Linux Plumbers Conference 2022



Contribution ID: 42

Type: not specified

SSDFS: ZNS SSD ready file system with zero GC overhead

Wednesday, 14 September 2022 10:25 (20 minutes)

The architecture of SSDFS is the LFS file system that can: (1) exclude the GC overhead, (2) prolong NAND flash devices lifetime, (3) achieve a good performance balance even if the NAND flash device's lifetime is a priority. The fundamental concepts of SSDFS: (1) logical segment, (2) migration scheme, (3) background migration stimulation, (4) diff-on-write. Every logical block is described by {segment_id, block_index_inside_segment, length}. This concept completely excludes block mapping metadata structure updates that results in decreasing the write amplification factor. Migration scheme implies that after erase block exhaustion every update of logical block results in storing new state in the destination erase block and invalidation of logical block in the exhausted one. Regular I/O operations are capable to completely invalidate the exhausted erase block for the case of "hot" data (no necessity in GC operations). SSDFS is using the migration stimulation technique as complementary to migration scheme. It means that if some LEB is under migration then a flush thread is checking the opportunity to add some additional content into the log under commit. SSDFS is using the inline techniques to combine metadata/data pieces into one I/O request of decreasing write amplification factor. SSDFS architecture is ZNS SSD friendly and it can run efficiently even with limited number of active/open zones (14 active zones, for example). Preliminary benchmarking and estimations of conventional SSDs has showed the ability of SSDFS to decrease write amplification 2x - 10x times and prolong SSD lifetime 2x - 10x for real life use-cases comparing with other file systems (ext4, xfs, btrfs, f2fs, nilfs2).

I agree to abide by the anti-harassment policy

Yes

Primary author: DUBEYKO, Viacheslav (ByteDance)

Presenter: DUBEYKO, Viacheslav (ByteDance)

Session Classification: Zoned Storage Devices (SMR HDDs & ZNS SSDs) MC

Track Classification: LPC Microconference: Zoned Storage Devices (SMR HDDs & ZNS SSDs) MC