



Data Caching for Enterprise-Grade Petabyte-Scale OLAP

Chunxu Tang and Bin Fan, *Alluxio*; Jing Zhao and Chen Liang, *Uber, Inc.*; Yi Wang and Beinan Wang, *Alluxio*; Ziyue Qiu, *Carnegie Mellon University and Uber, Inc.*; Lu Qiu, Bowen Ding, Shouzhuo Sun, Saiguang Che, Jiaming Mai, Shouwei Chen, Yu Zhu, and Jianjian Xie, *Alluxio*; Yutian (James) Sun, *Meta, Inc.*; Yao Li and Yangjun Zhang, *Uber, Inc.*; Ke Wang, *Meta, Inc.*; Mingmin Chen, *Uber, Inc.*

<https://www.usenix.org/conference/atc24/presentation/tang>

This paper is included in the Proceedings of the
2024 USENIX Annual Technical Conference.

July 10–12, 2024 • Santa Clara, CA, USA

978-1-939133-41-0

Open access to the Proceedings of the
2024 USENIX Annual Technical Conference
is sponsored by



Data Caching for Enterprise-Grade Petabyte-Scale OLAP

Chunxu Tang¹, Bin Fan¹, Jing Zhao², Chen Liang², Yi Wang¹, Beinan Wang¹, Ziyue Qiu^{3,2}, Lu Qiu¹, Bowen Ding¹, Shouzhuo Sun¹, Saiguang Che¹, Jiaming Mai¹, Shouwei Chen¹, Yu Zhu¹, Jianjian Xie¹, Yutian (James) Sun⁴, Yao Li², Yangjun Zhang², Ke Wang⁴, and Mingmin Chen²

¹Alluxio, Inc., ²Uber, Inc., ³Carnegie Mellon University, ⁴Meta, Inc.

Abstract

With the exponential growth of data and evolving use cases, petabyte-scale OLAP data platforms are increasingly adopting a model that decouples compute from storage. This shift, evident in organizations like Uber and Meta, introduces operational challenges including massive, read-heavy I/O traffic with potential throttling, as well as skewed and fragmented data access patterns. Addressing these challenges, this paper introduces the Alluxio local (edge) cache, a highly effective architectural optimization tailored for such environments. This embeddable cache, optimized for petabyte-scale data analytics, leverages local SSD resources to alleviate network I/O and API call pressures, significantly improving data transfer efficiency. Integrated with OLAP systems like Presto and storage services like HDFS, the Alluxio local cache has demonstrated its effectiveness in handling large-scale, enterprise-grade workloads over three years of deployment at Uber and Meta. We share insights and operational experiences in implementing these optimizations, providing valuable perspectives on managing modern, massive-scale OLAP workloads.

1 Introduction

With the exponential growth of data and the continuous evolution of use cases, modern online analytical processing (OLAP) systems have become crucial in many data-centric organizations, such as Uber, Meta, and Twitter. Historically, many interactive OLAP engines were designed with co-located compute and storage, necessitating the pre-loading of data into local SSDs, HDDs, or even memory to enhance querying performance. While this approach was effective in certain contexts, it led to significant challenges, including storage fragmentation, data duplication, and stranded resources. In response, there has been a fundamental shift to a more versatile architecture that decouples compute and data lake storage [3, 24]. With this disaggregation, data, metadata, and compute are allowed to scale independently, thereby enhancing efficiency and flexibility in handling diverse and large-scale workloads.




Figure 1: Compute-storage disaggregation with SSD cache.

This architectural evolution towards decoupled compute and data lake storage, however, introduces new operational challenges for large-scale OLAP systems like Presto [21], Redshift [7], and BigQuery [6]. A primary challenge lies in the significant strain on network resources or API throughput of the data lake, driven by the vast volumes of data accessed by these systems. Additionally, while the total amount of data is rapidly expanding, we observe pronounced skewness in access: a few hot files or data blocks account for the majority of the traffic. Moreover, the prevalence of columnar data formats in modern data lakes, while beneficial for interactive queries, leads to the need for small, disparate reads instead of sequentially reading large chunks of data. Our study of data analytics and storage traces at companies like Uber has identified these intertwined challenges – disproportionate read traffic, data access hot spots, and network resource strain – all contributing to the complex landscape of managing data in these evolved systems.

While the challenges are complex and multifaceted, introducing major software or hardware changes at the platform level, or fundamentally redesigning applications, query engines, or storage engines, can be a lengthy process. Such changes often take years and may significantly impact critical business operations, sometimes proving to be nearly impossible. Consequently, with joint efforts from Alluxio, Uber, and

Meta, we experimented and implemented incremental and isolated I/O optimizations within the architecture. In particular, we introduce the *Alluxio local (edge) cache*, an embeddable cache suitable for both compute and storage engines in a non-disruptive manner (Figure 1). Designed to leverage local SSD resources, the Alluxio local cache not only effectively minimizes slow network or hard-disk data retrieval associated with repetitive I/O requests but also supports both random and sequential data access in OLAP [22], thanks to its SSD-based, page-oriented structure optimized for high throughput and concurrency in petabyte-scale OLAP.

Being a collaboration of open-source communities across different organizations, this project intentionally kept design choices simple and incremental to facilitate rapid iteration and quick consensus among multiple parties. Surprisingly, our relatively straightforward designs and optimization significantly impact business-critical, enterprise-grade platforms at a petabyte scale. We observed that adding such a local cache to both compute and storage effectively leverages the strong temporal and spatial locality of modern OLAP workloads, significantly reducing request latency and improving data transfer efficiency between compute and storage layers. After over three years of deployment at Uber [15, 17] and Meta [11], Alluxio local cache has proven effective in handling Internet-scale, enterprise-grade workloads. Specifically, the integration of Alluxio local cache within Presto, known as *Presto local cache*, has decreased the 95th percentile (P95) of query latency by 49% in Meta’s production data analytics. Similarly, the implementation of Alluxio local cache in HDFS, referred to as *HDFS local cache*, has reduced the number of processes blocked due to I/O throttling by 86% in Uber’s production HDFS clusters.

Our contributions in this paper include:

- We provide an analysis of real-world trace data collected from Uber’s OLAP and storage platforms, focusing on I/O and their implications.
- We detail the requirements, design philosophy, and implementation of the Alluxio local cache, illustrating its role as an effective optimization for petabyte-scale OLAP and storage systems.
- We provide valuable insights stemming from our first-hand industrial experience in supporting and maintaining this optimized architecture over several years.

The remainder of this paper is organized as follows: Section 2 explores industrial data analytical and storage systems, outlining key workload characteristics derived from production traces that motivate the Alluxio local cache design. Section 3 details practical design requirements for the cache system. Section 4 elaborates on the architectural design and primary components of Alluxio local cache. Optimizations and enhancements are discussed in Section 5. Section 6 presents case studies on the application of Alluxio local cache in Presto

and HDFS. Sections 7 and 8 reflect on lessons learned and failure case studies from various developmental and operational contexts. Section 9 reviews related work in this domain, and Section 10 concludes the paper.

2 Background

This section provides an in-depth examination of the complexities and challenges associated with large-scale data analytics in an enterprise environment, enumerating Uber as a primary case study. It offers a detailed exploration of enterprise-grade data analytical and storage systems, emphasizing the specific characteristics and challenges posed by managing extensive datasets. Given the data privacy regulations adhered to by various organizations, our analysis primarily focuses on the real-world workloads observed in Uber’s data analytical and storage operations. Additionally, to provide a broader perspective and supplementary insights, the workloads and data handling strategies of other enterprises, such as Meta and Twitter, are also considered.

2.1 Uber’s Data Analytical & Storage Systems

As a prime example of a modern data-centric corporation, Uber’s reliance on analytics to inform its vast array of services cannot be overstated. Enumerating Uber’s data infrastructure as an example, we elaborate on typical enterprise-grade data analytical and storage systems.

2.1.1 Data Query

For Uber, data isn’t just instrumental—it’s the backbone of decision-making. At the heart of this data-driven model is Presto. Presto is an open-source distributed SQL query engine for OLAP. It employs a coordinator-worker architecture: a central coordinator node takes charge of parsing queries, formulating query plans, and distributing tasks to worker nodes; and multiple worker nodes (can be as high as thousands) handle data retrieval from storage sources, such as on-premises HDFS or cloud-based storage like GCS.

Presto has been adopted for a wide spectrum of use cases. At Uber, whether it is the operations team synthesizing real-time dashboards or the Uber Eats and marketing divisions making pricing calls, Presto’s influence permeates. Units like compliance and growth marketing also rely on insights gathered from Presto SQL queries. Besides Uber, Presto is embraced by a large number of companies: Meta uses Presto for diverse applications, including dashboard reporting, ad-hoc analysis, and data transformation. Meta engineers also combined Presto and Spark within the Presto on Spark project to use Presto as the single piece to serve interactive, ad-hoc, extract-transform-load (ETL), and graph processing jobs [23]. Presto dominated Twitter’s interactive analytics for

hate speech analysis, engagement data validation, and tool dashboard reporting [26].

At Uber, the sheer magnitude of Presto’s operation is evident: approximately 9,000 daily users, around 500,000 daily queries, and a management of 50PB of data. This expansive setup spans 2 data centers, 7,000 nodes, and 20 unique Presto clusters across two regions [14].

2.1.2 Data Storage

With one of the most formidable Hadoop Distributed File System (HDFS) deployments globally, Uber’s exabyte-spanning data across various clusters demands persistent scalability. Achieving a harmony between efficiency, reliability, and performance is critical and challenging. HDFS powers a wide range of use cases at Uber, including fraud detection, machine learning, and expected time of arrival (ETA) calculation. The data analytical requests dominate HDFS traffic, which is one of the motivations for implementing a cache system for both Uber’s data analytical and storage systems.

To scale Uber’s HDFS, Uber engineers instituted several enhancements, such as the adoption of View File System (ViewFs), NameNode garbage collection tuning, and an HDFS load management service [33]. For further cost efficiency, a transition to denser HDDs (16+TB SKUs) is underway, superseding the current 4TB HDD SKUs that dominate the HDFS clusters. This indicates that HDD capacities will expand by factors of 2x to 4x [15].

2.2 Industrial Data Analytical & Storage Workload Characteristics & Challenges

With a closer look at enterprise-grade workloads for interactive query processing and large-scale data storage, we observed a few crucial workload characteristics and challenges, which significantly motivated the work of Alluxio local cache.

Data-intensive Analytics at Massive-Scale Our analysis highlights the massive scale of traffic in enterprise-grade data analytics and storage workloads. For instance, Uber’s Presto system processes an astonishing volume of data, handling around 500 thousand queries and managing 50 petabytes (PB) of data daily [14]. Similarly, Twitter’s hybrid-cloud SQL federation system addresses tens of thousands of queries each day, involving roughly 10 PB of data [25]. Additionally, Meta’s extensive deployment of Presto across thousands of nodes processes hundreds of PBs of data on a daily basis [18].

Table 1 presents statistics from four HDFS DataNodes in high-activity production clusters, reflecting the scale of data transactions over a period of approximately 20 hours. These statistics underscore the substantial scale of operations and the need for advanced solutions to manage such large-scale data workloads. Notably, at Uber, HDFS traffic is primarily

driven by OLAP queries, with Presto queries alone accounting for approximately 90% of this traffic.

Host	Host1	Host 2	Host 3	Host 4
Total reads (M)	13.5	12.8	8.5	14.3
Total writes (K)	3.3	4.7	4.6	45
Reads / writes	4091.0	2723.4	1847.8	317.8
Read traffic on top 10K blocks	89%	94%	99%	99%

Table 1: Production traffic of Uber’s HDFS clusters.

Skewness in Data Access Patterns A deeper analysis of enterprise-grade workloads reveals a pronounced skewness in read data access distribution. For illustration, Figure 2 demonstrates that an average Presto node at Uber has a Zipfian factor of up to 1.39, indicating heavily skewed data access patterns. This means that some files are accessed with disproportionately high frequency, highlighting the potential efficiency gains achievable through caching.




Figure 2: Popularity rank and Zipfian distribution, indicating a heavy skewness in data access distribution.

Similarly, in Uber’s HDFS traffic, distinct hot spots are observed. Table 1 provides a summary of the read traffic for the top 10K data blocks across four HDFS DataNodes. This data indicates that up to 99% of the read traffic is concentrated on these top 10K blocks, further emphasizing the importance of effective data management strategies.

Limited I/O Bandwidth with High-Density Disks Adopting high-density disk SKUs presents challenges, particularly in terms of disk I/O bandwidth. For instance, while Uber’s HDD capacities are expanding by 2x to 4x, the I/O bandwidth of each HDD isn’t increasing at a comparable rate. This discrepancy results in I/O throttling within HDFS DataNodes, which is evident from the number of blocked processes on a node. In Uber’s production traffic, the count of blocked processes can reach up to several thousand within just one minute. Given the persistent and significant instances of slow reads, preventing performance degradation becomes crucial.

Fragmented Access to Columnar Files In modern data lakes, the data processed by distributed SQL engines like

Presto and Spark SQL is increasingly converging to columnar formats such as ORC [2] or Parquet [29]. These formats are favored for their efficiency in read operations. These engines have implemented various query optimization techniques, with predicate pushdown being a key example. This technique allows the system to move filtering conditions closer to the storage layer. While these optimizations lead to performance gains, they also often result in a high number of read requests for small portions of data files.

This shift is clearly reflected in Uber’s HDFS production traces. More than 50% of SQL requests on HDFS access less than 10KB of data, and over 90% involve less than 1 MB [27]. Such statistics underscore the importance of page-based cache optimizations and strategic management in fragmented data access, especially for small data chunks.

3 Design Requirements

In addition to the aforementioned workload characteristics and challenges that motivate the design and implementation of a cache system, there are a few other design requirements listed as follows.

An application-agnostic cache. The prevalence of the aforementioned characteristics across various data analytical and storage workloads necessitates the development of a cache system that is both flexible and extensible. Such a system should be capable of integrating with a range of applications and use cases, thus requiring an application-agnostic approach. We present the architectural design of Alluxio local cache in Section 4.

Minimal socket communication with applications. Fetching data over the network is slow for interactive analytics scenarios [3]. The cache system needs to minimize socket communication with the application leveraging the cache. Due to this requirement and resource limitations, we preferred a local cache system over a distributed cache. The local cache runs within the same JVM process as the application, interacting with the remote storage through a Hadoop-compatible interface. Further details are discussed in Section 4.1.

Non-volatile data storage. Given the substantial memory usage by modern OLAP systems, primarily for in-memory storage of intermediate results, the cache system is expected to have a minimal memory footprint. By design, we aim for single-node caching of up to 1 TB of data. The high volume of data processed, combined with memory constraints, renders in-memory caching of all data files impractical. We shed light on the storage medium for cache data storage in Section 4.2.

Capability to support both random access and sequential data access. Industrial data analytical and storage workloads are inclined to frequently access small data segments. Therefore, the cache system must efficiently support not only traditional sequential data access but also random access, which targets specific portions of a data file. A page-based

organization of cached data is an effective way to meet this need. We present details in Sections 4.3 and 4.4.

High performance. Traditionally, a cache system is expected to be highly performant, aiding in reducing the request latency through the storage of frequently accessed data in a quicker-access medium. This conventional cache design requirement remains essential. We discuss details on various cache optimizations and enhancements in Section 5.

Besides the above requirements, when an enterprise-grade large-scale data system adopts a local cache, the local cache will inevitably face classical distributed systems’ challenges, encompassing aspects like reliability, scalability, and availability. These challenges foster adaptations and further development of the local cache, such as soft-affinity scheduling, application-level metrics aggregation, and cache rate limiting strategies. We elaborate on concrete practical challenges and corresponding implementation solutions in Section 6.

4 Architectural Design

4.1 Overall Design

Figure 3 depicts the comprehensive workflow of handling data requests in Alluxio local cache, which extends the file system APIs from Alluxio [12, 13]. The local cache is primarily designed to enhance read performance by introducing a straightforward yet effective local cache layer. The workflow involves an *admission controller*, a pivotal component that identifies and prioritizes the most frequently accessed files for caching. The controller ensures that only data deemed essential for caching is admitted, while other data leads to a non-cache read path, directed to external *data sources*.

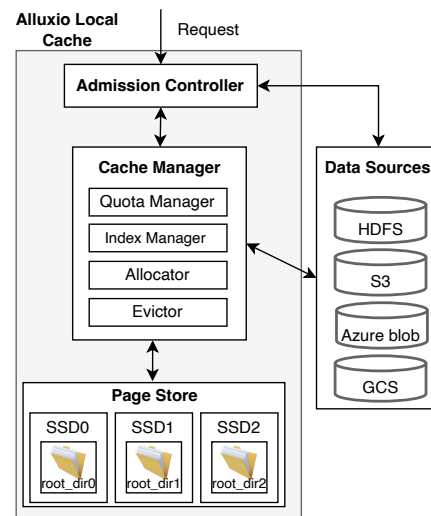


Figure 3: The workflow of Alluxio local cache.

Alluxio local cache transforms file-level read operations into more granular page-level operations through the *page*

store. The page store comprises multiple local cache directories, each with its designated cache capacity. Local SSDs are the recommended storage medium for cached data, considering their performance efficiency.

Central to the Alluxio local cache is the *cache manager*, which encompasses the system’s key functionalities. When a requested page is found in the cache (cache hit), the cache manager promptly returns the requested page to the client. Conversely, in cases of cache misses, the page is fetched from external data sources, cached locally, and served to the client (read-through strategy). The *quota manager* maintains a hierarchical resource control management, allowing for quota specifications at various levels, such as global, schema, table, and partition levels. The *index manager* organizes pages efficiently based on page metadata criteria, thereby optimizing the conditional page search to $O(1)$ complexity. The *allocator* is responsible for assigning cache pages to appropriate directories, considering factors like file identification, hash algorithms, directory capacity, and page affinity. In scenarios where a directory’s available space is insufficient or if a quota rule is breached, the system initiates an eviction process. The *evictor* component orchestrates multiple cache eviction strategies, such as FIFO, random, and LRU. It provides an interface for the integration of alternative policies if needed. Besides the conventional policy-based cache eviction, the local cache also incorporates a time-based eviction strategy by setting the time to live for each cached file. A periodic background job evicts expired data as needed. This feature is raised and adopted in response to the growing data privacy requirements.

4.2 Cache Data Storage

Currently, many cache systems prefer to store data in memory. The high speed of memory access, especially random data access, makes this design suitable for extremely latency-sensitive production scenarios. For example, Twitter developed Twemcache based on the open-sourced Memcached, which serves as an in-memory key-value cache to serve real-time user and tweet data [31]. Similarly, Meta scaled and enhanced Memcached to manage billions of requests per second in vast social media networks [19].

Nonetheless, the storage medium of the local cache is constrained by the available memory in the node. Modern enterprise-grade large-scale OLAP systems are known for intensive memory usage, leaving limited room for additional cache storage within the memory. Although scaling up the local node’s hardware resources is a viable solution to augment higher memory capacity for caching, this solution can incur substantially higher costs, particularly when scaling to serve large volumes of enterprise-grade data analytics.

SSDs emerged as an alternative solution for caching data locally. SSDs present an advantageous balance between cost-efficiency and high-speed access for both random and sequential data access. An analysis of resource utilization in data

analytical jobs at Uber and Meta reveals that local SSDs are usually underutilized, indicating a potential opportunity for cache optimization. Consequently, these insights guide our decision to leverage local SSDs for caching data, while maintaining the metadata still in memory to ensure fast access.

4.3 Page-Based Organization

Alluxio local cache organizes cached data in pages. Initially, the default page size was 64 MB, the same as the default block size in HDFS. However, from our operational experience on effective support of columnar files and SSD bandwidth, we adjusted the default page size to 1 MB (Section 7). In practice, a production data file, such as a Parquet file at Uber, is cached as hundreds of pages, providing finer-grained control on portions of the file. When the write operation of a page is completed, the cached page will be immediately available for subsequent read operations. We developed fine-grained locking mechanisms to support high-read concurrency. Page information is self-contained in page names and parent folders, so the application can identify whether the data it wants is cached by simply looking up the metadata information.




Figure 4: Organization of cached data files.

Additionally, cached data is organized in a multi-level hierarchy, as shown in Figure 4. Shared information is organized as folders. Top-level folders represent persistent global information that can be used in cache recovery. For example, the *page_size* folder is a top-level folder, since the page size information is required to calculate the page index, which is subsequently used together with file IDs to compute page IDs. Shared file information is stored as folders, such as full paths, and file version information. To avoid one directory containing too many sub-folders or files to downgrade the lookup performance, we 1) add buckets for an additional layer and 2) add file ID directories.

4.4 Index Management

Alluxio local cache supports common cache management operations such as read, write, and delete on different granularity levels. The basic levels include the page level, being the finest

level where each page can be addressed individually, and the file level composed of the pages belonging to the same file.

Logical levels can be leveraged by applications to perform bulk page operations efficiently. For example, in Presto, the data is organized in a partition-table-schema hierarchy. This hierarchy maps to a tree of nested scopes in Alluxio local cache: a global scope representing the entirety of the cache and scopes for each schema, table, or partition. Pages belonging to a particular data file are assigned with scopes of the partition, table, and schema of that file. Users can perform I/O operations on all pages of a specified scope efficiently. For example, in a table partitioned by a date column, a user may wish to delete all pages belonging to a certain outdated partition, to free up space for more up-to-date data. Without the partition scopes, the user would otherwise have to first find all files that belong to that specific partition, possibly by invoking a costly directory listing operation and filtering through the result to locate all the files.




Figure 5: Using indexed sets for organizing pages.

There is also a level that corresponds to the medium of the page store. This level allows users to address all pages stored in a particular storage device, mounted as a directory on the file system of the host. This is typically useful for reporting the cache usage on a particular device, or deleting all pages in a storage device, when the backing device becomes faulty.

Initially, all the page metadata was stored in two maps, indexed by the page ID and the ID of the files they are part of, respectively. This naive approach quickly proved insufficient as the need to support more diverse levels became prominent. To support efficiently accessing pages for a given level, without resorting to iterating over the entire page store, the pages need to be properly indexed with the different levels. We use indexed sets to store all pages' metadata. As Figure 5 illustrates, the universe set contains all pages that are currently stored in the cache. Each indexed set is a subset of the universe indexed by a certain property of the page's metadata. There are three sets indexed by the file ID A, B, and C, and three sets indexed by storage directories 1 and 2, respectively.

5 Cache Optimizations & Improvements

5.1 Cache Admission Strategies

Nowadays, a modern OLAP system is capable of handling petabyte-scale data daily, making storing all production datasets in the local cache unfeasible. As reported in Section 2.2, the frequency of data block access varies significantly, with certain hot data blocks accessed hundreds of thousands of times within a single day. Consequently, the admission decisions are governed by several strategies.

One such approach defines cache admission criteria as regular expressions or more versatile JSON-format expressions. This strategy has been effectively used in the Presto local cache (Section 6.1). For instance, the following code snippet illustrates a static filtering rule applied to the *table_bar* table. The *maxCachedPartitions* field specifies the upper limit of 100 on the number of data partitions of the table that can be retained within the cache. In production, the filtering rules are set by platform owners and infrequently updated. At Uber, after such filtering, less than 10% of requests require remote storage access [14].

```

1 {
2   "databases": [
3     {
4       "name": "database_foo",
5       "tables": [
6         {
7           "name": "table_bar",
8           "maxCachedPartitions": 100
9         }
10      ]
11    }
12  ]
13 }

```

An alternative solution involves the analysis of historical data access patterns in a defined time frame. This admission strategy leverages sliding windows to assess the hotness of data blocks. This approach proves its effectiveness in HDFS local cache (Section 6.2). For the requests which fulfill the admission policy, only around 1% of them require slower storage access.

5.2 Quota Management for Multi-Tenancy

The support of multi-tenancy is a common need in an enterprise-grade cache system. To enable fairness and efficient resource utilization, Alluxio local cache adopts flexible multi-layer quota management, preventing any single tenant from excessively consuming resources. It facilitates the allocation of quotas through a hierarchical tenant-based approach: 1) user tenants, where quotas are set for different teams, organizational units, or regions; 2) data tenants, where resource allocation is based on specific data entities like partitions or tables; 3) custom tenants, offering flexibility for bespoke

quota configurations based on any logical grouping, such as project-based or application-specific.

The quota verification process in Alluxio local cache is hierarchical, starting from the most detailed level (often partitions) and ascending through tables, schemas, and up to the global level. This quota check prevents any single tenant from monopolizing cache resources and starving others.

Our initial implementation restricted the total quota for a table's partitions to not exceed the table's quota. However, practical experience in a production environment revealed that this limitation hindered efficient resource sharing. Consequently, we evolved the design to allow the collective quota of partitions to surpass the quota of their parent table. For instance, if a table with a 1 TB quota has two partitions, each with an 800 GB quota, the system ensures that no partition exceeds its individual quota (800 GB).

In situations where a quota violation occurs, Alluxio local cache evicts data to re-establish quota compliance. There are two eviction strategies: 1) Partition-level eviction. If a partition exceeds its quota, data eviction is triggered for that specific partition. 2) Table-level sharing and eviction. If the overall usage at the table level exceeds its quota, the system performs random eviction across partitions. This randomization is particularly beneficial in cases where one partition demands significantly more cache space than others, ensuring efficient resource sharing and utilization.

6 Use Cases & Evaluation




Figure 6: Alluxio local cache use cases, which cover the compute, distributed cache, and storage layers.

This section discusses practical use cases and evaluation of Alluxio local cache. Thanks to years of efforts from open-source communities, nowadays, the original Alluxio local cache has evolved into a more mature, robust system that contains advanced features, such as caching filtering, scheduling, and quota management for multi-tenancy. Alluxio local cache

has been used in scenarios that are well beyond the initially designed boundaries of SQL query engines like Presto. It functions as a general-purpose local cache and finds applications across a diverse range of scenarios including distributed file systems, machine learning training, and distributed cache.

As depicted in Figure 6, at the compute layer, in the data analytics landscape, query engines like Presto use Alluxio local cache to enhance performance and cost-efficiency. Meanwhile, in the realm of machine learning, particularly in training phases, Filesystem in Userspace (FUSE) [16] utilizes the local cache to help improve training performance and GPU utilization. At the storage layer, HDFS DataNodes leverage Alluxio local cache within each node to bolster data retrieval performance and mitigate I/O bottlenecks. Bridging the compute layer and the storage layer is a distributed cache layer, where Alluxio local cache is integrated into each cache worker node to serve the traffic. This section primarily focuses on two key aspects: the local cache applied for data analytics (Presto local cache) and the local cache used for storage (HDFS local cache). A detailed discussion on the application of local cache in machine learning training and the distributed cache layer will be presented in a following paper.

6.1 Case Study - Presto Local Cache

6.1.1 Design Overview

The journey of integrating Alluxio local cache into Presto started in 2019. It is initially envisioned as a pluggable and lightweight cache module, aiming at decreasing query latency and reducing the data transfer between the compute and storage tiers. Contributions from open-source communities of Presto and Alluxio, alongside inputs from tech companies like Uber, Meta, and Twitter, have significantly evolved Presto local cache to fulfill the growing demand for data analytics at the enterprise-grade. Presto supports reading and parsing various file formats, particularly column-oriented file formats widely used in OLAP, such as Apache Parquet [29] and Apache ORC [2]. These column-oriented file formats inherently organize data files by columns and expedite data processing by segmenting data into row groups (stripes in ORC, row groups in Parquet). Presto employs such design and distributes splits among the worker nodes following the generation of a distributed query plan by the coordinator node.

As shown in Figure 7, Presto incorporates specific reader classes to handle diverse file formats. For example, the *ParquetReader* is responsible for reading and parsing Parquet files. These file reader classes access data cached in Presto local cache, where data is systematically organized in pages (1 MB by default). In cases where the data is already cached (warm cache), it is promptly returned to the reader. Otherwise, the local cache employs a read-through strategy, loading the data from external storage for the reader. Operational insights reveal that parsing complex column-oriented data files can

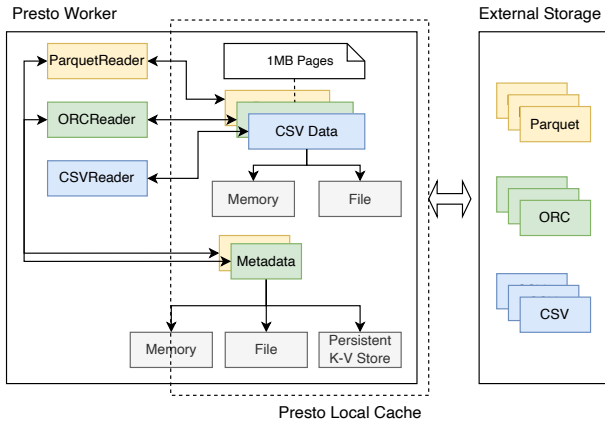


Figure 7: Caching data and metadata in Presto local cache.

consume as much as 30% of CPU resources [30]. To mitigate the issue, Presto local cache also caches file metadata, including indexes, stripe (ORC) and column (Parquet) metadata, and file-level metadata. Thanks to its extensible design, the metadata, represented as key-value pairs, can be stored in memory, files, or persistent key-value stores like RocksDB. In contrast, data can be stored in memory or files, but not in external key-value stores due to its organization in data chunks (pages). In enterprise-grade production environments, data is usually cached in files and metadata in memory or RocksDB.

To ensure cache coherence, Presto will always fetch the latest metadata of input files from persistent storage, before splitting the input files on different workers along with the metadata. In case an input file is changed, the stale copy in the cache will be invalidated based on the timestamp of file creation or modification stored in the cache.

6.1.2 Soft-Affinity Scheduling

The Presto coordinator node is responsible for distributing splits to worker nodes, which in turn retrieve data from external storage for subsequent processing. Conventionally, each data file comprises multiple splits, and the scheduler’s primary objective was to evenly distribute tasks by randomly assigning splits to workers. This approach, however, proved to be inefficient for caching as it led to frequent admission and eviction of data from each worker’s local cache. To address that, we implemented a soft-affinity scheduler in Presto to allocate splits from the same file to the same worker node with the best efforts, as depicted in Figure 8.

The soft-affinity scheduler uses the consistent hashing algorithm, with the file as the hashing input, to calculate the preferred worker node for a split. However, due to the potential for hot spots—where certain files are more frequently accessed than others—a designated worker node may be overwhelmed with an excessive number of splits. To gauge the

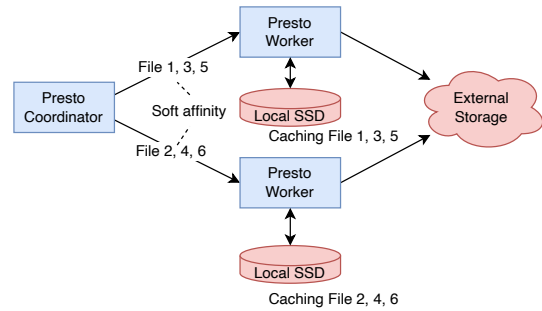


Figure 8: File-level soft-affinity scheduling.

workload of a node, the scheduler compares *max-splits-per-node* with *max-pending-splits-per-task*. If the initially chosen worker node is deemed busy, the scheduler opts for a secondary worker node from the hash ring. If the secondary node also lacks sufficient resources, it indicates a temporary inability to maintain soft-affinity, compelling the scheduler to assign the task to the least burdened worker in the cluster. This worker is instructed to fetch data directly from external storage, bypassing local caching, thus serving as a fallback.

6.1.3 Per-Query Metrics Aggregation

Our operational experience of Presto local cache underscores the importance of collecting and monitoring the metrics associated with frequently accessed data. In an enterprise-grade OLAP workload, access to data within a commonly used table is often unevenly distributed across its partitions, with some partitions experiencing substantially higher access frequencies. The local cache addresses this by identifying such hot partitions and aggregating per-query metrics to generate table-level insights. To facilitate this, we have developed a comprehensive metric system that integrates both Presto and Alluxio local cache metrics. Apart from utilizing Presto JMX exporters for logging JMX metrics, we have enhanced Presto *RuntimeStats* to reflect the runtime statistics of each query stage. Whenever Presto I/O operations engage the local cache, relevant metrics, such as cache hit rate and pages read, are recorded through the Alluxio system metrics. Concurrently, query-level runtime statistics are logged as in-memory metrics, which are periodically gathered for extensive monitoring and analytical purposes.

6.1.4 Evaluation

TPC-DS benchmark. To empirically validate the effectiveness of Presto local cache, we deployed a Presto cluster within the AWS Kubernetes environment, consisting of one coordinator node and four worker nodes. Each node has 32 CPU cores, 128 GB memory, and 150 GB SSD storage. The evaluation utilized the TPC-DS [5] benchmark with a scale factor of 100. The dataset, stored in Parquet format, was hosted on AWS





Figure 9: Comparison of query execution time of TPC-DS Query 81 to Query 99 without and with Presto local cache.

S3. We compared the query execution time between scenarios using the Presto local cache (warm cache) and direct reads from S3 (non-cache read). Figure 9 indicates a reduction in query execution times of Query 81 to Query 99, ranging from approximately 10% to 30% when data is pre-loaded into the cache. The full TPC-DS benchmark evaluation results are included in Appendix A.

ments were observed in Meta’s production data analytics. Specifically, in one of Meta’s internal use cases, the query latency P50 was reduced by around 33%, and P95 was reduced by around 49% [11]. These results highlight Presto local cache’s effectiveness in mitigating the query processing tail latency. Additionally, there was a 57% reduction in total data scanned from remote storage, indicating the local cache’s impact in reducing data transfer between compute and storage layers.



(a) 90th percentile (P90) of query time spent on reading files before and after enabling the cache. (b) 50th percentile (P50) of query time spent on reading files before and after enabling the cache.

Figure 10: Comparison of execution time spent on fetching data before and after enabling the cache.

Uber’s production results. In an evaluation within Uber’s enterprise-grade production environment, we focused on the performance enhancement in query processing attributable to the utilization of local caching. This assessment involved measuring the time spent on reading Parquet files and comparing the results obtained before and after the activation of caching. In Presto, this is quantified by the *inputWall* metric of the *ScanFilterProjectOperator*, a key internal phase within a Presto query, responsible for data input handling and initial filtering. Notably, at Uber, for the tables onboarded with Presto local cache, more than 97% of queries involve reading through a *ScanFilterProjectOperator* stage, underscoring the relevance of this metric. Figure 10 shows that the 90th percentile (P90) of the time spent on reading files was reduced by 67% and 50th percentile (P50) was decreased by 64%.

Meta’s production results. Similar performance enhance-

6.2 Case Study - HDFS Local Cache

6.2.1 Design Overview

The Alluxio local cache module is embedded into each HDFS DataNode to construct the HDFS local cache, as depicted in Figure 11. The local cache stores hot data in the local SSD and original datasets are stored in the HDD due to the capacity restrictions. Take Uber’s production HDFS clusters as an example, with 500 TB of data storage capacity on a DataNode, only a few TBs of SSD is available [15]. It is unfeasible to cache all data, which will cause frequent cache admission and eviction and lower cache hit rates. To ensure cache efficiency, when a requested data block is not loaded in the cache, the cache rate limiter will decide whether the data block should be loaded into the cache, based on historical data access patterns. If the data block is unsuitable for caching, the cache rate limiter will go to the normal non-cache read path. Otherwise, if the data block is determined as cache-worthy, this block will be cached for faster access in the future.

Similar to the implementation of Presto local cache, HDFS local cache employs the Hadoop-compatible file system APIs provided by Alluxio, which help with the seamless integration with the DataNode. Notably, the soft-affinity scheduling in Presto local cache does not apply here as the HDFS NameNode has already maintained a metadata table, recording the location of each data block.

HDFS DataNodes manage HDFS blocks, each contains a block file and a metadata file. The metadata file contains checksum info of the block file, and therefore, it has to match




Figure 11: The architectural design of HDFS local cache.

the corresponding block file all the time. Upon caching a block on disk, both the block and metadata files are stored separately on the local SSD, maintaining the same format as the original files. This design guarantees that either both the block and metadata files are read from the cache, or both are read from their original non-cache locations, but never any form of the mix. This block-level caching design simplifies reasoning and enhances reliability, forming the basis for subsequent design decisions discussed below.

6.2.2 Cache Rate Limiter

The cache rate limiter plays a crucial role in favoring the retention of frequently accessed data blocks in the cache while restricting those less frequently accessed. By evaluating a few design options against simulation tests based on production traffic, we created the rate limiter with the design we called *BucketTimeRateLimit*.

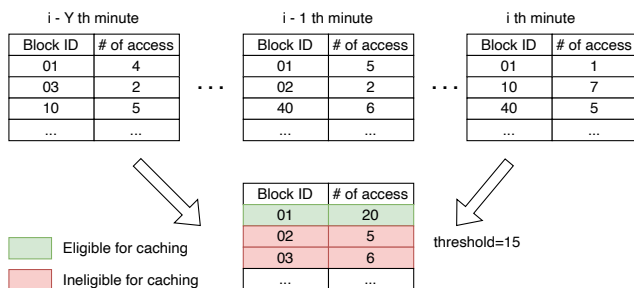


Figure 12: Data structures used in the cache rate limiter.

The algorithm functions by evaluating if a data block has been accessed more than X times in the past Y time interval. It incorporates an ordered list of buckets as a sliding window. Each bucket logs the access count for each block during its respective minute-long window. *BucketTimeRateLimit* keeps a constant number of active buckets and discards the oldest bucket every minute. To determine whether a block should be subjected to rate limiting, the algorithm aggregates the total number of accesses across all existing buckets, as shown in Figure 12. A block is classified as high-frequency and eligible

for caching if its access count surpasses a predefined threshold (e.g. 15 in the example).

6.2.3 HDFS Write Operations

Although the HDFS local cache stores only finalized block files, certain HDFS operations can still modify these blocks, necessitating updates in the cached data. Such HDFS operations include appending to a block (*append*), deleting a block (*delete*), and truncating a block (*truncate*). Since HDFS *truncate* is not used in our HDFS local cache use cases, we focus on the *append* and *delete* operations.

Append to a block. Within HDFS, the append operation allows a client to add new content to the end of an existing file. To distinguish between the original and modified blocks, HDFS employs a versioning system where each block is assigned a *generation stamp*. Each invocation of the append operation increments the block’s generation stamp, signaling a new version of the block.

HDFS local cache uses the generation stamp for snapshot isolation. When a block and its checksum are loaded into the cache, the associated generation stamp is retained and used in conjunction with the blockID to create a unique cache key. This approach ensures that during an ongoing append operation, cache readers access the original version of the block, as our design prohibits loading blocks that are currently being written into the cache. Post completion of the append operation, the updated block, identifiable by its new generation stamp, is considered a distinct cache entry. This allows the latest version of the block to be cached and made available for subsequent read operations.

Delete a block. The deletion of a block in HDFS requires the removal of its cache copy. Initially, our approach involved a time-based policy for periodically purging obsolete cache entries. To improve the efficiency and delete outdated cache entries more timely, we introduced an in-memory mapping within each DataNode. This mapping tracks the blocks currently cached, enabling the DataNode to identify and instruct the local cache to eliminate related page files upon block deletion. The in-memory mapping takes the form of $\langle \text{blockId} \rightarrow (\text{cacheId}, \text{fileLength}) \rangle$, where *fileLength* helps compute the relevant page files.

A caveat of this in-memory approach is the loss of mapping data when a DataNode restarts. To address this, the DataNode clears all local cached contents and rebuilds the cache from the ground up upon restarting. Although this approach may impact cache efficiency, we deem it a viable compromise, given that DataNode restarts are infrequent in production.

6.2.4 Evaluation

To validate the effectiveness of HDFS local cache, we deployed it within Uber’s production HDFS clusters. As depicted in Figure 13, the introduction of HDFS local cache led

to a substantial portion of read traffic being handled by the cache. This observation aligns with our earlier workload analysis in Section 2.2, which identified that a significant share of traffic is directed towards hot spots. In this DataNode, in the given time period of one hour, the rate of bytes read from the cache is, on average, threefold that of non-cache reads. More than 70% of total read bytes are serviced by the local cache.




Figure 13: The comparison of cache read rates and non-cache read rates in an HDFS DataNode.



Figure 14: The comparison of blocked processes (I/O throttling) in one DataNode enabling and disabling the local cache.

Section 2.2 highlights the I/O throttling issue in the production HDFS cluster, reflected in a high number of blocked processes in the DataNodes. Figure 14 compares the number of blocked processes in a DataNode, with and without the HDFS local cache. Upon disabling the cache at timestamp 70, there is a rapid increase in blocked processes, reaching up to approximately five thousand. During this one-hour period, the local cache reduces the number of blocked processes by an average of 86%.

7 Lessons Learned

This section recounts the experiences and qualitative lessons gleaned from the development and operation of Alluxio local cache in organizations like Alluxio, Uber, and Meta.

Enhancing OLAP system performance with data file metadata caching. Nowadays, column-oriented file formats, such as Apache Parquet and ORC, are widely used in enterprise-grade data analytical systems. These formats

involve extensive scanning and parsing of metadata by query systems, consuming up to 30% of CPU resources in practice, thereby impacting overall system performance. Our experience shows that caching metadata, like column metadata in Parquet files, can significantly alleviate this overhead. Specifically, caching deserialized metadata objects can reduce CPU usage by up to 40%, enhancing the efficiency of query systems [30].

The trade-off in cache page size between read amplification and the number of requests to remote storage. Our analysis, as detailed in Section 2.2, indicates that large-scale query systems like Presto are inclined to generate a large number of small read operations. A larger cache page size, while reducing the number of read requests to remote storage, increases read amplification. Conversely, smaller cache page sizes reduce data fetched but increase the metadata memory footprint and the number of storage requests.

In practice, either a high read amplification or a large number of requests to storage can have a non-negligible impact on the organizational-level infrastructure resource utilization and system performance, potentially leading to unexpectedly high monetary costs. The ideal page size is a function of traffic patterns, file formats, and desired metadata size. Our empirical findings, particularly in the context of Uber and Meta’s production traffic, indicated that a cache page size of 1 MB strikes an optimal balance. For practical cache page size tuning, we recommend starting with a page size of 1 MB or a number akin to the data file’s page size, followed by gradual adjustments to identify the optimal setting.

An aggregated metrics system is crucial for cache tuning and debugging. A modern enterprise-grade data analytical or storage system can contain thousands of nodes, leading to thousands of local cache deployments. An important design decision we made is implementing an aggregated metrics system, collecting both basic metrics from the local cache and aggregated metrics based on the applications. The metric system provides a centralized view of predefined and user-customized metrics and helps developers drill down to detailed metrics. We found such a system is helpful for cache setting tuning and debugging. Besides the common metrics, such as cache hit rate, cache eviction rate, and cache bytes read from storage, we also learned that error-related metrics, including error counts of different operations (put, get, delete, etc.) and breakdowns of concrete types of errors, are extremely helpful to identify root causes in debugging.

Keeping the seats for temporary offline nodes (lazy data movement). Nowadays, enterprise-grade large-scale systems are usually deployed in container orchestration platforms, such as Kubernetes [28] and Mesos [10]. Handling temporary offline nodes becomes a critical aspect of cache architecture design. These systems, due to various operational factors such as application failures, resource constraints, periodic maintenance, or software updates, may experience intermittent node restarts. To effectively manage such scenarios, we incorpo-

rated a strategy in our cache design that utilizes a timeout mechanism within the hashing ring. This approach is designed to allocate a waiting period for nodes that temporarily go offline. If these nodes come back online within the timeout period, the system avoids unnecessary data movement across the cache nodes. This method of 'lazy data movement' has proved to be efficient in maintaining cache stability and performance, mitigating the impact of frequent data movement in containerized deployment environments.

Considering data imbalance and fallback when setting the number of cache replicas. Hot spots are common in production data traffic. Some tables or data blocks are much more frequently accessed than others. This uneven distribution of data access stimulates careful consideration in the allocation of cache replicas. Our operational insights suggest that increasing the number of replicas can alleviate pressure on hot spots but may inadvertently lead to increased latency in locating an unoccupied cache node. In practice, to find an optimal balance, we adopted a strategy that limits the number of cache replicas to a maximum of two. In cases where both replicas are unavailable due to resource limitations, the system defaults to retrieving data from remote storage. This hybrid approach, which combines a limited replica count with a remote storage fallback, has demonstrated greater robustness and lower latency in practice compared to simply increasing the number of replicas. This method effectively addresses the challenges posed by data hot spots while optimizing for system performance and reliability.

Trade-offs between local cache and distributed cache. At Meta and Uber, implementing a local cache was favored given the stringent timeline while adding services in production is more challenging. Moreover, local caching offers a substantial return on investment (ROI) for decision-makers by efficiently utilizing available resources on compute nodes. For example, each of Meta's Presto node is capable of utilizing 1-2 TB of SSD resources for local caching. However, local caching also has certain limitations. It is constrained by the available local resources, lacks the ability to scale independently, and does not support data sharing across different computing engines, such as between Presto and Spark.

8 Failure Case Study

This section shares insights from unexpected failures encountered in the operation of the local cache within large-scale distributed systems, such as Presto and HDFS.

File read hanging. As we discussed in Section 4.2, Alluxio local cache employs the local SSDs to store the cached data. Accessing the cached data requires a *read_file* operation from the disk. From our operational experience, the local cache sometimes experiences read operation hangups, for as much as 10 minutes, probably caused by resource contention with other applications on the shared machine. Implementing a fallback to the remote storage system and with a default 10-

second timeout for *read_file* operations have proven to be effective in mitigating this issue. If the operation does not return in this time range, the local cache will retrieve the data from the remote storage.

Corrupted files. In the realm of SSD-based caching systems, such as the one used in Presto local cache, file corruption, although infrequent, remains a potential issue. Instances of corrupted files have been observed during our operations, leading to significant impacts such as the failure of SQL queries accessing these files and a notable decrease in cache hit ratios due to the cache's inability to write new data. To address this challenge, we adopted a proactive strategy by initiating early eviction of cache entries when the failure occurs, even before the cache reaches its full capacity. This approach has been instrumental in mitigating the adverse effects of file corruption. Furthermore, as highlighted in the previous section, a detailed breakdown of error types played a crucial role in identifying and understanding the root causes of such failures.

Insufficient disk capacity. Alluxio local cache has a configurable cache capacity that is determined during the cache start. This configuration defines the maximum disk space allocated for cache use. However, it is possible for the SSD storage to become fully utilized before reaching the set cache capacity during ongoing operations. Such scenarios of insufficient disk capacity can lead to intricate and demanding debugging processes. To address the issue, through catching the *No space left on device* exception, the local cache initiates early cache eviction, rather than waiting for the cache to reach its pre-configured full capacity.

9 Related Work

The growing need for an efficient cache system for large-scale data analytics and storage has stimulated a large body of research work. We highlight the most closely related work.

Cache for OLAP. Caching saves data access from remote storage in OLAP systems by keeping the frequently accessed data in memory or high-bandwidth fast storage. NoDB [1] recognizes raw data files as first-class citizens and uses caching to optimize query execution on external tables. Many also find caching helpful for bridging the performance gap caused by storage-compute disaggregation architecture in modern cloud databases. FlexPushdownDB by Yang et al. [32] aims to keep hot data in the compute layer and reduce the amount of data transferred between storage and compute. Similarly, Crystal by Durner et al. [4] provides a unified cache management system as a storage middleware to improve query latency and save bandwidth costs between compute and storage in cloud analytical databases. Apache Spark initially used Tachyon [13] (now rebranded as Alluxio) for built-in resilient distributed datasets (RDD) caching. Alluxio local cache was developed by the original Tachyon team, with a special focus on a simple, independent, and embedded caching system. A

similar caching framework is Rubix [20], which provides a lightweight data cache for big data engines. Compared with Rubix, Alluxio local cache operates without the need for daemon processes (*BookKeeper*) and employs soft-affinity scheduling for higher cache hit rate and performance.

Cache for distributed file systems. Caching is also crucial for distributed file systems to reduce latency by keeping hot data blocks or files in faster storage. HDFS implements a centralized cache management in which the NameNode instructs DataNodes to cache hot data blocks in off-heap memory [8]. This HDFS caching contributes to the construction of multi-tier storage in distributed file systems. For example, Herodotou et al. [9] presented a distributed tier-based storage system using HDFS cache as one of the tiers. Furthermore, as SSDs outperform conventional HDDs, similar to the HDFS local cache presented in this paper, Zhao et al. [34] created a unified caching middleware for distributed file systems by organizing the storage in both HDDs and SSDs.

10 Conclusion

Designed to tackle a range of challenges, including massive-scale data-intensive analytics, uneven data access distribution, limited I/O bandwidth, and fragmented access in column-oriented files, the relatively straightforward implementation of Alluxio local cache has proven its notable effectiveness in enterprise-grade petabyte-scale OLAP and storage systems. Its application as the Presto local cache has led to a 49% reduction of query latency P95 and a 57% decrease in data retrieval from remote storage within Meta’s production data analytics. HDFS local cache, adopted in Uber’s production HDFS clusters, resulted in an 86% reduction in blocked processes due to I/O throttling. The insights and lessons gathered from the development and operation of these cache systems are believed to be of substantial value for building enterprise-grade compute and storage systems.

Acknowledgment

The widespread adoption of Alluxio local cache highlights a successful collaborative effort involving numerous individuals. We would like to express our gratitude to Haoyuan Li, Amelia Wong, Xiaodan (Danica) Wang, Xuan Du, Adit Madan, Rongrong Zhong, Rohit Jain, Junyan Guo, Fengnan Li, and many others. We also thank the anonymous USENIX ATC reviewers for their insightful comments, which considerably improved this paper.

References

- [1] Ioannis Alagiannis, Renata Borovica, Miguel Branco, Stratos Idreos, and Anastasia Ailamaki. NoDB: Efficient query execution on raw data files. In *Proceedings of*

the 2012 ACM SIGMOD International Conference on Management of Data, pages 241–252, 2012.

- [2] Apache. Apache ORC, 2024. <https://orc.apache.org/>.
- [3] Biswapesh Chattopadhyay, Pedro Pedreira, Sameer Agarwal, Yutian Sun, Suketu Vakharia, Peng Li, Weiran Liu, and Sundaram Narayanan. Shared foundations: Modernizing Meta’s data lakehouse. In *The Conference on Innovative Data Systems Research (CIDR)*, 2023.
- [4] Dominik Durner, Badrish Chandramouli, and Yinan Li. Crystal: A unified cache storage system for analytical databases. *Proceedings of the VLDB Endowment*, 14(11):2432–2444, 2021.
- [5] Yuanan Feng. TPC-DS queries for Presto, 2016. <https://github.com/yuananf/tpcds-presto>.
- [6] Google. Google BigQuery, 2024. <https://cloud.google.com/bigquery>.
- [7] Anurag Gupta, Deepak Agarwal, Derek Tan, Jakub Kulesza, Rahul Pathak, Stefano Stefani, and Vidhya Srinivasan. Amazon Redshift and the case for simpler data warehouses. In *Proceedings of the 2015 ACM SIGMOD international conference on management of data*, pages 1917–1923, 2015.
- [8] Apache Hadoop. Centralized cache management in HDFS, 2024. <https://hadoop.apache.org/docs/stable/hadoop-project-dist/hadoop-hdfs/CentralizedCacheManagement.html>.
- [9] Herodotos Herodotou and Elena Kakoulli. Automating distributed tiered storage management in cluster computing. *arXiv preprint arXiv:1907.02394*, 2019.
- [10] Benjamin Hindman, Andy Konwinski, Matei Zaharia, Ali Ghodsi, Anthony D Joseph, Randy H Katz, Scott Shenker, and Ion Stoica. Mesos: A platform for fine-grained resource sharing in the data center. In *NSDI*, volume 11, pages 22–22, 2011.
- [11] Rohit Jain. Improving presto latencies with alluxio data caching, 2020. <https://prestodb.io/blog/2020/06/16/alluxio-datacaching/>.
- [12] Haoyuan Li. *Alluxio: A Virtual Distributed File System*. PhD thesis, EECS Department, University of California, Berkeley, May 2018.
- [13] Haoyuan Li, Ali Ghodsi, Matei Zaharia, Scott Shenker, and Ion Stoica. Tachyon: Reliable, memory speed storage for cluster computing frameworks. In *Proceedings of the ACM Symposium on Cloud Computing*, pages 1–15, 2014.

- [14] Chen Liang and Beinan Wang. Speed up Presto at Uber with Alluxio local cache, 2022. <https://www.uber.com/blog/speed-up-presto-with-alluxio-local-cache/>.
- [15] Chen Liang, Jing Zhao, Yangjun Zhang, Junyan Guo, and Fengnan Li. Optimizing HDFS with datanode local cache for high-density HDD adoption, 2023. <https://www.uber.com/blog/optimizing-hdfs-with-datanode-local-cache/>.
- [16] libfuse. libfuse, 2024. <https://github.com/libfuse/libfuse>.
- [17] Zhenxiao Luo. Engineering data analytics with Presto and Apache Parquet at Uber, 2017. <https://www.uber.com/blog/presto/>.
- [18] Zhenxiao Luo, Lu Niu, Venki Korukanti, Yutian Sun, Masha Basmanova, Yi He, Beinan Wang, Devsh Agrawal, Hao Luo, Chunxu Tang, et al. From batch processing to real time analytics: Running Presto® at scale. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*, pages 1598–1609. IEEE, 2022.
- [19] Rajesh Nishtala, Hans Fugal, Steven Grimm, Marc Kwiatkowski, Herman Lee, Harry C Li, Ryan McElroy, Mike Paleczny, Daniel Peek, Paul Saab, et al. Scaling memcache at Facebook. In *10th USENIX Symposium on Networked Systems Design and Implementation (NSDI 13)*, pages 385–398, 2013.
- [20] Rubix. Rubix, 2024. <https://github.com/qubole/rubix>.
- [21] Raghav Sethi, Martin Traverso, Dain Sundstrom, David Phillips, Wenlei Xie, Yutian Sun, Nezh Yegitbasi, Haozhun Jin, Eric Hwang, Nileema Shingte, et al. Presto: SQL on everything. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*, pages 1802–1813. IEEE, 2019.
- [22] Utku Sirin, Pinar Tözün, Danica Porobic, and Anastasia Ailamaki. Micro-architectural analysis of in-memory OLTP. In *Proceedings of the 2016 International Conference on Management of Data*, pages 387–402, 2016.
- [23] Yutian Sun, Tim Meehan, Rebecca Schlusel, Wenlei Xie, Masha Basmanova, Orri Erling, Andrii Rosa, Shixuan Fan, Rongrong Zhong, Arun Thirupathi, et al. Presto: A decade of SQL analytics at Meta. *Proceedings of the ACM on Management of Data*, 1(2):1–25, 2023.
- [24] Junjay Tan, Thanana Ghanem, Matthew Perron, Xiangyao Yu, Michael Stonebraker, David DeWitt, Marco Serafini, Ashraf Aboulnaga, and Tim Kraska. Choosing a cloud DBMS: Architectures and tradeoffs. *Proceedings of the VLDB Endowment*, 12(12):2170–2182, 2019.
- [25] Chunxu Tang, Beinan Wang, Zhenxiao Luo, Huijun Wu, Shajan Dasan, Maosong Fu, Yao Li, Mainak Ghosh, Ruchin Kabra, Nikhil Kantibhai Navadiya, et al. Forecasting SQL query cost at Twitter. In *2021 IEEE International Conference on Cloud Engineering (IC2E)*, pages 154–160. IEEE, 2021.
- [26] Chunxu Tang, Beinan Wang, Huijun Wu, Zhenxiao Wang, Yao Li, Vrushali Channapattan, Zhenxiao Luo, Ruchin Kabra, Mainak Ghosh, Nikhil Kantibhai Navadiya, et al. Serving hybrid-cloud SQL interactive queries at Twitter. In *European Conference on Software Architecture*, pages 3–21. Springer, 2022.
- [27] Chunxu Tang, Yi Wang, Bin Fan, Beinan Wang, Shouwei Chen, Ziyue Qiu, Chen Liang, Jing Zhao, Yu Zhu, Mingmin Chen, et al. Rethinking the cloudonomics of efficient I/O for data-intensive analytics applications. *arXiv preprint arXiv:2311.00156*, 2023.
- [28] Abhishek Verma, Luis Pedrosa, Madhukar Korupolu, David Oppenheimer, Eric Tune, and John Wilkes. Large-scale cluster management at Google with Borg. In *Proceedings of the Tenth European Conference on Computer Systems*, pages 1–17, 2015.
- [29] Deepak Vohra. Apache Parquet. In *Practical Hadoop Ecosystem*, pages 325–335. Springer, 2016.
- [30] Beinan Wang, Chunxu Tang, Rongrong Zhong, Bin Fan, Yi Wang, Jasmine Wang, Shouwei Chen, Bowen Ding, and Lu Zhang. Metadata caching in Presto: Towards fast data processing. In *2022 IEEE International Conference on Big Data (Big Data)*, pages 2737–2741. IEEE, 2022.
- [31] Juncheng Yang, Yao Yue, and KV Rashmi. A large-scale analysis of hundreds of in-memory key-value cache clusters at Twitter. *ACM Transactions on Storage (TOS)*, 17(3):1–35, 2021.
- [32] Yifei Yang, Matt Youill, Matthew Woicik, Yizhou Liu, Xiangyao Yu, Marco Serafini, Ashraf Aboulnaga, and Michael Stonebraker. Flexpushdowndb: Hybrid push-down and caching in a cloud DBMS. *Proceedings of the VLDB Endowment*, 14(11):2101–2113, 2021.
- [33] Ang Zhang and Wei Yan. Scaling Uber’s Apache Hadoop distributed file system for growth, 2018. <https://www.uber.com/blog/scaling-hdfs/>.
- [34] Dongfang Zhao and Ioan Raicu. Hycache: A user-level caching middleware for distributed file systems. In *2013 IEEE International Symposium on Parallel & Distributed Processing, Workshops and Phd Forum*, pages 1997–2006. IEEE, 2013.

A TPC-DS Evaluation




Figure 15: Comparison of query execution time of TPC-DS Query 1 to Query 49 without and with Presto local cache.




Figure 16: Comparison of query execution time of TPC-DS Query 50 to Query 99 without and with Presto local cache.