

# Kernels of Splitting TCP in the Clouds

## ABSTRACT

The Pathway project, by the Office of the CTO at VMware(OCTO), constructs a global overlay network across public clouds. Pathway interconnects geographically-dispersed corporate datacenters and branches. The project leverages the global geographical spread of public clouds and their vast compute and networking infrastructure. Clients using Pathway could see improvements in download times sometimes by orders of magnitude. The main contributor to improving network performance was shown to be the split TCP technique. Unfortunately, the available TCP split solutions incur a penalty on connection time, and actually hurt the performance of short flows. In addition, available TCP split implementations are in user-space, and thus have two expensive system calls per forwarded buffer. The redundant system calls hurt server performance and limit the scalability of such solutions.

This paper describes KTCP, a Linux Kernel module built to take full advantage of a cloud environment. The design of KTCP aims to minimize and effectively eliminate any overheads of establishing the split TCP connection. We define a theoretical model for optimal transmission and cover the methods we use to bring KTCP as close as possible to this ideal model. We demonstrate that KTCP is able to considerably increase the link utilization by TCP connections and reduce the connection *latency* close to its theoretical minimum. Furthermore, KTCP avoids the main performance penalties associated with splicing sockets i.e., system calls and memory copying.

## 1 INTRODUCTION

**Pathway.** The Pathway [20] project taps into an underutilized resource: the vast network infrastructure being built by the cloud providers [1, 16, 18]. A Pathway customer, the *client* connecting to any *server* would traverse multiple Pathway routers. In this paper, without loss of generality, we assume exactly two relays; in all figures the two assumed Pathway routers are depicted as  $R_C$ , for the router connected to the client, and  $R_S$ , for the router connected to the server. In reality, these routers are just commodity VMs hosted by the cloud provider.

**TCP split.** TCP split is a well known technique for optimizing link utilization with TCP flows. For the past 20 years, TCP split has been used for wireless networks, satellite transmissions, LAN and WAN optimizations [3, 7, 11, 12, 17, 21]. Now, as part of the the Pathway project, we are introducing TCP split into a new context of public clouds as a networking infrastructure. TCP split is standard in all WAN optimization products [22, 23, 30, 31].

### 1.1 Laying pipes in the clouds

The cloud provides an auto-scaling networking infrastructure with links of virtually-infinite capacity. As a result, flows in the cloud will rarely ever encounter congestion. In fact, like others [6], we find that in-cloud paths provide a more predictable performance than the public Internet, with a loss rate that is lower by an order of magnitude.

Given the particularly favorable conditions in the cloud, we start by exploring whether it is possible to achieve (or get closer to) an *ideal transmission* with minimal latency, which is currently quite out of reach in the Internet.

**Ideal (Figure 1a).** Figure 1a illustrates our fundamental model of an ideal transmission. Importantly, this model reflects a protocol-free, theoretical and ideal transmission; thus we are free to disregard overheads like the TCP three way handshake. In an ideal scenario, we would like to wait no more than one round-trip time (RTT), for the response to our request to start arriving. In this case the request would go through the Pathway routers directly triggering the transmission of all response packets. The time-to-first-byte (TTFB) would be just one RTT, the lowest possible TTFB.

**Real (Figure 1b).** Unfortunately, the classical end-to-end data transfer shown in Figure 1b falls short of this ideal model. In this example, we suppose that the client requests three MSS-sized packets using HTTP over TCP, that the initial TCP window size is one MSS, and that there are no losses. The end-to-end<sup>1</sup> HTTP transmission over TCP first requires the establishment of an end-to-end connection, adding one RTT to the ideal finish time. Waiting one RTT for the first ACK further delays the download.

**TCP Split (Figure 1c).** Now, when we use standard TCP-split logic in the Pathway routers, we also add new sources of delay. The basic TCP split mechanism divides the long control loop into several separate legs, each with a shorter control loop. This separation results in shorter download time for larger files. But a naive implementation also introduces additional overheads. Due to these overheads, TCP split may be detrimental for small file downloads. The main overheads stem from: (a) on-demand resource allocation on the VM, and (b) socket semantics, a server accepting connections will be notified of a new connection only once the TCP 3WHS is complete; meaning that the split connections can only be established successively.

In Section 2 we discuss TCP split overheads in greater detail. The most significant contribution of this paper is that we efficiently address the delays marked (1)-(4) in Figure 1c, and therefore reduce the connection latency to be very close to its theoretical minimum. However, if we can only change the TCP-split implementation in the relays, we also find that we cannot address the small leftover delays  $\Delta_c$  and  $\Delta_s$  on the client and server sides, respectively, and therefore cannot fully reach the ideal minimum latency. To do so, the best ways would be either by making the relays closer to the end-points (client and server), or by changing the congestion control in these end-points.

## 2 APPROXIMATING THE IDEAL PIPE

To approximate the ideal data transmission model, we introduce *KTCP*(K=Kernel-based TCP split). The goal of KTCP is to provide

<sup>1</sup>In this case the Pathway routers just forward packets, without any additional logic.

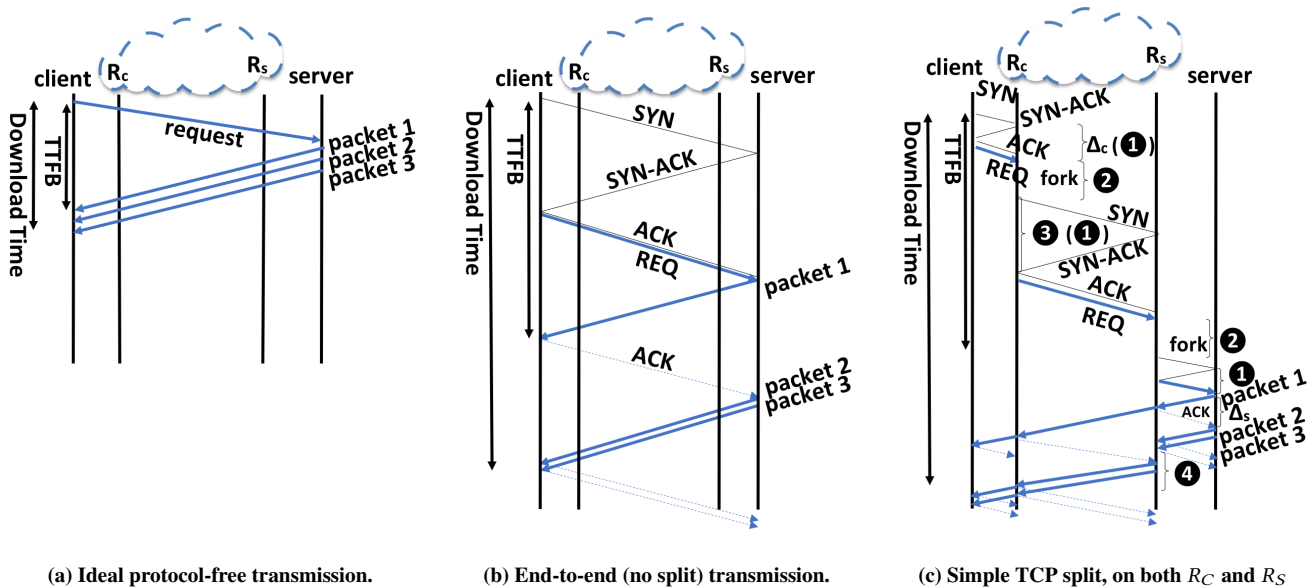


Figure 1: Illustrated comparison of the considered baseline data transmission methods.

an efficient, delay-free TCP optimization while utilizing commodity VMs and standard programming APIs. We introduce four improvements over the naive TCP split approach. The effects of these improvements are illustrated in Figure 2.

**Improvement 1: Early SYN (Figure 2a).** In Early-SYN [10, 17], the relay server sends a SYN packet to the next-hop server as soon as the SYN packet arrives, without waiting for the three-way handshake to complete. KTCP captures this first SYN packet and triggers the start of a new connection. This allows the proxy to establish the two legs of a split connection in parallel. Using Early-SYN, we can remove the SYN-ACK and ACK delays, marked as (1) in Figure 1c. Out of our 4 suggested improvements, this is the only one that was already known in the literature.

**Improvement 2: Thread pool (Figure 2b).** The creation of new processes or threads for each new split connection is time-consuming and adds greatly to the connection jitter. Some outliers may take tens of milliseconds, greatly hurting performance. For small files/objects, this jitter may even nullify the benefit of KTCP. To mitigate this problem, we create a pool of reusable threads. These are sleeping threads, awaiting to accept new tasks. Using a *thread pool* removes the delays marked as (2) in Figure 1c.

**Improvement 3: Reusable connections (Figure 2c).** This optimization aims to improve the performance of long-haul connections, *i.e.*, those where the RTT between the two cloud relays dominates. The goal is to negate the delay of the long three-way handshake. We achieve this goal by preemptively connecting to distant Pathway routers. In each Pathway router we create a pool of pre-established connections between each pair of distant Pathway routers. With a *connection pool*, delay (3) in Figure 1c is eliminated.

**Improvement 4: Turbo-Start TCP (Figure 2d).** Congestion is not an issue within the cloud, hence, there is essentially no need to

use TCP’s slow-start mechanism. It is redundant to probe the network when a connection is established between two Pathway routers within the same cloud provider. We thus configure a large initial congestion window (CWND) and large receive window (RWIN), thus eliminating delay (4) in Figure 1c. In addition, we increase the socket buffers for the relay machines, so that memory would not limit the performance of the intra-cloud flows. Note that we do not change the CWND used on any Internet-facing flows. We wish to remain friendly to other TCP flows potentially sharing a bottleneck link with our Pathway routers.

### 3 KTCP DESIGN

KTCP is a Linux kernel-based module, incorporating the four optimizations described in Section 2. In this section we discuss the various design and implementation details.

**Kernel mode.** We implemented KTCP as a kernel module. We rely on *procfs* [28] to control the behaviour of KTCP. *Procfs*, a virtual file system [32] provides a simple interface and facilitates easy scripting that allows communication with the module at run time. The decision to use kernel mode is a significant one. While developing in user space would have provided an easier development environment, implementing KTCP in the kernel allows us to (1) take advantage of resources only available in the kernel, such as *Netfilter* [27], which is crucial to our needs; and (2) avoid the penalties that stem from numerous *system calls* [13, 19]. By working in the kernel, we eliminate the redundant transitions to and from user space by avoiding gratuitous system calls. *Netfilter*, which provides hooks for callbacks on the network stack, is the standard option for capturing and processing a packet in various stages of its path through the network stack.

The decision to implement the components of KTCP in the kernel is further made easy by the fact that all socket APIs have kernel

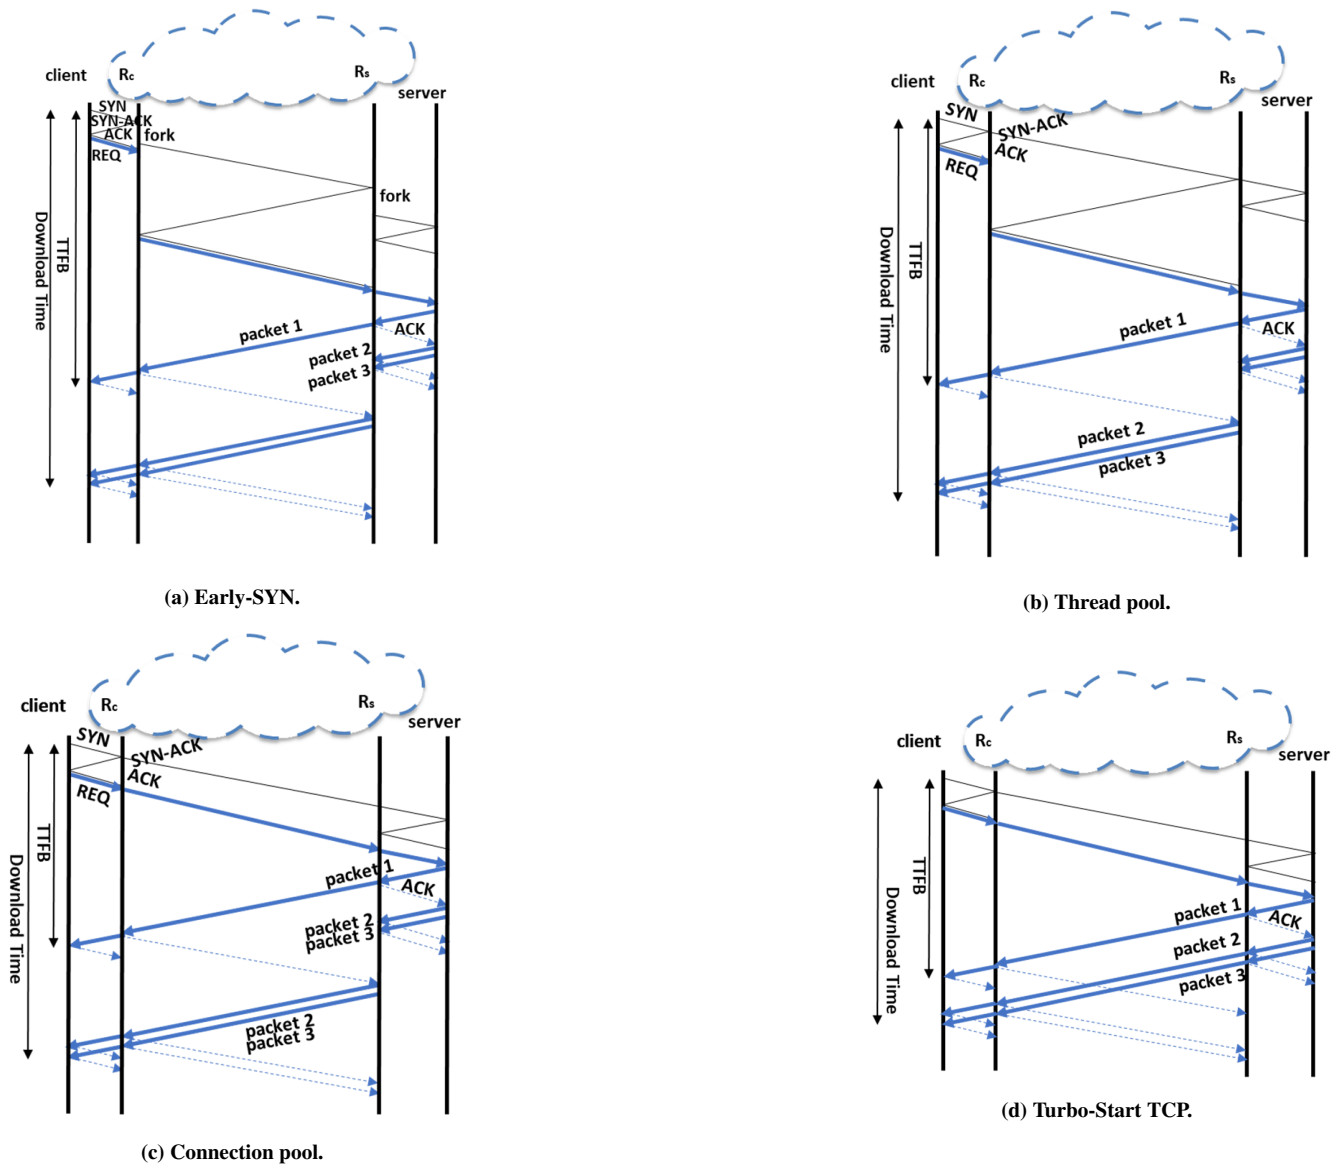


Figure 2: KTCP successive implementation improvements

counterparts.<sup>2</sup> Instead of POSIX threads we utilize kernel\_threads to handle each leg of the split connections.

**Cost of System Calls** To better gauge the cost of system calls that we are avoiding by working in the kernel, we measure the cost of context switching between kernel\_threads on our setup, assuming an n1-standard(GCP) VM with Intel Skylake Xeon CPU. Specifically, we measure the time it takes for two kernel\_threads to call schedule() 10 million times each. This experiment completes in under 3.2 seconds, resulting in 0.16  $\mu$ sec per context switch on average; by comparison, an analogous experiment with two processes runs for

<sup>2</sup>One limitation of an in-kernel implementation is that epoll [24], a scalable I/O event notification mechanism, has no kernel counterpart.

15.6 seconds and 12.9 seconds for two POSIX threads [29]. The experiment highlights the cost of using blocking system calls. The time it takes to send/receive a packet is in the range of several  $\mu$ sec [13]. These numbers hint at the direct impact system calls would have had on the performance of KTCP. The actual cost of sending and receiving packets varies, depending on hardware used, the size of packets, system configuration and system load.<sup>3</sup>

**Basic implementation.** The basic implementation of KTCP, on each Pathway router, has three components. (1) A socket listening for incoming connections. (2) Iptable [25] rules that redirect TCP connections to our proxy socket. (3) A second TCP socket is used to

<sup>3</sup>A load on a shared resource like the L3 cache or the memory controller can impact performance greatly [14].

connect to the next server and thus complete the second leg of the split connection. Once both connections are established, the bytes of a single stream are read from one socket, and then forwarded to its peer. This forwarding happens in both directions. When either connection is terminated via an error or a FIN packet, the other connection is shut down as well. This means that the bytes in flight (*i.e.*, not yet acked) will reach their destination, but no new bytes can be sent.

**Memory Footprint.** We found that the size of the buffer used to read and write the data is important. At first we used a 4KB buffer, and experienced degraded performance. Starting at 16KB buffer sizes the performance levels out. Each split connection has two sockets and two kernel\_threads [26], with each kernel\_thread taking 9KB in addition to 16KB for the forwarding in each direction. Each split connection takes 50KB of memory to maintain.

**Early-SYN.** As there is no standard API that enables the capture of the first SYN packet, we use Linux Netfilter [27] hooks. We add a hook that captures TCP packets, and then parse the headers for the destination IP and the SYN flag. With this information, KTCP launches a new kernel\_thread. It is impossible to create a new thread while in the Netfilter callback, which is an atomic context. We use our thread pool to launch kernel\_threads from atomic contexts. The “new” thread initiates a connection to the intended destination. Capturing the SYN allows the Pathway router to establish the two sides of a connection concurrently.

**Thread pool.** We use blocking send/receive calls with our sockets allowing for a simple implementation; this also means that we need a kernel\_thread per active socket. Unfortunately, the creation of a new kernel\_thread is costly. On our setup, a kernel\_thread creation takes about 12 $\mu$ sec, on average.<sup>4</sup> But an outlier may consume several milliseconds, resulting in a jittery behaviour.

To mitigate this problem and the problem of creating new kernel\_threads from atomic context, we create a pool of reusable threads. Each kernel\_thread in this pool is initially waiting in state TASK\_INTERRUPTIBLE (ready to execute). When the thread is allocated, two things happen: (1) a function to execute is set and (2) the task is scheduled to run (TASK\_RUNNING). When the function is finished executing, the thread returns to state TASK\_INTERRUPTIBLE and back to the list of pending threads, awaiting to be allocated once more. A pool of pre-allocated kernel threads thus removes the overhead of new kernel\_thread creation. A new kernel\_thread from the waiting pool can start executing immediately and can be launched from any context. KTCP attempts to keep the number of threads in a pool between two configurable watermarks. KTCP creates new threads when the number drops below the low mark and frees from the pool when the limit grows above the high mark. On a multi-core system, the heavy lifting of thread creation is offloaded to a dedicated core.

**Pre-established connections.** For pre-established connections, we have added a dedicated server thread that accepts connections from other Pathway routers that may create new pre-established connections to this server.<sup>5</sup> When established, these connections wait for

the target address to be sent from the initiating peer. The destination address is sent over the connection itself. This information is sent in the very first bytes, and all following bytes belong to the forwarded stream. Once the destination address is received by  $R_S$ , a connection to the target is initiated and the final leg of the split connection is created.

For example, a SYN captured on  $R_C$  will trigger a look-up in the pool of pre-established connections on  $R_C$ ; looking for an existing TCP connection between  $R_C$  and  $R_S$ . If a pre-established connection exists it will be used for the split connection.  $R_C$  will send the target information inside the connection. Upon receiving the target information from  $R_C$ ,  $R_S$  will establish a TCP connection with the target, and from that point on  $R_S$  will forward the stream from  $R_C$  to the target. When the connection terminates the socket associated with the connection will be freed and will not be reused.

On pre-established connection sockets Nagle’s Algorithm [15] is disabled. In our experiments, we have seen that the TTFB is increased by some 200 milliseconds, unless Nagle’s Algorithm is disabled.

**Proc.** The water marks of the thread-pool, the destination and number of pre-established connections are controlled via the procs [28] interface. These parameters can be modified at run time.

**Multi-core systems.** Shared resources that are used by multiple cores may result in degraded performance due to contention. In KTCP only the two pools, *i.e.*, the pool of pre-established connections and the thread pool, form a shared resource. To avoid contention, we have implemented a generic magazine [2] infrastructure; it is used to implement both pools.

The asynchronous creation of a new split connection initially results in two threads each holding only one socket. One thread has the socket with an established connection to the source, and the other holds the socket with an established connection to the next server. For the split connection to work, both threads need to find their peer. This look up is facilitated by a red-black tree (rb\_tree).<sup>6</sup> This rb\_tree holds elements with three fields, the two sockets and a 12 Byte key, we call these elements QP (queue pair). The 12 Byte key is simply the source/destination IP/port. To avoid contention over a single rb\_tree, we create an rb\_tree on each core. When conducting the look up, both threads are bound to run on a specific core; this core is determined by a xor on the 12 Byte key. Once the lookup is complete, both threads are free to be scheduled anywhere on the system, and the QP is removed from the rb\_tree.

**Kernel Zero Copy.** While KTCP is able to sidestep system calls and context switch costs by utilizing kernel threads; the cost of copying remains high. We expand the existing Linux TCP API with a tcp\_read\_sock\_zcopy for RX and add a new msg flag SOCK\_KERN\_ZEROCOPY for tcp\_sendmsg\_locked in TX. We base our new function tcp\_read\_sock\_zcopy on existing infrastructure *i.e.*, tcp\_read\_sock. It is used by tcp\_splice\_read to collect skbs from a socket. For TX, zero copy infrastructure already exists in the form of MSG\_ZEROCOPY[4]. When kernel memory is used for I/O, enabling zero copy is trivial when compared to zero copy from user space. The pages are already pinned in memory and there is not need for a notification on TX completion. The pages are reference counted, and can be freed by the device driver completion handler.

<sup>6</sup>A generic implementation already exists in the Linux Kernel

<sup>4</sup>By comparison a fork consumes more than 25 $\mu$ sec, while launching a POSIX pthread consumes around 13 $\mu$ sec.

<sup>5</sup>In order to keep the connection from closing before being used, the sockets are configured with KEEP\_ALIVE.




Figure 3: Time-To-First-Byte (TTFB) in seconds. Median results of 50 runs.

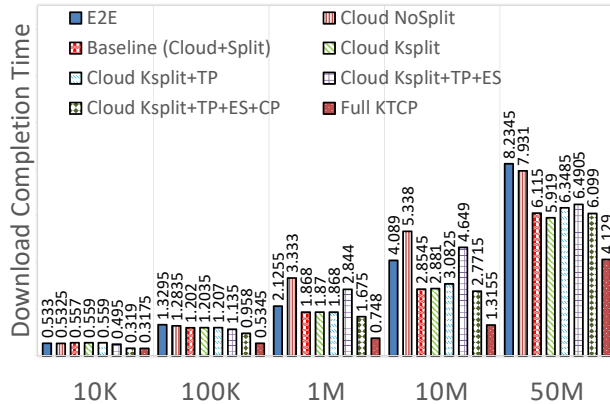


Figure 4: Download completion time in seconds. Median results of 50 runs.

## 4 EXPERIMENTAL EVALUATION

To evaluate the contribution of each of these improvements, we set up a server in Bangalore as a VM on a Digital Ocean infrastructure; and a client PC in San-Francisco, connected using a residential Internet service<sup>7</sup>. Our relays are two VMs on GCP’s public cloud: one in Mumbai, close to the server ( $R_S$ ), and another ( $R_C$ ) in Oregon, near the client. Both VMs run on Ubuntu 17.04 and use small (n1-standard-1) machines with a single vCPU and 3.75 GB memory. The client and server are Ubuntu 16.04 machines.

We set up an Apache web server on the Bangalore VM and evaluate the performance of each of the options by measuring both download times and time-to-first-byte (TTFB). We experiment with files of different sizes that the client is downloading from the server, using HTTP via the curl utility. The times reported in Figure 3 and Figure 4 are as returned by the curl utility.

<sup>7</sup>The ISP is Comcast.

The RTT (as measured by ICMP) between the client and  $R_C$  is 32.7ms, between  $R_C$  and  $R_S$  is 215ms, and between  $R_S$  and the server is 26ms.

We compare the performance of the following configurations: (i) simple End-to-End (*e2e*); (ii) routing through the cloud relays; using iptable’s DNAT, without splitting the TCP flow (Cloud NoSplit); (iii) splitting the TCP flow using SSH’s port forwarding feature (*Baseline*); (iv) TCP splitting using our KTCP kernel module, set up to use the improvements listed in section 2: thread pool only (KTCP+TP), thread pool and early-SYN (KTCP+TP+ES), a complete bundle of the three connection improvements including pre-established connection (KTCP+TP+ES+CP), and finally also configuring the intra-cloud TCP connection to use Turbo-Start (*KTCP*).

The benefit from the improvements is best observed by looking at the Time-To-First-Byte in Figure 3. We can see that the TTFB and total download time of the basic KTCP coincide with those of the Baseline. Our basic kernel-based implementation performs at least as well as the well-established ssh utility. We also note that KTCP+TP does not improve the median performance by a noticeable amount. However, we have noticed throughout our testing that the thread pool improves the stability of our results.

For all file sizes we notice an improvement of  $\approx 60ms$  when using KTCP+TP+ES. This is in line with Figure 1c, as Early-SYN eliminates one RTT on the (client  $\leftrightarrow R_C$ ) and another RTT of the ( $R_S \leftrightarrow$  server). The amount of time reduced, according to our RTT measurements is supposed to be 59ms, in line with our results. Adding pre-established connections to the mix should potentially reduce the TTFB by one RTT of the ( $R_C \leftrightarrow R_S$ ) leg. However, since the REQ cannot be forwarded with the SYN packet sent by the client (without any TCP extensions), we can only gain  $215 - 33 = 182ms$ . Indeed, the benefit of adding CP as evident in Figure 3 is of  $\approx 180ms$ . The addition of Turbo-Start does not reduce the TTFB, as it only influences the way packets are sent after the first bytes. The contribution of Turbo-Start is clearly evident when considering the total download time (Figure 4). We see considerable reduction of the file download time when using Turbo-Start for all file sizes, except that of 10 KB file. The default initial congestion window size for both Ubuntu 17.04 and 16.04 is 10 segments, so the entire file is sent in a single burst. Indeed, the results show that for 10 KB file the download completion time is about 1 ms after the first byte arrives. All other improvements contribute to the reduction of TTFB, and so reduce the total download time by roughly the same amount. This reduction is barely noticeable for large files, where the main benefits stem from splitting the TCP flow and using Turbo-Start. In this experiment we notice that the best performing implementation improvement (*i.e.*, KTCP) outperforms *e2e* file transfer by up to 3 times! (depending on the file size).

## 5 RELATED WORK

To our knowledge KTCP is the first fully featured in-kernel implementation of TCP split. One previous work [7] uses Netfilter hooks to forward packets but doesn’t maintain a TCP stack per connection, instead using its own bookkeeping code for packet re-transmission. A different approach [17] uses unikernels [9] and a modified lwip [5] as the basis of their proxies. Both approaches lack the proven stability, reliability and versatility of the Linux kernel, all needed for

a reliable product.

In parallel to our own work, a new API i.e., SOCKMAP [8, 23] was introduced. SOCKMAP aims to solve the performance problems of splicing two sockets, using BPF programs. But SOCKMAP its not flexible enough to facilitate all Pathway requirements (e.g., Early-SYN Fig.2a).

## REFERENCES

- [1] AWS. Announcing network performance improvements for amazon ec2 instances. <https://aws.amazon.com/about-aws/whats-new/2018/01/announcing-network-performance-improvements-for-amazon-ec2-instances/>. Published Jan 26, 2018.
- [2] BONWICK, J., AND ADAMS, J. Magazines and vmem: Extending the slab allocator to many cpus and arbitrary resources. In *USENIX Annual Technical Conference, General Track* (2001), pp. 15–33.
- [3] CHAKRAVORTY, R., KATTI, S., PRATT, I., AND CROWCROFT, J. Flow aggregation for enhanced TCP over wide area wireless. In *INFOCOM* (2003), pp. 1754–1764.
- [4] DE BRUIJN, W., AND DUMAZET, E. sendmsg copy avoidance with msg\_zerocopy.
- [5] DUNKELS, A. Design and implementation of the lwip tcp/ip stack. *Swedish Institute of Computer Science 2* (2001), 77.
- [6] HAQ, O., RAJA, M., AND DOGAR, F. R. Measuring and improving the reliability of wide-area cloud paths. In *WWW* (2017), pp. 253–262.
- [7] JAIN, R., AND OTT, T. J. *Design and implementation of split TCP in the linux kernel*. PhD thesis, New Jersey Institute of Technology, Department of Computer Science, 2007.
- [8] JOHN, F. Bpf: sockmap and sk redirect support. <https://lwn.net/Articles/731133/>, Aug 2017. Accessed Jan 2020.
- [9] KIVITY, A., COSTA, D. L. G., AND ENBERG, P. Os v—optimizing the operating system for virtual machines. In *Proceedings of USENIX ATC'14: 2014 USENIX Annual Technical Conference* (2014), p. 61.
- [10] LADIWALA, S., RAMASWAMY, R., AND WOLF, T. Transparent tcp acceleration. *Computer Communications 32*, 4 (2009), 691–702.
- [11] LE, F., NAHUM, E., PAPPAS, V., TOUMA, M., AND VERMA, D. Experiences deploying a transparent split TCP middlebox and the implications for NFV. In *HotMiddlebox* (2015), pp. 31–36.
- [12] LUGLIO, M., SANADIDI, M. Y., GERLA, M., AND STEPANEK, J. On-Board Satellite "Split TCP" Proxy. *IEEE Journal on Selected Areas in Communications 22*, 2 (2004), 362–370.
- [13] MARKUZE, A., MORRISON, A., AND TSAFRIR, D. True iommu protection from dma attacks: When copy is faster than zero copy. *ACM SIGARCH Computer Architecture News 44*, 2 (2016), 249–262.
- [14] MARKUZE, A., SMOLYAR, I., MORRISON, A., AND TSAFRIR, D. Damn: Overhead-free iommu protection for networking. In *Proceedings of the Twenty-Third International Conference on Architectural Support for Programming Languages and Operating Systems* (2018), ACM, pp. 301–315.
- [15] NAGLE, J. Congestion control in ip/tcp internetworks. RFC 896, RFC Editor, January 1984. Accessed: Jan 2018.
- [16] SILVA, G. Maximize your VM's performance with Accelerated Networking - Microsoft Azure. <https://azure.microsoft.com/en-us/blog/maximize-your-vm-s-performance-with-accelerated-networking-now-generally-available-for-both-windows-and-linux/>. Published Jan 5, 2018.
- [17] SIRACUSANO, G., BIFULCO, R., KUENZER, S., SALSANO, S., MELAZZI, N. B., AND HUICI, F. On the fly tcp acceleration with miniproxy. In *HotMiddlebox* (2016), ACM, pp. 44–49.
- [18] SLOSS, B. T. Expanding our global infrastructure with new regions and subsea cables - Google Cloud. <https://www.blog.google/topics/google-cloud/expanding-our-global-infrastructure-new-regions-and-subsea-cables/>. Published Jan 16, 2018.
- [19] SOARES, L., AND STUMM, M. Flexsc: Flexible system call scheduling with exception-less system calls. In *Proceedings of the 9th USENIX conference on Operating systems design and implementation* (2010), USENIX Association, pp. 33–46.
- [20] VMWARE. Project pathway. <https://research.vmware.com/projects/cloudified-business-network>.
- [21] Akamai technologies. [http://www.akamai.com/dl/feature\\_sheets/fsedgesuite\\_surerule.pdf](http://www.akamai.com/dl/feature_sheets/fsedgesuite_surerule.pdf).
- [22] Sisco. Sd/wan. [https://sdwan-docs.cisco.com/Product\\_Documentation/Software\\_Features/Release\\_18.2/Network\\_Optimization/02Configuring\\_TCP\\_Optimization](https://sdwan-docs.cisco.com/Product_Documentation/Software_Features/Release_18.2/Network_Optimization/02Configuring_TCP_Optimization).
- [23] Cloudflare. <https://blog.cloudflare.com/sockmap-tcp-splicing-of-the-future/>.
- [24] epoll. <http://man7.org/linux/man-pages/man7/epoll.7.html>. Accessed: Jan 2018.
- [25] The netfilter.org "iptables" project. <https://www.netfilter.org/projects/iptables/index.html>.
- [26] Kernel threads made easy. <https://lwn.net/Articles/65178/>.
- [27] The netfilter.org project. <http://www.netfilter.org/>. Accessed: Jan 2018.
- [28] Procs. <http://www.tldp.org/LDP/Linux-Filesystem-Hierarchy/html/proc.html>.
- [29] POSIX threads. <http://man7.org/linux/man-pages/man7/pthreads.7.html>.
- [30] Riverbed: Sd/wan. <https://www.riverbed.com/newsletter/how-wan-optimization-works.html>.
- [31] Silverpeak: Unity boost. <https://www.silver-peak.com/products/unity-edge-connect>.
- [32] A tour of the linux vfs. <http://www.tldp.org/LDP/khg/HyperNews/get/fs/vfstour.html>.