

## **ShardFS vs. IndexFS: Replication vs. Caching Strategies for Distributed Metadata Management in Cloud Storage Systems**

Lin Xiao, Kai Ren, Qing Zheng, Garth Gibson  
{lxiao, kair, zhengq, garth}@cs.cmu.edu

CMU-PDL-15-104

April 2015

**Parallel Data Laboratory**  
Carnegie Mellon University  
Pittsburgh, PA 15213-3890

**Acknowledgements:** This research was supported in part by the National Science Foundation under awards CNS-1042537 and CNS-1042543 (PROBE, [www.nmc-probe.org](http://www.nmc-probe.org)), the DOE and Los Alamos National Laboratory, under contract number DE-AC52-06NA25396 subcontracts 161465 and 153593 (IRHPIT), the Qatar National Research Foundation, under award 09-1116-1-172, a Symantec research labs graduate fellowship, and Intel as part of the Intel Science and Technology Center for Cloud Computing (ISTC-CC). We thank the DoD and Sterling Foster for support and problem formulation. We thank Robert Chansler, LinkedIn and Yahoo! for giving us a trace of its HDFS metadata workload. We thank Sage Weil and Plamen Jeliakov for helping us run and understand experiments on Ceph and Giraffa. We thank the members and companies of the PDL Consortium: Actifio, American Power Conversion, EMC Corporation, Facebook, Google, Hewlett-Packard Labs, Hitachi, Huawei Technologies Co., Intel Corporation, Microsoft Research, NetApp, Inc., Oracle Corporation, Panasas, Samsung Information Systems America, Seagate Technology, Symantec Corporation, VMware, Inc., and Western Digital for their interest, insights, feedback, and support.

**Keywords:** cloud storage, metadata management, caching, replication

## Abstract

*The rapid growth of cloud storage systems calls for fast and scalable namespace processing. While few commercial file systems offer anything better than federating individually non-scalable namespace servers, a recent academic file system, IndexFS, demonstrates scalable namespace processing based on client caching of directory entries and permissions (directory lookup state) with no per-client state in servers. In this tech report we explore explicit replication of directory lookup state in all servers as an alternative to caching this information in all clients. Both eliminate most repeated RPCs to different servers in order to resolve hierarchical permissions tests. Our realization for server replicated directory lookup state, ShardFS, employs a novel file system specific hybrid optimistic and pessimistic concurrency control favoring single object transactions over distributed transactions. Our experimentation suggests that if directory lookup state mutation is a fixed fraction of operations (strong scaling for metadata), server replication does not scale as well as client caching, but if directory lookup state mutation is constant as workload scales (weak scaling for metadata), then server replication can scale more linearly than client caching and provide lower 70 percentile response times as well.*

# 1 Introduction

Modern distributed storage systems commonly use an architecture that decouples metadata access from file read and write operations [20, 24, 44]. In such a file system, clients acquire permission and file location information from metadata servers before accessing file contents. Popular new distributed file systems such as HDFS [24] and the first generation of Google file system [20] have used centralized single-node metadata services and focused on scaling only the data path. However, the single-node metadata server design limits the scalability of the file system in terms of the number of objects stored and concurrent accesses to the file system [38]. Federating independent metadata services employs multiple server nodes but does not ensure load balancing among them. Moreover, all too often the data-to-metadata ratio is not high, since most files even in large file system installations are small [15, 43, 23]. Scalable storage systems should expect the numbers of small files to soon achieve and exceed billions, a known challenge for many existing distributed file systems [30]. While distributed file system metadata servers perform many functions (such as block management and data server fail-over), our focus in this tech report is on the management of the file system namespace.

Hierarchical file systems traditionally enforce transactional verification of hierarchical access control permission on any pathname argument; that is, access to an object specified by a pathname requires permission to lookup that object's name in its parent directory, and permission to lookup that parent directory's name in the grandparent directory, recursing as specified by the pathname back to either the file system's root or a directory currently open in the caller's process. We call this information directory lookup state. The multi-lookup resolution of a pathname is traditionally serializable with respect to all file system operations on pathnames (not necessarily for commands within open files such as data access) and is the basis of most users' understanding of file system access control. Because absolute pathnames all share the same root directory, pathname operations conflict most of the time, so pathname lookup is an important limitation on scaling the performance of distributed metadata processing in scalable cloud file systems.

In this tech report, we examine three design approaches that reduce pathname lookup overheads differently, including *replicated directories with sharded files*, *dynamically partitioned namespace with client caching*, and *table partitioned namespace*. The systems implemented these approaches are ShardFS, IndexFS and Giraffa respectively. They either replicate directory lookup state in servers or use coherent lease to cache the information in clients. These design approaches distribute metadata in different ways. Giraffa stores file system metadata as entries in a table. Metadata distribution and load balancing is delegated to the underlying table system, HBase [19]. Unlike Giraffa, ShardFS and IndexFS manages metadata distribution themselves by hashing or dynamic directory allocation. ShardFS replicates directory lookup states to pathname lookups. ShardFS reduces replication cost by specializing distributed transaction implementation based on categorizing metadata operations. Alternatively, IndexFS [35] follows traditional file system design and partitions the entire namespace with dynamically balanced directory distribution. To reduce the number of pathname lookups on metadata servers, IndexFS uses coherent client caching of pathname lookup directory attributes. IndexFS's clients observe higher latency due to cache misses on some directories while ShardFS's clients experience slow directory mutation distributed transactions. These systems are compared across several dimensions such as load balance and RPC amplification.

There are two common notions of scalability measurements: strong scalability and weak scalability. Strong scalability measures how well a system scales with increasing processing resources for a fixed workload, while weak scalability measures the scalability with a fixed workload per unit of resource. They are well defined for data dominant workloads. Consider, however, the impact of a weak scaling experiment on metadata workload. For example, a checkpoint workload may create a file for each processing unit in one directory, which produces the same number of directory operations independent of the system size. We define weak scalability for metadata workloads as linear speedup with increasing amount of processing resources under a fixed rate of directory operations but an increasing rate of file operations.

This tech report makes four contributions. The first contribution is a characterization of metadata workloads in cloud storage systems. Secondly, we analyze and pin-point the pathname lookup trade-offs in metadata service design. We compare the scale-out client caching and server-side replicated states for pathname lookups. Thirdly, ShardFS’s novel design uses directory lookup state replication and optimizes directory operations with optimistic concurrency control schemes to efficiently scale workloads that are intensive with file metadata operations. The code of ShardFS will be released here [37]. Finally, we propose definitions of strong and weak scalability for metadata workloads and devise a benchmark suites to systematically compare scalable metadata processing systems.

In the rest of this tech report, a detailed comparison of these approaches is presented. We evaluate design choices made by these systems using a set of micro-benchmarks and trace replay tests on a 256-node cluster.




Figure 1: Typical distributed file system architecture.

## 2 Background

Before comparing design of scalable metadata service, we present the overall file system architecture and the design requirements. We also examine file system traces generated by production clusters to demonstrate characteristics of real metadata workloads. Then we briefly discussed previous works on distributed metadata management.

### 2.1 System Architecture

The architecture used by both design approaches is similar to many modern distributed file systems [22, 20, 24]. It decouples metadata from data management by having separate metadata servers and data servers, as shown in Figure 1. In such an architecture, metadata servers are responsible for managing the file system namespace and metadata associated with files and directories. Clients first communicate with metadata servers to locate data and then subsequently interact directly with data servers. This allows efficient parallel data transfers between a large number of clients and data servers. To simplify the implementation and design, previous file systems mostly use only a centralized single-node metadata server with backup servers in case of failure, or statically partition the namespace, suffering when a workload applies all work to one server.

In this tech report, we focus on scaling out metadata management by distributing metadata across many servers. We assume that clients and servers are always connected, which represents the scenario inside a datacenter.

### 2.2 Metadata Workload Analysis

We analyze file system traces obtained from previously published work [43, 28] or collected from our colleagues in industry. These storage systems include desktop filesystems from Microsoft engineers [28],

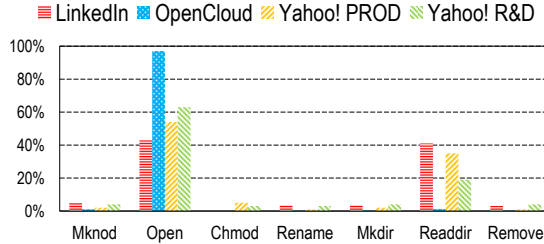


Figure 2: *Distribution of file system operations in LinkedIn, Yahoo!, and OpenCloud traces.*

commercial HPC customer installations from Panasas [43], and several Hadoop clusters (from Yahoo!, Facebook, LinkedIn, and our campus Hadoop cluster). The size of storage systems in the traces ranges from one machine to thousands of machines.

**File System Namespace Structure** File system namespaces exhibit a heavy-tailed distribution of directory sizes. Typically, a file system consists of many small directories and relatively few large directories. About 60 out of 70 file systems in our traces have nearly 90% of directories with size fewer than 128 entries. On the other hand, large directories continue to grow with the increasing size of the storage system. Many parallel computing applications (e.g. map-reduce, checkpoint) require concurrent access to large directories [4, 16]. File size has a similar distribution: about 64% of traced file systems have median file size smaller than 64KB, while the largest file size can be several terabytes. We found that the average file system tree depth does not increase much as the system grows: most file systems have 90% of their directories with a depth of 16 or less. Thus, a scalable metadata service should support a large amount of small objects, concurrent access to large directories, and can perhaps expect a short depth pathname.

**Skewed Distribution of Operations** We also studied the dynamic statistics of four Hadoop clusters: LinkedIn, Yahoo! Research, Yahoo! Production and our campus Hadoop cluster. The number of individual file system metadata operations are counted over the entire traces. Figure 2 shows the fraction of each type of metadata operation. In all traces one or two operations dominate. *open* is the most prevalent metadata operation, and the next is *readdir*. The skewed distribution of file operation types motivates us to explore design choices that trade the performance of some metadata operations for that of others.

### 2.3 Previous Work

Many popular distributed file systems such as HDFS [24] and the first generation of Google file system [20] store file metadata on a single server with limited scalability. Although the Lustre [27] parallel file system supports multiple metadata servers, only one of them can be active. Federated HDFS and PanFS [24, 44] further scale metadata management by using an approach where individually non-scalable namespace servers are responsible for an administrator selected subtree of the overall namespace. Such solutions lack the flexibility to automatically balance load across multiple metadata servers.

A few distributed file systems have added support for fully distributed file system namespace. PVFS [10] and Ceph [42] dynamically distribute the file system namespace at the granularity of a directory. Ceph uses an adaptive partitioning technique to split large directories, though users are advised to disable this feature when using multiple metadata servers [12].

The distributed directory service [17] in Farsite [3] uses a tree-structured file identifier for each file. It partitions the metadata based on the prefix of a file identifier, which simplifies the implementation of rename operations. Farsite employs field-level leases and a mechanism called a disjunctive lease to reduce false sharing of metadata across clients and mitigate metadata hotspots. However, it maintains state about

the owner of each lease at the server in order to later invalidate the lease and is not designed for highly concurrent mutations.

Both Frangipani [40] and GPFS [36] are built on shared disks. Frangipani uses distributed coarse-grained locks on entire files. GPFS uses token-based distributed fine-grained locking on ranges in files for high throughput and strong consistency among clients. GPFS uses whole directory locks to maintain consistency of inodes in a directory.

CalvinFS [41] leverages a distributed database system for metadata management. With full pathname as the key, file metadata are partitioned by hash of the pathname and replicated across a (shared-nothing) cluster of independent servers. To optimize for single-file operations (read and write), all permissions from the root are stored with each file or directory. A directory's files and subdirectories are stored as its value for *readdir*. This makes file creation a distributed transaction and hence does not scale well. Other metadata operations such as directory rename and permission changes need to recursively modify all entries under the affected subtree.

Transaction management used in both ShardFS and IndexFS share some similarities with database systems. The local storage backend LevelDB used by ShardFS and IndexFS use the multiversion concurrency control protocol (MVCC [6]) to manage concurrent data access. H-store [39] is a modern main memory databases that uses a sophisticated concurrency control strategy. H-store classifies transactions into multiple classes where single-sited or sterile transactions are executed without locks. For non-sterile and non-single-sited transactions, H-store uses an optimistic method to execute distributed transactions and dynamically escalates to more sophisticated strategy when the percentage of aborted transactions exceeds certain threshold. ShardFS follows a similar design methodology.

Object storage systems such as CohoData [14], Nutanix and Microsoft Azure [9] target different problems than metadata performance. They focus on providing high data throughput with little or no hierarchy in the namespace of objects. The address mapping for objects is usually stored in systems such as Zookeeper [25] or Cassandra [11].

### 3 Comparison of System Designs

Our comparison of file system metadata designs focuses on key factors that affect system performance such as throughput and latency of metadata operations. We choose to compare designs that dynamically distribute metadata with fine granularity. That's because static subtree partitioning of the namespace cannot provide good load balancing when the workload only accesses a small portion of the namespace. IndexFS dynamically assign directories to servers and splits large directories, which makes it susceptible to load imbalance only on small directories. ShardFS partitions at the file level to accommodate hot directories. Giraffa relies on the underlying distributed table for metadata distribution.

As indicated in the introduction, pathname lookup is an important factor restricting the scaling of distributed metadata management. Traditional file systems scale pathname resolution by coherently caching namespace information (structure, names, permissions) in each client under the protection of a (leased) lock. Our representative traditional file system, IndexFS, implements coherent client caching this way and simplifies server error handling logic by blocking all mutations until all leases have expired. Caching state under coherent leases is a replication strategy with replication costs proportional to the number of clients and to the size of the working set for the cache (number of directories frequently consulted in pathname lookup). Since the number of clients is rarely less than the number of servers, fully replicating the same information in each server can in principal be a less costly way to avoid multiple RPCs and locking congestion across multiple metadata servers. ShardFS is the example system of using full replication of directory lookup states. An alternative approach taken by Giraffa is to relax file system access control semantics and store the full pathname for each entry to reduce the number of lookups.

In the following sections, we explore the viability of server replication of namespace information, and discuss the performance implication brought by these three design choices.

### 3.1 Dynamically partitioned namespace (IndexFS)

In IndexFS, the namespace is partitioned dynamically at the directory subset granularity. As shown in Figure 3, each directory is assigned to an initial, randomly selected metadata server when it is created. The directory entries of all files in that directory are initially stored in the same server. This works well for small directories since storing directory entries together can preserve locality for scan operations such as *readdir*.

For the few directories that grow to a very large number of entries, IndexFS uses the GIGA+ binary splitting technique to distribute directory entries over multiple servers [30]. Each directory entry is hashed into a large hash-space that is range partitioned. A directory gets split incrementally in proportion to its size: it starts small on a single server that manages its entire hash-range. As the directory grows, IndexFS splits the hash-range into halves and assigns the second half of the hash-range to another metadata server. As these hash-ranges gain more directory entries, they can be further split until the directory expands to use all metadata servers. This splitting stops after the distributed directory is well balanced on all servers.

**Metadata Caching** POSIX semantics require many metadata operations to perform pathname traversal and permission checking across each ancestor directory. This access pattern is not well balanced across metadata servers because pathname components near the top of the file namespace tree are accessed much more frequently than those lower in the tree. IndexFS maintains a consistent cache of pathname components and their permissions (called the directory entry cache). IndexFS avoids invalidation storms by assigning short term leases to each pathname component accessed and delaying any modification until the largest lease expires. This allows IndexFS servers to only record the largest lease expiration time with any pathname component in its memory, pinning the entry in its memory and blocking updates to the entry until all leases have expired.

Any operation that wants to modify the server’s copy of a pathname component, which is a directory entry in the IndexFS server, blocks operations that want to extend a lease (or returns a non-cacheable copy of the pathname component information) and waits for the lease to expire. While this may incur high latency for these mutation operations, client latency for non-mutation operations, memory and network resource consumptions are greatly reduced. This method assumes the clocks on all machines are synchronized, which is achieved in modern data centers [8, 13].

**Implementation** IndexFS is distributed metadata middleware, written in C++, that can be layered on top of existing file systems such as HDFS to scale their metadata performance [35]. IndexFS uses the Thrift [1] RPC library for network communication. It stores metadata into LevelDB [26] which implements a write-optimized B-tree: a Log-structured Merged tree (LSM tree) [29]. LevelDB also stores its data in HDFS where the data is replicated to tolerate failures.

**Fault Tolerance** IndexFS’s primary fault tolerance strategy is to push states into the underlying object storage e.g. HDFS, that provides replication – metadata files and recent changes in write-ahead logs. IndexFS server processes are all monitored by standby server processes prepared to replace failed server processes. Each IndexFS metadata server maintains a separate write-ahead log that records mutation operations such as file creation and rename. Leases for client caching are not made durable. A standby server restarting from logs simply waits for the largest possible timeout interval.

For operations requiring a distributed transaction protocol such as directory splitting and rename operations, two phase distributed transactions are used in IndexFS. Server failure is protected by write-ahead logging in source and destination servers and eventual garbage collection of resource orphaned by failures. Currently IndexFS only implements a simplified version of the *rename* operation that supports renaming files and leaf directories.






Figure 3: *IndexFS* distributes its namespace on a per-directory partition basis. Path resolution is conducted at clients with a consistent lease-based directory lookup cache. Directories are not partitioned until they are large and partition is done by repeated binary splitting.

### 3.2 Table partitioned namespace (Giraffa)

As shown in Figure 4, Giraffa stores file system metadata inside a distributed table, HBase [19], which provides single-row transaction guarantees. Each file or directory in Giraffa namespace is stored as one row inside a table in HBase. In order to maintain the hierarchical namespace of the file system, Giraffa embeds its file system tree structure inside the row keys of all the file system objects. The default strategy is to use as the key a full pathname prefixed with the depth (of this file system object) in the namespace tree. This ensures that all entries within a same directory can share the same prefix of their row keys, which secures the necessary locality required to implement *readdir* efficiently. Giraffa translates metadata operations into a set of key-value operations to HBase, and reuses the load balancing techniques and persistence guarantees provided by both HBase [19] and HDFS [24].

**Implementation Details** The Giraffa metadata server is implemented as a “coprocessor” embedded in each HBase region server [19], which works like a “stored procedure” in a relational database [32]. The current implementation of Giraffa relies on the underlying HBase to dynamically partition and distribute the metadata table across all its region servers to achieve load balance. By default, HBase horizontally partitions its table as regions according to the size of existing regions. Since HBase is unaware of any semantic meaning of stored table contents, it will not deliberately partition a large directory or cluster small directories as *IndexFS* does. The default split threshold for a HBase region is as large as 10GB, which is much larger than the split threshold for directories in *IndexFS*. During our experiments, we found Giraffa can easily suffer a skewed distribution in its lexicographic key space, in part due to its default schema for generating row keys. In order to mitigate this problem, we modified Giraffa’s code by prepending a hash of the parent path to the original row key, and pre-split the namespace table in HBase at the beginning of each experiment. HBase allows users to pre-split tables to help better balance the system during the initial workload, provided that the key distribution is known beforehand. In our experiments, this trick allows Giraffa to immediately distribute the key space to all region servers as soon as the system starts up, achieving static load balance without the overhead of incremental data migration.

**Relaxed Operation Semantics** Giraffa relaxes semantics of several metadata operations. For access control, Giraffa does not check the permission information of every ancestor directory when accessing a file. This reduces the number of client-server communications and helps with performance. To support POSIX access control model, it could either adopt a metadata caching technique used by *IndexFS*, or use a schema like Lazy Hybrid (LH) [7] or CalvinFS that replicate a directory’s permission bits to all files nested beneath it. The row key schema used by Giraffa also affects directory rename operation. Since the directory name is part of its children’s row keys, rename requires read-modify-write on all of its children files. Because of the difficulty of supporting atomic rename of directories, it only supports rename of files in the same directory.




Figure 4: *Giraffa* stores its metadata in HBase, which partitions its table as a B-Tree. Each file or directory is mapped to an unique row at one HBase region server. The current implementation of *Giraffa* does not have hierarchical permission checking so no pathname resolution is performed.




Figure 5: *ShardFS* replicates directory lookup state to all metadata servers so every server can perform path resolution locally. File metadata and non-replicated directory metadata is stored at exactly one server determined by a hash function on the full pathname.

### 3.3 Replicated directories w/sharded files (ShardFS)

ShardFS scales distributed file system metadata performance by fully replicating directory lookup state across servers, insuring each file operation on a server is single-site [39], so it need not obtain locks from multiple servers. This strategy slows down mutations that affect this information (changes to directory names, directory access permissions, or namespace parent-child structure of pairs of directories) in order to speed up and load balance accesses to the objects reached by a successful pathname lookup as shown in Figure 5. Every metadata server contains a complete replica of this namespace information, so a single RPC to the appropriate (single-site) metadata server will be sufficient to complete pathname resolution without additional RPCs. Unlike directory lookup states, file metadata is stored only in one metadata server. They are distributed by a sharding function using pathname as input. The metadata server that a pathname shards to is defined as the primary metadata server for this pathname. When a pathname represents a file, its metadata is stored only on the primary metadata server.

Due to having fully replicated directory lookup states, *ShardFS* is based on pessimistic multi-server locking for all metadata operations. However, operations accessing metadata only on one metadata server are single-site transactions and can execute with one RPC to the primary metadata server (all metadata servers are internally serializable and transactional for metadata operations). Almost all operations on file metadata or non-replicated directory metadata such as timestamps are single RPC operations. The scaling benefit that these operations get from namespace replication is the primary motivation for this technique. Other operations are distributed transactions and use a two phase locking protocol, limiting overall scalability.

#### File System Specific Optimizations

*ShardFS* strives to reduce the latency for single RPC operations on files by not blocking them on locks taken by a replicated state mutation. Specifically, using analysis similar to what was used for H-store [39], *ShardFS* uses optimistic concurrency control for single RPC operation in as many transaction classes as possible and will fall back to retry with pessimistic two phase locking concurrency control when optimistic verification fails. With the limited number and semantics of file system metadata operations, *ShardFS* does not compare the entire read and write sets for concurrent operations. Instead, file system operation semantics

are used to detect optimistic verification failures, causing single site transactions to abort and retry with full pessimistic locking.

The file system operations that are not single RPC are:

1. increasing permitted operations on replicated state: create a directory (*mkdir*) or increase permitted operations (*chmod +mode*, or *chgrp +group*)
2. decreasing permitted operations on replicated state: remove a directory (*rmdir*) or decrease permitted operations (*chmod -mode*, or *chgrp -group*)
3. increasing and decreasing permitted operations on replicated state: changes that both allow and deny operations (rename, mode changes that add and delete permissions, changes to ownership).

For the first two, monotonic changes in the set of operations that will be permitted, we design our distributed transactions so that concurrent single RPC operations that might detect inconsistent replicated states will all fail, allowing the calling client to recognize potential optimistic concurrency control failure and retry with pessimistic locking.

For the third class of operations, non-monotonic changes in the set of operations that will be permitted, we design our distributed transactions to be maximally defensive—they serialize all operations (even optimistic single RPC operation) with conflicting scope by taking locks at every server. As slow as these operations can be, they are quite rare in our traces and do not significantly impact performance results.

For multi-server transactions causing monotonic changes in permitted operations, ShardFS first serializes distributed transactions with conflicting scope. This could be done by taking locks on any single metadata server, but we instead use a separate server because it will also serve as the redo log server for clients that fail during a distributed transaction. This ensures that among all concurrent transactions with conflicting scope there is either one non-monotonic distributed transaction running serially or one monotonic distributed transaction running concurrently with multiple single RPC transactions.

Any monotonic distributed transaction changes one object of the replicated state, which may be the target object or on the pathname to the target object of a concurrent single RPC operation. Fortunately pathname lookup determines only the success or failure of each pathname component’s existence and permission check, so concurrent single RPC operations that do not fail only learn that this operation is permitted. So concurrent single RPC operations conflicting on an ancestor directory can only observe the success of their pathname lookup; this will not expose inconsistency if all single RPC operations targeted at the conflicting object cannot concurrently observe a value inconsistent with the success of the succeeding pathname lookups.

There are two cases: monotonically increasing permitted operations (type 1: *mkdir*, *chmod+mode*) and monotonically decreasing permitted operations (type 2: *rmdir*, *chmod -mode*). The procedures to implement these two types of operations are presented in Procedure 1 and 2.

---

**Procedure 1** type 1 distributed transaction *op*

---

I1: Log and take a lock on lock server

I2: Execute *op* on primary server

I3: If I2 succeeds [commit point], execute *op* in parallel on all non-primary servers

I4: Log and unlock

---

For type 1, monotonically increasing permitted operations, concurrent single RPC operations referring to the target object during pathname lookup can only see inconsistency during step I3 in Procedure 1, after the commit point, so concurrent single RPC operations targeting the same object will see successes as consistent because these targeted operations go only to the primary server, which has committed the transactions that only increase permitted operations.

---

**Procedure 2** type 2 distributed transaction *op*

---

D1: Log and take a lock on lock server

D2: In parallel execute *op* on all non-primary servers

D3: If every *op* in D2 succeeds, execute *op* on primary

D4: If every *op* in D2&D3 succeeds, continue [commit point]; otherwise rollback *op* on non-primary servers

D5: Log and unlock

---

For type 2, monotonically decreasing permitted operations, the only inconsistency possible would be between an single RPC referencing the target object only in its path, and succeeding, relative to a single RPC referencing the targeted object directly on the primary server and succeeding in some way that makes it believe some pathname reference should have failed because of the decreasing permissions done by the type 2 distributed transaction. Without loss of generality the conflicting case is a type 2 distributed transaction removing permission or a directory that other single PRC operations reference in pathname lookup. For succeeding single RPC operation referring the changing object in pathname lookup to be inconsistent they have to occur after a single RPC operation directly accessing the targeted object sees permissions as having been take away. Since such a directly targeted single RPC operation addresses the primary server for this object it cannot see the permission as taken away until step D3 finishes successfully, after which no pathname lookup can see an inconsistent replica because these were all changed in step D2.

In summary, any concurrent single RPC operation that might observe inconsistent state because of replicated pathname lookup will experience failure, then retry with pessimistic concurrency control.

**Implementation** ShardFS is implemented as a new HDFS client library written in JAVA on top of standalone IndexFS servers [33, 35] with modifications to support server side pathname resolution. Sharing the same IndexFS code base provides a fair comparison between these two approaches. The namespace replication with optimized directory lookup state mutation operations is implemented on the client side. We run a simple lock server to resolve races among all directory lookup state mutation operations. Locking a pathname is equivalent to acquire a write lock on the last component of the path and read locks on all ancestors. We also implemented a HDFS client library with HDFS’s Namenode as its metadata server. However, since a single Namenode’s performance turns out to be inferior to a single IndexFS server, we only show ShardFS on top of IndexFS servers in the evaluation section.

**Fault Tolerance** ShardFS relies on high availability of its underlying metadata servers. IndexFS tolerates failures by replicating data in its underlying storage (HDFS in our experiments) as described in section 3.1. As described before, the lock server also keeps logs of outstanding distributed directory transactions to tolerate client failures. When a client fails, later lock acquisition on the same path will trigger recovery process for the operation. To tolerate lock server failure, we can run a backup lock server. The active lock server only grants a lock after the request has been replicated to the backup lock server.

### 3.4 Comparison Summary

Table 1 summarizes the design difference among the three metadata systems. To analyze the performance implication of these designs under various workloads, we discuss a few aspects of the three systems:

**RPC amplification of metadata operations** RPC amplification is defined as the the number of RPCs sent for an operation. One major source of RPC amplifications comes from path resolution. POSIX semantics require accessing each ancestor directory on the path to check permissions. The path components may be stored on different servers with a distributed namespace, which require multiple RPCs to fetch. ShardFS replicates directory metadata to every metadata server for server local pathname resolutions, complicating the protocol of directory metadata mutation operations. In IndexFS, both clients and servers maintain a consistent cache of path attributes to reduce RPCs, which can be limited by cache effectiveness. Giraffa

	<b>Replicated directories with sharded files</b>	<b>Dynamically partitioned namespace</b>	<b>Table partitioned namespace</b>
Example system	ShardFS	IndexFS	Giraffa
Metadata distribution	Replicated dir lookup states; sharded files	Partitioned into directory subsets	Partitioned by HBase
Metadata addressing	hash(pathname)	parent dir's inode num + hash(filename)	hash(path prefix) + depth + pathname
File operation (e.g. <i>stat</i> , <i>chmod</i> )	1 RPC for all operations	Number of RPCs depends on cache locality	1 RPC for <i>stat</i> & <i>chmod</i> 2 RPCs for <i>mknod</i>
Directory metadata mutations (e.g. <i>mkdir</i> )	Optimized distribution for monotonic ops	1 RPC but waiting for lease expiration	2 RPCs, similar to file operations
Concurrency control	Optimistic locking on the central server	Locking at directory partition level	Serialized by each tablet server
Client caching	Only cache server configuration	Cache path prefix ACLs & directory partition location	Cache the location of tablets
Load balancing	Load balance file access by static hashing	Dynamically assign dirs; split large directory by size	Pre-split & dynamically split tablets by size

Table 1: Summary of design choices made by three metadata services.

abandons the access control model of POSIX by only checking permission bits of the final component. By using the full pathname as part of the primary key, most single file metadata operations in Giraffa require only one RPC.

Another source of RPC amplifications is the partitioning and replication of metadata. For directory mutation metadata operations, the ShardFS client contacts all metadata servers to execute the distributed transactions, adding to RPC amplification. For IndexFS, attribute modification operations such as *chmod* and *chown* require only one RPC but may need to wait for lease expiry. Moreover, *mkdir* and *splitting* involve two servers to perform transactions. And *rmdir* checks each partition to see if the directory is empty. The RPC amplification for *readdir* is proportional to the number of partitions and the overall size of each partition.

**Metadata operation latencies** Both Giraffa and ShardFS try to keep the latency of file metadata operations low as one round-trip RPC. In this case, latency is mostly affected by server load. For most directory metadata mutations in ShardFS, multiple RPCs are issued to all metadata servers in parallel. Thus its latency is sensitive to the slowest RPC. For IndexFS, the latency of metadata operations is affected by the hit ratio of the directory entry cache. Directory metadata mutations such as *chmod* and *chown* are also sensitive to the lease duration in IndexFS. Giraffa will have problems similar to IndexFS if it someday supports POSIX semantics for access control.

**Consistency model for metadata operations** All three metadata systems guarantee serializability for file metadata operations. The accuracy of access and modification time stamps for directories are relaxed for better performance in all three systems. *readdir* behavior under concurrency is not well defined in POSIX, which makes it flexible for system designers. All three systems provide an isolation level called “read committed” [5]. A *readdir* will always reflect a view of the directory at least as new as the beginning of the operation. It reflects all mutations committed prior to the operation, but may or may not reflect any mutation committed after the issue of the operation. This allows the systems to implement *readdir* as multiple independent requests to each directory partition without locking the entire directory.

**Load balanced across metadata servers** An important scalability factor is the balance of load across metadata servers. Overall performance is limited by the slowest server. IndexFS uses a dynamic growth policy such that a newly created directory starts with one server and is dynamically split as it grows. While this policy maintains locality for small directories, it may experience higher load variance when a directory is close to splitting. Servers containing the top of the namespace may get higher load due to client cache misses or renewal. Giraffa will experience similar problems as it also splits the tablet according to the size. Since the split is handled by HBase, a directory in Giraffa is not necessarily clustered nor well balanced across servers. ShardFS, on the contrast, maintains both capacity and load balance by sharding files and replicating directory metadata at all times. It has a more balanced load across servers, but its directory metadata mutation is slower. Currently none of the systems implement strategies for the case where a few popular files dominate a workload [18].

**Scalability** Both IndexFS and Giraffa scale in throughput as more servers are added to the system for most workloads. In contrast, ShardFS’s directory metadata mutation operations get slower with more servers due to replication. In the extreme, if the workload only contains directory metadata mutation operation, ShardFS with more servers gets slower. However, when the ratio between file and directory metadata mutation operation scales as the number of servers (see section 4.3 on weak scaling metadata workload), the performance scales in ShardFS. As the namespace grows larger, both systems face the challenges of maintaining an effective directory entry cache to avoid unnecessary RPCs and disk requests.

## 4 Evaluation

In this section, we report experiments done on IndexFS, ShardFS, and Giraffa. To run these experiments, we have used up to 256 machines partitioned into up to 128 server nodes and up to 128 client nodes. Running client and server processes on separate machines facilitates better resource isolation and simpler performance measurement. All machines are part of the NSF PROBE Kodiak cluster [21] described in Table 2.

<b>CPU</b>	2x AMD Opteron 252 (2.6GHz)
<b>Memory</b>	8GB
<b>NIC</b>	1Gbps Ethernet
<b>Storage</b>	1TB 7200 RPM HDD
<b>OS</b>	64-bit Ubuntu 12.04.5 with Linux 3.2
<b>Software</b>	Hadoop Release 2.4.1, Giraffa Release 0.0.1

Table 2: PROBE Kodiak testbed configuration.

All systems were built on top of HDFS, which served as an underlying object storage infrastructure holding three copies of file data and file system metadata logs. Giraffa additionally depended on HBase to manage its file system metadata. For IndexFS, the threshold for directory splitting (Section 3.1) was set to 2,000 entries in all tests. In addition, all systems were configured to asynchronously commit (sync) metadata log entries to the underlying storage (HDFS) every 5 seconds. All experiments were repeated at least three times and showed insignificant variance in observed performance. The benchmark framework that we used is a modified version of YCSB++ [31, 34] that supports multi-client coordination.

### 4.1 Microbenchmarks

In this section, we executed a set of microbenchmarks on the 3 targeted file systems to study the tradeoffs among their designs. To evaluate the impact of different namespace distribution techniques, we prepared 3

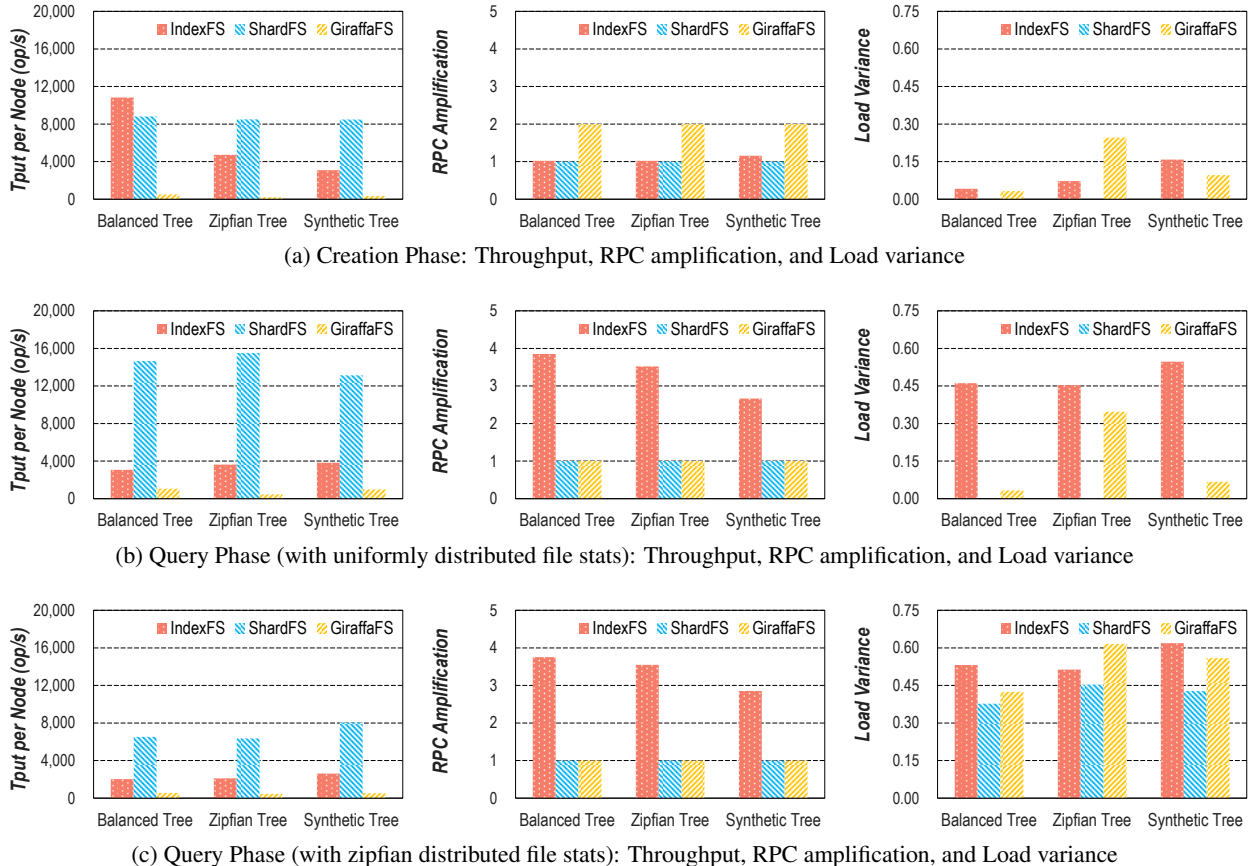


Figure 6: Performance comparison among IndexFS, ShardFS, and Giraffa creating and stating zero-byte files with 64 server machines and 64 client machines.

distinct file system images representing different namespace tree structures.

**#1. Balanced Tree:** In a balanced tree, each internal directory has 10 sub-directories, and each leaf directory has 1,280 children files. The height of the tree is 5. In total there are 111 K directories and 128 M files.

**#2. Zipfian Tree:** Zipfian trees share the same internal tree structure as a balanced tree, except that the sizes of its leaf directories are randomly generated according to a zipfian distribution with an exponent parameter of 1.8. There are 111 K directories and approximately 128 M files inside the tree.

**#3. Synthetic Tree:** This tree is generated by a workload generator named Mimesis [2]. Mimesis models existing HDFS namespace images based on several statistical characteristics, such as the distribution of the number of files per directory and the number of sub-directories per parent directory. For this tree, we used Mimesis to model and scale a HDFS trace extracted from a Yahoo! cluster. The original Yahoo! namespace had 757 K directories and 49 M files. Mimesis expanded this file system image to contain 1.9 M directories and 128 M files with the same statistical characteristics as the original Yahoo! namespace.

To microbenchmark the metadata performance of each namespace tree, we designed a simple three-phase workload. The first phase creates all internal directories. The second phase populates the namespace with empty files. During the third phase, each client performs *stat* on files randomly selected from the namespace. To model different access patterns, files to be accessed are chosen either uniformly or following a zipfian distribution with an exponent parameter of 1.8. In all cases, the number of client threads is selected to saturate the target file system at its maximal throughput (an external tuning knob for YCSB). We measured three different metrics: *average throughput per server*, *RPC amplification*, and *load variance*. The RPC

amplification is reported as the total number RPCs over the total number of application-level file system operations. The load variance is measured as the coefficient of variation of the number of RPC requests received by each metadata server. All microbenchmark experiments were run with 128 machines with 64 configured as servers and 64 as clients.

Figure 6 shows the experimental results for the file creation and query phases. In general, Giraffa appears to be much slower than both IndexFS and ShardFS. According to our profiling results, we believe the main reason lies in the less optimized code in Giraffa’s implementation, such as inefficient memory copies, overhead of communicating with HBase, as well as the use of global locks. As a result, in this section, we will mainly focus on RPC amplification and load variance when comparing with Giraffa.

In the file creation phase, IndexFS achieves its highest throughput in the balanced tree workload, for the other two workloads have a few very large directories. This gives rise to a set of hot servers performing necessary background activities to spread those large directories to multiple servers and balance the system for future operations. IndexFS also shows higher load variance in the later two workloads. This is because populating files for the zipfian and synthetic trees produces namespace lookup requests that are imbalanced by nature. Fortunately, as files are created with depth-first order preserving path locality, the RPC amplification of IndexFS during the entire file creation phase is relatively low compared to that observed in the query phase. Unlike IndexFS, the performance of ShardFS tends to be stable for all three workloads. In fact, with files uniformly distributed across all of its servers, ShardFS can often achieve good load balance for file creates regardless of the actual file system tree structure. Different from both IndexFS and ShardFS, a large directory in Giraffa can easily be held entirely by a single HBase region server and become a performance bottleneck. In addition to this vulnerability to load imbalance, Giraffa also shows higher RPC amplification. As Giraffa has to check the existence of the parent directory when creating a new file, there is an additional RPC for almost all file creation operations, for parent directories are very likely to be distributed to a remote region server according to Giraffa’s current namespace partitioning strategy. In fact, according to POSIX semantics, Giraffa should have consulted all ancestor directories before it can ever create a new file. Unfortunately, if enforced, this can only lead to even severer RPC overhead.

**Lesson#1:** ShardFS shards each directory to all servers, maximizing the insertion throughput for every directory. IndexFS incrementally partitions directories as they grow up, which can slow down the system when directories are small. Giraffa rarely splits individual directories and can suffer high load variance in practice.

For the query phase with uniform file selection, ShardFS shows excellent load balance, lower RPC amplification, and higher throughput. This is because ShardFS always distributes files evenly to all of its metadata servers and each metadata server can perform pathname lookups locally without contacting peer servers. Unlike ShardFS, IndexFS’s performance is largely limited by its RPC amplification, which can in part be attributed to this random read access pattern. These requests make IndexFS’s client-side lookup cache relatively useless, causing more cache misses and forcing IndexFS clients to frequently fetch lookup states from servers. However, since IndexFS is able to dynamically split large directories, it doesn’t get bottlenecked on large directories even with a skewed namespace such as the zipfian and synthetic trees. In fact, IndexFS performs better in these namespaces as its client-side lookup cache becomes more effective. This leads to lower RPC amplification, higher throughput, albeit higher load variance. For Giraffa, due to its inability to split large directories, its load variance is largely determined by the shape of the namespace tree. For example, it shows higher load variance when it comes to the Zipfian tree. In addition, since Giraffa does not actually perform pathname resolution like IndexFS and ShardFS, the RPC amplification for Giraffa is always one for all file stat operations. Finally, all the 3 file systems show lower performance and higher load variance when files are selected following the zipfian distribution.

**Lesson#2:** IndexFS’s file *stat* performance is mainly a function of its cache effectiveness. ShardFS is able to deliver deterministic fast file *stat* performance by replicating directory lookup state. Giraffa often suffers load imbalance even without performing pathname resolution.



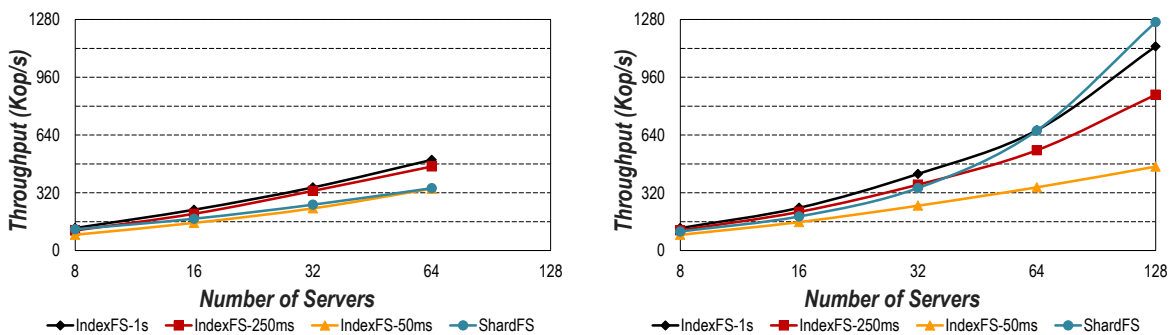


Figure 7: Aggregated throughput replaying an one-day trace collected from LinkedIn versus number of machines using either strong scaling (Section 4.2) or weak scaling (Section 4.3). Figure 7(a) on the left. Figure 7(b) on the right.

## 4.2 Trace Replay and Scaling Experiment

In this section, we compare the performance of IndexFS and ShardFS using workloads constructed from a real-world trace. In fact, we also performed experiments on Giraffa but decided not to show its results due to limited space as well as the fact that its throughput is much lower than IndexFS and ShardFS. For this macrobenchmark, we used an one-day trace collected from a LinkedIn cluster and generated a two-phase trace-replay workload. The namespace in this trace consists of 1.9 M directories and 11.4 M files. We set all files to be empty since we only focus on metadata scalability in this tech report. The first phase of the workload recreates all the files and directories referenced in the trace. In the second phase, each client independently replays a set of non-overlapping portions of the trace, which has been pre-divided into a series of sub-traces each containing 200 consecutive operations.

We used this trace-replaying workload to test the strong scaling performance of both file systems. By strong scaling, we mean that the workload remains the same as we increase the size of the system. In addition to the default client cache expiry of 1 second, we also tested two other shorter cache expiration times for IndexFS, 250ms and 50ms. This allows us to demonstrate the performance of IndexFS under requirements for shorter directory mutation latencies. Note that the performance of IndexFS with a fixed 1-second lease expiration time will be comparable to IndexFS with adaptive leases determined by both directory mutation frequencies as well as the depth of pathnames[35].

Figure 7(a) shows the aggregated throughput during the replay phase with scaling number of servers and a matching number of clients. In general, both systems demonstrate sub-linear scalability. As the number of clients increases, each client gets fewer operations to replay. For IndexFS, this inevitably reduces cache effectiveness and introduces higher RPC amplification. Since all clients will have to refresh lookup states for entries at the top of the tree when cache is cold, metadata servers storing these directory entries will eventually become performance bottlenecks, affecting the scalability of IndexFS. On the other hand, ShardFS suffers from the increasing overhead of directory lookup state mutation operations as system resources grow. Though each metadata server gets fewer file operations in a bigger system, the number of directory lookup state mutations remains the same. Besides, directory lookup state mutation operations in ShardFS can experience more blocking than IndexFS because there is more state updates in each ShardFS server, leading to more LevelDB compactions. With strong scaling, directory lookup state mutation operations can dominate ShardFS servers' load and strongly inhibit its scalability.

Figure 8(a) and 8(b) show the latency distribution for *stat* (read) and *chmod* (mutation) on directories respectively. For *stat*, since ShardFS completes the request in a single RPC, its latency is better than that in IndexFS even at the 99th percentile. IndexFS with a shorter expiry time sees longer latency for both 50th and

70th percentile latency due to more cache misses and renewals. For *chmod* on directories, because ShardFS always waits for the slowest *chmod* to finish, the latency is much higher than that in IndexFS. For IndexFS, the 70th percentile latencies are comparable among all IndexFS configurations (and are much smaller than the lease expiration time) because lookup states not cached by any IndexFS clients can simply be modified without waiting for leases to expire. The tail latency, however, does reflect the actual lease expiration setting used in IndexFS.

**Lesson#3:** Pathname lookup is a key bottleneck limiting file system metadata scalability. IndexFS suffers from load imbalance caused by cache renewals at top of the namespace tree. ShardFS is throttled at the write path with increasing overhead of replication.

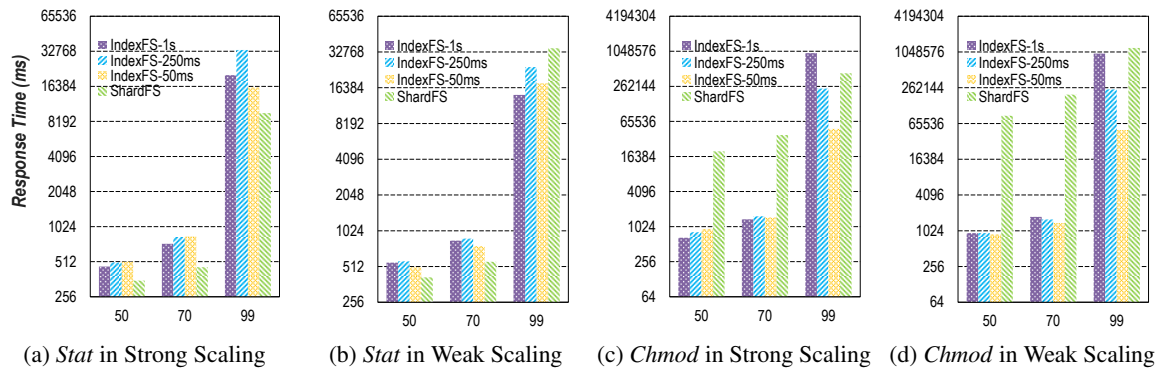


Figure 8: 50th, 70th, and 99th percentile latency for *stat* operations, and *chmod* operations on directories.

### 4.3 Weak Scaling Workload Experiment

Because of unavoidable non-scalability of strong scaling experiments, we define weak scaling for metadata workloads. Strong scalability measures speedup with increasing amount of resources under a fixed amount of work. The experiment in the previous section 4.2 is a strong scaling workload.

**Weak Scaling Metadata Workload Conjecture:** In workloads like HPC checkpoint creation the number of files created is linear in the number of processing tasks, which is linear in the number of compute cores; more cores are used by more tasks each accessing a separate file (or chunk of a file). However the number of directories used is often not linear in the number of tasks—there may be a fixed number of directories per job. We define weak scalability as linear speedup with increasing amount of processing resources under a fixed number of total directory operations but with an increasing number of total file operations. As Amdahl’s law predicts, strong scalability is rarely achievable; weak scaling workloads are a better way to measure if an application and the file systems it uses are scalable.

We enhanced our benchmark tool with support for weak scaling workloads. We first create file variants with names formed by appending a suffix. This allows us to scale the number of files in each directory as the system size grows. File operations are also generated against these files to mimic the effect of more files as the system scales. Directories and their operations remain the same with increasing scale, but both the number of files and file operations increases linearly with the size of the system. For each experiment, we adjust the number of YCSB threads in each client to achieve the best performance.

Figure 7(b) shows a weak scaling version of the experiment in the last section. All systems show the same performance as in Figure 7(a) with 8 servers because this is the same workload. With more servers, all systems scale better than with a strong scaling workload. IndexFS-50ms starts to flat out as it reaches 128 servers. IndexFS-1s performs 17% worse than ShardFS at 128 servers even though it has 12% better performance at 8 servers, a relative shift down by 25%. With a weak scaling workload, the number of requests sent to each ShardFS metadata server is expected to be the same. The load is balanced across all

metadata servers when there are no extremely popular files in the workload, so ShardFS scales well with weak scaling workload. On the contrary, IndexFS experiences higher RPC amplification and thus results in higher load variance as the system grows. This is due to more client requests for the directories at the top of the namespace.

Figure 8(b) and 8(d) show the latency distribution for *stat* and *chmod* on directories operations with weak scaling workload. Without extra RPC requests for *stat* operations, ShardFS completes 70th percentile operations faster than IndexFS. ShardFS sees higher *stat* latencies in the tail than in strong scaling workload because more *stat* operations can be affected by compactions. Since directory mutations are replicated on each metadata server generating more changes in the state on each server, compactions are triggered more frequently in the ShardFS metadata servers, competing with request execution for CPU resources. This can be alleviated with optimizations of individual ShardFS metadata servers by separating file and directory metadata to reduce size of data in compactions. For *chmod* on directories, most operations in IndexFS finish much before the timeout. Because IndexFS tracks the last time any client cached a directory entry, IndexFS can complete a directory mutation operation on cold directories without waiting for cache to expire. In ShardFS, clients not only need to acquire locks for the operation, but also need to wait for the operation to complete on all metadata servers machines. A slow *chmod* slows the whole operation and results in higher latency in ShardFS.

**Lesson#4:** ShardFS is able to drive more scalability and facilitate better latencies when the frequency of metadata updates is fixed and independent of the size of the system.

## 5 Conclusion

Many cloud storage systems lack a general-purpose scalable metadata service that distributes both namespace and directories. This tech report identifies key design trade-offs that need to be made for delivering scalable and parallel metadata performance. A novel system design called ShardFS is proposed specifically for weak scaling metadata workloads. ShardFS replicates directory lookup state for load balance and optimizing file metadata operation latency. Our experiments systematically compare the metadata operation performance of ShardFS with other alternative designs under a variety of common cloud storage workloads. We found that IndexFS-like solutions highly rely on the cache locality to reduce pathname lookup overhead and achieve load balancing, although they can provide better performance for strong scaling metadata workload and certain metadata mutation operations like *chmod* on directories. Giraffa sacrifices file system semantics for reducing network round trips. Because of its layering approach, the underlying table software may not provide ideal load balancing for metadata workload, especially concurrent accesses to large directories. ShardFS with server replication can scale more linearly and provide lower average response time than systems with client caching for weak scaling and read intensive metadata workloads.

## References

- [1] Apache thrift. <http://thrift.apache.org>.
- [2] ABAD, C. L., LUU, H., ROBERTS, N., LEE, K., LU, Y., AND CAMPBELL, R. H. Metadata traces and workload models for evaluating big storage systems. In *Proceedings of the 2012 IEEE/ACM Fifth International Conference on Utility and Cloud Computing (UCC)* (2012), IEEE Computer Society.
- [3] ADYA, A., BOLOSKY, W. J., CASTRO, M., CERMAK, G., CHAIKEN, R., DOUCEUR, J. R., HOWELL, J., LORCH, J. R., THEIMER, M., AND WATTENHOFER, R. P. FARSITE: Federated, available, and reliable storage for an incompletely trusted environment. In *Proceedings of the 5th symposium on operating systems design and implementation (OSDI)* (2002).

- [4] BENT, J., GIBSON, G., GRIDER, G., MCCLELLAND, B., NOWOCZYNSKI, P., NUNEZ, J., POLTE, M., AND WINGATE, M. PLFS: a checkpoint filesystem for parallel applications. In *Proceedings of the conference on high performance computing networking, storage and analysis (SC)* (2009).
- [5] BERENSON, H., BERNSTEIN, P., GRAY, J., MELTON, J., O'NEIL, E., AND O'NEIL, P. A critique of ANSI SQL isolation levels. In *ACM SIGMOD Record* (1995), vol. 24.
- [6] BERNSTEIN, P. A., AND GOODMAN, N. Concurrency control in distributed database systems. *ACM Computing Survey* 13 (1981).
- [7] BRANDT, S. A., MILLER, E. L., LONG, D. D., AND XUE, L. Efficient metadata management in large distributed storage systems. In *Proceedings of the 20th IEEE / 11th NASA Goddard Conference on Mass Storage Systems and Technologies (MSST)* (2003), IEEE Computer Society.
- [8] BURBANK, J., MILLS, D., AND KASCH, W. Network time protocol version 4: Protocol and algorithms specification. *Network* (2010).
- [9] CALDER, B., WANG, J., OGUS, A., NILAKANTAN, N., SKJOLSVOLD, A., MCKELVIE, S., XU, Y., SRIVASTAV, S., WU, J., SIMITCI, H., ET AL. Windows Azure Storage: a highly available cloud storage service with strong consistency. In *Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles (SOSP)* (2011), ACM.
- [10] CARNS, P. H., III, W. B. L., ROSS, R. B., AND THAKUR, R. PvfS: A parallel file system for linux clusters. In *Proceedings of the 4th annual Linux showcase and conference* (2000), pp. 391–430.
- [11] CASSANDRA. Apache cassandra database. <http://cassandra.apache.org/>.
- [12] CEPH. 'benchmark ceph mds' in user mailing list. <http://ceph-users.ceph.narkive.com/cTkbQ2C3/benchmarking-ceph-mds>.
- [13] CORBETT, J. C., DEAN, J., EPSTEIN, M., FIKES, A., FROST, C., FURMAN, J. J., GHEMAWAT, S., GUBAREV, A., HEISER, C., HOCHSCHILD, P., HSIEH, W., KANTHAK, S., KOGAN, E., LI, H., LLOYD, A., MELNIK, S., MWAURA, D., NAGLE, D., QUINLAN, S., RAO, R., ROLIG, L., SAITO, Y., SZYMANIAK, M., TAYLOR, C., WANG, R., AND WOODFORD, D. Spanner: Google's globally-distributed Database. In *Proceedings of the 10th USENIX conference on operating systems design and implementation (OSDI)* (2012).
- [14] CULLY, B., WIRES, J., MEYER, D., JAMIESON, K., FRASER, K., DEEGAN, T., STODDEN, D., LEFEBVRE, G., FERSTAY, D., AND WARFIELD, A. Strata: High-performance scalable storage on virtualized non-volatile memory. In *Proceedings of the 12th USENIX Conference on File and Storage Technologies (FAST)* (2014), pp. 17–31.
- [15] DAYAL, S. Characterizing HEC storage systems at rest. *Carnegie Mellon University PDL Technique Report CMU-PDL-08-109* (2008).
- [16] DEAN, J., AND GHEMAWAT, S. MapReduce: Simplified Data Processing on Large Clusters. In *Proc. of the 6th Symposium on Operating System Design and Implementation (OSDI)* (Berkeley, CA, USA, 2004), pp. 137–150.
- [17] DOUCEUR, J. R., AND HOWELL, J. Distributed directory service in the farsite file system. In *Proceedings of the 7th symposium on operating systems design and implementation (OSDI)* (2006).

- [18] FAN, B., LIM, H., ANDERSEN, D. G., AND KAMINSKY, M. Small cache, big effect: Provable load balancing for randomly partitioned cluster services. In *Proceedings of the 2nd ACM Symposium on Cloud Computing (SOCC)* (2011), ACM, p. 23.
- [19] GEORGE, L. Hbase: The definitive guide. In *O'Reilly Media* (2011).
- [20] GHEMAWAT, S., GOBIOFF, H., AND LEUNG, S.-T. The google file system. In *Proceedings of the 19th ACM symposium on operating systems (SOSP)* (2003).
- [21] GIBSON, G., GRIDER, G., JACOBSON, A., AND LLOYD, W. Probe: A thousand-node experimental cluster for computer systems research. *USENIX; login* 38, 3 (2013).
- [22] GIBSON, G. A., NAGLE, D. F., AMIRI, K., BUTLER, J., CHANG, F. W., GOBIOFF, H., HARDIN, C., RIEDEL, E., ROCHBERG, D., AND ZELENKA, J. A cost-effective, high-bandwidth storage architecture. In *ACM SIGPLAN Notices* (1998), vol. 33, ACM.
- [23] HARTER, T., BORTHAKUR, D., DONG, S., AIYER, A. S., TANG, L., ARPACI-DUSSEAU, A. C., AND ARPACI-DUSSEAU, R. H. Analysis of hdfs under hbase: a facebook messages case study. In *Proceedings of the 12th USENIX conference on file and storage technologies (FAST)* (2014), pp. 199–212.
- [24] HDFS. Hadoop file system. <http://hadoop.apache.org/>.
- [25] HUNT, P., KONAR, M., JUNQUEIRA, F. P., AND REED, B. Zookeeper: Wait-free coordination for internet-scale systems. In *USENIX Annual Technical Conference* (2010).
- [26] LEVELDB. A fast and lightweight key/value database library. <http://code.google.com/p/leveldb/>.
- [27] LUSTRE. Lustre file system. <http://www.lustre.org/>.
- [28] MEYER, D. T., AND BOLOSKY, W. J. A study of practical deduplication. In *Proceedings of the 9th USENIX conference on file and storage technologies (FAST)* (2011).
- [29] ONEIL, P., CHENG, E., GAWLICK, D., AND ONEIL, E. The log-structured merge-tree. *Acta Informatica* 33, 4 (1996), 351–385.
- [30] PATIL, S., AND GIBSON, G. Scale and concurrency of GIGA+: File system directories with millions of files. In *Proceedings of the 9th USENIX conference on file and storage technologies (FAST)* (2011).
- [31] PATIL, S., POLTE, M., REN, K., TANTISIRIROJ, W., XIAO, L., LÓPEZ, J., GIBSON, G., FUCHS, A., AND RINALDI, B. YCSB++: Benchmarking and performance debugging advanced features in scalable table stores. In *Proceedings of the 2nd ACM Symposium on Cloud Computing (SOCC)* (2011).
- [32] RAMAKRISHNAN, R., AND GEHRKE, J. *Database management systems*. Osborne/McGraw-Hill, 2000.
- [33] REN, K., AND GIBSON, G. TableFS: Enhancing metadata efficiency in the local file system. *USENIX annual technical conference (ATC)* (2013).
- [34] REN, K., LÓPEZ, J., AND GIBSON, G. Otus: resource attribution in data-intensive clusters. In *Proceedings of the second international workshop on MapReduce and its applications (MapReduce)* (2011), ACM.

- [35] REN, K., ZHENG, Q., PATIL, S., AND GIBSON, G. Scaling file system metadata performance with stateless caching and bulk insertion. In *Proceedings of the conference on high performance computing networking, storage and analysis (SC)* (2014).
- [36] SCHMUCK, F., AND HASKIN, R. GPFS: A Shared-Disk File System for Large Computing Clusters. In *Proceedings of the 1st USENIX conference on file and storage technologies (FAST)* (2002).
- [37] SHARDFS. Shardfs code release. <http://www.pdl.cmu.edu/ShardFS/index.shtml>.
- [38] SHVACHKO, K. V. Hdfs scalability: The limits to growth. *USENIX ;login* 35, 2 (2010), 6–16.
- [39] STONEBRAKER, M., MADDEN, S., ABADI, D. J., HARIZOPOULOS, S., HACHEM, N., AND HEL-  
LAND, P. The end of an architectural era: (it’s time for a complete rewrite). In *Proceedings of the 33rd International Conference on Very Large Data Bases (VLDB)* (2007).
- [40] THEKKATH, C., MANN, T., AND LEE, E. Frangipani: A scalable distributed file system. In *ACM SIGOPS Operating Systems Review* (1997), vol. 31, ACM, pp. 224–237.
- [41] THOMSON, A., AND ABADI, D. J. CalvinFS: consistent wan replication and scalable metadata man-  
agement for distributed file systems. In *Proceedings of the 13th USENIX Conference on File and Storage Technologies (FAST)* (2015).
- [42] WEIL, S. A., BRANDT, S. A., MILLER, E. L., LONG, D. D. E., AND MALTZAHN, C. Ceph: A Scalable, High-Performance Distributed File System. In *Proceedings of the 7th symposium on operating systems design and implementation (OSDI)* (2006).
- [43] WELCH, B., AND NOER, G. Optimizing a hybrid ssd/hdd hpc storage system based on file size distributions. *Proceedings of 29th IEEE conference on massive data storage (MSST)* (2013).
- [44] WELCH, B., UNANGST, M., ABBASI, Z., GIBSON, G. A., MUELLER, B., SMALL, J., ZELENKA, J., AND ZHOU, B. Scalable performance of the panasas parallel file system. In *Proceedings of the 6th USENIX conference on file and storage technologies (FAST)* (2008).