



# Efficient Data Retrieval from a Secure, Durable, Append-Only Log

Paul Bramsen, Andrew Chen, Sam Kumar, and John D. Kubiatowicz

## Motivation

- The Global Data Plane (GDP) is a service that provides secure data storage and distribution over untrusted hardware
  - The basic primitive provided by the GDP is a secure, single-writer log
  - Allows replication with relatively little effort
- Providing complex abstractions on top of the GDP is nontrivial
  - Tradeoff between strong guarantees and rich semantics
- Is it feasible to build a system on top of the GDP that provides mutable data semantics and efficient data retrieval?

## The Global Data Plane File System (GDPFS)

- The GDPFS stores all data in GDP logs
  - Each file lives within its own GDP log
  - Each log entry stores some metadata such as the offset and size of that entry's data
- GDP log metadata includes signer's public key and a log's name is the hash of its metadata
- The integrity of an entire filesystem can be verified given only the name of the root log

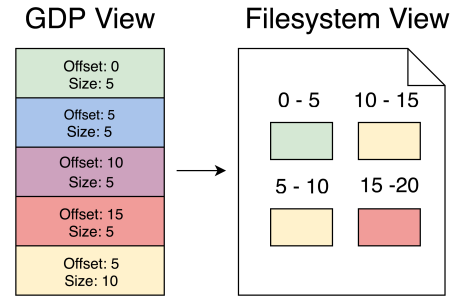
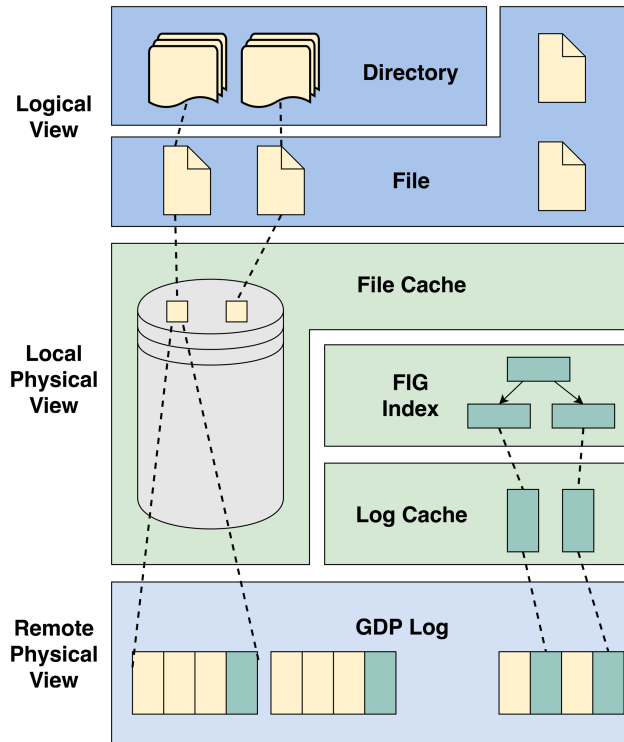


Figure 1: Basic file in the GDPFS

## Caching Strategies and Optimizations

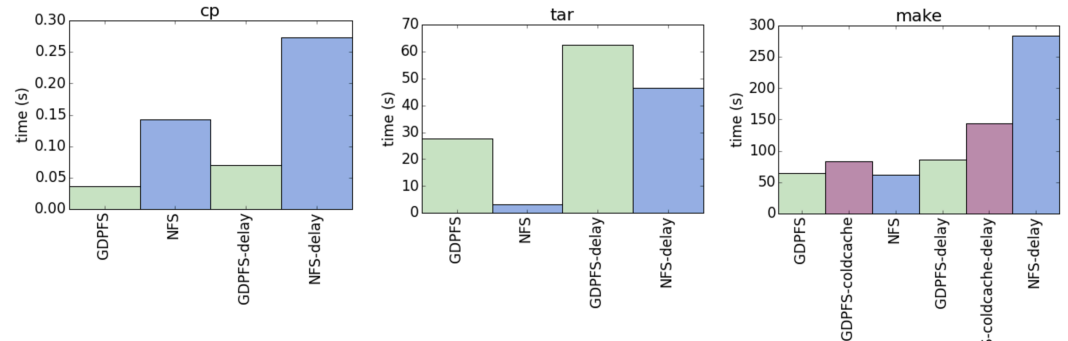
- Cache files on the local filesystem to eliminate network overhead and file reconstruction costs
- Periodically checkpoint files with a File Indexed Group (FIG) Tree to allow efficient fetching of relevant log entries on reads
- Locally cache portions of FIG Trees so that the index can usually be accessed without hitting the network

Figure 2: GDPFS structure



## Performance Results

Figure 3: Redis cp/tar/make



- Benchmarked GDPFS against NFS on workload compiling Redis
  - Redis is a popular key-value store
  - Comprised of 3 steps: cp, tar, make
- AFS benchmark also run (results omitted due to space constraints)
- Tested with no added latency as well as 10±2ms simulated latency
- Tested make after wiping cache to test effectiveness of indices

## Related Work

- Global Data Plane (Log-Based Data Substrate)
- Kafka Messaging Service (Log-Based Messaging)
- Ocean Store (Similarity in Index)
- Log-Structured File System
- B-Tree File System
- Tahoe File System (Secure Distributed File System)

