InftyDedup: Scalable and Cost-Effective Cloud Tiering with Deduplication

Iwona Kotlarska
9LivesData, LLC

Andrzej Jackowski
9LivesData, LLC

Krzysztof Lichota
9LivesData, LLC

Michal Welnicki
9LivesData, LLC

Cezary Dubnicki
9LivesData, LLC

Konrad Iwanicki
University of Warsaw

Abstract

Backup solutions increasingly offer cloud tiering, that is, moving selected data from on-premise storage to the cloud. As subsequent backups usually contain repeating data, it is reasonable to pair cloud tiering with deduplication to significantly reduce the cloud storage utilization, and hence the associated costs. However, solutions for cloud tiering with deduplication that would harness the scaling potential of the cloud tier and minimize the total expenditures due to this tier are essentially still lacking. This paper aims to bridge this gap. First, it introduces InftyDedup, a novel system for cloud tiering with deduplication, which aims to maximize scalability by utilizing cloud services not only for storage but also computation. Accordingly, it performs deduplication in the cloud, using distributed batch algorithms, in effect allowing for processing multi-petabyte backups for a couple of dollars. Second, the paper presents algorithms for InftyDedup that employ multiple types of cloud storage to further reduce the costs of the cloud tier. They take into account the characteristics of each data chunk to efficiently select between cloud services providing hot and cold data stores, thereby reducing the overall costs by up to 26%–44%. The solutions are implemented in a state-of-the-art commercial backup system and evaluated in the cloud of a hyperscaler.

1 Introduction

Managing the surging volumes of data that require protection or long-term retention increasingly necessitates novel backup strategies [14]. A popular approach is employing cloud-based solutions. For instance, according to Veeam, the number of organizations adopting them is expected to rise from 60% in 2020 to 79% in 2024 [70]. In particular, in a survey by ESG, 72% of the participants confirmed utilizing *tiering* techniques to move colder data (e.g., older backups and archives) from on-premise storage to cloud storage [19].

In this context, data *deduplication* can become effective. Since consecutive backups often contain repeating data [44, 47], this technique reduces storage utilization on the order of tens of times [61]. As a result, it has been a core feature of several storage systems for on-premise backup applications [26, 53, 80]. In this light, for backup use cases, it is sensible to consider *cloud tiering with deduplication*, that is, moving data from a local tier (e.g., an on-premise backup appliance) to a cloud tier (e.g., an object store of a hyperscaler), so that ultimately the data kept in the cloud tier are deduplicated.

However, implementing solutions for cloud tiering with deduplication poses two major problems. First, the state-of-the-art cloud storage systems provided by hyperscalers, notably Amazon, Google, and Microsoft, have been designed for more general applications and do not offer deduplication as core functionality for their clients. Consequently, custom deduplication mechanisms suitable for cloud tiering have to be developed. Second, there is a large variety of available cloud storage service types, notably with respect to pricing models. Whereas initially a lower cost of storage implied a longer retrieval time, as in AWS Glacier, nowadays systems like AWS Glacier Instant Retrieval [73] offer the same performance as other cloud storage services. The trade-off is in turn that with a decreased per-byte monthly storage fee, the costs of data retrieval and the minimal data storage period are increased. Therefore, algorithms have to be devised to decide what type of service to use and for which data, specifically taking into account the peculiarities due to deduplication.

As we discuss in more detail further in the paper, despite some research progress, these two problems are largely open. In short, regarding the first problem, although a few backup applications [51, 68] and backend appliances [27, 34] with deduplication offer mechanisms for cloud tiering, they heavily rely on and are implemented mainly at the local tier. In effect, global deduplication (i.e., of data written via different local tiers) is not supported and, more importantly, the entire process is fundamentally limited by the resources of the local tier. In other words, despite the possibilities offered by the hyperscalers, the actual scalability of the cloud tier in such solutions is severely limited, proportionally to what is offered by an instance of the local tier. As to the second

problem, in turn, although the diversity of the service models offered by the hyperscalers (e.g., S3 vs. Glacier in AWS) can also be exploited in some solutions [49], this has to be configured manually or, at best, through policies depending on data collection ages. However, deduplication typically entails chunking data collections into smaller pieces that may be referenced multiple times, thereby having possibly different access patterns. This calls for finer-grained and more automated approaches to storage type selection.

In this paper, we address both these problems, introducing solutions for scalable and cost-effective cloud tiering with deduplication. Accordingly, our contributions are twofold.

First, we present InftyDedup, a novel system for cloud tiering with deduplication. Like the existing tiering-to-cloud backup solutions, InftyDedup moves selected data stored in a local-tier system to cloud storage, based on customer-specific backup policies. However, its operation aims to maximize scalability by exploiting cloud services—not only for storage but also computation. Therefore, rather than relying on deduplication methods for on-premise solutions, InftyDedup deduplicates data periodically in batches and using the cloud infrastructure but before actually transferring them to the cloud, which, among others, enables leveraging cloud mechanisms such as dynamic resource allocation. Other necessary functionalities, notably garbage collection of deleted data, are supported in the same way. We integrate InftyDedup with HYDRAstor [26], our commercial backup system with deduplication, and evaluate its performance in the cloud of Amazon, demonstrating, in particular, that multiple petabytes of data can be deduplicated for a couple of dollars. All in all, being highly independent of the local tier, InftyDedup overcomes the limitations of the state of the art and offers unprecedented scalability. To the best of our knowledge, this is the first application of such solutions to multi-tier backup systems.

The second contribution is an algorithm for decreasing the financial cost of storing deduplicated data in the cloud tier. It extends InftyDedup by allowing it to move deduplicated data chunks between cloud services dedicated to hot and cold storage. Whereas existing solutions do not address the problem at all or enable some optimizations at the level of data collections (e.g., backups or files), the fact that chunks are deduplicated between backups/files makes them arguably a better unit for optimizations. In InftyDedup, they are moved based on their metadata, notably reference counts due to deduplication and terse information provided by system administrators on their data collections. Our empirical evaluation of the algorithm shows that mixing storage types can reduce the total financial cost of cloud tiering with deduplication by up to 26–44%.

The rest of the paper is organized as follows. Sec. 2 gives the necessary background. Sec. 3 describes the overall architecture and specific algorithms comprising InftyDedup. Sec. 4 discusses the algorithm for exploiting cold cloud storage for cost minimization. Sec. 5 presents the experimental results. Sec. 6 surveys related work. Finally, Sec. 7 concludes.

2 Background

This section reviews the characteristics of deduplication storage, backups, and cloud services, which are essential to InftyDedup architecture.

2.1 Deduplication Storage

Deduplication is a data reduction technique that avoids writing the same data twice. For data with many duplicates, deduplication reduces the storage capacity requirements of the system [61], increases throughput, and decreases network traffic [3]. Typically, deduplication is implemented in the following steps [75]. Firstly, the data stream is chunked into small immutable blocks of size from 2 KB to 128 KB [62]. Secondly, each block receives a fingerprint, for instance, by computing SHA-256 hash of the block's data. Finally, the fingerprint is compared with other fingerprints in the system, and if the fingerprint is unique, the block's data is written.

The deduplicated blocks are typically organized in a directed acyclic graph. Each file has its *root block*, which corresponds to a vertex that keeps references to other blocks. The blocks with actual data are leaves of the DAG and keep no references. Blocks with data corresponding to a particular file form a subgraph of vertices reachable from the root block representing that file. Therefore, the movement of a deduplicated file to a different tier is effectively the movement of a subset of leaves that are reachable from the root block of the file.

A block can be removed after it is migrated to another system. However, reclaiming storage capacity in the system with deduplication is nontrivial, as the system must ensure that there are no other references to the removed block. Therefore, complex garbage-collecting algorithms that can process blocks' metadata for hours are implemented [32, 60].

The most natural use case for deduplication is backup storage, as most data does not change in consecutive backups. In our research, we leverage the characteristics and lifecycle of backups to decrease the total storage cost.

2.2 Lifecycle of Backups

Typically, numerous copies of the backup data are created based on assigned retention policies [56]. From the perspective of our research, there are two important constraints regarding the timing and life cycle of protected data.

On the one hand, the data should become quickly available and up-to-date in case of a disaster. For instance, Zerto reports [78] that their customers achieve Recovery Point Objectives of seconds, and Recovery Time Objectives of minutes. To achieve such ambitious objectives, recent data is kept as closely as possible to the infrastructure which is recovered.

On the other hand, older versions of backups need to be stored for weeks, months, or even years [69]. As the objective points for older data differs, backups are often moved

to cheaper storage after a specific time [72, 79]. Cloud is often chosen to keep the older backups for numerous reasons, including storing data in a different physical location. The pricing model of cloud storage is also appealing but, as we describe in the next section, many factors influence the total cost of storage in the cloud.

2.3 Cloud Storage

The majority of the cloud storage market is shared between three hyperscalers (Amazon Web Services, Microsoft Azure, and Google Cloud) [65], so in our considerations we assume services offered by the three as a market standard.[1] The portfolio of hyperscalers consists of numerous storage and computing products: from databases, queues, and distributed filesystems to simple storage primitives such as objects or blocks. Our goal is to minimize the storage cost of backups, so our research focuses on the most affordable products. The lowest price per stored gigabyte is offered by cold archival object stores, which are orders of magnitude cheaper than block devices, as shown in Tab. 1. However, the total cost consists of many factors, including fees per request or IO, charges for removing data before meeting the minimal storage duration, and the cost of data transfer. Accessing data in some types of the coldest storage takes additional time (e.g., 12 hours) but every hyperscaler offers cold storage with instant access [20, 22, 73].

	Amazon Web Services	Microsoft Azure	Google Cloud
Block Storage [$/GB]	0.05	0.15	0.04
Object Storage [$/GB]	0.021	0.0166	0.02
Archival Object Storage [$/GB]	0.004	0.01	0.004
Coldest Archival Object Storage [$/GB]	0.00099	0.00099	0.0012

Table 1: Sample[2] monthly costs of storing blocks and objects in public clouds [6, 30, 46].

Uploading data to the cloud is usually free, whereas cost of downloading data once a month can outweigh the cost of monthly data storage. In either case, the network throughput to the cloud is a major concern. Hyperscalers offer connecting data centers to the cloud dircetly (e.g., with 100 GbE) [7, 11] but the availability of such a high throughput networks is limited to specific regions. Alternatively, physical devices can be used for the quick movement of data [9], but it is rather for peculiar applications. Therefore, moving terabytes to the cloud can take up days.

[1] However, there are numerous innovative services offered by other providers. For instance, the latest trend to decentralize the cloud [57, 58] can help to implement InftyDedup efficiently.

[2] The price of storage products depends on many factors, including region. Moreover, each cloud provides numerous products, for instance, each provider offers more than one cold object store. The prices between providers cannot be compared directly, however, the point is that there are several categories of cloud storage products similar to the order of magnitude of the price.

2.4 Cloud Computing

The product portfolio of cloud computing services is also versatile. There are virtual machines (e.g., AWS EC2), containers (e.g., AWS ECS), and other services, such as event-driven function execution (e.g., AWS Lambda). Some of the products are prepared for specific use cases, including machine learning [31] and databases [4, 29].

The pricing model of computation services is typically based on the cost of the lower-level resource billing. For instance, ECS allows running containers on EC2 instances, so the cost of container execution depends on the amount and size of virtual machines which host the containers [5]. The billing model enables using a large number of nodes (e.g., a hundred servers) momentarily which costs next to nothing.

What is important for cost reduction, hyperscalers offer so-called *spot instances*, which are virtual machines with a discounted price up to 90% but can be interrupted at any moment . The exact price of a spot instances depends on multiple factors (e.g., the momentary demand), but historical data shows that achieving both very low risk of termination and significant cost reduction is possible [28]. Virtual machines (including spot instances) can have their own local storage (e.g., SSD drives), which is cheaper than network-attached drives but has limited durability as the data are lost if the machine is destroyed or fails. To minimize the costs of computations, we considered all of these cloud attributes in InftyDedup architecture, which we will describe shortly.

3 InftyDedup Architecture

InftyDedup moves selected data from local tier systems (which are on-premise backup appliances implemented as in Section 2.1) to the cloud tier. Local tier stores data not selected for tiering or before data is moved to the cloud, so it is expected to have their own deduplication and to be hardware-failure resistant to some degree (e.g., by implementing erasure codes or RAID). As shown in Fig. 1, the cloud tier keeps deduplicated data of files moved to the cloud with necessary persistent metadata and occasionally executes highly optimized *batch algorithms*.

Before we describe the details of the structures and algorithms, we discuss our study of cloud characteristics (Sec. 3.1) and the assumptions we made based on them (Sec. 3.2). After that, we describe the structure of in-cloud data and metadata (Sec. 3.3), model of communication between tiers (Sec. 3.4), algorithms for deduplication (Sec. 3.5), garbage collection (Sec. 3.6), and data restore (Sec. 3.7).

3.1 Cloud Cost Considerations

We studied the pricing of cloud storage and computing products to design InftyDedup architecture in line with the current trends. First, we chose product types which are common for

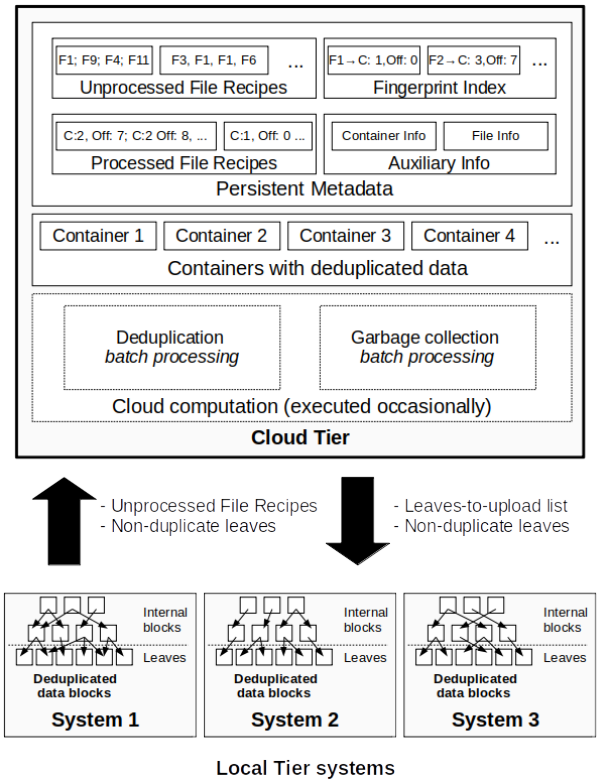

Figure 1: InftyDedup architecture.

all vendors and compared the pricing models and capabilities of each product with other products of the same cloud. We did not compare pricing between vendors, as our goal was to design an architecture that is cost efficient for any regular cloud, not choosing a particular vendor.

Keeping 1 PB of non-deduplicated data in a standard cloud object storage costs between $\$16,600$ and $\$21,000$ per month, and between $\$3,600$ and $\$10{,}000$ in the case of archival object storage with instant access. Therefore, the overall cost of storing data with deduplication, including additional storage for deduplication metadata and costs of computations, must be lower than that to bring any financial benefit.

Considering the deduplication block of 8 KB, 10:1 deduplication, and 20 bytes per fingerprint, 1 PB of data requires 262 GB of fingerprints. If new backups of similar size are written each week, over 496 billion queries whether a fingerprint exists in the cloud are required each month.

Modern architectures of inline deduplication often keep the fingerprint index (or its parts) on SSDs [3,23,45]. Considering a naive approach in which each deduplication query requires a read IO from an SSD drive, at least 190k IOPS are required to perform the necessary queries each month. To estimate the cost, let us consider AWS as an example. The monthly cost of EBS gp3 block storage which provides such amount of IOPS is $978, and EC2 instances (m5.large) capable of utilizing the IOPS cost $3827. With the total cost of nearly $5000 monthly for just handling deduplication queries, there is still room for a cost benefit from deduplication (depending on usage patterns and deduplication ratio), but the price is significant in comparison to the cost of storage without deduplication.

These calculations led us to our conclusion that, despite the fact that SSDs generally provide a high number of random read IOs, relying on random read intensive fingerprint index stored in the cloud environment *is not negligibly cheap*. There are techniques that reduce the number of read IOs for traditional sequential workloads [80], but their efficiency is decreased for modern non-sequential workloads, which need to be handled just as classic sequential workloads, as explained by Y.Allu et al. [2]. Moreover, the efficiency of methods that rely on the data locality (like SISL [80]) decreases when data is highly fragmented.[3] Finally, these methods often are not prepared to update block information during deduplication, which is a necessary part of our algorithms for cold storage.

On the other hand, transferring data within the cloud is free of charge, and even the cheapest instance can transfer hundreds of gigabytes per hour [10]. Having the possibility of dynamically scaling resources between zero and hundreds of servers, processing the fingerprint index sequentially with a batch job can be more cost-effective than keeping the fingerprint index online $24/7$ or relying on short-lived lambdas [8]. In particular, considering up to 10 times cheaper computation using the aforementioned spot instances. This key observation was used when designing the InftyDedup architecture based on assumptions explained in the next section.

3.2 Assumptions and Design Decisions

Despite the connections between tiers, our principal assumption is that *our cloud tiering duplication must be processed outside of the local tier* to prevent resource restrictions and enable functionalities like deduplication between many local tier systems. Therefore, all metadata required for deduplication must be stored and processed outside the local tier.

As the network throughput between the tiers is limited, *the data movement between the tiers should be minimal*. Therefore, only non-duplicate data must be uploaded to the cloud tier. When restoring data, it must be possible to download only the data which is not already present in the local tier. However, for efficient disaster recovery, *quick and granular backup restores must be possible*, even when the local tier is unavailable.

The next major assumption is that *batch processing is preferred over streaming processing*. Therefore, the algorithms are executed occasionally (e.g., once a day or week for deduplication and even less frequently for garbage collection). There are multiple reasons for that. Firstly, as our cost analysis of public clouds shows, being prepared for data deduplication

[3] Fragmentation also concerns restore throughput [37,41] but in the case of cloud storage, the read performance scales, and even with random 8 KB reads the egress traffic cost is equal to per request fee of such small reads.

24/7 is not negligible cheap. Secondly, as explained in Section 2.2, the backups are typically moved to the cloud after a specified period, so batch processing can be done without disrupting the data lifecycle. Finally, tiering to cloud with deduplication requires steps that take a significant time: uploading data to the cloud, and garbage collection to reclaim data in the local tier. Therefore, performing costly inline deduplication brings few benefits in practice, and we decided to use cheaper batch processing which is executed occasionally.

Garbage collection in the cloud tier must be cost-aware to ensure that the cost of data removal is not higher than keeping data for a longer period. Similarly, storing frequently accessed data in cold cloud storage actually increases the costs, so the deduplication and garbage collection algorithms must be extendable with intelligent storage type selection.

Finally, our solution is meant to be suitable for a variety of cloud platforms and providers. Although in our description and evaluation we focus on the most popular hyperscalers, our architecture can be easily adapted to others. In particular, we paid special attention to private clouds which ensure privacy and compliance. Therefore, our solution was also entirely verified in our private environment.

3.3 Data and Metadata in Cloud

Based on the aforementioned assumptions, we designed in-cloud structures of InftyDedup as follows. The largest structure contains the blocks with deduplicated data, and these blocks are grouped into *containers*. Selecting the size of containers depends on the cloud pricing, as writing and reading larger containers requires fewer requests but increases rewriting cost when reclaiming space after garbage collection.

The largest metadata structure contains *file recipes*, which are effectively a list of consecutive blocks as they appear in each file. If one block exists in a file multiple times, it also occurs multiple times in its file recipe. There are two types of file recipes. Firstly, there are *unprocessed file recipes* (UFR in short), which are provided by the local tier. UFRs contain the fingerprint of each block, as the local tier does not know the block's cloud location. Later, during deduplication processing, each entry of UFR receives a cloud address of the block it references, so the file recipe is converted to *processed file recipes* (PFR in short). PFRs can be a simple list of cloud addresses or have a tree structure to enable the deduplication of PFR's parts. In the latter case, fingerprints of PFR chunks are added to the fingerprint index, which is described shortly.

The second largest metadata structure is a fingerprint index which contains a mapping from the deduplication fingerprint of each block to its cloud location. The index is expected to be smaller than file recipes, as it contains only one entry per fingerprint. The fingerprint index is bucketed [63] rather than sorted, meaning the fingerprints are divided into thousands of buckets based on a hash function. Such data representation enables optimization of distributed fingerprint index processing, as each bucket is small enough to fit into server memory.

There are also a few smaller structures that keep information per file or per container, which are orders of magnitude smaller than the previous two. The metadata structures are compressed to reduce space and network usage.

3.4 Communication between Tiers

The exchange of data between the tiers is bidirectional but kept to a minimum as a network connection between the tiers can easily become a bottleneck. Two types of information are sent from the local tier to the cloud. For each file selected for cloud tiering, the local tier system generates a UFR, which contains a list of fingerprints of all blocks in the file. The UFR is later used as an input to batch deduplication, which generates in return a *leaves-to-upload list* that is, in fact, a list of containers. Each container consists of unique blocks that have not been uploaded to the cloud tier yet. Based on the list, the local tier uploads the blocks to the cloud. The blocks can be later downloaded from the cloud tier, which happens during the file restore operation.

Therefore, the cloud tier has minimal requirements on the interface of the local tier. It is sufficient that the local tier can generate a UFR and later upload blocks based on the list of fingerprints. The local tier can be composed of multiple systems if each system uses coherent chunking and fingerprinting methods.

3.5 Batch Deduplication

Batch deduplication (BatchDedup in short) is our distributed method of performing block deduplication in the cloud. It is expected to be performed periodically, in harmony with schedule of backups and garbage collection in local tier systems. Each execution of BatchDedup is a distributed, fault-tolerant computation that ultimately modifies persistent structures kept in the cloud storage. The computations are divided into steps, and each of the steps consists of smaller jobs that are parallelized and can be repeated in an event of node failure. For instance, our implementation uses Hive, which relies on YARN [67] to schedule jobs, and HDFS [59] for reliable storage of temporary data. The jobs can be executed on spot instances as proposed in the AWS best practices guide [16] because even if the processing is interrupted, the valid version of metadata is always kept in the cloud storage.

In short, BatchDedup takes UFRs as an input, specifies new containers with blocks to be uploaded, waits until the local tier upload the blocks, and updates persistent metadata. The UFRs are expected to be uploaded to the cloud before BatchDedup is started (UFRs partially uploaded do not participate in the process). The steps are as follows:

Step #1: UFR processing selects blocks that need to be uploaded to the cloud by comparing fingerprints in UFRs with the fingerprint index. The fingerprint index and UFRs

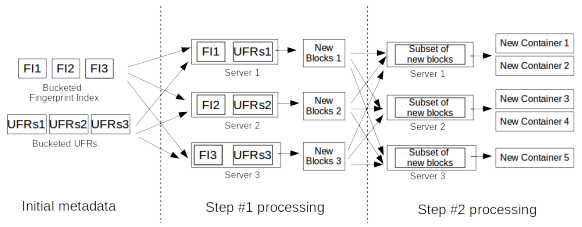

Figure 2: First two steps of BatchDedup processed in a distributed manner.

are bucketed based on fingerprints, and the buckets are distributed across multiple servers. After that, the fingerprints are compared in batches that are small enough to fit in memory. As the size of UFRs is expected to be comparable to the size of the fingerprint index, the cost of fingerprint index distribution does not dominate the cost.

Step #2: Container generation divides blocks selected in *Step #1* into containers to generate descriptions for the local tier. Each server processes a subset of blocks, and the blocks are distributed based on their original file (so blocks from the same file can be placed in the same container). The blocks are sorted by the order (offsets) in their original files,[4] as preserving the original order makes the later step of uploading the container easier, and reduces the number of requests for garbage collection and data restores for non-fragmented data.

Step #3: PFR update is conducted after the first two steps, when the block location (its container and offset) is finally known for both new and old blocks. Based on that information, each newly written file receives its PFR.

Step #4: Blocks upload is initiated by the local tier systems. The local tier systems first download the descriptions of containers (which blocks should be uploaded to which container). After that, each of the local tier systems uploads the actual data. When the uploads are completed successfully, the in-cloud metadata structures are updated to mark the new files as ready in the cloud.

The first two steps of BatchDedup are depicted in Fig. 2. Similar techniques are used to perform the remaining steps of BatchDedup and garbage collection in scale.

Overall, the process is expected to take time: BatchDedup is executed periodically, the computation in Steps #1-#3 is expected to take from minutes to hours, and the block upload in Step #4 can even take days, depending on the data volume and network bandwidth. As Step #4 is inevitable in any cloud-tiering solution, cloud tier alone is not suitable for providing very short RPO (e.g., below a minute). However, for files moved to the cloud tier after a given period, it is possible to schedule all steps in periods that will not violate the timing constraints of the backup policy.

[4] A block is expected to exist in multiple files or to be repeated within one file. In such a case, only the first appearance is stored in a container.

3.6 Batch Garbage Collection

Batch garbage collection (BatchGC in short) is expected to be executed periodically but less frequently than BatchDedup. Its purpose is to identify blocks that are no longer referenced by any PFR and reclaim free space in the containers. However, deciding whether a container should be modified to remove unreferenced data is nontrivial, as rewriting a container in the cloud has a significant cost, so we propose different strategies to decide whether a container should be rewritten.

PFRs keep addresses of containers, so rewriting a container require modifications in PFRs. The cost of processing PFRs is discouraging, as PFRs can be many times larger than fingerprint index. However, garbage collection is done occasionally, so even if it is few times more expensive than BatchDedup, the overall cost of InftyDedup is not affected that much. Therefore, our main goal is ensuring scalability which enables meeting the time constraints of other garbage collection algorithms for deduplication storage [24, 60].

BatchGC consists of the following steps:

Step #1: File removal processes non-removed PFRs to find blocks that are still referenced by at least one file.

Step #2: Container verification checks how many blocks in each container are live. Based on one of the strategies, a set of containers that will be removed or rewritten is selected.

Step #3: Container metadata are updated based on the decision taken in *Step #2*. New metadata for modified containers are calculated. In particular, some of the blocks may receive a new address, so new versions of the fingerprint index and PFRs are computed.

Step #4: Containers are rewritten to actually reduce space usage. When all newly generated containers are written, the metadata computed in *Step #3* take effect, and old containers are deleted.

We investigated three strategies which decide whether a particular container should be rewritten:

GC-Strategy #1: Reclaim only empty containers. As in most cloud services sending a request to remove an entire container is free of charge, the strategy brings cost reduction (as less capacity needs to be stored) with no additional cost. However, the strategy does not remove containers in which only a fraction of data has been deleted.

GC-Strategy #2: Reclaim containers if the rewrite pays for itself after T days. To determine whether rewriting a container will bring a cost benefit, the following ratio can be calculated for each container:

$$gc_r = \frac{COST_{rewrite}}{T_{days} * CAPACITY_{to\_be\_reclaimed} * COST_{byte\_per\_day}} \quad (1)$$

Only if $gc_r < 1.0$ rewriting container is cheaper than storing its data for T_{days}. However, picking the right value of T_{days} is nontrivial. For instance, if T_{days} is the time left until the next BatchGC, the containers are rewritten only if it brings

financial benefit before the next chance to remove any data. In many cases, such T_{days} value is too small and will prevent from rewriting some containers, despite the fact that rewriting the container would bring a financial benefit in the long run. On the other hand, large T_{days} value implies often rewriting, which can lead to exceeding *Strategy #1* costs.

GC-Strategy #3: Reclaim containers based on file expiration dates. *GC-Strategy #2* can be improved, if files contain information about their expiration date. Such information can be provided by the local tier systems in UFRs, if the expiration date results from the backup configuration. Therefore, for each container, T_{days} can be calculated as the maximal expiration date of its blocks (aligned up to the BatchGC schedule). The expiration date is expected to increase in time,[5] as new files with later expiration dates will be sotred. However, even with the constantly increasing expiration dates, the cost never exceeds *GC-Strategy #1*, as a non-empty container is rewritten only when it is beneficial.

3.7 File Restore

The cloud metadata format supports straightforward file restores. Each file has its own object, with the key based on the local tier system identifier and file path. Therefore, object storage interface features such as ACLs and per-prefix listings can be used for convenient file management. Based on the content of PFR, which stores the container address and data offset, the file can be read without any interaction with the local tier systems. As PFRs are updated during BatchGC, the movement of data between containers during GC does not spoil the reads.

However, egress traffic is a major cost, so the restore can be integrated into the local tier system for cost reduction. If the block is available locally, there is no need to download it from the cloud. If the block is not available locally, it is downloaded, and optionally be stored in the local tier system, as some workloads require reading data again in the near future (e.g., restoring multiple similar VMs). Implementing such local-tier assisted reads requires storing fingerprints in PFRs, which increases the storage cost for metadata. The fingerprints can be added and removed from PFRs on-demand during BatchDedup or BatchGC.

4 Cold Storage Utilization

To reduce the cost of storing data in the cloud, InftyDedup can be extended with an algorithm that selects whether a block should be stored in hot or cold cloud storage. Our primary goal was to utilize services that offer different pricing with comparable durability and latency [22, 73], to prevent situations in which the movement of data to cold storage negatively

[5]The expiration date for a container can also decrease if someone deletes a file before the expiration date. We find such case rather marginal. In particular, enabling WORM protection [50] prevents such removals.

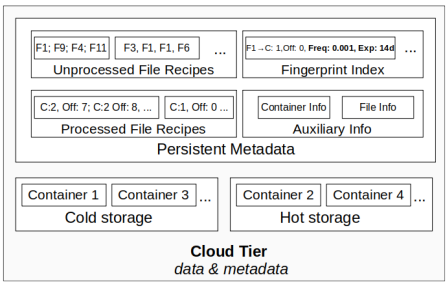

Figure 3: Architecture of data and metadata with two types of data storage. Fingerprint index is extended.

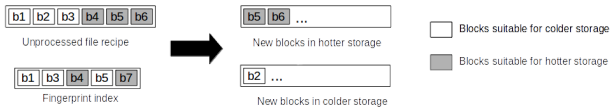

Figure 4: Writing blocks to more than one storage type. Block *b5* is written to hotter storage, despite the fact it is already available in colder storage if it brings a cost benefit due to frequent restores of *b5*.

affects the recovery time.[6] Therefore, we focused on colder storage which offers a reduced price of storing data but increases the price of restores, and demands a minimal storage period (e.g., 90 days). To utilize the storage effectively, we rely on two additional pieces of information provided with each file (in UFRs):

1. **Current expiration date**, as in *GC-Strategy #3*.
2. **Rough, expected frequency of file restore.**

As explained earlier, the expiration time is typically known. The restore frequency is unknown in advance but assessing the read frequency of file is a common practice for data kept in the cloud. For instance, Amazon explicitly recommends different storage classes for data accessed "once per quarter" and "1-2 times per year". In the specific case of backups, assessing restore frequency should be relatively easy, as a study of a large number of backup jobs [12] suggests that backup domains are divided into these with very frequent restores, sporadic restores, and virtually no restores. Moreover, particular backup policies clearly influence the restore frequency [54], and an upper bound on the restores can be calculated based on restore SLAs. Finally, modern backup software already implements tools that allow viewing historical data on the restore frequency of selected resources [71].

The persistent data and metadata structures are organized as shown in Fig. 3. The process of container writing during BatchDedup and BatchGC is extended, so each black can be stored in an appropriate cloud storage type, as shown in Fig. 4.

[6]Usage of our algorithms with the coldest storage services which lengthens the retrieval process is also possible. However, in such case providing an additional information regarding the required retrieval time of each file is necessary.

Each block is stored in a storage type for which the following formula has lower value:

$$t = COST_{insert} + (COST_{B/day} + COST_{restore} * FREQ_{restore}) * EXP_{time} \quad (2)$$

In the formula, the $COST_{insert}$ depends on cloud pricing, as well as the sizes of the block and its container, as the amortized cost of data insertion is included. The $COST_{B/day}$ describes the storage cost of the block. The $COST_{restore}$ depends on the data locality, as many blocks can be read with one request, so the upper bound for the $COST_{restore}$ can be calculated as *one request per block* or assessed with a heuristic. The $FREQ_{restore}$ and EXP_{time} are inherited from each file referencing block, and stored with each block in the fingerprint index.

However, further adjustments to $FREQ_{restore}$ and EXP_{time} are required. That is because the first decision about storage type must be taken when the block is stored for the first time, when blocks' $FREQ_{restore}$ and EXP_{time} are understated, as more references will come in the future. For instance, a block can be initially stored in cold storage but soon it receives more references and its actual restore frequency increases significantly. Vice-versa, data with short EXP_{time} can be kept in hot storage, despite the fact a reference with a larger EXP_{time} will come soon.

Therefore, both $FREQ_{restore}$ and EXP_{time} should be heuristically modified. A heuristic that worked very well in our experiments relies on block reference counts. First, we select a number R of expected references for each block (e.g., a hardcoded value 5 or value calculated from the system state). Then, we modify $FREQ_{restore}$ and EXP_{time} for blocks that have not reached the expected number based on the formula (e.g., we multiply it by $R - r$, where r is the actual number of references). In the end, the $FREQ_{restore}$ and EXP_{time} for newly written blocks are more similar to their future values.

In justified cases, a block can be stored in multiple storage types (e.g., when a block stored in cold storage receives a reference with high $FREQ_{restore}$), but BatchGC will eventually remove the unnecessary copies. Similarly, BatchGC can move a block from one type of storage to another (e.g., a reference with high restore frequency has been deleted). Generally, during BatchGC, a formula for calculating whether a container should be rewritten, takes into account the potential cost reduction caused by a change of the storage type. A decision on whether rewriting a particular container is profitable must be made for the whole container because rewriting the container also introduces costs. Nevertheless, blocks from one container can be moved to containers in various tiers (Fig. 5).

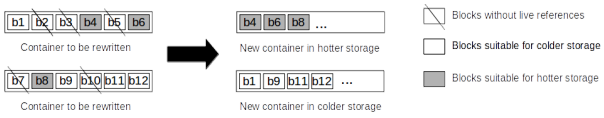

Figure 5: Rewriting containers to multiple types of storage.

5 Evaluation

We performed multiple experiments to evaluate InftyDedup and we present them in two parts. First, we evaluate the performance and cost of our implementation executed in a public cloud. Secondly, we evaluate our garbage collection and storage type selection strategies under various workloads.

5.1 Performance Evaluation

To evaluate the performance, we implemented InftyDedup using Apache Hive [21], which we selected as a possible approach to provide portability between different public and private clouds. We present results from the implementation of our batch algorithms, as uploading containers and restoring data are straightforward object storage operations in which the bottleneck is expected mostly on the network to the cloud (even a naive implementation can saturate 1 GbE network with uploads and restores using a single core).

Our batch algorithms are much different from the state-of-the-art tiering to cloud with deduplication techniques, therefore comparing with existing solutions was not possible. Instead, we just present the results using publicly available hardware. The evaluation was conducted in AWS using m5d.xlarge instances with 4x vCPU and 16 GiB of RAM. We aimed to use the smallest possible instances (to maximize the horizontal scaling) but in our workloads the technological stack of Apache Hive did not utilize the limited memory of the smallest instances efficiently. The selected instance type has 1x 150 GB NVMe which costs less than network attached EBS.

Presented experiments used synthetic data with the following characteristics. Each file contained approximately 51 GB (as backup files typically have tens of gigabytes or more [74]) chunked into blocks of approximately 64 KB (the target block size of the deduplication system for which we prepared InftyDedup). The content of the files is described with each experiment. We decided to present results with synthetically generated data, as our algorithms mostly distribute the data (e.g., based on fingerprints) and later sort the the data in small portions, so the exact characteristic of the data (e.g., the initial order of block) does not affect the performance much.

5.1.1 Batch Deduplication Processing

We evaluated BatchDedup in configurations varying in size. Each experiment consisted of two steps. In the first (initial) step, a large number of files without duplicates was processed to resemble a situation in which new backups were uploaded to the cloud. In the second (incremental) step, a dataset 3x smaller than the initial backup is uploaded to the cloud (as typically incremental backups are smaller than their full backups [12]), where 90% of the blocks are duplicates (which matches the expected average daily deduplication ratio [12]).

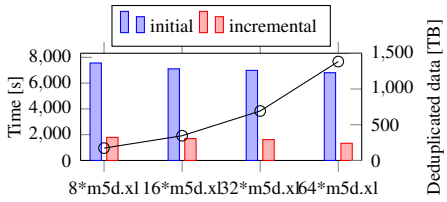

Figure 6: BatchDedup performance.

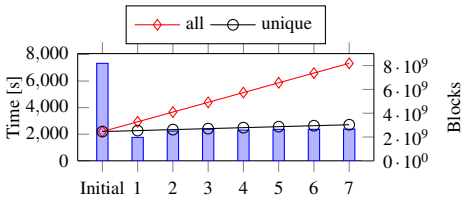

Figure 7: BatchDedup with growing data.

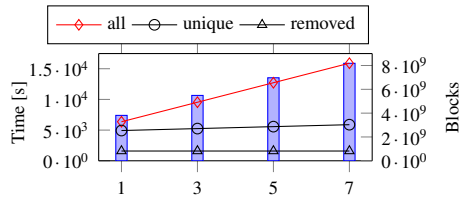

Figure 8: BatchGC performance.

The smallest configuration (8 instances) uploads 3072 files in the first step and 1024 in the second step. In larger configurations, the amount of data to be processed is scaled linearly with the system size. Therefore, the smallest test processed metadata of 208 TB data and the largest one of 1.66 PB.

In all configurations, the first step takes up between 1h25m and 2h10m (Fig. 6), and the second step takes up to 30m. Overall, the performance scales close to linearly. We analyzed the resource utilization, and the main bottleneck is the CPU, as most of the time its usage is above 95%. The network and the local NVMe drive are underutilized, with peak per-node usage of respectively 350 MB/s of network bandwidth and 6% of disk utilization. We expect that the performance can be further optimized but the computation cost is already marginal compared to the cost of data storage. For instance, in the experiment with 32 instances, the second stage eliminates 250 TB of duplicates and costs below $1, which is less than 0.1% of monthly savings on storage. Similarly, the costs of accessing in-cloud metadata during processing were marginal, as both steps required roughly 250K GETs ($0.1), 20K PUTs ($0.1), and transfer within one zone is free of charge.

We also conducted a different experiment with multiple steps of incremental uploads in one configuration (8 instances). As shown in Fig. 7, the computation time increases close to linearly with the amount of non-duplicate data which is added to the fingerprint index in each experiment.

5.1.2 Batch Garbage Collection Processing

First, we evaluated BatchGC by removing a fraction of data uploaded in the experiments described in Section 5.1.1. Specifically, we removed the data uploaded in the first step to resemble removing the oldest version of a backup. In each verified configuration, the processing took between 61 and 65 minutes.

BatchGC, unlike BarchDedup, reads all PFRs, so we also verify that the processing time increases close to linearly with the size of both fingerprint index and filerecipes as (Fig. 8). The results confirm that for data with many duplicates BatchGC is more expensive than BatchDedup. However, BatchGC is expected to be executed less frequently, so both algorithms will have comparable total execution costs.

5.2 Strategies Evaluation

We evaluated how our garbage collection and storage type selection strategies behave in a large number of workload simulations. The strategies optimize the costs of storing data for months and years, so we could not conduct these experiments in the public cloud, as it would take too long. Instead, we ran some initial experiments to confirm that we understand the pricing model and features of the cloud, and based on the results, we implemented a simulator. The simulator calculates costs based on cloud pricing of storage, requests, transfer, and other factors like the minimal storage duration.

Each experiment was conducted in numerous configurations of workload characteristics and system parameters. We present aggregated (minimal, maximal, and average) results, with values normalized to the result with the minimal cost.

5.2.1 Workload Characteristics

Our simulator allowed specifying following factors to evaluated various backup workloads:

Data source was selected from the following two sets. Firstly, we generated synthetic workloads in which a given fraction of data was *modified* and *deleted* for each day. Both types of modifications were applied in variable length *stream-contexts* (of size from 1 to 1024 blocks), so a given number of consecutive blocks was modified at once. Introduction of the stream-contexts was necessary, as data modified in small contexts is more fragmented, so the number of requests required to read is increased. Secondly, FSL traces [64] were used, as they are real-world datasets that contain information on how the data of multiple users change over the years.

Retention policy specifies how long each file (backup) is be stored. We analyzed a large number of documents and guidelines related to retention policies [1, 25, 66] to generate realistic policies. Typically, each type of backup is stored for a longer time than its backup period (e.g., weekly backups are kept for four weeks). In our experiments, daily backups are kept for one week, weekly backups are kept for a month, monthly backups are kept for a year, and yearly backups are kept for five years. Based on that, we came up with three different policies: *keepAll* policy in which all types of backups are stored in the cloud, *dailyExcluded* in which daily backups are excluded (so only backups stored for at least a month are kept in the cloud), and *dailyOnly* in which only daily backups

are kept in the cloud. In all experiments, the data were written for a period of 5 years.

Read patterns remarkably affect the total cost of ownership of data in the cloud. Unlike for writing data, we have not found any collected read traces for backup data. Similarly, there are no precise guidelines that describe typical backup read patterns. Therefore, we adapted a model in which each file is read with a given probability, and verified the full spectrum of possible values.

5.2.2 Garbage Collection Strategies Evaluation

To evaluatd how the proposed garbage collection strategies perform in different workloads, we conducted a large number of experiments with the pricing model of AWS S3 standard as *hot* storage and Glacier Instant Retrieval as *cold* storage.[7] We denoted the experiments in which the storage types are mixed as *mixed*. We verified our three strategies for garbage collection. Strategy #1 is denoted as *onlyEmpty*, *less*{*25, 50, 75, 99*} denotes Strategy #2 with the T parameter such that the behavior is equivalent to reclaiming space when less than *25, 50, 75, 99* percent of container capacity is used by live data, and Strategy #3 is denoted as *costBased*.

As shown in Fig. 9, *onlyEmpty* strategy achieved the worst results. In general, keeping data in hot storage was more expensive than in cold / mixed storage, which is expected since there were no reads in the experiments. For cold / mixed storage, *costBased* strategy gave significantly better results (on average 1.4%-23%), whereas for hot storage (where the rewrite cost is marginal) it gave similar results to *less99*.

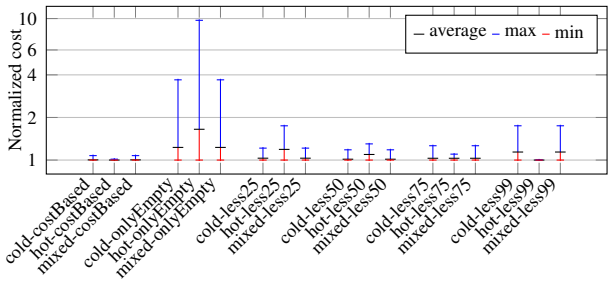

Figure 9: Garbage collection with different strategies.

In another set of experiments (Fig. 10), which included reads, there are more differences between the strategies. On average, *costBased* strategy is only 2.2% better, but comparing the worst cases, the difference is 24%. The analysis of the number of containers that are rewritten, deleted empty, or remain live at the end of the test, confirms that *onlyEmpty* has the largest number of containers that are live (Fig. 11).

[7] At the moment of writing, cold storage had 4/25x more expensive PUT/GET requests, 5.25x times cheaper storage costs, the minimum storage duration was 90 days, and an additional per-gigabyte retrieval cost for cold storage was equal fee for 3000 GET requests.

The analysis of garbage collection strategies led to the question of how container sizes affect the costs, as smaller containers increase the probability of removing the entire container but also increase the number of PUT requests needed to store data initially or during container rewriting. As shown in Fig. 12, for the *costBased* strategy, the lowest average cost is with 16 MB containers (4 MB and 64 MB are respectively 4.5% and 2% more expensive). The smallest, 1 MB containers were the most expensive, even with the *onlyEmpty* strategy, because with such small size the costs of the initial container creation prevail (Fig. 13). In case of cold storage, cost of PUT requests is significant, so storing the data in small containers in cold / mixed storage increases the cost up to 40%.

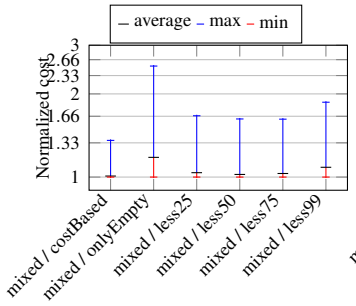

Figure 10: Garbage collection strategies with reads.

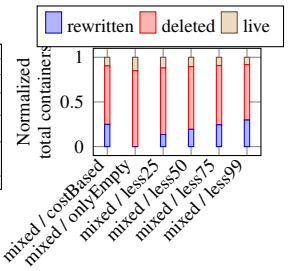

Figure 11: Normalized number of containers rewritten, deleted empty and live.

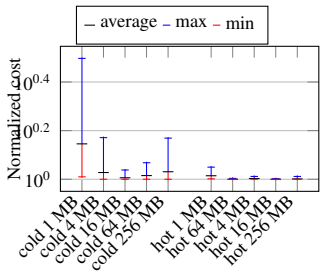

Figure 12: Garbage collection with varying container sizes.

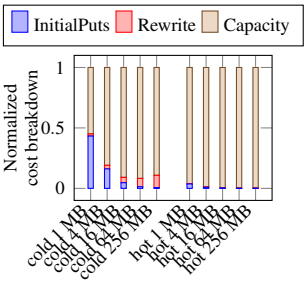

Figure 13: Cost breakdown with different container sizes.

5.2.3 Storage Storage Type Selection

We evaluated our storage type selection strategies in workloads with varying read frequency. For each experiment, there are 4 synthetically generated sets of files, and each set has a different read frequency: once a month, once a year, once a year with 1% probability, and once a year with 0.1% probability. All 4 sets were written together, just as in a storage system that keeps files with varying read frequencies. The experiments were conducted in series, and in each series the read frequency was scaled by a factor from 0.001 to 10. Therefore, cases in which reads are virtually not existent, in which reads dominates the total cost, and cases in-between were evaluated. A real-world ratio between backup and recovery

jobs is typically 100 : 1 [13] but varies depending on the system [12]. In our tests, the ratio of backups to recoveries for scale factor 0.01 is 70 – 700 : 1 (mean=216 : 1) depending on the retention policy, therefore we expect results with scale factors 0.01 and 0.1 to reflect a typical use-case.

As shown in Fig. 14, on average *mixed* strategy gives 55% cost savings in comparison to *cold* if there are many reads and 70% in comparison to *hot* if there are hardly any reads. The breakdown of newly created containers (Fig. 15) confirms that data ends up in cold storage when there are hardly any reads, and in hot storage when there are frequent reads. The cost breakdown (Fig. 16) confirms that our strategy for mixing storage types balances the costs between the expensive storage of hot data and the expensive reads of cold data.

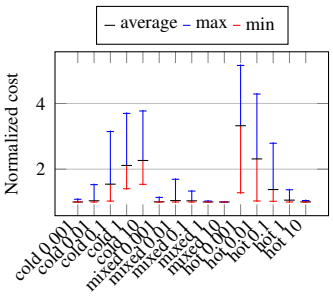

Figure 14: Storage type selection depending on read frequency.

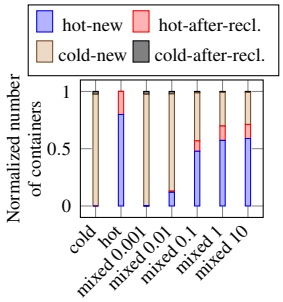

Figure 15: Number of containers created initially and after reclamation.

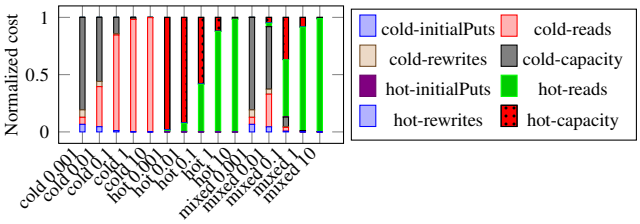

Figure 16: Cost breakdown with varying read frequencies.

We also evaluated how predicting the number of references in the future affects the cost. Fig. 17 presents the normalized cost, depending on the selection of expected references number. Mixing the storage types without predicting that more references will come, the cost is higher on average by 11% (worst case 277%) compared to predicting 5-10 references. For 3-10 references, the results are very similar (the lowest average cost was achieved with 10 and the lowest worst-case scenario with 5). Therefore, we confirmed that predicting the number of references brings significant cost reduction. However, in general the results are not very sensitive for slight changes in the expected number of references.

The mixed strategy depends on the expected frequency of reads, which may be incorrectly assessed. We conducted experiments with a significant prediction error (value underestimated and overestimated ten times). Even such large estimation error, the results are close to perfect (Fig. 18). Therefore, in all other experiments we assumed the perfect estimation, to facilitate studying the remaining experimental parameters.

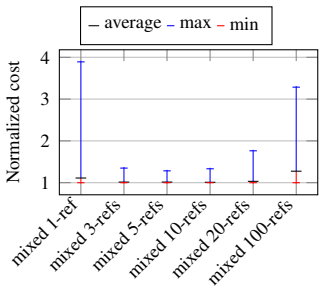

Figure 17: Cost of storing data depending on expected number of references.

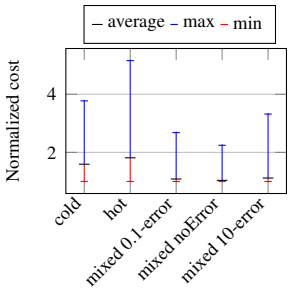

Figure 18: Cost of storing data depending on the error of frequency prediction.

5.2.4 Different Public Clouds

To confirm that our strategies are generally applicable to public clouds, we repeated most of the experiments with the pricing model of Google Cloud and Microsoft Azure. As our evaluation shows, mixing cold and hot storage reduces the costs for all three major providers (Fig. 19). The noticeable differences in gain between the cloud providers follow from the different ratios of costs, especially the cost of storing data and egress traffic. On average, keeping data only in hot storage is 66% more expensive, and keeping data only in cold storage is 30% more expensive than using the mixed strategy.

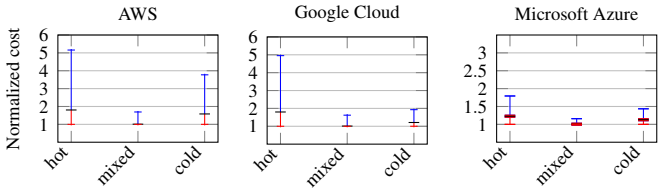

Figure 19: Comparison of gain from mixing storage types with different public cloud pricing.

5.2.5 FSL Traces

Finally, we verified our strategies using the FSL traces [64]. Specifically, we used the available homes snapshots dataset with 64 KB chunking. The traces contain metadata of files chunked during writing, but they have no information about the read pattern. Therefore, for each user, we verified how our storage type selection works with the varying number of reads (restoring each backup with a frequency from 0.0001 to 1 time a month). As shown in Fig. 20, at the extreme read frequencies the mixed strategy keeps almost all the data in

the cheapest of the two storage types. However, if the number of reads is in between, the mixed strategy works better than keeping data in a single type of storage, as depending on the data characterization a different decision should be taken for each block. In particular, the characterization of the reference number of each block is important, as frequently referenced blocks are accessed more often. Therefore, mixing the storage type can outperform keeping the data in one storage type, decreasing the cost by 26%-44%. This result shows that even when the restore frequency of each file is known in advance, relying on the selection of one storage type can be significantly more expensive than using our mixed strategy.

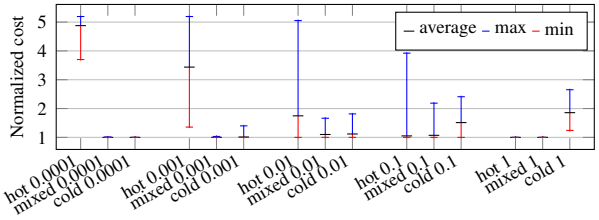

Figure 20: Storage cost in experiments with FSL traces.

6 Related Work

Hierarchical storage is widely adapted, as storage devices offer a trade-off between cost, capacity, and performance [52]. Systems with storage tiers are actively researched and, in recent years, a large number of publications refer to tiering in the cloud [36, 43, 55]. Hsu et al. [35] propose an AI-based prediction model for the classification of whether data is cold or hot. Liu et al. [42] describe an online algorithm for two-tier cloud storage which works without any prior knowledge of future access frequencies. However, less attention is given to tiering techniques in the cloud in the context of deduplication.

MUSE [76] is a framework focused on providing SLA for deduplicated data focused on the primary storage use case, which is a different use case than storing backups. DD Tier [27] is a tiering with deduplication that performs its computation in the local tier, hence imposing fundamental restrictions and limitations. Firstly, deduplication of data from more than one local tier system is not possible, as each system performs the deduplication on its own. Furthermore, all or at a least large fraction of metadata is needed locally to operate. Therefore, metadata are stored in both tiers, which not only increases storage capacity usage, but also forces downloading a large amount of metadata to recover even a single file. Moreover, the resources for metadata storage and processing of the local tier are limited. As locally stored metadata can consume hundred of terabytes of local storage, the size of the cloud tier is limited (to 2x the size of the local tier). Alike, deduplication and garbage collection algorithms cannot overuse scarce local resources, especially RAM, CPUs, and disk I/Os. Therefore, perfect hashing is used to decrease memory requirements below 3 bits per fingerprint, so extending it with techniques similar to our storage type selection is very difficult.

DD Tier introduces a technique for estimating how much space will be freed from the local tier after moving data to the cloud, and in recent years, significant research attention has been paid to the problem of selecting files for efficient data removal and migration in systems with deduplication [33, 39, 48]. As long as such methods do not require storing additional metadata locally, they can be used with InftyDedup.

A large number of publications explore the topic of security threats of deduplication in the cloud. Therefore, several methods of preventing particular attack types were proposed [18,38,40,77]. Alike, side channels leaking information from deduplication storage are studied [15, 17]. The majority of threats arise from the situation in which a public cloud provider implements the deduplication between users. InftyDedup is meant to be used by a single organization, and writing to InftyDedup requires accessing the local tier, so the situation is much different. Still, some organizations might find the deduplication side-channels as a threat within the organization, and adding security mechanisms to InftyDedup can be required. Nevertheless, users of InftyDedup may not trust the cloud provider, so the local tier can encrypt the blocks with actual data stored in InftyDedup. The structure of the data (information on which blocks are referenced by which files) is still exposed to allow the computations in the computations, but the situation is pretty much the same in any tiering with deduplication, as restoring blocks reveals the structure of files.

7 Conclusions

We presented InftyDedup, a novel, cloud-native approach to tiering to cloud for a storage system with deduplication. Compared to the state of the art, our architecture does not impose any limit on the size of the cloud tier and supports deduplication from multiple first-tier systems. We implemented InftyDedup for a commercial storage system (HYDRAstor) and evaluated it in a public cloud (AWS). The evaluation confirmed the desired scalability of deduplication handling: our batch algorithms, designed to reduce cloud costs and harness dynamic resource allocation, were able to process metadata of multi-petabyte data collections for a couple of dollars.

To further decrease the cost of cloud storage, we proposed an extension to InftyDedup, which moves chunks between hot and cold cloud stores based on their anticipated access patterns. Its evaluation with real-world traces showed that our deduplication-specific heuristic for adjusting the expected read frequency, which takes into account block reference counts, decreased the costs on average by 11%, and the overall solution achieved 26%–44% reductions. The algorithm requires minimal input from a system administrator and was demonstrated to retain its cost benefits even when the administrator's estimations were under- or over-estimated.

Acknowledgments

We sincerely thank our shepherd Philip Shilane and the anonymous reviewers for helping us improve our paper significantly.

References

[1] Acronis. Retention rules: how and when they work. 2022. https://kb.acronis.com/content/68304.

[2] Yamini Allu, Fred Douglis, Mahesh Kamat, Ramya Prabhakar, Philip Shilane, and Rahul Ugale. {Can't} we all get along? redesigning protection storage for modern workloads. In *2018 USENIX Annual Technical Conference (USENIX ATC 18)*, pages 705–718, 2018.

[3] Yamini Allu, Fred Douglis, Mahesh Kamat, Philip Shilane, Hugo Patterson, and Ben Zhu. Backup to the future: How workload and hardware changes continually redefine data domain file systems. *Computer*, 50(7):64–72, 2017.

[4] Amazon. Amazon aurora - fully mysql and postgresql compatibile managed database service. 2022. https://aws.amazon.com/rds/aurora/?did=ap_card&trk=ap_card.

[5] Amazon. Amazon elastic container service pricing. 2022. https://aws.amazon.com/ecs/pricing/.

[6] Amazon. Amazon s3 pricing. 2022. https://aws.amazon.com/s3/pricing/.

[7] Amazon. Aws direct connect locations. 2022. https://aws.amazon.com/directconnect/locations/.

[8] Amazon. Aws lambda - faqs. 2022. https://aws.amazon.com/lambda/faqs/.

[9] Amazon. Aws snow family faqs. 2022. https://aws.amazon.com/snow/faqs/.

[10] Amazon. General purpose instances - network performance. 2022. https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/general-purpose-instances.html#general-purpose-network-performance.

[11] Amazon. Link aggregation groups. 2022. https://docs.aws.amazon.com/directconnect/latest/UserGuide/lags.html.

[12] George Amvrosiadis and Medha Bhadkamkar. Identifying trends in enterprise data protection systems. In *2015 USENIX Annual Technical Conference (USENIX ATC 15)*, pages 151–164, 2015.

[13] George Amvrosiadis and Medha Bhadkamkar. Getting back up: Understanding how enterprise data backups fail. In *2016 USENIX Annual Technical Conference (USENIX ATC 16)*, pages 479–492, 2016.

[14] Associate Research Director Andrew Smith, Research Manager; Archana Venkatraman. Enterprise data growthandadoption of cloud applications challenge traditional data protection strategies. 2021. https://afi.ai/r/US48310921.pdf.

[15] Frederik Armknecht, Colin Boyd, Gareth T Davies, Kristian Gjøsteen, and Mohsen Toorani. Side channels in deduplication: Trade-offs between leakage and efficiency. In *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security*, pages 266–274, 2017.

[16] AWS. Best practices for cluster configuration. 2022. https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-plan-instances-guidelines.html.

[17] Andrei Bacs, Saidgani Musaev, Kaveh Razavi, Cristiano Giuffrida, and Herbert Bos. DUPEFS: Leaking data over the network with filesystem deduplication side channels. In *20th USENIX Conference on File and Storage Technologies (FAST 22)*, pages 281–296, Santa Clara, CA, February 2022. USENIX Association.

[18] Mihir Bellare, Sriram Keelveedhi, and Thomas Ristenpart. Message-locked encryption and secure deduplication. In *Annual international conference on the theory and applications of cryptographic techniques*, pages 296–312. Springer, 2013.

[19] Christophe Bertrand. Esg research report: The evolution of data protection cloud strategies. 2021. https://www.esg-global.com/research/esg-research-report-the-evolution-of-data-protection-cloud-strategies.

[20] Sriprasad Bhata. Introducing azure cool blob storage. 2022. https://azure.microsoft.com/en-us/blog/introducing-azure-cool-storage/.

[21] Jesús Camacho-Rodríguez, Ashutosh Chauhan, Alan Gates, Eugene Koifman, Owen O'Malley, Vineet Garg, Zoltan Haindrich, Sergey Shelukhin, Prasanth Jayachandran, Siddharth Seth, et al. Apache hive: From mapreduce to enterprise-grade big data warehousing. In *Proceedings of the 2019 International Conference on Management of Data*, pages 1773–1786, 2019.

[22] Google Cloud. Storage classes. 2022. https://cloud.google.com/storage/docs/storage-classes#descriptions.

[23] Biplob Debnath, Sudipta Sengupta, and Jin Li. {ChunkStash}: Speeding up inline storage deduplication using flash memory. In *2010 USENIX Annual Technical Conference (USENIX ATC 10)*, 2010.

[24] Fred Douglis, Abhinav Duggal, Philip Shilane, Tony Wong, Shiqin Yan, and Fabiano Botelho. The logic of physical garbage collection in deduplicating storage. In *15th USENIX Conference on File and Storage Technologies (FAST 17)*, pages 29–44, 2017.

[25] Druva. What is backup retention policy? how is it implemented? 2022. https://docs.druva.com/Knowledge_Base/inSync/Client/010_FAQ/What_is_Backup_Retention_Policy%3F_How_is_it_implemented%3F.

[26] Cezary Dubnicki, Leszek Gryz, Lukasz Heldt, Michal Kaczmarczyk, Wojciech Kilian, Przemyslaw Strzelczak, Jerzy Szczepkowski, Cristian Ungureanu, and Michal Welnicki. HYDRAstor: A scalable secondary storage. In *FAST*, volume 9, pages 197–210, 2009.

[27] Abhinav Duggal, Fani Jenkins, Philip Shilane, Ramprasad Chinthekindi, Ritesh Shah, and Mahesh Kamat. Data domain cloud tier: Backup here, backup there, deduplicated everywhere! In *2019 USENIX Annual Technical Conference (USENIX ATC 19)*, pages 647–660, 2019.

[28] Nnamdi Ekwe-Ekwe and Adam Barker. Location, location, location: exploring amazon ec2 spot instance pricing across geographical regions. In *2018 18th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID)*, pages 370–373. IEEE, 2018.

[29] Google. Google cloud platform managed relational database. 2022. https://cloud.google.com/sql.

[30] Google. Google cloud storage pricing. 2022. https://cloud.google.com/storage/pricing#north-america.

[31] Google. Vertex ai - google cloud's unified ml platform. 2022. https://cloud.google.com/vertex-ai.

[32] Fanglu Guo and Petros Efstathopoulos. Building a high-performance deduplication system. In *2011 USENIX Annual Technical Conference (USENIX ATC 11)*, 2011.

[33] Danny Harnik, Moshik Hershcovitch, Yosef Shatsky, Amir Epstein, and Ronen Kat. Sketching volume capacities in deduplicated storage. In *17th USENIX Conference on File and Storage Technologies (FAST 19)*, pages 107–119, Boston, MA, February 2019. USENIX Association.

[34] HPE. Hpe cloud bank storage: A data protection solution you can bank on. 2017. https://community.hpe.com/t5/Around-the-Storage-Block/HPE-Cloud-Bank-Storage-A-Data-Protection-Solution-You-Can-Bank/ba-p/6965903.

[35] Ying-Feng Hsu, Ryo Irie, Shuuichirou Murata, and Morito Matsuoka. A novel automated cloud storage tiering system through hot-cold data classification. In *2018 IEEE 11th International Conference on Cloud Computing (CLOUD)*, pages 492–499, 2018.

[36] Ryo Irie, Shuuichirou Murata, Ying-Feng Hsu, and Morito Matsuoka. A novel automated tiered storage architecture for achieving both cost saving and qoe. In *2018 IEEE 8th International Symposium on Cloud and Service Computing (SC2)*, pages 32–40. IEEE, 2018.

[37] Michal Kaczmarczyk, Marcin Barczynski, Wojciech Kilian, and Cezary Dubnicki. Reducing impact of data fragmentation caused by in-line deduplication. In *Proceedings of the 5th Annual International Systems and Storage Conference*, pages 1–12, 2012.

[38] Sriram Keelveedhi, Mihir Bellare, and Thomas Ristenpart. DupLESS: Server-Aided encryption for deduplicated storage. In *22nd USENIX Security Symposium (USENIX Security 13)*, pages 179–194, Washington, D.C., August 2013. USENIX Association.

[39] Roei Kisous, Ariel Kolikant, Abhinav Duggal, Sarai Sheinvald, and Gala Yadgar. The what, the from, and the to: The migration games in deduplicated systems. In *20th USENIX Conference on File and Storage Technologies (FAST 22)*, pages 265–280, Santa Clara, CA, February 2022. USENIX Association.

[40] Jingwei Li, Chuan Qin, Patrick P. C. Lee, and Jin Li. Rekeying for encrypted deduplication storage. In *2016 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, pages 618–629, 2016.

[41] Mark Lillibridge, Kave Eshghi, and Deepavali Bhagwat. Improving restore speed for backup systems that use inline {Chunk-Based} deduplication. In *11th USENIX Conference on File and Storage Technologies (FAST 13)*, pages 183–197, 2013.

[42] Mingyu Liu, Li Pan, and Shijun Liu. To transfer or not: An online cost optimization algorithm for using two-tier storage-as-a-service clouds. *IEEE Access*, 7:94263–94275, 2019.

[43] Yaser Mansouri and Abdelkarim Erradi. Cost optimization algorithms for hot and cool tiers cloud storage services. In *2018 IEEE 11th International Conference on Cloud Computing (CLOUD)*, pages 622–629, 2018.

[44] Dirk Meister and André Brinkmann. Multi-level comparison of data deduplication in a backup scenario. In *Proceedings of SYSTOR 2009: The Israeli Experimental Systems Conference*, pages 1–12, 2009.

[45] Dirk Meister and André Brinkmann. dedupv1: Improving deduplication throughput using solid state drives (ssd). In *2010 IEEE 26th Symposium on Mass Storage Systems and Technologies (MSST)*, pages 1–6. IEEE, 2010.

[46] Microsoft. Microsoft azure storage pricing. 2022. https://azure.microsoft.com/en-us/pricing/details/storage/blobs/#pricing.

[47] Jaehong Min, Daeyoung Yoon, and Youjip Won. Efficient deduplication techniques for modern backup operation. *IEEE Transactions on Computers*, 60(6):824–840, 2010.

[48] Aviv Nachman, Gala Yadgar, and Sarai Sheinvald. GoSeed: Generating an optimal seeding plan for deduplicated storage. In *18th USENIX Conference on File and Storage Technologies (FAST 20)*, pages 193–207, Santa Clara, CA, February 2020. USENIX Association.

[49] Netapp. Step 7. configure long-term retention. 2021. https://docs.netapp.com/us-en/cloud-manager-tiering/pdfs/fullsite-sidebar/Cloud_Tiering_documentation.pdf.

[50] Veritas NetBackup. About netbackup worm storage support for immutable and indelible data. 2020. https://www.veritas.com/support/en_US/doc/25074086-143197427-0/v143250065-143197427.

[51] Veritas NetBackup. Veritas netbackup™ deduplication guide. 2021. https://www.veritas.com/support/en_US/doc/25074086-146020141-0/v145698641-146020141.

[52] Junpeng Niu, Jun Xu, and Lihua Xie. Hybrid storage systems: A survey of architectures and algorithms. *IEEE Access*, 6:13385–13406, 2018.

[53] Myoungwon Oh et al. Design of global data deduplication for a scale-out distributed storage system. In *2018 IEEE 38th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 2018.

[54] Oracle. Backing up file-system data. 2022. https://docs.oracle.com/cd/E91325_01/OBADM/osb_filesystem_backup.htm.

[55] Ajaykrishna Raghavan, Abhishek Chandra, and Jon B Weissman. Tiera: Towards flexible multi-tiered cloud storage instances. In *Proceedings of the 15th International Middleware Conference*, pages 1–12, 2014.

[56] Santhosh Rao, Nik Simpson, Michael Hoeck, and Jerry Rozeman. Magic quadrant for enterprise backup and recovery software solutions. 2021. https://www.gartner.com/en/documents/4003661.

[57] Meet Shah, Mohammedhasan Shaikh, Vishwajeet Mishra, and Grinal Tuscano. Decentralized cloud storage using blockchain. In *2020 4th International conference on trends in electronics and informatics (ICOEI)(48184)*, pages 384–389. IEEE, 2020.

[58] Pratima Sharma, Rajni Jindal, and Malaya Dutta Borah. Blockchain-based decentralized architecture for cloud storage system. *Journal of Information Security and Applications*, 62:102970, 2021.

[59] Konstantin Shvachko, Hairong Kuang, Sanjay Radia, and Robert Chansler. The hadoop distributed file system. In *2010 IEEE 26th Symposium on Mass Storage Systems and Technologies (MSST)*, pages 1–10, 2010.

[60] Przemyslaw Strzelczak, Elzbieta Adamczyk, Urszula Herman-Izycka, Jakub Sakowicz, Lukasz Slusarczyk, Jaroslaw Wrona, and Cezary Dubnicki. Concurrent deletion in a distributed {Content-Addressable} storage system with global deduplication. In *11th USENIX Conference on File and Storage Technologies (FAST 13)*, pages 161–174, 2013.

[61] Zhen Sun, Geoff Kuenning, Sonam Mandal, Philip Shilane, Vasily Tarasov, Nong Xiao, et al. A long-term user-centric analysis of deduplication patterns. In *2016 32nd Symposium on Mass Storage Systems and Technologies (MSST)*, pages 1–7. IEEE, 2016.

[62] Zhen "Jason" Sun, Geoff Kuenning, Sonam Mandal, Philip Shilane, Vasily Tarasov, Nong Xiao, and Erez Zadok. Cluster and single-node analysis of long-term deduplication patterns. *ACM Transactions on Storage (TOS)*, 14(2):1–27, 2018.

[63] Liyin Tang and Namit Jain. Join strategies in hive. *Hive Summit*, 2011.

[64] Vasily Tarasov, Amar Mudrankit, Will Buik, Philip Shilane, Geoff Kuenning, and Erez Zadok. Generating realistic datasets for deduplication analysis. In *2012 USENIX Annual Technical Conference (USENIX ATC 12)*, pages 261–272, 2012.

[65] Lionel Sujay Vailshery. Cloud infrastructure services vendor market share worldwide from 4th quarter 2017 to 4th quarter 2021. 2022.

https://www.statista.com/statistics/967365/worldwide-cloud-infrastructure-services-market-share-vendor/.

[66] Global Data Vault. Data backup: Developing an effective data retention. 2022. https://www.globaldatavault.com/blog/data-retention-policy-and-scheduled-backups/.

[67] Vinod Kumar Vavilapalli, Arun C Murthy, Chris Douglas, Sharad Agarwal, Mahadev Konar, Robert Evans, Thomas Graves, Jason Lowe, Hitesh Shah, Siddharth Seth, et al. Apache hadoop yarn: Yet another resource negotiator. In *Proceedings of the 4th annual Symposium on Cloud Computing*, pages 1–16, 2013.

[68] Veeam. Cloud object storage deep dive – part two, implementation. 2021. https://www.veeam.com/blog/cloud-object-storage-implementation.html.

[69] Veeam. Step 7. configure long-term retention. 2021. https://helpcenter.veeam.com/docs/backup/vsphere/backup_job_gfs_vm.html?ver=110.

[70] Veeam. 2022 data protection trends strategies. 2022. https://go.veeam.com/wp-data-protection-trends-2022.

[71] Veeam. Restore operator activity. 2022. https://helpcenter.veeam.com/docs/one/reporter/restore_operator_activity.html?ver=110.

[72] Veritas. Aws cloud storagewith veritas netbackup. 2022. https://www.veritas.com/content/dam/www/en_us/documents/white-papers/WP_aws_cloud_storage_with_netbackup_long_term_retention_solution_V1259.pdf.

[73] Marcia Villalba. Amazon s3 glacier is the best place to archive your data – introducing the s3 glacier instant retrieval storage class. 2022. https://aws.amazon.com/blogs/aws/amazon-s3-glacier-is-the-best-place-to-archive-your-data-introducing-the-s3-glacier-instant-retrieval-storage-class/.

[74] Grant Wallace, Fred Douglis, Hangwei Qian, Philip Shilane, Stephen Smaldone, Mark Chamness, and Windsor Hsu. Characteristics of backup workloads in production systems. In *FAST*, volume 12, pages 4–4, 2012.

[75] Wen Xia, Hong Jiang, Dan Feng, Fred Douglis, Philip Shilane, Yu Hua, Min Fu, Yucheng Zhang, and Yukun Zhou. A comprehensive study of the past, present, and future of data deduplication. *Proceedings of the IEEE*, 104(9):1681–1710, 2016.

[76] Jianwei Yin, Yan Tang, Shuiguang Deng, Bangpeng Zheng, and Albert Y. Zomaya. Muse: A multi-tierd and sla-driven deduplication framework for cloud storage systems. *IEEE Transactions on Computers*, 70(5):759–774, 2021.

[77] Haoran Yuan, Xiaofeng Chen, Jin Li, Tao Jiang, Jianfeng Wang, and Robert H Deng. Secure cloud data deduplication with efficient re-encryption. *IEEE Transactions on Services Computing*, 15(1):442–456, 2019.

[78] Zerto. Maximize recovery achieve your best rtos and rpos. 2020. https://www.zerto.com/wp-content/uploads/2020/08/Fastest-RTO-and-RPO-in-the-Industry_Guide.pdf.

[79] Zerto. Deploy & configure zerto long-term retention amazon s3. 2022. https://www.zerto.com/page/deploy-configure-zerto-long-term-retention-amazon-s3/.

[80] Benjamin Zhu, Kai Li, and R Hugo Patterson. Avoiding the disk bottleneck in the data domain deduplication file system. In *Fast*, volume 8, pages 269–282, 2008.