An Oracle White Paper
December 2013

SAS Application Performance on the Oracle M5-32 SPARC Server

Introduction

Today, information technology solutions play a critical role in helping organizations achieve their primary mission. In many cases, business success is dependent on the continuous availability and rapid response time for IT services such as Oracle Database and SAS application solutions. All of these come together in providing a data center infrastructure that can support mission-critical applications like SAS application solutions in delivering results rapidly to key decision makers in an organization. Meeting these challenges requires a server infrastructure that provides features such as mainframe-class reliability, performance-based throughput and virtualization-enabled agility.

To demonstrate that the SAS application solution running on Oracle's SPARC M5-32 server can meet these challenges, Oracle and SAS performed a series of performance related tests. This joint testing initiative of the Oracle and SAS team involved deploying SAS 9.4 application solution software on the SPARC M5-32 server. This paper highlights the specific outcomes and describes how Oracle's M5-Series server technology made a difference within the testing constructs that are detailed in Figure 1.

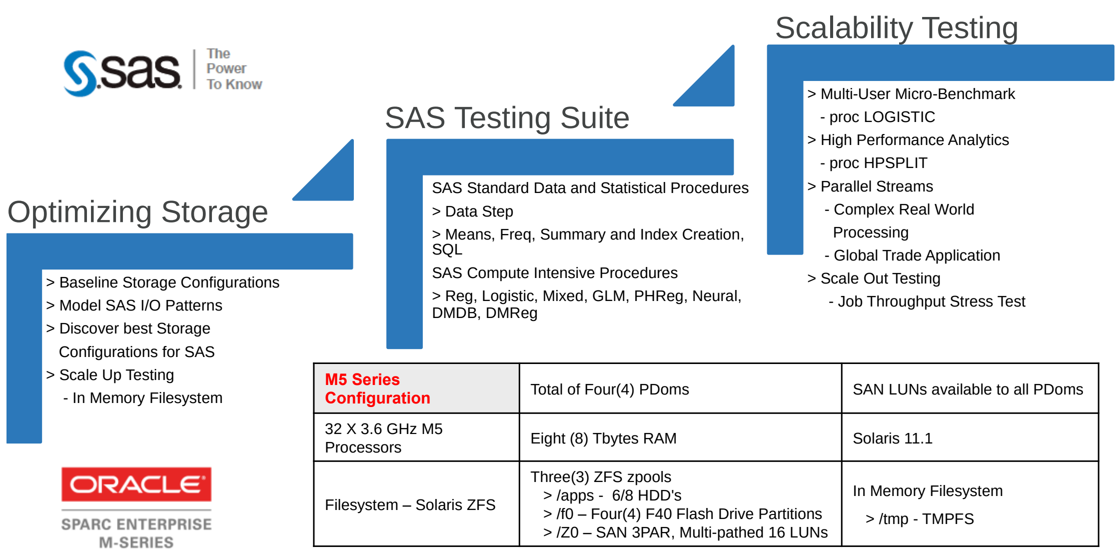

M5 Series Configuration	Total of Four(4) PDoms	SAN LUNs available to all PDoms
32 X 3.6 GHz M5 Processors	Eight (8) Tbytes RAM	Solaris 11.1
Filesystem – Solaris ZFS	Three(3) ZFS zpools > /apps - 6/8 HDD's > /f0 – Four(4) F40 Flash Drive Partitions > /Z0 – SAN 3PAR, Multi-pathed 16 LUNs	In Memory Filesystem > /tmp - TMPFS

Figure 1. SPARC M5-32 server and SAS application solution beta test information.

SAS Application Solutions Exploit Oracle's SPARC Technology

SAS application solutions are often similar to what was exercised in both of the SAS test suites and engineered to be well balanced. When examining the profile of the SAS testing suites, the numerous data steps and SAS procedures (a.k.a. PROCs) are not only data intensive but also compute and memory intensive. This means that SAS applications can respond well to an accelerated computing infrastructure with hardware compute and memory efficiencies. Oracle's SPARC servers in general, and especially the SPARC M-Series servers, allow SAS applications to exploit the multithreading chip

Workload Content	Workload Description
Standard SAS Data/Procedures	**SAS Data Step and Procedures**
Data Step	Ingest Data into SAS Datasets and Oracle Databases
Proc Sort	Transform ingested data into ordered datasets witin SAS Datasets or Oracle Databases
Proc Datasets	Create SAS Indexes to improve storage performance
Proc Reg	SAS Linear Regression
Proc Logistic	SAS Logistic Regression
Proc Mixed	Creates mixed linear models
Proc GLM	Creates general linear models using a least squares fit
Proc PHREG	Performs regression analysis of survival data
Proc Neural	Neural Network Procedure often used in Data Mining
Proc DMDB	Creates Data Mining Databases (aka Data Marts)
Proc DMREG	Fits data to both linear and logistic regression models
Proc SQL	Used to retrieve, update, and report on data from SAS data sets and Oracle Databases
Proc Summary	Summarizes statistically all data in a SAS Dataset or Oracle Database
Proc Freq	Creates Frequency distributions of selected data
Proc Means	Provides basic statistiscs of selected data in a SAS Dataset or Oracle Database
SAS High Performance Analytics	**New SAS 9.4 Multi-Threaded High Performance Procedure(s)**
Note: All SAS workloads work with data in SAS Datasets and/or Oracle Databases as applicable	

Figure 2. Typical SAS Workload Types used during SAS Testing.

design of each of the SPARC processors. Since the SPARC M5-32 server scales up to 32 SPARC M5 processors, each with six cores, up to 192 cores are supported. Each processor in the SPARC M5-32 server contains six SPARC S3 3.6 Ghz cores that support chip-level multithreaded computing (CMT), supporting up to 48 strands of executable code simultaneously. The Oracle Solaris operating system exploits the SPARC M5-32 server's computing ability by distributing SAS application solution executable code streams across the processor, core (CPU), and CMT architecture to dramatically improve application computing efficiencies, resulting in accelerated execution and faster elapsed times.

In addition, traditional processor designs have singularly focused on improving application throughput by increasing the speed of execution. With the SPARC M-Series server(s), Oracle has also focused on improving memory speeds and the amount of memory available to applications like SAS. While the SPARC M5 server used at SAS in testing had 8 terabytes of memory, up to 32 terabytes of memory is supported. With this massive amount of memory and dramatically improved memory speed, the SPARC M5-32 processor design mitigates performance issues associated with excessive memory demand and/or memory latency, thereby improving application throughput performance. As shown in these elapsed time results, by redesigning the cores within each processor, designing a new floating-point pipeline, improving memory speeds, and further increasing the internal computing network bandwidth, Oracle has dramatically improved the execution times of SAS application solutions when executed on the SPARC M5-32 server computing platform.

To summarize, SAS application solution workloads can exercise the ability of Oracle Solaris to multi-process workloads across up to 32 processors, 192 cores, and associated threads present in the computing architecture—all while potentially referencing memory across processor and/or domain boundaries. These SAS job elapsed time results reveal the beneficial effects of the SPARC M-Series server computing speeds, context switching between threads, and memory speeds on job execution times and job turnaround revealing the advantages of deploying SAS application solutions on the Oracle M5-32 and getting key business results rapidly.

Managing SAS Workloads on the SPARC M-Series Server

Computing requirements change over time so customers need to come to conclusions as to how their particular SAS workloads would behave with respect to anticipated growth or consolidation on a server such as the SPARC M5-32 server. As such, the data center manager can deploy multiple workloads as

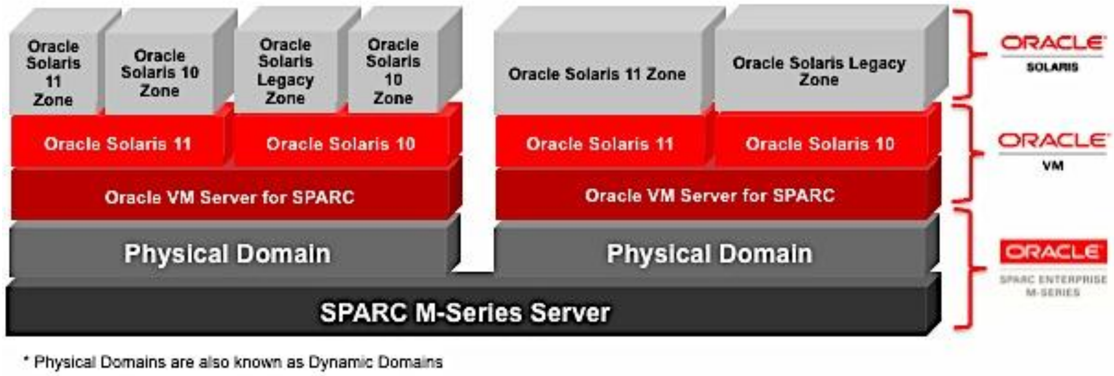

Figure 3. Oracle virtualization deployment on the SPARC M5-32 server.

well as accommodate growing and consolidated workloads easily while preserving mission-critical performance on the SPARC M5-32 server by utilizing its advanced virtualization features.

Advanced Virtualization and the SPARC M5-32 Server

The SPARC M5-32 server has advanced virtualization capabilities not present in other enterprise-class servers available today. These features empower the data center manager to easily manage and consolidate workloads within the IT infrastructure while simultaneously providing high availability. The advanced virtualization features of Oracle's SPARC M5-32 server can easily accommodate consolidated SAS application solution workloads as well as efficiently utilize the powerful computing and memory architecture. Built-in virtualization technologies include Dynamic Domains, a feature of Oracle's SPARC Enterprise M-Series servers; Oracle VM Server for SPARC (logical domains or LDoms); and Oracle Solaris Zones. With this rich set of virtualization technologies, the data center manager can easily configure a server of virtually any size and configuration up to the maximum size of the SPARC M5-32 server to support SAS compute, memory, and I/O demand, making it the ideal platform for consolidating SAS workloads.

- **Dynamic Domains (physical domains or PDoms).** The SPARC M-32 server scales up to 32 processors and up to 32 terabytes of memory. Dynamic Domains (PDoms) on the SPARC M5-32 server provide IT organizations with the ability to divide this single large system into multiple, fault-isolated servers, each running independent instances of the Oracle Solaris operating system. With proper configuration, hardware or software faults in one domain remain isolated and unable to impact the operation of other domains. Each domain within a single server platform can deploy a different version of Oracle Solaris, making this technology extremely useful for pre-production testing of new or modified applications.

  While the maximum number of PDoms that can be deployed on the SPARC M5-32 server is four, the Dynamic Domains feature empowers data center managers to deploy computing resources across these domains as required using either bounded PDoms or regular PDoms. Regular PDoms can grow from 4 to 32 SPARC M5 processors. A bounded PDom is restricted to four or eight SPARC M5 processors. Note that the advantage of a regular PDom is that it can grow to include all available memory and processors, while the advantage of a bounded PDom is that it has a performance advantage due to lower intra-processor and memory latency. In addition, with each PDom, servers can be easily configured more granularly by deploying LDoms and/or Oracle Solaris Zones, giving data center managers the maximum agility in being able to deploy servers with the key compute, memory, and I/O attributes required for each workload. For more information on Dynamic Domains, see the technical white paper entitled, "Oracle's SPARC M5-32 Server Architecture" on the Oracle Technology Network. See the section titled, "For More Information" in this document to learn how to access this aforementioned document.

- **Oracle VM Server for SPARC**. Oracle VM Server for SPARC is a built-in firmware-based Hypervisor that supports multiple virtual machines, called logical domains (LDoms), on a single system. In the case of the SPARC M5-32 server, multiple LDoms also can be defined within a single dynamic domain. The Hypervisor allocates subsets of system resources (memory, I/O, networking, and CPU) to each logical domain, isolating each Oracle Solaris instance and SAS workload to a virtual machine with dedicated resources. Built-in virtual machine snapshot and cloning capabilities help to speed virtual machine configuration and migration, enabling faster provisioning when growth in SAS application solutions occurs or when consolidating SAS application and server workloads.
- **Oracle Solaris Zones**. Using flexible, software-defined boundaries, Oracle Solaris Zones (previously known as Oracle Solaris Containers), a lightweight virtualization technology, creates multiple private execution environments within a single Oracle Solaris instance. SAS applications running within zones are completely isolated, preventing processes in one zone from affecting processes running in another. Oracle Solaris Zones supports fault isolation, features extremely fast boot times, and can be configured to instantly restart SAS applications. Because zones make it easy to prioritize applications and adjust resource allocations, they are ideal for consolidated SAS application solution workloads.

Configuration of the SPARC M5-32 Server for SAS Testing

The deployed configuration of the SAS application solution testing on the SPARC M5-32 server is illustrated in Figure 4. There were four PDoms, each running Oracle Solaris 11 in a fault-isolated environment. The application solution workload layer varied depending on the nature of the test.

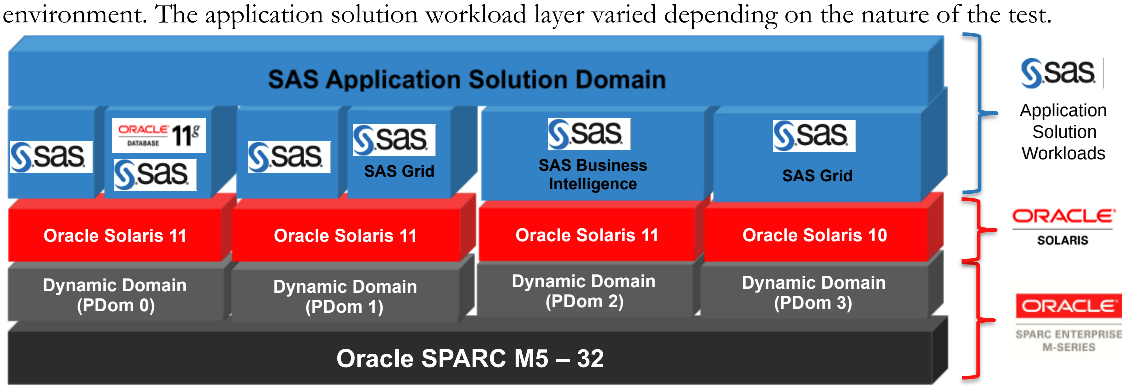

Figure 4. Potential configuration of the SPARC M5-32 server's advanced virtualization during SAS testing.

The workloads annotated in the diagram are notional in order to illustrate a possible workload configuration for SAS on the SPARC M5-32 server.

Compute and Memory Intensive SAS Application Workloads

With the SPARC M5-32 server, large application workload configurations are possible and can make a dramatic difference for mission-critical applications. Often, mission-critical applications that require rapid turnaround time and in some cases require real time response, execute concurrently with large application workloads. Rapid turnaround is often required in order to sustain business response and preserve the business bottom line. The SPARC M5-32 server can sustain large numbers of simultaneous workloads in memory while sustaining an in-memory file system to meet these business requirements.

SAS can exploit large memory systems such as the SPARC M5-32 server in numerous ways as illustrated in Figure 5. In this paper what will be described is how SAS High Performance Analytics (HPA) procedures and allocation of filesystems in-memory with the SPARC M5-32 server can accelerate SAS application solutions. This is an option for the SAS user because of the huge memory footprint as well as the massive processor and core computing resources available on the SPARC M5-32. SAS HPA application solutions can exploit memory and computing resources by simultaneously distributing processing and where possible, parallelizing SAS workloads and taking advantage of data resident in memory

- SAS 9.4 LASR Server
- SAS High Performance Analytic (HPA) Procedures
- Allocation of performance sensitive filesystems to In-Memory (TMPFS)
- SAS In-Line Memory Settings such as MEMSIZE/SORTSIZE/UTILSIZE
- SASFILE primitive for "pinning" datasets In-Memory

Figure 5. Examples of SAS how can be deployed use to exploit large memory systems.

Complex SAS Global Trade Application Solution Testing

During testing, a complex resource-intensive large global trade SAS application solution workload was deployed on the SPARC M5-32 server. Multiple parallel streams of SAS workloads were simultaneously executed utilizing data stored on an in-memory file system rather than on traditional storage. Each SAS job workload stream utilized an 80 GB data set that was streamed to SAS working storage configured in the in-memory file system (i.e., /saswork).

As illustrated in Figure 6, when the workloads were deployed in a single PDom, average SAS step times did increase, but only marginally as multiple instances of the SAS global trade application were executed on the SPARC M5-32 server. The results, although compelling, became even more significant when all four of the available PDoms on the SPARC M5-32 server were engaged with this identical workload in multiple streams.

Average SAS step times were collected, and as can be seen, there were no significant differences between average SAS step times on a single PDom across the job workload mix compared to the identical job mix streamed across multiple PDoms. From a results perspective, this translates into roughly 256 times the work being accomplished in less than twice the amount of time of a single execution of a SAS global trade application.

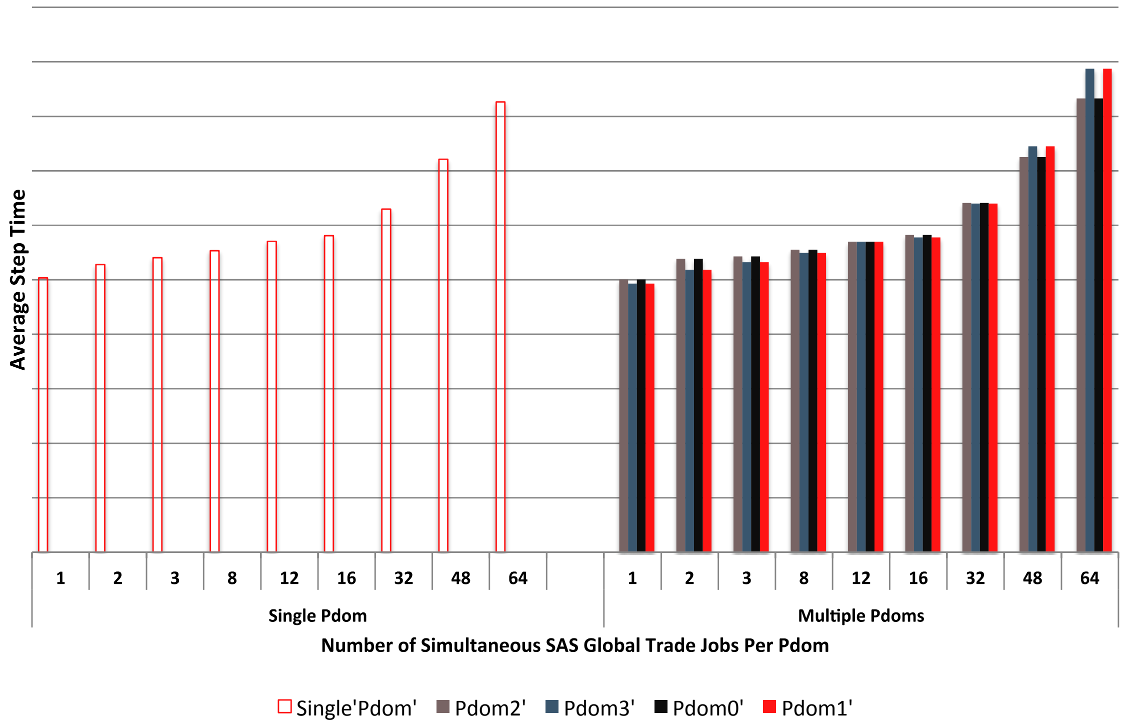

Figure 6. I/O and memory intensive complex SAS Global Trade Application performance

All this comes together on the SPARC M5-32 server in proving it's ability to consolidate and process large workloads utilizing the inherent large compute resources and in-memory storage efficiencies. This effectively produces quick turnaround of large data problems in multiple streams often used in SAS application solution architectures.

SAS High-Performance Analytics and the SPARC M5-32 Server

SAS HPA high-performance analytics (HPA) procedures perform statistical modeling and model selection by exploiting computing resources, whether they are in a single machine (a.k.a. SMP) or in a distributed computing environment. An important feature of all HPA procedures is that they are engineered to solve analytic tasks in distributed or parallel fashion by using concurrently scheduled threads that exploit the multiple processors that are ubiquitous in current-generation computers such as the SPARC M5-32 server.

In addition to distributed and to the extent possible, parallel processing, SAS HPA procedures with the SPARC M5-32 can run with data in memory which can further improve step times and job elapsed times such that they can approach, depending on the application, real time processing.

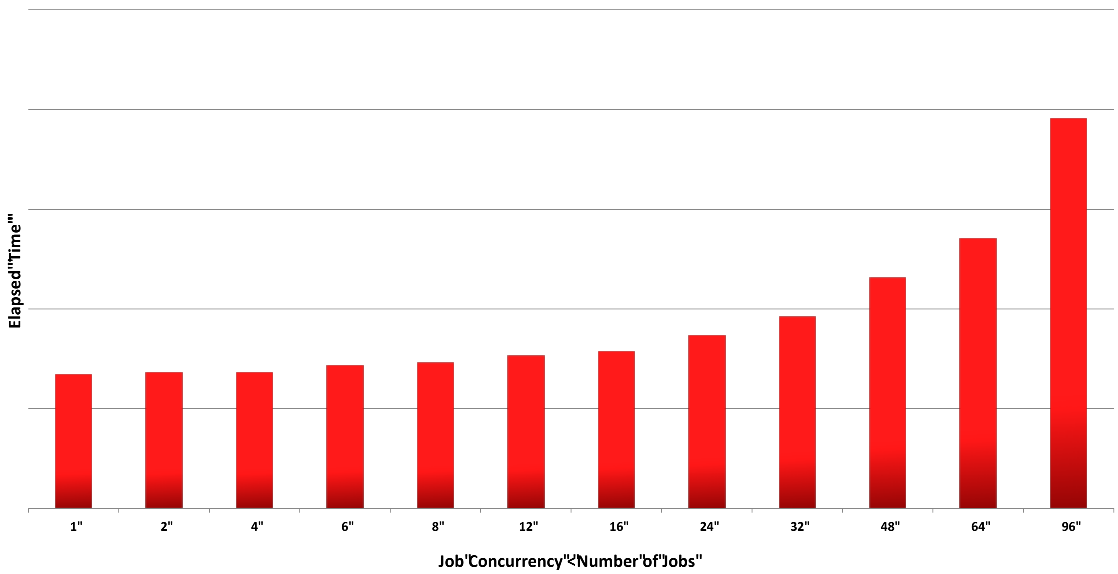

Figure 7. Testing Results of oversubscribing a single PDom by 2X with SAS HPA simultaneous job execution.

The testing approach was to execute a single job performing this type of workload and subsequently add identical jobs while observing overall job elapsed time to completion with the added job mix factor of oversubscribing by 2X a single PDom containing 48 cores with up to 96 simultaneous SAS High Performance Analytics jobs. As shown in Figure 7, a single SAS HPA procedure was executed on a single PDom and subsequently; additional simultaneous jobs were submitted within a single PDom to determine the effects on job elapsed time(s). As the number of simultaneous jobs (a.k.a. threads) increased, the pronounced effects of the multiple thread efficiencies of SAS HPA coupled with CMT on the SPARC M5 become apparent. A dramatic key indicator in the outcome is that when 96jobs are executed simultaneously, about 96 times the work can be accomplished—all things being equal (identical jobs)—in only three times the amount of time of a single identical job.

Conventional SAS Workloads and Job Performance

SAS workloads were tested on the SPARC M5-32 server in a variety of different ways. The intent was to show improvement in job elapsed times as more computing and memory resources are dedicated in the form of PDoms. While large scale numerically- and memory-intensive SAS workloads were tested on the SPARC M5-32, more conventional workloads were also executed to show job concurrency and scalability results within a single PDom, as well as how workloads scaled out as the configuration of available computing resources grew up to four PDoms.

The testing approach was to execute a single job performing this type of workload and subsequently add identical jobs while observing overall job elapsed time to completion. The results of the job concurrency test are illustrated in Figure 8. The results are revealing in that within a single PDom, a single job tabulated an elapsed time of about 10 minutes, while in the same PDom, 48 identical jobs executing simultaneously completed in an elapsed time of about 19 minutes. This is tantamount to saying that over 48 times the work and associated results can be achieved on the SPARC M5-32 server in about 19 minutes, which is slightly less than twice the time that a single job would execute.

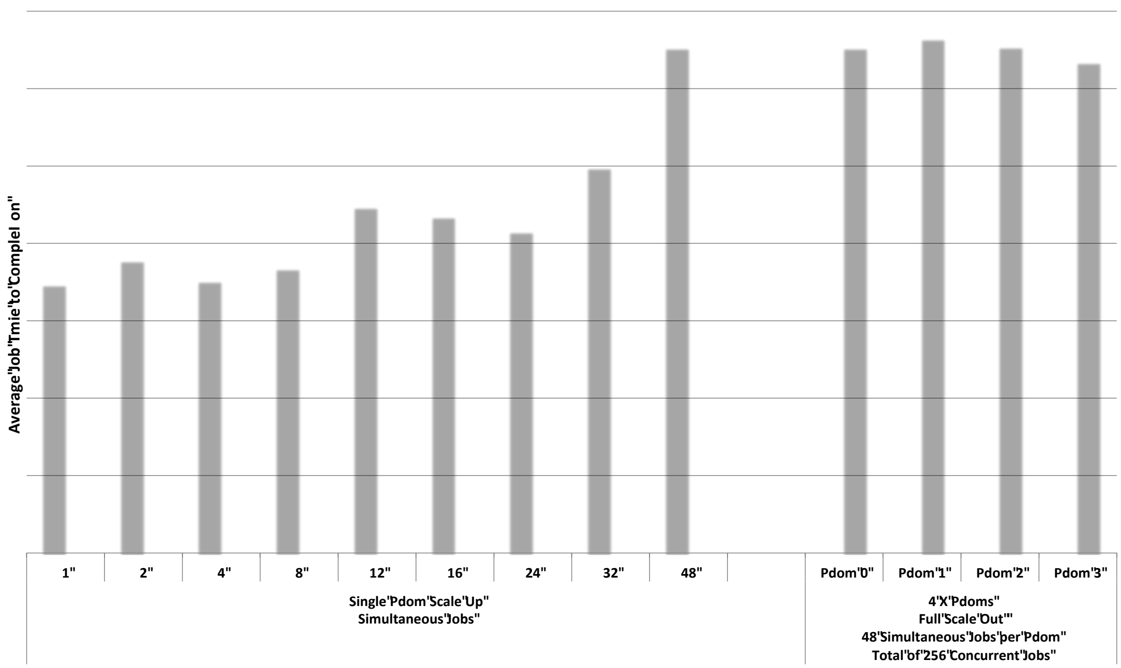

Figure 8. Conventional SAS Workload job performance.

The results, although compelling, become even more significant if all four of the available PDoms on the SPARC M5-32 server are engaged with this identical workload. They show similar results with more than 192 times the work completed in about twice the elapsed time a single job would take to execute. This illustrates the computing efficiencies of the SPARC M5-32 server as it engaged computing- and memory-intensive SAS workloads running in the Solaris 11 environment and was able to distribute workloads across the computing architecture at the processor, core, and CMT level. The server thereby improved job elapsed times and achieved more work by orders of magnitude.

SAS I/O Performance on the SPARC M5-32 Server

As discussed earlier, the SPARC M5-32 server features a balanced, highly scalable SMP design that utilizes the latest generation of SPARC M5 processors interconnected to memory and I/O by a high-speed, low-latency system interconnect that delivers exceptional throughput for applications. To prove this, SAS application solutions were tested in diverse PDom and storage configurations in order to show not only the advanced I/O performance of the SPARC M5-32 server but also the advantages in configuring storage to maximize I/O performance.

Testing was accomplished using high-end Storage Area Network (SAN) storage to model performance heuristics with different numbers of LUNs and multi-pathing configurations. The SAN configuration included a combination of four or eight LUN sets utilizing multiple paths with two or four port HBA connectivity through either two or four controllers. All testing was accomplished using the ZFS file system and throughout the testing scenario(s) scalability increased as expected. As the amount of workloads increased and the number of spindles included within each of the LUNs remained constant, the limits of the LUN config, which matched the expected throughput numbers for multi-stream, sequential, write access, became apparent. It should be noted that this represents the more typical worst case I/O scenario for multi-user SAS application solution scenarios.

During the testing, three simple best practices for all Soalris 11.1 SAN configuration were found. These include:

- Use 3 or more LUNs regardless of how many spindles back the LUNs – this is due to a SCSI spec limitation which limit the per LUN queuing heuristics
- Deploy the following ZFS file system tunings in /etc/system
    - set zfs:zfs_vdev_max_pending=35 (default is 10)
    - set zfs:zfs_prefetch_disable=1

In addition to I/O performance modeling, performance heuristics were examined when leveraging the large memory capacity of the Oracle M5-32 SPARC server.

SAS and Leveraging M5-32 Large Memory for SASWORK

Historically, there have been a number of pros and cons when considering whether to assign SASWORK to /tmp or TMPFS – a memory based file system. While performance can be very good, the generic recommendation was to use a ZFS based file system for SASWORK. ZFS achieves excellent performance results because of inherent caching properties. Deploying ZFS with the large memory capacity and performance of the Oracle M5-32 SPARC server for SASWORK would proximate the alternative to using TMPFS.

However, with the large memory configurations available on the Oracle M5-32 SPARC server, using TMPFS for SASWORK becomes a much more viable option as long as the lack of persistence of TMPFS between reboots is an acceptable condition. While testing SAS application solutions on the SPARC M5-32 server the objective was to analyze using SASWORK on diverse storage area network (SAN) configurations while also deploying SASWORK on TMPFS. The objective being to leverage the SPARC M5-32 large memory configuration by deploying SASWORK in a TMPFS filesystem in-memory to determine if it can mitigate known I/O limitations. The testing approach was to execute a single job performing a SAS workload and subsequently to add identical jobs while observing overall SAS job step times as they each relate to each storage configuration deployed. As jobs were added to the job mix the object was to significantly oversubscribe the 48 core PDom compute and associated storage configuration. As shown in Figure 9 below, different storage configurations were tested in a single PDom with SAS 9.4.

As can been seen, a traditional approach was used to improve storage area network (SAN) I/O performance. As more LUNs, paths, and controllers were added to the SAN Fibre Channel configuration, spreading the I/O workload across available storage assets, I/O performance was sustainable.

The storage configurations with multiple paths were able to offer relatively sustained performance until 16 SAS application solution jobs were executing simultaneously. At this inflexion point, despite the

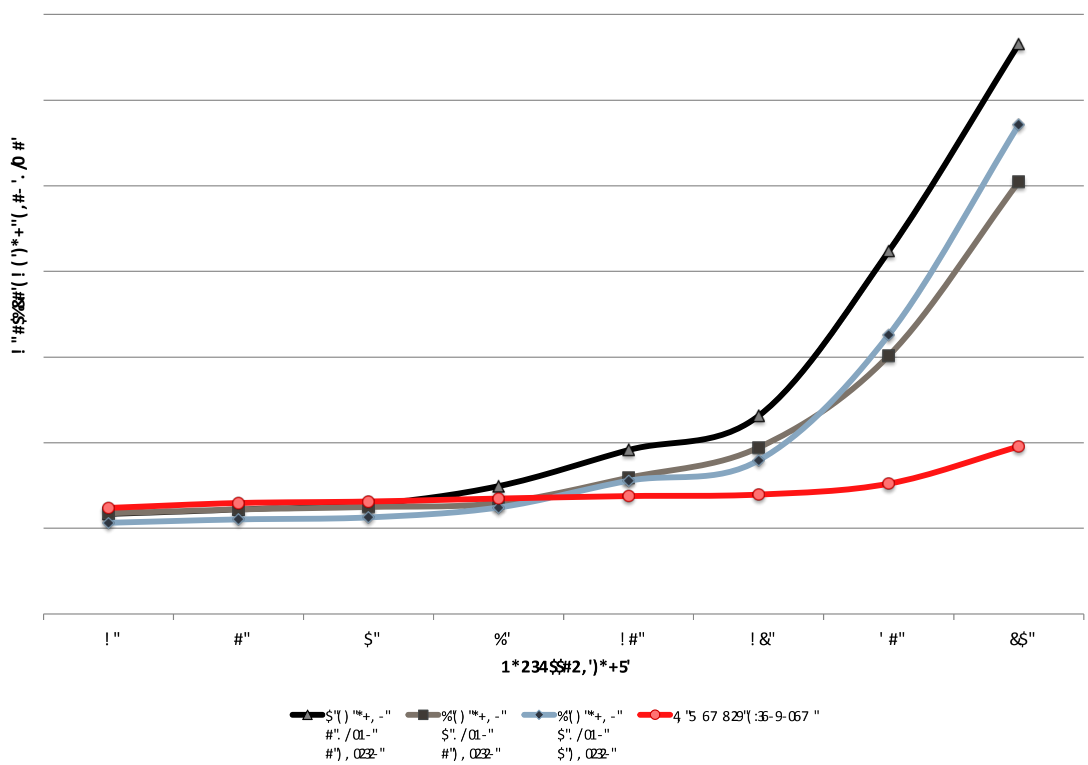

Figure 9. I/O Scalability with different SAN configurations and TMPFS

multiple path configurations, the number of hard disk drive spindles was not sufficient to sustain I/O performance, and storage latency effects began to dramatically affect SAS job step times. On the other hand, by using TMPFS for SASWORK the SAS job step times are dramatically sustainable and remain relatively constant thereby demonstrating excellent scalability under resource intensive workload conditions.

Conclusion

Organizations with significant SAS application solution workloads continue to seek ways to manage and improve SAS performance through IT modernization. As they seek to make these decisions a critical key performance indicator is realizing in advance the advantages associated with deploying a server like the SPARC M5-32 server. To that end, this paper illustrated job concurrency testing that showed how SAS and SAS HPA application solutions exploited the technology of the Oracle M5-32 server demonstrating the ability to scale up/scale out workloads and significantly improvement SAS application solution job performance. In-memory job execution efficiencies highlighted the SPARC M5-32 server computing and memory architecture advantages for SAS workloads while large complex SAS application solution workloads were shown to exploit the extensive computing architecture of the Oracle SPARC M5-32 server .

Oracle continues to deliver technology excellence through engineered systems such as the SPARC M5-32 that can make the difference for organizations by delivering SAS application solution results rapidly thereby providing superior customer value.

For More Information

TABLE 1. ADDITIONAL RESOURCES

ORACLE TECHNOLOGY WITH SAS	
Oracle and SAS partnership	Oracle.com/SAS
SAS Marketing Automation	SAS Marketing Automation
“Scaling SAS Data Access to Oracle Database”	“Scaling SAS Data Access to Oracle Database”
“How to Create a Robust Platform for SAS Grid Computing with Sun Blade 6000 Modular System and the Sun ZFS Storage Appliance”	“How to Create a Robust Platform for SAS Grid Computing with Sun Blade 6000 Modular System and the Sun ZFS Storage Appliance”
ORACLE SPARC M5 SERIES WHITE PAPERS	
“Oracle's SPARC M5-32 Server Architecture”	“Oracle's SPARC M5-32 Server Architecture”
“Maximizing Application Reliability and Availability with SPARC M5-32”	“Maximizing Application Reliability and Availability with SPARC M5-32”
“Consolidation Using the SPARC M5-32 High End Server”	“Consolidation Using the SPARC M5-32 High End Server”
ORACLE	
Oracle Database	Oracle Database
Oracle Optimized Solutions	Oracle Optimized Solutions
Oracle SPARC SuperCluster	Oracle SPARC SuperCluster
Oracle’s SPARC M5 Series	Oracle’s SPARC M5 Series
Oracle Solaris	Oracle Solaris
Oracle’s Sun ZFS Storage Appliance	Oracle’s Sun ZFS Storage Appliance

December 2013
Author: Chris Martin and Maureen Chew

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com