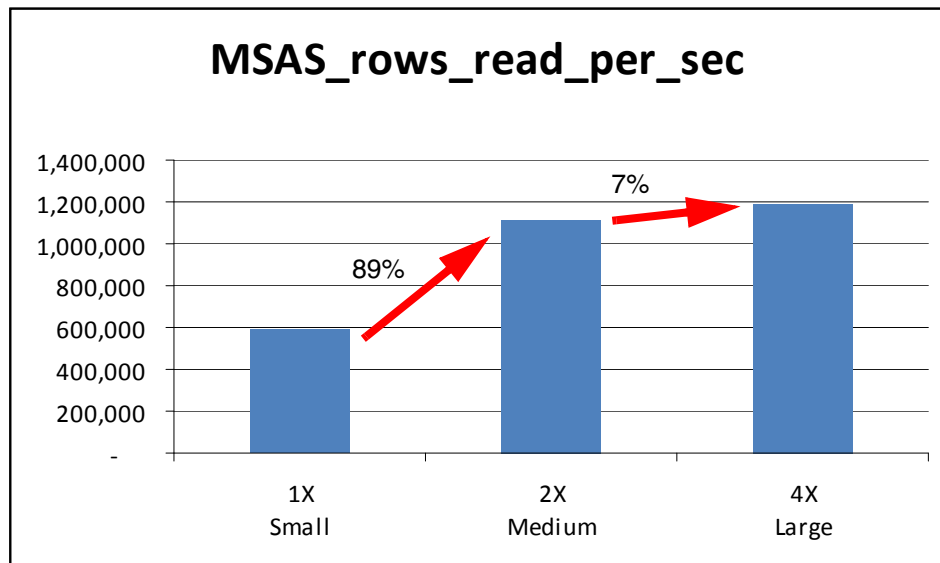


Figure 5 Rows read per second from the DW

Rows written per second to the cube

Figure 6 shows the number of rows written to the cube while building the cube. Higher is better. The rate of rows written to the cube increases in proportion to the increased workload.




Figure 6 Rows written per second to the cube

Query performance

After building the cube, we tested 50 concurrent users, each running five queries against the cube, and recorded the response time in seconds. We did this for each of the cube sizes:

- ▶ 1X
- ▶ 2X
- ▶ 4X

Figure 7 shows the 90th percentile response time in seconds for 50 concurrent users, each running the same five queries. This chart shows that even with 50 concurrent users running queries against the cube, the server provides consistent response times for all three cube sizes (1X, 2X and 4X).

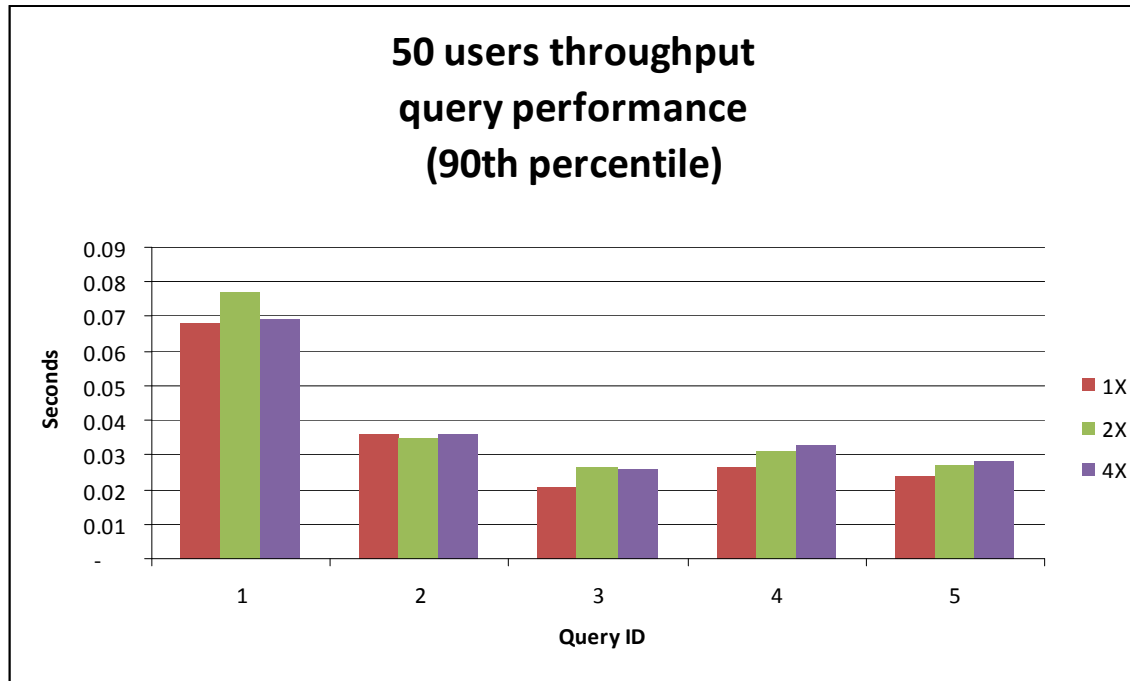


Figure 7 50 concurrent users running five queries at 1X, 2X, and 4X cube sizes

Conclusion

In this paper 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, results show the configurations to be well-suited for Microsoft SQL DW and BI applications:

- ▶ These configurations show excellent, nearly 100% scaling when going from a 1.7 TB DW on a two-socket 32-core x3850 X5 to a 3.6 TB DW on a four-socket 64-core x3850 X5.
- ▶ These configurations show that the IBM XIV Storage System provides excellent performance and ease of use for this business intelligence workload.
- ▶ Complex cube processing and querying is performed in a timely manner using the x3850 X5 server, which supports the extraction of actionable insights from a company's transactional data, the goal of business intelligence solutions.

Appendix

Hardware and software specifications

The specifications were the same for all configurations and are listed in Table 3.

Table 3 Hardware and software specifications

| Component | Description |
|------------------------------------|---|
| System | IBM System x3850 X5, 4x Intel Xeon X7560 processors 8C, 2.26 GHz |
| Server local drives: OS and SQL | 1 RAID-1 array, 2 drives, 146 GB, 10K rpm, SFF |
| Memory | DDR3 10667 MHz |
| SAN disks | 1 TB, 7200 RPM, SATA |
| SQL Server | Microsoft SQL Server 2008 R2 – 10.50.1600.1 (x64) Enterprise Edition (64-bit) |
| Windows | Windows Server 2008 R2 Enterprise Edition x64 Build 7600 |

Tuning parameters

The tuning parameters listed in Table 4 were used in each of the reference configurations.

Table 4 Tuning parameters

| | Small (1X) | Medium (2X) | Large (4X) |
|---|------------|-------------|------------|
| SQL Server | | | |
| Tempdb files | 64 | 64 | 64 |
| sp_configure: max server memory | 30% | 30% | 30% |
| sp_configure: affinity mask | 0xFFFF | 0xFFFFFFFF | 0xFFFFFFFF |
| sp_configure: Affinity64 mask | 0x0 | 0xFFFFFFFF | 0xFFFFFFFF |
| sp_configure: max degree of parallelism | 1 | 1 | 1 |
| Analysis Services | | | |
| Msmdsrv.ini: OLAP\Process\AggregationMemoryLimitMin | 1% | 1% | 1% |
| Msmdsrv.ini: OLAP\Process\AggregationMemoryLimitMax | 5% | 5% | 5% |
| Msmdsrv.ini: Memory\TotalMemoryLimit | 50% | 50% | 50% |
| Msmdsrv.ini: Memory\LowMemoryLimit | 30% | 30% | 30% |
| Msmdsrv.ini: Memory\PreAllocate | 35% | 35% | 35% |
| Msmdsrv.ini: ExternalCommandTimeout | 3,600 | 3,600 | 3,600 |
| Create cube xmla: MaxActiveConnections | 32 | 64 | 64 |
| Process cube xmla: Parallel MaxParallel | 32 | 64 | 64 |

References

See the following sources for more information:

- ▶ 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://msdn.microsoft.com/en-us/library/ms130214.aspx>
- ▶ SQL CAT (Customer Advisory Team) Analysis Services resources
<http://sqlcat.com/tags/Analysis+Services/default.aspx?PageIndex=1>

The team who wrote this guide

Vinay Kulkarni has a bachelor's degree in Computer Engineering from Shivaji University in India. He has a 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 onsite at the Microsoft Redmond campus. He has been working with Microsoft for 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 co-author of the three other IBM Redguides in this series.

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 the 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 three other IBM Redguides in this series.




Figure 8 The authors (l-r): Mike and Vinay

We would like to acknowledge the Microsoft SQL team for providing the Web Tracking DW and the associated workloads.

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

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:
<http://www.facebook.com/IBMRedbooks>
- ▶ 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. Certain 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites 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-4721-00, was created or updated on January 4, 2011.




Trademarks

IBM, the IBM logo, and [ibm.com](http://www.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 Storage® |
| IBM® | Redguides™ | System x® |
| Redbooks® | Redbooks (logo)  ® | XIV® |

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.