

# SAS and IBM BI Reference Test – Why System p for SAS Business Intelligence



Recently SAS and IBM proved the power of their alliance by performing a SAS® Business Intelligence Benchmarking Scenario test suite that leveraged the power of IBM's technology. This Reference Architecture was developed by SAS as a means to show customers the scalability of their solutions on various hardware platforms.

The IBM System p family of products brings the technology, products and services to develop and deploy SAS® Business Intelligence solutions successfully and at a very competitive return on investment.

Business Intelligence solutions have unique functional and operational infrastructure requirements. These include performance, availability, maintainability, and scalability across n-tier architectures. BI solutions differ from traditional transactional types of workloads. BI solutions are typically very I/O intensive. They invariably are about gathering a large number of very small records, resulting in a smaller number of large transactions and disseminating them.

Business Intelligence solutions stress the limits of platform functionality. The workload characteristics are typically so varied that all the resources of the system's architecture are exercised. The amount of data being processed and the associated vast demand for raw processing power is huge. The amount of data be processed is also growing at a phenomenal

rate. These put scalability and performance at the top of the criteria for selecting a Business Intelligence solution.

## **SAS The power to Know**

Clients trust SAS to help them drive profitable growth, unlock the value of their existing business data, and help ensure the security and integrity of their global data warehouse

SAS provides complete Business intelligence solutions. IBM provides industry leading technologies and products that help maximize the use of SAS software systems in your business today.

## **IBM System p – Systems to fit every requirement**

Any BI solution is going to require a robust infrastructure made up of servers. The IBM System p family of products, in conjunction with the AIX 5L operating system, provides unparalleled functionality and performance for implementing Business Intelligence solutions of all sizes - from small business operations through large-scale enterprise-wide data warehouses. The IBM System p product line encompasses:

- Enterprise servers designed for a broad range of applications required by medium and large companies.
- Entry servers delivering affordable power, reliability and scalability to departments and small businesses.
- Workstations delivering a broad range of features for various commercial needs.

The performance of the 64-bit IBM POWER5™ platform, enabled by simultaneous multi-threading and the IBM AIX 5L™ operating system, is ideal for all Business Intelligence processing needs. IBM's continued investment in the AIX 5L system and unique advanced virtualization technologies like the IBM Micro-Partitioning™ capability enable UNIX system users alike to do more with a single system.

The IBM System p servers are designed for agility to enable you to quickly respond to changes in UNIX processing requirements, so you can meet anticipated and unanticipated business demands. System p servers provide the ability to dramatically increase system utilization in a consolidated server environment. IBM System p servers are designed to provide a highly secure UNIX processing platform that helps safeguard data from unauthorized access. System p systems also make migrations and upgrades easy, so anyone, including existing IBM clients and users of competing systems, can take advantage of these game-changing capabilities designed to provide excellent overall value today—and tomorrow

### **The IBM TotalStorage™ Difference**

IBM provides organizations what they need to make the right storage decision for both now and the future.

IBM provides a focus on open standards, support for heterogeneous environments and continuous leverage of innovative systems technologies, including the IBM Power Architecture™ platform. This storage system approach provides the following key benefits:

- TotalStorage solutions that are designed to provide a lower total cost of ownership.
- Logical partitioning (LPAR) functionality on select combined with leading IBM POWER performance, permit new efficiencies of scale.
- IBM TotalStorage SAN Volume Controller and TotalStorage SAN File System (SFS) software enable the virtualized management of heterogeneous storage resources.

IBM TotalStorage solutions can help provide industry-leading business continuity capabilities, to help keep your information readily accessible when needed for critical business decisions.

### **Web serving**

An organization's web serving needs are ideally supplied through the IBM System p5 520 and p5 550 product offerings. These systems are

packaged as a drawer for mounting in a 19-inch rack enclosure. They also provide exceptional processing power, high throughput and low latency with a minimal footprint at a very attractive total cost of ownership and return on investment.

The System p5 520 and p5 550 provide enhanced processor performance and I/O capability. While the System p5 520 is ideally suited to high density web server farms and to smaller web server solutions, the System p5 550 provides unparalleled performance as a large capacity web server for enterprise environments that also require high availability, reliability and maintainability.

### **Application Serving**

The application server is a key factor to turning the data residing in the data warehouse into information to be viewed by the end-user. The IBM System p family of products is ideally suited to the demands placed on it by the requirements of the Business Intelligence application server. For a system to be successful as an application server it must exhibit industry leading performance, provide robust reliability, availability and serviceability criteria, and be cost effective. The IBM mid-range and enterprise-level System p UNIX servers provide all of this and are an ideal platform for even the most demanding and strenuous of enterprise-wide Business Intelligence scenarios.

- The System p5 550, eServer p5 570 and eServer p5 575 are the mid-range server offerings.
- The eServer p5 590 and p5 595 are ideally suited to large enterprise-wide application server needs. The high-end eServer p5 595 can provide up to a maximum of 64 physical POWER5 processors in a single system and a maximum of 2 TB of RAM. The system also provides the capability of being partitioned into a maximum of 254 physically discrete partitions.

### **Database Serving**

The database server and its inherent performance are important factors in the deployment of an effective Business Intelligence

solution. Not only must this server execute and manage database transactions, it also must monitor and control the communications systems to the application server to insure optimal performance, reliability and scalability. The key performance factor of a database server is its I/O bandwidth and memory. The IBM System p family of products provides this functionality and performance at all price points from the entry-level System p5 520, through the mid-range eServer p5 570 up to the enterprise-wide high-end eServer p5 595.

These systems have demonstrated outstanding performance in key database benchmarks including TPC-C and TPC-H. (Source: TPC.org)

### IBM System p - Performance and Price/Performance

Recently IBM and SAS achieved outstanding performance running the SAS BI Benchmarking Scenario test suite to help demonstrate the power of SAS® Enterprise BI Server. This test suite was developed by SAS as a means to show customers the scalability of their solutions on various hardware platforms.

The test suite simulates a team performing retail related business activities. One such activity is viewing reports. Reports are viewed via web browser and include static relational reports, dynamic relational reports, and dynamic OLAP reports. Also, SAS® Enterprise Guide® users are simulated by having SAS jobs typical of users run in background processes. In addition, SAS® Add-in for Microsoft Office users are simulated by background processes.

During testing the test suite emulates a morning of business activity. The number of users increases over a forty minute or longer period reaching a constant level of use for several hours, and then gradually logging off the BI platform. Users log on and consistently follow a cycle of report navigation, launching and viewing reports for 45 to 90 seconds. When these tasks are completed, the user logs off the system. Delays occur at appropriate intervals. Once a user cycle is completed a new user cycle is started. The system maintained 1496 concurrent users.

### User Base

- 21 SAS Enterprise Guide software client users.
- 175 users accessing SAS Stored Processes simulating SAS Add-In for Microsoft Office use.
- 175 users building dynamic reports based on relational data via SAS Web Report Studio.
- 175 users viewing OLAP cubes via SAS Web Report Studio.
- 950 users viewing static reports through SAS Information Delivery Portal, HTTP and SAS Web Report Studio.

### SAS BI Benchmarking Scenario Hardware Configuration

- 28 POWER5 Processors (8 partitions)
- 112 gigabytes (GB) RAM
- DS8100 TotalStorage with 12 controllers
- Data, OS and SAS on RAID5
- SASWORK on RAID10
- Dedicated 1Gb Ethernet network



Figure 1

### Application Environment

- SAS 9.1.3 with Service Pack 2
- SAS® Web Report Studio 2.1
- IBM Websphere as the middle-tier Web application server
- Xythos® as the WebDAV server with four CPUs
- SAS Workspace Server with eight CPUs
- SAS OLAP Server with four CPUs
- SAS middle tier server with eight CPUs
- SAS® Metadata Server with four CPUs

### Data Definition

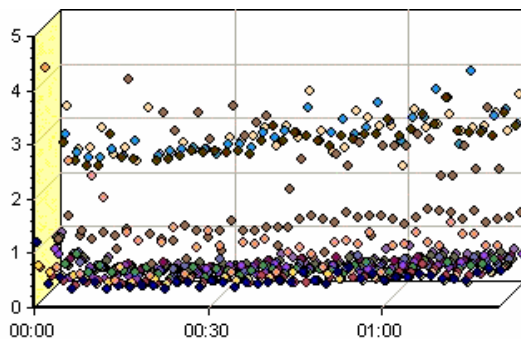
Data storage totals are approximately 160 GBs of uncompressed data including:

- Transaction or fact tables totaling 4.9 GBs, with an average of 36 million rows.
- Dimension tables of 252 and 42,273 rows that are directly joined to the fact tables.
- Fact and support data.
- Four cubes sized 1.5, 5.3, 5.3 and 8.8 GB.

Cardinality was as high as the 42,273 row table but more typically was 252 rows or less.

### Performance Results

A total of 1496 concurrent users were sustained on the tuning environment. The average user experienced less than 5 seconds in this scenario. There was less than 1 second to log into SAS applications such as Web Report Studio and averaged less than 1 second to navigate to reports. Figure 2 shows the average response time for the multiple reporting groups.



**Figure 2: Response time for individual reporting/user groups**

	Number of CPUs	CPU Utilization
<b>Middle Tier</b>	8 CPUs <ul style="list-style-type: none"> <li>• Four partitions of 2 CPUs each</li> </ul>	75%
<b>SAS Metadata Server</b>	4 CPU Partition	75%
<b>SAS Server</b>	12 CPUs <ul style="list-style-type: none"> <li>• Eight CPU partition for SAS Workspace Server</li> <li>• Four CPU partition for SAS OLAP Server</li> </ul>	75%
<b>WebDav Server</b>	4 CPU Partition	75%

**Table 1: SAS BI Benchmarking Scenario CPU utilization**

These outstanding test results confirm for clients the superior performance advantages of IBM systems when used to run critical business intelligence applications in real business situations. The SAS BI Benchmarking Scenario test suite offers clients a look into the performance potential of IBM systems running SAS® Enterprise BI Server applications under the most demanding of circumstances.

### For more information

To learn more about IBM Systems and offerings for SAS, please contact your IBM sales representative, or visit: [ibm.com/servers](http://ibm.com/servers)

Product data is subject to change without notice. This information could include technical inaccuracies or typographical errors. The performance data contained herein was obtained in a controlled, isolated environment. Actual results that may be obtained in other operating environments may vary significantly. While IBM has reviewed each item for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead. It is the user's responsibility to evaluate and verify the operation of any non-IBM product, program or service.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice and represent goals and objectives only

The information provided in this document is distributed "AS IS" without any warranty, either express or implied. IBM EXPRESSLY DISCLAIMS any warranties of merchantability, fitness for a particular purpose OR non-INFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. IBM is not responsible for the performance or interoperability of any non-IBM products discussed herein

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 in connection with this publication 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.

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both: AIX, AIX 5L, DS8000, eServer, iSeries, i5, Micro-Partitioning, OpenPower, POWER4 POWER5, pSeries, p5, System p, System p5, Tivoli, and TotalStorage.