



Contribution ID: 206

Type: **not specified**

Filesystem atomic writes / O_ATOMIC

Wednesday 11 September 2019 10:40 (15 minutes)

seems like the patches proposed by Fusion-io devs for general O_ATOMIC support within Linux kernel are in stand-by since 6 years.. – any plans to address it ?. What is the main reason to not guarantee atomicity of O_DIRECT writes on flash drives? – seems like most of flash storage vendors are able to provide atomic writes support on HW level, and just SW level (kernel/FS/etc.) is missed.. The main benefit for MySQL/InnoDB is to get a rid of “double write” to protect from data corruption (partially written pages) – so, every page is written twice, increasing IO write traffic + doubling page write latency + reducing by half flash drive life expectation..

I agree to abide by the anti-harassment policy

Yes

Primary author: KRAVTCHUK, Dimitri

Presenter: KRAVTCHUK, Dimitri

Session Classification: Databases MC