

EXT4 and ZFS Filesystems and Their Countermeasures

- ▶ Author: Lars Thoms
- ▶ Type: Master's Thesis
- ▶ Date: 2020-11-27
- ▶ Reviewers: Prof. Dr. Thomas Ludwig, Jun.-Prof. Dr. Michael Kuhn
- ▶ Supervisors: Jun.-Prof. Dr. Michael Kuhn
- ▶ Download: PDF

Filesystems, like almost everything, are subject to an aging process. The effects of this are critical in terms of performance. In most cases, depending on the specific filesystem, long-term use leads to a massive fragmentation of the entire storage system. This master thesis quantifies these performance losses using the benchmark tool IOR and compares common countermeasures. For this purpose, different storage media and filesystems (EXT4 and ZFS) are artificially aged using Geriatrix. Afterward, they are defragmented using various methods like the defragmentation tool of EXT4 or several copy mechanisms: file-based by using `rsync` or streamed with ZFS standard tools. After each step, the files' fragmentation is determined, and the respective read and write performance. For this purpose, a new tool called `fraggy` was developed, which counts the filesystem's file fragments by using `fiemap`. Additionally, a patch extended ZFS to offer the `fiemap` interface. Furthermore, Geriatrix had to be modified to operate on ZFS. The results of this thesis show that aging effects occur in the form of fragmentation. The direct consequences are losses in the read and write performance of new data sets. The defragmentation tool of EXT4 is not suitable for long-term use and worsens the result even more. The best solution is actually to copy the data to a new filesystem. However, file-based transfers are not suitable for ZFS; the better solution is to transfer a complete snapshot to a new pool via its standard streaming tools.