

CSE 451: Operating Systems

Spring 2017

Module 16

Berkeley Log-Structured File System

John Zahorjan


LFS inspiration

- Suppose, instead, what you wrote to disk was a log of changes made to files
 - log includes modified data blocks and modified metadata blocks
 - buffer a huge block (“segment”) in memory – 512K or 1M
 - when full, write it to disk in one efficient contiguous transfer
 - right away, you’ve decreased seeks by a factor of $1\text{M}/4\text{K} = 250$
- So the disk contains a single big long log of changes, consisting of threaded segments
- Reminds you of journaling?
 - Yes, except that there is no “home location” for data or metadata
 - All there is is the log



LFS basic approach

- Use the disk as a *log*
 - A log is a data structure that is written only at one end
- If the disk were managed as a log, there would be effectively no seeks
- If you write enough data at a time, you can come closer to the transfer speed of the device
- New data and metadata (i-nodes, directories) are accumulated in the buffer cache, then written all at once in large blocks (e.g., segments of .5M or 1M)
 - If you write enough data at once, you can achieve closer to the transfer rate of the device
- Sounds simple – but complicated under the covers

LFS vs. UNIX File System or FFS



Unix File System



Log-Structured File System

Blocks written to create two 1-block files: dir1/file1 and dir2/file2, in UFS and LFS

LFS Challenges

- Locating data written in the log
 - FS/FFS place inodes in a well-known location, LFS writes data “at the end of the log”
- Even locating i-nodes!
 - In LFS, i-nodes too go into the log!
- Managing free space on the disk
 - Disk is finite, and therefore log must be finite
 - So cannot just keep appending to log, ad infinitum!
 - need to recover deleted blocks in old part of log
 - need to fill holes created by recovered blocks
- (Note: Reads are the same as FS/FFS once you find the i-node – and writes are a ton faster!)

LFS: Locating data and i-nodes

- LFS uses i-nodes to locate data blocks, just like FS/FFS
- LFS appends i-nodes to end of log, just like data
 - makes them hard to find
- Solution:
 - use another level of indirection: **i-node maps**
 - i-node maps map i-node #s to i-node location
 - so how do you find the i-node map?
 - after all, changes to it must be appended to the log
 - location of i-node map blocks are kept in a **checkpoint region**
 - checkpoint region has a fixed location
 - cache i-node maps in memory for performance

LFS: File reads and writes

- Reads are no different than in FS/FFS, once we find the i-node for the file
 - The i-node map, which is cached in memory, gets you to the i-node, which gets you to the blocks
- Every write causes new blocks to be added to the tail end of the current “segment buffer” in memory
 - When the segment is full, it’s written to disk

LFS: Free space management

- Writing segments to the log eats up disk space
- Over time, segments in the log become fragmented as we replace old blocks of files with new blocks
 - i-nodes no longer point to blocks, but those blocks still occupy their space in the log
 - Imagine modifying a single block of a file, over and over again – eventually this would chew up the entire disk!
- Solution: Garbage-collect segments with little “live” data and recover the disk space

LFS: Segment cleaning

- Log is divided into (large) segments
- Segments are threaded on disk
 - segments can be anywhere
- Reclaim space by cleaning segments
 - read segment
 - copy live data to end of log
 - now have free segment you can reuse!
- Cleaning is an issue
 - costly overhead, when do you do it?
- A cleaner daemon cleans old segments, based on
 - utilization: how much is to be gained by cleaning?
 - age: how likely is the segment to change soon?

LFS summary

- As caches get big, most reads will be satisfied from the cache
- No matter how you cache write operations, though, they are eventually going to have to get back to disk
- Thus, most disk traffic will be write traffic
- If you eventually put blocks (i-nodes, file content blocks) back where they came from, then even if you schedule disk writes cleverly, there's still going to be a lot of head movement (which dominates disk performance)

- Suppose, instead, what you wrote to disk was a log of changes made to files
 - log includes modified data blocks and modified metadata blocks
 - buffer a huge block (“segment”) in memory – 512K or 1M
 - when full, write it to disk in one efficient contiguous transfer
 - right away, you’ve decreased seeks by a factor of $1\text{M}/4\text{K} = 250$
- So the disk is just one big long log, consisting of threaded segments

- What happens when a crash occurs?
 - you lose some work
 - but the log that's on disk represents a consistent view of the file system at some instant in time
- Suppose you have to read a file?
 - once you find its current i-node, you're fine
 - i-node maps provide a level of indirection that makes this possible
 - details aren't that important

- How do you prevent overflowing the disk (because the log just keeps on growing)?
 - segment cleaner coalesces the active blocks from multiple old log segments into a new log segment, freeing the old log segments for re-use
 - Again, the details aren't that important

Tradeoffs

- LFS wins, relative to FFS
 - metadata-heavy workloads
 - small file writes
 - deletes

(metadata requires an additional write, and FFS does this synchronously)
- LFS loses, relative to FFS
 - many files are partially over-written in random order
 - file gets splayed throughout the log
- LFS vs. JFS
 - JFS is “robust” like LFS, but data must eventually be written back “where it came from” so disk bandwidth is still an issue

LFS history

- Designed by Mendel Rosenblum and his advisor John Ousterhout at Berkeley in 1991
 - Rosenblum went on to become a Stanford professor and to co-found VMware, so even if this wasn't his finest hour, he's OK
- Ex-Berkeley student Margo Seltzer (faculty at Harvard) published a 1995 paper comparing and contrasting LFS with conventional FFS, and claiming poor LFS performance in some realistic circumstances
- Ousterhout published a "Critique of Seltzer's LFS Measurements," rebutting her arguments
- Seltzer published "A Response to Ousterhout's Critique of LFS Measurements," rebutting the rebuttal
- Ousterhout published "A Response to Seltzer's Response," rebutting the rebuttal of the rebuttal

- Moral of the story
 - If you're going to do OS research, you need a thick skin
 - *Very* difficult to predict how a FS will be used
 - So it's hard to generate reasonable benchmarks, let alone a reasonable FS design
 - *Very* difficult to measure a FS in practice
 - depends on a HUGE number of parameters, involving both workload and hardware architecture