



Contribution ID: 65

Type: **not specified**

Enabling Large Block Size devices in Linux

Tuesday, 14 November 2023 10:15 (45 minutes)

Increasing block sizes in storage devices will be one of the keys to support larger capacities, more density, and higher cost-effective SSDs in the future. Although R&D on this topic has been discussed in the Linux community for 16 years recent advances in Linux are making support for larger block sizes more easily attainable and we may soon be able to start leveraging support for it.

512 bytes logical/physical sectors were the de-facto standard for a long time in the storage industry. Considering the rate at which the storage density was increasing, the industry started to realize 512 bytes sector is too small and settled on 4096 bytes (4k) sector as the new standard. Known technical historic advocacy for supporting large block sizes have been reducing fsck times, reducing IO for larger data sizes (writing 1 TB requires 256 million 4k IOs today on x86), cross architecture compatibility (mounting a filesystem with block size > 4k on x86), but more recently there has also been interest to leverage the ability to do larger atomic writes to reduce database latencies due to large IO journaling requirements.

This talk explores the ongoing challenges and effort to support LBS greater than the page size in Linux based on recent advancements with folio adoption in the page cache and with iomap. This talk will also discuss existing known potential advantages of large LBS devices to enable a new generation of higher cost-effective SSDs with larger capacities.

Primary author: CHAMBERLAIN, Luis (Samsung)

Co-author: RAGHAV, Pankaj (Samsung)

Presenters: CHAMBERLAIN, Luis (Samsung); RAGHAV, Pankaj (Samsung)

Session Classification: LPC Refereed Track

Track Classification: LPC Refereed Track