

Paper 229-2011

## Performance Benchmark with SAS<sup>®</sup> 8.2, 9.1.3, and 9.2 on UNIX SPARC Architectures

Edmond Cheng, U.S. Bureau of Labor Statistics, Washington, DC, USA

### ABSTRACT

In the rapidly growing pace of information technology advancement, making the right decisions about software upgrades and hardware purchases becomes an important factor in maximizing one's SAS investment. With a wide range of SAS products and infrastructure platforms available, one challenge faced by SAS administrators is choosing the optimal bundle in software selection and hardware configuration.

The case study presented in this paper evaluates the performance level of Base SAS product versions, SAS<sup>®</sup> 8.2, SAS<sup>®</sup> 9.1.3, and SAS<sup>®</sup> 9.2, on selected UNIX SPARC server architectures. In a controlled environment, a series of performance simulations were executed to quantify input/output time, CPU time, and other process statistics among various SAS versions and UNIX configurations. The benchmark measurement provides first-hand tangible insights for administrators and managers in making informed SAS purchases and upgrade decisions to meet specific needs.

### INTRODUCTION

With the wide range of hardware and software products for enterprises, system administrators and managers face complex decisions for acquisition on information technology related resources. It is obviously the goal to maximize the return on investment for these acquisitions. To do this, SAS system administrators must choose the most suitable operating system platform and hardware configuration for the multiple offices they support. SAS system administrators must also test software upgrades to identify potential issues and make informative recommendations to managers.

My experiences with the team I work with involve running frequently scheduled production survey processing and statistics computation with SAS in UNIX platforms. SAS is the core of the production system. Besides using SAS for statistics computation and data intensive processing, it also is harnessed for data warehousing, web applications, and communication across the network which is compromised of UNIX and Windows server.

Over time our production processes have become much more complex and resource intensive. This highlights the need for wisely planning future system infrastructure upgrades and changes. Before making decisions on the next hardware and software upgrade, we needed to recognize the baseline performance level of different system architectures available from the vendors. Relevant to our operational needs, the key satisfactory criterion is an UNIX platform system providing high performance levels measured by processing time. Hardware configuration factors such as cost, storage, multi-threading, expandability, scalability, technical support, and other factors were considered. For the software components, as SAS version 9.2 became available for evaluation at the time of system assessment, we installed it on selected UNIX servers for benchmark testing.

### BENCHMARK SETUP

The best approach for generating realistic SAS performance benchmark levels is by simulating typical production operations in each UNIX hardware configurations. The same production setup is implemented with all three SAS versions, SAS 8.3, 9.1.3, and 9.2. The key performance measurement is the amount of execution time under each hardware and software combination.

### HARDWARE

The server specifications for hardware configurations are shown in Table 1.

Performance Benchmark with SAS 8.2, 9.1.3, and 9.2 on UNIX SPARC Architectures, continued

**Table 1. Hardware configuration for the UNIX servers**

Hardware Information	BASESUN1	BASESUN2	TESTSUN3	TESTSUN4	TESTSUN5
<b>Server Model</b>	Sun Microsystems Sun Fire V480R	Sun Microsystems Sun Fire V490	Sun Microsystems Sun Fire T5240	SPARC Enterprise M4000	SPARC Enterprise M5000
<b>Processor</b>	UltraSPARC III+	UltraSPARC IV+	UltraSPARC T2+	SPARC64 VII	SPARC64 VII
<b>Specs</b>	4 CPUs, 1.2 GHz, single-core	4 CPUs, 1.5 GHz, dual-core	2 CPUs, 1.2 GHz, 6-core	2 CPUs, 2.4 GHz, quad-core	2 CPUs, 2.4 GHz, quad-core
<b>Thread</b>	4	8	96	16	16
<b>Ram</b>	16GB	16GB	16GB	32GB	32GB

In the **control server group**, the BASESUN1 V480R and BASESUN2 V490 are the current UNIX servers in service. These two Sun Fire servers are marketed by Sun Microsystems as mid-range servers.

- The **V480R** server consists of 4 **UltraSPARC III+** single-core CPUs at 1.2GHz and implements the 64-bit SPARC V9 architecture.
- The **V490 UltraSPARC IV+** runs 4 dual-core CPUs at 1.5GHz. The architecture of the newer UltraSPARC IV is basically two UltraSPARC III CPUs manufactured into one CPU.

In the **test server group**, three advanced microprocessors were available for benchmarking.

- The TESTSUN3 **T5240** is designed with a very different microprocessor architecture. The server runs on 2 **UltraSPARC T2+** 6-core CPUs with a maximum of 96 concurrent threads. The multicore and multithreading architecture is advantageous for hosting server applications such as web server, network storage, and databases.
- Lastly, our team received two high-end servers TESTSUN4 and TESTSUN5 for additional benchmark testing. These are the SPARC Enterprise **M4000** and **M5000**. These systems run on the more advanced **SPARC64 VII+** architecture. The specifications include 2 quad-core CPUs running at 2.4 GHz with multithreading capability and double the RAM of the other. We have high expectation for the benchmark performance for these M-series machines.

## SOFTWARE

On the software side, both **SAS 8.2** and **SAS 9.1.3** are installed on the Sun Fire V480R, V490, and T5240 servers. With the availability of **SAS 9.2** for evaluation, the newer SPARC Enterprise M4000 and M5000 servers are tested with both SAS 8.2 and SAS 9.2.

**Table 2. SAS versions installed on the UNIX servers**

SAS Version	BASESUN1	BASESUN2	TESTSUN3	TESTSUN4	TESTSUN5
<b>SAS 8.2</b>	✓	✓	✓	✓	✓
<b>SAS 9.1.3</b>	✓	✓	✓		
<b>SAS 9.2</b>				✓	✓

The settings in the SAS configuration file are kept consistent across the multiple versions and servers. All servers have Solaris 10 operating system installed.

## TEST METHODOLOGY

The objective of the benchmark testing is to measure the performance level of a production simulation by recording execution time. Other performance factors such as CPU workload, system I/O, single/multiple threading, disks load, bottlenecking, and other considerations are observed manually where possible. The two versions of SAS are installed

Performance Benchmark with SAS 8.2, 9.1.3, and 9.2 on UNIX SPARC Architectures, continued

on each server in the control and test groups. Two tests are created from SAS jobs selected from the existing production system. The selected SAS processes are good representation of routine and production operations.

**Table 3. Benchmark Test Setup**

Benchmark Test Setup #1	
Item	Description
<b>Job</b>	Perform transposing and post-processing of survey responses
<b>Data attributes</b>	4.25GB (data) w/ 2.2GB (index), 3.95 million observations, 21 columns
<b>Process</b>	A DATA step intensive process to transpose the 3.95 million observation data set from vertical format to a horizontal format, while using the INPUT function to convert text fields to numeric fields
Benchmark Test Setup #2 - #5	
Item	Description
<b>Job</b>	Construct time series database and calculate summary statistics
<b>Data attributes</b>	0.85GB, 8.5million observations, 41 columns
<b>Process</b>	Use DATA step and PROC SQL to merge 13 consecutive quarterly data sets by key identifiers, apply logic rules, perform computation, and create multilevel summary statistics

**Benchmark Test Setup #1** provides a first cut baseline performance comparison.

**Benchmark Test Setup #2 - #5** demands more processing time working with multiple databases and computations are at the observation level. The procedure involves merging (13) 0.85GB data sets to construct one large time series database. At the end, the process calculates summary level statistics for multiple levels.

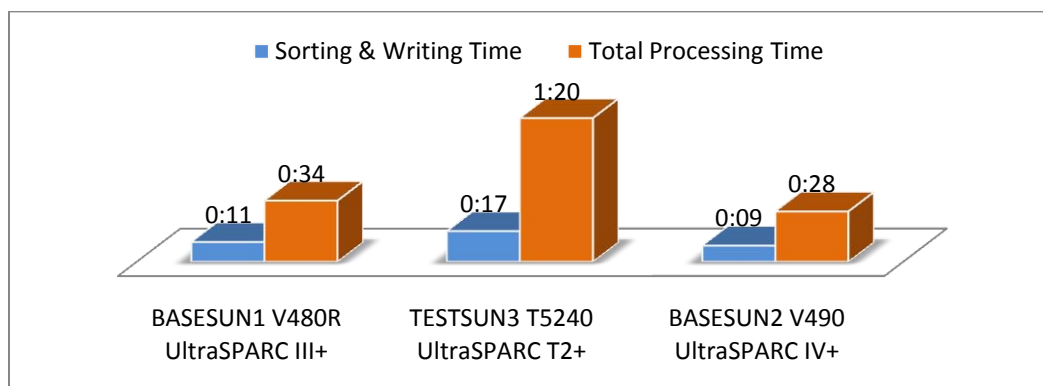
## BENCHMARK RESULTS

For measuring performance levels, the processing time for each SAS test job is recorded from SAS. Each test program includes codes for tracking the starting/ending time and total processing time for job completion. To ensure consistency among observed results, two test iterations are performed and the average time is recorded as final results. Testing was conducted under a controlled environment with no other activities occurring in the servers.

### BENCHMARK #1 HANDLING OF LARGE SIZE DATA SET

The first benchmark tests compare the performance between the Sun Fire V480R UltraSPARC III+, V490 UltraSPARC IV+, and T5240 UltraSPARC T2+ servers. An identical test job running on SAS version 8.2 is used for this benchmarking. The test setup processes an indexed SAS data set containing 3.95 millions observations and 21 variables. The procedures consist series of ARRAY () statements, IF...THEN...ELSE logics, and INPUT () functions to construct a new data table.

**Figure 1. Performance Benchmark with SAS 8.2 (in hh:mm)**



As Figure 1 shown, the server running UltraSPARC T2+ requires the highest amount of processing time to complete the SAS job. The processing time (1:20) doubles the time recorded for servers running UltraSPARC III+ (0:34) and UltraSPARC IV+ (0:28). The UltraSPARC IV+ records the fastest processing time as expected with noticeable performance advantage over the older UltraSPARC III+ microprocessor. Our initial assessment on the UltraSPARC

Performance Benchmark with SAS 8.2, 9.1.3, and 9.2 on UNIX SPARC Architectures, continued

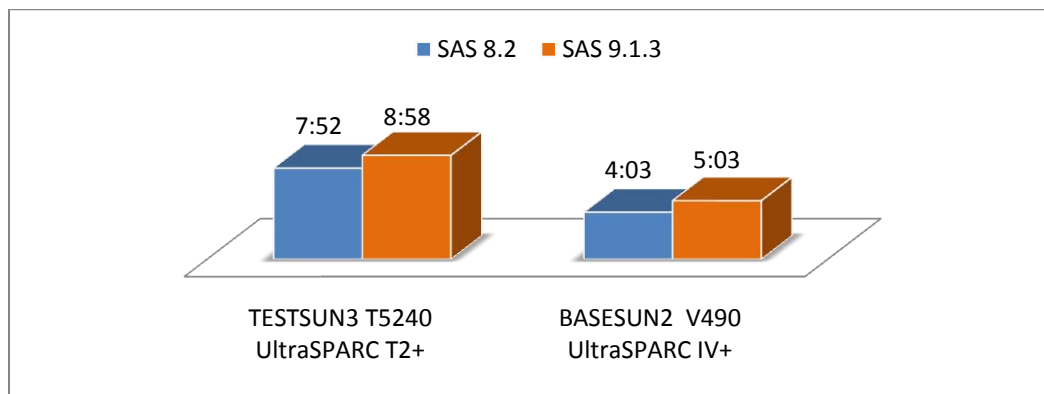
T2+ expects the architecture design to be less advantageous for intensive CPU operations. It is our understanding the multicore and multithreading capability exceeds in performing tasks with multiple concurrent threads, such as serving web applications and handling database queries for large pool of users.

## BENCHMARK #2 LOGIC AND ARITHMETIC COMPUTATION

The test methodology for Benchmark #2 to #4 uses the identical test setup. The test servers are the V490 UltraSPARC IV+ and T5240 UltraSPARC T2+ servers with SAS 8.2 and 9.1.3 installed. The selected SAS test job merges 13 non-indexed SAS data sets with into one single data set by key IDs and linkages rules. Sequentially, the program processes logics statements, user defined formulas, and mathematical computations to produce multilevel summary statistics.

The processing time requires to complete each test job is recorded. Two iterations were executed and the average time is recorded.

**Figure 2. Performance Benchmark with SAS 8.2 and SAS 9.1.3 (in hh:mm)**



Similar to the results recorded from Benchmark #1, the server running UltraSPARC IV+ completes the single-threaded SAS twice as fast against the UltraSPARC T2+. With SAS 9.1.3, the UltraSPARC IV+ completes in 5:03 (hh:mm), while the UltraSPARC T2+ takes 8:58. In an actual production environment, the matter of differences would determine whether completing production operations and delivering outputs to customers on time; or causing schedule delay.

The results from Figure 2 show an unexpected surprise. The processing times recorded on version 9.1.3 were 15% - 20% slower than version 9.2. Given the information we know about the UNIX servers and SAS configuration setups, we have no obvious explanations to interpret the longer processing time observed under version 9.1.3. Opportunity was not available to conduct additional benchmarking to validate whether the results are feasible. This question will be researched in the future.

## BENCHMARK #3 CPU TIME ON SELECTED PROCESS

Continuing with previous results, Benchmark #3 takes closer details at the recorded CPU timer for selected DATA steps and PROC procedures parsed from the system generated SAS logs. The objective is to determine evidence (if any) as whether the two microprocessor architectures might be more efficient at performing certain DATA steps or PROC procedures.

**Figure 3. CPU Time on Selected Process (in mm:ss)**

Process	TESTSUN3 UltraSPARC-T2+ T5240		BASESUN2 UltraSPARC-IV+ v490	
	SAS 8.2	SAS 9.1	SAS 8.2	SAS 9.1
	Read-in LDB input	2:00	4:06	0:46
Array processing	8:08	8:26	2:41	2:54
Sorting	6:25	11:51	3:02	5:55
Merging	4:42	4:56	2:12	2:56
Computation	6:00	6:16	2:13	2:22

Performance Benchmark with SAS 8.2, 9.1.3, and 9.2 on UNIX SPARC Architectures, continued

As shown in Figure 3, the CPU time at the individual process level are consistent with results collected at the job level from Benchmark #2. Under both SAS versions, the identical individual processes such as reading in data input, sorting, and merging procedures takes about two times longer running on UltraSPARC T2+ than UltraSPARC IV+.

For array processing and statistics computation, the UltraSPARC T2+ takes an average of three times longer than the UltraSPARC IV+. The UltraSPARC IV+ is expected to be more efficient at performing computational intensive activities than the UltraSPARC T2+. Perhaps, the test results would be contrasting under multiprocessing scenario.

#### BENCHMARK #4 PARALLEL MULTIPLE SAS SESSIONS

The test objective for Benchmark #4 examines the performance level running on parallel SAS sessions. Only SAS version 8.2 is used for testing. The multithreading capability is tested for the UltraSPARC IV+ and UltraSPARC T2+ processors. To simulate concurrent multithreading scenarios, we simultaneously executed four identical single-threaded test job used in Benchmark #2 in separate SAS sessions. The SAS test programs themselves do not utilize multiprocessing programming techniques or communication services for parallel sessions.

Figure 4. Multiple SAS Sessions Processing with SAS 8.2 (in hh:mm)

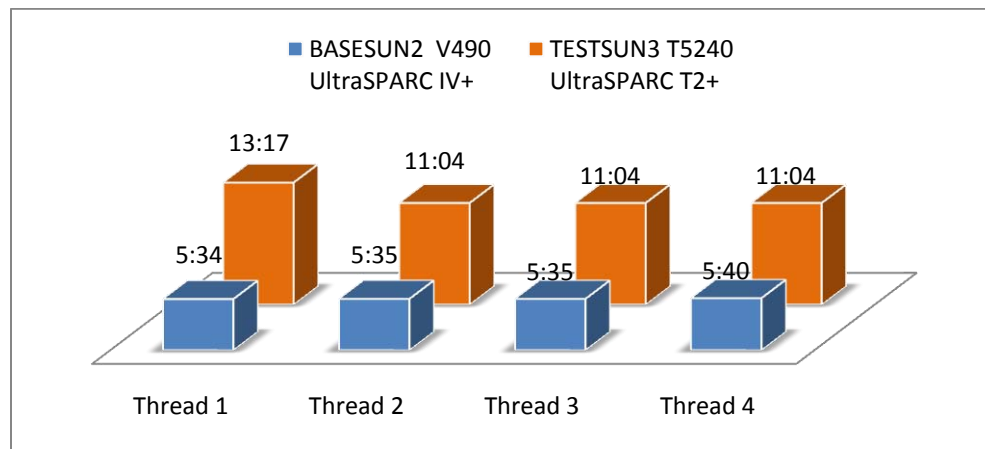


Figure 5. PRSTAT Statistics on Multithreading (4) Parallel SAS Sessions (hh:mm)

```

---BASESUN2---
  PID USERNAME  SIZE  RSS STATE PRI NICE   TIME    CPU PROCESS/NLWP
 16849 root      5208K 3752K cpu2   0   0 27:07:58 14% prstat/1
 10640 chenge    14M   10M cpu3   0   0 0:26:58 12% sas/1
 10625 chenge    14M 9968K cpu0   0   0 0:23:49 9.8% sas/1
 10636 chenge    14M   10M cpu19  0   0 0:24:37 7.9% sas/1
 10632 chenge    14M   10M cpu17  0   0 0:24:34 7.9% sas/1

---TESTSUN3---
  PID USERNAME  SIZE  RSS STATE PRI NICE   TIME    CPU PROCESS/NLWP
 21788 chenge    15M   12M cpu48  0   0 0:24:12 1.0% sas/1
 21780 chenge    63M   61M cpu96  20  0 0:28:11 0.1% sas/1
 21792 chenge    15M   12M cpu24  10  0 0:24:24 0.1% sas/1
 21784 chenge    15M   12M cpu16  10  0 0:24:26 0.1% sas/1

```

The multithreading results provide few insights. Both servers completes all four parallel SAS sessions without conflicting resources. The UltraSPARC IV+ takes an average time of 5:36 to complete each job, which is about 37% more than running a single identical single threaded job (4:03). The UltraSPARC T2+ took about an average time of 12:22, almost 55% longer than a single job (7:52). Both servers show parallel sessions contributes significant 'total time' saving versus running in singles. Take results from UltraSPARC IV+ U as an example, running individual SAS sessions to complete four jobs would require 16:00 hours versus 5:35 hours for running four jobs parallel with each other.

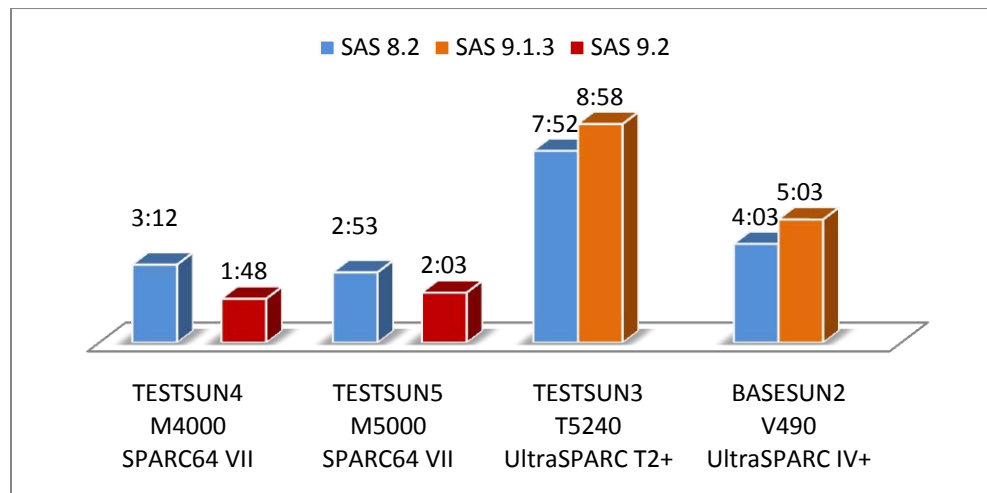
The multicore and multithreading design in UltraSPARC T2+ suggests unrealized potential in handling multiple concurrent sessions as the number of processing threads increases. Hypothetically, if UltraSPARC T2+ has the resources to process its maximum limit of 64 concurrent threads at once, the performance could exceed the UltraSPARC IV+ design, which has a limit of 8 concurrent threads. This will be left for further research in the future.

Performance Benchmark with SAS 8.2, 9.1.3, and 9.2 on UNIX SPARC Architectures, continued

## BENCHMARK #5 SAS VERSION 9.2

Benchmark #5 measures the performance levels on two SPARC Enterprises SPARC64 VII servers with both SAS 8.2 and 9.2 installed. These two models are M4000 and M5000 series. The single threaded SAS test job is used for this final benchmarking.

Figure 6. Performance Benchmark with SAS 8.2, 9.1.3, and 9.2 (in hh:mm)



The benchmark results as shown in Figure 6 concur with our initial assessment on the high performance expectation to come from these two high-end servers. The SPARC64 VII series completes the single threaded test job with an impressive 3:00 under SAS 8.2 and 2:00 under SAS 9.2. Comparing to the earlier UltraSPARC IV+ benchmark, the new series has an advantage of being 25% faster. Comparing to the UltraSPARC T2+, the new series has an advantage of more than 125% faster. The more advanced SPARC64 VII processor architecture shows itself as the clear winner for handling intensive CPU bound processing.

Another encouragement from the Benchmark #5 results is the significant reduction in processing time recorded under SAS version 9.2. The difference is approximately 1:00 against SAS version 8.2, which equates to 33% execution time saving. The result provides much comforting results from the unexpected increase in processing time recorded under SAS 9.1.3 over 8.2 during Benchmark #2.

## FUTURE WORKS

In the future, it is desirable to conduct additional benchmarking on different hardware platforms and SAS configuration settings. The list of future works includes:

- Comparing single vs. multiple threading performance advantage thru UNIX job queue commands, multiprocessing programming techniques, or SAS CONNECT services
- Configuring system I/O such as different disks array and system RAM
- Tuning SAS configuration settings such as adjusting SAS MEMSIZE allocation size. However, early attempts in modifying the SAS MEMSIZE allocation size between 128MB to 1GB shows no noticeable performance improvement under our environment
- Quantifying benchmark statistics by using UNIX commands to capture system statistics or enabling SAS FULLSTIMER system option to parse performance statistics.

## SUMMARY

This paper presents series of SAS performance benchmark results completed on different combinations of SPARC UNIX servers and SAS versions. The results shows the SPARC64 VII processor architecture design offers the fastest performance than the earlier UltraSPARC III+ / IV+ design, approximately twice as fast. The UltraSPARC T2+ processor bears the slowest processing time. Yet, the multicore and multithreading architecture design has the potential to run 96 concurrent processes at the limit. On SAS versions performance level, SAS version 9.2 exceeds 8.2 and 9.1.3 by as much as 33% during conducted benchmark testing. Together, the benchmark results provide other SAS system administrators and managers in evaluating UNIX platform hardware configuration decisions and planning for SAS versions upgrades.

Performance Benchmark with SAS 8.2, 9.1.3, and 9.2 on UNIX SPARC Architectures, continued

## REFERENCES

Kearns, William and Keefer, Tom. (2003). "Performance Tuning & Sizing Guide for SAS Users and Sun System Administrators". Proceedings of the 28<sup>th</sup> SAS Users Group International. Available at <http://www.sas.com/partners/directory/sun/sugi28.pdf>.

Hetherington, R., Hendel, A., Spracklen, L., and Vajapey, S. (2007) "UltraSPARC T2 Processor System On a Chip." *Sun Wiki*. Available at [http://wikis.sun.com/download/attachments/31400118/N2\\_Announce\\_Breakout\\_final.pdf](http://wikis.sun.com/download/attachments/31400118/N2_Announce_Breakout_final.pdf).

## ACKNOWLEDGMENTS

Thanks Steven Holmes, G&B Solutions, Inc., for providing the system infrastructure and coordination to make the benchmarking tasks possible.

## DISCLAIMER

The contents of this paper are the work of the author and do not necessarily represent the opinions, recommendations, or practices of U.S. Bureau of Labor Statistics. The benchmark results of this paper are based on the hardware specifications, software configurations, and test procedures in our test environment. The presented information do not represent actual performance of the mentioned UNIX platforms, SPARC processors, and SAS software.

## CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author at:

Edmond Cheng  
U.S. Bureau of Labor Statistics  
2 Massachusetts Avenue, NE  
Washington, DC 20212-0001  
(202) 691-5458  
[cheng.edmond@bls.gov](mailto:cheng.edmond@bls.gov)

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.

Other brand and product names are trademarks of their respective companies.