



Hewlett Packard Research and Development Update March 2002

Scaling SAS[®] applications on the HP Superdome computing platform

Statement of Opportunity

Both HP and SAS agree the focus in today's markets is for the offering of integrated solutions, such as SAS's Marketing Automation (MA), Customer Relationship Management (CRM) and Supplier Relationship Management (SRM). For effective delivery of such enterprise -class solutions, we must take a serious look at the underlying software and computer technologies.

Scalability, both up and out, continues to be a growing challenge for SAS application environments. Large data volumes can develop out of the integration a broader range of data sources in a B-I environment, or grow out of the huge data sources driven by e-Commerce storefront applications. The movement, management, and analyses of data volumes at this level are serious matters, and require tools and strategies that can meet the associated scaling requirements.

Our proposal is a study of these environments, first, by developing the SAS technology solutions strategies, and the by proofing such strategies on HP's largest computer/mass storage environments.

SAS and HP have collaborated on the development of SAS Version 9 and the implementation on HP-UX. To assess the performance and configuration aspects of SAS Version 9, we launched a joint project to run various test suites of SAS Version 9 development code on a 64-way HP Superdome enterprise server. This project focused primarily on analysis of high-performance enhancements to the base SAS software, multithreaded analytical procedures, foundational server technology, and SAS's Scalable Performance Data Server.

History and Background

For over 25 years, SAS software has been a leader the Decision Support Systems (DSS) marketplace, providing robust, end-to-end solutions for customers worldwide. As the industry has evolved from Decision Support to Data Warehousing to Business Intelligence to Intelligence Architecture, SAS continues to lead the market with cutting-edge, high quality software solutions.

SAS has delivered parallel technologies that interface to the SAS system, beginning with the Scalable Performance Data Server (SPDS) in Version 6 as a way to speed up data services by exploiting SMC servers. Advances in the SAS/CONNECT technology at

Version 8 allowed the creation and control of multiple SAS sessions from a single parent SAS session. This functionality, called MP-CONNECT, allows for course-grained parallelism on a multi-CPU server platform, or even across a network of computers. With the development of SAS Version 9, SAS introduces the SAS Scaleable Architecture (SSA). For the first time, partitioned I/O and parallel processing capabilities will be available to multiple areas of the SAS System.

The Importance of Being Early

We deliberately started this project while SAS Version 9 was still under active development. This was done for several reasons:

- We wanted not only to detect bottlenecks, but also to make the changes needed to insure optimum performance and scalability.
- We wanted to gain an ‘early look’ at SAS Version 9 performance in a large systems setting.
- We wanted to establish performance and scalability expectations on a large enterprise-class server complete with enterprise-class storage.

While this strategy provided us with an advantage in identifying and resolving potential problems very early in the development cycle, it also presented some challenges. We had to identify tests for which the development was far enough along to provide a reasonable estimate of performance. For example, much of the testing was done with certain chosen analytic routines, not necessarily because of their business or performance priority, but because the application was more fully developed than others.

It thus turned out that we were limited somewhat arbitrarily to specific modules, and were not able to test complete solution environments, such as Marketing Automation (MA). However, this did not prove to be an obstacle, as we were nevertheless able to gather valuable insight with the available code.

Focus of the Test Plan

Due to the diverse and complex nature of the needed testing, testing was divided into 5 “tracks”.

- Track 1: Sort testing – verifying the proper operation and performance of a parallel sort operation.
- Track 2: Summary testing – verifying the proper operation and performance of a parallel summarization engine.
- Track 3: EM server – verifying the proper operation and performance of specific Enterprise Miner (EM) operations.
- Track 4: SPD server functionality – verifying proper operation of the Scalable Performance Data Server (SPDS).
- Track 5: Scalability – Verifying proper scaling, both in concurrent users (up) and quantity of data (out).

The testing was scheduled from late November 2001 thru mid-January 2002. During this time, the testing team had exclusive access to the HP Superdome and associated storage.

Implementing SAS Version 9 on the HP Superdome

Implementing and conducting SAS Version 9 and associated tests on the HP Superdome called for consideration of the HP Superdome architecture to optimally set up and run the test suites.

The HP Superdome used for the testing is located in the HP Development Alliances Lab (DAL) facility in Cupertino, CA. The HP Superdome was configured as follows:

- 2 cabinet HP Superdome server
- 64 PA 8600 552MHz processors, 1.5 MB on-chip cache
- 128 GB ECC Memory
- 92 PCI I/O Slots (32 are 66MHz capable)
- XP 512 Storage array, providing over 3 TB of RAID 0/1 storage, 16 Fibre Channel ports

The two cabinet systems were arranged as 16 cells, with each cell containing 4 CPUs and 8GB of memory. The cell is the basic building block of the HP Superdome system. Each cabinet can contain up to eight cell boards, which are plugged into the backplane of the cabinet.

In order to allow concurrent testing, the server was split into two 32-way servers using HP's hardware partitioning. This allowed for complete hardware and software isolation between the two servers, while optimizing overall system utilization. The XP 512 storage array was divided as well between the two partitions. A HP VA7400 Storage array was added to provide a different type of storage for testing.

The server was prepared and administered throughout the testing period by the DAL staff. The system was accessed both locally by SAS and HP development staff, as well as through secure Internet connections between the DAL and SAS World Headquarters in Cary, NC.

Results

Working together, HP and SAS have improved the scalability and performance of SAS Version 9. As noted by Keith Collins, the Chief Technology Officer for SAS:

"The next release of SAS software, Version 9, is a significant step towards a scalable, high performance architecture. The performance analysis of SAS V9 on HP Superdome turned out to be an important contribution.

As a result of the close work of the teams at SAS and HP, we've achieved performance and scalability gains for SAS Version 9. Discoveries made during the course of the testing have been incorporated into SAS.

Though the work was preliminary, testing of SAS V9 also provided HP initial insight into strategies for the optimal organization of storage (via interactions with the XP512). That insight that will prove crucial for HP and SAS's handling of NAS/SAN storage optimally in V9 and beyond.

This project is yet another example of the excellent R&D relationship between SAS and HP, and proves the value of such collaboration."

Measuring the effects of CPU and cache behaviour became important to establishing performance baselines. Development tools were used, and in some cases modified, to better understand the behaviours of the execution paths of SAS processing and evaluate its impact on the server. These learnings were then reviewed with the developers, who in turn made modifications based on the information gathered.

Another area of particular interest was the memory utilization patterns of the Sort and Summary routines and the associated parallel algorithms. Using one of the two partitions on the HP Superdome, SAS developers were able to construct a detailed performance study on the existing algorithms, as well as develop alternatives and integrate them into the study almost as they were being developed. This 'near-time' access proved to be very beneficial in evaluating alternatives and selecting the proper methods for performance and scalability.

Based on information gathered during this project, a different approach to optimizing memory access patterns was developed and implemented to allow the optimum use of larger memory regions.

Next Steps

To extend our successful collaboration with SAS R&D teams, we will continue to jointly analyze the data collected in order to enhance our understanding and exploiting the HP Superdome server with the SAS architecture.

To Learn More

There is much more to this story than what we have included in this R&D Update. We invite you to contact us to learn more about this project and the many other activities jointly being done by HP and SAS that can benefit you. We can be reached on the web at: <http://www.hp.com/go/sas>.

SAS and all other SAS Institute Inc. product or service names are registered trademarks of trademarks of SAS Institute Inc. in the USA and other countries. ® Indicates USA registration. HP and all other HP product or service names are registered trademarks or trademarks of Hewlett-Packard Company.

Technical information contained in this document is subject to change without notice.

© Copyright Hewlett-Packard Company 2002.

All Rights Reserved. Reproduction, adaptation, or translation without prior written permission is prohibited except as allowed under copyright laws.