

P01 - IRON File Systems

Review

Aiman Priester

Summary

- The research paper discusses an extended implementation of an already existing filesystem. The IRON filesystem boasts added features and notifications for the end user in lieu of filesystem errors. It builds off ext3 and completely removes the Dzero detection style which was found to be not effective and essentially adds error codes in replacement. A stand out added implementation is the Metadata replication. This allows for quicker data consistency check via checksumming the files. The paper ends by showcasing its increased performance in different data storage use cases.

Strengths

- The IRON file system implementation would allow for more checks, which in turn will cause the errors to be detected and/or resolved in a quicker and more efficient manner.
- Letting the user know via propagation will avoid long term issues if the user deems the risk of continued use of that particular drive is low.

Weaknesses

- Storage size of drive for the ixt3 filesystem would need to be much larger than its counterpart to handle the metadata replication.
- For large scale data solutions, the cost per GB would increase enough that it might be cheaper for corporations to carry the risk of the less “checked” ext3 filesystem.

Unsolved issues in paper

- It is acknowledged in 4.1 that not every code path was stress tested. With additional testing the possibility of uncovering faults in the logic of ixt3 might be uncovered.
- This filesystem does not account for mechanical errors within the drive itself. At what point does the filesystem know if the underlying issue is infact a mechanical one? Would the metadata be overwritten twice and cause checksumming to pass when it shouldn't?

Discussion

- One of the main issues with ixt3 is its overhead. More space is required to implement this method. RAID is mentioned in the paper. Would the combination of ixt3 and RAID, specifically RAID 0, become the best possible solution for the best of both worlds?