

Progress on FileBench

[FileBench](#) is a framework of parameterizable file system workloads for measuring and comparing file system performance, originally developed at Sun Microsystems. FileBench uses a high-level workload language to model the I/O behavior and other characteristics of desired applications. Using the workload modeling language, FileBench can benchmark local file systems with I/O that is either typical of those applications, or spans a target space of workloads.

FileBench continues to evolve with new features aimed at network file server performance measurement. Bugs continue to be fixed, and our user community seems to be experiencing fewer problems when running FileBench. Meanwhile, several key features have been added since last year's progress report:

- Random Variables.
- Composite Flowops.
- Multi-client support.
- Indexed selection of files.
- Framework for network file system client plug-ins.
- NFSv3 plug-in nearing completion.

These changes, along with new workload models that use them, will greatly extend the range of application scenarios that can be modeled and tested against a local or remote file system.