subject: **SPEC SFS Release 1.1 Reporting Rules**

date: **November 28, 1994**

from: **SPEC Steering Committee**

ABSTRACT

This paper provides the rules for reporting results of official runs of the SPEC SFS Release 1.1 Benchmark suite according to the norms laid down by the SPEC SFS subcommittee and approved by the SPEC Open Systems Steering committee. This is a companion paper to "SPEC SFS Release 1.1 Run Rules", which provides rules to follow for all submitted or reported runs of the SPEC System File Server (SFS) Benchmark suite. These papers can be found in files RUN_RULES and RPT_RULES in the $SPEC directory on the release tape.

Memorandum to
Performance Analysts

subject: **SPEC SFS Release 1.1 Reporting Rules** date: **November 28, 1994**

from: **SPEC Steering Committee**

1. SPEC SFS RELEASE 1.1 REPORTING RULES

SPEC SFS Release 1.1 is a maintenance update for the 097.LADDIS benchmark. Following the release of SPEC SFS 1.1, all future SFS testing will use SPEC SFS Release 1.1.

The 097.LADDIS benchmark progressively stresses an NFS file server by increasing the NFS operation request rate (NFS load) of the NFS clients used to generate load on the server.

The performance metric is defined as the average NFS operation response time measured at a specified NFS load (NFS operations per second).

The NFS server's performance is characterized in terms of a complete average NFS operation response time versus NFS throughput curve for a given server configuration. Also, the NFS capacity of the server in terms of NFS operations per second is reported at no higher than an average NFS operation response time of 50 milliseconds.

The reporting rules detailed in the following sections, as stipulated by the SPEC Open Systems Steering Committee (OSSC), are mandatory.

1.1 Reporting Guidelines

This section describes the standard SPEC reporting format that must be used when reporting SPEC SFS Release 1.1 results.

1.1.1 Metrics And Reference Format

The performance metric is the average NFS operation response time, in terms of milliseconds, for a given NFS load, in terms of NFS operations per second. The results of a benchmark run, comprised of several NFS load levels, are plotted on a performance curve on the results reporting page. The data values for the points on the curve are also enumerated in a table.

When referencing any point on the performance curve, the format "XXX SPECnfs_A93 NFS operations per second at YY Milliseconds average response time" must be used. If an abbreviated format is required, the format "XXX SPECnfs_A93 NFS ops./sec. @ YY Msec." must be used.

While all SPEC members agree that a full performance curve best describes a server's performance, the need for a single figure of merit is recognized.

The SPEC SFS single figure of merit is a triple which specifies:

1. NFS throughput at an average NFS response time no greater than 50 milliseconds.
2. The number of "SPECnfs_A93 USERS" determined by the algorithm:

$$\text{Number of SPECnfs\_A93 USERS} = \frac{\text{Maximum NFS Throughput @ <= 50 milliseconds}}{\text{10 NFS Operations/Second per SPECnfs\_A93 USER}}$$

The SPEC SFS single figure of merit is reported as a boxed item on the LADDIS performance graph on the reporting page:

<NFS Throughput> @ <Response Time>	==> XXX SPECnfs_A93 USERS
where:	
<NFS Throughput>	- is stated as "ZZZ SPECnfs_A93 NFS Ops./Sec."
<Response Time>	- is stated as "YY Msec."
==>	- is the logical implication operator symbol
XXX SPECnfs_A93 USERS	- is printed in a bold typeface

1.1.2 Reporting Format

1.1.2.1 Table Format

A table, from which the server performance graph is constructed, consists of a number of data points which are the result of a single run of the benchmark. The table consists of two columns, NFS Throughput in terms of SPECnfs_A93 NFS operations per second rounded to the nearest whole number on the left, and Average NFS Response Time in terms of Milliseconds rounded to the nearest tenth on the right. The data points are selected based on the criteria described in Section 10.1.2.2.

1.1.2.2 Graphical Format

NFS server performance is depicted in a plot with the following format:

1. Average Server Response Time in units of milliseconds is plotted on the Y-axis with a range from 0 to x milliseconds, where x obeys the run rule in 10.1.1, item 1.
2. The plot must consist of a minimum of 10 data points uniformly distributed across the range of the maximum server load. Additional points beyond these 10 uniformly distributed points also can be reported.
3. All data points of the plot must be enumerated in the table described in Section 10.1.2.1.
4. No data point within 25% of the maximum reported throughput may be reported whose "Actual NFS Mix Pcnt" versus "Target NFS Mix Pcnt" differs by more than 10% for any operation.

1.1.3 System Configuration

The system configuration information that is required to duplicate published performance results must be reported.

This list is not intended to be all-inclusive, nor is each feature in the list required to be described. The rule of thumb is: if it affects performance or the feature is required to duplicate the results, describe it.

1.1.3.1 Hardware

1.1.3.1.1 Server

The following server hardware components must be reported:

1. Vendor's (Benchmark User's) name
2. System model number, main memory size, number of CPUs
3. Critical customer-identifiable firmware or option versions such as network and disk controllers, write caches, or other accelerators
4. Number, type, and model of disk controllers
5. Number, type, model, and capacity of disk drives
6. Relationship among disk controllers and disk drives
7. Relationship among disk drives and filesystems
8. Number, type, and model of filesystem/NFS accelerators
9. Number, type, and model of network (Ethernet/FDDI) controllers
10. Number of networks and type
11. Number, type, model, and relationship of external network components to support server (e.g., external routers)
12. Alternate sources of stable storage including un-interruptible power supply systems (UPS), battery-backed caches, etc.

1.1.3.1.2 Load Generators

The following load generator hardware components must be reported:

1. System model number, main memory size, number of CPUs
2. Compiler used to compile benchmark
3. Number, type, model, and relationship of external network components

1.1.3.2 Testbed Configuration

A brief description of the system configuration used to achieve the benchmark results is required. The minimum information to be supplied is:

1. Relationship of load generators, load generator type, network, filesystem, and filesystem mount point options.
2. If the configuration is large and complex, added information should be supplied either by a separate drawing of the configuration or by a detailed written description which is adequate to describe the system to a person who did not originally configure it.

1.1.3.3 Software

The following software components must be reported:

1. Shipping OS version or pre-release OS version, deliverable within six months
2. Other clarifying information as required to reproduce benchmark results (e.g. number of NFS daemons, server buffer cache size, disk striping, non-default kernel parameters, etc.)
3. Number of load generators, number of processes per load generator, server filesystems targeted by each load generator
4. Number of BIOD_MAX_READ and BIOD_MAX_WRITEs used.

1.1.3.4 Notes/Summary of Tuning Parameters

This section is used to document:

1. Single or multi-user state
2. System tuning parameters other than default
3. Process tuning parameters other than default
4. Background load, if any
5. ANY changes made to the individual benchmark source code including module name, line number of the change.
6. Additional information such as compiler options may be listed here.
7. Additional important information required to reproduce the results, which do not fit in the space allocated above must be listed here.
8. A full description of the definition of tuning parameters used should be included as an auxiliary document similar to the tuning notes included with SPECint and SPECfp CPU benchmark results.

1.1.3.5 Other Required Information

The following additional information is also required to appear on the results reporting page for SPEC SFS Release 1.1 results:

1. General Availability of the System Under Test. All the system, hardware and software features are required to be available within 6 months of the date of test.
2. The date (month/year) that the benchmark were run
3. The name and location of the organization that ran the benchmark
4. The SPEC license number