

Hewlett-Packard Research & Development Update May 2001

Performance gains for 64-bit SAS[®] applications

More Power, Scotty!

When we are looking for more power out of a computer system for our applications, we traditionally look into three areas: faster CPU speeds, more system memory and faster disc input/output (I/O). Large SAS applications like we see in data warehousing, business and electronic intelligence, and data mining demand best-of-breed performance in all these areas. In this update we take a closer look at how HP and SAS R&D staff have worked together to make performance gains for SAS applications using HP's 64-bit architecture for HP-UX.

Benchmarking vs. Testing

One of the challenges we have in characterizing SAS applications is many SAS users develop their own custom applications using SAS software as a state-of-the-art 4th generation development environment. Thus, unlike pre-defined integrated SAS applications, there is usually not a one-size-fits-all approach we can take to predict the different ways that customers use SAS software. This comes down to learning to be careful about making broad performance statements for a specific customer situation. If time and resources permit, we always recommend you build a representative sample of your particular SAS work load and run a benchmark on different systems to make the best determinations.

However, over the years we have been able to characterize many different types of customer-developed SAS applications and we have developed basic SAS software test suites. These test suites give us insight into the different aspects of system performance based on SAS jobs that make sense to SAS customers. For the research results reported here, we use these test suites to evaluate the performance of SAS software in three configurations on HP-UX 11.0:

- SAS 6.12 (V6), which was used as a performance baseline,
- SAS 8.2 32-bit (V8-32),
- SAS 8.2 64-bit (V8-64).

For all results, performance is defined as time-to-completion of single batch jobs.

Test Suites

In order to evaluate the impact of the 64-bit architecture the broadest way possible we drew from multiple SAS application environments with several different SAS programs from each. The following table identifies these test suites and the respective SAS programs (Test Codes). We also include a short description of what the test program is doing using either a functional explanation or the name of the predominant SAS Procedure in the test program.

Test Suite	Test Code	Description
Computational A	CA1	Stepwise logistic regression
	CA2	Survival analysis – PHREG
	CA3	PROC MIXED
Computational B	CB1	PROC MIXED
	CB2	MEANS, FREQ, TABULATE
	CB3	Big GLM
	CB4	IML - Matrix
Data Mining	DM1	DMDB, NEURAL, DMREG
Input/Output	IO1	Create SAS datasets from flat files
	IO2	SORT, index datasets
	IO3	FREQ, summary
	IO4	UNIVARIATE, SORT
	IO5	MDDB

The test suites are of two general types: compute intensive and input/output. Here is how the different test suites are organized:

- **Computational A-** Simulates much of the activity typically performed in a data warehouse environment,
- **Computational B-** Simulates some of the compute-intensive processing customers typically do with SAS Software,
- **Data Mining-** Simulates the server-side processing done by the SAS Enterprise Miner application,
- **Input/Output Suite-** Simulates the activities typically found in a data warehousing environment.

The Testing Environment

The work on this project was done in the Corporate Technology Center (CTC) at SAS' worldwide headquarters in Cary, North Carolina. The testing environment consisted of:

Computer system:

- HP 9000/L2000 server,
- Four 440 MHz PA8500 CPUs,
- 4 Gbyte ECC memory,
- Two dual-channel SCSI controllers with four 9 Gbyte Seagate 10K RPM HDD.

Operating system:

- HP-UX 11.0 (June 2000 Quality Pack),
- Default kernel settings for HP-UX,
- Address space set at 2 Gbyte for all tests.


SAS Software versions:

- SAS 6.12 TS060,
- SAS 8.2 TS2M0 (32-bit),
- SAS 8.2 TS2M0 (64-bit).

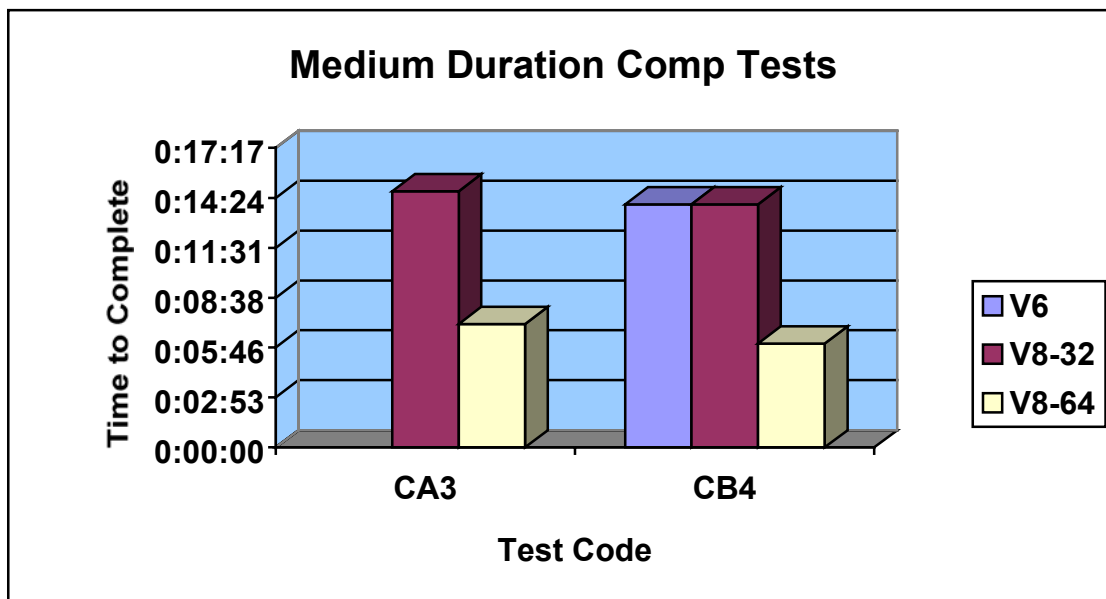
Results

Results are shown in the following series of bar charts where the time-to-completion of each of the programs is plotted, in group fashion, with each version of SAS Software. In some tests the V6.12 case was not run but we still show the benefits of 64-bit V8.2 versus 32-bit V8.2.

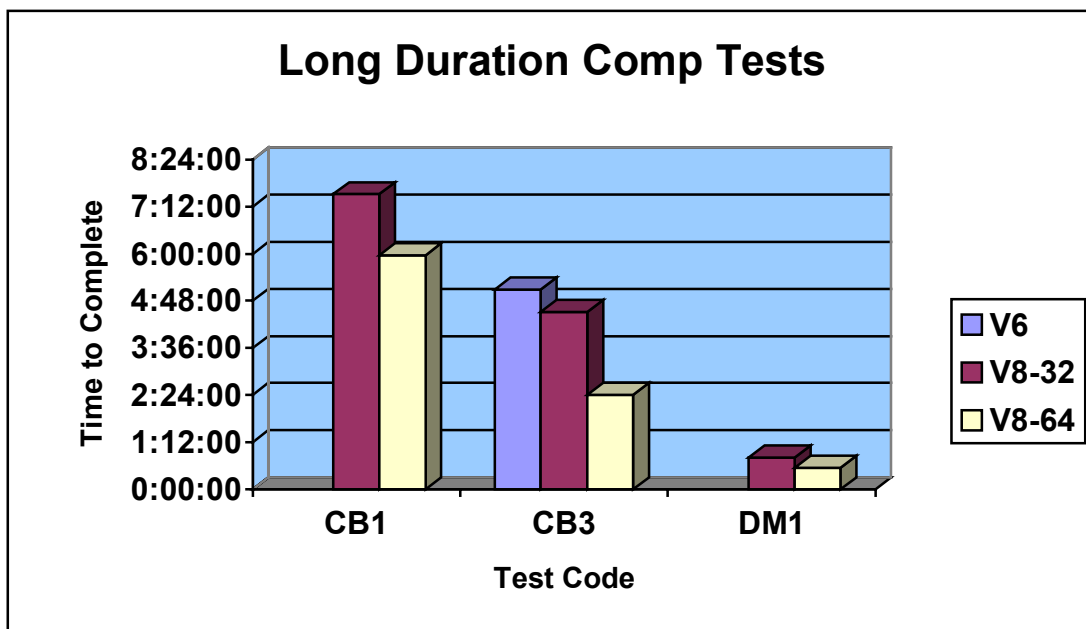
In the following chart we plot several of the computation tests together in what we call *short duration tests*. These are programs that complete in a few minutes and represent the building blocks of larger SAS programs with iterative procedures like those in survival analyses. In all cases there is improvement with V8 versus V6. In addition, we show an 11% to 18% improvement in performance from using the 64-bit SAS build versus the 32-bit version. The logistic regression in CA1 did not benefit as much in the 64-bit build as we would expect because of earlier tuning HP and SAS had done specifically related to the VSUBS work (Please see our April 2001 R&D Update for details).



In the next chart we show the results for two *medium duration computational tests*. Test program CA3 involves the PROC MIXED model and CB4 is matrix mathematics using SAS/IML. We achieved a 52% improvement in CA3 and a 57% improvement in CB4 with 64-bit V8 versus the 32-bit version. There was no change in time-to-completion for 32-bit V8 versus V6 for test program CB4. This draws even more attention to the conclusion that the 64-bit version of SAS provides the real improvements in this SAS program.




In the following chart we look at *long duration computational tests*. These are programs that run for several hours and place a long runtime load on the CPU and memory management aspect of the computer system. CB1 in another PROC MIXED program and we again achieve a significant improvement of 21%. CB3 is a larger linear equation model solution using GLM and we see a 47% performance increase from the 32-bit to the 64-bit versions. In DM1, the data mining suite, we see a 31% increase in performance from 32- to 64-bits.



In the final series of tests we ran several programs known to have extensive input/output operations. This *I/O Suite* is similar to the typical data warehouse update and report operations found throughout the industry. This suite includes creating SAS datasets from flat files, doing sorting, indexing, running summary statistics, MDDDB creation, etc. The following chart shows the results of these tests. We see general performance improvements from V6 to V8 owing to improved I/O work by SAS' staff for V8. The performance gains from 32-bit to 64-bit V8 are very modest. The work HP and SAS staff has done here at the compiler and build level for 64-bit

V8 isn't the sort that we would expect to significantly improve I/O. However, we must always check to make sure we don't gain performance in one area at the expense of another.



How We Did It

We have shown that with the 64-bit version of SAS V8.2 on HP-UX we can make significant performance gains in SAS compute intensive routines – from 11% to 57% - without degradation of I/O performance. These outstanding results come mainly from the compile options that were used to build the 64-bit SAS code, taking advantage of the 64-bit features found in HP's PA 2.0 architectures. The HP and SAS R&D staff worked closely together to deciding what compiler options to use in which module build situation to maximize the benefit and minimize problem areas.

Just as with the VSUBS work and MLIB routines accessed from SAS software, we reported on in earlier R&D updates, this 64-bit compiler build is part of the SAS code development process, so SAS users do not need to do anything special for improved performance on HP systems.

If you are ready to look at going to 64-bits for your current and future SAS solution needs HP is ready. Contact your SAS account manager to get the latest information or reach HP's SAS Alliance Team at the web address below.

Acknowledgements

We would like to acknowledge Leigh Ihnen and Margaret Crevar in SAS' Corporate Technology Center for their assistance in providing the test suites and with running these tests. We would also like to thank Clarke Thacher and Bill Cooper for providing their expertise in SAS software development on HP-UX.

To Learn More

We invite you to contact us to learn more about this project and the other activities jointly being done by SAS and HP that can benefit you. We can be reached on the web at:

www.unixservers.hp.com/partners/sas.

###

SAS and all other SAS Institute Inc. product or service names are registered trademarks or 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 2001.

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