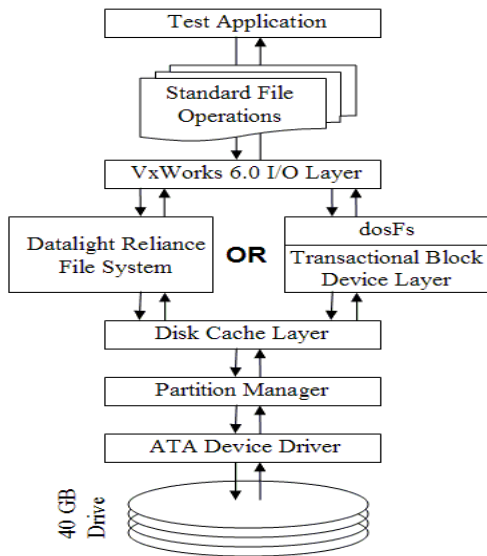




## Performance Comparison, Datalight's Reliance, Wind River VxWorks dosFs with and without a Transactional Layer



### Overview

The traditional view about performance and reliability is that you can only have one or the other, and developers will at some point have to make a conscious trade-off. Datalight is working to make that trade-off decision a thing of the past by making our highly reliable file system as fast as, or faster than the competition.

To illustrate our work, we've run two suites of tests. The first presents an "apples to apples" comparison of Datalight's Reliance file system versus VxWorks' dosFs both with and without their transactional layer. In that test, all three file systems are configured in exactly the same way.

In the second test suite, the adjustable features of Reliance such as transaction rate, transaction rules, and cluster/block size have been configured to optimize it for a specific application. This second set of data provides developers with an idea of the flexibility and file system optimization they can achieve with Reliance.

### "Test Summary"

Datalight characterized the performance of the Datalight Reliance file system and the Wind River VxWorks 6.0 dosFs file system with and without a transactional layer. Each file system used the VxWorks 6.0 ATA block device driver interfacing to a Western Digital 40 GB hard drive in a 1.8 Ghz Pentium Celeron system. The block diagram for the test setup is shown to the left. Additional information about the test can be found in Datalight's White Paper: [Transactional File System Performance Characteristics: Datalight Reliance File System and Wind River dosFs File System](#) available at [www.datalight.com](http://www.datalight.com)

### Test Results – "Apples to Apples"

The performance tests consist of both sequential and random I/O, as well as a comparison of boot times. The summary data shown below is for cluster/block<sup>1</sup> size of 8192 bytes<sup>2</sup> and write/read sizes of 8192 bytes.

Performance Test	Datalight Reliance	dosFs	dosFs w/ Trans. Layer
Sequential Writes	2.04 MB/sec	1.47 MB/sec	0.04 MB/sec.
Sequential Reads	2.70 MB/sec	2.56 MB/sec	2.04 MB/sec
Random Writes	0.91 MB/sec	0.53 MB/sec	0.02 MB/sec
Random Reads	0.83 MB/sec	0.67 MB/sec	0.59 MB/sec
General File Operations	1393 / minute	1934 /minute	418/ minute
Boot Time Comparison <sup>3</sup>	0.05 second	0.20 second	7.3 seconds

<sup>1</sup> The dosFs file system uses the term "cluster" to mean smallest unit of storage that can be allocated by the file system. The Reliance file system uses the term "block" to mean the smallest unit of storage that can be allocated by the file system.

<sup>2</sup> Reliance offers a configurable block size of between 128 bytes and 64 KB. The cluster size of dosFs is set automatically depending on the volume size. The volume size for this test was 335 MB. For this summary, dosFs had a cluster size of 8192 bytes, so Reliance was configured for a block (analogous to cluster size) size of 8192 bytes.

<sup>3</sup> Boot time is a function of the media size and system configuration. These clean mount measurements are for a 1 GB volume size.



The **sequential performance tests** simulated an application writing to and reading from a 100 MB data file in varying byte sizes. At a cluster/block size of 8192 bytes using a write/read size of 8192 bytes, Reliance wrote data at 2.04 MB/sec which was faster than dosFs with and without a transactional layer. Reliance read data at 2.70 MB/second faster than dosFs alone at 2.56 MB/sec. and dosFs with a transactional layer at 2.04 MB/seconds.

The **random performance tests** simulated an application writing to or reading from a 10 MB file representing a database in varying byte sizes. At a cluster/block size of 8192 bytes using a write/read size of 8192 bytes, Reliance and dosFs wrote at rates of 0.91 MB/second and 0.50 MB/second, respectively. Both were faster than dosFs with a transactional layer. Reliance and dosFs read at rates of 0.83 MB/second and 0.67 MB/second.

The **general file operations test** measured the number of random file and directory operations that could be performed over a 5 minute period. dosFs performed slightly more operations per minute (1934/min. versus 1393/min.) than Reliance and both Reliance and dosFs performed significantly more operations than dosFs (418/min.) with a transactional layer.

The **boot time comparison test** measured the time required for the file system to mount after a power loss event occurred. Reliance and dosFs both booted in less than 1 second while dosFs with a transactional layer required over 7 seconds.

### Test Results – Reliance Configured for High Performance

The performance tests consist of both sequential and random I/O, as well as a comparison of boot times. The summary data shown below is for cluster/block<sup>4</sup> size of 8192 bytes<sup>5</sup> and write/read sizes of 128 Kbytes. This configuration of Reliance would be most used by many types of Consumer Devices where typically large files are written to the flash under predictable conditions and the flash is mostly used for reading operations. Such devices are portable media players, streaming media devices, and MP3 players.

Performance Test	Datalight Reliance	dosFs	dosFs w/ Trans. Layer
Sequential Writes	2.95 MB/sec	1.47 MB/sec	0.04 MB/sec.
Sequential Reads	2.81 MB/sec	2.56 MB/sec	2.04 MB/sec
Random Writes	2.93 MB/sec	0.53 MB/sec	0.02 MB/sec
Random Reads	2.47 MB/sec	0.67 MB/sec	0.59 MB/sec
Boot Time Comparison <sup>6</sup>	.05 seconds	0.20 second	7.3 seconds

The **sequential performance tests** simulated an application writing to and reading from a 100 MB data file in varying byte sizes. At a cluster/block size of 8192 bytes using a write/read size of 128 Kbytes, Reliance wrote data at 2.95 MB/sec which was faster than dosFs with and without a transactional layer. Reliance read data at 2.81 MB/second faster than dosFs alone at 2.56 MB/sec. and dosFs with a transactional layer at 2.04 MB/seconds.

The **random performance tests** simulated an application writing to or reading from a 10 MB file representing a database in varying byte sizes. At a cluster/block size of 8192 bytes using a write/read size of 128 Kbytes, Reliance and dosFs wrote at rates of 2.93 MB/second and 0.50 MB/second, respectively. Both were faster than dosFs with a transactional layer. Reliance and dosFs read at rates of 2.47 MB/second and 0.67 MB/second.

The **boot time comparison test** measured the time required for the file system to mount after a power loss event occurred. Reliance and dosFs both booted in less than 1 second while dosFs with a transactional layer required over 7 seconds.

<sup>4</sup> The dosFs file system uses the term “cluster” to mean smallest unit of storage that can be allocated by the file system. The Reliance file system uses the term “block” to mean the smallest unit of storage that can be allocated by the file system.

<sup>5</sup> Reliance offers a configurable block size of between 128 bytes and 64 KB. The cluster size of dosFs is set automatically depending on the volume size. The volume size for this test was 335 MB. For this summary, dosFs had a cluster size of 8192 bytes, so Reliance was configured for a block (analogous to cluster size) size of 8192 bytes.

<sup>6</sup> Boot time is a function of the media size and system configuration. These clean mount measurements are for a 1 GB volume size.