STANDARD PERFORMANCE EVALUATION CORPORATION

spec newsletter

Volume 5, Issue 1 *March 1993*

IN THIS ISSUE

RESULTS SUMMARIES

SPEC Announces System File Server Benchmark Suite – SFS Release 1.0

LADDIS: The Next Generation in NFS Server Benchmarking

by Bruce E. Keith
Digital Equipment Corporation
Maynard, Mass.

On Thursday, March 18, 1993, at the UniForum 1993 Conference in San Francisco, Calif., the Standard Performance Evaluation Corporation (SPEC) announced the availability of its first release of a new system-level file server (SFS) benchmark suite, SFS Release 1.0. The suite is built around the 097.LADDIS NFS file server benchmark.

The SFS Release 1.0 suite provides users and the industry with the first fair, consistent and complete method for measuring and reporting NFS file server performance.

The 097.LADDIS benchmark generates a synthetic NFS workload based on a workload abstraction of an NFS operation mix and an NFS operation request rate. The benchmark measures NFS file server performance in terms of NFS operation response time and throughput. SPEC's run rules for 097.LADDIS provide a consistent method of measuring NFS file server performance. SPEC's reporting rules for the benchmark provide a uniform method for reporting NFS file server performance results.

The workload generated by the initial offering of the 097.LADDIS benchmark emulates an intensive software development environment at the NFS protocol level, based on several independent NFS file server performance studies collected from the NSF industry and academia.

SPEC expects that future releases of the SFS suite will contain additional workload abstractions for other NFS environments such as CAD/CAM, imaging, animation, commercial applications, and others. To this end, SPEC strongly encourages the workload characterization of actual NFS environments to serve as the basis for the development of future workload abstractions for the 097.LADDIS benchmark.

Background

Motivated by the need to provide customers with a unified, standard method of comparing NFS file server performance, the LADDIS Group was formed in February 1990. The group consisted of technical representatives from Legato Systems (Robert Lyon), Auspex Systems (Bruce Nelson),

See **SFS RELEASE 1.0,** *page 3*

SFS RELEASE 1.0, *from page 1*

Data General (Mark Wittle), Digital Equipment Corp. (Bruce Keith), Interphase (Vincent Lefebvre and Peter Oleinick), and Sun Microsystems (John Corbin, Maneesh Dhir and Brian Pawlowski). The group developed the initial version of LADDIS that was submitted to SPEC in August 1991. After the initial code submission to SPEC, LADDIS group members became participants of SPEC's SFS subcommittee. Since the submission of LADDIS to SPEC, the SPEC membership has worked vigorously to critique, refine, and enhance the benchmark in addition to porting the benchmark to the various members' platforms. The SPEC PRE-LADDIS beta test program allowed the benchmark to be ported to non-member platforms as well. The 097.LADDIS benchmark has been ported to AIX, BSD, SVR3, SVR4, and OSF/1 environments.

Written in C, the 097.LADDIS benchmark is loosely based on Legato's 1989 nhfsstone benchmark program. 097.LADDIS maintains and enhances nhfsstone's NFS workload abstraction of an NFS operation mix and an NFS operation request rate. However, 097.LADDIS overcomes other limitations in nhfsstone, specifically, its NFS client platform sensitivity, the accuracy of generated load, and single NFS client operation.

097.LADDIS minimizes NFS client platform sensitivity and offers improved accuracy of its generated load by executing the NFS protocol directly within the benchmark. Thus, NFS protocol implementation issues such as file attribute and data caching are normalized across different vendors' platforms running the 097.LADDIS benchmark. This normalization allows 097.LADDIS to measure file server performance independent of the NFS client platform used to execute the benchmark. Hence, 097.LADDIS is a measure of NFS file server performance and is not a measure of NFS client performance.

The multiple-client capability of 097.LADDIS allows collections of NFS clients to be used to generate an aggregate load on an NFS server using one or more TCP/IP networks.

The multiple-client capability of 097.LADDIS allows collections of NFS clients to be used to generate an aggregate load on an NFS server using one or more TCP/IP networks. This is important since a single client or network may saturate before a given server saturates. Further, the use of multiple NFS clients provides a more realistic environment with respect to network contention.

New functionality in 097.LADDIS scales the size of the server file set targeted by the benchmark with increasing NFS loads. The number of files targeted on the server by the benchmark increases as the NFS load is increased in terms of NFS operations per second. This allows the benchmark to provide a greater degree of realism, given that as NFS load increases due to the addition of more users or the inclusion of more powerful NFS clients or applications, more files are accessed on the server.

Operation

To benchmark an NFS server, the server and two or more NFS client systems are configured on one or more isolated networks. One of the NFS clients is designated the "LADDIS Prime Load Generator." The prime load generator controls the execution of the 097.LADDIS NFS load-generating code on the remaining NFS clients, which are called "LADDIS Load Generators." Figure 1 illustrates a 097.LADDIS testbed comprised of an NFS server, the LADDIS Prime Load Generator, and several LADDIS Load Generators.

The prime load generator typically executes the NFS load-generating code in addition to controlling the benchmark. However, the prime load generator is not required to execute the NFS load-generating code. This permits an external "test manager" to control the benchmark and coordinate other distributed system performance monitoring processes outside the scope of the 097.LADDIS benchmark. Therefore, the benchmark can be used as a server characterization tool as well as a benchmark.

The individual conducting the test uses a single user interface on the 097.LADDIS prime load generator to control the benchmark. All test parameters are defined in a single control file, called **laddis_rc**. The **laddis_mgr** control script is invoked to start the test. The script executes successive runs of the benchmark at increasingly higher NFS load levels while holding the NFS operation mix constant.

Similar to other SPEC benchmark suites, the **runsfs** script seamlessly guides the user through the build and execution of the benchmark by prompting the user to supply required information and to perform required setup tasks.

See **SFS RELEASE 1.0,** *page 4*

SFS RELEASE 1.0, *from page 3*

Metrics

The 097.LADDIS benchmark measures NFS server response at the NFS protocol level on the LADDIS load generators for each NFS load level. For each run of the benchmark, the server's average NFS operation response time is measured for the applied NFS load.

At the conclusion of a series of runs, a graph of server response time versus NFS load (throughput) for the particular mix of NFS operations is constructed. Average NFS operation response time in units of milliseconds is plotted on the y-axis while average NFS throughput in terms of operations per second is plotted on the x-axis.

Different servers or different configurations of a given server can be compared by plotting their 097.LADDIS NFS response time versus NFS throughput curves on a single graph. NFS server performance is most accurately depicted by the NFS response time versus throughput curve. The curve clearly illustrates the rate of server response degradation for increased server load. SPEC recognized the need for a single figure of merit, and the SPEC SFS Release 1.0 reporting rules for 097.LADDIS require that NFS throughput (operations per second) be reported at an average response time of 50 milliseconds. The single figure of merit contains a mapping of the NFS throughput achieved at this response level to a number of users.

The mapping of NFS throughput at the 50 millisecond reference response level to number of users provides an arbitrary, conservative, and consistent method of quantifying NFS server performance. When selecting an NFS server based on SPEC SFS Release 1.0 performance results, users should incorporate all aspects of the SFS Release 1.0 single figure of merit into their decision: NFS response time, NFS throughput and derived number of users.

As with other SPEC benchmarks, users are highly encouraged to relate SPEC SFS Release 1.0 results to their specific computing environments before selecting an NFS file server configuration.

As with other SPEC benchmark suites, SFS Release 1.0 reporting rules require the disclosure of a graph of NFS server performance and a table of associated data points. The rules further require disclosure of a detailed summary of the tested NFS server configuration as well as the testbed used to generate the server's performance results. Identification of hardware, software, network components, and non-default system parameters used to generate the performance results is also required.

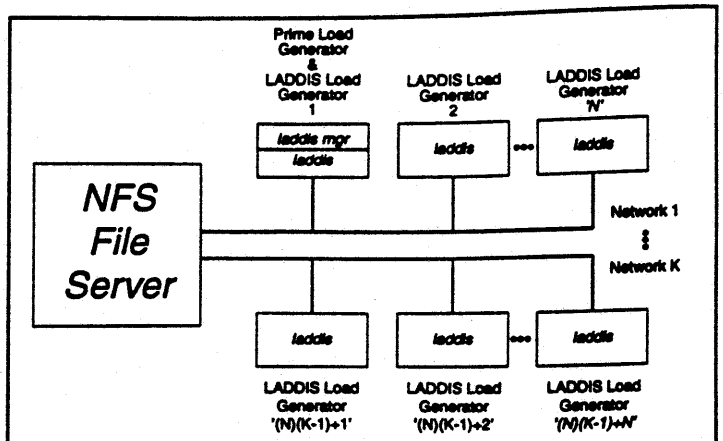

Figure 1: LADDIS Testbed Scenario

Obtaining SFS Release 1.0 Suite

To obtain the SPEC SFS Release 1.0 Benchmark Suite, please complete and mail the SPEC order form found in this newsletter, or contact SPEC at:

Standard Performance Evaluation Corporation (SPEC)
c/o NCGA
2722 Merrilee Drive, Suite 200
Fairfax, VA 22031-4499
Telephone: 703-698-9600
Fax: 703-560-2752
E-mail: spec-ncga@cup.portal.com

Bruce Keith is the NFS performance project leader at Digital Equipment Corporation in Maynard, Mass. He is the release manager for SPEC's System-level File Server (SFS) Benchmark Suite and is the SPEC Project Leader for the 097.LADDIS benchmark. Bruce is also Digital's representative in the LADDIS Group.

Further Reading

Keith, Bruce, "Perspectives on NFS File Server Performance Characterization," USENIX Summer Conference Proceedings, pp. 267-277, June 1990.

LADDIS Group, "LADDIS - A Vendor-Neutral Standard NFS Benchmark," INTEROP 1991 Fall Conference Proceedings, October 9, 1991.

Sandberg, Russel, et al., "Design and Implementation of the Sun Network Filesystem," USENIX Summer Conference Proceedings, pp. 119-130, June 1985.