

BPF-oF: Storage Function Pushdown Over the Network

Ioannis Zarkadas^{*,1}, Tal Zussman^{*,1}, Jeremy Carin¹, Sheng Jiang¹, Yuhong Zhong¹, Jonas Pfefferle³, Hubertus Franke³, Junfeng Yang¹, Kostis Kaffes¹, Ryan Stutsman², and Asaf Cidon¹

¹Columbia University, ²University of Utah, ³IBM, * denotes equal contribution

Abstract

Storage disaggregation, wherein storage is accessed over the network, is popular because it allows applications to independently scale storage capacity and bandwidth based on dynamic application demand. However, the added network processing introduced by disaggregation can consume significant CPU resources. In many storage systems, logical storage operations (e.g., lookups, aggregations) involve a series of simple but dependent I/O access patterns. Therefore, one way to reduce the network processing overhead is to execute dependent series of I/O accesses at the remote storage server, reducing the back-and-forth communication between the storage layer and the application. We refer to this approach as *remote-storage pushdown*. We present BPF-oF, a new remote-storage pushdown protocol built on top of NVMe-oF, which enables applications to safely push custom eBPF storage functions to a remote storage server.

The main challenge in integrating BPF-oF with storage systems is preserving the benefits of their client-based in-memory caches. We address this challenge by designing novel caching techniques for storage pushdown, including splitting queries into separate in-memory and remote-storage phases and periodically refreshing the client cache with sampled accesses from the remote storage device. We demonstrate the utility of BPF-oF by integrating it with three storage systems, including RocksDB, a popular persistent key-value store that has no existing storage pushdown capability. We show BPF-oF provides significant speedups in all three systems when accessed over the network, for example improving RocksDB’s throughput by up to $2.8\times$ and tail latency by up to $2.6\times$.

1 Introduction

Storage disaggregation enables accessing block storage devices over the network and scaling independently of the application [18, 26, 30, 31, 37, 40, 41]. As datacenter networks become faster ($\sim 10\ \mu\text{s}$ latency), accessing storage devices quickly over the network becomes feasible. Similarly, datacenter networks offer ample bandwidth to saturate the I/O bandwidth of storage devices [30, 31, 40]. Consequently, a variety

of data-intensive datacenter applications (e.g., transactional databases [25, 45], data analytics [13], key-value stores [21], and data warehouses [46]) use storage disaggregation.

NVMe-oF (NVMe over Fabrics) is the standard networked storage protocol for storage disaggregation. NVMe-oF allows an application to transparently access an NVMe block device on a remote server. NVMe-oF supports TCP and RDMA, is fully supported by Linux, and can be used without specialized hardware. The major downside of NVMe-oF (compared to accessing a device locally via NVMe) is the added network latency and the CPU cost of processing the network packets. The CPU cost is especially substantial in the case of TCP, incurring a nearly 50% throughput decrease compared to local performance in our experiments (§2.2). We show that even in the case of RDMA, where this cost can be partially mitigated by offloading network processing to NICs, network processing still incurs a significant toll. With the goal of reducing this cost, we observe that in modern storage systems, many queries, such as key-value lookups and aggregations, involve a series of dependent I/O accesses (e.g., a B-tree index lookup), which translate into multiple round-trip accesses in the disaggregated setup. Therefore, one approach to amortize the network processing cost of storage disaggregation is to “pushdown” dependent I/O accesses to the remote storage server, thereby eliminating the need to process multiple round-trip network requests for each logical query. This idea itself is not new, and there are many examples of academic and commercial systems that push down functions over the network to operate closer to the data [5, 16, 32, 33, 36, 52]. However, these systems all implement this capability as an *application-specific feature*. Restructuring a storage system or database to allow it to run its operations on a remote server may entail many intrusive modifications, and it requires developers to contend with how to implement such a capability while maintaining security and concurrency guarantees.

We take a cue from recent work [32, 51, 54, 55] that has explored the use of eBPF (extended Berkeley Packet Filter) as a general-purpose framework for safely executing user-defined and kernel-verified storage functions closer to the

storage device within the Linux NVMe driver [54, 55] or a programmable storage device [51], thereby reducing the CPU overhead of traversing the kernel storage stack. However, these prior works assume pushdown occurs locally (i.e. on the same server) or on simple storage systems.

Two significant challenges arise when trying to adopt a similar approach in a disaggregated storage setting with a modern storage system. First, due to the design of block-level protocols like NVMe-oF, the programmable function that executes on the remote server operates below the file system layer (i.e. at the NVMe driver). However, user-defined storage functions, such as index traversals, operate on files. Prior work on local eBPF storage pushdown [54] has solved this by synchronously synchronizing the file-to-block mappings between the file system and NVMe driver. However, such an approach would be very slow for a networked setting, as all storage functions would block until the metadata synchronizes between the client and server. Second, in modern storage systems, data is stored in large data structures (e.g., LSM-trees or B-trees), and parts of these data structures may be partially (or fully) cached in memory and are frequently updated. However, in a networked setting, these in-memory structures would not be accessible at the remote server, which runs the I/O functions. Thus, if we use remote-storage pushdown, we may lose the benefit of these in-memory caches, hurting performance.

We solve these challenges by designing BPF-oF, a new *general-purpose* remote-storage pushdown protocol for TCP and RDMA that is built on top of NVMe-oF, allowing a client to submit custom eBPF storage functions and safely trigger them at a remote storage server. BPF-oF’s design relies on a key observation: modern storage systems often know in advance which files they might access to satisfy a storage operation. Thus, we tackle the first challenge by requiring the application to specify upfront which files it will access throughout its storage operation. BPF-oF uses this information to perform the file-descriptor-to-inode translation at the client. To keep the file-to-block mappings up to date, BPF-oF updates them asynchronously, thereby avoiding blocking concurrent storage functions. To avoid conflicts when storage functions execute concurrently with metadata synchronization, BPF-oF applies a version to each inode’s file-to-block mapping, and it checks for modified versions before issuing and after completing a request. We implement BPF-oF as a new NVMe command that is triggered by a system call from userspace, which also registers the required eBPF functions for the storage pushdown.

We tackle the second challenge using a new technique called *query splitting*, which divides queries into an in-memory portion and an on-disk portion. Accessing in-memory data structures is often orders of magnitude cheaper than a single storage I/O. Therefore, the in-memory portion of the query speculatively checks all the potentially-required in-memory (cached) data structures at the client (some of which may eventually not be needed), and then, based on the

results, executes the remaining on-disk portion of the query at the remote server. While this approach may result in “unnecessary” accesses to in-memory data structures (e.g., looking up bloom filters for files that will never be used), it makes the on-disk portion of the query much more efficient by grouping all I/O accesses together. Another caching-related problem is that the client cache may become stale, as it is not updated with the data that is accessed by the remote function. One potential solution to this problem is to simply return the intermediate data accessed by the remote storage query, but this would increase the network consumption between the client and remote server, negating much of the benefit of BPF-oF. Instead, we solve this problem using a lightweight technique of *cache sampling*, where a small percentage (by default 1%) of storage functions do not get offloaded, and fully update the client cache, keeping it fresh.

Our experiments yield another surprising result: for some workloads, maintaining a client data cache with BPF-oF actually *hurts* performance, because BPF-oF is very CPU efficient and the cost of maintaining the client cache is higher than its benefit. For example, when we run BPF-oF with a uniform workload on fast storage devices, it is in fact better to use BPF-oF with no cache (beyond indices) at the client!

We integrate BPF-oF with three storage systems: RocksDB [11] a popular and performance-optimized storage system used by Meta and many other companies [12], BPF-KV [54], a toy eBPF-friendly key-value store, and WiredTiger [15], a simple storage engine used by MongoDB. Our integration required a relatively modest change to RocksDB and WiredTiger: about 1,900 and 500 lines of code, respectively. BPF-oF provides significant performance boosts to all three systems: under the vast majority of workloads BPF-oF accelerates RocksDB by up to $2.8\times$ and reduces 99th-percentile latency by up to $2.6\times$; it accelerates BPF-KV by up to $8\times$; and it improves WiredTiger’s throughput by up to 30%. These gains are most apparent in read-heavy workloads and workloads where the working set does not fit in cache, but they translate to other workloads as well. However, as expected, the utility of storage pushdown is lower in write-heavy workloads or when almost all requests are serviced from in-memory data structures. We will make BPF-oF and our storage system integrations open source upon publication. Our major contributions are:

- BPF-oF, an efficient mechanism to safely enable storage pushdown for storage applications running on disaggregated storage. BPF-oF is the first implementation of storage pushdown in a general networked storage protocol, and its key ideas are an efficient and asynchronous file metadata synchronization mechanism and a versioning scheme for inode mapping to protect concurrent access to files. BPF-oF greatly reduces the CPU and network overhead of NVMe-oF while requiring no custom hardware.
- Our novel query splitting and cache sampling techniques, which allow BPF-oF to support local in-memory caching,

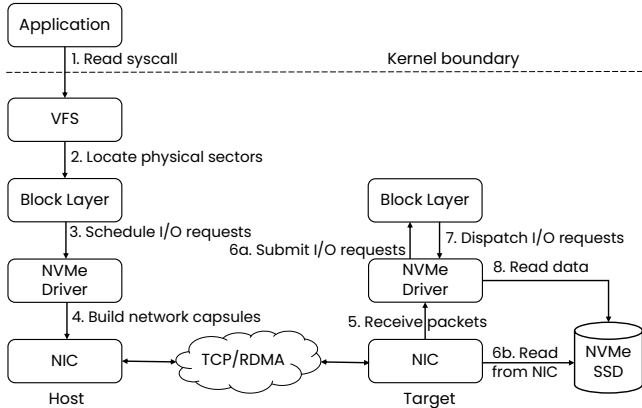


Figure 1: NVMe-oF overview.

address a significant challenge in storage pushdown. We integrate BPF-oF with modern storage systems, including the highly-optimized RocksDB.

2 Background and Motivation

This section first provides a background on NVMe-oF. It then discusses the CPU overhead of NVMe-oF, providing the motivation for remote storage pushdown. Finally, we review eBPF as a framework for storage pushdown.

2.1 NVMe-oF Primer

NVMe-oF has emerged as the main networked storage protocol for disaggregated storage, replacing iSCSI due to its better performance [26, 29]. NVMe-oF allows an application to directly access a block storage device connected to a remote server using NVMe. Figure 1 shows the flow of an NVMe-oF read request. The *host* (i.e. the client, left) initiates the request, and the *target* (i.e. the server, right) contains the SSD and performs the actual disk I/O. To initiate an NVMe-oF request, an application at the host issues a storage system call, such as READ (step 1), which then traverses its local OS storage stack (steps 2-3). The request is treated as a regular NVMe request until it reaches the local NVMe driver. The *host* and *target* drivers maintain I/O queues for exchanging the NVMe-oF capsule, which is the unit of communication between the host and the target (i.e. a capsule is to NVMe-oF what a frame is to Ethernet). The NVMe driver handles the request by constructing an NVMe-oF command within a capsule and submitting it to an NVMe I/O queue. The capsule is then forwarded to the relevant network stack (step 4) depending on the fabric type (TCP, RDMA, Fiber Channel), and it is transmitted to the target. At the target (step 5), after the driver extracts the NVMe-oF command, it generates the block layer request and submits it to the block layer for I/O scheduling (step 6a). The target’s NVMe driver, at last, receives the I/O request from the block layer (step 7) and reads the user’s data from the local NVMe SSD (step 8), which is then sent back to the host via the fabric-specific mechanism (not shown in the figure). Major NIC model lines (e.g., NVIDIA ConnectX, Broadcom Stingray, Intel IPU) support offloading the NVMe-oF target




Figure 2: RocksDB run on local disk compared to NVMe/RDMA (with offloading) and NVMe/TCP on a NAND SSD.

datapath to the NIC when the underlying protocol is RDMA, and some NICs even support TCP offload [1]. This allows the NIC to directly read from or write to the NVMe device, bypassing the target’s CPU (step 6b).

2.2 Improving NVMe-oF CPU Efficiency

A significant drawback of NVMe-oF is the CPU overhead of processing the packets, both at the host and the target [26]. This is especially acute in the case of NVMe/TCP, due to the relatively high CPU cost of processing TCP packets [28, 31, 53]. High CPU consumption can adversely affect the performance of data-intensive applications, especially when they are CPU-bound.

To demonstrate the network processing cost of NVMe-oF, we compare running two workloads on a 100 GB RocksDB instance with a NAND SSD locally to NVMe/RDMA (with NIC offloading) and NVMe/TCP (Figure 2). Both workloads request objects by keys drawn uniformly at random: the first is read-only, and the second is 50% reads, 50% updates. Each RocksDB instance uses 1 GB of cache and 3 cores locally at the host. The remote target is also allocated 3 cores, which are only partially used by the NVMe/TCP setup. Our experimental setup on CloudLab is fully described in §6.1.

By running on a local NAND SSD, RocksDB achieves a 1.5–1.6× throughput improvement compared to running on the same NAND SSD accessed via NVMe/TCP. While using RDMA is more CPU-efficient than TCP, as expected, NVMe/RDMA is still 12–15% slower than the local disk. The primary cause of NVMe-oF’s slowdown is the CPU overhead of processing the packets at the host, as RocksDB is often bottlenecked by CPU. In addition, NVMe-oF/TCP consumes additional CPU cycles at the target. Similarly, running RocksDB with NVMe/TCP incurs a significant latency overhead. While NVMe-oF provides significant operational advantages, it incurs a non-negligible performance overhead due to the extra network processing cost and communication time.

Motivation for storage pushdown. A promising way to reduce the CPU and network consumption of NVMe-oF is to “push down” computation to the target storage server. This reduces potential back-and-forth communication between the target and host. Storage function pushdown has been explored extensively in the database community as “predicate pushdown” [27, 38, 50], in the systems community as near-storage compute [16, 33, 36, 51, 52] and, more recently, as user-defined

functions in the kernel [17, 32, 54, 55]. For more details, see §7. However, to the best of our knowledge, no prior work has combined storage function pushdown with a general-purpose networked storage protocol such as NVMe-oF.

Key-value stores such as RocksDB are amenable to such an approach for two reasons. First, many RocksDB read operations are composed of a sequence of dependent I/O operations. For example, in the previous experiment, RocksDB issues on average 3.8 I/Os and 4 I/Os per read request, for the read-only and 50-50 read-write workload, respectively. Most read requests must access multiple files, since files are spread across multiple levels of RocksDB’s on-disk tree data structure. Second, RocksDB only writes data to disk in large immutable files. This significantly simplifies synchronization between potentially-conflicting simultaneous reads and writes. We will expand on this point later in the paper. Since many SSD-optimized storage systems exhibit these two properties [6, 13–15, 35, 39, 43, 44], pushing computation closer to the device can alleviate the overheads of accessing storage devices over the network for a wide variety of deployments.

2.3 Local eBPF-based Storage Pushdown

The idea of storage pushdown can be applied locally within a single host. XRP [54] is an eBPF framework that adds a hook to the Linux kernel’s NVMe driver, which allows installation of application-defined functions that can intercept NVMe I/O completions directly in the kernel’s NVMe interrupt handler. This allows these application-defined functions to immediately process the data returned after each block I/O access. These functions can process, aggregate, or filter the results before they are returned to the application. These functions can also efficiently trigger additional block I/Os by recycling the NVMe I/O descriptor to read additional blocks based on the data they observed in earlier I/O completions. For example, application-defined functions can traverse a B-tree stored in a file and return a requested key-value pair with a single system call, despite the multiple, data-dependent I/Os needed for the traversal. Using these application-defined functions, XRP showed that data-dependent operations like index traversal in two simple key-value stores (BPF-KV and WiredTiger) led to significant speedups even in a single host. However, these speedups were demonstrated with simple applications: locally-run key-value stores where all I/O operations involve one file with one cache and static fixed-sized data structures on disk.

Why eBPF? BPF-oF also uses eBPF to enable applications to run custom storage pushdown functions. By running the application-defined storage functions as eBPF functions, one can ensure that they are isolated from one another and from sensitive kernel state. In order to safely allow user functions to execute in the kernel, eBPF functions are statically verified by the kernel to ensure several safety properties. For example, functions cannot contain an unbounded number of instructions, and they cannot access memory addresses outside the

sandbox without the help of special, trusted helper functions that the kernel makes available to eBPF functions. eBPF functions are efficient partly because they can eliminate the need to switch to userspace and back to run user-defined code, and because eBPF functions are JIT-compiled to native machine code by the kernel when they are registered with a specific kernel hook site (such as the aforementioned NVMe hook).

3 Challenges

In this section, we describe the main research challenges in allowing storage functions to be triggered over the NVMe-oF protocol on a modern storage system.

Challenge 1: Safely access files in the target. As described in §2.1, in the NVMe-oF protocol, once the request reaches the target, it only knows how to access blocks, and does not have any file system information, which resides at the host. Therefore, an eBPF function that submits storage I/O for a file on the target must determine the location of the next block without having access to the file system metadata mapping files to blocks. In addition, our system must support resubmission across different file descriptors (e.g., when traversing an LSM-tree) to support modern storage applications. Thus, the target must be able to translate a file descriptor to the corresponding inode and then find the block offset corresponding to the requested file offset. It must do so in a safe manner, without the risk of reading and returning invalid data.

A simpler version of this problem also existed in XRP’s setting, where functions are resubmitted at the local NVMe driver, which sits below the file system. XRP solved this problem by synchronously maintaining a “metadata digest”, which is an RCU-protected copy of the file system’s files-to-blocks mapping [54]. The problem becomes more complex in the networked setting, because the “metadata digest” has to be maintained over the network, resubmissions can target different files, and BPF-oF must guarantee that resubmissions read and return valid data in the face of concurrent requests.

Challenge 2: Integrate in-memory data structures. Integrating a production storage system with BPF-oF presents a number of interesting challenges, primarily owing to the interactions between the system’s in-memory data structures, which sit locally at the host, with the remote function execution, which now occurs remotely at the target. First, when executing a storage operation, storage systems typically do not produce a “pure” chain of I/O accesses. Instead, some of the intermediate data structures in the operation path may be cached in memory. For example, in LSM-tree storage systems [6, 11, 19, 42, 43] each file has its own index, and these indices are often cached in memory, while the actual data mostly resides on disk. Therefore, a naive pushdown operation that traverses multiple files will either ignore these cached indices, issuing many more storage I/Os than necessary, or return back to the host for each file, negating the benefit of storage pushdown. In order to achieve optimal speedup with BPF-oF, we need to ideally chain the I/O storage accesses

together in one network traversal.

Second, assuming we offload such a chain of I/O accesses, there is still the question of how to keep the host’s in-memory cache up-to-date. While the final result of the query will always be returned, the “intermediate” I/O accesses (e.g., accessing a file in the middle of the LSM-tree lookup) would not. A strawman approach would be to return the results of *all* the I/O accesses back to the host and then appropriately update the application’s in-memory data structures. However, this would negate much of the benefit of storage pushdown, as these intermediate results are not actually needed by the application (only the final query result is needed), and would lead to approximately the same network bandwidth consumption as running the requests with no storage pushdown.

4 BPF-oF Design

In this section we present BPF-oF’s design principles (§4.1), BPF-oF’s general architecture (§4.2), and discuss how it addresses the first challenge from §3. Finally, we discuss the limitations of our current design in §4.3.

4.1 Principles

We design BPF-oF with the following principles in mind.

1. **Wide applicability.** Require minimal modifications to the NVMe-oF protocol, Linux kernel, and application that uses BPF-oF. In addition, the system should be immediately deployable on public clouds, and it should not require any specialized hardware support.
2. **CPU efficiency.** Increase throughput per core and reduce latency.
3. **Concurrent access.** BPF-oF must not return invalid data in the face of concurrent requests that may modify file system metadata.
4. **Flexible resubmissions.** Support variable-sized resubmissions across different files, to enable data structure traversal in modern storage systems.
5. **Target SSD-optimized applications.** We target SSD-optimized storage applications [6, 13, 15, 35, 39, 43, 44], wherein incoming updates are written to in-memory buffers and are periodically flushed to disk as large immutable files. BPF-oF primarily focuses on supporting read-only eBPF functions, since writes are background operations. It is also beneficial that such applications do not issue in-place updates to disk, otherwise updates would interfere with pushdown operations. In addition, BPF-oF initially only targets applications that implement their own userspace cache (i.e. no support for the OS page cache), which is the case for many SSD-optimized systems.

4.2 Architecture

BPF-oF supports storage pushdown with eBPF (Figure 3) over NVMe-oF by adding three mechanisms. First, it adds

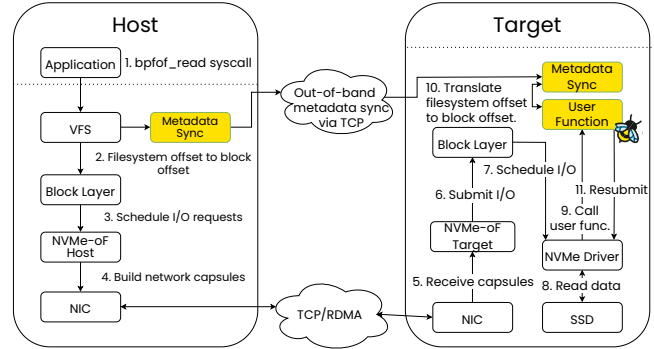


Figure 3: BPF-oF architecture.

a mechanism to asynchronously update the target’s file-to-block mapping. Second, it adds a versioning scheme that allows eBPF functions to safely access storage blocks at the target concurrently with file metadata updates. These two mechanisms jointly enable the third mechanism: safely running user-defined functions at the target server’s NVMe driver. Applications interact with BPF-oF using two interfaces: a new NVMe command that sends a “pushdown” NVMe-oF request and a userspace API that allows applications to install remote eBPF functions at the target. Before diving into these mechanisms and interfaces, we walk through an example of a BPF-oF request that traverses a B-Tree index by pushing down the traversal to the target. We match the steps of the example to the steps shown in Figure 3.

End-to-end example of a BPF-oF request. At initialization, the application installs the remote eBPF index traversal function using BPF-oF’s API. In our current implementation, the application on the host machine transfers the eBPF program over SCP to the target machine, where it is compiled, verified and loaded. For each query, the application first stores the query key in a buffer, which is called the *scratch buffer* [54]. The scratch buffer is allocated by the application and stores the query key and the query result from the remote eBPF function. In addition, the application specifies a list of file descriptors that it might access remotely. The first file descriptor in this list will be read first. The application also specifies the number of bytes and offset to read from the first file descriptor, and the eBPF function ID to call on the target.

Then, the application calls our `read_bpf_of` system call (step 1). The system call obtains a reference for each passed file descriptor, ensuring that the files’ data blocks are not overwritten if the file is deleted while there are outstanding requests. In addition, it creates a file-descriptor-to-inode mapping, which lets us avoid synchronizing the application’s entire file-descriptor table by using inode versions. As described in §4.2.1, an inode’s version changes whenever the file-to-block mapping (e.g., extent tree) for that inode is updated. Then, the application’s file system request is translated to a block request (step 2) and reaches the NVMe-oF host driver (steps 3, 4). There, the host constructs a new `read_bpf_of` NVMe command. The command transfers the scratch

buffer and its size to the target, along with the file-descriptor-to-inode mapping, the offset to read from, the number of bytes to read, and the eBPF function ID to call. Before submitting the request, the host driver checks that the inode versions have been synchronized, as described in Algorithm 1.

Next, the NVMe-oF target driver receives the request (step 5), and prepares a block request and submits it (step 6). The block request is translated to an NVMe read request and submitted to the target’s NVMe device (steps 7, 8). The target’s NVMe driver receives the request result and calls the user’s eBPF function (step 9), passing in the result and the scratch buffer. The eBPF function processes the request result and decides what happens next. In this example, the returned data represents an internal node of the B-tree; hence, the eBPF function parses the node to find the file offset of the child node that it should read to continue the point query. The eBPF function indicates to BPF-oF’s *resubmission mechanism* that it should read this additional block, so BPF-oF resubmits the read request to the target’s NVMe driver with the new offset (steps 10, 11). The resubmitted read’s file descriptor is translated with the previously sent file-descriptor-to-inode mapping, and the offset is translated using the inode’s file-to-block mapping, which the host synchronized with the target.

When the eBPF function finds the correct key, it writes the value to the scratch buffer and indicates that the result should be sent to the host. The scratch buffer is then returned to the host, where BPF-oF will again verify that the inode versions remain in sync, as shown in Algorithm 2. Finally, BPF-oF returns the result to the application.

In this way, BPF-oF implements index traversal at the target without requiring back-and-forth communication with the host. This workflow contains a versioning scheme, which allows BPF-oF to function correctly even in the face of concurrent reads and writes, which can lead to stale remote metadata. We now provide more details on these mechanisms and interfaces. For simplicity, in the description below, we assume an integration with the ext4 file system.

4.2.1 Mechanisms

Versioning. BPF-oF guarantees safety by implementing versioning on the file system metadata. To do so, we make a minor change to the file system in order to add a version to each file’s extent tree and increment it when the tree changes. With that functionality in place, we design our versioning algorithm, described by Algorithms 1 and 2.

The versioning algorithm runs only on the host, before request submission and after completion. Before submission, the host driver gets the latest inode ID and version for the given requests’ file descriptors. For all of them, it checks if that version of the inode’s extent tree has been synchronized with the target. If not, it aborts the request. This ensures that the target always has the latest version of the file system metadata. This check is actually enough to guarantee that the file-to-block mapping will be correct on the target, except for

Algorithm 1: Host - Check Before Submission

```

1: for  $idx = 0$  to  $num\_fds - 1$  do
2:    $fd \leftarrow bpf\_request.fds[idx]$ 
3:    $curr\_version \leftarrow get\_fd\_version(fd)$ 
4:    $sent\_version \leftarrow$  latest version sent to target
5:   if  $curr\_version \neq sent\_version$  then
6:     abort request
7:   end if
8:    $request.f\_versions[idx] \leftarrow curr\_version$ 
9: end for

```

Algorithm 2: Host - Check After Completion

```

1: for  $idx = 0$  to  $num\_fds - 1$  do
2:    $fd \leftarrow bpf\_response.fds[idx]$ 
3:    $curr\_version \leftarrow get\_fd\_version(fd)$ 
4:    $used\_version \leftarrow$ 
      $get\_sent\_request\_fd\_version(bpf\_response)$ 
5:   if  $curr\_version \neq used\_version$  then
6:     abort request
7:   end if
8: end for

```

one case: when the extent tree is remapped, i.e. an existing file offset points to a new disk offset. This can happen during truncation and defragmentation. We handle this case by adding a final version check when the response is received from the target: if the version has changed, the host aborts the request. It also wipes the scratch buffer to ensure that the target will not return stale data, even if the extent tree is remapped during the request.

Metadata synchronizer. The next challenge is synchronizing the file system metadata (i.e. file-to-block mappings) from the host to the target. The metadata synchronizer consists of two kernel modules running on the host and the target. The host module watches for file system changes by installing *fs-notify* kernel hooks. For each change, it schedules a metadata synchronization for the specified inode. A kernel workqueue processes these requests and checks if the file-to-block mapping has changed from the last version it synchronized, using the versioning scheme described above. Then, the synchronizer copies the inode’s extent tree and transfers it to the target module, where it is used for resubmissions. Metadata synchronization is done asynchronously over a TCP connection and does not block the data path. For simplicity, we implemented the synchronizer as an out-of-band TCP client and server, rather than creating a new NVMe command for it.

eBPF execution. We design the eBPF execution and resubmission mechanism at the target on top of XRP [54]. We add support for (a) resubmitting eBPF functions across multiple files instead of a single file and (b) specifying a different request size for each resubmission. We provide more details on these changes in §5.

```
int bpf_of_read(int fd,
               size_t count,
               off_t offset,
               void *scratch_buf,
               size_t scratch_buf_count,
               unsigned int bpf_id)
```

Figure 4: BPF-oF read system call.

4.2.2 Interfaces

To integrate with NVMe-oF, BPF-oF implements a new NVMe command following the specifications for NVMe and NVMe-oF [7,9]. The command transfers the scratch buffer, which contains the query key and result (on completion), from the host to the target and back. The command’s arguments are the same as a read request. We transfer the additional arguments needed by BPF-oF, as described above, with NVMe-oF’s data transfer mechanism (scatter-gather lists).

BPF-oF implements the new `read_bpf_of` system call (Figure 4) with the following arguments: a list of file descriptors it might read from, the number of bytes and offset to read from the first file descriptor, the scratch buffer and its size, and the eBPF function ID to call on the target.

4.3 Limitations

We view the current system as the first step in enabling push-down in disaggregated storage settings. BPF-oF’s current design has several limitations, which we discuss below, and plan to address in future work.

Files known in advance. BPF-oF assumes that all files accessed in an I/O chain are known in advance. This assumption holds true in many storage systems (e.g., LSM-tree storage systems), but may not hold in others.

Block device support. Currently, BPF-oF runs eBPF functions within the target’s NVMe driver. From a performance standpoint, this is optimal, as it allows BPF-oF to almost entirely bypass the target kernel’s storage stack, thereby reducing CPU consumption at the target. However, this design choice has several disadvantages: most notably, it doesn’t support complex block device configurations (e.g., RAID, LVM) but only single partition drives. However, BPF-oF could run the eBPF functions at a higher layer of the kernel (i.e. in the target’s block layer) in exchange for a small performance hit.

File system support. We currently only support the ext4 file system. We anticipate the integration for additional file systems will be very similar to ext4.

Limitations of eBPF. While eBPF has several attractive properties, including its usage in popular systems [3,4], its increasing adoption in standards such as NVMe for smart storage [8], and its security guarantees, it also has major known downsides. These include the difficulty of programming in eBPF and documented security vulnerabilities [23,34,48]. We choose to implement remote-storage pushdown using eBPF primarily due to its widespread adoption, but our design

principles would hold for other privileged execution sandbox frameworks, such as WASM or Privbox [34].

5 Supporting In-Memory Data Structures

A significant challenge in implementing remote-storage push-down is utilizing a storage system’s in-memory data structures, which are stored at the client, with the on-disk lookup, which occurs remotely. Our solutions to this challenge apply generally to any storage system that needs to use storage pushdown, but in this section, we focus on how we tackled this problem in the integration with RocksDB, as it utilizes a wide variety of continuously-updated in-memory caches (e.g., index and data block caches, Bloom filters) common to many systems [6,19,35,43].

5.1 RocksDB Architecture

LSM-tree. RocksDB is an LSM-tree based storage engine [42]. To avoid SSD write amplification and wear out, incoming writes to RocksDB are buffered in memory, in a data structure called the MemTable. After they fill up, MemTables are asynchronously flushed to disk as sorted-string table (SST) files, which are composed of sorted key-value pairs. SST files are divided into data and metadata blocks. Metadata blocks include index blocks, which map keys to data blocks, and filter blocks, which contain Bloom filters (these are disabled by default in RocksDB, but are supported by BPF-oF).

These SST files are organized into hierarchical levels (L_0, L_1, \dots, L_N), where the “upper” levels store the latest versions of each key-value pair (e.g., L_0 is “higher” than L_1). The immutable MemTables are flushed into L_0 , which stores files with overlapping key ranges, whereas files in the lower levels (L_1, \dots, L_N) hold non-overlapping ranges. RocksDB uses a background thread that periodically scans SST files from adjacent levels (e.g., L_2 and L_3), and combines them into a single file, which is written to the lower level (e.g., L_3). In the process, RocksDB removes deleted and stale versions of keys, freeing up space. It has been shown that to obtain good performance in LSM-trees [42], each level contains a multiple of (e.g., $10\times$) the capacity of the previous level, whereby the lowest level (L_N) makes up most of the LSM-tree’s capacity.

Workflow of a read. Since more updated versions of key-value pairs sit at the “upper” levels of the tree, when a read request arrives, RocksDB searches for the object in the levels of the tree from top to bottom. It first checks if the key-value pair is in RocksDB’s MemTables. If not, it checks in L_0 ’s files, followed by L_1 , etc. For each file, RocksDB first reads the filter blocks to check if the key might be in that file (with a high probability). If so, it checks the index block to find the relevant data block and searches for the key in the data block. These blocks may or may not be cached in memory by RocksDB: typically many (or all) of the index and filter blocks are cached. Therefore, as it is locating the key within a file, BPF-oF may consult an in-memory data structure (e.g., a

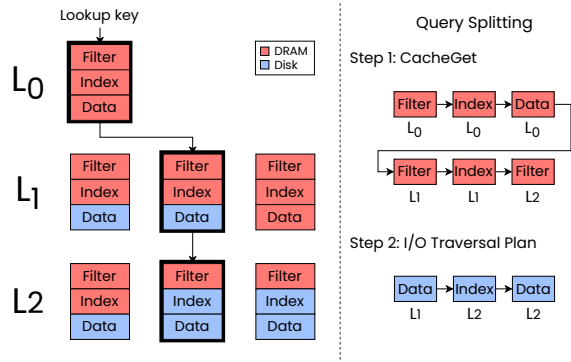


Figure 5: RocksDB LSM-tree traversal and query splitting.

cached index block), followed by disk I/O (e.g., a data block). This can create an interleaved pattern of in-memory and disk lookups, which is challenging for BPF-oF, since it only provides a significant speedup for a *chain* of storage I/Os.

To facilitate lookups, LSM-trees cache the key ranges each file contains. We make a key observation, which we rely on in our design, that this cache can be used to determine all the files whose range contains a particular key, and therefore *might* be needed to be accessed to satisfy a read operation.

5.2 Integration with BPF-oF

In integrating with RocksDB, we only focus on accelerating the read path, since writes are already heavily optimized and are not in the critical path (all writes are buffered in memory and written sequentially to disk). Initially, our prototype only accelerates single-key point lookups. In integrating RocksDB with BPF-oF we make the following two key contributions: (a) a new technique, *query splitting*, that splits the in-memory accesses from the storage accesses to maximize the effectiveness of storage pushdown, and (b) periodic cache sampling to keep the storage engine’s cache fresh. We also describe some of the challenges of converting components of a C++ commercial storage engine to eBPF. Overall, our RocksDB integration with BPF-oF required ~ 700 LoC in changes to RocksDB itself and $\sim 1,200$ LoC to reimplement RocksDB’s parser function in eBPF.

5.2.1 Query Splitting

In order to maximize the efficiency of storage lookups, ideally the remotely-executed function would involve a chain of I/O requests that exclusively access the disk, rather than in-memory data structures. To this end, we split the query into two phases: first, the query accesses all the in-memory data structures that *may* be involved in a key lookup, and it then triggers all the required I/O as a series of dependent I/O functions at the remote node. Recall that due to the design of LSM-trees, the host can determine all the files that may be needed to complete the operation. We are guided by the observation that accessing RocksDB’s in-memory data structures is orders of magnitude less costly than storage I/O (especially when that I/O is over the network). Therefore, there is little cost in potentially accessing more in-memory data structures

than would have been necessary otherwise.

To this end, we implement a new function in RocksDB, CACHEGET. CACHEGET is based on RocksDB’s existing per-SST file GET function, but it only reads data from that file’s cached blocks without triggering any storage I/O. Our prototype runs CACHEGET on all files that may contain the key, across all levels. Based on the results of this function (e.g., whether a key is found in a particular data block), our prototype creates an *I/O traversal plan*, which only fetches uncached blocks that may be required in the I/O traversal. The traversal plan skips files that have both index and data blocks cached, those that are filtered by a per-file Bloom filter, and it starts I/O at the data block if the index block is cached.

We illustrate this process in Figure 5. In this example, three files across L_0 , L_1 , and L_2 have key ranges that contain the requested key. All of these files may need to be read to return the key. As the figure shows, these files’ blocks are partially or fully cached. CACHEGET retrieves the blocks in the potential path of the key that are cached in DRAM. The remaining uncached blocks form the on-disk I/O traversal plan. If the key is not found in the cached blocks, our RocksDB prototype issues a `read_bpf_of` system call, with the files from the I/O traversal plan. Note that `read_bpf_of` does not blindly follow the I/O traversal plan to completion, and it terminates once the key is found. For example, if a key exists in L_1 , the parser will not read data from L_2 , since the value in L_1 will be more up-to-date than other versions lower in the tree. As such, CACHEGET will often be called on more files than necessary, since RocksDB does not know the *minimal* set of files that will be involved in the traversal in advance.

5.2.2 Cache Sampling

In RocksDB, all data blocks that were accessed as part of a query are cached by default. However, BPF-oF returns only the final query result, not the intermediate blocks. With BPF-oF, in some cases RocksDB’s cache might be “missing” intermediate data blocks. This problem would exist in any system with storage pushdown that caches intermediate results. A naive solution would be to return all of the data blocks accessed in the query, but this would negate much of the benefit of storage pushdown, by increasing network bandwidth (and CPU costs) significantly. As such, we seek another way to keep the cache fresh. We are inspired by prior work on cache sampling [47], where a low percentage of the randomly-sampled requests update the cache with recently-accessed intermediate data blocks (as well as with the end result). We implement this idea by sending a small percentage of requests through the slower, non-offloaded “normal” read path, which caches all the intermediate results. Empirically, we find this technique is quite effective with sampling rates as low as 0.1-1% (see §6.4).

5.2.3 Converting RocksDB’s I/O Lookups to eBPF

The final practical obstacle in integrating RocksDB was converting RocksDB’s I/O lookups into eBPF functions. We converted RocksDB’s SST file parsers, which parse SST index and data blocks, to eBPF. To the best of our knowledge, our parser is the first example of such complex logic implemented in eBPF. Our eBPF parser first uses the generated I/O traversal plan, which is stored as an array in the scratch buffer, to parse each file at the correct file offset and parsing stage (i.e. index or data block). If the requested key is not found in the parsed data, the eBPF parser will resubmit a new I/O request for the next file in the plan. This process continues until either the key is found or all relevant files have been traversed. If the eBPF parser fails, we fall back to the regular RocksDB read path. This occurs infrequently when index blocks are pinned, as such a request will only fail when a RocksDB data block (by default $\sim 4\text{KB}$) is split between ext4 extents, which tend to be significantly larger than 4KB for (immutable) SST files.

Converting RocksDB’s parser to eBPF frequently ran into the limitations of the eBPF programming framework, due to the need to rewrite RocksDB’s C++ I/O functions into C. It also required adding various memory access and loop iteration checks to pass verification. Due to eBPF’s instruction-count and jump limits, many logical RocksDB functions have to be split up into smaller functions so they can be verified successfully. We did so using *function-by-function verification* [2] by splitting the logical functions into “isolated” components that are independently verified. Our eBPF parser demonstrates the viability of implementing complex storage functionality in eBPF, despite its restrictions.

6 Evaluation

This section answers the following questions. Questions 1–4 are answered with RocksDB, and Q5 targets the other systems (WiredTiger and BPF-KV):

- Q1:** How does BPF-oF perform under different network protocols, storage devices, and cache sizes compared to the baseline NVMe-oF and local environments?
- Q2:** What is BPF-oF’s impact on CPU and network utilization and energy consumption?
- Q3:** How do different sampling rates affect performance?
- Q4:** What is the performance impact of version mismatches?
- Q5:** How does BPF-oF generalize to other storage systems?

6.1 Experimental Setup

We conduct our experiments in three different CloudLab storage configurations [22], using NAND and Optane SSDs. We measure an average roundtrip latency of $30\mu\text{s}$ and $18\mu\text{s}$ for TCP and RDMA, respectively.

NAND configuration. This is a typical enterprise NAND SSD setup. Both host and target are c6525-100g machines in CloudLab. Each machine has a 24-core AMD EPYC 7402P CPU, 128 GB of RAM, two 1.6 TB NAND SSDs (Dell Ent

NVMe AGN MU U.2), and a Dual-port Mellanox ConnectX-5 NIC with a 100 Gbps network. The NAND SSD’s read latency is $90\mu\text{s}$ and can sustain 700K read IOPS.

Optane configuration. This setup uses a low-latency storage device. The host is a c6525-100g machine, while the target is d750 with Intel Optane SSD. The target has two 16-core Intel Xeon Gold 6326 CPUs, 128 GB of RAM, a 400 GB Optane 5800X SSD, and a Quad-port BCM57504 NetXtreme-E on a 25 Gbps network. The Optane SSD has a $3\mu\text{s}$ 512B read latency and can sustain 5M read IOPS.

RDMA offloading configuration. This setup uses RDMA NIC offloading. Both host and target are r6525, with 32-core AMD EPYC 7543 CPUs, 256 GB of RAM, one 1.6 TB NAND SSD (Dell Ent NVMe AGN MU U.2), and a Dual-port Mellanox ConnectX-6 100 GB NIC on a 100 Gbps link.

Software configuration. To measure the throughput per core, we increase the number of threads per core up until saturation, and use CPU pinning to bind each thread to a specific core, and flow-steering to associate each flow with a specific host and target core. To measure energy, we use AMD’s μProf tool. To make our results reproducible, we disable hyperthreading and address space randomization. We run Ubuntu 20.04 and Linux 5.12.0.

Evaluated systems. We implement BPF-oF on three systems: RocksDB, WiredTiger and BPF-KV. RocksDB is a complex and popular key-value store, and we test it with BPF-oF to show whether a pushdown approach would benefit state-of-the-art storage applications. WiredTiger is a very simple production-grade key-value store used primarily by MongoDB. BPF-KV is a bare-bones academic key-value store custom-designed for running eBPF functions [54]. All systems bypass the OS page cache with direct I/O.

RocksDB setup. We build upon RocksDB 7.7.3 with a 100 GB database in 5 layers ($L_0 - L_4$) and 3 cores (which saturate our enterprise SSD); we use 24 and 20 threads-per-core for NAND and Optane SSD, respectively. Unless noted, all index blocks are cached and Bloom filters are disabled (RocksDB’s default). We run RocksDB with both TCP and RDMA; our results focus on TCP, as it is more prevalent. By default, we use a cache sampling rate of 1% (i.e. a random 1% of reads go through the regular non-BPF-oF read path).

Workloads. To test BPF-oF, we run the YCSB [20] benchmark suite under Zipfian (0.99) key distributions, along with uniform read and read-write workloads. We do not run YCSB E, because our RocksDB integration does not yet accelerate scans. We also test BPF-oF with a production trace taken from clusters 19 and 37 of the Twitter cache workloads [49]. These traces are read-dominant (75% and 63%, respectively) and have low skew ($\alpha = 0.7$ and 0.4 , respectively), ensuring that their working set does not fit entirely in memory.

6.2 BPF-oF With Different Setups (Q1)

TCP with NAND SSD. Figure 6 shows RocksDB’s throughput with BPF-oF with NAND and TCP, compared




Figure 6: RocksDB (a) uniform read throughput-latency and throughput (b) without and (c) with data blocks cached of BPF-oF vs. NVMe/TCP on NAND SSD.

to NVMe/TCP and to local RocksDB (not over the network).¹ Without a data cache, BPF-oF can achieve 1.4-2.2 \times higher throughput than NVMe/TCP, and in many cases (including the Twitter cache traces, which are real-world workloads) matches or exceeds the throughput of running on local disk. This is because BPF-oF eliminates repeated traversals of the kernel I/O stack on the target, which in some cases outweighs the added network processing cost. This trend continues when adding a data cache, where BPF-oF provides 1.2-1.7 \times higher throughput than NVMe/TCP.

The relative speedup of BPF-oF with a data cache is generally lower, because some lookups may be cached, thereby reducing the average number of roundtrips saved by BPF-oF. The only cases where BPF-oF falls slightly short are YCSB A, D, and F, with data blocks cached. Since these workloads are highly-cacheable, they issue very few I/Os per request, and BPF-oF’s slightly higher overhead for in-memory requests (e.g., its preemptive access of all possible in-memory caches in each query) leads to a slight slowdown. However, with just index blocks cached, BPF-oF provides a speedup in all workloads except YCSB D, and it even beats the baseline with cache for a majority of the workloads. Figure 6 also shows that BPF-oF reduces tail latency by up to 2 \times .

RDMA with NAND SSD. Figures 7a and 7b show the throughput of RocksDB with BPF-oF with NAND and RDMA, compared to *NIC offloaded* NVMe/RDMA and to local RocksDB. Even though BPF-oF does not take advantage of offloading the RDMA requests to the NIC, BPF-oF is up to 1.8 \times faster than offloaded NVMe/RDMA. Since RDMA is very CPU-efficient it achieves an up to 1.6 \times speedup even compared to the local disk baseline. In our experiments, there was no noticeable difference in the baseline’s throughput between offloaded and non-offloaded RDMA. The reason for this is that while offloading the RDMA protocol to the NIC saves CPU at the target, the target’s CPU is usually not a bottleneck with RocksDB.

TCP with Optane SSD. We also test BPF-oF with Optane, providing a glimpse of BPF-oF’s benefits as SSDs become faster. BPF-oF provides even greater speedups with Optane

(Figures 7c and 7d), providing 2-2.8 \times higher throughput. This is because Optane can support a high I/O bandwidth, while in the NAND SSD experiments I/O bandwidth may become a bottleneck. BPF-oF also provides a consistent improvement (up to 2.6 \times) in tail latency in all workloads. In addition, we see a surprising result: BPF-oF without a cache yields higher throughput than BPF-oF with a cache. We discuss this behavior below.


Caching and Storage Pushdown. Caching has an interesting and unexpected interaction with BPF-oF. In certain cases, running RocksDB with a cache can actually slow down the system, as seen in Figure 7d. In the baseline NVMe-oF case, RocksDB transfers all the intermediate blocks over the network until it finds the final result. Thus, maintaining a cache for these blocks incurs little additional cost. However, when using storage pushdown with BPF-oF, RocksDB only transfers the final result over the network, and we employ cache sampling to keep the cache fresh with intermediate results. In some workloads, the cost of keeping the cache fresh outweighs its benefits. This balance depends on the workload type, the cache miss rate, and the I/O cost.

Non-Default Configurations. By default, we evaluated RocksDB with indices pinned and without Bloom filters, as these are the default configuration options in RocksDB. For completeness, we also tested (1) with Bloom filters enabled and (2) with indices not pinned. The results were not very different than with our default configurations: BPF-oF provides significant speedups in both cases, with up to 2 \times improvement with Bloom filters and 1.8 \times improvement with unpinned indices (for more details, see §10.1).

6.3 CPU, Network, and Energy Savings (Q2)

BPF-oF provides significant CPU, network, and energy (and carbon emissions) savings, which are directly correlated with performance increases. To measure these savings, we monitor the CPU time spent *on both the host and target*, network bytes exchanged, and energy spent while running a benchmark. We present these results (normalized per-request) for uniform reads with a 1 GB data cache, in Table 1. For TCP, BPF-oF consumes 37% fewer CPU cycles (measured in mCPUs, i.e. 1/1000th of a CPU) and generates 23% less network traffic, resulting in a roughly 70% throughput improvement. It does

¹The TCP results in Figure 2 were obtained on the offloading machine in order to compare them to RDMA with offloading. As such, the TCP results in Figure 6 are not identical to those in Figure 2.



(a) NVMe/RDMA with NAND SSD without data blocks cached. (b) NVMe/RDMA with NAND SSD with data blocks cached. (c) NVMe/TCP with Optane SSD without data blocks cached. (d) NVMe/TCP with Optane SSD with data blocks cached.

Figure 7: RocksDB throughput under different configurations.

so while consuming 36% less energy per request. In the offloaded RDMA setting, BPF-oF consumes 29% and 19% less CPU and network traffic per request, respectively, resulting in a 65% throughput improvement. Energy measurements were not available in the RDMA machines. Additionally, BPF-oF should can yield even lower carbon emissions due to the decreased reliance on energy-hungry DRAM.

Table 1: Resource consumption with uniform read and 1 GB data cache.

Setup	Throughput (req/s)	CPU (mCPU / req)	Network (kB / req)	Energy (mW / req)
NVMe-oF/TCP	48279	0.10	12.07	3.63
BPF-oF/TCP	69935	0.06	9.27	2.33
NVMe-oF/RDMA (offload)	77553	0.04	12.3	-
BPF-oF/RDMA	120991	0.03	10.02	-

6.4 Sampling Rate (Q3)

As discussed in §5.2.2, our RocksDB integration sends a small portion of reads through the normal read path in order to keep the cache warm. We evaluate the impact of this sampling rate on RocksDB’s performance. Figure 8 shows that throughput peaks at a sample rate between 0.1% and 10%, except for YCSB D, which is highly cacheable and thus benefits from a higher sampling rate. Similarly, we see a slight decrease in performance in YCSB A and F at higher sampling rates. These workloads are write-heavy and highly skewed, which means that as the popular keys are written repeatedly, they reappear in the higher levels of the LSM tree and shorten the resubmission chains, leading to a smaller performance improvement. Based on these experiments, we run BPF-oF with a 1% sampling rate by default.

6.5 Version Mismatches (Q4)

As discussed in §4.2.1, BPF-oF versions file extent metadata in order to guarantee safe concurrent access to files for the offloaded applications. A version mismatch occurs when the synchronized file system metadata in the target becomes stale. A version mismatch results in a failed request, which is then retried through the non-accelerated read-path. Thus, version mismatches are very costly: they approximately decrease throughput by 1 request and more than double the requests’ latency. Fortunately, they are very rare. Across all benchmarks, we observed only 20 mismatches, out of millions of requests per benchmark. The only benchmark that had a slightly higher number of mismatches was YCSB D, due to

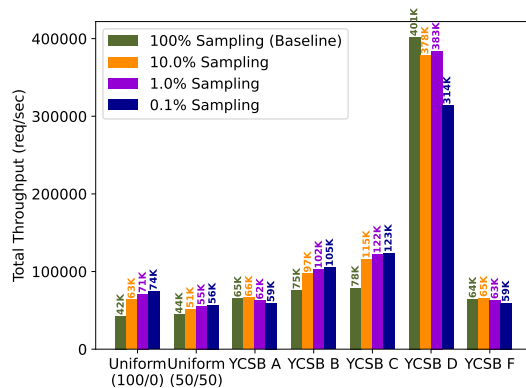


Figure 8: RocksDB throughput with different sampling rates with cache under BPF-oF using TCP and NAND SSD.

its write-heaviness, but they still occurred in only 0.03% of requests. Thus, the performance impact of version mismatches is negligible in our integration with RocksDB.

6.6 WiredTiger (Q5)

We run WiredTiger in its LSM-tree version, wherein data is split into levels, and each level is stored in a single file. Each file uses a B-tree index where the data is stored as the leaf nodes. We configure the B-tree page size to be equal to 512B. The database contains 1G items with a 16B object size. We test WiredTiger with BPF-oF using the YCSB benchmark on 4 cores, the minimal number of cores to eliminate CPU contention between foreground and background threads. Figure 9a shows that for read-write workloads such as YCSB A, BPF-oF has 20% higher throughput. For read-heavy workloads such as YCSB B, YCSB C and uniform, the improvement reaches 30% when using an Optane SSD, as seen in Figure 11 in the Appendix. The performance gain of WiredTiger is more modest than that of RocksDB (and BPF-KV), since WiredTiger issues far fewer I/Os, and therefore reducing the time (and CPU consumption) of I/O leads to a more limited speedup.

6.7 BPF-KV (Q5)

BPF-KV [54] is a toy high-performance key-value store optimized for eBPF functions that was introduced by the XRP paper. It organizes its data in a B⁺-tree, can cache the top-most levels of the tree, and uses fixed-sized 512B key-value




Figure 9: WiredTiger and BPF-KV evaluations with TCP and NAND SSD.

pairs. The leaves of the tree contain a pointer to the values, which are stored in an unsorted log. We use a 67 GB database file that has 6 index levels and 1 level for the log. We conducted a full suite of experiments with BPF-KV to evaluate BPF-oF’s performance under different network protocols, storage devices, and caching strategies. We present a subset of the results here, with further results presented in §10.3.

Point queries. We run a simple microbenchmark using a BPF-KV instance containing 8B keys with 64B values stored on the NVMe-oF target. Figure 9b compares the results of running BPF-oF vs. regular NVMe-oF on NAND SSDs on three cores. BPF-oF consistently provides about $2.6\times$ higher throughput than NVMe/TCP when the upper levels of BPF-KV are cached. Note that when no cache is used, the access pattern limits the SSD to 120K IOPS (which translates to 30K requests per second at 7 resubmissions per request), because the upper level nodes become read hot-spots in the NVMe queues. Note that for NVMe/TCP on Optane, BPF-oF provides more than $8\times$ higher throughput when BPF-KV is run without a cache. This speedup is due to BPF-oF saving significant CPU cycles in processing TCP packets, which can instead be used to handle additional I/O requests.

Range queries. BPF-oF can support more complex operations than single-key lookups. Figure 9c shows the results over TCP on NAND with range queries, where a range of keys is fetched sequentially from the leaves of BPF-KV’s B^+ -tree. BPF-oF can achieve over $5\times$ higher throughput.

6.8 Takeaways

We find that BPF-oF can provide significant performance benefits for disaggregated storage systems, of up to $2.8\times$ higher throughput and $2.6\times$ lower tail latency with RocksDB. In particular, BPF-oF benefits workloads that are read-heavy, and whose working set does not fit in the cache, as these necessitate longer resubmission chains which BPF-oF is able to accelerate. BPF-oF also provides significant CPU and network savings, which are directly correlated with performance increases. Notably, we find that caching with storage push-down is a mixed bag. When workloads are highly skewed a cache avoids issuing unnecessary network traffic, but in some workloads maintaining a cache can be detrimental due to the

high CPU and network costs of regularly updating it.

7 Related Work

Decades of prior work have explored pushing functions close to data [16, 17, 27, 32, 33, 38, 50–52, 54, 55]. Most of these pushdown mechanisms are application-specific, i.e. they allow clients of a particular system (e.g., a database) to pushdown specific functions (e.g., SQL filters).

Systems that use eBPF. Several recent systems utilize eBPF to push computation closer to storage. Kourtis et al. propose using eBPF to safely execute user-defined functions on a remote NVM storage server [32], by running functions in userspace at the remote server, without integrating with a networked storage protocol. Their paper presents preliminary work and does not provide a full system design, implementation, and evaluation. BMC [24] implements an eBPF-based cache at a remote server’s kernel to service clients’ memcached requests to minimize kernel-userspace crossings at the server. Unlike BPF-oF, BMC does not support custom user functions at the server. XRP [54] and ExtFUSE [17] allow users to bypass some of the kernel and file system’s codepath, thus reducing overhead and eliminating kernel-userspace context switches by running user-defined eBPF functions. Similarly, λ -IO [51] provides a BPF framework to execute custom functions within a computational storage device. These systems do not support a disaggregated storage use case. Finally, Electrode [56] uses eBPF extensions in the kernel to accelerate distributed protocols like Multi-Paxos. Electrode’s acceleration of coordination protocols is complementary to BPF-oF.

Other frameworks. Splinter [33], ASFP [16], and Kayak [52] are in-memory key-value stores that allow users to ship lightweight Rust functions that are executed as coroutines on a remote in-memory server. Redis [10] provides a similar capability with Lua. Unlike BPF-oF, these systems implement pushdown in the context of a specific system (an in-memory key-value store), and are not a general-purpose networked storage protocol. In addition, since they operate in an in-memory setting they do not contend with the issues that arise in storage (e.g., file metadata synchronization). Some storage systems provide application-specific pushdowns.

Amazon S3 Select [5] allows users of Amazon S3 to apply specific SQL filters to data when it is fetched from S3 to save network bandwidth. Ceph [36] provides similar capabilities with Ceph “extensions”. These systems also do not provide a general-purpose storage pushdown capability that can be plugged into any storage system.

8 Conclusion

This work explores how to design a networked storage protocol that allows the client to push user-defined storage functions to the server over the network, thereby significantly improving both client and server CPU efficiency. We believe there are significant future research challenges in this area, such as adding support for commonly-needed storage capabilities, such as disk arrays, RAID, and encrypted storage. Another promising direction is to allow the functions to fully execute on a smartNIC. Furthermore, our insights into the interaction of in-memory and fast remote disks, indicate the potential for designing efficient nearly memory-less disaggregated storage systems.

9 Acknowledgments

We would like to thank Tanvir Khan and Haoyu Li for their feedback. We also thank the CloudLab team for their help in supporting our experiments. This work was supported by gifts from IBM and Intel, and by NSF award CNS-2143868. Tal Zussman was supported by NSF award DGE-2036197. Ioannis Zarkadas is an Onassis Foundation scholar.

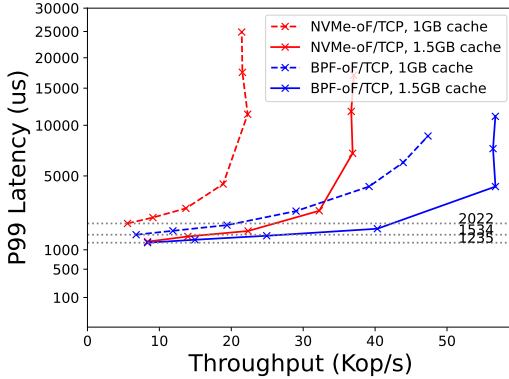
References

- [1] 100g kernel and user space NVMe/TCP using Chelsio offload. <https://www.chelsio.com/wp-content/uploads/resources/t6-100g-nvmetcp-offload.pdf>.
- [2] BPF: introduce function-by-function verification. <https://lore.kernel.org/bpf/2020109063745.3154913-4-ast@kernel.org/>.
- [3] Cilium. <https://github.com/cilium/cilium>.
- [4] Cloudflare architecture and how BPF eats the world. <https://blog.cloudflare.com/cloudflare-architecture-and-how-bpf-eats-the-world/>.
- [5] Filtering and retrieving data using Amazon S3 Select. <https://docs.aws.amazon.com/AmazonS3/latest/userguide/selecting-content-from-objects.html>.
- [6] LevelDB. <https://github.com/google/leveldb>.
- [7] NVMe base specification. https://nvmexpress.org/wp-content/uploads/NVM-Express-1_4b-2020.09.21-Ratified.pdf.
- [8] NVMe computational storage - an update on the standard. <https://www.snia.org/educational-library/nvme-computational-storage-update-standard-2022>.
- [9] NVMe over fabrics specification. <https://nvmexpress.org/wp-content/uploads/NVMe-over-Fabrics-1.1a-2021.07.12-Ratified.pdf>.
- [10] Redis functions. <https://redis.io/docs/manual/programmability/functions-intro/>.
- [11] RocksDB. <https://rocksdb.org/>.
- [12] RocksDB users. <https://github.com/facebook/rocksdb/blob/main/USERS.md>.
- [13] Rockset: real-time analytics at cloud scale. <https://rockset.com/>.
- [14] SQLite pluggable storage engine. <https://sqlite.org/src4/doc/trunk/www/storage.wiki>.
- [15] WiredTiger storage engine. <https://docs.mongodb.com/manual/core/wiredtiger/>.
- [16] Ankit Bhardwaj, Chinmay Kulkarni, and Ryan Stutsman. Adaptive placement for in-memory storage functions. In *2020 USENIX Annual Technical Conference (USENIX ATC 20)*, pages 127–141. USENIX Association, July 2020.
- [17] Ashish Bijlani and Umakishore Ramachandran. Extension framework for file systems in user space. In *2019 USENIX Annual Technical Conference (USENIX ATC 19)*, pages 121–134, Renton, WA, July 2019. USENIX Association.
- [18] Laurent Bindschaedler, Ashvin Goel, and Willy Zwaenepoel. Hailstorm: Disaggregated compute and storage for distributed LSM-based databases. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS ’20*, page 301–316, New York, NY, USA, 2020. Association for Computing Machinery.
- [19] Alexander Conway, Abhishek Gupta, Vijay Chidambaram, Martin Farach-Colton, Richard Spillane, Amy Tai, and Rob Johnson. SplinterDB: Closing the bandwidth gap for NVMe key-value stores. In *2020 USENIX Annual Technical Conference (USENIX ATC 20)*, pages 49–63, 2020.
- [20] Brian F Cooper, Adam Silberstein, Erwin Tam, Raghu Ramakrishnan, and Russell Sears. Benchmarking cloud serving systems with YCSB. In *Proceedings of the 1st ACM symposium on Cloud computing*, pages 143–154, 2010.

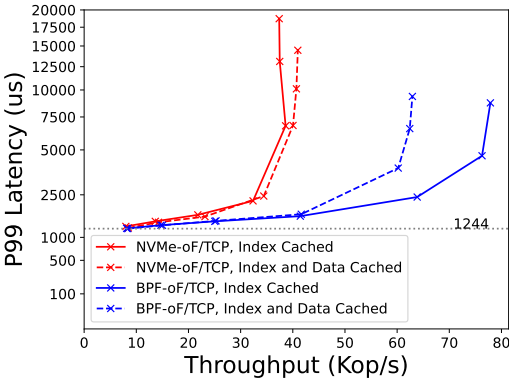
- [21] Siying Dong, Andrew Kryczka, Yanqin Jin, and Michael Stumm. Evolution of development priorities in key-value stores serving large-scale applications: The RocksDB experience. In *19th USENIX Conference on File and Storage Technologies (FAST 21)*, pages 33–49. USENIX Association, February 2021.
- [22] Dmitry Duplyakin, Robert Ricci, Aleksander Maricq, Gary Wong, Jonathon Duerig, Eric Eide, Leigh Stoller, Mike Hibler, David Johnson, Kirk Webb, Aditya Akella, Kuangching Wang, Glenn Ricart, Larry Landweber, Chip Elliott, Michael Zink, Emmanuel Cecchet, Snigdhaswin Kar, and Prabodh Mishra. The design and operation of CloudLab. In *2019 USENIX Annual Technical Conference (USENIX ATC 19)*, pages 1–14, Renton, WA, July 2019. USENIX Association.
- [23] Elazar Gershuni, Nadav Amit, Arie Gurfinkel, Nina Narodytska, Jorge A. Navas, Noam Rinetzky, Leonid Ryzhyk, and Mooly Sagiv. Simple and precise static analysis of untrusted Linux kernel extensions. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2019*, page 1069–1084, New York, NY, USA, 2019. Association for Computing Machinery.
- [24] Yoann Ghigoff, Julien Sopena, Kahina Lazri, Antoine Blin, and Gilles Muller. BMC: Accelerating memcached using safe in-kernel caching and pre-stack processing. In *18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21)*, pages 487–501. USENIX Association, April 2021.
- [25] Zhihan Guo, Xinyu Zeng, Kan Wu, Wuh-Chwen Hwang, Ziwei Ren, Xiangyao Yu, Mahesh Balakrishnan, and Philip A. Bernstein. Cornus: Atomic commit for a cloud DBMS with storage disaggregation. *Proc. VLDB Endow.*, 16(2):379–392, oct 2022.
- [26] Zvika Guz, Harry (Huan) Li, Anahita Shayesteh, and Vijay Balakrishnan. Performance characterization of NVMe-over-Fabrics storage disaggregation. *ACM Trans. Storage*, 14(4), dec 2018.
- [27] Joseph M Hellerstein and Michael Stonebraker. Predicate migration: Optimizing queries with expensive predicates. In *Proceedings of the 1993 ACM SIGMOD international conference on Management of data*, pages 267–276, 1993.
- [28] Jaehyun Hwang, Qizhe Cai, Ao Tang, and Rachit Agarwal. TCP \approx RDMA: CPU-efficient remote storage access with i10. In *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*, pages 127–140, Santa Clara, CA, February 2020. USENIX Association.
- [29] J Kim. iSCSI - is it the future of cloud storage or doomed by NVMe-oF. <https://www.snia.org/sites/default/files/news/iSCSI-Future-Cloud-Storage-Doomed-NVMe-oF.pdf>.
- [30] Ana Klimovic, Christos Kozyrakis, Eno Thereska, Binu John, and Sanjeev Kumar. Flash storage disaggregation. In *Proceedings of the Eleventh European Conference on Computer Systems, EuroSys '16*, New York, NY, USA, 2016. Association for Computing Machinery.
- [31] Ana Klimovic, Heiner Litz, and Christos Kozyrakis. ReFlex: Remote flash \approx local flash. In *Proceedings of the Twenty-Second International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS '17*, page 345–359, New York, NY, USA, 2017. Association for Computing Machinery.
- [32] Kornilios Kourtis, Animesh Trivedi, and Nikolas Ioannou. Safe and efficient remote application code execution on disaggregated NVM storage with eBPF. *arXiv preprint arXiv:2002.11528*, 2020.
- [33] Chinmay Kulkarni, Sara Moore, Mazhar Naqvi, Tian Zhang, Robert Ricci, and Ryan Stutsman. Splinter: Bare-Metal extensions for Multi-Tenant Low-Latency storage. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*, pages 627–643, Carlsbad, CA, October 2018. USENIX Association.
- [34] Dmitry Kuznetsov and Adam Morrison. Privbox: Faster system calls through sandboxed privileged execution. In *2022 USENIX Annual Technical Conference (USENIX ATC 22)*, Carlsbad, CA, July 2022. USENIX Association.
- [35] Avinash Lakshman and Prashant Malik. Cassandra: a decentralized structured storage system. *ACM SIGOPS operating systems review*, 44(2):35–40, 2010.
- [36] Jeff LeFevre and Carlos Maltzahn. SkyhookDM: Data processing in Ceph with programmable storage. *USENIX login*., 45(2), 2020.
- [37] Sergey Legtchenko, Hugh Williams, Kaveh Razavi, Austin Donnelly, Richard Black, Andrew Douglas, Nathanael Cherière, Daniel Fryer, Kai Mast, Angela Demke Brown, Ana Klimovic, Andy Slowey, and Antony Rowstron. Understanding Rack-Scale disaggregated storage. In *9th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage 17)*, Santa Clara, CA, July 2017. USENIX Association.
- [38] Alon Y Levy, Inderpal Singh Mumick, and Yehoshua Sagiv. Query optimization by predicate move-around. In *VLDB*, pages 96–107, 1994.

- [39] Lanyue Lu, Thanumalayan Sankaranarayanan Pillai, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau. WiscKey: Separating keys from values in SSD-conscious storage. In *14th USENIX Conference on File and Storage Technologies (FAST 16)*, pages 133–148, Santa Clara, CA, February 2016. USENIX Association.
- [40] Jaehong Min, Ming Liu, Tapan Chugh, Chenxingyu Zhao, Andrew Wei, In Hwan Doh, and Arvind Krishnamurthy. Gimbal: Enabling multi-tenant storage disaggregation on SmartNIC JBOFs. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference, SIGCOMM '21*, page 106–122, New York, NY, USA, 2021. Association for Computing Machinery.
- [41] Mihir Nanavati, Jake Wires, and Andrew Warfield. Decibel: Isolation and sharing in disaggregated Rack-Scale storage. In *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*, pages 17–33, Boston, MA, March 2017. USENIX Association.
- [42] Patrick O’Neil, Edward Cheng, Dieter Gawlick, and Elizabeth O’Neil. The log-structured merge-tree (LSM-tree). *Acta Informatica*, 33(4):351–385, 1996.
- [43] Pandian Raju, Rohan Kadekodi, Vijay Chidambaram, and Ittai Abraham. PebblesDB: Building key-value stores using fragmented log-structured merge trees. In *Proceedings of the 26th Symposium on Operating Systems Principles*, pages 497–514, 2017.
- [44] Rebecca Taft, Irfan Sharif, Andrei Matei, Nathan VanBenschoten, Jordan Lewis, Tobias Grieger, Kai Niemi, Andy Woods, Anne Birzin, Raphael Poss, et al. CockroachDB: The resilient geo-distributed SQL database. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, pages 1493–1509, 2020.
- [45] Alexandre Verbitski, Anurag Gupta, Debanjan Saha, Murali Brahmadesam, Kamal Gupta, Raman Mittal, Sailesh Krishnamurthy, Sandor Maurice, Tengiz Kharatishvili, and Xiaofeng Bao. Amazon Aurora: Design considerations for high throughput cloud-native relational databases. In *Proceedings of the 2017 ACM International Conference on Management of Data, SIGMOD '17*, page 1041–1052, New York, NY, USA, 2017. Association for Computing Machinery.
- [46] Midhul Vuppalapati, Justin Miron, Rachit Agarwal, Dan Truong, Ashish Motivala, and Thierry Cruanes. Building an elastic query engine on disaggregated storage. In *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*, pages 449–462, Santa Clara, CA, February 2020. USENIX Association.
- [47] Carl Waldspurger, Trausti Saemundsson, Irfan Ahmad, and Nohhyun Park. Cache modeling and optimization using miniature simulations. In *2017 USENIX Annual Technical Conference (USENIX ATC 17)*, pages 487–498, Santa Clara, CA, July 2017. USENIX Association.
- [48] Qiongwen Xu, Michael D. Wong, Tanvi Wagle, Srinivas Narayana, and Anirudh Sivaraman. Synthesizing safe and efficient kernel extensions for packet processing. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference, SIGCOMM '21*, page 50–64, New York, NY, USA, 2021. Association for Computing Machinery.
- [49] Juncheng Yang, Yao Yue, and K. V. Rashmi. A large scale analysis of hundreds of in-memory cache clusters at twitter. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*, pages 191–208. USENIX Association, November 2020.
- [50] 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. *Proc. VLDB Endow.*, 14(11):2101–2113, jul 2021.
- [51] Zhe Yang, Youyou Lu, Xiaojian Liao, Youmin Chen, Junru Li, Siyu He, and Jiwu Shu. λ -IO: A unified IO stack for computational storage. In *21st USENIX Conference on File and Storage Technologies (FAST 23)*, pages 347–362, Santa Clara, CA, February 2023. USENIX Association.
- [52] Jie You, Jingfeng Wu, Xin Jin, and Mosharaf Chowdhury. Ship compute or ship data? why not both? In *18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21)*, pages 633–651. USENIX Association, April 2021.
- [53] Irene Zhang, Amanda Raybuck, Pratyush Patel, Kirk Olynyk, Jacob Nelson, Omar S. Navarro Leija, Ashlie Martinez, Jing Liu, Anna Kornfeld Simpson, Sujay Jayakar, Pedro Henrique Penna, Max Demoulin, Piali Choudhury, and Anirudh Badam. The demikernel datapath os architecture for microsecond-scale datacenter systems. In *Proceedings of the ACM SIGOPS 28th Symposium on Operating Systems Principles, SOSP '21*, page 195–211, New York, NY, USA, 2021. Association for Computing Machinery.
- [54] Yuhong Zhong, Haoyu Li, Yu Jian Wu, Ioannis Zarkadas, Jeffrey Tao, Evan Mesterhazy, Michael Makris, Junfeng Yang, Amy Tai, Ryan Stutsman, and Asaf Cidon. XRP: In-Kernel storage functions with eBPF. In *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)*, pages 375–393, Carlsbad, CA, July 2022. USENIX Association.

- [55] Yuhong Zhong, Hongyi Wang, Yu Jian Wu, Asaf Cidon, Ryan Stutsman, Amy Tai, and Junfeng Yang. BPF for storage: An exokernel-inspired approach. In *Proceedings of the Workshop on Hot Topics in Operating Systems*, HotOS '21, page 128–135, New York, NY, USA, 2021. Association for Computing Machinery.
- [56] Yang Zhou, Zezhou Wang, Sowmya Dharanipragada, and Minlan Yu. Electrode: Accelerating distributed protocols with eBPF. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*, 2023.



(a) Indices unpinned.



(b) Bloom filters enabled.

Figure 10: Tail latency and throughput of BPF-oF vs. NVMe-oF/TCP on RocksDB with NAND SSD for a uniform read workload for non-default configurations.

10 Appendix

10.1 RocksDB Non-Default Configurations

When running with indices not pinned, we did not run a no-cache configuration, since RocksDB is not designed to operate with neither index blocks nor data blocks cached, leading to very low throughput rates. We also run a configuration with a 1.5 GB cache, as RocksDB must cache both index and data blocks in its block cache. This minor increase allows for increased caching while also ensuring that some index blocks must be read from disk. Figure 10a compares these workloads across BPF-oF and NVMe-oF. As expected, the 1.5 GB configuration scales better than the 1 GB configuration. BPF-oF leads to a $1.8\times$ throughput improvement while maintaining lower tail latency.

To test performance with Bloom filters, we ran with both a 1 GB cache and no cache. We did not run with filter blocks not pinned, as RocksDB’s filter blocks do not provide a meaningful reduction in disk accesses if they are not stored in memory. Interestingly, enabling Bloom filters did not lead to a significant performance benefit when compared to the default configuration, whether using BPF-oF or NVMe-oF. However, BPF-oF maintained its performance improvement over NVMe-oF in this configuration as well.

10.2 WiredTiger Evaluation

We present further WiredTiger results using an Optane SSD and four cores in Figure 11.

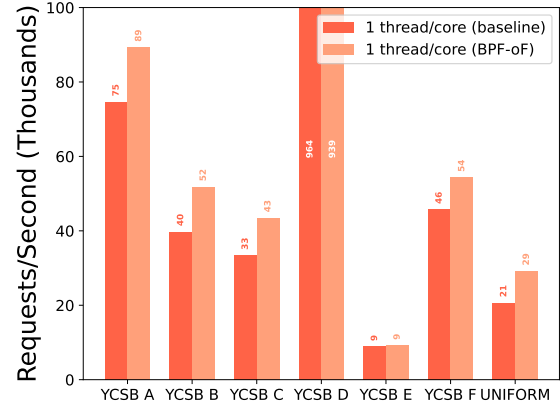


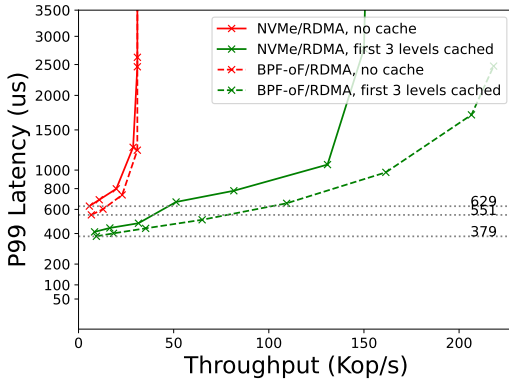
Figure 11: Throughput of WiredTiger with TCP and Optane SSD.

10.3 BPF-KV Evaluation

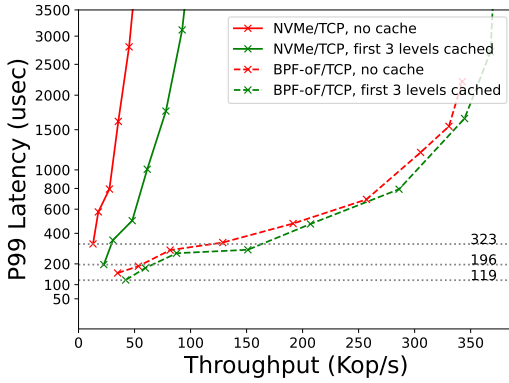
To demonstrate the performance benefit of BPF-oF, we experiment with a simple microbenchmark using a BPF-KV instance containing 8B keys with 64B values stored on the NVMe-oF target. The application on the host issues lookups of keys at uniform random; the offered load is varied by the host. Figure 9b compares the results of running BPF-oF vs. regular NVMe-oF on a NAND SSD with TCP. Beyond three cores, BPF-KV bottlenecks on the NAND disk’s I/O bandwidth. The results show that on TCP, BPF-oF consistently provides about $2.6\times$ higher throughput than NVMe/TCP when the upper levels of BPF-KV are cached, and more than $8\times$ higher throughput when BPF-KV is run without a cache. These results are consistent when we run with different numbers of cores. Note that when no cache is used, the access pattern limits the SSD to 120K IOPS (which translates to 30K requests per second at 7 resubmissions per request), because the upper level nodes become read hot-spots in the NVMe queues. The primary reason for the significant speedup is that BPF-oF saves significant CPU cycles in processing TCP packets, which can be used instead to generate, submit, and process additional I/O requests.

The results also show that BPF-oF provides a more modest 20-25% improvement in unloaded 99th-percentile latency. The tail latency improvement is more modest than the throughput improvement because NAND SSD latency (90 μ s) is significantly higher than TCP latency for our CloudLab nodes (30 μ s). Therefore, since BPF-oF reduces the network roundtrips, but does not reduce the number of storage I/O requests, the unloaded latency improvement is less significant. Still, as is evident from the graphs, BPF-oF keeps the tail latency under control as the load gets higher, in contrast to the NVME-oF baseline.

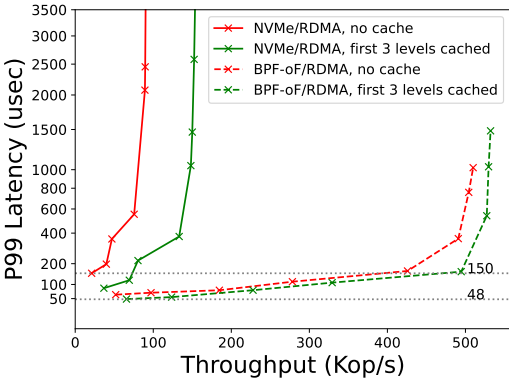
Figure 12a shows the results with RDMA and NAND. As



(a) 3 cores with RDMA, NAND SSD.



(b) 3 cores with TCP, Optane SSD.



(c) 3 cores with RDMA, Optane SSD.

Figure 12: Tail latency and throughput of BPF-oF vs. NVMe-oF on BPF-KV with random 512B reads.

expected, BPF-oF can accelerate NVMe-oF with RDMA, although the throughput gains are generally less significant than with TCP, since RDMA’s CPU processing cost is less expensive than TCP’s. The throughput peaks at around 220K because it is saturating the SSD’s bandwidth.

Figures 12b and 12c similarly compare BPF-oF to NVMe-oF using Optane SSD. In this case, the throughput improvement of BPF-oF is even more significant with TCP (a speedup consistently higher than 5 \times) since the CPU overhead of TCP is even more dominant than with NAND due to Optane’s much higher IOPS (5M reads per second). Similarly, the tail latency improvement is also more significant, and is reduced

by up to 3.1 \times . Similar to our experiments with RocksDB, in the case of BPF-oF with Optane SSD, there is almost no difference between running BPF-KV with a cache and without a cache for a uniform workload. The reason for this is that since Optane SSD is so fast (the device itself has a latency of only 3 μ s), the cached version of BPF-KV reduces the number I/O resubmissions by 3, saving a total of about 9 μ s and negligible CPU overhead. This is significantly lower than the latency (and CPU overhead) of the network, even with RDMA.