

WildFire File System

Abstracted Performance Goals and
Test Results from a Paper Prepared for SDC 2019,
Discussing How to Build an All-Flash
File System with Superior Performance and Data Integrity.

Douglas Dumitru
CTO, WildFire Storage
doug@easyco.com
(1) 610-237-2000 x 43

The Performance Goal

- 5 uSec file read and create overhead.
- Single IO file read
 - If not cached
- Fractional IO file create
 - Assuming a cached read

The Efficiency Goal

- Use storage space for actual file contents
 - No “allocation units”
- Low write amplification even for very small files.
- Flash “perfect” write patterns

The Feature Set Goal

- Device Redundancy
 - Built in array support using erasure codes
 - Up to 250 drives in a “write set”
 - Up to 8 parity drives
 - Flexible “pool” of drives
 - Add and remove drives on the fly
 - Fast parity reconstruction
 - Automatic size expansion / reduction
- Data Validation
 - End to end CRC32 on every 4K block

The Scalability Goal

- One to hundreds of drives
- 64 GB to 16 PB target array size
- > 16 trillion file support in a single directory
- > PB file size

The Performance Goal

- 5 uSec latency on small file read or create
- Excellent parallelism
- No layout degradation for “huge” directories
- Minimal overhead for housekeeping ...
 - Access times and other directory “extras” all have small impact on performance.

Early Benchmarks – File Create Tests

- Small File Create:
 - 100 byte files in a single directory
 - Create ‘n’ files
 - Sync at end
 - All times are uSec per file.

	WFFS	EXT4	XFS
100K	3.55 uSec	12.6 uSec	24.20 uSec
1M	4.88 uSec	24.72 uSec	93.28 uSec
2M	5.97 uSec	34.78 uSec	
4M	6.63 uSec	53.90 uSec	258.14 uSec
8M	6.68 uSec		

Early Benchmarks – File Create Tests

- Larger File Create:
 - 1K, 3K, 4K, 5K, and 64K byte files in a single directory
 - Create 'n' files
 - Sync at end
 - All times are uSec per file.

WFFS	1K	3K	4K	5K	64K
100K	5.51	6.48	6.95	8.05	40.50
1M	9.59	8.69	9.17	13.13	

EXT4	1K	3K	4K	5K	64K
100K	10.09	9.00	9.07	10.38	127.40
1M	23.13	27.63	28.16	38.80	

XFS	1K	3K	4K	5K	64K
100K	25.62	22.95	23.44	24.11	128.97
1M	93.86	92.69	92.43	102.38	

Early Benchmarks – File Create Tests

- Small File Create:
 - 100 byte files
 - Create 10,000 files
 - Sync **every** file
 - All times are uSec per file

	WFFS	EXT4	XFS
10K	63.2 uSec	8359.4 uSec	6605.8 uSec