

Linux Plumbers Conference 2022

>> Dublin, Ireland / September 12-14, 2022



EROFS

David Anderson (dvander@google.com)



Introduction to EROFS

- EROFS is a read-only filesystem that supports deduplication and LZ4 compression.
- EROFS graduated from staging in Linux 5.4.
- EROFS consistently outperforms other filesystems in our real-world tests and in our microbenchmarks.
- EROFS supports in-place decompression. This works by storing compressed data at the end of each page, and streaming the decompressed bytes into the start of the same page. From the original whitepaper, this strategy works on more than 99% of blocks.



Partition Sizes

- With LZ4 compression enabled, EROFS reduces partition sizes by around 24% on average.
- With tuning, this can be higher - we've seen up to 45% compression on our devices.
- In absolute terms, a 3.9GB EXT4 image becomes a 2.5GB EROFS image.
- EROFS also supports uncompressed data and deduplication. In this case, it is on par with EXT4 in terms of space requirements.



OTA Packages

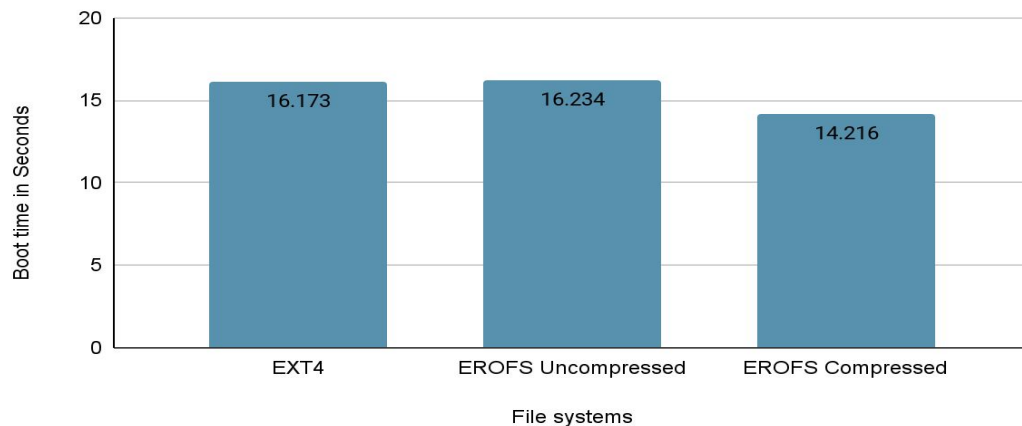
- The OTA update generator will decompress LZ4 streams to intelligently diff EROFS images.
- OTA deltas with EROFS images are roughly the same size - within 10% - as with EXT4.



Boot Time performance on Pixel 6

Boot Time

Lower is better

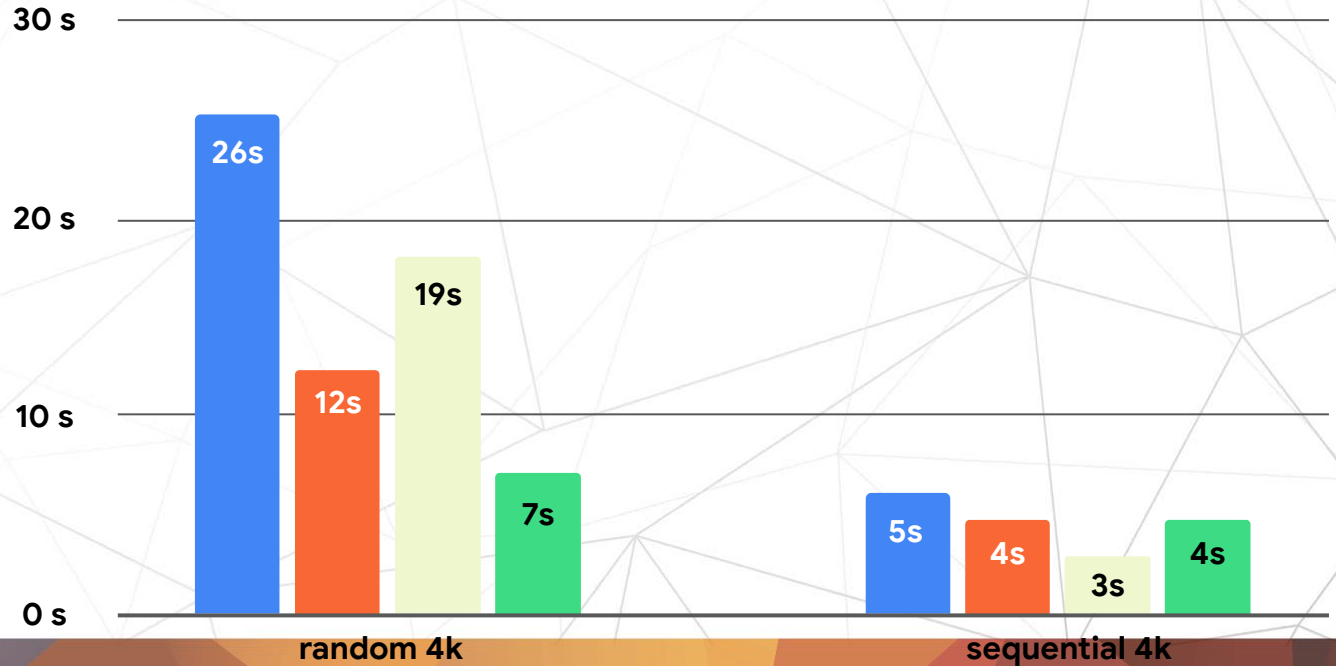




Raw Read Performance

Read Blocks Stress Test (lower is better)

- ext4
- erofs default
- erofs no-comp
- erofs pc=256k

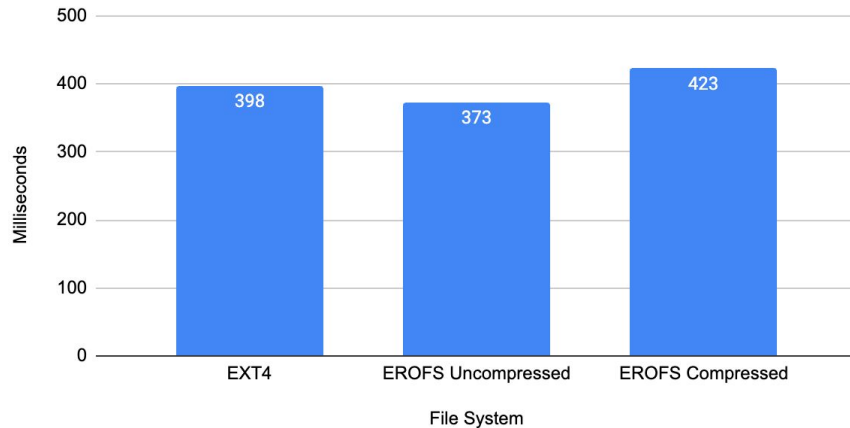




Performance on Pixel 6


Mean Photos App launch time

Lower is better



Median Photos App launch time

Lower is better



- Photos app launch time are almost similar or better with EROFS Uncompressed compared to EXT4
- With EROFS Compressed we see up to 10% performance penalty



Performance with CPU Contention

Decompression has noticeable impact during heavy load.

Scheduling latency as a factor is currently not known.





Linux
Plumbers
Conference 2022

>> Dublin, Ireland / September 12-14, 2022

Discussion