

Up to 50 X SAS® Performance Gains on Large Data Volumes Using Scalable Parallel Computing

Author: Jesse Davis, Consultant

Co-Author: Helen Spade, MasPar Computer Corporation, Sunnyvale CA.

ABSTRACT

Turning large volumes of historical data into valuable information to support the decision making process of management is a growing need of corporations in the 90's. However, the ability to perform fast data reduction on large data volumes is limited by traditional computer systems which can take hours of CPU time and days of elapsed time to process multi-gigabyte datasets.

This paper describes how one of my customers in the health care industry addressed the volume issue in health decision support processing and improved the end user's toolsets. The paper describes the problems encountered while trying to achieve their business goals, and the hardware and software solution implemented to resolve these problems. Actual performance numbers are given for Tabulate, Select, 3-way Join, Nested Select Queries and a multiple user test. With up to 50 X faster query turn-around times on large data volumes, productivity was increased by:

- 1) Allowing users to extract valuable information with queries that could not be run before.
- 2) Reducing the cost of data analysis.
- 3) Providing users more elaborate reporting and greater versatility due to SAS and the new client/server environment.
- 4) Improving overall business with more timely, accurate and cost-effective results.

INTRODUCTION

This customer provides management and leadership of the health care system on behalf of the government. Through a broad range of programs, services and public funding, they are responsible for maintaining high quality, accessible, affordable health care for the residents. Their primary business goal is to promote and provide for the physical, mental, and social well-being of all residents.

PROBLEM

There are three issues in Decision Support – Volume, Volume, and Volume. For this customer, data volume was driving up processing costs, interfering with operational systems, and slowing responses to client's information


requests. They are working with more than 30 Gigabytes of historical medical claims filed against the Province's insurance plan with more than two million transactions being added to this historical database bi-weekly. This volume of data stems from the number of health claims being produced and the amount of information captured on each claim. Summarizing data according to areas of known interest only addresses a small percentage of the questions and could not deal with the requirement to study individual client utilization patterns and other questions requiring detailed data. The best use of traditional computing methods resulted in slow and restricted data analysis in a clumsy multi-platform environment. Please refer to Diagram 1 below for an overview of the old computing network using the traditional computing methods.

Diagram 1

Hardware and Software Configuration using Traditional Computing Methods

Mainframe OLTP and Decision Support Server

- IBM 9021-962, 1.5 GB CSTORE, 2.0 GB ESTOR,
- 4 x MVS LPARS and 1.5 TB DASD
- DB2 Database, tapes
- Reports Generated via QMF/DB2, PL/I Syncsort



SOLUTION

In May 1993 the customer began investigating potential dedicated DSS platforms. A dedicated system would relieve stress on the OLTP environment and improve service to the DSS area. The cost of the alternative platform, however, was not to exceed the current operating budget. In simple terms, the cost of buying the new hardware and software was required to be equal to, or less than, the anticipated mainframe decision support operational costs. Containment of escalating costs, however, was only half the picture. The other half of the solution was the introduction of better tools for data analysis.


Since the required price/performance levels could not be met by traditional computer systems, they looked toward parallel

processing solutions such as Oracle on a cluster of workstations and the SAS System integrated with a scalable parallel computer. After reviewing numerous proposals, it became obvious that the scalable parallel processing architecture in conjunction with a SAS server environment far exceeded other technical options. The reasons for this are examined in the next section.

Utilizing their existing hardware and software and SAS integrated with a scalable parallel server, they designed the Alternate Platform Environment shown in Diagram 2, intending to improve service to both OLTP and Decision Support users:

Diagram 2

Alternate Platform Environment



Hardware/Software Details

Mainframe OLTP Server

- IBM® 9021-962, 1.5 GB CESTOR, 2.0 GB ESTOR, 4 x MVS LPARS and 1.5 TB DASD
- DB2 with SAS/ACCESS®
- SAS products on M/F include: SAS/BASE®, SAS/STAT®, SAS/GRAPH®, SAS/ASSIST®, SAS/AF and SAS/FSP
- FTP to transfer data
- Data load initiation via scheduler and TCP/IP REXEC call to MasPar server

MasPar Decision Support Server

- 256 MB Distributed Memory, 40 GB RAID 3 Parallel Disk Array, Massively Parallel - 4096 Processors and 6000 MIPS
- SAS/CONNECT® interface to PC's and M/F
- Ultrix/SAS with SAS/BASE, SAS/STAT, SAS/GRAPH, SAS/ASSIST, SAS/AF and SAS/FSP
- TABLE-MAKER™ software and TM/ACCESS® interface

SAS Server

- 5 GB SCSI DISK
- SAS, SAS/STAT, SAS/GRAPH, SAS/ASSIST, SAS/AF and SAS/FSP
- SAS/CONNECT and FTP to transfer data to PC's
- TM/ACCESS interface

IBM PCs

- Microsoft® Windows, XWindows and Telnet
- SAS/AF query generator
- SAS and SAS/CONNECT
- Excel® using the DDE interface with SAS
- X Windows drivers for SAS sessions
- TM/ACCESS interface

WHY MasPar?

There are larger databases, however, it is hard to imagine data with a wider scope than that managed in the Health Care business. It is impossible to cover all the bases with partitions, indexes, summaries, and samples. All fall short when the client asks....

"How many 13-19 year old native females from Poskapa are receiving services outside of their health care region? And by the way, could you classify this by the servicing region and the specialty of the practitioner?"

No matter what type of system, you end up scanning every record to see if it is of interest.

Traditional computing architectures, including mainframes, examine one record from one tape or drive at a time. Workstation clusters or other parallel systems may use 4 to 20 processors with independent drives and intelligent networks for an extraordinary increase in speed. The client's MasPar system uses 4,096 processors with a specialized array of disks attached via an I/O channel that can achieve over 1 GB/Sec data transfer rates. The MasPar system reviews 4,096 records at a time. The performance of the MasPar approach was about 20x faster than the other parallel approaches and up to 50x faster than the mainframe. Why? For queries against a specific year, the workstation clusters were not effective because of the way they distribute their data. Furthermore, the I/O limitations of the workstation cluster solutions restricted their performance results. Clusters are great for OLTP processing but for combing through data the MasPar architecture could not be matched.

WHY TABLE-MAKER?

TABLE-MAKER is the software operating on the MasPar and the client machines that provides data access, storage, and management functionality. TABLE-MAKER storage

techniques also contribute to the solution of the volume problem.

In the environment using traditional computing methods as shown in Diagram 1, the entire 256 byte record had to be retrieved to determine the sex for a specific claim. In the relational database world, special indexes are sometimes constructed to get around this problem. An index is created that contains all the information needed to allow index only access. This results in redundant, "skinny" files managed by the RDBMS - a trade-off of increased storage requirements for reduced parsing of data to increase performance gains.

The TABLE-MAKER software on the MasPar stores all elements of a record in independent files. Each file is much like a "skinny" index file. TABLE-MAKER manages the synchronization of these files to maintain record integrity. This data storage strategy means that if you want the paid amount of claims with specialty 30, you parse only the specialty file (2 bytes wide) and the paid amount file (3 bytes wide), rather than the complete 256 byte record - resulting in a 50 fold decrease in the amount of data that needs to be parsed.

WHY SAS?

The Alternative Platform solution integrates a variety of specialized tools to meet the needs of the Research officers and to contain costs. These tools include software, servers, mainframes, and a MasPar. Early in the investigation of alternate platforms, SAS was identified as the glue that would assist in integrating all these tools, including the MasPar. Development in the old environment was shifting to SAS to facilitate the eventual conversion to the new platform. SAS allowed us to cut conversion costs significantly.

The SAS System data analysis functions and the ability to use SQL and pass data requests to other platforms also weighed heavily in it's selection as the primary development and analysis toolset.

THE PROJECT APPROACH

The MasPar solution is relatively straight forward from a technical viewpoint. The hardware installation and the set-up of the PCs with SAS and communication software required minimal effort. Very early in the project user acceptance was identified as the single most critical project success factor. User acceptance would affect productivity, quality, morale, staffing, costs and just about every aspect of the workplace that makes a successful department. The end user tends to be a distinct type of individual with a mind of his or her own. Trying to force users to move from the old to the new would be a fools game.

To address this issue, the project was handed over to the client's end users and system staff. Rather than bring in a team to set-up the Alternate Platform Environment, they made

the decision to have the users do it. The systems integration role was essentially reduced to that of a single project manager – making the end user responsible for not only their own success, but the success of co-workers as well. This resulted in a number of benefits to the project:

- 1) User acceptance was virtually guaranteed.
- 2) Experts within the client area were built before the system was put in place. This seeded the environment for quality help.
- 3) Implementation issues such as procedures, standards, quality education, testing and many other "soft" issues were cost effectively addressed.
- 4) Development and testing was thorough and not constrained by preset contractual limitations to scope and effort.

RESULTS

Please refer to Table 1 below for details on the benchmark results of DB2 on the mainframe versus the SAS System integrated with a MasPar scalable parallel server. Following Table 1 is more detailed data on single user and multiple user benchmarks.


Table 1
Benchmark Results

Query	MVS/DB2 Elapsed Time	MasPar Elapsed Time	# Rows
Tabulate	3 Minutes	21 Seconds	2 Million
Tabulate	16 Hours	19 Minutes	260 Million
Select	4 Hours	12 Minutes	36 Million
Select	24 Hours*	49 Minutes	148 Million
Tabulate	1 Minute	19 Seconds	2 Million
Tabulate	10 Seconds	3 Seconds	130,000
Tabulate	1 Second	6 Seconds	53,000
3-way Join	10 Seconds	3 Seconds	8,000
Nested Select	2 Minutes	6 Seconds	46,000

*Includes Tapemounts

Single User/Single Query Test

- Query: Distribution of long term and extended care client groups for a calendar year of service.
- Total Data set size = 14 GB
- Total number of records = 123M
- Verb: Tabulate




16 Hours
ES/9000-962 (6 processor)
MVS/DB2



13 Minutes
MP-1204
MasPar/SAS

Six User/30 Query Test

- Total Data Set Size = 43 GB
- Total Number of Records = 381 M
- Both systems dedicated to multiple user test.
- Verbs: Tabulate and Extract



Elapsed Time
73.5 Hours
ES/9000-962 (6 processor)
MVS/DB2



Elapsed Time
1.8 Hours
M-1204
MasPar/SAS

With performance results such as these, it became clear that they could not only resolve their immediate problem, but could significantly improve their business by:

- 1) Extracting valuable information with queries that could not be run before due to technology limitations.
- 2) Offering wide-spread, desk top access to any size data set – empowering the end user.

Furthermore, they were able to reduce mainframe processing costs by roughly \$1.5M over 3 years for an overall savings of approximately \$300K after including new system hardware and software cost, residual MVS costs and system integration contract cost.

BUSINESS BENEFITS

Adding new abilities to clients often results in unexpected benefits. The Alternative Platform now allows the client to access detailed data without incremental cost and with a response time that allows interactive processing of requests (a cornerstone of Decision Support success). This fundamentally changes the way in which work is done and the benefits will be uncovered as the area grows into its new abilities. Some immediate benefits are:

- Ability to meet the increased demand for information without sending system costs through the ceiling.
- Ability to answer urgent detailed questions in minutes rather than days.
- Ability to meet the challenges of changes to business such as monitoring Clinical Guidelines and Protocols.
- Ability to perform detailed audits without needing costly, and, often difficult to get programmers.
- Potential to let end users access information on their own, thereby freeing up MIS resources to work on other initiatives.

SUMMARY

The project team prepared a five month project plan including three one month pilot phases. The project came in on schedule and migration of the end users began in the fifth month.

In summary, they believe that the SAS System integrated with a MasPar scalable parallel server will significantly benefit both their overall business and the end-user (e.g. analyst) by:

- Improving query turn-around times, even for large queries that could not be run before.
- Reducing cost of data analysis.
- Offering more elaborate reporting and greater versatility due to SAS and the new client-server environment.
- Opening up new business possibilities.
- Implementing a solution that is applicable to many other areas of the business.
- Enabling them to get the job done.

SAS, SAS/ACCESS, SAS/AF, SAS/ASSIST, SAS/BASE SAS/CONNECT, SAS/FSP, SAS/GRAPH, AND SAS/STAT Software are registered trademarks or trademarks of SAS Institute, Inc. in the USA and other countries. IBM, DB2 and SQL are registered trademarks of International Business Machines Corporation. ® indicates USA registered

Other brand and product names are registered trademarks or trademarks of their respective companies