

Linux Plumbers Conference 2024



Contribution ID: 73

Type: **not specified**

Data Placement at Scale: Landscape, Trade-Offs, and Direction

Wednesday, 18 September 2024 12:45 (45 minutes)

Data Placement has been one of the major sources of innovation in storage. Specifically in NAND Memory, technologies such as Open-Channel SSDs, Key-Value SSDs, Multi Stream, Zoned Namespaces (ZNS), and lately Flexible Data Placement (FDP) have attempted at covering different use-cases. While there is overlap among several of these technologies, they exhibit significant differences when it comes to the way they can be adopted at the system level.

In this talk, we will talk about the evolution of these data placement technologies with a focus on the Linux ecosystem support, and the big impact that vertical system integration is having in the adoption of these technologies in enterprise and hyper-scale environments. Specifically, we will cover: (i) read/write model, where we will detail the changes needed in the I/O path for each technology and the effects this has in the OS; (ii) guarantees to reduce Write Amplification (WAF), where we will cover how each technology is able to reduce end-to-end WAF in different File-System and Application setups; and (iii) ecosystem complexity, where we will comment on the Kernel, library, and application modifications that each technology imposes. Here, we will focus on the ongoing work to make Linux ready for FDP NVMe SSDs. In the process, we will provide our vision on where each technology fits in the NVMe landscape when it comes to enable different types of memory.

Primary author: GONZALEZ, Javier (Samsung Electronics)

Presenter: GONZALEZ, Javier (Samsung Electronics)

Session Classification: LPC Refereed Track

Track Classification: LPC Refereed Track