

SAS SPDS on Sun Fire X4500 A Performance Study

Nick Welke, Zencos Consulting
SAS Quality Partner

SAS Scalable Performance Data Server (SPDS) is a fully integrated and seamless way to access large volumes of detail data. It serves large numbers of concurrent users by utilizing parallel processing and data server capabilities, offering both vertical (users and queries) and horizontal (data volume) scalability.

Based on our experience, most SAS customers deal with large volumes of de-normalized data; de-normalized SAS data is most efficiently stored in SPDS; and SPDS performed very well in our tests on the Sun Fire X4500.

The Sun Fire X4500 is a new 'hybrid' server introduced in the summer '06. By integrating state of the art server and storage technologies, the Sun Fire X4500 Server delivers the remarkable performance of a four-way x64 server and the highest storage density available, with up to 24 TB in 4U of rack space. This system also delivers incredibly high data throughput for about half the cost of traditional solutions. Our performance study was done with SAS 9 & SPDS 4.4 on Solaris 10, 6/06 with the Solaris ZFS file systems. Solaris ZFS was introduced in Solaris 10 (6/06).



Sun Fire X4500
2x2.6 GHz Opteron (dual core)
16 GB RAM
12/24 TB in 4U
\$2/GB



The configuration that was tested consisted of the following:

- Disk 0, 1 for Solaris Logical Volume Mgr mirror for root (/)
- Four (4) disks configured for spares, alternate holding area
- Forty-two (42) remaining disks configured as seven pools, where each pool contained a disk from each of the six internal controllers, i.e.: /z0 - /z6. SPDS was configured to use these seven file systems.

SPDS provides high-performance data storage of large SAS data tables. It is particularly well suited for use as part of a data warehousing solution, where the SAS system is being used to construct, manage, and analyze enterprise-wide data marts. SPDS offers 100 percent compatibility with the SAS language and products, and it is designed to scale well in performance when handling large tables. SPDS also offers a high availability design that includes a centralized name server for mapping logical database names to real database paths. KERBEROS-type token passing validates and authenticates user logins, and Access Control Lists (ACLs) provide table security. SPDS exploits symmetric multi-processing (SMP) capabilities to achieve maximum performance with targeted threads to perform parallel and overlapped processing.

During preliminary benchmarking of the Sun Fire X4500 several tests were conducted to characterize the performance of SPDS on this system. Baseline SPDS tests compared three SPDS domain configurations; index creation tests compared adding indexes to tables and building tables with pre-defined indexes; concurrency tests simulated multiple users; and append tests characterized the performance of adding records to a table.

Three types of configuration were used in SPDS baseline tests. The simplest configuration, z0, is a single file system, and effectively removes any efficiency that SPDS could offer. The second configuration, BIG1, is a set of five file systems in a simple sequential array. The third configuration, SPDSTEST, uses a zigzag array configured to more evenly distribute the utilization of its component file systems. The primary advantage of the zigzag array is that partitions are being written to the file systems from left to right and from right to left. In a sequential array, such as BIG1, all partitions are written from left to right, resulting in a skewed data distribution (and consequently a skewed I/O workload distribution) that tends to utilize the file systems at the left or beginning of the array more than at the right or end.

Baseline SPDS tests compared the creation of a 32GB table and subsequent searches with various degrees of concurrency. The chart below shows the results of these tests:

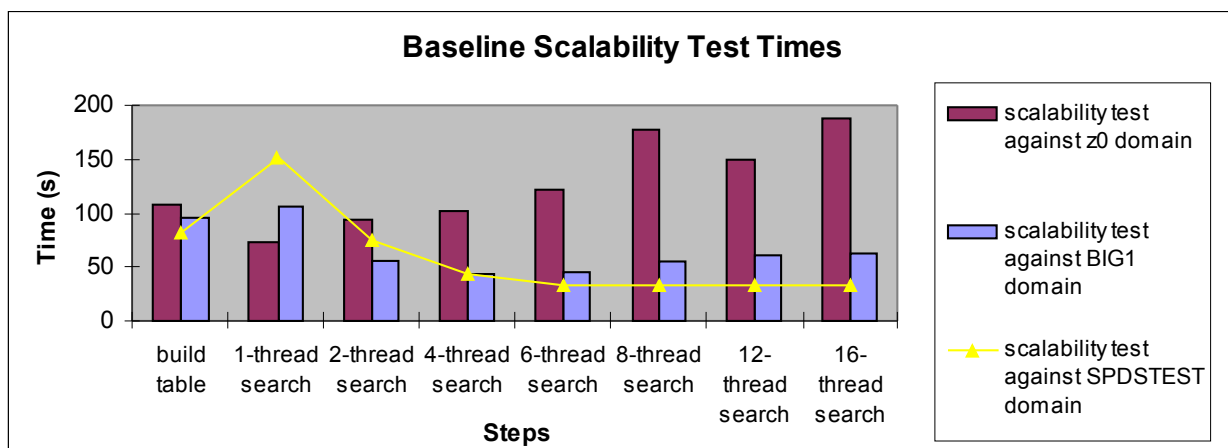


Figure 1. SPDS Test Domain Outperforms Other Domain Configurations

Index creation test compared the time required to build a table and subsequently add indexes to the time required to build a table with indexes built at table creation.

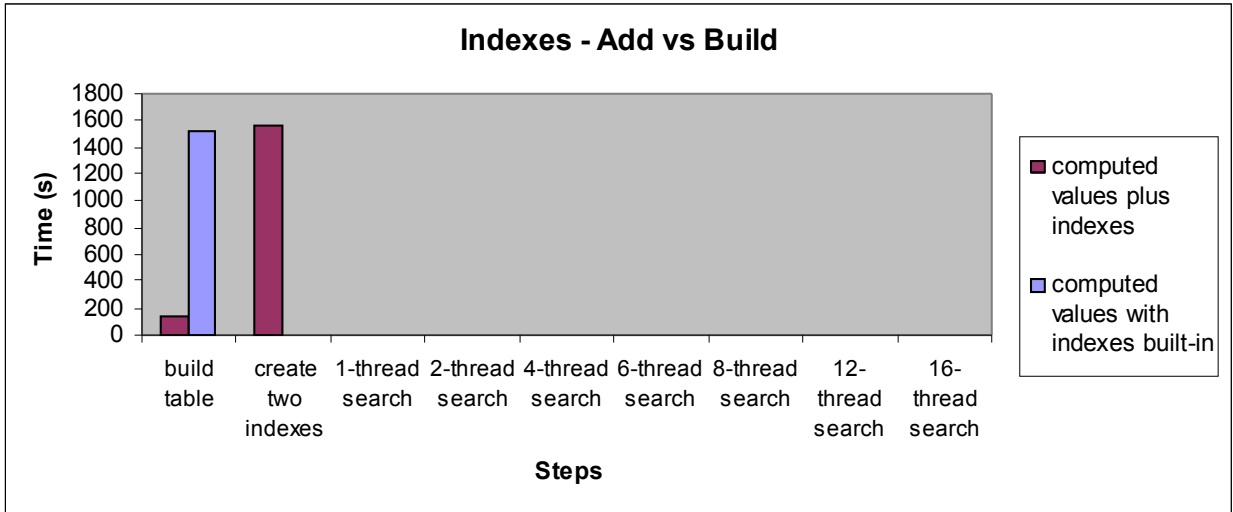


Figure 2. Inline Index Creation (Build) is Faster than Adding Index to Existing Table (Add)

Concurrency tests simulated ten users retrieving a million records each from one table. The append tests characterized the times required to append 32GB to a 32GB table and run the same search against the new 64GB table.

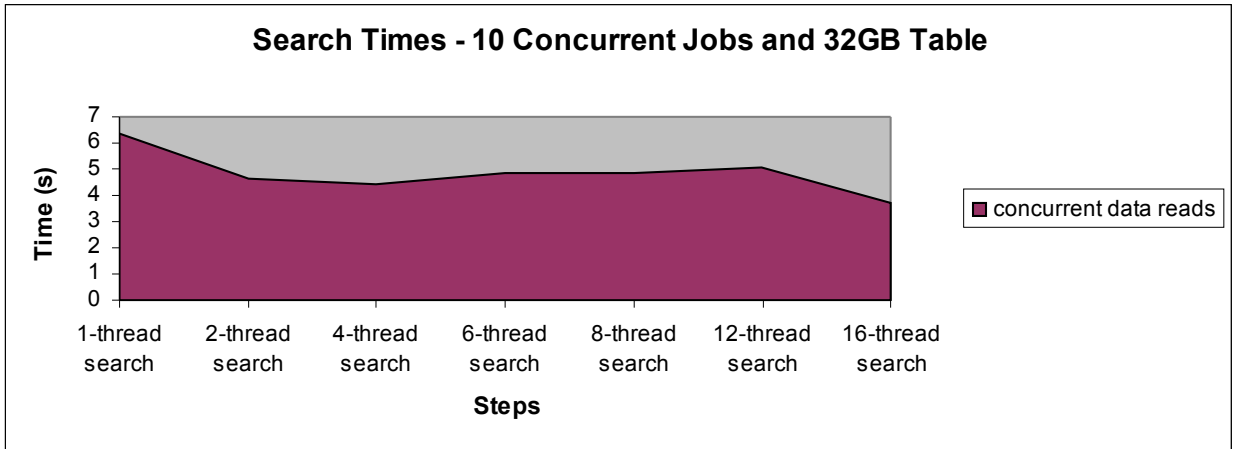


Figure 3 Simulating 10 Users, Average Time to Search 32GB Under 5 Seconds

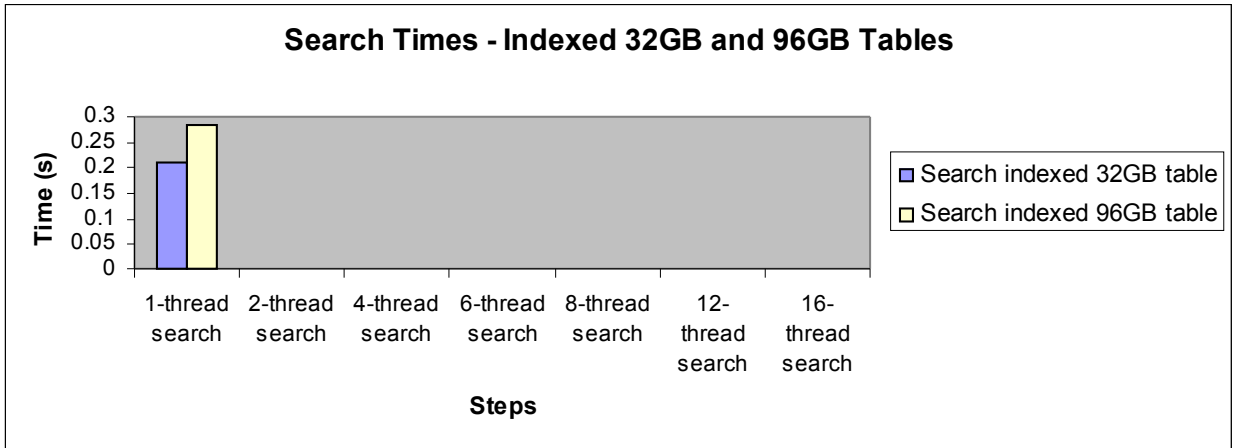


Figure 4 Search of Indexed Tables does not Linearly Increase With Data Size

Based on our experience with SAS SPDS the Sun Fire X4500 unit is an ideal platform for optimal SAS data storage - it provides parallel access to multiple drives while giving SPDS the ability to precisely control how those drives are used. Unlike larger (and much more expensive) storage solutions, the Sun Fire X4500 doesn't add multiple layers of management between SPDS and the spindles. This direct control allows the SPDS administrator to configure SPDS domains with a precision we haven't been seen before in a system with comparable capacity. Our testing showed excellent performance, with read times in the range of 1GB/second. Combined with the unit's cost, the precise control and exemplary performance make the Sun Fire X4500 an SPDS "dream machine".

For more information contact:

Brian McDonald
Consultant
Zencos Consulting LLC
brian.mcdonald@zencos.com

Maureen Chew
Staff Engineer
Sun Microsystems
maureen.chew@sun.com

References

Sun Fire X4500 Home page
<http://www.sun.com/servers/x64/x4500/>