



Contribution ID: 131

Type: **not specified**

Object caching on Zoned Storage

Wednesday, 14 September 2022 11:05 (20 minutes)

Object caching is a great use case for SSDs but it comes with a big device write amplification penalty - often as much as 50% [1] of the capacity is reserved for over-provisioning to reduce the impact of this on conventional SSDs.

There is a great opportunity to address this problem using zoned storage, as the garbage collection can be co-designed with the cache eviction policy.

Objects stored in flash caches have a limited life time and a common approach to eviction is to simply throw out the oldest objects in the cache to free space. Conventional drives have no notion of how old objects are and are not allowed to just throw out data out of erase units on the drive to reclaim space. If the garbage collection of the drive data is done cooperatively with a ZNS cache FTL on the host however, objects can be chosen to be evicted instead of relocated when space is reclaimed.

To get there, we will need an ZNS cache FTL and an interface between the FTL and the cache implementation for indicating which LBAS that should be re-located or invalidated to minimize write amplification.

How could this be implemented? What options do we have? A ZNS Cache userspace library? A cache block device?

The user/integration point of this I have in mind would be Cachelib [2], what other potential users do we have?

This is a great opportunity to work together on a common solution for several use cases and vendors, pushing the eco-system forward!

[1] <https://research.facebook.com/file/298328535465775/Kangaroo-Caching-Billions-of-Tiny-Objects-on-Flash.pdf>

[2] <https://github.com/facebook/CacheLib>

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