NetKernel: Making Network Stack Part of the Virtualized Infrastructure

Zhixiong Niu, Qiang Su[®], Peng Cheng[®], Yongqiang Xiong[®], Dongsu Han[®], *Member, IEEE*, Keith Winstein[®], Chun Jason Xue[®], *Senior Member, IEEE*, and Hong Xu[®], *Senior Member, IEEE, ACM*

Abstract—This paper presents a system called NetKernel that decouples the network stack from the guest virtual machine and offers it as an independent module. NetKernel represents a new paradigm where network stack can be managed as part of the virtualized infrastructure. It provides important efficiency benefits: By gaining control and visibility of the network stack, operators can perform network management more directly and flexibly, such as multiplexing VMs running different applications to the same network stack module to save CPU cores, and enforcing fair bandwidth sharing. Users also benefit from the simplified stack deployment and better performance: For example mTCP can be deployed without API change to support nginx natively, and shared memory networking can be readily enabled to improve performance of colocated VMs. Testbed evaluation using 100G NICs shows that NetKernel preserves the performance and scalability of both kernel and userspace network stacks, and provides the same isolation as the current architecture.

Index Terms—Network stack, virtualization, VM, re-architecting.

I. INTRODUCTION

WIRTUAL machines (VMs) are the predominant virtualization form in today's cloud due to their strong isolation guarantees. VMs allow customers to run applications in a wide variety of operating systems (OSes) and configurations. VMs are also heavily used by cloud operators to deploy internal services, such as load balancing, proxy, VPN, etc., both in a public cloud for tenants and in a private cloud for various business units of an organization.

Manuscript received September 11, 2020; revised June 8, 2021; accepted November 12, 2021; approved by IEEE/ACM TRANSACTIONS ON NET-WORKING Editor S. Rao. The work of Keith Winstein was supported in part by the Sloan Research Fellowship, in part by the National Science Foundation (NSF) under Grant CNS-1909212 and Grant CNS-1763256, in part by Google, in part by Huawei, in part by VMware, in part by Dropbox, in part by Amazon, and in part by Facebook. The work of Hong Xu was supported in part by the Hong Kong Research Grants Council (RGC) General Research Fund (GRF) under Grant 11210818 and Grant 11209520; and in part by The Chinese University of Hong Kong (CUHK) under Grant 4937007, Grant 4937008, Grant 5501329, and Grant 5501517. Part of the work has appeared in USENIX Annual Technical Conference (ATC) 2020 [49] [DOI: 10.5555/3489146.3489156]. (*Corresponding author: Hong Xu.*)

Zhixiong Niu, Peng Cheng, and Yongqiang Xiong are with the Networking Research Group, Microsoft Research Asia, Beijing 100080, China.

Qiang Su and Chun Jason Xue are with the Department of Computer Science, City University of Hong Kong, Hong Kong.

Dongsu Han is with the School of Electrical Engineering, Korea Advanced Institute of Science and Technology, Daejeon 34141, South Korea.

Keith Winstein is with the Department of Computer Science, Stanford University, Stanford, CA 94305 USA.

Hong Xu is with the Department of Computer Science and Engineering, The Chinese University of Hong Kong, Hong Kong (e-mail: henryxu@ cse.cuhk.edu.hk).

Digital Object Identifier 10.1109/TNET.2021.3129806

VM based virtualization largely follows traditional OS design. In particular, the TCP/IP network stack is encapsulated inside the VM as part of the guest OS as shown in Figure 1(a). Applications interface with the network stack, which is transparent for operators, using the virtual NIC abstraction. This architecture preserves the familiar hardware and OS abstractions so a vast array of workloads can be easily moved into the cloud. It also provides high flexibility for applications to customize the entire network stack.

We argue that the current division of labor between application and network infrastructure is becoming increasingly inadequate, especially in a private cloud setting. The central issue is that the network stack is controlled solely by individual guest VM; the operator has almost zero visibility or control. This leads to efficiency problems that manifest in various aspects of running the cloud network. Firstly, the operator is unable to orchestrate resource allocation at the end-points of the network fabric, resulting in low resource utilization. It remains difficult today for the operator to meet or define performance SLAs despite much prior work [15], [33], [40], [54], [55], as she cannot precisely provision resources just for the network stack or control how the stack consumes these resources. Further, resources (e.g. CPU) have to be provisioned on a per-VM basis based on the peak traffic; it is impossible to coordinate across VM boundaries. This degrades the overall utilization of the network stack since in practice traffic to individual VMs is extremely bursty. Also, many network management tasks like monitoring, diagnosis, and troubleshooting have to be performed in an extra layer outside the guest VMs, which requires significant efforts in design and implementation [21], [58], [59]. They can be done more efficiently if the network stack is opened up to the operator.

Even the simple task of maintaining or deploying a network stack suffers from much inefficiency today. Numerous new stack designs and optimizations ranging from congestion control [12], [17], [48], scalability [32], [41], zerocopy datapath [4], [32], [53], [64], NIC multiqueue scheduling [63], etc. have been proposed in our community. Yet the operator, with sufficient expertise and resources, could not easily deploy these solutions in a virtualized cloud to improve performance and reduce overheads because it does not own or control the network stack. As a result, our community is still finding ways to deploy DCTCP in the public cloud [18], [29], [34]. On the other hand, applications without much knowledge of the underlying network or expertise on networking are forced to juggle the deployment and maintenance details. For example

1558-2566 © 2021 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.



(a). Existing architecture (b). Decoupling network stack from the guest

Fig. 1. Decoupling network stack from the guest, and making it part of the virtualized infrastructure.

if one wants to deploy a new stack like mTCP [32], he faces a host of problems such as setting up kernel bypass, testing with kernel versions and NIC drivers, and porting applications to the new APIs. Given the intricacy of implementation and the velocity of development, it is a daunting task for individual users, whether tenants in a public cloud or first-party services in a private cloud, to maintain the network stack all by themselves.

To address these limitations, we advocate the *separation* of network stack from the guest OS as a new paradigm, in which the network stack is coupled into the virtualized infrastructure of cloud. Operators can directly deploy, manage, optimize the network stack, and comprehensively promote the efficiency without affecting applications' execution. As the heavy-lifting is taken care of, applications can just use network stack as a basic service of the infrastructure and focus on their business logic.

More concretely, as shown in Figure 1(b), we propose to decouple the VM network stack from the guest OS. We keep the network APIs such as BSD sockets intact, and use them (instead of vNIC) as the abstraction boundary between application and infrastructure. Each VM is served by an external network stack module (NSM) that runs the network stack chosen by the user, *e.g.*, the kernel-bypass stack mTCP or the improved kernel stack FastSocket [41]. Application data are handled in the NSM, which is designed and implemented by the operator. Various network stacks can be provided as different NSMs to ensure applications with diverse requirements can work. This new paradigm does not necessarily enforce a single transport design, or trade off such flexibility of the existing architecture.

We make three specific contributions in this paper.

- We design and implement NetKernel that demonstrates our new approach is feasible on existing KVM virtualization platforms (§3–§5). NetKernel provides transparent BSD socket redirection so existing applications can run directly.
- We present NetKernel's benefits by showcasing novel use cases that are difficult. For example, we show that NetKernel enables multiplexing: One NSM can serve multiple VMs at the same time and save over 40% CPU cores without degrading performance using traces from a production cloud.

• We conduct a comprehensive testbed evaluation with commodity 100GbE NICs to show that NetKernel achieves the same scalability and isolation as the current architecture: Using a Linux kernel stack NSM, it reaches 100G send throughput with 3 cores; using an mTCP NSM it achieves 979K RPS with 8 cores which almost doubles that of the kernel stack NSM.

NetKernel's official website is https://netkernel.net.

II. MOTIVATION

Decoupling the network stack from the guest OS and making it part of the infrastructure marks a clear departure from the way networking is provided to VMs nowadays. In this section we elaborate why this is a better architectural design by presenting its benefits and tradeoffs, and contrasting it with alternative solutions.

A. Benefits and Tradeoffs

We highlight key benefits of our vision with several new use cases that we experimentally realize with NetKernel in §7.

Better efficiency in management for the operator. Gaining control over the network stack, the operator can now perform network management more efficiently. For example it can orchestrate the resource provisioning strategies more flexibly: For mission-critical workloads, it can dedicate CPU resources to their NSMs to offer performance SLAs in terms of throughput and RPS (requests per second) guarantees. For elastic workloads, on the other hand, it can consolidate their VMs to the same NSM (if they use the same network stack) to improve its resource utilization. The operator can also directly implement management functions as an integral part of user's network stack, compared to doing them in an extra layer outside the guest OS.

Use case 1: Multiplexing (§7.1). Utilization of network stack in VMs is very low most of the time in practice. Using a real trace from a large cloud, we show that NetKernel enables multiple VMs to be multiplexed onto one NSM to serve the aggregated traffic and saves over 40% CPU cores for the operator without performance degradation.

Use case 2: Fair bandwidth sharing (§7.2). TCP's notion of flow-level fairness leads to poor bandwidth sharing in data centers [59]. We show that NetKernel can facilitate VM-level congestion control [59] in a NSM, achieving fair sharing regardless of the number of flows.

Deployment and performance gains for users. Making network stack part of the virtualized infrastructure is also beneficial for users in both public and private clouds. Various kernel stack optimizations [41], [64], high-performance userspace stacks [9], [16], [32], [53], and even designs using advanced hardware [6], [7], [42] can now be deployed and maintained transparently without user involvement or application code change. For instance, DCTCP can now be deployed across the board easily in a public cloud. Since the BSD socket is the only abstraction exposed to the applications, it is now feasible to adopt new stack designs independent of the guest kernel or the network API. Our vision also opens up new design space by allowing the network stack to exploit visibility of the infrastructure for performance benefits.

TABLE I

COMPARISON OF DIFFERENT NETWORK STACK ARCHITECTURES DEPENDING ON WHERE THE STACK IS. THE CURRENT ARCHITECTURE IS A GUEST-BASED PARADIGM WHERE THE NETWORK STACK IS PART OF THE GUEST OS OF A VM

Paradigm	Scenario	Multiplexing	Multiplexing New Stack Deployment	
Guest-based	VM	X	Require user effort	X
Host-based	Container	✓	Limited by host OS	1
Application-based	Library OS	X	Require user effort	X
NetKernel	VM + NSM	✓	✓	✓

Use case 3: Deploying mTCP without API change (§7.3). We show that NetKernel enables unmodified applications in the VM to use mTCP [32] in the NSM, and improves performance greatly due to mTCP's kernel bypass design. mTCP is a userspace stack with new APIs (including modified epoll/kqueue). During the process, we also find and fix a compatibility issue between mTCP and our NIC driver, and save significant maintenance time and effort for users.

Use case 4: Shared memory networking (§7.4). When two VMs of the same user are colocated on the same host, NetKernel can detect this and copy their data via shared memory to bypass TCP stack processing and improve throughput. This is difficult to achieve today as VMs have no knowledge about the underlying infrastructure [39], [65].

Tradeoffs. We are conscientious of the tradeoffs our approach brings about. For example, due to the removal of vNIC and redirection from the VM's own network stack, some networking tools like netfilter are affected. This is acceptable since most users wish to focus on their applications instead of tuning a network stack. If they wish to gain maximum control over the network stack they can still use VMs without NetKernel. Also, additional fate-sharing may be introduced by our approach say when multiple VMs share the same NSM. We believe this is not serious because cloud users already have fate-sharing with the vSwitch, hypervisor, and the complete virtual infrastructure. The efficiency benefits of our approach as demonstrated outweigh the marginal increase of fate-sharing; the success of cloud computing these years is another strong testament to this tradeoff. NetKernel enforces another level of indirection in order to achieve flexibility which does not cause performance degradation in most cases as we will show in §6, and part of it can run on hardware for more efficiency (see §8). Lastly, one may have security concerns with using the NSM to handle tenant traffic. Most of the security protocols such as HTTPS/TLS work at the application layer and are not affected. One exception is IPSec. Due to an issue with exchanging certificates, IPSec does not work in our approach. However, in practice IPSec is implemented at dedicated gateways instead of end-hosts [61]. Thus we believe the impact is not serious. More discussion on security can be found in §8.

B. Alternative Solutions

We now discuss several alternative architectures depending on where the network stack resides, and why they are inadequate compared to NetKernel as summarized in Table I. Note that none of them provides all four key benefits as NetKernel does. **Host-based containers.** The first alternative is the container scenario in the cloud, where the network stack runs on the host machine (*a.k.a.* host-based). A container is essentially a process with namespace isolation: It shares the host's network stack in the hypervisor. Therefore containers can achieve some of NetKernel's benefits, i.e., multiplexing and performance optimization with infrastructure, since the operator can access the hypervisor. However, container has tight coupling with the host OS which makes the stack deployment difficult. A Windows application in a container cannot use the Linux-based mTCP, unless the operator ports mTCP to Windows. With NetKernel no such porting is needed: mTCP can run in a Linux-based NSM and serve a Windows user as the only coupling is the BSD socket APIs.

We also note that currently containers have performance isolation problems [37] and as a result are usually constrained to be deployed inside VMs in production settings. In fact we find that all major public clouds [1], [2], [5] require users to launch containers inside VMs. Thus, our work is centered around VMs that cover the majority of usage scenarios in a cloud. NetKernel readily benefits containers running inside VMs as well.

Application-based. Another alternative is to move the network stack upwards by taking an application-based paradigm. A representative scenario is library OS including unikernels [20], [43] and microkernels [25], where many OS services including the network stack are packaged as libraries and compiled with the application in userspace. Similar to the guest-based paradigm, users who are unwilling to know any details of the network stack still have to deploy the network stack by themselves though the I/O performance can be improved with unikernels [45] and microkernels. Unexpected security risks also exist as these stacks reside in the application address space and bypass the kernel check. In addition, application-based paradigm is a clean-slate approach and requires radical changes to both the virtualization software and user applications. NetKernel can flexibly decouple the network stack from the guest without re-writing existing applications or hypervisor.

III. DESIGN PHILOSOPHY

NetKernel answers three fundamental design questions regarding the separation of network stack from the guest OS:

- 1) How should socket API calls be transparently redirected without changing applications?
- 2) How should the socket semantics be transmitted between the VM and NSM whose implementation of the stack may vary?

3) How to ensure high performance with semantics transmission (e.g., 100 Gbps)?

These questions touch upon a largely uncharted territory in the design space. Thus our main objective in this paper is to demonstrate feasibility of our approach on existing virtualization platforms and showcase its potential. Performance and overhead are not our primary goals. It is also not our goal to improve any particular network stack design.

In answering the questions above, NetKernel's design has the following highlights.

1) Transparent socket API redirection. NetKernel needs to redirect BSD socket calls to the NSM instead of the tenant network stack. This is done by inserting into the guest a library called GuestLib. The GuestLib provides a new socket type called NetKernel socket with a complete implementation of BSD socket APIs. It replaces all TCP and UDP sockets when they are created with NetKernel sockets, effectively redirecting them without changing applications.

2) A lightweight semantics channel. Different network stacks may run as different NSMs, so NetKernel needs to ensure socket semantics from the VM work properly with the actual NSM stack implementation. For this purpose NetKernel builds a lightweight socket semantics channel between VM and its NSM. The channel relies on small fix-sized queue elements as intermediate representations of socket semantics: Each socket API call in the VM is encapsulated into a queue element and sent to the NSM, who would effectively translate the queue element into the corresponding API call of its network stack.

3) Scalable lockless queues. As NIC speed in cloud evolves from 40G/50G to 100G [22] and higher, the NSM has to use multiple cores for the network stack to achieve line rate. NetKernel thus adopts scalable lockless queues to ensure VM-NSM socket semantics transmission is not a bottleneck. Each core services a dedicated set of queues so performance is scalable with number of cores. More importantly, each queue is memory shared with a software switch, so it can be lockless with only a single producer and a single consumer to avoid expensive lock contention [31], [32], [41].

Switching the queue elements offers important benefits beyond lockless queues. It facilitates a flexible mapping between VM and NSM: An NSM can support multiple VMs without adding more queues compared to binding the queues directly between VM and NSM. In addition, it allows dynamic resource management: CPU cores can be readily added to or removed from an NSM, and a user can switch her NSM on the fly. The CPU overhead of software switching can be addressed by hardware offloading [22], [26], which we discuss in §6.4 in more detail.

VM based NSM. Lastly we discuss an important design choice regarding the NSM. The NSM can take various forms. It may be a full-fledged VM with a monolithic kernel. Or it can be a container or module running on the hypervisor, which is appealing because it consumes less resource and offers better performance. NetKernel is compatible with both hypervisor- and VM-based NSMs, and we currently choose to use VM-based NSMs. VM based NSM readily supports existing kernel and userspace stacks from various OSes. VMs





Fig. 2. NetKernel design overview.

also provide good isolation and we can dedicate resources to an NSM to guarantee performance. VM based NSM is the most flexible: We can run various network stacks independent of the hypervisor.

IV. DESIGN

Figure 2 depicts NetKernel's architecture. The BSD socket APIs are transparently redirected to a complete NetKernel socket implementation in GuestLib in the guest kernel (§4.1). Network stacks are implemented by the operator on the same host as Network Stack Modules (NSMs), which are individual VMs in our current design. Inside the NSM, a ServiceLib interfaces with the network stack. The NSM connects to the vSwitch, be it a software/hardware switch, and then the pNICs. Thus our design also supports SR-IOV.

All socket operations and their results are translated into NetKernel Queue Elements (NQEs) by GuestLib and ServiceLib (§4.2). For NQE transmission, GuestLib and ServiceLib each has a NetKernel device, or NK device in the following, consisting of one or more sets of lockless queues. Each queue set has a send queue and receive queue for operations with data transfer (e.g. send()), and a job queue and *completion queue* for control operations without data transfer (e.g. setsockopt()). Each NK device connects to a software switch called CoreEngine, which runs on the hypervisor and performs actual NQE switching (§4.3). The CoreEngine is also responsible for various management tasks such as setting up the NK devices, ensuring isolation among VMs, etc. (§4.4) A unique set of hugepages are shared between each VM-NSM tuple for application data exchange. A NK device also maintains a hugepage region that is memory mapped to the corresponding application hugepages as in Figure 2 (§4.5). Note that as the socket API that copies data is preserved, misbehaving applications cannot pose security risks on NetKernel, this is the same as original kernel design. We discuss additional security implications of NetKernel in §8.

For ease of presentation, we assume both the user VM and NSM run Linux, and the NSM uses the kernel stack.

A. Transparent Socket API Redirection

We first describe how NetKernel's GuestLib interacts with applications to support BSD socket semantics transparently.

NetKernel socket API. GuestLib creates a new type of sockets—SOCK_NETKERNEL, in addition to TCP (SOCK STREAM) and UDP (SOCK DGRAM) sockets.

TABLE II NetKernel Socket Implementation

syscall	inet_stream_ops	netkernel_pro
bind	inet_bind()	nk_bind()
connect	inet_connect()	nk_connect()
accept	inet_accept()	nk_accept()
poll	tcp_poll()	nk_poll()
ioctl	inet_ioctl()	nk_ioctl()
listen	inet_listen()	nk_listen()
shutdown	inet_shutdown()	nk_shutdown()
setsockopt	<pre>sock_common_setsockopt()</pre>	nk_setsockopt()
recvmsg	tcp_recvmsg()	nk_recvmsg()
sendmsg	tcp_sendmsg()	nk_sendmsg()

It registers a complete implementation of BSD socket APIs as shown in Table II to the guest kernel. When the guest kernel receives a socket() call to create a new TCP socket say, it replaces the socket type with SOCK_NETKERNEL, creates a new NetKernel socket, and initializes the socket data structure with function pointers to NetKernel socket implementation in GuestLib. The sendmsg() for example now points to nk_sendmsg() in GuestLib instead of tcp_sendmsg().

Kernel space API redirection. There are essentially two approaches to redirect BSD socket calls to NSM, each with its unique tradeoffs. One is to implement it in userspace using LD_PRELOAD for example. The advantages are: (1) It is efficient without syscall overheads and performance is high [32]; (2) It is easy to deploy without kernel modification. However, this implies each application needs to have its own redirection service, which limits the usage scenarios. Another way is kernel space redirection, which naturally supports multiple applications without IPC. The flip side is that performance may be lower due to context switching and syscall overheads.

We opt for kernel space API redirection to support most of the usage scenarios, and leave userspace redirection as future work. This is feasible by distributing images of para-virtualizated guest kernels to users, a practice operators are already doing nowadays. Note that kernel space redirection follows the asynchronous syscall model [60] to get better performance.

Guest change. The GuestLib can be deployed as a kernel patch as the kernel space redirection is implemented as a kernel module. NetKernel supports zero code change for the guest applications, and an application can use NetKernel by only setting the socket type as SOCK_NETKERNEL when invoking the socket() syscall.

B. A Lightweight Semantics Channel

Socket semantics are contained in NQEs and carried around between GuestLib and ServiceLib via their respective NK devices.

NQE and socket semantics translation. Figure 3 shows the structure of a NQE with a fixed size of 32 bytes. Translation happens at both ends of the semantics channel: GuestLib encapsulates the socket semantics into NQEs and sends to ServiceLib, which then invokes the corresponding API of its network stack to execute the operation; the execution result is again turned into a NQE in ServiceLib first, and then translated by GuestLib back into the corresponding response of socket APIs.

1B	1B	1B	4B	8B	8B	4B	5B
op type	VM ID	Queue set ID	VM socket ID	op_data	data pointer	size	rsved

Fig. 3. Structure of a NQE. Here socket ID denotes a pointer to the sock struct in the user VM or NSM, and is used for NQE transmission with VM ID and queue set ID in §4.3; op_data contains data necessary for socket operations, such as IP address for bind; data pointer is a pointer to application data in hugepages; and size is the size of pointed data in hugepages.



Fig. 4. NetKernel socket implementation in GuestLib redirects socket API calls. GuestLib translates socket API calls to NQEs and ServiceLib translates results into NQEs as well (not shown here).

For example in Figure 4, to handle the <code>socket()</code> call in the VM, GuestLib creates a new NQE with the operation type and information such as its VM ID for NQE transmission. The NQE is transmitted by GuestLib's NK device. The <code>socket()</code> call now blocks until a response NQE is received. After receiving the NQE, ServiceLib parses the NQE from its NK device, invokes the <code>socket()</code> of the kernel stack to create a new TCP socket, prepares a new NQE with the execution result, and enqueues it to the NK device. GuestLib then receives and parses the response NQE and wakes up the <code>socket()</code> call. The <code>socket()</code> call now returns to application with the NetKernel socket file descriptor (fd) if a TCP socket is created at the NSM, or with an error number consistent with the execution result of the NSM.

We defer the handling of application data to §4.5.

Queues for NQE transmission. NQEs are transmitted via one or more sets of queues in the NK devices. A queue set has four independent queues: A *job queue* for NQEs representing socket operations issued by the VM without data transfer, a *completion queue* for NQEs with execution results of control operations from the NSM, a *send queue* for NQEs representing operations issued by VM with data transfer; and a *receive queue* for NQEs representing events of newly received data from NSM. Queues of different NK devices have strict correspondence: The NQE for socket() for example is put in the job queue of GuestLib's NK device, and sent to the job queue of ServiceLib's NK device. Note that this won't increase the number of queues that are serviced by the network stacks (*e.g.*, mTCP).

We now present the working of I/O event notification mechanisms like epoll with the receive queue. Figure 5 depicts the details. Suppose an application issues epoll_wait() to monitor some sockets. Since all sockets are now NetKernel sockets, the nk_poll() is invoked by epoll_wait() and checks the receive queue to see if there is any NQE for this socket. If yes, this means there are new data received,



Fig. 5. The socket semantics channel with epoll as an example. GuestLib and ServiceLib translate semantics to NQEs, and queues in the NK devices perform NQE transmission. Job and completion queues are for socket operations and execution results, send queues are for socket operations with data, and receive queues are for events of newly received data. Application data processing is not shown.

epoll_wait() then returns and the application issues a recv() call with the NetKernel socket fd of the event. This points to nk_recvmsg() which parses the NQE from receive queue for the data pointer, copies data from the hugepage directly to the userspace, and returns.

We adopt an interrupt-driven polling method (§4.6) to mitigate the CPU overhead. If nk_poll() does not find any relevant NQE, it sleeps until CoreEngine wakes up the NK device when new NQEs arrive to its receive queue. GuestLib then parses the NQEs to check if any sockets are in the epoll instances, and wakes up the epoll to return to application. An epoll_wait() can also be returned by a timeout.

C. NQE Switching Across Lockless Queues

We now elaborate how NQEs are switched by CoreEngine and how the NK devices interact with CoreEngine.

Scalable queue design. The queues in a NK device is scalable: There are one dedicated queue set per vCPU for both VM and NSM, so NetKernel performance scales with CPU resources. Each queue set is shared memory with the CoreEngine, essentially making it a single producer single consumer queue without lock contention. VM and NSM may have different numbers of queue sets.

Switching NQEs in CoreEngine. NQEs are load balanced across multiple queue sets with the CoreEngine acting as a switch. CoreEngine maintains a connection table as shown in Figure 6, which maps the tuple $\langle VM \text{ ID}, \text{queue set ID}, \text{ socket ID} \rangle$ to the corresponding $\langle NSM \text{ ID}, \text{queue set ID}, \text{ socket ID} \rangle$ and vice versa. Here a socket ID corresponds to a pointer to the sock struct in the user VM or NSM. We call them VM tuple and NSM tuple respectively. NQEs only contain VM tuple information.

Using the running example of the socket () call, we can see how CoreEngine uses the connection table. The process is also shown in Figure 6. (1) When CoreEngine processes the socket NQE from VM1's queue set 1, it realizes this is a new connection, and inserts a new entry to the table with the VM tuple from the NQE. (2) It checks which NSM should handle it,¹ performs hashing based on the three tuple

¹A user VM to NSM mapping is determined either by the users/operator offline or some load balancing scheme dynamically by CoreEngine.



Fig. 6. NQE switching with CoreEngine.

to determine which queue set (say 2) to switch to if there are multiple queue sets, and copies the NQE to the NSM's corresponding job queue. CoreEngine adds the NSM ID and queue set ID to the new entry. (3) ServiceLib gets the NQE and copies the VM tuple to its response NQE, and adds the newly created connection ID in the NSM to the op_data field of response NQE. (4) CoreEngine parses the response NQE, matches the VM tuple to the entry and adds the NSM socket ID to complete it, and copies the response NQE to the completion queue 1 of VM1 as instructed in the NQE. Later NQEs for this VM connection can be processed by the correct NSM connection and vice versa. ServiceLib pins its connections to its vCPUs and queue sets, so processing the NQE and sending the response NQE are done on the same CPU.

The connection table allows flexible multiplexing and demultiplexing with the socket ID information. For example one NSM can serve multiple VMs using different sockets. CoreEngine polls all queue sets to maximize performance.

D. Management With CoreEngine

CoreEngine acts as the control plane of NetKernel and carries out many control tasks beyond NQE switching.

NK device and queue setup. CoreEngine allocates shared memory for the queue sets and sets up the NK devices accordingly when a VM or NSM starts up, and de-allocates when they shut down. Queues can also be dynamically added or removed with the number of vCPUs.

Isolation. CoreEngine sits in an ideal position to carry out isolation among VMs. In our design CoreEngine polls each queue set in a round-robin fashion to ensure the basic fair sharing. Operator can implement other isolation mechanisms to rate limit a VM in terms of bandwidth or the number of NQEs (i.e. operations) per second, which we show in §6.3. Note that CoreEngine isolation happens for egress; ingress isolation at the NSM is more challenging and may have to use physical NIC queues [19].

Busy-polling. The busy-polling design of CoreEngine requires a dedicated core per machine which is an inherent overhead of our design. We resort to this simple design as we focus on showing feasibility and potential of NetKernel in this work, and prior work also used dedicated cores for software polling [39]. One can explore hardware offloading using FPGAs for example to eliminate this overhead [21], [22].

E. Processing Application Data

We now discuss the last missing piece of NetKernel design: How application data are actually processed in the system. Sending data. Data is transmitted by hugepages shared between the VM and NSM. Their NK devices maintain a hugepage region that is mmaped to the application hugepages. For sending data with send(), GuestLib copies data from userspace directly to the hugepage, and adds a data pointer to the send NQE. It also increases the send buffer usage for this socket similar to the send buffer size maintained by the kernel. The send() now returns to application. ServiceLib invokes tcp_sendmsg() provided by the kernel stack upon receiving the send NQE. Data are obtained from hugepages, processed by the network stack, and sent to the vNIC. A new NQE is generated with the result of send by the NSM and sent to GuestLib, who then decreases the send buffer usage.

Receiving data. Now for receiving packets in the NSM, a normal network stack would send received data to userspace applications. In order to send received data to the user VM, ServiceLib then copies the data chunk to hugepages and create a new NQE to the receive queue, which is then sent to the VM. It also increases the receive buffer usage for this connection, similar to the send buffer maintained by GuestLib described above. The rest of the receive process is already explained in §4.2. Note that application uses recv() to copy data from hugepages to their own buffer.

ServiceLib. As discussed ServiceLib deals with much of data processing at the NSM side so the network stack works in concert with the rest of NetKernel. One thing to note is that unlike the kernel space GuestLib, ServiceLib should live in the same space as the network stack to ensure best performance. We have focused on a Linux kernel stack with a kernel space ServiceLib here. The design of a userspace ServiceLib for a userspace stack is similar in principle. ServiceLib busy-polls its queues for maximum performance.

F. Optimization

We present several optimizations employed in NetKernel.

Pipelining. NetKernel applies pipelining between VM and NSM for performance. For example on the VM side, a send() returns immediately after putting data to the hugepages, instead of waiting for the actual send result from the NSM. Similarly the NSM would handle accept() by accepting a new connection and returning immediately, before the corresponding NQE is sent to GuestLib and then application to process. Doing so does not break BSD socket semantics. Take send() for example. A successful send() does not guarantee delivery of the message [11]; it merely indicates the message is written to socket buffer successfully. In NetKernel a successful send() indicates the message is written to buffer in the hugepages successfully. As explained in §4.5 the NSM sends the result of send back to the VM to indicate if the socket buffer usage can be decreased or not.

Interrupt-driven polling. We adopt an interrupt-driven polling design for NQE event notification to GuestLib's NK device. This is to reduce the overhead of GuestLib and user VM. When an application is waiting for events e.g. the result of the socket() call or receive data for epoll, the device will first poll its completion queue and receive queue. If no new NQE comes after a short time period $(20\mu \text{ s in our experi$ $ments})$, the device sends an interrupt to CoreEngine, notifying that it is expecting NQE, and stops polling. CoreEngine later wakes up the device, which goes back to polling mode to process new NQEs from the completion queue. This is similar in spirit to busy-polling sockets in Linux kernel [3], [8]. Interrupt-driven polling presents a favorable trade-off between overhead and performance compared to pure polling based or interrupt based design. It saves precious CPU cycles when load is low and ensures the overhead of NetKernel is very small to the user VM. Performance on the other hand is competent since the response NQE is received within the polling period in most cases for blocking calls, and when the load is high polling automatically drives the notification mechanism. As explained before CoreEngine and ServiceLib use busy polling to maximize performance. In addition, the software polling overhead can also be alleviated through a hardware/software notification mechanism [24], [62].

Batching. Batching is used in many parts of NetKernel for better throughput. CoreEngine uses batching whenever possible for polling from and copying into the queues. The NK devices also receive NQEs in a batch.

V. IMPLEMENTATION

Our implementation is based on QEMU KVM 2.5.0 and Linux kernel 4.9 for host and guest, with over 11K LoC.

GuestLib. We add the SOCK_NETKERNEL socket to the kernel (net.h), and modify socket.c to rewrite the SOCK_STREAM to SOCK_NETKERNEL during socket creation. We implement GuestLib as a kernel module with two components: Guestlib_core and nk_driver. Guestlib_core is mainly for Netkernel sockets and NQE translation, and nk_driver is for NQE communications via queues. Guestlib_core and nk_driver communicate with each other using function calls.

ServiceLib and NSM. We also implement ServiceLib as two components: Servicelib_core and nk_driver. Servicelib_core translates NQEs to network stack APIs, and the nk_driver is identical to the one in GuestLib. For the kernel stack NSM, Servicelib_core calls the kernel APIs directly to handle socket operations without entering userspace. We create an independent kthread to poll the job queue and send queue for NQEs to avoid kernel stuck. Some BSD socket APIs cannot be invoked in kernel space directly. We use EXPORT_SYMBOLS to export the functions for ServiceLib. Meanwhile, the boundary check between kernel space and userspace is disabled. We use per-core epoll_wait() to obtain incoming events from the kernel stack.

We also port mTCP [10] as a userspace stack NSM. It uses DPDK 17.08 for packet I/O. For simplicity, we maintain its two-thread model and per-core data structure. We implement the NSM in mTCP's application thread at each core. The ServiceLib is essentially an mTCP application: Once receiving a NQE from its send queue, it accesses data from the shared hugepage by the data pointer in the NQE and sends it using mTCP with DPDK. For receiving, the received data is copied into the hugepage, and ServiceLib encapsulates the data pointer into a NQE of the receive queue. The per-core application thread (1) translates NQEs polled from the NK device to mTCP socket APIs, and (2) responds

IEEE/ACM TRANSACTIONS ON NETWORKING

NQEs to the tenant VM based on the network events collected by mtcp_epoll_wait(). Since mTCP works in non-blocking mode for performance, we buffer send operations at each core and set the timeout parameter to 1ms in mtcp_epoll_wait() to avoid starvation when polling NQE requests.

Queues and hugepages. The hugepages are implemented based on QEMU's IVSHMEM. The page size is 2 MB and we use 128 pages. The queues are ring buffers implemented as much smaller IVSHMEM devices. Together they form a NK device which is a virtual device to the VM and NSM.

CoreEngine. The CoreEngine is a daemon with two threads on the KVM hypervisor. One thread listens on a pre-defined port to handle NK device (de)allocation requests, namely 8-byte network messages of the tuples $\langle ce_op, ce_data \rangle$. When a VM (or NSM) starts (or terminates), it sends a request to CoreEngine for registering (or deregistering) a NK device. If the request is successfully handled, CoreEngine responds in the same message format. Otherwise, an error code is returned. The other thread polls NQEs in batches from all NK devices and switches them as described in §4.3.

VI. MICROBENCHMARKS

We present microbenchmarks of crucial aspects of NetKernel: Performance and multicore scalability in §6.2; isolation of multiple VMs in §6.3; and system overhead in §6.4.

A. Setup

Each of our testbed servers has two Xeon E5-2698 v3 16-core CPUs clocked at 2.3 GHz, 256 GB memory at 2133 MHz, and a Mellanox ConnectX-4 single port 100G NIC. Hyperthreading is disabled. We compare to the status quo where an application uses the kernel TCP stack in its VM, referred to as Baseline in the following. We designate NetKernel to refer to the common setting where we use the kernel stack NSM in our implementation. When mTCP NSM is used we explicitly mark the setting in the results. The same TCP parameter settings are used for both systems. The NSM uses the same number of vCPUs as Baseline since CPU is used almost entirely by the network stack in Baseline. NetKernel allocates 1 more vCPU for the VM to run the application and ServiceLib throughout the evaluation. Its CPU utilization is usually low: We report the actual CPU overheads of NetKernel in §6.4. The throughput results are measured by iperf and the RPS results are measured by ab, unless stated otherwise. The throughput results are averaged over 5 runs each lasting 30 seconds.

B. Performance and Scalability

We now look at NetKernel's basic performance.

NQE switching and memory copy. NQEs are transmitted by CoreEngine as a software switch. It is important that CoreEngine offers enough horsepower to ensure performance at 100G and higher. We measure CoreEngine throughput which is defined as the number of 32-byte NQEs copied from GuestLib's NK device queues to the ServiceLib's NK device queues. Table III shows the results with varying batch sizes. CoreEngine achieves ~8M NQEs/s without batching. With a small batch size of 4 or 8 throughput reaches 41.4M NQEs/s



Fig. 7. Send throughput of 8 TCP streams with varying numbers of vCPUs, 8KB messages.



Fig. 8. Recv throughput of 8 TCP streams with varying numbers of vCPUs, 8KB messages.



Fig. 9. Performance of TCP short connections with multiple vCPUs. Message size 64B.

and 65.9M NQEs/s, respectively, which is sufficient for most applications.

We also measure the memory copy throughput between GuestLib and ServiceLib via hugepages. A memory copy in this experiment includes the following: (1) application in the VM issues a send() with data; (2) GuestLib gets a pointer from the hugepages; (3) copies the message to hugepages; (4) prepares a NQE with the data pointer; (5) CoreEngine copies the NQE to ServiceLib; and (6) ServiceLib obtains the data pointer and puts it back to the hugepages. Thus it measures the effective application-level throughput using NetKernel (including NQE transmission) without network stack processing.

We observe from Table IV that NetKernel delivers over 100G throughput with messages larger than 4KB: with 8KB messages 144G is achievable. Thus NetKernel provides enough raw performance to the network stack and is not a bottleneck to the 100G deployment in production.

Single TCP Stream. We benchmark the single stream TCP throughput with different message sizes. The send and receive throughput are shown in Figure 10 and Figure 11. NetKernel

NIU et al.: NetKernel: MAKING NETWORK STACK PART OF VIRTUALIZED INFRASTRUCTURE

TABLE III CoreEngine Switching Throughput Using a Single Core With Different Batch Sizes

Batch Size (B)	1	2	4	8	16	32	64	128	256
NQEs per second ($\times 10^6$)	8.0	14.4	22.3	41.4	65.9	100.2	119.6	178.2	198.5

TABLE IV Message Copy Throughput via Hugepages With Different Message Sizes



Fig. 10. Single TCP stream send throughput with the kernel stack NSM. The NSM uses 1 vCPU.



Fig. 11. Single TCP stream receive throughput with the kernel stack NSM. The NSM uses 1 vCPU.

performs on par with Baseline in all cases, this is because a single TCP stream is bounded by the single-core kernel stack processing. Note that send throughput reaches 30.9Gbps and receive throughput tops at 13.6Gbps in NetKernel. Receive throughput is much lower because the kernel stack's RX processing is much more CPU-intensive with interrupts. If the other cores of the NUMA node are not disabled, soft interrupts (softirq) may be sent to those cores instead of the one assigned to the NSM (or VM), thereby inflating the receive throughput.²

Multiple TCP Streams. We examine throughput performance using the kernel stack NSM and 8 TCP streams with 8KB messages. Figures 7 and 8 show respectively the send and receive throughput with varying number of vCPUs. NetKernel achieves the same throughput performance and scalability with Baseline. The single-core send and receive throughput reaches 48Gbps and 17Gbps, respectively. Receive throughput is much lower because the kernel stack's RX processing is much more CPU-intensive with interrupts. Note that if the other cores of the NUMA node are not disabled, soft interrupts (softirq) may be sent to those cores instead of the one assigned to the

TABLE V

THROUGHPUT SCALING AND SHORT CONNECTIONS WITH VARYING NUMBERS OF NETKERNEL WITH KERNEL STACK NSM EACH WITH TWO VCPUS

# of 2-vCPU NSMs	1	2	3	4
Send throughput (Gbps)	85.1	94.0	94.1	94.2
Receive throughput (Gbps)	33.6	61.2	91.0	91.0
Requests per sec ($x10^3$)	131.6	260.4	399.1	520.1

NSM (or VM), thereby inflating the receive throughput. Both systems achieve the line rate of 100G using at least 3 vCPUs for send throughput as in Figure 7. For receive, both achieve 91Gbps using 8 vCPUs as in Figure 8.

Short TCP connections. We also benchmark NetKernel's performance in handling short TCP connections using a custom server sending a short message as a response. The server runs multiple worker threads that share the same listening port. Each thread runs an epoll event loop. Our workload generates 10 million requests in total with a concurrency of 1000. The connections are non-keepalive. The message size is 64B. Socket option SO_REUSEPORT is always used for the kernel stack. Figure 9 shows that NetKernel has the same multicore scalability as Baseline: Performance increases from \sim 71K RPS with 1 vCPU to \sim 400K RPS with 8 vCPUs, i.e. 5.6x the single core performance. To demonstrate NetKernel's full capability, we also run the mTCP NSM with 1, 2, 4, and 8 vCPUs.³ NetKernel with mTCP offers 167K, 313K, 562K, and 979K RPS respectively, and shows better scalability than the kernel stack.

The results here show that NetKernel preserves the performance and scalability of network stacks, including high performance stacks like mTCP since our scalable queue design can ensure NetKernel is not the bottleneck and the contention is not severe in this situation.

NetKernel scalability. We now examine the scalability of NetKernel's overall design. With the same epoll server, we use 8 connections and 8KB messages for throughput experiments and 10 millions of requests with 64B messages for short connections experiments. Each kernel stack NSM now uses 2 vCPUs. The servers in different NSMs listen on different ports and does not share an accept queue. We vary the number of NSMs to serve this 1-core VM.

Table V shows the throughput scaling results. Throughput for send is already 85.1Gbps with 2 vCPUs (recall Figure 7), and adding NSMs does not improve it beyond 94.2Gbps. Throughput for receive shows almost linear scalability on the other hand. Performance of short connections also exhibits near linear scalability: One NSM provides 131.6K RPS, 2 NSMs 260.4K RPS, and 4 NSMs 520.1K RPS which is 4x better. The results indicate that NetKernel's design is highly scalable.

C. Isolation

Isolation is important to ensure co-located users do not interfere with each other, especially in a public cloud. It is different from fair sharing: Isolation ensures a VM's performance guarantee is met despite network dynamics, while

²We observe 30.6Gbps receive throughput with 16KB messages in both NetKernel and Baseline when leaving the other cores on.

³Using other numbers of vCPUs for mTCP causes stability problems even without NetKernel.

IEEE/ACM TRANSACTIONS ON NETWORKING



Fig. 12. VM 1 is capped at 1Gbps, VM2 at 500Mbps, and VM3 uncapped. All VMs use the same kernel stack NSM. The NSM is assigned 10Gbps bandwidth. NetKernel isolates VM1 and VM2 successfully while allowing VM3 to obtain the remaining capacity.

TABLE VI

DISTRIBUTION OF RESPONSE TIMES (MS) FOR 64B MESSAGES WITH 5 MILLION REQUESTS AND 1K CONCURRENCY

	Baseline	NetKernel	mTCP baseline	NetKernel, mTCP NSM
Min	0.6	0.8 (+33.3%)	0.7	3.1 (+ 342.9%)
Mean	16.0	16.4 (+2.5%)	4.0	4.2 (+ 5.0%)
Median	2.1	2.2 (+4.8%)	3.9	4.0 (+ 2.6%)

fairness ensures a VM obtains a fair share of the bottleneck capacity which varies dynamically. We conduct an experiment to verify NetKernel's isolation guarantees. As discussed in §4.4, CoreEngine uses round-robin to poll each VM's NK device for basic fairness. In addition, to achieve isolation we implement token buckets in CoreEngine to limit the bandwidth of each VM, taking into account varying message sizes. There are 3 VMs now: VM1 is rated limited at 1Gbps, VM2 at 500Mbps, and VM3 has unlimited bandwidth. They arrive and depart at different times. They are colocated on the same host running a kernel stack NSM using 1 vCPU. The NSM is given a 10G VF for simplicity of showing work conservation.

Figure 12 shows the time series of each VM's throughput, measured by our epoll server at 100ms intervals. VM1 joins the system at time 0 and leaves at 25s. VM2 comes later at 4.5s and leaves at 21s. VM3 joins last and stays until 30s. We can observe that NetKernel throttles VM1's and VM2's throughput at their respective limits correctly despite the dynamics. VM3 is also able to use all the remaining capacity of the 10G NSM: It obtains 9Gbps after VM2 leaves and 10Gbps after VM1 leaves at 25s. Therefore, NetKernel is able to achieve the same isolation in today's clouds with bandwidth caps.

D. Overhead

Latency. One may wonder if NetKernel with the NQE transmission would add delay to TCP processing, especially in handling short connections. Table VI shows the latency statistics when we run ab to generate 1K concurrent connections to our epoll server for 64B messages. A total of 5 million requests are used. We can see NetKernel incurs less than 5% latency slowdown in terms of the mean and median values, and the latency overhead is mostly introduced by CoreEngine. Even for the mTCP NSM, NetKernel preserves its low latency due to the much simpler TCP stack processing and various optimization [32]. Note that NetKernel doesn't perform well in the minimum latency, which is not a primary concern here.

CPU. Now to quantify NetKernel's CPU overhead, we use the epoll server at the VM side, and run clients from a different machine with fixed throughput or requests per second for both NetKernel and Baseline with kernel TCP stack. We disable

TABLE VII

OVERHEAD FOR THROUGHPUT. THE NSM RUNS THE LINUX KERNEL TCP STACK. WE USE 8 TCP STREAMS WITH 8KB MESSAGES. NETKERNEL'S CPU USAGE IS NORMALIZED OVER THAT OF BASELINE

Throughput	20Gbps	40Gbps	60Gbps	80Gbps	100Gbps
Normalized CPU usage	1.14	1.28	1.42	1.56	1.70

TABLE VIII

OVERHEAD FOR SHORT TCP CONNECTIONS. THE NSM RUNS THE KERNEL TCP STACK. WE USE 64B MESSAGES WITH A CONCURRENCY OF 100

Requests per second (RPS)	100K	200K	300K	400K	500K
requests per second (10.5)	1001	2001	0001	1001	0001
Normalized CPU usage	1.06	1.05	1.08	1.08	1.09
i wormanzed er o usage	1.00	1.00	1.00	1.00	1.07

all unnecessary background system services in both the VM and NSM, and ensure the CPU usage is almost zero without running epoll servers. During the experiments, we measure the total number of cycles spent by the VM in Baseline, and that spent by the VM and NSM together in NetKernel. We then report NetKernel's CPU usage normalized over Baseline's for the same performance level in Tables VII and VIII.

We can see that to achieve the same throughput, NetKernel incurs relatively high overhead especially as throughput increases. To put things into perspective, we measure CPU usage when the client runs in a docker container with the bridge networking mode. Docker incurs 13% CPU overhead compared to Baseline to achieve 40 Gbps throughput whereas NetKernel's is 28%. The overhead here is due to the extra memory copy from the hugepages to the NSM. It can be optimized away by implementing zerocopy between the hugepages and the NSM, which we are working on currently.

Table VIII shows NetKernel's overhead with short TCP connections. We can observe that the overhead ranges from 5% to 9% in all cases and is mild. As the message is only 64B here, the results verify that the NQE transmission overhead in NK devices is small.

Lastly, throughout all experiments of our evaluation we dedicated one core to CoreEngine, which is another overhead. As we focus on showing feasibility and potential of NetKernel in this paper, we resort to software NQE switching which attributes to the polling overhead. Fortunately, this overhead can be mitigated and even removed by offloading the progress onto hardwares like FPGAs. Therefore CoreEngine does not consume CPU for the majority of the NQEs: Only the first NQE of a new connection needs to be handled in CPU (direct to a proper NSM as in §4.3).

To quickly recap, current NetKernel implementation incurs CPU overheads especially for extra data copy and CoreEngine. We believe, however, they are not inevitable when separating the network stack from the guest OS. They can be largely mitigated using known implementation techniques which is left as future work. In addition, as shown in §7.1 multiplexing can be used in NetKernel's current implementation to actually save CPU compared to dedicating cores to individual VMs.

VII. NEW USE CASES

We present some new use cases that are realized using our prototype to demonstrate the potential of NetKernel.



Fig. 13. Traffic of three most utilized application gateways (AGs) in our trace. They are deployed as VMs.



Fig. 14. Per-core RPS comparison. Baseline uses 12 cores for 3 AGs, while NetKernel with multiplexing only needs 9 cores.

Details of the performance and overhead microbenchmarks are presented in §6.

A. Multiplexing

Here we describe a new use case where the operator can optimize resource utilization by serving multiple bursty VMs with one NSM.

To make things concrete we draw upon a user traffic trace collected from a large cloud in September 2018. The trace contains statistics of tens of thousands of application gateways (AGs) that handle tenant (web) traffic in order to provide load balancing, proxy, and other services. The AGs are internally deployed as VMs by the operator. We find that the AG's average utilization is very low most of the time. Figure 13 shows normalized traffic processed by three most utilized AGs (in the same datacenter) in our trace with 1-minute intervals for a 1-hour period. We can clearly see the bursty nature of the traffic. Yet it is very difficult to consolidate their workloads in current cloud because they serve different customers using different configurations (proxy settings, LB strategies, etc.), and there is no way to separate the application logic with the underlying network stack. The operator has to deploy AGs as independent VMs, reserve resources for them, and charge customers accordingly.

NetKernel enables multiplexing across AGs running distinct services, since the common TCP stack processing is now separated into the NSM. Using the three most utilized AGs which have the least benefit from multiplexing as an example, without NetKernel each needs 4 cores in our testbed to handle their peak traffic, and the total per-core requests per second (RPS) of the system is depicted in Figure 14 as Baseline. Then in NetKernel, we deploy 3 VMs each with 1 core to replay the trace as the AGs, and use a kernel stack NSM with 5 cores which is sufficient to handle the aggregate traffic. Totally 9 cores are used including CoreEngine, representing a saving of 3 cores in this case. The per core RPS is thus improved by 33% as shown in Figure 14. Each AG has exactly

TABLE IX NetKernel Multiplexes More AGs and Saves Over 40% Cores

	Total Cores	NSM	CoreEngine	AGs
Baseline	32	0	0	16
NetKernel	32	2	1	29

the same RPS performance without any packet loss. Note that when the baseline is assigned with 3 cores for each VM, the peak RPS performance becomes 75 for AG1 and AG3 in Figure 13 while the per-core RPS performance does not change in Figure 14.

In the general case multiplexing these AGs brings even more gains since their peak traffic is far from their capacity. For ease of exposition we assume the operator reserves 2 cores for each AG. A 32-core machine can host 16 AGs. If we use NetKernel with 1 core for CoreEngine and a 2-core NSM, we find that we can always pack 29 AGs each with 1 core for the application logic as depicted in Table IX, and the maximum utilization of the NSM would be well under 60% in the worst case for ~97% of the AGs in the trace. Thus one machine can run 13 or 81.25% more AGs now, which means the operator can save over 40% cores for supporting this workload. This implies salient financial gains for the operator: According to [22] one physical core has a maximum potential revenue of \$900/yr.

B. Fair Bandwidth Sharing

TCP is designed to achieve flow-level fairness for bandwidth sharing in a network. This leads to poor fairness in a cloud where a misbehaved VM can hog the bandwidth by establishing excessive TCP flows. Distributed congestion control at an entity-level (VM, process, etc.) such as Seawall [59] has been proposed and implemented in a non-virtualized setting. Yet using Seawall in a cloud has many difficulties: The provider has to implement it on the vSwitch or hypervisor and make it work for various guest OSes. The interaction with the VM's own congestion control logic makes it even harder [29]. NetKernel enforces VM-level fair bandwidth sharing by implementing schemes like Seawall in a new NSM. As all the flows from the VMs go through this NSM, the operator can fully control and manage them. Therefore, NetKernel naturally facilitates implementation of such fair sharing.

To prove NetKernel's efficiency, we run a simple VM-level congestion control setting in the NSM: One VM maintains a global congestion window shared among all its connections to different destinations. Each individual flow's ACK advances the shared congestion window, and when sending data, each flow cannot send more than 1/n of the shared window where n is the number of active flows. We then run experiments with 2 VMs: A well-behaved VM that has 8 active flows, and a selfish VM that uses varying number of active flows. As shown in Figure 15, NetKernel with the VM-level congestion control NSM is able to realize an equal share of bandwidth between two VMs regardless of number of flows. We leave the implementation of a complete general solution such as Seawall in NetKernel as future work.

C. Deploying mTCP Without API Change

We now focus on use cases of deployment and performance benefits for users.



Fig. 15. By sharing a unified congestion window to same destination, an NSM can achieve VM-level fairness.

TABLE X

PERFORMANCE OF NGINX USING ab WITH 64B HTML FILES, A CONCURRENCY OF 100, AND 10M REQUESTS IN TOTAL. THE NSM AND VM USE THE SAME NUMBER OF VCPUS

# vCPUs	1	2	4
Kernel stack NSM	71.9K	133.6K	200.1K
mTCP NSM	98.1K	183.6K	379.2K

Most userspace stacks use their own APIs and require applications to be ported [4], [9], [32]. For example, in mTCP an application has to use mtcp_epoll_wait() to fetch events [32]. The semantics of these APIs are also different from socket APIs [32]. These factors lead to expensive code changes and make it difficult to use the stack in practice. The lack of modern APIs also makes it difficult to support complex web servers like nginx. mTCP also lacks some modern kernel TCP features such as advanced loss recovery, small queue, DSACK, *etc.*

With NetKernel, applications can directly take advantage of userspace stacks without any code change. To show this, we deploy unmodified nginx in the VM with the mTCP NSM we implement, and benchmark its performance using ab. Both VM and NSM use the same number of vCPUs. Table X depicts that mTCP provides 1.4x–1.9x improvements over the kernel stack NSM across various vCPU setting.

NetKernel also mitigates the maintenance efforts required from users. We provide another piece of evidence with mTCP here. When compiling DPDK required by mTCP on our testbed, we could not set the RSS (receive side scaling) key properly to the mlx5_core driver for our NIC and mTCP performance was very low. After discussing with mTCP developers, we were able to attribute this to the asymmetric RSS key used in the NIC, and fixed the problem by modifying the code in the DPDK mlx5 driver. We have submitted our fix to mTCP community. Without NetKernel users would have to deal with such technical complication by themselves. Now they are taken care of transparently, saving much time and effort for many users.

D. Shared Memory Networking

Inter-VM communication is well-known to suffer from high overheads [57]. A VM's traffic goes through its network stack, then the vNIC and the vSwitch, even when the other VM is on the same host. It is difficult for users and operator



Fig. 16. Using shared memory NSM for NetKernel for traffic between two colocated VMs of the same user. NetKernel uses 2 cores for each VM, 2 cores for the NSM, and 1 core for CoreEngine. Baseline uses 2 core for the sending VM, 5 cores for receiving VM, and runs TCP Cubic. Both schemes use 8 TCP connections.

to optimize for this case, because a VM has no information about where the other endpoint is. The hypervisor cannot help either as the data has already been processed by the TCP/IP stack. With NetKernel the NSM is part of the infrastructure, the operator can easily detect the on-host traffic and use shared memory to copy data for the two VMs. We build a prototype NSM to demonstrate this idea: When a socket pair is detected as an internal socket pair by the GuestLib, and the two VMs belong to the same user, a shared memory NSM takes over their traffic. This NSM simply copies the message chunks between their hugepages and bypasses the TCP stack processing. As shown in Figure 16, with 7 cores in total, NetKernel with shared memory NSM can achieve 100Gbps, which is \sim 2x of Baseline using TCP Cubic and same number of cores.

VIII. DISCUSSION

How can I do netfilter now? Due to the removal of vNIC and redirection from the VM's own TCP stack, some networking tools like netfilter are affected. Though our current design does not address them, they may be supported by adding additional callback functions to the network stack in the NSM. When the NSM serves multiple VMs, it then becomes challenging to apply netfilter just for packets of a specific VM. We argue that this is acceptable since most users wish to focus on their applications instead of tuning a network stack. NetKernel does not aim to completely replace the current architecture. Tenants may still use the VMs without NetKernel if they wish to gain maximum flexibility on the network stack implementation.

What about troubleshooting performance issues? In current virtualized environment, operators cannot easily determine whether a performance issue is caused by the guest network stack or the underlying infrastructure. With NetKernel operators gain much visibility of the guest network stack, which potentially facilitates debugging the performance issues. For example operators can closely monitor their NSMs to detect problems with the network stack; they can also deploy additional mechanisms in the NSMs to monitor their datacenter network [27], [47], all without disrupting users at all.

Does NetKernel increase the attack surface? It is well-known that shared memory design might suffer from

side-channel attacks where malicious tenants could temper with other tenants' data on the hugepages. In this regard, NetKernel limits the visibility of NK devices into the hugepage for guest VMs: Each device can only access its own address space, which is guaranteed by enforcing the address allocation and isolation control at CoreEngine. Therefore, NetKernel keeps the same attack surface as the virtual NIC and virtual switch abstractions, meaning that the attack consequence is the same as before without NetKernel. It is also possible to exploit hardware for memory copying: guest VM and NSM has dedicated memory regions, in which data can be copied by hardware DMA engines (*e.g.*, FPGA and Intel QuickData), thus avoiding side-channel attacks.

How about supporting stacks with non-socket API? There are many fast network stacks with non-socket API such as PASTE [30], Seastar [9], and IX [16]. As NetKernel keeps the socket API, the central challenge to support these stacks (as NSMs) is how to resolve the semantic differences. While this requires case-by-case porting efforts, in general the ServiceLib should take care of the semantic transformation between the APIs.

Future directions. We outline a few future directions that require immediate attention with high potential: (1) Isolation. When multiple guest VMs share the same NSM, fine-grained performance isolation [37] is imperative. It is also necessary and interesting to design charging policies that promote fair use of the NSM and CoreEngine. For example, tenants may be charged according to how many requests are processed. In addition, as multiple guest VMs may share the same NSM, complete resource isolation becomes crucial for multi-tenant scenarios, which NetKernel considers as future work. (2) Resource efficiency. Various aspects of NetKernel's design can be optimized for efficiency and practicality. The CPU overhead of CoreEngine, mostly to poll the shared memory queues for NQE transmission, can be optimized by offloading to hardware like FPGA and SoC. For example, a FPGA-based CoreEngine can interact with NK devices via a DMA engine, and it only transmits the NQEs so that the DMA bus does not become a performance bottleneck. Besides, such offloading saves precious CPU cores that could benefit user applications. (3) API. NetKernel currently provides BSD socket semantics, but a guest VM might need more network functionalities, such as access control and firewall. Therefore, identifying the minimal API set for guests to efficiently interact with the network stack becomes another future direction. Additionally, applications based on netmap [56], RDMA and eRPC [35] libraries cannot benefit from NetKernel at present. NetKernel aims at supporting these network libraries in the future.

IX. RELATED WORK

We discuss related work besides those mentioned in §2.2.

There are many novel network stack designs to improve performance. The kernel stack continues to receive optimization in various aspects [41], [51], [64]. Userspace stacks based on fast packet I/O are also gaining momentum [9], [32], [39], [44], [53]. Beyond transport layer, novel flow scheduling [14] and end-host based load balancing schemes [28], [36] are developed to reduce flow completion times. These proposals are targeting specific problems of the stack, and can be potentially deployed as NSMs in NetKernel. This paper takes on a broader and more fundamental issue: How can we properly re-factor the network stack, so that new designs can be easily deployed, and operating them in cloud can be more efficient?

Snap [46] is a microkernel networking framework that implements a range of network functions in userspace motivated by the need of rapid development and high performance packet processing in a private cloud. As NetKernel's design space and design choice are significantly different, it achieves many advantages that Snap does not target, such as multiplexing, porting a network stack across OSes or from kernel to user space, enforcing different network stack for different VMs, etc.

Lastly, our earlier papers [49], [50] present the vision of NetKernel and part of the work. In this paper we extend the conference version [49] with a bandwidth sharing use case and more benchmarks to demonstrate the performance and scalability.

X. CONCLUSION

We have presented NetKernel, a system that decouples the network stack from the guest, therefore making it part of the virtualized infrastructure in the cloud. NetKernel improves network management efficiency for operator, and provides deployment and performance gains for users. We experimentally demonstrated new use cases enabled by NetKernel that are otherwise difficult to realize in the current architecture. Through testbed evaluation with 100G NICs, we showed that NetKernel achieves the same performance and isolation as today's cloud.

We focused on efficiency benefits of NetKernel in this paper since they seem most immediate. The idea of separating network stack from the guest VM applies to public and private clouds as well, and brings additional benefits that are more farreaching. For example, it facilitates innovation by allowing new protocols in different layers of the stack to be rapidly prototyped and experimented. It provides a direct path for enforcing centralized control, so network functions like failure detection [27] and monitoring [38], [47] can be integrated into the network stack implementation. It opens up new design space to more freely exploit end-point coordination [23], [52], software-hardware co-design, and programmable data planes [13], [42]. We encourage the community to explore these opportunities in the future.

REFERENCES

- Amazon EC2 Container Service. Accessed: Jan. 10, 2020. [Online]. Available: https://aws.amazon.com/ecs/details/
- [2] Azure Container Service. Accessed: Jan. 10, 2020. [Online]. Available: https://azure.microsoft.com/en-us/pricing/details/container-service/
- [3] Busy Polling: Past, Present, Future. Accessed: Jan. 10, 2020. [Online]. Available: https://netdevconf.info/2.1/papers/BusyPollingNextGen.pdf
- [4] F-Stack: High Performance Userspace Stack Based on FreeBSD Stable. Accessed: Jan. 10, 2020. [Online]. Available: http://www.f-stack.org/
- [5] Google Container Engine. Accessed: Jan. 10, 2020. [Online]. Available: https://cloud.google.com/container-engine/pricing
- [6] Mellanox Smart Network Adaptors. Accessed: Jan. 10, 2020. [Online]. Available: http://www.mellanox.com/page/programmable_network_ adapters?mtag=programmable_adapter_cards

- [7] Netronome. Accessed: Jan. 10, 2020. [Online]. Available: https://www.netronome.com/
- [8] Open Source Kernel Enhancements for Low Latency Sockets Using Busy Poll. Accessed: Jan. 10, 2020. [Online]. Available: http://caxapa.ru/ thumbs/793343/Open_Source_Kernel_Enhancements_for_Low-.pdf
- [9] Seastar. Accessed: Jan. 10, 2020. [Online]. Available: http://www.seastar-project.org/
- [10] (Mar. 2018). mTCP. [Online]. Available: https://github.com/eunyoung14/ mtcp/tree/2385bf3a0e47428fa21e87e341480b6f232985bd
- [11] The Open Group Base Specifications Issue 7, 2018 Edition, IEEE Standard 1003.1-2017. 2018. [Online]. Available: http://pubs.opengroup. org/onlinepubs/9699919799/functions/contents.html
- [12] M. Alizadeh et al., "Data center TCP (DCTCP)," in Proc. ACM SIGCOMM Conf. (SIGCOMM), 2010, pp. 63-74.
- [13] M. T. Arashloo, M. Ghobadi, J. Rexford, and D. Walker, "HotCocoa: Hardware congestion control abstractions," in *Proc. 16th ACM Workshop Hot Topics Netw.*, Nov. 2017, pp. 108–114.
- [14] W. Bai, L. Chen, K. Chen, D. Han, C. Tian, and W. Sun, "PIAS: Practical information-agnostic flow scheduling for data center networks," in *Proc. 13th ACM Workshop Hot Topics Netw.*, Oct. 2014, pp. 1–7.
- [15] H. Ballani, P. Costa, T. Karagiannis, and A. Rowstron, "Towards predictable datacenter networks," in *Proc. ACM SIGCOMM Conf.* (*SIGCOMM*), 2011, pp. 242–253.
- [16] A. Belay, G. Prekas, A. Klimovic, S. Grossman, C. Kozyrakis, and E. Bugnion, "IX: A protected dataplane operating system for high throughput and low latency," in *Proc. USENIX OSDI*, 2014, pp. 49–65.
- [17] N. Cardwell, Y. Cheng, C. S. Gunn, S. H. Yeganeh, and V. Jacobson, "BBR: Congestion-based congestion control," *Commun. ACM*, vol. 60, no. 2, pp. 58–66, 2017.
- [18] B. Cronkite-Ratcliff *et al.*, "Virtualized congestion control," in *Proc.* ACM SIGCOMM Conf., Aug. 2016, pp. 230–243.
- [19] M. Dalton *et al.*, "Andromeda: Performance, isolation, and velocity at scale in cloud network virtualization," in *Proc. USENIX NSDI*, 2018, pp. 373–387.
- [20] D. R. Engler, M. F. Kaashoek, and J. O'Toole, "Exokernel: An operating system architecture for application-level resource management," in *Proc.* 15th ACM Symp. Operating Syst. Princ. (SOSP), 1995, pp. 251–266.
- [21] D. Firestone, "VFP: A virtual switch platform for host SDN in the public cloud," in *Proc. NSDI*, 2017, pp. 315–328.
- [22] D. Firestone *et al.*, "Azure accelerated networking: SmartNICs in the public cloud," in *Proc. USENIX NSDI*, 2018, pp. 51–66.
- [23] P. X. Gao, A. Narayan, G. Kumar, R. Agarwal, S. Ratnasamy, and S. Shenker, "PHost: Distributed near-optimal datacenter transport over commodity network fabric," in *Proc. 11th ACM Conf. Emerg. Netw. Experiments Technol.*, Dec. 2015, pp. 1–12.
- [24] H. Golestani, A. Mirhosseini, and T. F. Wenisch, "Software data planes: You can't always spin to win," in *Proc. ACM Symp. Cloud Comput.*, Nov. 2019, pp. 337–350.
- [25] D. B. Golub *et al.*, "Microkernel operating system architecture and Mach," in *Proc. USENIX Workshop Micro-Kernels Other Kernel Architectures*, 1992, pp. 11–30.
- [26] A. Greenberg, "SDN in the cloud," in *Proc. ACM SIGCOMM*, 2015. Accessed: Jan. 10, 2020. [Online]. Available: https://www.cl.cam.ac.uk/ teaching/1718/R02/papers/keynote.pdf
- [27] C. Guo et al., "Pingmesh: A large-scale system for data center network latency measurement and analysis," in *Proc. ACM Conf. Special Interest Group Data Commun.*, Aug. 2015, pp. 139–152.
- [28] K. He, E. Rozner, K. Agarwal, W. Felter, J. Carter, and A. Akella, "Presto: Edge-based load balancing for fast datacenter networks," in *Proc. ACM Conf. Special Interest Group Data Commun.*, Aug. 2015, pp. 465–478.
- [29] K. He *et al.*, "AC/DC TCP: Virtual congestion control enforcement for datacenter networks," in *Proc. ACM SIGCOMM Conf.*, Aug. 2016, pp. 244–257.
- [30] M. Honda, G. Lettieri, L. Eggert, and D. Santry, "PASTE: A network programming interface for non-volatile main memory," in *Proc. USENIX NSDI*, 2018, pp. 17–33.
- [31] J. Hwang, K. K. Ramakrishnan, and T. Wood, "NetVM: High performance and flexible networking using virtualization on commodity platforms," in *Proc. USENIX NSDI*, 2014, pp. 445–458.
- [32] E. Jeong et al., "MTCP: A highly scalable user-level TCP stack for multicore systems," in Proc. USENIX NSDI, 2014, pp. 489–502.
- [33] V. Jeyakumar, M. Alizadeh, D. Mazieres, B. Prabhakar, C. Kim, and A. Greenberg, "Eyeq: Practical network performance isolation at the edge," in *Proc. USENIX NSDI*, 2013, pp. 297–311.

- [34] G. Judd, "Attaining the promise and avoiding the pitfalls of TCP in the datacenter," in *Proc. USENIX NSDI*, 2015, pp. 145–157.
- [35] A. Kalia, M. Kaminsky, and D. Andersen, "Datacenter RPCs can be general and fast," in *Proc. USENIX NSDI*, 2019, pp. 1–16.
- [36] N. Katta, M. Hira, A. Ghag, C. Kim, I. Keslassy, and J. Rexford, "CLOVE: How i learned to stop worrying about the core and love the edge," in *Proc. 15th ACM Workshop Hot Topics Netw.*, Nov. 2016, pp. 155–161.
- [37] J. Khalid *et al.*, "Iron: Isolating network-based CPU in container environments," in *Proc. USENIX NSDI*, 2018, pp. 313–328.
- [38] A. Khandelwal, R. Agarwal, and I. Stoica, "Confluo: Distributed monitoring and diagnosis stack for high-speed networks," in *Proc. USENIX NSDI*, 2019, pp. 421–436.
- [39] D. Kim et al., "FreeFlow: Software-based virtual RDMA networking for containerized clouds," in Proc. USENIX NSDI, 2019, pp. 113–126.
- [40] K. LaCurts, J. C. Mogul, H. Balakrishnan, and Y. Turner, "Cicada: Introducing predictive guarantees for cloud networks," in *Proc. USENIX HotCloud*, 2014, pp. 1–6.
- [41] X. Lin et al., "Scalable kernel TCP design and implementation for short-lived connections," in Proc. 21st Int. Conf. Architectural Support Program. Lang. Operating Syst., Mar. 2016, pp. 339–352.
- [42] Y. Lu et al., "Multi-path transport for RDMA in datacenters," in Proc. USENIX NSDI, 2018, pp. 357–371.
- [43] A. Madhavapeddy et al., "Unikernels: Library operating systems for the cloud," in Proc. 18th Int. Conf. Architectural Support Program. Lang. Operating Syst. (ASPLOS), 2013, pp. 461–472.
- [44] I. Marinos, R. N. M. Watson, and M. Handley, "Network stack specialization for performance," in *Proc. ACM Conf. SIGCOMM*, Aug. 2014, pp. 175–186.
- [45] J. Martins *et al.*, "ClickOS and the art of network function virtualization," in *Proc. USENIX NSDI*, 2014, pp. 459–473.
- [46] M. Marty et al., "Snap: A microkernel approach to host networking," in Proc. 27th ACM Symp. Operating Syst. Princ., Oct. 2019, pp. 399–413.
- [47] M. Moshref, M. Yu, R. Govindan, and A. Vahdat, "Trumpet: Timely and precise triggers in data centers," in *Proc. ACM SIGCOMM Conf.*, Aug. 2016, pp. 129–143.
- [48] A. Narayan et al., "Restructuring endpoint congestion control," in Proc. Appl. Netw. Res. Workshop, Jul. 2018, pp. 30–43.
- [49] Z. Niu et al., "NetKernel: Making network stack part of the virtualized infrastructure," in Proc. USENIX ATC, 2020, pp. 143–157.
- [50] Z. Niu, H. Xu, D. Han, P. Wang, and L. Liu, "Netkernel: Network stack as a service in the cloud," in *Proc. ACM HotNets*, 2017, pp. 65–71.
- [51] S. Pathak and V. S. Pai, "ModNet: A modular approach to network stack extension," in *Proc. USENIX NSDI*, 2015, pp. 425–438.
- [52] J. Perry, A. Ousterhout, H. Balakrishnan, D. Shah, and H. Fugal, "Fastpass: A centralized 'zero-queue' datacenter network," in *Proc. ACM SIGCOMM*, 2014, pp. 307–318.
- [53] S. Peter *et al.*, "Arrakis: The operating system is the control plane," in *Proc. USENIX OSDI*, 2014, pp. 1–16.
- [54] L. Popa, A. Krishnamurthy, S. Ratnasamy, and I. Stoica, "FairCloud: Sharing the network in cloud computing," in *Proc. 10th ACM Workshop Hot Topics Netw. HotNets*, 2011, pp. 187–198.
- [55] L. Popa, P. Yalagandula, S. Banerjee, J. C. Mogul, Y. Turner, and J. R. Santos, "ElasticSwitch: Practical work-conserving bandwidth guarantees for cloud computing," in *Proc. ACM SIGCOMM Conf.* (*SIGCOMM*), Aug. 2013, pp. 351–362.
- [56] L. Rizzo, "netmap: A novel framework for fast packet I/O," in Proc. USENIX ATC, 2012, pp. 101–112.
- [57] L. Rizzo, G. Lettieri, and V. Maffione, "Speeding up packet I/O in virtual machines," in *Proc. Architectures Netw. Commun. Syst.*, 2013, pp. 47–58.
- [58] A. Saeed, N. Dukkipati, V. Valancius, V. T. Lam, C. Contavalli, and A. Vahdat, "Carousel: Scalable traffic shaping at end hosts," in *Proc. Conf. ACM Special Interest Group Data Commun.*, Aug. 2017, pp. 404–417.
- [59] A. Shieh, S. Kandula, A. Greenberg, C. Kim, and B. Saha, "Sharing the data center network," in *Proc. USENIX NSDI*, 2011, pp. 23–36.
- [60] L. Soares and M. Stumm, "FlexSC: Flexible system call scheduling with exception-less system calls," in *Proc. USENIX OSDI*, 2010, pp. 33–46.
- [61] J. Son, Y. Xiong, K. Tan, P. Wang, Z. Gan, and S. Moon, "Protego: Cloud-scale multitenant IPsec gateway," in *Proc. USENIX ATC*, 2017, pp. 473–485.
- [62] B. Stephens, A. Akella, and M. Swift, "Loom: Flexible and efficient NIC packet scheduling," in *Proc. USENIX NSDI*, 2019, pp. 33–46.

- [63] B. Stephens, A. Singhvi, A. Akella, and M. Swift, "Titan: Fair packet scheduling for commodity multiqueue NICs," in *Proc. USENIX ATC*, 2017, pp. 431–444.
- [64] K. Yasukata, M. Honda, D. Santry, and L. Eggert, "StackMap: Low-latency networking with the OS stack and dedicated NICs," in *Proc. USENIX ATC*, 2016, pp. 43–56.
- [65] D. Zhuo et al., "Slim: OS kernel support for a low-overhead container overlay network," in Proc. USENIX NSDI, 2019, pp. 331–344.



Zhixiong Niu received the B.E. degree in network engineering from Dalian Maritime University (DMU) in 2012, the M.Sc. degree in computer science from The University of Hong Kong (HKU) in 2014, and the Ph.D. degree from the Department of Computer Science, City University of Hong Kong, in 2019. He is a Senior Researcher at Microsoft Research Asia.



Qiang Su received the B.S. degree in computer science and technology from Northeastern University, China, in 2018. He is currently pursuing the Ph.D. degree in computer science with the City University of Hong Kong, China. His research interests include systems and networking.



Peng Cheng received the B.S. degree in software engineering from Beihang University in 2010 and the Ph.D. degree in computer science and technology from Tsinghua University in 2015. He was a visiting Ph.D. student with UCLA from 2013 to 2014. He is a Principal Researcher at Microsoft Research Asia. His research interests are computer networking and networked systems.



Dongsu Han (Member, IEEE) received the B.S. degree in computer science from the Korea Advanced Institute of Science and Technology (KAIST) in 2003 and the Ph.D. degree in computer science from Carnegie Mellon University in 2012. He is currently an Associate Professor with the School of Electrical Engineering and the Graduate School of Information Security, KAIST. His research interests include networking, distributed systems, and networks/systems security (more details about his research can be found at http://ina.kaist.ac.kr).

Keith Winstein received the B.S., M.S., and Ph.D. degrees in computer science from the Massachusetts Institute of Technology in 2004, 2005, and 2014, respectively. He is an Assistant Professor of computer science and, by courtesy, of electrical engineering at Stanford University. His research interest includes networked systems.





Yongqiang Xiong received the B.S., M.S., and Ph.D. degrees from Tsinghua University, Beijing, China, in 1996, 1998, and 2001, respectively, all in computer science. He is currently a Principal Researcher and the Research Manager at Microsoft Research Asia and leads the Networking Research Group. His research interests include systems and networking, as well as network security.



Hong Xu (Senior Member, IEEE) received the B.Eng. degree from The Chinese University of Hong Kong in 2007 and the M.A.Sc. and Ph.D. degrees from the University of Toronto in 2009 and 2013, respectively. From 2013 to 2020, he was with the City University of Hong Kong. He is currently an Associate Professor with the Department of Computer Science and Engineering, The Chinese University of Hong Kong. His research areas are computer networking and systems, particularly big data systems and data center networks. He is a

Senior Member of ACM. He was a recipient of an Early Career Scheme Grant from the Hong Kong Research Grants Council in 2014. He received three best paper awards, including the IEEE ICNP 2015 Best Paper Award.