



Deploying SAS[®] Scoring Accelerator for IBM DB2 on IBM AIX

Alfredo Mendoza, STG
Matthias Nicola, SWG
Xinghua (Michael) Hu, SWG
Vishwanathan Krishnamurthy, STG

ISV Business Strategy and Enablement
August 2010



Table of contents

Abstract	1
Introduction	1
SAS Scoring Accelerator overview	2
Using the DB2 Database Partitioning Feature for scalability	3
Deployment example	5
Storage configuration.....	7
Scalability results	8
Scalability results in charts	8
Best practices	11
Summary	13
Resources	14
Acknowledgements.....	15
About the authors	15
Trademarks and special notices	16



Abstract

SAS Scoring Accelerator can leverage the processing capabilities of IBM DB2 to reduce complexity, latency and storage-space consumption in the scoring process. By using the Data Partitioning Feature (DPF) of DB2 9.5 or 9.7, the database can be configured to provide a massively parallel environment (MPP) that is capable of providing scalable performance. This white paper shows a reference architecture for the deployment of SAS Scoring Accelerator for DB2 9.5 or 9.7 running on IBM AIX 6.1 and IBM POWER7 processor-based IBM p780.

Introduction

Businesses today want accurate and fast access to relevant data. Each day, the data warehouse grows, accumulating data at an unprecedented rate. The need to process and analyze the data is critical to making confident business decisions – decisions that can impact the company's bottom line.

SAS® Scoring Accelerator for IBM® DB2® combines the modeling methods available through SAS Enterprise Miner with the scalability and processing speed of DB2. This combination provides an integrated analytical environment that can support the needs of the most demanding organizations by automating model-scoring processes with the database.

SAS Scoring Accelerator overview

Figure 1 shows the SAS Scoring Accelerator process using a database backend. The SAS Scoring Accelerator improves scoring performance by transferring work to an enterprise data warehouse. Models from the SAS Enterprise Miner (1) are translated into user defined functions (UDFs) and deployed inside DB2 (3 and 4). Once the UDFs are in place, SAS clients (5) are able to access these functions and do the work within DB2. This allows customers to take advantage of DB2 features that can optimize scoring function performance. For more information on the SAS Scoring Accelerator please see the fact sheet at <http://www.sas.com/resources/factsheet/sas-scoring-accelerator-factsheet.pdf>.

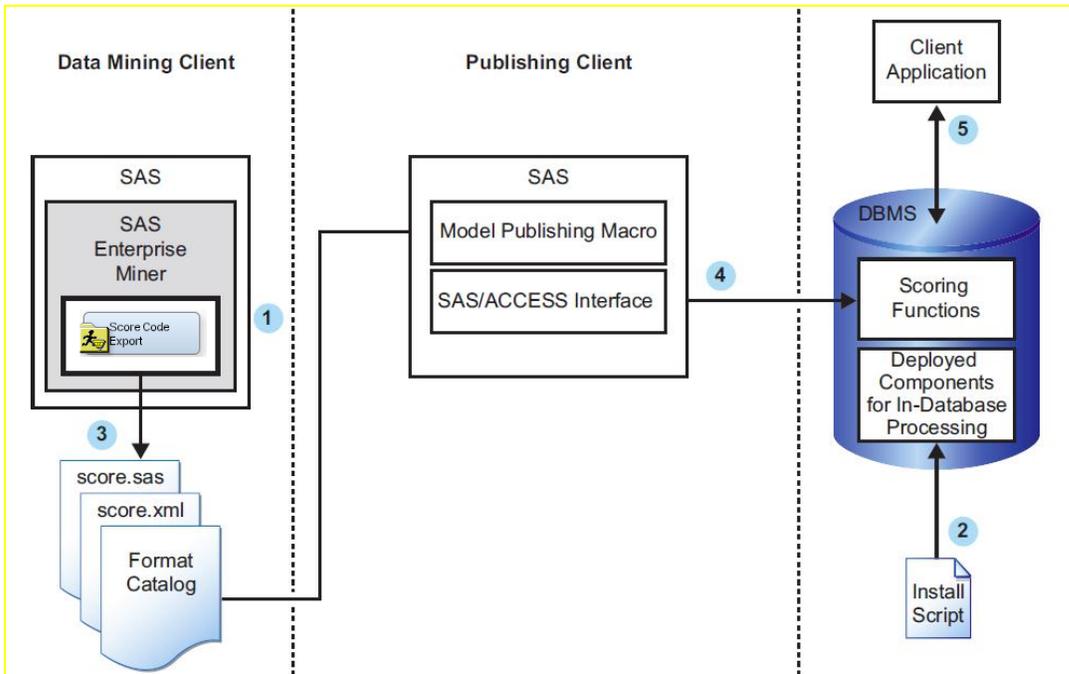


Figure 1. SAS Scoring Accelerator process from SAS Enterprise Miner to a database

Using the DB2 Database Partitioning Feature for scalability

When using a database backend for online-transaction processing (OLTP) or enterprise-resource planning (ERP), the typical database configuration uses a single database partition. However, as data volumes increase, as in the case of data warehouses, a multipartition database might be required. The DB2 Database Partitioning Feature (DPF) allows a database to be distributed across multiple database partitions to enable parallel processing of large database volumes. The rows of large database tables are distributed through a hash function across the database partitions. The partitions can then work in parallel, each processing a subset of the table (divide and conquer).

The DB2 system architecture has the flexibility to be scaled in a *vertical* or *horizontal* fashion. Figure 2 shows DB2 configured using multiple partitions on a single multiprocessor server. In this sample configuration, the DB2 database is distributed across several DB2 partitions installed on the same server, but using separate logical disks. The logical disks can be on a fast storage device such as the IBM DS8100 storage system. As the need to add another DB2 partition arises, it is possible to add more disks, memory and processor resources to the single server. This method of scaling is known as *vertical scaling*.

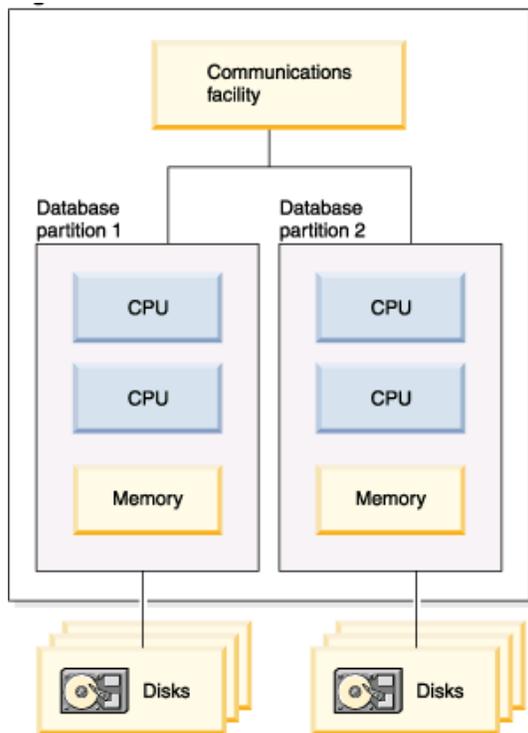


Figure 2. DB2 multipartition configuration on a single multiprocessor server

Figure 3 shows DB2 configured using multiple partitions installed on a cluster of multiprocessor servers. In this configuration, the DB2 database is again spread across the DB2 partition, with each partition placed on its own logical disk on separate and unique servers. As the need to add

more DB2 partitions arises, more servers can be added to the cluster. This method of scaling is known as *horizontal scaling*. Figure 3 shows that the DB2 Database Partitioning Feature can be deployed in a shared-nothing architecture to scale up to hundreds of database partitions. A shared-nothing (massively parallel processing (MPP)) database architecture is possible because each database partition is an autonomous database system. It has responsibility for its own (and only its own) data, logs, data locking and other essential elements that make up a database. At the same time, the database partitions employ a coordinator mechanism to distribute and coordinate the work they are performing.

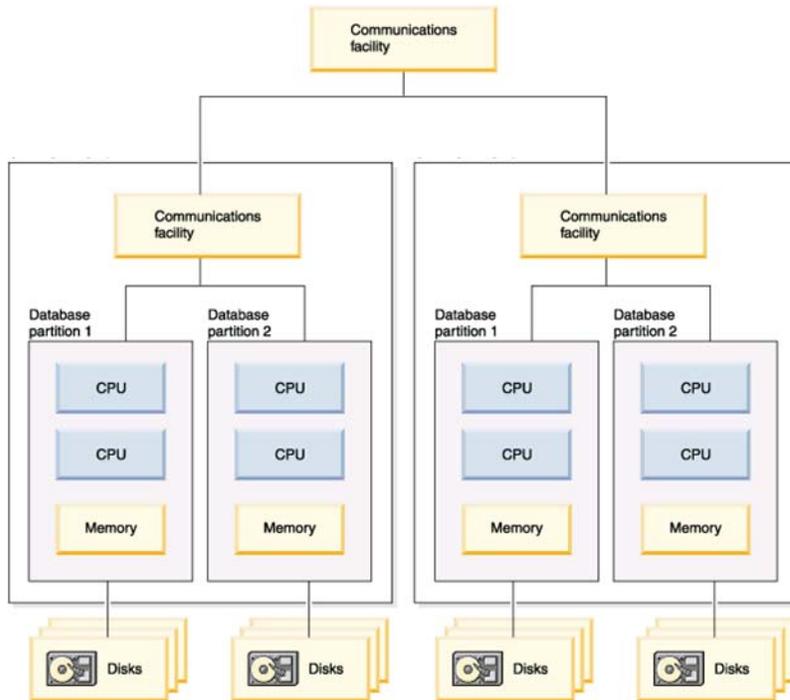


Figure 3. DB2 multipartition configuration on a multiple-server cluster

For the purpose of this scalability exercise, this test environment was configured with DB2 using DPF on a single multiprocessor system. The test environment details are described in the next section.



Deployment example

The focus of these tests is to verify the scalability of SAS Scoring Accelerator for DB2 using varying numbers of database partitions. An IBM AIX® server is configured to use a proportionate number of DB2 database partitions (relative to the number of system processors, size of memory, and size of the file system) as the data volumes are scaled.

Table 1 and Figure 4 show how the database and system configuration are managed during scaling by using different DB2 partitioning configurations.

Number of DB partitions	Number of processors	Memory (GB)	Number of RAID5 devices	Initial data volume** (rows)
1	1	7.5	1	4 million
4	4	30	4	16 million
8	8	60	8	32 million
16	16	120	16	64 million

Table 1. Configurations for DB2 scalability test

Processor and memory resources are controlled through the configuration of dynamic logical partitions (DLPARs) as the test scales the number of database partitions. The data volume is scaled proportionately to the number of DB2 partitions. For example, from 4 partitions to 16 partitions the number of partitions quadrupled as well as the hardware resources. Although the amount of data per DB2 partition remains constant, the amount of data to be processed increases as database partitions are added to the test configuration. In running the tests, the expectation was to see that, as the number of DB2 partitions increase, the number of rows processed per second would increase proportionately, exhibiting linear scalability. The row size is 256 bytes, so one million rows are 244 MB, and 60 million rows are 14 GB in size.

For the purpose of this test, SAS provided sample data models to run the scalability tests. These models are provided through the SAS Enterprise Miner. Two types of data models are used, namely a *neural-network* model and a *regression* model. Test cases for each model run from the SAS client.

The test environment also includes a DS8100 storage system configured with RAID5 (7+1). The IBM AIX® operating system sees each logical array of the storage system as a logical disk. Sixteen logical disks with a total size of 2 terabytes (TB) are configured for DB2 to use.

The SAS client is installed with SAS version 9.20M3. The AIX version is 6.1.

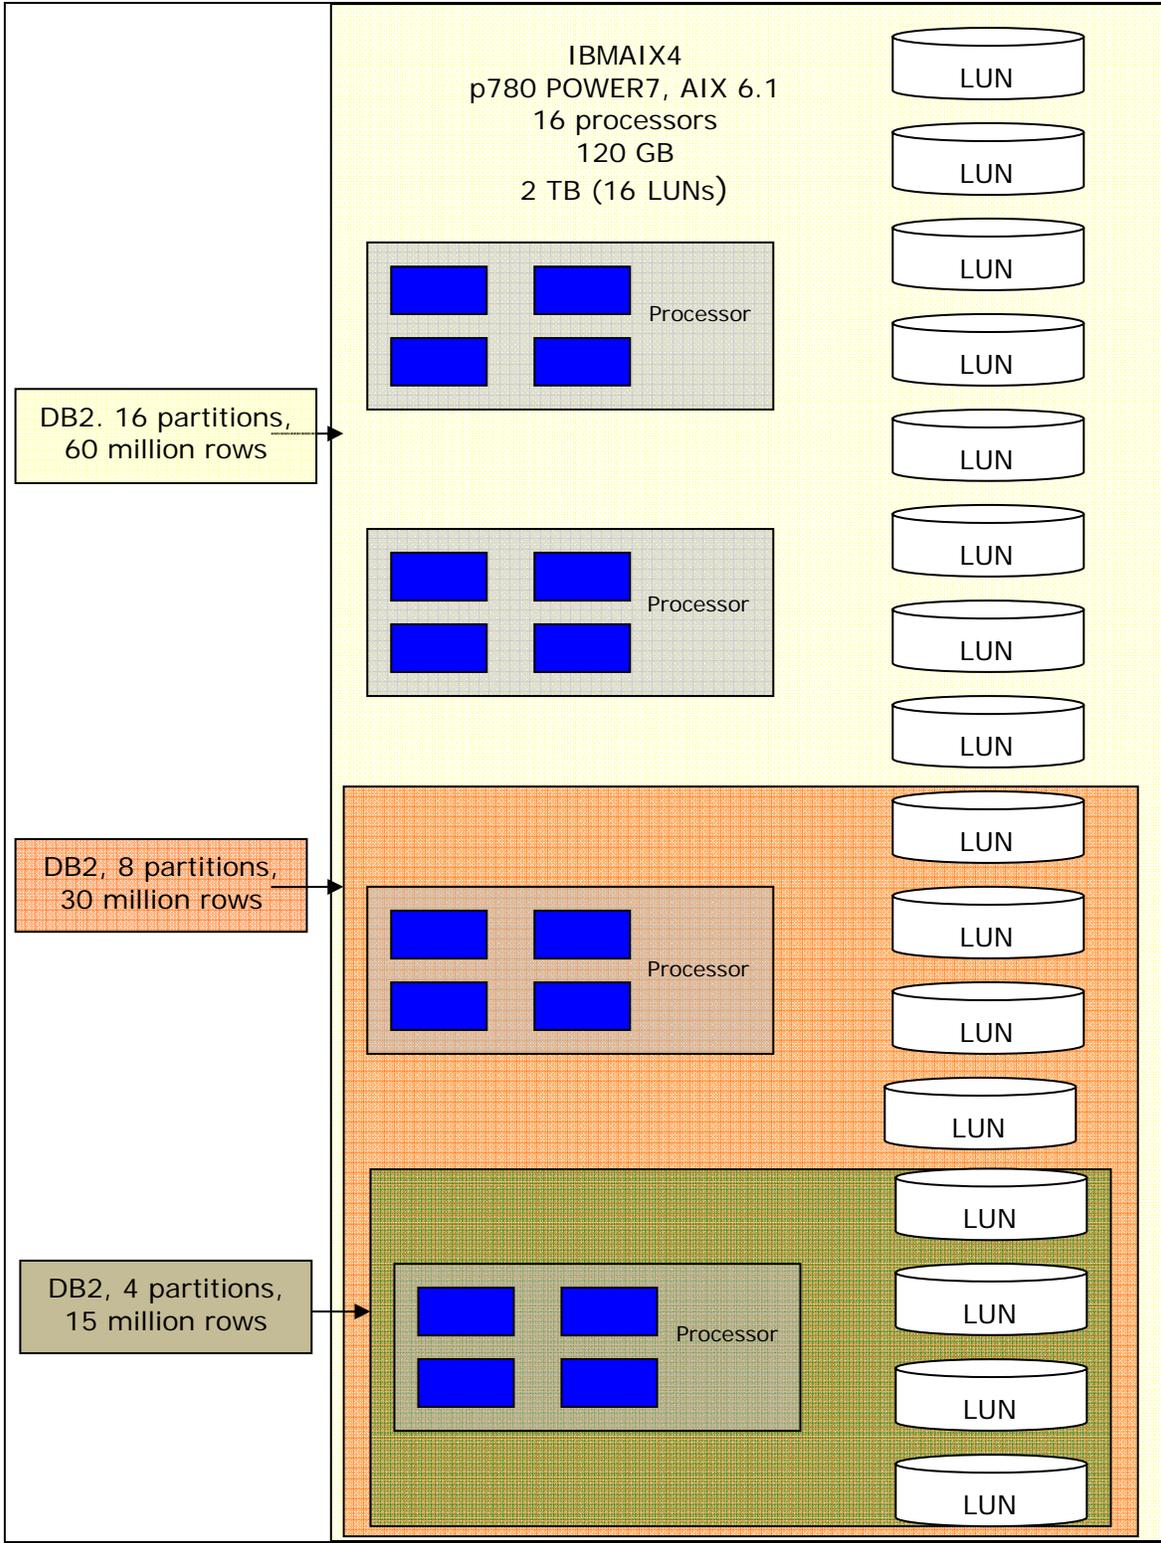


Figure 4. Four, eight and sixteen DB2 partition configuration used for scalability testing

Storage configuration

The IBM DS8100 storage system consists of 128 physical disks. To strike a balance between disk I/O performance and economies of scale, the storage system is configured as RAID5. The minimal unit, also known as a *logical unit (LUN)*, on the system that a piece of storage can be created is called an *array*. Each array consists of eight physical disks. Figure 5 shows a representation of a LUN consisting of an array with eight disks.

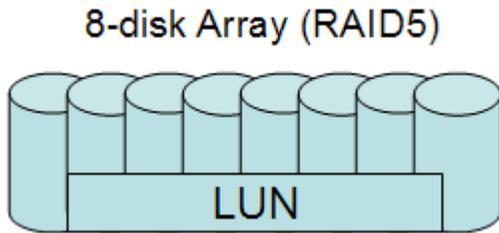


Figure 5. Representation of a LUN that consists of an array with eight disks

With 128 physical disks, the disks are grouped into arrays of eight disks, each resulting in 16 arrays. To simplify the AIX disk management, a LUN is created from each array and then is presented to AIX (see Figure 6).

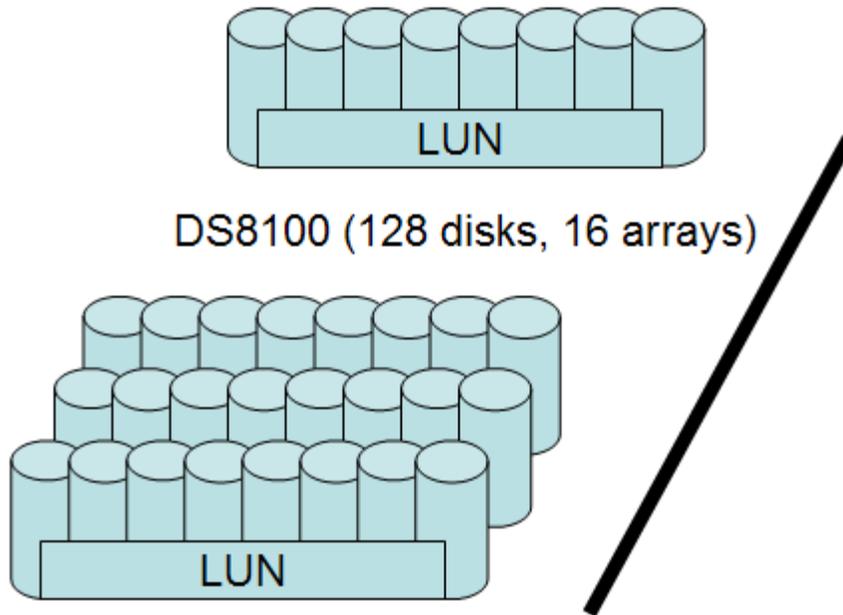


Figure 6. A 128-disk subsystem configured with 16 arrays

From the point of view of the receiving AIX partition, the LUNs are seen as 16 logical disk drives (hdisks) from which a volume group is created. In these tests, each DB2 partition uses one LUN.

Scalability results

This section summarizes the scalability results for SAS Scoring Accelerator for DB2, when DPF is used with up to 16 partitions.

All results described in this white paper use the following:

- Up to 16 DB2 partitions running on up to 16 cores and 16 LUNs
- One of two models:
 - Neural-network model (NN)
 - Regression model (REG)
- 4-, 16-, 32- and 64-million rows
- Unfenced UDFs - unless otherwise noted.

Scalability results in charts

This section describes the scalability of the in-database scoring functions when the number of DB2 partitions varies. Figure 7 and Figure 8 show the number of rows processed per second for 1, 4, 8 and 16 DB2 partitions.

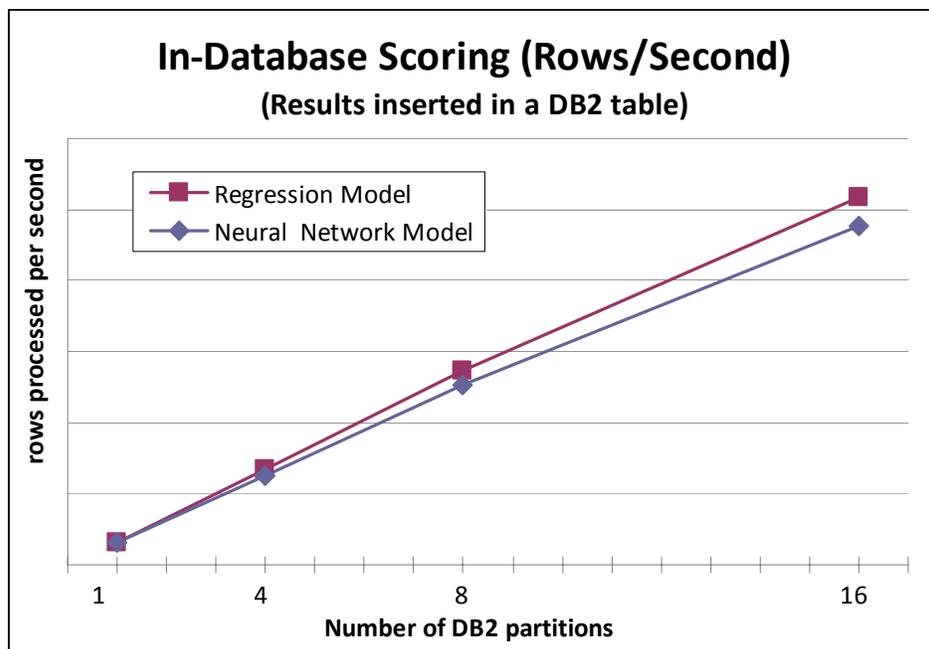


Figure 7. Scoring results written to DB2

The tests in Figure 7 write the scoring results into a DB2 table without sending the results to the SAS client. The graphs show that the number of rows processed scales linearly in the number of DB2 partitions, as desired. Additional tests (not shown here) used larger data volumes and confirmed that the scoring performance scales linearly, even with 640 million rows.

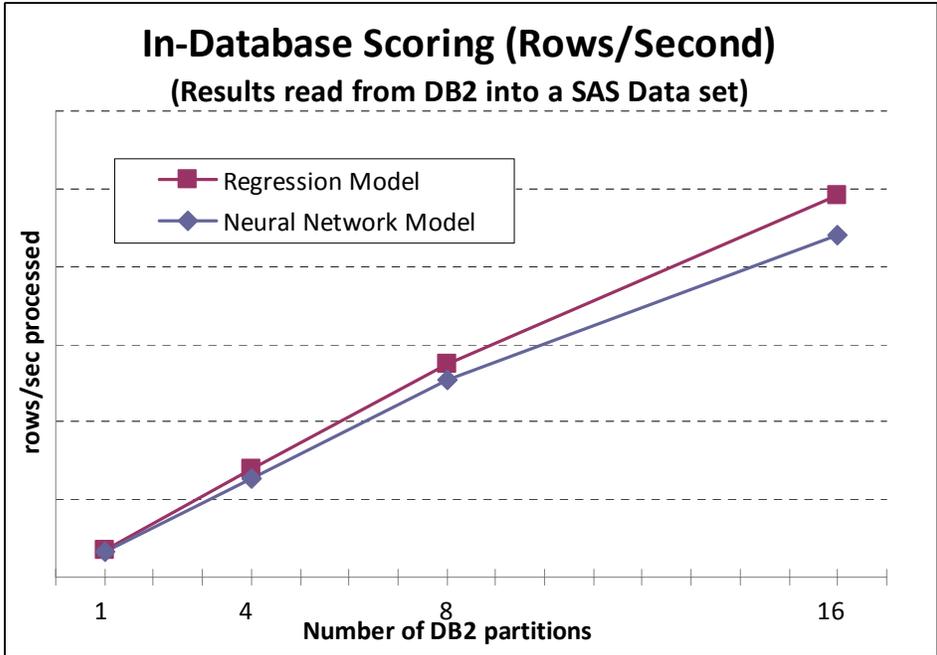


Figure 8. Scoring results written to a SAS dataset

The tests in Figure 8 do *not* write the results into a DB2 table but send them into a SAS data set in the SAS client. In both cases, the scalability is close to linear. The scalability of writing results into a DB2 table (which is fully parallelized) is slightly better than retrieving results into the SAS client (and that retrieval is not parallelized). Figure 9 compares the scalability of writing results to a SAS data set instead of a DB2 table for the neural-network scoring model.

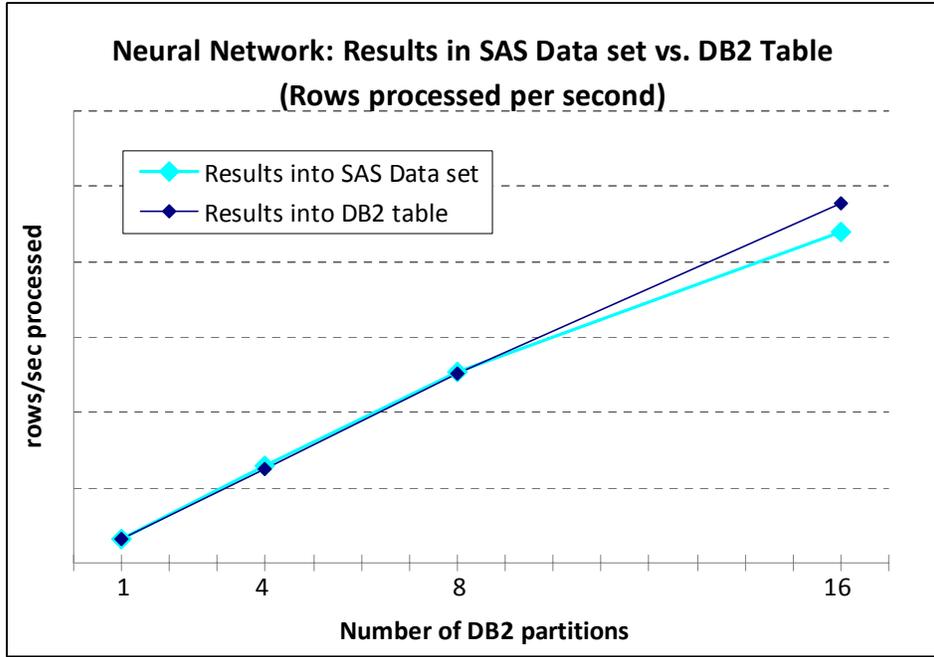


Figure 9. Writing results to a DB2 table vs. SAS dataset (NN)

Figure 10 compares the scalability of writing results to a SAS data set versus a DB2 table for the regression-scoring model.

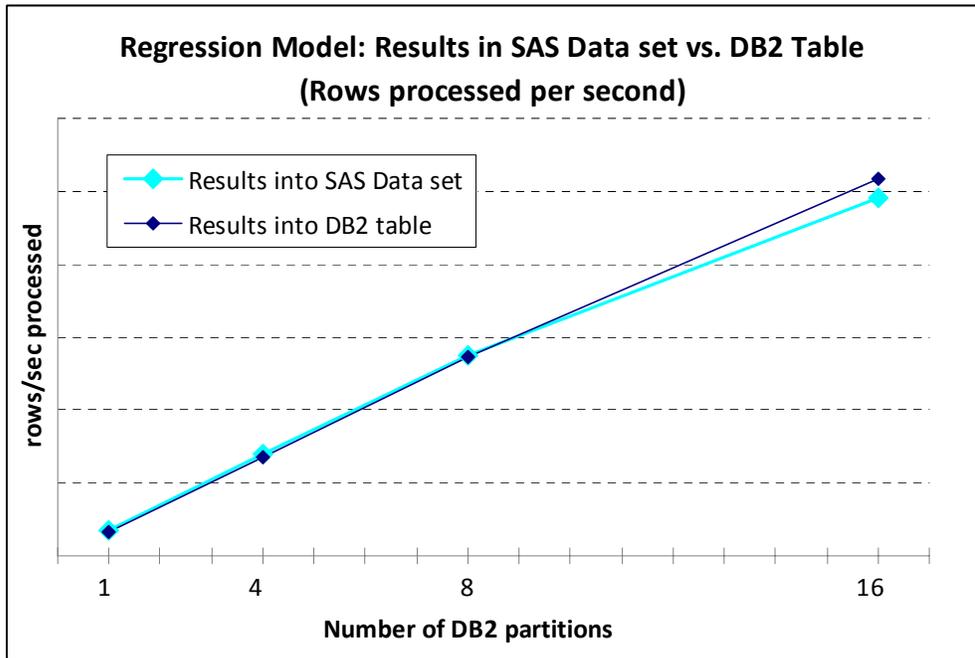


Figure 10. Writing results to a DB2 table vs. SAS dataset (Regression scoring)

In these tests, the number of rows processed is proportional to the number of DB2 partitions, that is, four million rows per partition. For 16 partitions with 64 million rows total, the retrieval of 64 million results from the DB2 server into the SAS client causes a slight dent in the scalability.

Best practices

DB2 has the ability to run user-defined functions (UDFs) as fenced (outside of the database-manager process) or unfenced (within the database-manager process). In the previous sections, all results use *unfenced* UDFs. Several individual tests with *fenced* UDFs consistently show about 40-percent overhead, regardless of the number of DB2 partitions or rows processed (see Figure 11).

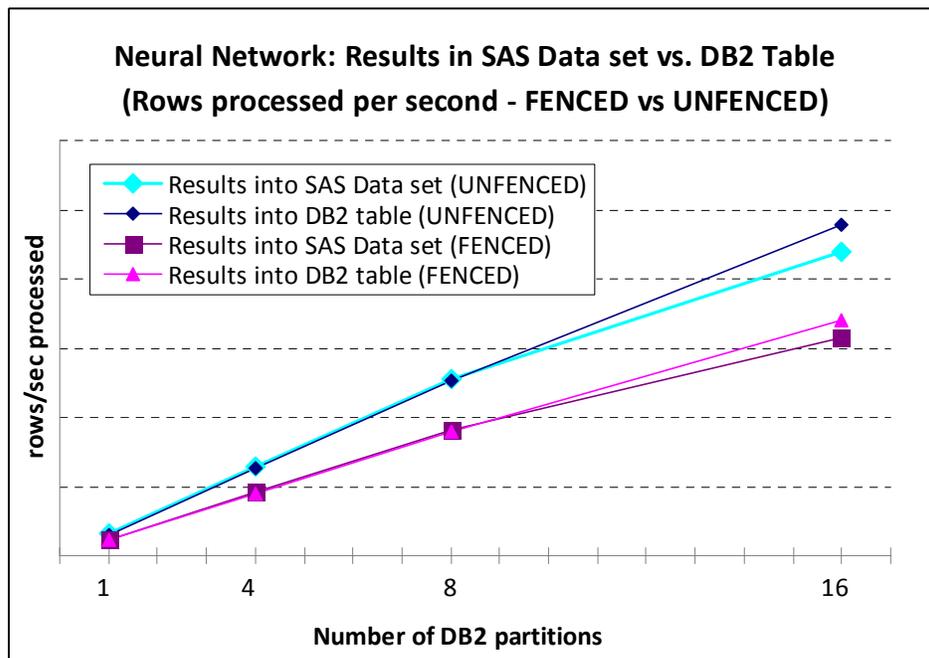


Figure 11. Performance of unfenced UDFs compared to fenced UDFs

Figure 11 shows how the 40-percent overhead affects the performance of the neural-network model. Although scalability is still close to linear, the absolute performance is lower than with unfenced UDFs.

- Use fenced-mode UDFs for the initial testing of a new in-database scoring model. After confirming that the UDFs run with stability, redeploy them in unfenced mode for best performance.
- The use of in-database scoring UDFs tends to make the system rather processor-bound than I/O-bound. However, if inserting scoring results into a target table suffers from high logging costs or *log full* errors, consider disabling logging for the DB2 target table in the same unit of work:

```
ALTER TABLE <tablename> ACTIVATE NOT LOGGED INITIALLY;
```

- If data from a source table is read for scoring, and if the scoring results are inserted into another database table, then the target table needs to be partitioned in the same manner as the source table (collocated). This avoids the shipping of scoring results across DB2 partitions.
- Prior to DB2 9.5 Fixpack 7 or DB2 9.7 Fixpack 3, avoiding ORDER BY clauses in the queries that use scoring UDFs ensures best performance.
- If the results of an in-database scoring operation need to be stored in another database table, explicitly pass an SQL statement of the form *INSERT... INTO... SELECT... FROM...* through *EXECUTE()* function of the SAS SQL Proc to ensure the most direct and efficient flow of data. Do not use the SAS client to specify *db2 libref* as the target for the scoring results, as this requires results to be fetched from DB2 into the SAS client only to be sent back to DB2 for insertion, which is suboptimal (see Figure 12).

Figure 12 shows a comparison of elapsed times using different methods for handling results sets. If a user wants to keep a results table inside DB2, the best practice is to use SAS SQL to perform an INSERT inside DB2 to populate the target table first, then pull the target table into a SAS temporary dataset for further processing.

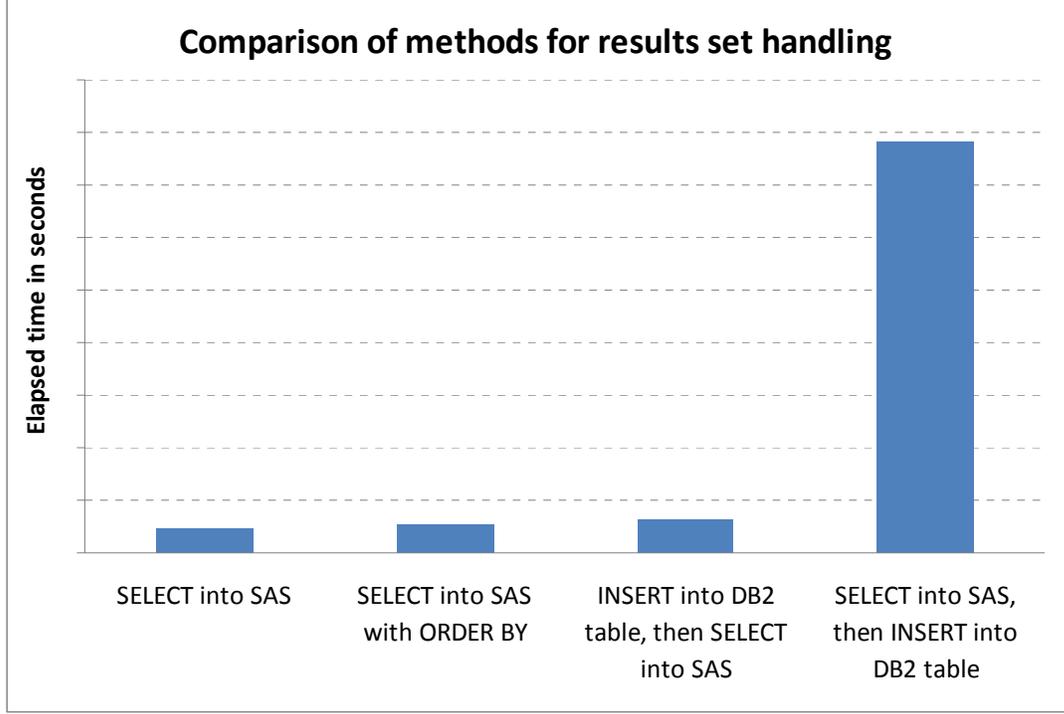


Figure 12. Shows elapsed times of different methods for handling results sets from the SAS client

Summary

Many customers need to analyze data within their data warehouses in order to make critical business decisions. Mainstream practices for analytic solutions suggested that organizations build databases that were analytic and subject-focused, which may not have taken advantage of the available in-database performance and flexibility, often create challenges, such as:

- Time constraints in moving large volumes of relevant data
- Management and provisioning of data
- Proliferation and governance of data

The partnership between IBM and SAS addresses these challenges by taking advantage of IBM DB2 in-database capabilities by moving certain analytic scoring tasks directly into the DB2 database resulting in a greatly improved integration between SAS9 and IBM DB2 data warehouse infrastructure capability included with InfoSphere Warehouse and the IBM Smart Analytics System.

This white paper documents a reference architecture that provides performance and the flexibility to scale when a customer's data warehouse grows. This solution also provides greater control of sensitive information within the data warehouse. Running SAS Scoring Accelerator with DB2 on IBM POWER7® 780 with AIX 6.1 and IBM System Storage® DS8100 provides the infrastructure for customers to gain a competitive advantage.

Resources

These Web sites provide useful references to supplement the information contained in this document:

- IBM System p Information Center
<http://publib.boulder.ibm.com/infocenter/pseries/index.jsp>
- System p on IBM PartnerWorld®
ibm.com/partnerworld/systems/p
- AIX on IBM PartnerWorld®
ibm.com/partnerworld/aix
- IBM Systems on IBM PartnerWorld
ibm.com/partnerworld/systems/
- IBM InfoSphere Data Warehouse (based on DB2 with DPF)
<http://www.ibm.com/software/data/infosphere/warehouse>
- IBM Smart Analytics System
<http://www.ibm.com/software/data/infosphere/smart-analytics-system>
- DB2 Best Practices
<http://www.ibm.com/developerworks/data/bestpractices>
- DB2 Information Center
<http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/index.jsp>
- IBM Publications Center
www.elink.ibm.com/public/applications/publications/cgibin/pbi.cgi?CTY=US
- IBM Redbooks®
ibm.com/redbooks
 - PowerVM Live Partition Mobility on IBM System p, Harding, M., et al. (2007)
www.redbooks.ibm.com/abstracts/sg247460.html
- IBM developerWorks®
ibm.com/developerworks
- Hayes-Hall, E, et al., The Deployment of a SAS Enterprise Business Intelligence Solution in a large IBM POWER6 Environment, IBM, 2006
ibm.com/support/techdocs/atmastr.nsf/WebIndex/TD103051
- Jain, V., Virtual Networking on IBM AIX 5L, IBM, 2004
ibm.com/systems/resources/systems_p_os_aix_whitepapers_pdf_aix_vn.pdf

Acknowledgements

Thanks go to the following people who contributed to the scalability testing and this paper:

- SAS: Kim Sherrill, Robert Ray, Douglas Liming, Phil Mohr, Pat Buckley and Rachel Wynne
- IBM: Armando Paniagua, Gopi Duddi, Frank Battaglia, Frank Bartucca and Leamon Hall

About the authors

Freddie Mendoza is a consultant in IBM Systems and Technology Group ISV Enablement. He has more than 20 years experience in the IT Industry. He currently works with several software vendors to enable their enterprise applications on the latest IBM operating platforms. He can be reached at mendoza1@us.ibm.com.

Matthias Nicola is a senior software engineer at the IBM Silicon Valley Lab, in San Jose, California, USA. He focuses on DB2 performance, scalability and benchmarks. Matthias works closely with the IBM DB2 development team, as well as with IBM clients and business partners to help them design, optimize and implement solutions that are based on DB2. Previously, Matthias worked on data-warehouse performance at Informix Software. He can be reached at mnicola@us.ibm.com and www.matthiasnicola.de.

Xinghua (Michael) Hu is a software engineer at the IBM Silicon Valley Lab, working on the development and integration testing of IBM data-warehouse and IBM Smart Analytics System offerings, based on DB2 and the database-partitioning feature. His email is hu@us.ibm.com.

Vishwanathan (Vish) Krishnamurthy is a technical consultant in IBM Systems and Technology Group. He has more than 10 years experience working in the IT industry and specializes in testing on IBM Power Systems with PowerVM virtualization. His e-mail is vishwanathan.k@in.ibm.com.



Trademarks and special notices

© Copyright IBM Corporation 2010. All rights Reserved.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. 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 "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

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

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

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

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

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.