SAS® Forecast Server
Performance Brief
SunFire X4600, Solaris 10

Executive Summary

The Sun Scalable BIDW (Business Intelligence/Data Warehousing) platform on the Solaris 10 Operating Environment provides a powerful platform for a wide variety of tasks such as SAS high end analytics as well as SAS data integration. This performance brief summarizes a performance modeling exercise using the SAS High Performance Forecasting (HPF) Benchmark, Version 1.0. This demonstrates the number of forecast and model selection time series that can be processed per hour on a Sun Fire X4600 using Dual-Core AMD Opteron™ 885 Processors.

Configuration:

- Sun Fire X4600, 8x 2.6 GHz Dual-Core AMD Opteron™ 885 processors, 64 GB of RAM
- Sun StorageTek 6920 FC Array
- Solaris 10, Update 3
- SAS® Forecast Server 1.2
- SAS 9.1.3

Table 1 – Model Selection Throughput in Series per Hour

Workload Size	Small	Medium	Large
Total server throughput in series/hr	360,721	708,661	718,850
Elapsed time (HH:MM:SS)	0:08:19	0:42:20	4:10:24
Workspace used	6 GB	57 GB	324 GB
Memory used	19 MB	51 MB	109 MB

Large: 16 jobs, 3000000 time series processed
Medium: 16 jobs, 500000 time series processed
Small: 8 jobs, 50000 time series processed

Table 2 – Model Fitting and Forecasting Throughput in Series per Hour

Workload Size	Small	Medium	Large
Total server throughput in series/hr	1,855,670	3,180,212	3,430,750
Elapsed time (HH:MM:SS)	0:01:37	0:09:26	0:52:28
Workspace used	12 GB	111 GB	648 GB
Memory used	18 MB	46 MB	107 MB

Large: 16 jobs, 3000000 time series processed
Medium: 16 jobs, 500000 time series processed
Small: 8 jobs, 50000 time series processed

Background

SAS Forecast Server predicts future activity – a forecast – based on time series data. It can detect underlying trends in data as well as adjust for seasonality. For example, typical forecasts requested might be next year's total sales or sales for the next holiday shopping season. The SAS HPF Benchmark runs SAS Forecast Server against several problem sizes: Small, Medium, and Large, which illustrate the type of computing resources required to solve the problem. The benchmark is predominantly CPU-intensive throughout and the time series are distributed across concurrent jobs. To take advantage of the

multiple processors available on the server, each forecasting project is split into smaller single threaded jobs that executing concurrently. The Small scenario uses 8 of the 16 available cores while the medium and large scenarios use all 16 cores.

Because not all features of SAS Forecast Server are used by the workload, this is not meant to suggest that the performance characteristics in this test scenario apply to all possible SAS Forecast workloads.

SAS HPF Benchmark Scenario

The test scenario consists of 3 years of daily sales data sized as Small, Medium, and Large, each with 50,000, 500,000, and 3,000,000 time series respectively. Time series are simply ordered observations spaced across specific time intervals. Three procedures are used to customize the analytical models. The Autoregressive Integrated Moving Average (ARIMA) model specification is selected with the HPFARIMASPEC procedure. ARIMA is known for as one of the most powerful and well-known models for time series forecasting. Next, another powerful model known as the Exponential Smoothing Model (ESM) is selected with the HPFESMSPEC procedure. Both ARIMA and ESM adjust for seasonality, which is particularly important with retail sales data. A variety of simple to complex model options are chosen with each type of model, so there are 6 ARIMA and 7 ESM models used. Finally, the HPFSELECT procedure is used to create lists of candidate model selections. The bulk of computations are performed by PROC HPFENGINE, a procedure which provides large-scale automatic forecasting of time series or transactional data. It automatically evaluates the candidate models and chooses the best one for each time series.

For the Small test scenario, the workload of 50,000 time series was distributed evenly across 8 concurrent jobs. For the Medium test scenario, the workload of 500,000 time series was distributed evenly across 16 concurrent jobs. For the Large test scenario, the workload of 3,000,000 time series was distributed evenly across 16 concurrent jobs. As applied by Sun in this instance, all 16 processor cores of the Sun Fire X6400 were available to the Small, Medium, and Large scenarios. Therefore in the Small case, only 8 cores are significantly utilized because the 8 jobs are single-threaded. In the Medium and Large scenarios, all 16 processor cores are significantly utilized by the workload.

The first phase in the forecasting process is model selection, in which SAS Forecast Server automatically determines the time series model that works best with the data. The second phase is Fit and Forecasting, where the model parameters are estimated against the data, and then finally a forecast is produced. Model fit and forecasting has approximately 5 times the throughput (the metric here is time series per hour) than model selection in all of the Small, Medium, and Large test scenarios. Another way to look at this is to realize that model selection is something that takes the most time, but the recommendation is to do this less frequently. Though the forecast might be run daily or weekly, model selection may be more appropriate to be run on a monthly or quarterly basis as new sales data comes in.

Performance Results - Small Workload (50,000 series)

Figure 1 shows the throughput of the model selection phase in each of 8 concurrent jobs in time series per hour for the Small workload. The aggregate throughput for all jobs is summarized in Table 3.

.

Figure 1 – Model Selection Job Throughput in for Small Workload

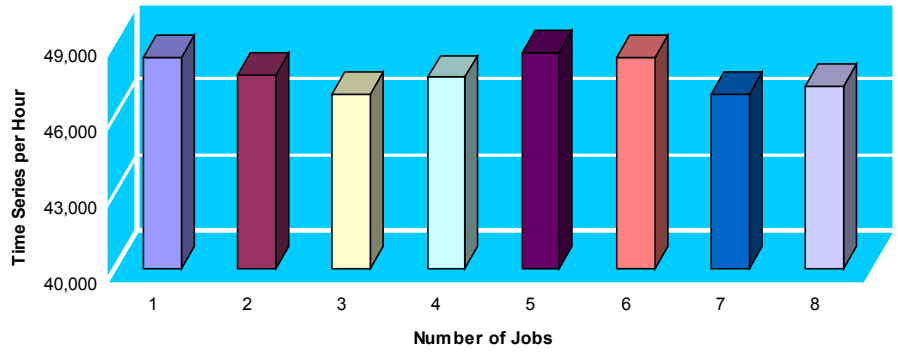

Table 3 – Total Model Selection Throughput for Small Workload

Model Selection	
Total server throughput	360,721 series/hr
Elapsed time (HH:MM:SS)	0:08:19
Workspace used	6 GB
Memory used	19 MB

Figure 2 shows the throughput of the model fit and forecasting phase in each of 8 concurrent jobs in time series per hour for the Small workload. The aggregate throughput for all jobs is summarized in Table 4.

Figure 2 – Fit and Forecast Job Throughput for Small Workload

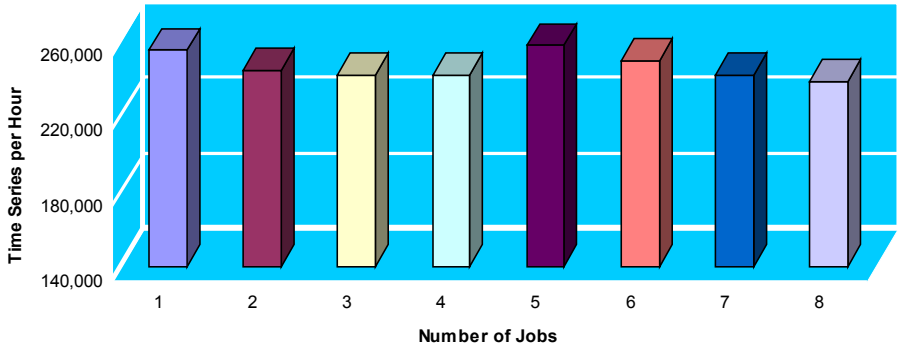

Table 4 – Total Fit and Forecasting Throughput for Small Workload

Fit and Forecasting	
Total server throughput	1,855,670 series/hr
Elapsed time (HH:MM:SS)	0:01:37
Workspace used	12 GB
Memory used	18 MB

Performance Results – Medium Workload (500,000 series)

Figure 3 shows the throughput of the model selection phase in each of 16 concurrent jobs in time series per hour for the Medium workload. The aggregate throughput for all jobs is summarized in Table 5.

Figure 3 – Model Selection Job Throughput for Medium Workload

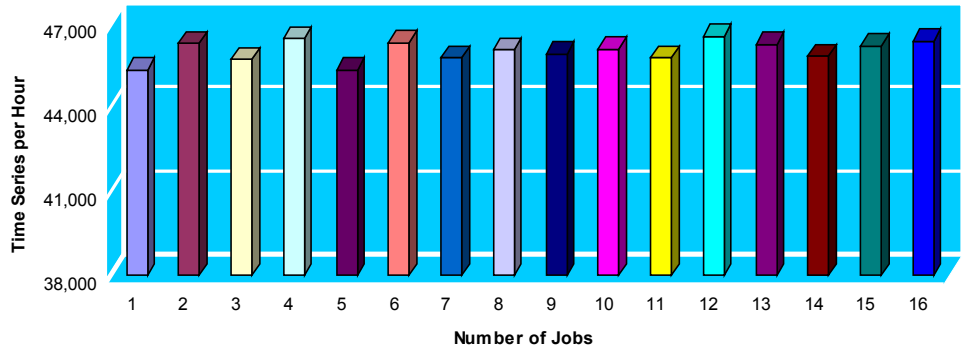

Table 5 – Total Model Selection Throughput for Medium Workload

Model Selection	
Total server throughput	708,661 series/hr
Elapsed time (HH:MM:SS)	0:42:20
Workspace used	57 GB
Memory used	51 MB

Figure 4 shows the throughput of the model fit and forecasting phase in each of 16 concurrent jobs in time series per hour for the Medium workload. The aggregate throughput for all jobs is summarized in Table 6.

Figure 4 – Fit and Forecast Job Throughput for Medium Workload

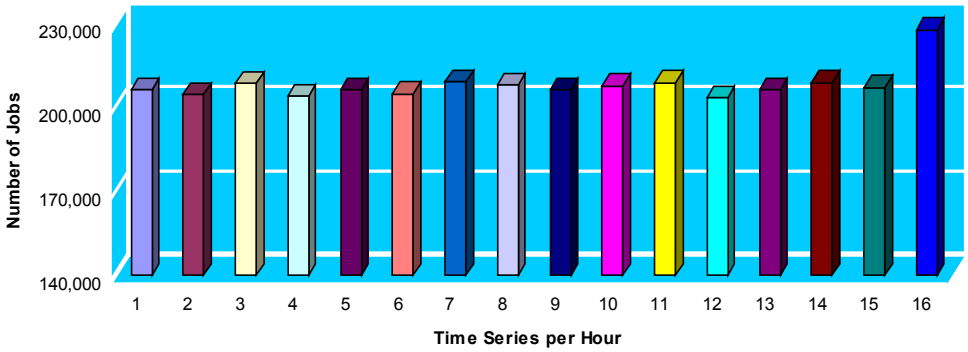

Table 6 – Total Fit and Forecasting Throughput for Medium Workload

Fit and Forecasting	
Total server throughput	3,180,212 series/hr
Elapsed time (HH:MM:SS)	0:09:26
Workspace used	111 GB
Memory used	46 MB

Performance Results – Large Workload (3,000,000 series)

Figure 5 shows the throughput of the model selection phase in each of 16 concurrent jobs in time series per hour for the Large workload. The aggregate throughput for all jobs is summarized in Table 7.

Figure 5 – Model Selection Job Throughput for Large Workload

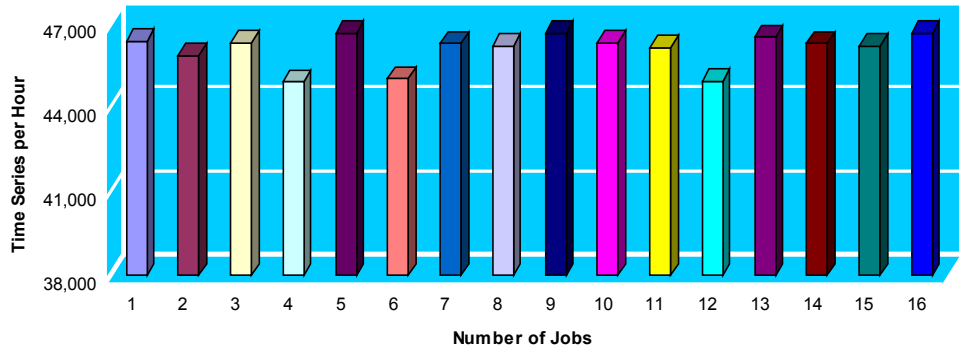

Table 7 – Total Model Selection Throughput for Large Workload

Model Selection	
Total server throughput	718,850 series/hr
Elapsed time (HH:MM:SS)	4:10:24
Workspace used	324 GB
Memory used	109 MB

Figure 6 shows the throughput of the model fit and forecasting phase in each of 16 concurrent jobs in time series per hour for the Large workload. The aggregate throughput for all jobs is summarized in Table 8.

Figure 6 – Fit and Forecast Job Throughput for Large Workload

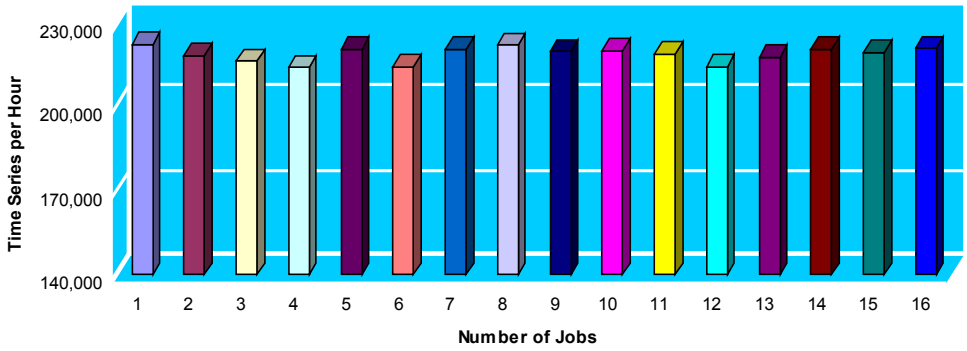

Table 8 – Total Fit and Forecasting Throughput for Large Workload

Fit and Forecasting	
Total server throughput	3,430,750 series/hr
Elapsed time (HH:MM:SS)	0:52:28
Workspace used	648 GB
Memory used	107 MB

Configuration Specifics

Sun Scalable Business Intelligence/Data Warehousing Platform

Sun Fire X4600
8 Dual-Core AMD Opteron 885
Processors
(2.6 GHz, 1 MB L2 per core)
64 GB DDR RAM

Sun StorageTek
6920 FC Array
Solaris 10 ZFS RaidZ

The Sun Fire X4600 is a large-scale shared memory system available in a variety of configurations ranging from 2 to 8 AMD Opteron™ processors. The Solaris 10 x64 release builds on more than 20 years of experience to provide customers with a powerful, secure, and highly available environment by design.

The Sun Fire X4600 used was configured with 8 Dual-Core AMD Opteron 885 processors and 64 GB of DDR RAM. The disk configuration consisted of a Sun StorageTek 6920 FC Array with 10k RPM 73 GB HDDs configured with Solaris 10 ZFS RaidZ volumes.

Because the X4600 uses Dual-Core AMD Opteron processors, there's enough CPU capability to easily handle up to 16 concurrent jobs in the Medium and Large workloads. It is important to realize that for the Small workload, only half the system is utilized and there are 8 remaining cores available to do useful work.

Summary

One of today's chief computing challenges is scaling cost-effectively to meet business demand. The Sun Scalable BIDW Platform is an excellent match for the compute-intensive and I/O-intensive demands of advanced SAS analytics. Combined with SAS Forecast Server software, the Sun platform gives the capability to produce millions of high-quality forecasts quickly and automatically to improve planning and decision making at all levels of business.