

Tutorial: File System Internals

Presenter: Ahmed Amer (University of Pittsburgh)

The filesystem is probably the most prevalent and unobtrusive abstraction in computer systems but, going beyond the familiar interface that allows so many uses, we will delve into the internals of filesystem implementations. The interface to the filesystem can be used to interact with anything from persistent storage to running processes or physical devices, and for this tutorial we will be focused on the filesystem's role as a means to store and access data. For this tutorial we will cover the internals of a varied range of example systems, running the gamut from basic filesystems that do little more than provide a means to access named files, to more complex systems. Our goal will be to understand the tradeoffs and decisions behind different optimizations and architectures. We will cover typical filesystems such as FFS, ext3, NTFS, HFS+, and NTFS, but we will also discuss the design choices and decisions behind a number of distributed and cluster filesystems, concluding with a discussion of recent research in filesystems.