

## The New and Improved FileBench

[FileBench](#) is a framework of 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, allowing testing of local file systems with file I/O typical of those application.

While FileBench continues to be used by performance analysts, development efforts stopped in 2005, and remained halted until recently. Yet the need for a powerful, flexible file system benchmarking tool is as strong as ever. With that in mind, Sun committed resources to further the development of FileBench about eight months ago so that it can reach its full potential. The renewed work has created fresh interest in the use and further development of FileBench among the larger FileBench community.

Much of the new work has involved clean up and maintenance. Specific accomplishments to date include:

- Fixing known bugs in both the source code and workload models.
- Improving usability
- Cleaning up and liberally commenting the source code.
- Greatly enhancing the documentation
- Integrating FileBench into OpenSolaris, so that is now part of the distribution.
- Porting FileBench to Linux by members of the FileBench open source community.

In addition to all these changes which have significantly enhanced FileBench, the implementation of the following additional features is planned or in progress:

- Random Variables.
- Multi-client support.
- NFSv3, NFSv4 and CIFS plug-in client modules.

There is also interest in the FileBench community for Windows™, OSX™ and freebsd support which is expected to happen eventually.

These changes, along with new workload models that use them, will greatly extend the range of application scenarios that can be modeled and tested.