

Intel® Architecture for Large Scale SAS® Computing: Sizing and Performance Considerations for the Unisys® ES7000 Enterprise Server Family

WHITE PAPER

Whether you're developing a new SAS solution or simply want to reduce costs of your current, mixed SAS environment, the Unisys ES7000 enterprise server family offers a potent combination of mainframe-level power, reliability and manageability with industry-standard agility, simplicity and cost savings provided by Intel® architecture. This guide summarizes SAS sizing and performance test results that can help you put the ES7000 to work on economically and incrementally meeting the growth needs of your large scale SAS deployments.



Contents

The Best of Both Worlds for High Performance SAS	2
Large Scale SAS Computing	2
Sizing and Performance Considerations	4
Four Key Findings	4
Summary	6
More Information	6
Review the Full Paper	6

The Best of Both Worlds for High Performance SAS

In today's highly competitive, customer-oriented business climate, turning volumes of data into in-depth understanding and actionable information can spell the difference between thriving and failing. It's no wonder companies are paying increasing attention to their business intelligence solutions, asking increasingly sophisticated questions, processing ever larger volumes of data and demanding results faster.

Given these demands, many companies are turning to the Unisys ES7000 family, a line of enterprise servers that combine Unisys' mainframe heritage with the performance, flexibility, scalability and cost advantages of Intel® architecture. Based on Unisys' unique Cellular MultiProcessing (CMP) architecture, the ES 7000 family allows companies to deploy from 4 to 32 Intel® Xeon™ processor MP and Intel® Itanium™ 2 processors in a single scalable, powerful and highly manageable platform. The resulting system offers new options to reduce complexity, risk and cost in SAS deployments with large data sets, many users, compute-intensive processing and/or numerous concurrent SAS jobs.

In the Fall of 2002, Intel and Unisys evaluated sizing and performance considerations for SAS solutions on the ES7000. Their results offer information that may be helpful to businesses in determining how to use the ES7000 platform to enhance the performance, robustness and cost-effectiveness of their large scale SAS deployments.

Large Scale SAS Computing

Large scale SAS computing encompasses a wide range of scenarios, all of which are well suited to the ES7000's flexible architecture. These scenarios are determined by factors such as:

- Number of simultaneous users or jobs.** As more jobs are processed simultaneously, a more powerful system configured with more CPUs can provide optimum performance. The ES7000 platform can handle multiple jobs per CPU and more users and jobs per overall system.
- Number of calculations per data point.** As the calculations per data point increase, the demands on the processor(s) rise. Extract-Transform-Load (ETL), Query-Reporting (Q&R) and data mining workloads in particular may put a significant load on the processors.
- Dataset size.** As dataset sizes increase from megabytes to gigabytes and terabytes, you'll benefit from the ES7000's fast system bus and a processor such as the Intel Itanium, which provides virtually unlimited addressability.
- Memory bandwidth requirements.** Tasks such as the statistical analysis of a large dataset can be particularly memory bandwidth-intensive and therefore likely to benefit greatly from the ES7000's fast memory bus and the Itanium processor's memory bandwidth.
- Amount of threading.** Highly threaded applications like general linear model (GLM) analysis benefit from higher performing, multi-CPU platforms. In addition, the Intel Xeon processor includes Hyper-Threading Technology (HT Technology) that enables each processor to execute two threads simultaneously.

ES7000 Advantages for Large Scale SAS Deployments

ES7000 Feature	Benefits for Large Scale SAS
Performance comparable to proprietary RISC platforms in a single machine at a fraction of the cost	<ul style="list-style-type: none"> • Robust, manageable platform with significant reduction in complexity • Better price/performance, scalability and simplicity than commodity, “small box” approaches • Accelerated return on investment
Architecture based on Intel® processors and optimized for high performance and availability	<ul style="list-style-type: none"> • Outstanding capacity, performance and throughput • Significant reductions in total cost of ownership (TCO) • Widespread availability of trained experts • Assurance of future performance advances
Ability to start with 4 Intel processors and grow in same box to 32 processors	<ul style="list-style-type: none"> • Flexibility and power by scaling up to add capacity as needed • Cost savings resulting from pay as you grow scalability • Investment protection and reduction in future spending
Ability to combine Intel® Xeon™ processor MP and Intel® Itanium™ 2 processors in a single system	<ul style="list-style-type: none"> • Mix and match flexibility • Ability to scale up by upgrading to next-generation processors • Investment protection • Enhanced performance by running SAS routines on the processor architecture to which they're best suited
Microsoft Datacenter Edition operating system	<ul style="list-style-type: none"> • Benefits of Windows standardization extended into the data center • Enhanced ease of use and simplicity • Significant reductions in TCO
Sophisticated management tools	<ul style="list-style-type: none"> • Increased manageability and cost savings from the ability to consolidate multiple SAS solutions and mixed SAS environments on a single, powerful platform • Easy ability to partition the system to support mixed environments
Single-source service and support from global solutions leader	<ul style="list-style-type: none"> • Minimal risk and complexity • Faster deployments through expert consulting • Advanced optimization and tuning services available

- **Business urgency.** Business intelligence results must be timely to be useful. The ES7000 is an outstanding choice when maximum throughput for complex jobs is a necessity.

- **Hybrid SAS computing.** If a company is running several different SAS applications, it can be very cost effective to consolidate them onto a single ES7000 system.

- **Heavy use of SAS across a network.** Since SAS workloads often involve large datasets, network delays can become an important consideration. Reducing network delays for the temporary work area is one effective way to optimize performance. The ES7000's high speed internal bus helps enhance the throughput of networked solutions.

- **SAS version.** Many SAS System 9 workloads run much faster on a larger system with more CPUs. Some analytical operations along with sorting and aggregation routines are threaded in SAS System 9; if you're using these routines, you may want to use SAS System 9 on a larger system for faster job completion.

Some companies find that their workload is high enough on just one of these factors to push them into the large scale category.

Others say it's a combination of factors that drives them toward deploying a large scale SAS solution. In either case, companies where SAS analysis is mission critical are well advised to consider the ES7000 as a way to gain the performance their most demanding SAS solutions require, as well as to consolidate multiple SAS solutions—such as data marts, a data warehouse, or analytics and reporting—onto a single platform.

Sizing and Performance Considerations

Given the need to deploy a large scale SAS solution, a number of further questions arise, including how large a system, the preferred number of users per processor, and whether to use the Intel Itanium 2 processor or Intel Xeon processor MP. To provide guidance in answering these questions, Intel and Unisys tested several ES7000 configurations on a range of SAS routines, including both single and multi-user usage models. The results revealed several performance trends that can help you configure a system to meet your particular needs. These results do not eliminate the need to benchmark your own solutions, but they do provide some general guidance. Because your environment and solution set are unique, we recommend that you work with your sales team to configure a system that suits your particular needs. In addition, Intel and SAS have created the SAS Customer Knowledge Practice to help you choose and optimize your system. Unisys, Intel and Intel® Solution Services also provide professional sizing and configuration services.

Test Specifics

ES7000 configurations	16 Itanium 2 processors 1 GHz 1.5 MB L3 cache 64 GB DDR memory Windows .NET Datacenter Server RC1 build 3663 Tested in 4-way, 8-way and 16-way partitions
	8 Intel Xeon processor MP, 1.60 GHz 4 GB SDRAM Windows 2000 Datacenter Server, build 2195 Service Pack 2 Hyper-Threading disabled Tested in 4-way and 8-way partitions
SAS software	SAS System 8.2 SAS System 9
Texts—Single user	Analytical tests (stepwise linear regression, general linear model and others) Enterprise Miner* to perform data mining tests Data manipulation on a large volume of input data Extract Transform and Load tests Query and reporting tests
Tests—Multi-user	As above, plus a mix of usage models—10% analytical, 60% Query and reporting, 30% data manipulation

Four Key Findings

The results are organized around four key findings:

- 1. Scaling.** System scaling has a different effect on performance in single-user versus multi-user environments, as well as on non-threaded versus threaded applications.
- 2. Threading.** Threads executing simultaneously can contend for processor time, affecting system performance.
- 3. Memory bandwidth.** Memory bandwidth can have a significant effect on performance, depending on the application and dataset.

4. Simultaneous users/jobs. The number of simultaneous users provides a good gauge of the optimum number of processors.

Scaling

System scaling has a different effect on performance in single-user versus multi-user environments, as well as on non-threaded versus threaded applications.

Using SAS System 9, a single user can take advantage of multiple CPUs to complete a complex job. This benefit is best realized when using procedures that are threaded internally. Scaling to four or more CPUs can provide significantly better performance for threaded procedures such as REG (regression), GLM (general linear regression), DMINE (data mining), LOESS,

DMREG, SORT, SUMMARY/MEANS, and SQL. The relative scalability of these procedures varies, since they depend on other system resources that may not increase with the number of CPUs.

Individual jobs that thread well can continue to improve in performance up to 16 CPUs and possibly beyond. The primary example from our test suite was the general linear model (GLM) using a large dataset.

This job ran the fastest on a 16-CPU system, indicating that large scale multiprocessing can greatly improve application performance.

Threading

Threads executing simultaneously can contend for processor time, affecting system performance.

ES7000 configurations with up to 32 CPUs offer outstanding performance for highly threaded applications such as general linear model analysis of very large datasets; for example, an application analyzing consumer preferences for a given product category across a large retail chain.

Managing multiple threads does incur a performance penalty, particularly when a single job runs across more than four to eight processors. You can limit the number of processors allocated to an individual job by setting the SAS thread count (CPU-COUNT) or dividing the system into partitions. Either approach modularizes the server's processing power, so that no single job attempts to distribute operations across all the available processors at once. By minimizing conflicts between jobs, this strategy ensures maximum throughput for all applications and users.

Memory Bandwidth

Memory bandwidth can have a significant effect on performance, depending on the application and dataset.

When your application must process a large volume of data in a single pass, memory bandwidth can be the chief determinant of overall performance. For example, statistical analysis using large datasets can benefit greatly from an ES7000 configuration based on the wide-open data paths of an Itanium 2-based system.

Extremely complex processing tasks, such as data mining and statistical analysis on large datasets, benefit from the high memory bandwidth of the Intel Itanium 2 processor. On the other hand, CPU-intensive applications that place less demand on memory bandwidth may provide better price/performance on multi-CPU servers based on the Intel Xeon processor MP. For example, data analysis using relatively smaller datasets may benefit from the extremely quick ramp-up time of the Xeon processor, without incurring a significant penalty in memory throughput. With both processors, the ES7000 offers a very wide data bus and high memory bandwidth.

As the analytical tests increased in size, performance was best on the Itanium 2 processor, which has the faster memory architecture.

Simultaneous Users/Jobs

The number of simultaneous users provides a good gauge of the optimum number of processors. As the number of simultaneous users rises, so should the number of CPUs in the system.

Testing on the Itanium 2-based systems reflected optimum price/performance at an average of 8–10 users per CPU, while the Intel Xeon processor MP-based platform performed better with 2–4 users per CPU. We reached this conclusion by examining the length of the processor queue and disk queue at various user loads. Ideally, each queue should have no more than one process at a time, indicating that the processor and disk I/O are keeping up with the tasks to be completed. Queues grow longer when the processor or disk are unable to keep up with the workload, indicating a drop in system efficiency. To recommend the ideal number of simultaneous users per CPU, we determined which workloads produced processor and disk queue lengths at or near 1.0.

Throughput of a multi-user workload increases as you raise the number of CPUs. At a load of eight simultaneous jobs per CPU on an Itanium 2-based system, the average processor queue and disk queue length remain at 1.0 tasks, indicating minimal competition for bandwidth. Tests also showed that even when the ES7000 is busy with multiple long jobs, short jobs can be submitted and completed quickly, with relatively little impact on overall performance.

Note that the recommended number of users per CPU is only an estimate. This figure is influenced by a number of factors, including whether the application is especially processor-intensive (keeping each CPU busier), as well as the other resources available on the system. For example, a system with insufficient memory is unlikely to benefit from more CPUs in the system.

Summary

Large ES7000 configurations provide outstanding performance, headroom and throughput for large SAS implementations, analysis of large datasets, multiple instances of SAS, large user bases, threaded applications such as general linear model analysis, and full solution stack implementations. The ES7000 family allows SAS users to start small and easily, economically grow. It also offers significant business value through its exceptional price/performance, simplicity and manageability. Configurations can be based on the Intel Itanium 2 processor for tasks with demanding memory bandwidth requirements and 8-10 users or jobs per CPU, or on the Intel Xeon processor MP for tasks that are more CPU than memory intensive and have 2-4 users or jobs per CPU. A single ES7000 can be configured with a mix of both processors, for outstanding flexibility and investment protection and maximum performance and throughput.

More Information

Unisys Center of Excellence for SAS

To serve SAS and Unisys customers across the world, Unisys has established a Center of Excellence at the SAS world headquarters in Cary, North Carolina. In the Center, a team of Unisys technical experts works with SAS research and development and technical consulting organization to provide dedicated technical and sales support for running SAS on the Unisys ES7000 family of servers. Expert teams from SAS and Unisys collaborate to optimize our joint solutions and also provide sizing, configuration, and ongoing technical support to our field sales organizations around the world.

Visit www.sas.com/partners/directory/intel to learn more about SAS, Intel and our continuing efforts to improve and quantify the performance of SAS solutions on Intel architecture.

SAS Customer Knowledge Practice from SAS and Intel

If you need support to install, tune or answer questions about SAS on Intel-based systems, Intel and SAS have built a practice to help. The SAS Customer Knowledge Practice is provided by Intel® Solution Services in conjunction with SAS. Staffed by Intel engineers trained in SAS by SAS, the practice provides a knowledge base and support for SAS installations and offers a single point of technical support, helping remove barriers to Intel architecture-based deployments of SAS. Please contact your local SAS or Intel representative for additional information.

Intel Solution Services

Intel Solution Services provides a worldwide consulting and services R&D organization, specializing in distributed solutions and enterprise e-Business infrastructure. Intel Solution Services can help your company accelerate the adoption and application of the latest technology advances on Intel architecture-based platforms. To learn more, visit www.intel.com/internetservices/intelsolutionservices

Or e-mail us at solution-services-questions@intel.com

Review the Full Paper


See Sizing and Performance Considerations for Intel® Architecture-Based SAS* Solutions, December, 2002. Visit www.intel.com/ids and click on Strategy & Solutions in the left-hand navigation bar.

Intel, the Intel logo, Intel Xeon and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
Unisys is a registered trademark of Unisys Corporation.

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

Copyright © 2003 Intel Corporation and Unisys Corporation. All rights reserved.

Printed in the U.S.A./0303/600/JM/KG/HOP • Order number: 252753-001

 Printed on recycled paper with soy-based inks.