# Tuning SAS for better performance on IBM Netfinity Servers running Windows NT 4.0

## *Tuning SAS for better performance on IBM Netfinity Servers running Windows NT 4.0*

SAS jobs generally stress the system in two ways. They can be I/O intensive (data extraction), computationally intensive (data modeling) or both. It then follows that tuning the I/O subsystem for better performance should allow the I/O intensive jobs to complete sooner, while increasing the processor speed should improve the performance of the computationally intensive jobs. In this publication, we will focus primarily on tuning the I/O subsystem, since small adjustments there will frequently provide a dramatic increase in the performance of most SAS jobs.

The details of our base hardware configuration are given in Appendix A. The system is an IBM Netfinity 7000M10 server with 4 Pentium III Xeon processors, 2GB of RAM and an IBM ServeRaid adapter.

Our first test was to run the five scripts listed below. Some of these scripts are memory or computationally intensive. Since the actions needed to improve performance of these scripts would be to add memory or increase processor speed, these are included simply to illustrate the minimal effect of our I/O subsystem tuning on these jobs.

After running the first test to establish a baseline, we ran them again after each of the following changes to our base configuration:

**Configuration A**: Disable ServeRaid adapter read ahead cache.
**Configuration B**: Reduce the number of physical drives per array from 5 to 2.
**Configuration C**: Move the saswork directory to a single SCSI device.

These changes are not cumulative. We returned to the base configuration before making the next adjustment and re-running the tests.

## SAS Test Scripts

The US Census Bureau makes available to the public its Public Use Microdata Sample. We have obtained this data as flat text files, and our suite of test scripts will be manipulating this data in various ways.
The current test suite consists of:

census_test1: Reads data from various flat files to create two SAS data sets and then validates the information in each SAS data set to make sure there is no incorrect information. (I/O intensive test)

census_test2: Sort the data sets created with test 1 and create various indexes so that they can be merged together if necessary. (memory intensive)

census_test3: Summarizes and collects frequency information on the two tables. (I/O intensive test)

census_test4: Collects statistical information on the fields in the files. (I/O and CPU intensive test with a memory intensive part)

census_test5: Creates several Multidimensional Database cubes (memory intensive)

## *Tuning the I/O subsystem for better performance*

In this section we will be considering how to tune the storage subsystem for better performance while running I/O intensive SAS jobs.

The base configuration was chosen as a somewhat optimal balance between the I/O and memory requirements of the test scripts. The time to completion for each of the scripts run under the base configuration is plotted graphically in Table 1. Table 2 illustrates the percentage of improvement or degradation of performance of the scripts for each configuration based on time to completion. We used the base run time to completion as the basis for comparison for the remaining configuration runs.

**Table 1**

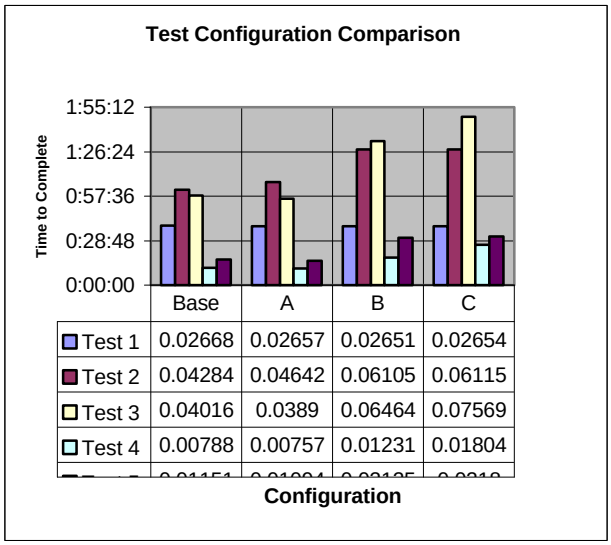


| | Base | A | B | C |
|---|---|---|---|---|
| Test 1 | 0.02668 | 0.02657 | 0.02651 | 0.02654 |
| Test 2 | 0.04284 | 0.04642 | 0.06105 | 0.06115 |
| Test 3 | 0.04016 | 0.0389 | 0.06464 | 0.07569 |
| Test 4 | 0.00788 | 0.00757 | 0.01231 | 0.01804 |

**Table 2**

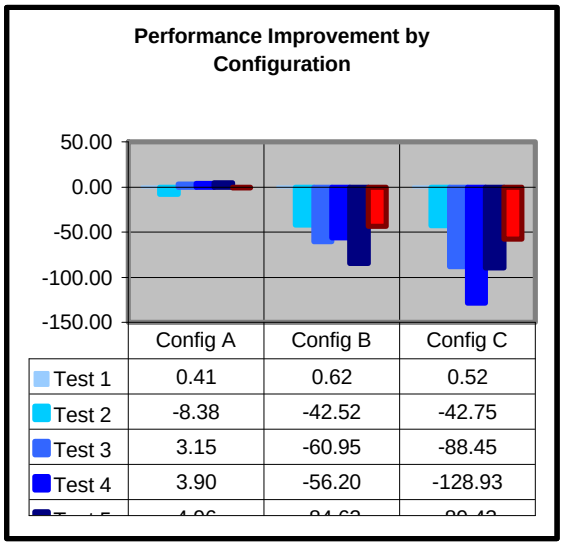


| | Config A | Config B | Config C |
|---|---|---|---|
| Test 1 | 0.41 | 0.62 | 0.52 |
| Test 2 | -8.38 | -42.52 | -42.75 |
| Test 3 | 3.15 | -60.95 | -88.45 |
| Test 4 | 3.90 | -56.20 | -128.93 |

Configuration A

For Configuration A, we chose to disable the ServeRAID adapter read-ahead cache to see what effect this would have on the performance of the test scripts. Interestingly, census_test1.sas showed a slight increase in performance. We had expected a decrease in performance resulting from disabling read ahead cache in the ServeRAID configuration. The performance of census_test2.sas actually decreased by about 8.4%. We attribute this decrease to the saswork drive not being able to take advantage of the read-ahead cache. The remaining tests, census_test3.sas, census_test4.sas and census_test5.sas also showed an increase in performance with the adaptive read-ahead cache disabled. The increase was small – 3, 4 and 5% respectively. Overall, performance was decreased by about 1% with the adaptive read ahead cache disabled.

## Configuration B

For configuration B, we reduced the number of physical hard drives in each RAID 0 logical drive from five to two in order to see the limiting effects of the individual drive transfer rates. Again, it is interesting to note that the first test, census_test1.sas showed a slight improvement in performance – about 0.62%. The second test, however, demonstrated the advantages of using multiple smaller disks in a RAID configuration, rather than fewer larger drives. The performance degradation was 42.5%! Tests 3, 4 and 5 also showed performance decreases of 61, 56 and 85% respectively. This is most likely due to the operating system waiting on the storage subsystem to transfer the data between the media and the ServeRAID SCSI bus. In the base configuration, reads and writes were being performed in parallel across five physical devices, thus reducing the time spent waiting for data to be actually written to the hard disk media.

## Configuration C

For configuration C, we used the base configuration from Appendix A, but changed the location of the saswork directory from a RAID 0 logical drive attached to the ServeRAID adapter to a single 9.1GB SCSI 2 drive attached to the Netfinity integrated SCSI adapter. The first test, census_test1.sas, again showed a slight performance increase of less than one percent. The second test, census_test2.sas, suffered an almost 46% decrease in performance, taking 1hour and 10 minutes to complete. The third, fourth and fifth tests also showed decreases in performance of 88, 129, and 89% respectively. Certainly, the location of saswork appears to have the most dramatic effect on SAS job performance.

## *Conclusion*

For our tests, we found that the ServeRaid adaptive read-ahead cache has little effect on the overall performance of our tests. This may be due in part to the fact that we had optimized the hardware configuration and physical data placement prior to beginning the tests. We also found that more is better when it comes to physical disk devices. Striping data across many physical disks by using RAID 0 dramatically increases the performance of our tests. The most important step to take when working to resolve your SAS I/O constraints is certainly to improve the throughput of your saswork volume. For peak performance, the saswork directory should always be placed on a RAID 0 array containing as many physical disks as possible.

## *Appendix A*

Base Test Configuration

NT System Configuration

The system used for the tests was a Netfinity 7000M10 with four 550MHz Pentium Xeon processors, 2MB L2 cache. The operating system was Windows NT 4.0, SP6a. The system contained 2G of system RAM. Disk subsystem consisted of a ServeRaid 3H adapter connected to an EXP-15 external SCSI enclosure.

Storage configuration

The EXP-15 was in a split bus configuration, with each of the two external RAID adapter connections attached to separate hot swap backplanes each containing five 9.1GB hard drives. The 9.1GB SCSI hard drive containing the operating system was attached to the system's internal hot swap drive backplane which, in turn, was connected to one of the system's integral Adaptec SCSI controllers.

The IBM ServeRaid 3H adapter, with its attached 9.1GB hard drives, was initially configured with two RAID 0 arrays. The first array consisted of two physical hard drives on SCSI channel 1 and three physical drives on SCSI channel 2. The second array consisted of three physical hard drives on SCSI channel 1 and two physical drives on SCSI channel 2. These arrays were presented to the operating system as two logical drives, X: and Y:, of approximately 45GB each. The stripe unit size for the adapter was set to 8, write cache was set to 'write through', and read ahead cache was set to Adaptive.

SAS Configuration

The test scripts were run using SAS version 8.0 The saswork directory was pointed to Y:\

## For questions, comments or more information, contact:

Brian Smyla
IBM International Competency Center at SAS Institute
(919) 531-1163
bsmyla@us.ibm.com

Torre DeVito
IBM International Competency Center at SAS Institute
(919) 531-1420
torre@us.ibm.com