



Contribution ID: 130

Type: **not specified**

Improving data placement for Zoned Linux File systems

Wednesday, 14 September 2022 10:45 (20 minutes)

In this talk I'll present what I've learned from building ZenFS, a user-space RocksDB file system for zoned block devices, and what features could be transferable to kernel file systems.

I'll go over the goals and high-level design of zenfs, focusing on the extent allocator, present what performance gains we've measured[2] and what the trade-offs are when constructing a file system for zoned block devices.

Finishing up, i'd like to open a discussion on how to enable similar levels of performance in posix-compliant, general purpose file systems with zone support. BTRFS would be a good first target but bcachefs could also benefit from this.

Unless we do data separation (separating files into different zones) we will not reap the full benefits of zoned storage.

[1] <https://github.com/westerndigitalcorporation/zenfs/>

[2] <https://www.usenix.org/conference/atc21/presentation/bjorling>

I agree to abide by the anti-harassment policy

Yes

Primary author: HOLMBERG, Hans

Presenter: HOLMBERG, Hans

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

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