



# **Performance Comparisons of Intel® Itanium® 2 Processor-Based Servers Running SAS® 9**

**at the Cornell Institute for Social and  
Economic Research (CISER)**

---

---

# Contents

- Executive Summary ..... 3**
- Background of CISER ..... 3**
- Early Experience with a 4-Way  
Intel® Itanium® 2 Processor-based Server ..... 4**
- Enter the UNISYS ES7000 Enterprise Server ..... 5**
- Upgrade to Second-Generation Itanium 2  
Processors with 6 MB Cache ..... 9**
- Comparing (Early Adopter) SAS® 9.1 with SAS® 9.0 ..... 10**
- Summary of Results ..... 11**
- Unisys ES7000 Enterprise Server Meets  
32-bit and 64-bit Needs ..... 12**
- Appendix ..... 12**

## Executive Summary

The Cornell Institute for Social and Economic Research (CISER) set out to increase the performance of some of its processing- and memory-intensive SAS® applications. CISER had to acquire the necessary resources to run 64-bit SAS® 9.0 when it became available for the Microsoft Windows Server\* 2003 operating system. Yet more than 600 researchers still needed broad-ranging 32-bit support. CISER needed to find an economical way to manage both 64-bit and 32-bit user communities requiring more power.

This white paper details CISER's findings from testing done with SAS 9.0 and SAS 9.1 on Intel® Itanium® 2-based servers. For the most part, these findings reinforce the findings of previous testing by Intel, Unisys, and SAS. Specifically, the latest versions of SAS provide performance increases for all servers tested, but the Intel Itanium architecture becomes increasingly advantageous as particular resource demands increase. Those demands include the need for multi-threaded processing, multiple simultaneous tasks (by one or more users), memory bus load, and data set size. As these resource demands increase, additional processors lead to better overall performance.

By moving to a more powerful multiprocessing server, CISER will also realize lower management costs. CISER will be able to phase out 12 smaller servers by replacing them with one 24-processor server. The consolidation to a single server will streamline the deployment of applications, operating systems, and upgrades and greatly simplify server management. With funding tight in university environments, it's important for CISER to be able to manage an increasing amount of computing power without increasing headcount. With its move to a single Unisys\* ES7000 server, CISER expects to spend 30 to 50 percent less time installing and maintaining software.

---

## Background of CISER

CISER was founded in 1981 to provide research support services tailored to the needs of social scientists at Cornell University. It is committed to enhancing the environment for social and economic research, and it encourages inquiry that cuts across traditional academic departments. CISER is supported jointly by seven colleges and is comprised of more than 400 faculty and senior researchers who represent a broad spectrum of interests, from basic theory to direct application.

CISER's support structure offers a range of services to CISER members and their graduate students and research staff, including proposal preparation, data archival, consulting, and computing resources. In the year 2000, CISER created the Cornell Restricted Access Data Center (CRADC), which provides a secure computing facility. CRADC is currently partnering with the U.S. Census Bureau, the Urban Institute, and the University of Maryland in a \$4.1 million National Science Foundation Social Data Infrastructure grant.

Many of the social science researchers supported by CISER rely on sophisticated software packages that are well designed to meet their needs. In general, these applications are designed for 32-bit architecture and provide a graphical user interface for the Windows\*

operating system, which is preferred by the majority of these researchers. There are times, however, when the combination of the available software and the 32-bit architecture does not provide the overhead needed for working with extremely large data sets or performing highly CPU-intensive operations. The CRADC researchers, for example, were regularly hitting the memory and processor limits of their 32-bit multi-processor servers using SAS 8.2.

In early 2002, Intel Corporation introduced the Itanium 2 processor, which was the first to feature the Explicitly Parallel Instruction Computing (EPIC) architecture. This new architecture accommodates 64-bit data and addressability in addition to new levels of parallelism. It allows software to address more data simultaneously. Running its software on a larger memory system like this would enable CISER to overcome former memory and processor limits, potentially increasing performance dramatically.

## Early Experience with a 4-way Intel® Itanium® 2 Processor-based Server

When SAS led the statistical software industry by being the first to announce its intention to release SAS 9.0 for the 64-bit Windows Server 2003 operating system, the CRADC group was ready with code to be ported to just such a system but without the resources to do so.

Initially, CISER staff arranged with Intel and SAS to beta test a combination of hardware and software. This combination consisted of an early adopter release of SAS 9.0, a beta release of Windows Server 2003, and Intel's latest generation of Itanium 2 processors (1 GHz with 3 MB cache). Intel provided CISER with a 4-way Itanium 2-based test system (see the appendix for system configuration details).

CISER installed this system behind a firewall at CRADC, and the machine was turned over to the CRADC researchers, who immediately ported their code and restructured it to take advantage of the increased resources available. Right away, these researchers reported success in processing larger data sets in less time than they were able to do previously. Encouraged by this experience, the CISER support staff set out to do more methodical testing to sort out to the extent possible what the primary contributing factors were for the observed performance increase.

A related issue of concern for CISER was whether, and to what extent, the Itanium EPIC architecture might serve its 600+ general users with equally important, but less resource-demanding, applications. The staff understood that the overhead required for a 64-bit system could actually slow down less processor- or memory-intensive tasks. They needed to determine the point at which use of these resources becomes practical, as well as to examine the ability of the system to handle multiple users with a broad range of applications.

The structure of CISER's tests was heavily shaped by the resources available, including not only hardware and software resources but also the time and expertise of the people doing the testing. In particular, the earliest comparisons are muddled by differences in processor speed, available RAM, and disk configurations on the test machines available (see Table 1 in the Appendix).

CISER researchers used elapsed time as their measure of performance, reported by SAS as "real time." The test code was run with the default system options on every machine and was not maximized for efficiency on any machine. Each set of tests was run only once. Although this limits the conclusions that can be drawn from the tests, it does represent a very common real-

world scenario, and CISER'S findings are remarkably similar to those derived from the more robust testing done by Intel and SAS software engineers.

CISER worked with the CRADC staff to extract critical pieces of SAS code from their research and modified it to use with carefully simulated data provided by CRADC. The data for the tests maintained the statistical properties of the actual data but contained no confidential information. These data sets were comparable to the underlying data in length (i.e. number of observations) but not in width, since the number of variables per observation was considerably smaller than that of the underlying data. The final test code used the MP CONNECT\* feature of SAS to generate from one to 10 simultaneous "remote" SAS sessions that sorted and merged up to 10 sets of three related files. Each set of three files contained synthetic integrated labor and product market information for 16 million jobs. Among other things, the test code reflected the heavy I/O requirements of the actual research data and included several procedures that were expected to be multi-threaded beginning in SAS 9.0.

CISER started by comparing its existing Intel® Pentium® III Xeon™ processor-based resources running SAS 8.2 with the new Intel system containing four first-generation Itanium 2 processors and SAS 9.0. The new system managed to handle up to 13 simultaneous sessions, while the older system only managed seven with the test code. In addition, the Itanium system ran multiple simultaneous jobs in much less time, depending on how many sessions were being run. In general, the Itanium-based jobs were completed in only half the time of the Intel Xeon processor-run jobs (and in some cases in only one-quarter of the time). It is assumed that this combined upgrade of both processors and SAS versions mimics to some extent the big jump in performance experienced by the CRADC researchers.

Later testing with the same code demonstrated that a large portion of the increased ability to manage multiple sessions might reasonably be attributed to the upgrade in SAS. In those tests, a 2-way Intel Pentium III Xeon processor-based server with SAS 9.0 successfully ran 17 simultaneous sessions. The newer, 4-way Itanium 2 processor-based (6 MB) server ran 19 sessions.

Because the CRADC research team had been working with "wider" data sets and had been known to reach the memory limits of its own Pentium III Xeon processor-based servers, CISER staff continued to investigate the capabilities of 64-bit architecture as well as examine the differences of the individual steps within its test code. They found that the comparative advantages of the different

test machines varied according to the task, and that for the more memory-intensive tasks in the test code, the Itanium 2-based server showed clear advantages.

To review the testing conditions: CISER’s code was not optimized for use on any server, and the full code included a wide mix of data steps and procedures. Some of these tasks were memory-bound, some CPU-bound, and some I/O-bound. The test data sets fell into the “fuzzy” range of neither small (measured in MB) or very large (measured in TB). So the testing scenario did not clearly fit into any scheme adhering strictly to available sizing guides, but it very much represented a common type of academic research application. The fact that CISER serves many users with a broad range of skills and computing needs – with individual needs varying over time – makes the flexibility of a mix of resources even more crucial to its success.

### Enter the Unisys\* ES7000 Enterprise Server

After testing the 4-way, 3 MB Itanium 2 processor-based system loaned by Intel, CISER gained the opportunity to test its code on a Unisys ES7000 Enterprise Server with 16 Itanium 2 processors (also 3 MB cache) loaned by the Cornell Theory Center (CTC). Again, SAS provided a license for the testing, and CISER had the opportunity to explore the effect of increasing the number of processors from four to 16.

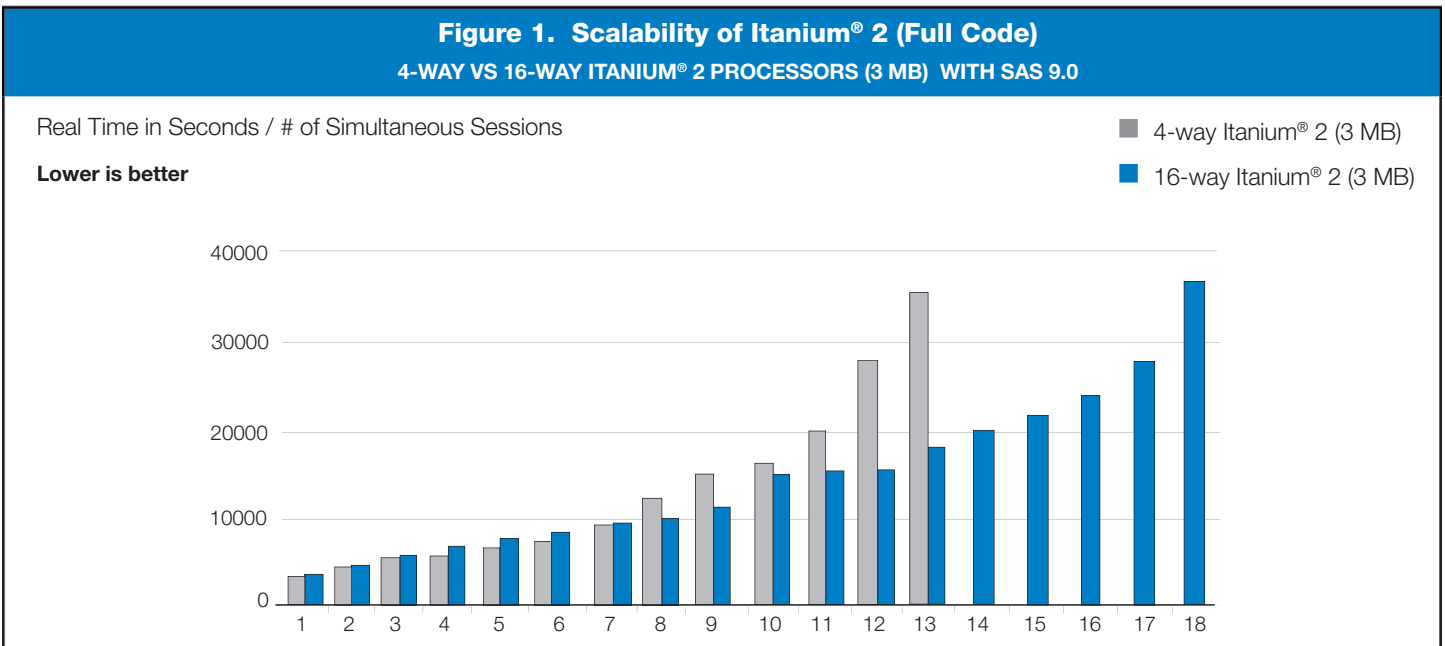
The Unisys ES7000 family of servers offers a variety of possible configurations, including the unique combination of both 32-bit Intel Xeon processors MP and EPIC architecture-based Itanium 2 processors in an adaptive infrastructure managed as a single operating environment.

The test code CISER used was the same code used with the Unisys system with the 4-way system, though the number of simultaneous sessions increased. Each session sorted and merged a separate set of three related files containing synthetic information about individuals, employers, and jobs. Each merged file represented integrated labor and product market information for 16 million jobs.

#### SCALING TEST OF 4-WAY VS 16-WAY ITANIUM 2-BASED SERVERS

For the scaling tests, the systems were configured to match the size of the server. For example, the 16-way ES-7000 server included significantly more disk space and memory, as one would expect of a larger server. The additional disk space and memory ensured that items did not become a bottleneck in the scaling study. Table 1 in the Appendix shows the system configurations.

For seven or fewer sessions, the 4-way Itanium 2-based pre-production server processed the full test code the fastest. However, once the number of simultaneous sessions exceeded seven, the 16-way server was faster. In addition, the 16-way server supported more simultaneous sessions. (See Figure 1.)




**Figure 2. Performance with Full Test Code**

16-WAY UNISYS VS 4-WAY INTEL® ITANIUM® 2-BASED SERVER (3 MB) WITH SAS 9.0

Real Time in Seconds / # of Simultaneous Sessions

Lower is better



The 4-way Itanium processor (3 MB) out-performed the 16-way Itanium processor (3 MB) in the first seven simultaneous sessions. However, for eight sessions and above, the 16-way was quite a bit faster – 14% faster at eight sessions and up to 47% faster at 13 sessions. In addition, the 16-way system was able to process at least 18 sessions compared to only 13 for the 4-way server. (See Figure 2.)

**RANK PROCEDURE FOR 16-WAY VS 4-WAY ITANIUM 2 PROCESSORS**

CISER compared several key steps within the full test code as well. One was the memory-intensive but non-multi-threaded RANK Procedure. For this procedure, the 4-way system was more

efficient for seven through 11 sessions, after which its performance deteriorated dramatically before failing, possibly due to memory or disk space limitations. After 11 sessions there was a very clear advantage to having more processors. (See Figure 3.)

**SUMMARY PROCEDURE FOR 16-WAY VS 4-WAY ITANIUM 2-BASED SERVER**


For the multi-threaded Summary Procedure, the 16-way Itanium 2 processor-based (3 MB) system was dramatically more efficient than the 4-way. At 13 sessions, the 16-way server showed a 92% increase in performance; the procedure ran on the 16-way server in

**Figure 3. Average Time for Procedure Rank (Memory-Intensive)**

4-WAY VS 16-WAY ITANIUM® 2 PROCESSORS (3 MB) WITH SAS 9.0

Real Time in Seconds / # of Simultaneous Sessions

Lower is better




**Figure 4. Average Time to Complete Summary Procedure (Multi-Threaded)**

4-WAY VS 16-WAY ITANIUM® 2 PROCESSORS (3 MB) WITH SAS 9.0

Real Time in Seconds / # of Simultaneous Sessions

Lower is better

■ 4-way Itanium® 2 (3 MB)  
■ 16-way Itanium® 2 (3 MB)



only 8% of the time it took on the 4-way server. Soon thereafter, the 4-way failed, possibly due to memory limitations. The 16-way server continued to perform very evenly through 17 simultaneous sessions. (See Figure 4.)

Not addressed here is the extent to which simultaneous sessions running identical multi-threaded procedures creates competition for resources. Determining the effect on overall performance would require testing by limiting the number of threads used per procedure. Even without such optimization, the additional processors clearly offer an advantage for multiple sessions per user or multiple users.

**I/O-INTENSIVE TEST: DATA STEP WITH MERGE FOR 16-WAY VS 4-WAY ITANIUM 2 PROCESSOR**

The performance differences in this data step were not nearly as clear as with the procedure steps shown above with CISER's code. The 4-way system was somewhat more efficient at up to 11 simultaneous sessions. CISER suspects that I/O bottlenecks may account for the generally poor performance of both machines in its data steps, with "poor" indicated by the increasingly long time it took for either machine to complete one session as the number of sessions increased. (See Figure 5.)


**Figure 5. Time to Complete Data Step with Merge (I/O-Intensive)**

4-WAY VS 16-WAY ITANIUM® 2 PROCESSORS (3 MB) WITH SAS 9.0

Real Time in Seconds / # of Simultaneous Sessions

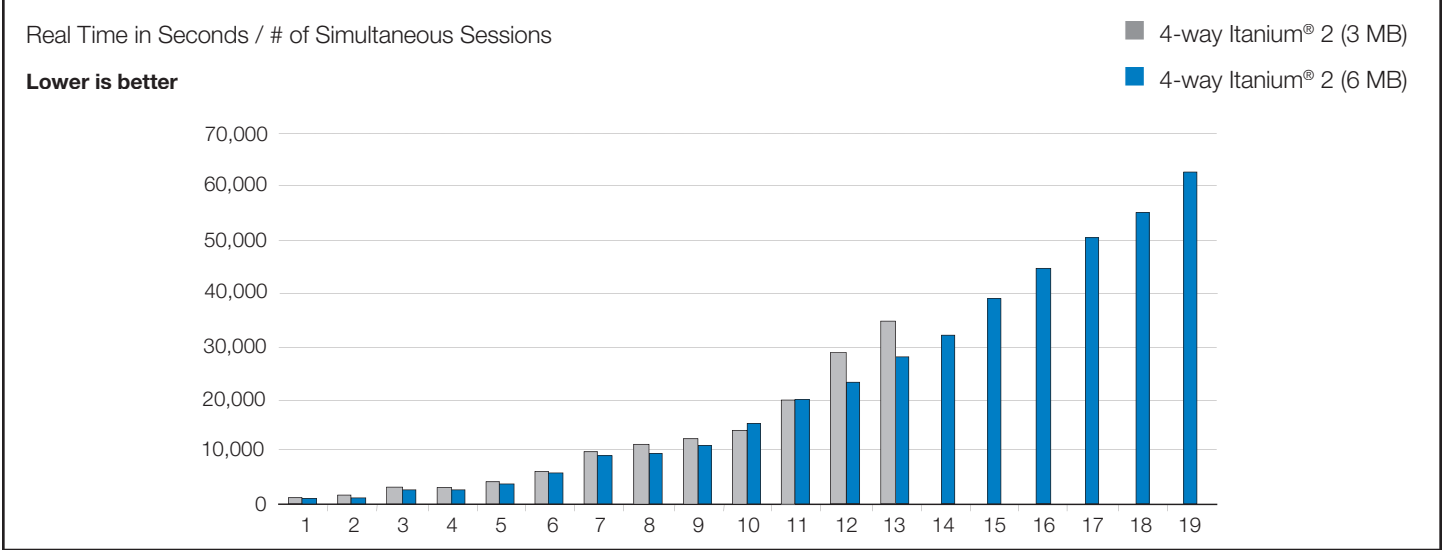
Lower is better

■ 4-way Itanium® 2 (3 MB)  
■ 16-way Itanium® 2 (3 MB)



**Figure 6. Maximum Time to Complete Full Code**

3 MB VS 6 MB ITANIUM® 2 PROCESSORS WITH SAS 9.0



It is this factor which appears to account for the fact that in many cases CISER's test code actually took longer to run with simultaneous sessions generated by MP CONNECT than it would have to run each session in sequence. In addition, there are multiple ways to make this particular test code more efficient that are not addressed in this paper.

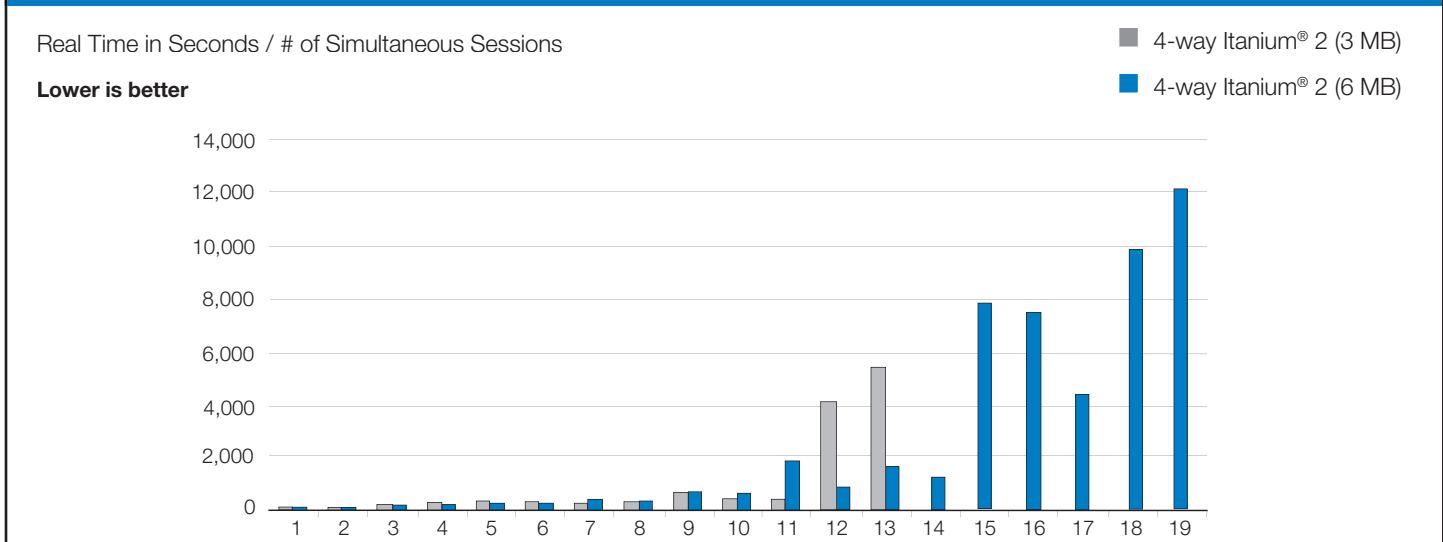
**RESULTS SUMMARY FOR 16-WAY VS 4-WAY ITANIUM 2 PROCESSOR (3 MB)**

Overall, CISER's testing of the 16-way versus the 4-way Itanium 2-based systems (with 3 MB cache) showed a clear advantage

of more processors as the number of simultaneous sessions increased, especially with memory-intensive and multi-threaded procedures. The first data step in the code seems to be the limiting factor in how well both machines perform overall, perhaps related to I/O bottlenecks. Readers should remember that all of these tests were performed without making any adjustments in the number of threads or processors dedicated to each session or in finding ways to make the code more efficient.

**Figure 7. Average Time for Procedure Rank (Memory-Intensive)**

3 MB VS 6 MB ITANIUM® 2 PROCESSORS WITH SAS 9.0




**Figure 8. Average Time for Procedure Summary (Multi-Threaded)**

3 MB VS 6 MB ITANIUM® 2 PROCESSORS WITH SAS 9.0

Real Time in Seconds / # of Simultaneous Sessions

Lower is better

■ 4-way Itanium® 2 (3 MB)  
■ 4-way Itanium® 2 (6 MB)



**Upgrade to Second-Generation Itanium 2 Processors with 6 MB Cache**

About this time in CISER's testing, Intel made available to CISER a new 4-way test server based on beta versions of the second-generation Itanium 2 processor. The newer Itanium 2 processor has 6 MB of L3 cache (versus 3 MB in the previous processor) and runs at 1.50 GHz (versus 1.00 GHz for the previous processor). CISER ran the same battery of tests to compare the 4-way Itanium 2 processor with 6 MB cache server with the 4-way Itanium 2-based server with 3 MB cache.

**SCALING TEST SUITE FOR 3 MB VS 6 MB ITANIUM 2 PROCESSOR**

The 6 MB Itanium 2-based server out-performed the 3 MB Itanium 2-based server consistently and was able to handle approximately 40% more simultaneous sessions with CISER's code. In this case, the difference in the number of sessions handled could be directly attributed to the upgrade in processors, since the available physical memory and disk space were identical for each series of tests. A look at some of the specific steps inside the test code (Figures 6 to 9) show that the performance of the 6 MB Itanium 2 processor was consistently better than the 3 MB Itanium 2 processors in memory-intensive, multi-threaded, and heavy I/O steps as well.


**Figure 9. Full Time to Complete Data Step with Merge (I/O Intensive)**

3 MB VS 6 MB ITANIUM® 2 PROCESSORS WITH SAS 9.0

Real Time in Seconds / # of Simultaneous Sessions

Lower is better

■ 4-way Itanium® 2 (3 MB)  
■ 4-way Itanium® 2 (6 MB)



## RESULTS SUMMARY 6 MB VS 3 MB ITANIUM 2 PROCESSOR

CISER's tests showed that the newer Itanium 2 processor with 6 MB cache provided consistently better performance as the number of simultaneous sessions increased, especially with memory-intensive and multi-threaded procedures. In every case, the 6 MB Itanium 2 processors were generally faster, and were able to sustain more simultaneous sessions than the 3 MB Itanium 2 processors. This was especially important to CISER, since they were using the number of sessions handled as a rough proxy for number of users.

### Comparing (Early Adopter) SAS 9.1 with SAS 9.0

In September, 2003, CISER obtained an early adopter version of SAS 9.1 for the 4-way (6 MB) Itanium 2-based server and was able to make direct comparisons between SAS 9.1 and 9.0 with its test code. Because the tests were run on the same server, all differences in performance could be attributed directly to the change in SAS versions.

#### COMPARING SAS 9.1 TO SAS 9.0 WITH THE FULL CODE

Improvements in performance were evident in the comparison of real time to complete the full test code, as shown in Figure 10.

#### THE RANK PROCEDURE SAS 9.1 VS SAS 9.0

The results from the relatively memory-intensive (each procedure used roughly 170 MB memory) but not multi-threaded RANK Procedure were fairly dramatic. The total real time to complete the

procedure with SAS 9.1 showed a speedup of between 25% and 97% for up to 13 simultaneous sessions. After 13 simultaneous sessions, the performance becomes erratic and deteriorates with both versions of SAS (more for SAS 9) on this 4-way machine; this is not shown here.

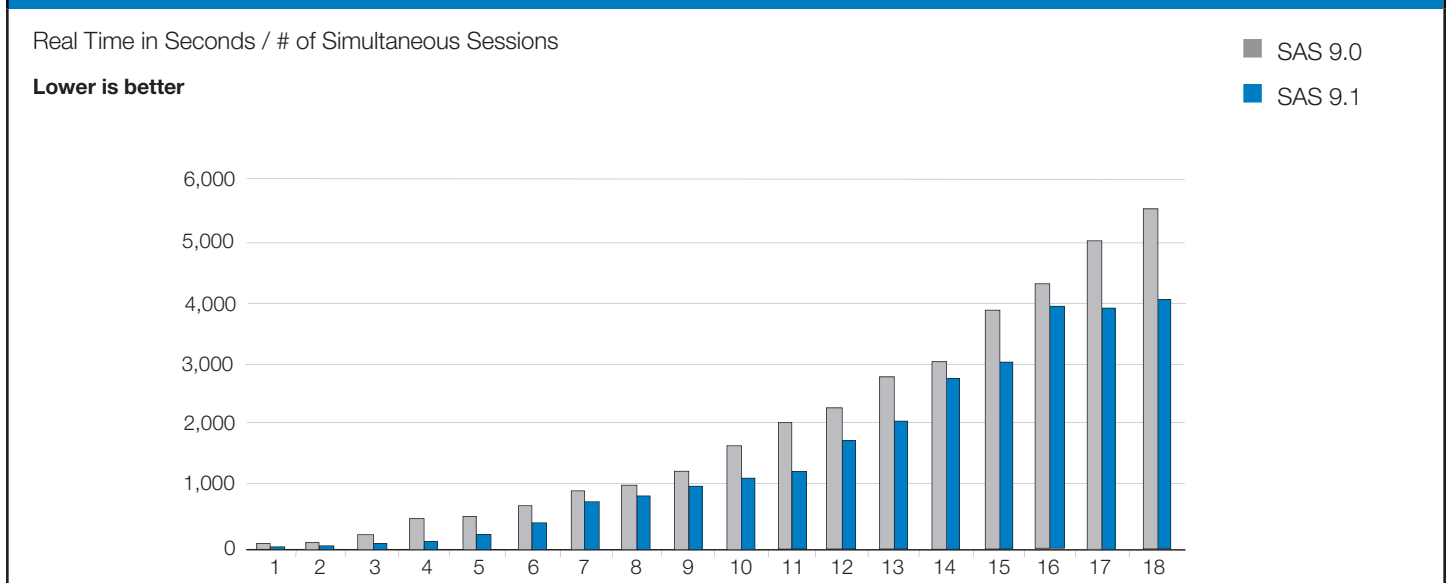
One of SAS' goals for SAS 9.1 was to increase performance. To achieve this goal, the SAS Intel Advanced Research Center (SIARC), a joint effort between SAS and Intel, tuned the compiler and math routines to reduce runtimes. PROC RANK is one procedure to benefit from these improvements. In addition, SAS developers put a specific effort into improving memory usage for PROC RANK. Together these changes contribute to the significant performance improvements seen in this test.

#### THE SUMMARY PROCEDURE (MULTI-THREADED) SAS 9.1 VS SAS 9.0

Figure 12 below shows the very dramatic increase in real-time performance of SAS 9.1 over SAS 9.0 for the multi-threaded Summary Procedure. In spite of the fact that these multi-threaded processes were occurring simultaneously across all processors and most likely competing for resources, the upgrade in software appears to greatly improve performance, even with unoptimized test code. In fact, the speedup in performance is so greatly dramatic (ranging from 92 to 99%) that the SAS 9.1 times barely show on Figure 12!

The compiler and math routine optimizations incorporated by SAS likely contributed to the improved PROC Summary performance.

**Figure 10. Maximum Time to Complete Full Test Code**  
4-WAY 6 MB ITANIUM® 2 PROCESSORS, SAS 9.0 VS SAS 9.1




**Figure 11. Average Time for Procedure Rank (Memory-Intensive)**

4-WAY 6 MB ITANIUM® 2 PROCESSORS, SAS 9.0 VS SAS 9.1

Real Time in Seconds / # of Simultaneous Sessions

Lower is better



In addition, SAS developers put a specific effort into improving the threading behavior of PROC Summary. Together these changes contributed to the dramatic improvement seen in this test.

**RESULTS SUMMARY FOR SAS 9.1 VERSUS SAS 9.0**

CISER’s comparison of SAS 9.1 and SAS 9.0 on the same 4-way Itanium 2-based server showed an obvious and substantial increase in performance for the full code as well as various steps within the code. (See Figures 10 to 12.)

**Summary of Results**


Overall it appears that there is not one ideal server configuration but that the “best” configuration depends on the demands placed on the servers in terms of memory, CPU, and I/O requirements. In addition, the number of users and the number and type of tasks (data steps and procedures) can also be crucial in determining the optimal server design. Within the Itanium 2 architecture, it seems clear that the newer Itanium 2 processor with 6 MB cache processors are consistently faster than the earlier 3 MB Itanium

**Figure 12. Average Time to Complete Procedure Summary (Multi-Threaded)**

4-WAY 6 MB ITANIUM® 2 PROCESSORS, SAS 9.0 VS SAS 9.1

Real Time in Seconds / # of Simultaneous Sessions

Lower is better



processors and that more processors are generally better as the demands on the system increase. It is also clear that SAS 9.1 managed multiple multi-threaded procedures and heavy I/O data steps more handily than SAS 9.0.

### UNISYS ES7000 ENTERPRISE SERVER MEETS 32-BIT AND 64-BIT NEEDS

Based on CISER's test results and its success in using the Cornell Theory Center's ES7000, Unisys has joined with Intel and SAS in helping CISER design a system that will meet the widely varying needs of all of CISER's users. The Unisys ES7000 offers a unique

opportunity to combine both 32-bit Intel Xeon processors and 64-bit Itanium 2 processors in one integrated system running the Windows Server 2003 operating system. CISER has recently purchased a Unisys ES7000 with sixteen (16) 32-bit Xeon processors and eight 64-bit Itanium 2 processors attached to a 5 TB EMC SAN\*. This new system should allow CISER to provide the latest generation of resources for social science users and software applications that depend on IA 32 architecture along with the newest 64-bit processing capabilities for the more high-end applications and users.

## Appendix

Table 1. CISER Test System Configurations		
First Intel server at CISER	Unisys server at Cornell Theory Center	Second Intel server at CISER
Intel® Server Board SR870BN4-based server	Unisys ES7000* Orion Model 130	Intel Corporation S870BN4A
4-way 1 GHz/3 MB cache Itanium® 2 processors	16-way 1 GHz/3 MB cache Itanium 2 processors	4-way 1.5 GHz/6 MB cache Itanium 2 processors
SAS 9.0* Microsoft Windows.NET* Enterprise Edition, 64-bit RC 1	SAS 9.0 Microsoft Windows.NET Datacenter Server, Version 5.2.3663	SAS 9.0, then SAS 9.1 Microsoft Windows Server* 2003 , Enterprise Edition, Version 5.2.3790
16 GB RAM 60 GB Disk	64 GB RAM 136 GB Disk	16 GB RAM 60 GB Disk

Solution provided by



For more information, visit  
<http://intel.com/go/itanium2>

Intel, and Intel Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

\* Other names and brands may be claimed as the property of others.

Servers provided by

