Galleon: Reshaping the Square Peg of NFV

Jianfeng Wang
University of Southern California

Tamás Lévai
Budapest University of Technology and Economics

Zhuojin Li
University of Southern California

Marcos A. M. Vieira
Universidade Federal de Minas Gerais

Ramesh Govindan
University of Southern California

Barath Raghavan
University of Southern California

Abstract

Software is often used for Network Functions (NFs)—such as firewalls, NAT, deep packet inspection, and encryption—that are applied to traffic in the network. The community has hoped that NFV would enable rapid development of new NFs and leverage commodity computing infrastructure. However, the challenge for researchers and operators has been to align the square peg of high-speed packet processing with the round hole of cloud computing infrastructures and abstractions, all while delivering performance, scalability, and isolation.

Past work has led to the belief that NFV requires custom approaches that deviate from today's abstractions. To the contrary, we show that we can achieve performance, scalability, and isolation in NFV judiciously using mechanisms and abstractions of FaaS, the Linux kernel, NIC hardware, and OpenFlow switches. As such, with our system Galleon, NFV can be practically-deployable today in conventional cloud environments while delivering up to double the performance per core compared to the state of the art.

1 Introduction

Software is often used for Network Functions (NFs)—such as firewalls, NAT, deep packet inspection, and encryption—that are applied to traffic in the network. For the last eight years, the hope has been that NF Virtualization (NFV) would enable rapid development of new NFs by multiple software vendors and leverage the power and economics of commodity computing infrastructure. The challenge for the systems community and for network operators has been to align the square peg of high-speed packet processing with the round hole of cloud computing infrastructures and abstractions.

To this end, many NFV solutions have been proposed to provide performance, scalability, and isolation, each with novel mechanisms and inherent tradeoffs. Some have explored how to integrate hardware and software to achieve high performance and scalability [9, 13, 34] but ignore NF isolation; others have solved isolation but require programming language support [22]; still others have used appropriate cloud abstractions like FaaS but have not addressed scalability [27].

The research community has believed that NFV is different enough that it requires novel, custom approaches that deviate from today's norms. To the contrary, we show that NFV may be a round peg after all—that we can achieve performance, scalability, and isolation using the standard mechanisms and abstractions of today's cloud computing environments—building upon Function-as-a-Service (FaaS), the Linux kernel, standard NIC hardware, and OpenFlow switches. As such, with Galleon, NFV can be deployable today in conventional cloud environments.

Performance, scalability, and isolation are key for NFV to be production ready. NFV workloads involve the deployment of chains of one or more NFs from different vendors, in a multi-tenant environment; for deployment and management ease, NFV needs to conform to the infrastructure rather than the other way around. Offloading some NF processing to specialized hardware, for example, entails operational complexity. NF vendors implement NFs as they see fit and make them available as containers or VMs, giving operators little language or hardware choice. Indeed, customization of the infrastructure undercuts a key focus of NFV: to eliminate the operational headaches of the hardware-based middlebox era.

NFV, as an important workload for the edge, is often involved in rack-scale deployments at telco central offices or ISP PoPs. Further, with limited commodity hardware, such a platform often puts an efficiency requirement for scaling up and down both NFV chains and other edge applications quickly. Further, it prefers not to split the infrastructure into smaller subsets that serve one type of dedicated applications within each subset. This not only increases the management overhead of the platform, but also prevents it from multiplexing the infrastructure among all deployed applications, which further reduces the system efficiency.

We posit that NFV must embrace a dirty-slate approach. In this paper, we ask: is it possible to achieve performance, scalability, and isolation in NFV on general-purpose infrastructure? We believe that three key enabling pieces are missing from today's NFV literature: first, the ability of deploying third-party NFs with an efficient isolation mechanism without

losing generality; second, the ability to run NF chains with an efficient yet standard scaling mechanism; third, the ability to execute NF chains that coexist with other edge applications while meeting SLO targets.

This paper makes the following contributions:

FaaS compatibility. Galleon minimally extends FaaS abstractions to enable them to support NFV workloads. As in existing FaaS implementations, each NF is a separate container. NFs are, however, triggered on packet arrival, not on application level request arrival. Galleon builds upon OpenFaaS [20] and leverages Kubernetes so that this optimized FaaS infrastructure maintains its compatibility with other applications while meeting the needs of NFV.

Performance-aware scheduling. Galleon dedicates cores to NF chains and uses kernel bypass to deliver packets to NFs. It uses standard OS interfaces to cooperatively schedule NFs in a chain, to mimic the *run-to-completion* [9] strategy that has proven to be essential for high NFV performance. Run-to-completion processes a batch of packets; Galleon selects batch sizes that satisfy SLOs while minimizing context switch overhead.

Efficient scaling with SLO adherence. In response to changes in traffic, Galleon auto-scales NF chain instances in a manner that minimizes core usage while preserving latency SLOs. This flexibility allows tenants to trade-off latency for lower cost, a capability present in other NFV systems [31].

High-performance spatiotemporal packet isolation. Galleon's use of containerized NFs, together with NIC virtualization, ensures that an NF chain can only see its own traffic. Galleon also ensures a stronger form of packet isolation: an NF in a chain can process a packet only after its predecessor NF. It does this by spatially isolating the first NF from the others in the chain using a packet copy. Subsequent NFs can process packets in zero-copy fashion, with temporal isolation enforced by scheduling. This approach is general and transparent to NF implementations, and requires neither language support nor high-overhead mechanisms.

We find Galleon achieves up to $1.93\times$ the per-core throughput of state-of-the-art NFV systems [22, 24] that use alternative isolation mechanisms. Under highly dynamic loads, Galleon achieves zero packet losses and is able to satisfy tail latency SLOs. Compared to a highly-optimized NFV system that does not provide packet isolation and is not designed to satisfy latency SLOs (but is designed to minimize latency) [9], Galleon uses slightly more CPU cores (15%) while achieving isolation *and* satisfying latency SLOs.

2 Background

2.1 NFV's promise

NFV's promise was to replace bulky hardware middleboxes with virtualized, easier-to-manage software-based NFs. NF functionality varies widely, from simple (*e.g.,* VLAN tun-

neling) to complex (*e.g.,* traffic inference). NFs are chained together to process packets in sequence to meet a network operator's needs; a canonical chain employs a firewall, deep packet inspection, and encrypted tunneling.

NFV workloads have demanding performance requirements, as they must forward packets at hundreds of Gbps while meeting stringent per-chain latency SLOs. It is also important to achieve high resource efficiency for deploying such workloads so that the NFV service providers can serve traffic with less resources, and further reduce the cost of NFV.

2.2 Galleon's Goal

While NFV initially promised to adopt cloud computing abstractions, the move from specialized hardware to commodity servers running canonical software stacks has been in name only. Indeed, the *state of the art in NFV employs custom interfaces, run-times and control planes* [9, 11, 21, 22, 24], breaks abstraction boundaries in the name of performance [9, 21], places a greater burden upon the programmer to ensure key properties such as isolation [22], and leverages specialized hardware to achieve high performance [34]. These *customized NFV platforms* have been absolutely essential for the intellectual advancement of the area, but they are insufficient to enable general, widespread NFV offerings in practice.

Our goal with Galleon is to align NFV with ubiquitous cloud abstractions without sacrificing performance, scalability, or isolation. In particular, our aim is to build upon the abstractions and mechanisms available with Function-as-a-Service (FaaS) platforms. FaaS closely aligns with NFV's needs: to execute a modular piece of code (*e.g.,* an NF) over discrete units of data (*e.g.,* packets). Recent prior work ServerlessNF [27, 28] aims to achieve these goals on unmodified FaaS platforms. In contrast, Galleon proposes minimal additional abstractions for packet processing that enable it to provide performance and feature sets comparable to customized NF processing platforms (*e.g.,* Metron [9]).

FaaS is a good candidate for deploying NFV. NF developers can focus on developing new NFs and leave the orchestration and management overhead to FaaS providers. With FaaS's fine-grained billing, NF developers can also benefit from a cost-efficient deployment when serving dynamic traffic. Developers write and upload either a piece of code or an executable, which is then executed in response to incoming requests. FaaS providers hide the underlying infrastructure from developers, and handle provisioning, executing, scheduling, scaling, and resource accounting of user-defined code. FaaS comes with built-in scalability, and pursues extreme efficiency such that deployments do not pay for resources that they do not use.

While these benefits sound plausible, today's FaaS platforms cannot serve NFV workloads with demanding performance requirements. In fact, most existing FaaS platforms do not offer any performance guarantees for underlying applications. They employ non-deterministic scheduling on top of

shared resources. This makes them mostly useful for latency-insensitive applications. While executing modular functions, FaaS platforms lack an efficient inter-function communication mechanism and often rely on third-party cloud storage services for inter-function communications. This introduces new sources of latency, and can be extremely cost-inefficient when almost all packets must be transmitted among NFs.

Galleon's aims to provide a FaaS platform suitable for deploying NF chains in a multi-tenant environment. As such, it must satisfy four requirements. The first of these is unique to Galleon, while the remaining requirements ensure that Galleon is comparable to custom NFV platforms.

Abstractions and Interfaces. Current FaaS systems provide REST-ful interfaces inappropriate for packet processing [27]. Galleon must 1) introduce minimal new abstractions for NFV while maintaining performance comparable to standalone platforms like Metron [9], NetBricks [22], or NFVnice [11], and 2) must do so without impacting the core FaaS abstractions that permit application-level functions. In developing these abstractions, the following challenges arise: How do NF developers interact with the system to develop new NFs? What information must be collected via this interface to specify the requirement of deploying NF chains? How does this interface co-exist with existing cloud interfaces?

Performance-aware scheduling. Each incoming packet triggers the execution of an NF chain. Thus, ingress traffic needs to be distributed to server CPU cores that are tasked with executing the target NF chain. There must be mechanisms for 1) routing packets so that they are processed by the right chain, 2) scheduling NFs to process incoming packets on underlying hardware resources (i.e., CPU cores), and 3) meeting per-chain SLO requirements. Prior approaches break abstraction boundaries and assume programmer involvement in performance tuning. Galleon cannot do this; fortunately, however, it can use techniques (*e.g.,* kernel bypass, virtualized NICs) developed for low-latency cloud communication.

Dynamic scaling. Dynamic scaling needs to quickly adapt to load changes and provision server resources. Such scaling is already built into FaaS offerings, but Galleon needs to define fine-grained low-overhead monitoring techniques, as well as policies for scaling NF chains. Unlike NFV platforms, Galleon must minimize core usage while scaling (so cluster resources can be used by other non-NFV FaaS workloads).

Isolation. Many NFV platforms provide *packet isolation* that Galleon must also support: an NF should not be able to access a packet until its predecessor NF in the chain has finished processing the packet, and NFs in an NF chain should not be able to see packets destined to another chain. This requirement arises in multi-tenant settings where each chain may consist of NFs from multiple vendors, and each chain may be responsible for processing a specific customer's traffic.

In §3, §4.3, and §5, we show how Galleon achieves these.

3 Galleon Overview

Figure 1 illustrates the key Galleon control-plane components. Galleon reuses the basic FaaS architecture, making (a) minimal modifications to the FaaS interface, and (b) augmenting existing FaaS components with mechanisms, instrumentation, and algorithms to support fast packet processing.

3.1 Galleon Interface

Galleon must provide a general abstraction for writing user-defined NFs. Most FaaS platforms today employ a RESTful API to handle network events, usually via HTTP. They allow developers to write a customized function as the event handler that takes an event struct as input, parsed from a HTTP request's payload. This abstraction is not ideal for processing raw packets. Unlike dedicated requests to the FaaS's ingress server (via a public IP), Galleon must permit each packet to be processed by one or more NFs. To do this without modifying FaaS's programming model, one may, at the ingress, encapsulate each packet within a HTTP request payload. This can introduce significant extra overhead [27].

Galleon modifies the FaaS programming and deployment model minimally to support packet processing. NF developers can accept a raw packet struct (a pointer to the packet) as input. A Galleon operator (*e.g.,* an ISP) can assemble an NF chain using these NFs. The customer then submits to Galleon the NF chain, a traffic filter specification for what is to be processed by the chain, and a per-packet latency SLO.

3.2 Galleon Design

Galleon extends FaaS functionality as shown in Figure 1. It assumes commodity servers and OpenFlow-enabled switches. Components that are unique in Galleon, hidden within the infrastructure, include: the Galleon controller, the Galleon NF runtime, the Galleon ingress, the Galleon scheduler, and a per-worker Galleon agent that acts as a representative of the centralized Galleon controller.

Galleon re-uses the existing FaaS worker subsystem for managing workers and deploying services. The worker subsystem (implemented using Kubernetes in many FaaS implementations [20]) manages the system resources, *i.e.,* a pool of worker machines. Each worker machine executes functions encapsulated in containers; in Galleon, each container hosts an NF. Customers provide container images for each NF in a chain. The centralized Galleon controller manages the deployment of FaaS services by interacting with this worker subsystem to deploy Galleon components (an ingress, per-worker scheduler and agents) prior to startup.

At runtime, the Galleon controller uses the worker subsystem to deploy NFs as containers. It collects NF performance statistics reported from each agent, serves queries from the ingress to make scheduling decisions, and pushes load balancing decisions to the ingress which enforces it by modifying a ToR's flow table. The Galleon ingress is where NFV traffic enters and leaves the system. It works with an Openflow-
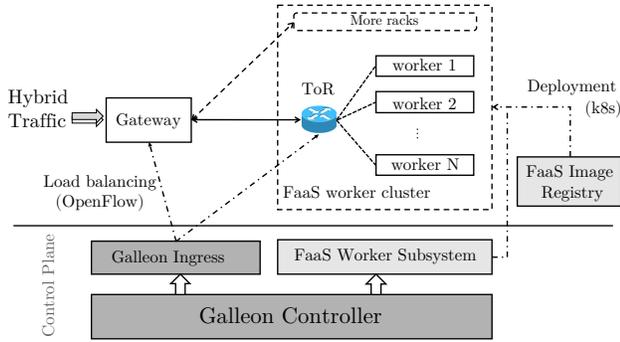
(a) Galleon's control plane.
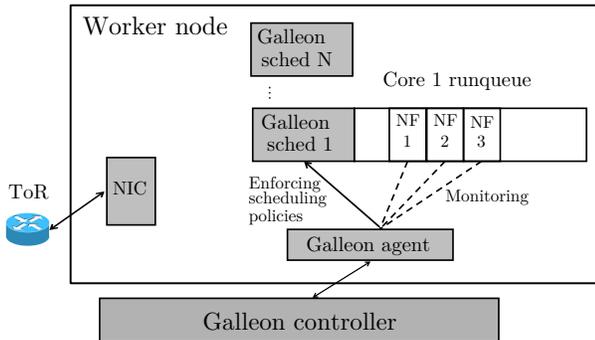
(b) Galleon worker.

Figure 1: Galleon reuses a FaaS worker subsystem to run both system components and function instances. Galleon controller interacts with Galleon ingress and FaaS worker subsystem to deploy NFs in serving packet processing functions. Unshaded boxes are existing FaaS components, lightly shaded ones are FaaS components that Galleon modifies, and darker ones represent new components specific to Galleon hidden within the infrastructure.

enabled switch to enforce the Galleon controller's workload assignment strategies. When a new flow arrives at the switch, packets are forwarded to the Galleon ingress that determines the target logical NF chain for this flow. To do so, it queries the Galleon controller to pick a chain for this flow, then applies the Galleon controller's load balancing for this logical chain to pick a deployed chain. Then it installs a flow rule to offload the flow dispatching task to the switch hardware. This design reduces the traffic sent to the software ingress module, and thus decreases CPU time spent routing packets.

In such a setup, a ToR switch only passes customer-specified aggregates to Galleon's ingress. Traffic that needs application-level processing will reach a normal FaaS ingress service that schedules requests as a normal FaaS system does. In this way, Galleon coexists with normal FaaS processing and does not impact its performance.

4 Galleon's Runtime

Galleon's main challenge is FaaS compatibility. Next we describe the design of Galleon's runtime to achieve performance, scaling, and isolation through careful allocation of NIC and CPU resources, and by careful CPU scheduling.

4.1 NF Execution Models

Prior work has explored different NF execution models that dictate how NFs share packet memory, how the runtimes steer packets to NFs, and how they schedule NF execution.

Memory model. A FaaS instance may reside in memory for some time before it times out. During its lifetime, packets and NF-internal state are loaded into memory for processing. For NFs running on the same worker machine, they (1) may share NF state memory, and packet buffer memory (e.g., in Metron [9] and NetBricks [22]), (2) do not share NF state memory, but share packet buffer memory globally (*e.g.,* in E2 [21] and NFVNice [11]), or (3) do not share either NF state memory or packet buffer memory (*e.g.,* in EdgeOS [24]).

Network I/O model. Packets must be sent to a specific NF running on a specific server core. In many NFV platforms, such as E2 [21], NFVNice [11], and EdgeOS [24], a hardware switch (presumably with Openflow) forwards packets to specific worker machines. Once packets arrive at the server's NIC, a virtual switch forwards traffic locally. In a multi-tenant environment, the vSwitch has read and write access to each individual NF's memory space, and copies packets when forwarding them from an upstream NF to its downstream. The vSwitch is a core service and can become the bottleneck for both intra- and inter-machine traffic. To scale it up, a runtime can add CPU cores for vSwitches, but cores can generate revenue, so this strategy is undesirable. On our test machine, a CPU core can only achieve a 6.9 Gbps throughput when forwarding 64-byte packets (or 13.5 Mpps). Consider a chain with 4 NFs running on a server with a 10 Gbps NIC. The aggregate traffic volume can reach 40 Gbps at peak, which requires at least 7 CPU cores to run vSwitches (more if traffic is not evenly distributed across the vSwitches).

An alternative approach is to offload packet switching to the ToR switch and the NIC's internal switch. Both switches coordinate to ensure packets arrive at the target machine's/target process's memory. When a packet hits the ToR, the switch not only forwards the packet to a dedicated machine, but also facilitates intra-machine forwarding via L2 tagging. This approach eliminates the need to run a vSwitch. However, it can only ensure that packets are received by the first NF in a chain. Metron and NetBricks take this approach, but rely on a strong assumption: that all NFs can be compiled and run in a single process. However, many popular NFs are only available in a closed-source form from commercial software developers, and they cannot be compiled with other NFs to form a single binary that runs the NF chain. Even if that were possible, the packet isolation requirement constrains flexibility significantly, since it can only then be achieved using language-level memory isolation (*e.g.,* by using Rust).

CPU scheduling model. Memory and network I/O models also impact CPU resource allocation and scheduling of NFs and NF chains. When NF chains run in a single process (as in Metron and NetBricks), those runtimes can dedicate a core

to an entire chain. When NFs run in separate processes (as in E2 or NFVnice), runtimes must decide whether to allocate one or more cores to a chain, and how to schedule each NF.

4.2 Galleon's Execution Model

To ensure minimal changes to existing FaaS architectures, Galleon chooses an execution model that sits in a different point in the design space. Galleon deploys each NF in an NF chain as a single container. From the perspective of the underlying FaaS platform, an NF chain represents a "function" (a unit of invocation). NFs can share packet buffers (as in Metron or NetBricks), but packet isolation is enforced through OS protection, careful scheduling and packet copying (unlike NetBricks which relies on language level isolation). Galleon uses NIC I/O virtualization and kernel bypass to reduce packet steering overhead. The rest of the section describes some of the details of Galleon's packet I/O and memory models. The next section describes CPU allocation and scheduling.

Packet I/O. Galleon uses DPDK for fast userspace networking to handle packet I/O for NFs. However, because Galleon must co-exist with other FaaS functions that may be executing on the same CPU, it must use userspace networking only for NF chains. To do so, Galleon uses the hardware virtualization technique Single-root Input/Output Virtualization [30] (SR-IOV) to virtualize the NIC hardware.[1] On a worker machine, the Galleon agent (Figure 1) manages the virtualized devices via kernel APIs through the PF. Once the Galleon controller allocates CPU core to an NF chain, the Galleon agent sets up a VF to steer NF-chain traffic to that core.

Memory. In Galleon, a runtime on behalf of an NF chain initializes a file-backed dedicated memory region that holds fixed-size packet structures for incoming packets. It also creates a ring buffer that holds packet descriptors that point to these packet structures. To receive packets from the virtualized NIC, the NF runtime passes this ring buffer to its associated VF so that the NIC hardware can perform DMA directly to the NF runtime's memory.

NF State Management. For stateful NFs, packet processing depends on both the packet itself and the NF's current internal state. Prior work (*e.g.,* statelessNF [8], S6 [33], SNF [28]) has proposed maintaining NF global state in an external store (such as Redis [25]); they adopt various caching schemes to mitigate the latency overhead of pulling state from the external store. In Galleon, we adopt a similar scheme by maintaining per-NF global state remotely and maintaining general NF state in a flow table locally at each NF runtime instance. At the front-end, the programming model provides a set of simple APIs for writing a stateful NF: `update(flow,`

---

[1]SR-IOV is a PCIe specification that allows a PCIe device to appear as many physical devices, providing a hardware method for sharing a hardware resource with isolation. With SR-IOV, NIC hardware generates one Physical Function (PF) that controls the physical device, and many Virtual Functions (VFs) that are lightweight PCIe functions with the necessary hardware access for receiving and transmitting data.
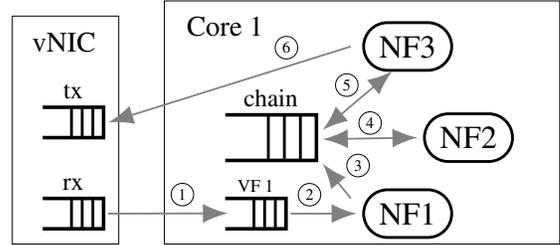
Figure 2: Timeline of packets on a Galleon worker. A packet is tagged at the ingress. 1) NIC's L2 switch sends it to NIC VF associated with the destined chain. NIC VF DMAs packets to the first NF's memory space. 2) NF 1 processes the packet. 3) After NF 1's packet processing function returns, the packet is copied to the chain's pktbuf by the NF runtime if there are other NFs. This is necessary to ensure packet isolation as the NIC's pktbuf should only been seen by NF 1. 3)-5) A per-core cooperative scheduler controls the execution sequence of NFs to ensure temporal packet isolation. 6) Final NF asks VF to send the packet out.

`val)` and `read(flow, val)`, where flow corresponds to a BPF matching rule. The NF runtime makes the process transparent and interacts with the external store to pull state when necessary and synchronize state periodically. Galleon can incorporate prior efforts on optimizing state access latency because the state management model remains similar and the optimizations do not require modified APIs.

4.3 Core Allocation and Scheduling

Userspace I/O and shared memory can reduce overhead significantly, but to be able to process packets at high throughput and low latency, Galleon must have tight control over NF chain execution. As discussed earlier, customized NF platforms use two different approaches. One approach bundles NFs in an NF chain into a single process to *run to completion* (in which each NF in the chain processes a batch of packets before moving onto the next batch (as in Metron or NetBricks). This approach ensures performance predictability and high performance by amortizing overhead over a packet batch. To achieve packet isolation, NetBricks relies on language isolation. The second approach, which NFVNice uses, is to run each NF in a separate process, which ensures isolation with copying packets, but can require careful allocation of CPU shares, and orchestration of process execution on the underlying scheduler (*e.g.,* CFS in NFVNice).

In contrast to these two approaches, we introduce *spatiotemporal packet isolation* in which NF chains operate on 1) spatially isolated packet memory regions (as opposed to the typical model in run-to-completion software switches such as BESS, in which all NF chains on a machine run in the same memory) and 2) are temporally isolated through careful sequencing of their execution, which proceeds in a run-to-completion fashion *across processes* and uses *cooperative scheduling* mechanisms to hand off control and the natural execution boundary of packet batch handoff. This isolation ensures that NF chains (which may process different customers' traffic) cannot see each others packet streams or state,

and even within a chain each NF maintains private state and only gets to execute (and thus access packet memory) when it is expected to perform packet processing in the chain.

In Galleon, NFs are individual containers deployed in a Kubernetes cluster. Galleon dedicates a core to a chain that actively serves traffic. This choice is appropriate for our setting, since we expect that NF chains will handle large traffic aggregates, so that multiple cores might be necessary to process traffic for a given chain (§5). The Galleon controller manages all NFs via Kubernetes APIs to control the allocation of memory, CPU share, and disk space.

4.3.1 Enforcing run-to-completion scheduling

Galleon relies on a per-core cooperative scheduler to schedule NFs on each core. All NF containers that are part of a chain are assigned to a single core. Each NF container runs an NF Runtime in a multi-process way for executing an user-defined NF. The NF process runs in a single-thread for processing traffic. The other process is transparent to NF authors. It runs a RPC server for controlling the NF's behaviour, and a monitoring thread that collects performance statistics.

Threads are transparent to NF authors, and are handled by Galleon completely. NF authors just develop a standard NF with Galleon's interface. Galleon's runtime takes the NF code and executes it in the NF process, while the monitoring thread (§5) collects performance statistics. A Galleon agent in the NF runtime communicates these statistics to the Galleon controller, and handles NF chain instantiation as well as NF-chain scaling (§5). To avoid interfering with packet processing, the monitoring thread runs on a separate core.

To tightly coordinate NF chain execution, Galleon uses Linux's real-time (RT) scheduling support, and manages NF threads real-time priorities and schedules them using a FIFO policy. We use this policy to emulate, as described below, NF chain run-to-completion execution in which each NF in the chain processes a batch of packets in sequence.

Scheduling model. In Galleon's cooperative scheduling, an upstream NF runs in a loop to process individual packets of a given batch, and then yields the core to its downstream NF. This is transparent to the NFs: once the user-defined NF finishes processing, it invokes the NF runtime to transmit the batch to the downstream NF; the runtime invokes yield.[2]

For this, the cooperative scheduler has to bypass the underlying scheduler (CFS in our implementation) and take full control of a core. Internally, the scheduler maintains two FIFO queues: a run queue that contains runnable NF processes, and a wait queue that contains all idle NF processes. It offers a set of APIs that the NF runtime can use to transfer the ownership of NF processes of a chain from CFS to the cooperative scheduler. These APIs are used by the Galleon agent, which runs as a privileged process. NF threads are not aware of these APIs and cannot interact with coopsched to change scheduling policy or core affinity.

Once a chain is deployed, all NF threads are managed by the Cooperative Scheduler, and are placed in the scheduler's wait queue as *detached*. Once an NF chain switches into the *attached* state, the Cooperative Scheduler pushes NFs of this chain into its run queue and ensures that the original NF dependencies are preserved in the run queue. To detach a chain, the Cooperative Scheduler waits for the chain to finish processing a batch of packets, if any, and then moves these NF processes back to the wait queue.

How scheduling works. The Galleon controller deploys multiple NF runtime instances that contain NF processes in a chain. Once an NF process starts, Galleon's NF runtime reports its thread ID (tid) to the Galleon agent running on the same worker. Once all NF processes are ready, the Galleon agent *registers* them as a scheduling group (sgroup) to the Cooperative Scheduler. After that, the cooperative scheduler takes full control of NF processes.

When the Galleon controller assigns flows to an NF chain (§5), the Cooperative Scheduler *attaches* the chain to the core. When the monitoring thread sees no traffic has arrived for the chain, the Scheduler *detaches* the chain, so the Galleon controller can re-assign the core.

To effect attach and detach operations, and to schedule NF chain execution, the Cooperative Scheduler has a master thread for serving scheduling requests and runs one enforcer thread on each managed core. The scheduler utilizes a key feature of Linux FIFO thread scheduling: 1) high-priority threads preempt low-priority threads; 2) a thread is executed once it is at the head of run queue, and is moved to the tail after it finishes. An enforcer thread is raised to the highest priority when enforcing scheduling decisions. When an NF chain is instantiated on a core, the enforcer thread registers the corresponding NF processes as low-priority FIFO threads so that they are appended to the wait queue. When attaching the NF chain, it moves NF processes to the run queue by assigning them a higher priority, and vice versa when detaching a sgroup. Operations are done in the sequence that NFs are positioned in the NF chain, so when an NF yields, the CPU scheduler automatically schedules the next NF in the chain.

In this model, each worker machine splits CPU cores into two groups. One group is managed by the Cooperative Scheduler, while the other runs with normal threads managed by CFS. NF processes run in a normal during it startup time. We use a standard kernel and support different schedulers on different cores. This enables running running NF and non-NF workloads on the same machine.

Adaptive batch size. The Cooperative Scheduler introduces N context switches for a chain with N NFs. Without an appropriate batch size, the core may spend a significant portion of time on these context switches. Galleon applies adaptive batch scheduling to bound the context switch overhead within

---

[2]To deal with non-responsive NFs, the runtime terminates chain execution if an NF fails to yield after a conservative timeout.

a threshold of total CPU execution time. This optimization largely applies to long but lightweight chains. In some cases, deploying such a lightweight long chain with a low latency requirement prevents us from achieving large batch sizes at execution time. However, a lightweight chain is often executed with high per-core throughput. Such chains easily serve tens of Gbps of traffic with just a few cores. As such, the absolute number of extra CPU cores remains small.

For an NF chain v with N NFs, suppose $B_v \in \mathbb{N}$ is the max packet batches that the first NF should process before it sends it downstream. We estimate the packet rate R_v that v can achieve at runtime as follows: $Freq$ is the CPU core's clock rate. $T_{v,i}$ is the ith NF's service time in CPU cycles, T_{ctx} is the profiled time duration of a context switch in seconds. b_v is the average number of packets per batch at the input of a chain. b_v depends on both the traffic arrival pattern and the chain service pattern, and is limited by the hardware, i.e. $b_v \in [1, b_m]$ where $b_m \in \mathbb{N}$ is the maximum number of packets in a batch.

$$R_v = \frac{Freq}{\sum_{i=1}^{N} T_{v,i} + N \cdot T_{ctx} \cdot Freq \cdot \frac{1}{B_v \cdot b_v}} \quad (1)$$

Suppose $\hat{R}_v$ is the packet rate when running v in a single thread. The goal is to bound the downgraded maximum per-core packet rate within a ratio p of the maximum rate $\hat{R}_v$. The batch size for chain v is given by:

$$min\ B_v \quad (2)$$

s.t.

$$R_v >= p \cdot \hat{R}_v, \quad where \quad \hat{R}_v = \frac{Freq}{\sum_{i=1}^{N} T_{v,i}} \quad (3)$$

Simple algebraic manipulations suffice to compute B_v.

4.4 Spatiotemporal Packet Isolation

Current FaaS offerings constrain functions to use distributed storage for transmitting data between functions. This can significantly increase latency of packet processing; *e.g.,* Redis client to server communication, even when the client and server are on the same rack, incurs an average latency of 310 μs, and a worst-case latency of 47 ms. Across multiple NFs in a chain, the resulting latency can be substantial.

Galleon can use standard IPC mechanisms or shared memory to communicate between NFs. To permit zero-copy transfer, Galleon uses shared memory, but its design must ensure packet isolation (§2); a chain should not be able to see packets belonging to another chain, and an NF should not be able to access a packet before its upstream NF has processed it or after it has completed its processing.

Figure 2 describes how Galleon achieves this packet isolation. Each distinct NF chain is allocated a separate virtual NIC queue with SR-IOV; moreover, different NF chains run in different containers. This ensures spatial packet isolation across NF chains. Ensuring packet isolation between NFs within a single chain is harder. Ideally, the virtual NIC could directly DMA packets to shared memory so all NFs can process packets without copies. However, this would violate isolation because a downstream NF could access shared memory while the NIC writes to it.

To avoid this, Galleon gives only the first NF in the chain access to the NIC queue, to process a batch of packets and copy them to shared memory. Thus, the first NF's packet memory is *spatially* isolated from other chains and from downstream NFs. Cooperative scheduling ensures NFs run in the order they appear in the chain, so *temporal* packet isolation is also preserved: even though it has access to shared memory, a downstream NF cannot access a batch that has not been processed by an upstream NF since it will not be scheduled. At the same time, this permits zero-copy packet transfer between all NFs except the first.

For a chain with only one NF, Galleon omits the unnecessary packet copying and cooperative scheduling. The Galleon NF runtime also applies an optimization that prefetches packet headers into the L1 cache before calling the user-defined NF for processing tasks. As shown in 7.1, this optimization improves performance by 11% when running a canonical NF.

5 Auto-scaling in Galleon

Galleon *auto-scales* (automatically adapts resources allocated to) NF chains in response to changes in traffic volume. In a canonical FaaS platform [20, 32], a FaaS controller interacts with a global ingress service and worker machines to coordinate auto-scaling. The global ingress service maintains a reception queue for incoming requests. The FaaS controller maintains a group of function instances on the workers. For each user-defined function, as new requests arrive, the ingress service dispatches each request to an idle instance. The FaaS controller scales up and down the number of instances to handle dynamic traffic while achieving cost efficiency.

Galleon employs a similar architecture (Figure 1), modified to satisfy performance and isolation goals. Unlike canonical FaaS workloads, NFV workloads typically have an SLO that specifies the target latency [31, 34]. Galleon's controller is designed to serve dynamic traffic while meeting stringent latency SLOs. The controller performs two functions: 1) balance load across the workers; 2) create/delete new instances of NF runtimes, and schedule/un-schedule NF threads among the workers. In this, it is aided by a per-worker Galleon agent that monitors NF's performance and interacts with the cooperative scheduler to enforce scheduling policies. The rest of the section describes components of the controller and the algorithms that it uses for auto scaling.

5.1 Monitoring

Monitoring is critical for scaling NF chains. At each NF runtime, an NF monitoring function collects performance statistics, including the instantaneous packet rate, NIC queue length, and the per-batch execution time. The packet rate is measured as the average processing rate of the NF chain.

NIC queue length is the number of packets reported by the NIC hardware. It also estimates per-batch execution time by recording the global CPU cycle counter at the beginning and the end of sampled executions.

To avoid interfering with data-plane processing, the NF monitoring function runs in a separate thread and is not scheduled on a core running NF processing threads. Each NF runtime maintains statistics locally, and sends updates to the Galleon controller only when significant events occur (to minimize control overhead). Examples of such events include queue lengths or rates exceeding a threshold.

5.2 Galleon ingress

In a typical Galleon deployment (*e.g.,* at an ISP edge), the traffic aggregate hitting an NF chain is likely to exceed the capacity of a single core. In general, then, multiple instances of an NF chain might be deployed across the workers. The load balancer component interacts with the Galleon ingress to split flows among the deployed NF chains.

Consider a rack deployment of Galleon: all worker machines are on a rack and connected to a ToR switch. External traffic enters the rack at this ToR switch. When a new flow hits the ToR switch, it does not match any rules, so the ToR forwards the flow to Galleon's ingress (a software component running on one of the workers). The ingress then queries the Galleon controller to select an NF chain to assign to this flow. It then installs a flow entry on the ToR switch to route subsequent packets of this flow to the worker running the NF chain instance. To do this, Galleon relies on L2 tagging to forward the flow to the worker's virtual NIC. Prior work [28] proposes to make such assignments at a flowlet (sub-flow) level to better balance load in the face of different flow sizes. Galleon can incorporate this optimization, which we have left this to future work.

Galleon could have performed this reactive rule installation using an OpenFlow controller. We found that this approach resulted in packet losses caused by high flow-entry-installation latency. Our ingress is implemented as a service and can buffer packets until the flow entry is installed, thereby avoiding packet loss.

The Galleon controller makes load balancing decisions based on statistics received from each worker (§5.1). For each flow, it always assigns the NF chain with the highest traffic load (the current packet rate / the maximum packet rate) among all non-overloaded chains. To determine a load threshold for setting a chain as overloaded, Galleon applies a profiling-based solution to pick the highest threshold that results in a latency bounded by the SLO requirement.

5.3 Scaling of NF Chains

Galleon attempts to schedule NF chains on the fewest number of CPU cores that can serve traffic without violating SLOs. It uses a combination of mechanisms to achieve this.

First, to avoid cold start latency and the resulting packet

```
def handle_flows_at_ingress(flow,
                            chains: NFChainsList):
  # Each logical chain is associated with
  # a list of active NF chains v.sgroups,
  # and a list of idle NF chains v.idle_sgroups
  for v in chains:
   if v.should_serve(flow):
    selected=v.
     pick_highest_load_chain_without_overload()
    if selected is None:
     selected=v.
     pick_idle_chain_from_lowest_load_worker()
    selected.set_active()
    update_routing(match=flow,
                   action=(set_port=selected.port,
                           set_l2=selected.l2_tag))
    while len(v.idle_sgroups) < scale_out_thresh:
     scale_out(v)
    while len(v.idle_sgroups) > scale_in_thresh:
     scale_in(v)

def worker_scheduler_loop(worker: GalleonWorker):
 while True:
  for sg in worker.sgroups:
   if sg.is_sched and not sg.is_active:
    worker.detach_sgroup(sg)
   # sg is set active by the above ingress
   if sg.is_active and not sg.is_sched:
    idle_core=worker.pick_idle_core()
    worker.attach_sgroup_to_core(sg, idle_core)
```

Listing 1: Pseudocode for scaling NF chains in a Galleon cluster

loss, Galleon maintains a pool of pre-deployed NF chains; it leverages prior work for reducing the startup time of serverless instances [17, 19]. Pre-deployed NF threads start out in the detached state, so they do not consume CPU resources.

Second, Galleon profiles the NF chain to determine the maximum packet rate achievable on a single core. When the traffic assigned to an NF chain approaches that rate, Galleon decides to allocate new incoming flows to a new NF chain instance. This is accomplished by the Galleon ingress (§5.2).

Finally, when an NF thread becomes idle (all flows previously assigned to it have completed), Galleon reclaims the assigned core. Galleon could have, instead, migrated flows away from underutilized NF chains, but this would have complicated state management for stateful NFs. We have left this optimization to future work.

Listing 1 describes the algorithms used at the ingress to assign flows to NF chain instances, and the algorithms used at the worker's cooperative scheduler to dynamically attach and detach cores from NF chains in response to traffic dynamics. Together, these algorithms achieve auto scaling in Galleon.

6 Implementation

Galleon is built upon OpenFaaS [20]. OpenFaaS consists of an infrastructure layer and an application layer. The former uses three components: Kubernetes, Docker, and the Container Registry. Galleon reuses these APIs for managing and deploying NFs. At the application layer, OpenFaaS has its own gateway service for triggering function execu-

tions. Galleon adds a separate ingress component. Incoming traffic is split at the system gateway; requests for normal applications are forwarded to OpenFaaS's gateway, while NFV traffic is forwarded to the Galleon ingress. OpenFaaS uses a function runtime that maintains a tunnel to the FaaS gateway, and hands off requests to user-defined functions. Galleon uses a different NF runtime by using the above mechanisms for receiving traffic from the Galleon ingress (§5). Galleon reuses OpenFaaS's general monitoring framework, but relies on a per-worker agent for NF performance monitoring and enforcing scheduling policies.

Galleon controller. This multi-threaded controller is implemented in Go and adds 9K LOC to OpenFaaS. It controls the Galleon ingress and the worker cluster subsystem via RPCs and Kubernetes APIs. On each worker, this controller deploys an agent that runs as a RPC server for collecting NF performance statistics, updating critical performance events, and enforcing scheduling policies. The controller also includes a UI subsystem. Operators may use this UI system to deploy NF chains, adjust cooperative scheduling and batch scheduling policies, and monitor NF chains by clearly seeing which chain runs at which core in the cluster with its instantaneous packet rate, load and NIC queue length.

Galleon ingress. This is implemented in 800 LOC of C++ as a BESS [6] module that receives Galleon controller's load balancing decision and enforces it by modifying the hardware switch table via OpenFlow.

Galleon Cooperative Scheduler. This is implemented in C as a multi-threaded process in 5K LOC. It offers a set of APIs that can be used by agents.

Galleon NF Runtime. This is implemented in C++ by reusing some BESS libraries (with 2.5K LOC including an RPC server and a driver that works with Galleon's network I/O). It runs two processes: one for RPC server and one for NF. The NF runtime uses an external datastore (i.e. a Redis cluster) for maintaining global NF states.

7 Evaluation

Experiment setup. We use a Cloudlab [2] cluster of 10 servers, each with dual-CPU 16-core 2.4GHz Intel Xeon(R) CPU E5-2630 v3 (Haswell) with 128GB ECC memory (DDR4 1866 MHz). To reduce jitter, all CPU cores have hyperthreading and CPU frequency scaling disabled and run at a fixed 2.4 GHz. Each server has one dual-port 10 GbE Intel X520-DA2 NIC. Both are connected to an experimental LAN for data-plane traffic. Each machine has one 1 GbE Intel NIC for control and management traffic. Servers connect to a Cisco C3172PQs ToR switch with 48 10 GbE ports and Openflow v1.3 support. The traffic generator and the Galleon ingress run on dedicated machines.

Methodology and Metrics. Our Galleon experiments employ end-to-end traffic. We deploy three representative chains from light to heavy in their CPU cycle cost, from documented

use cases [12]. Chain 1 implements a lightweight L2/L3 pipeline for tunneling: Tunnel→IPv4Forward; Chain 2 implements a complex chain with DPI and encryption NFs: ACL→UrlFilter→Encrypt; Chain 3 is a state-heavy chain that requires connection consistency: ACL→NAT. Key performance metrics include: the end-to-end latency distribution and packet loss rate, and the time-average and maximum CPU core usage for the test duration. The traffic generator uses BESS to generate flows with synthetic test traffic.

7.1 Executing NF chains

We explore how Galleon's performance changes when running an NF chain of various lengths. We compare the per-core max throughput achieved by Galleon and other NFV systems that make different isolation choices. For simplicity, we use a test NF (a BPF [15] module with a table of 200 BPF rules) and run chains with a sequence of the same NF.

Isolation via copying. EdgeOS [24] supports isolation using data copying. We emulate EdgeOS on top of a reimplementation of NFVNice [11] and apply the same packet copying technique to the master module that acts as the message bus for transmitting packets between processes. The master module runs as a multi-threaded process with one RX thread for receiving packets from the NIC, one TX thread for transmitting packets among processes, and one wake-up thread for notifying a process that a message has arrived at its message buffer. All three threads run on dedicated cores.

Isolation via safe languages. NetBricks [22] uses compile-time language support from Rust to ensure isolation among NFs plus a run-time array bound check.

Results. Figure 3 shows the throughput of different isolation approaches for different length chains. Galleon outperforms both NetBricks (1.28-1.44× perf) and NFVNice w/ packet copying (1.05-1.93× perf). NFVNice with packet copying has similar performance (95%) with a single-NF chain, but as chain length increases its throughput decreases despite its 3 extra CPU cores for transmitting packets among NFs.

We also implemented a variant labeled Galleon (single thread) that runs all NFs in a single thread. This variant offers no isolation because all NFs run in the same process, with neither language support for isolation nor our spatiotemporal packet isolation mechanism. Compared to this unsafe-but-fast variant, Galleon has an overhead that remains at the same level regardless of the chain length: Galleon achieves a 89.9%-95.2% per-core throughput when deploying a multi-NF chain while providing isolation. For a chain with one NF, Galleon achieves slightly better performance as it applies the prefetch-into-L1 optimization described earlier.

7.2 Scaling Experiments

Next we evaluate Galleon's performance when deploying NF chains with latency SLOs.
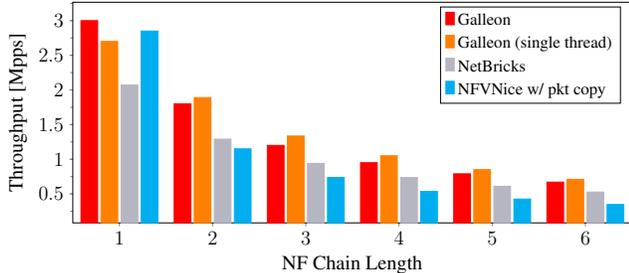
Figure 3: Throughput with increasing chain length for running an NF chain on a single core when using 80-byte packets.

7.2.1 Controlling NF chain latency

Prior to the deployment of an NF chain, Galleon profiles the chain to get its latency profile under different per-core load thresholds. When deploying a chain, the Galleon controller considers the user-specified NF chain latency SLO requirement, and uses the per-core load threshold as a knob to control the end-to-end delay for packets being processed by NF chains deployed in the cluster.

Results. Figure 4 shows profiling results for NF chains used in the cluster-scale experiments. As shown in the figure, the NF chain's tail latency is reduced when the Galleon controller gradually decreases the per-core load threshold when running a chain. According to the profiling results, we see the cause of this behaviour: by reducing the per-core load threshold for a chain, the chain runs at an operating point where the maximum queue length is lower. This is because the Galleon controller considers a chain as overloaded if its current traffic load exceeds the threshold, and stops assigning new flows to this chain. Therefore, this design limits the amount of traffic destined to a chain, and further controls the maximum queue length observed by the chain, and therefore controls the chain's end-to-end latency.

This feature, of course, aligns with the trade-off between latency and efficiency: a smaller tail latency means a smaller per-core throughput. When serving the same amount of traffic, it means that the Galleon controller has to devote more CPU cores for running this chain. This unique feature is also very important in the FaaS context: Galleon is able to use the right level of system resources to meet the latency SLO when deploying NFV chains.

7.2.2 The end-to-end evaluation

We compare Galleon against a customized, high-performance NFV platform, Metron [9]. Like Galleon, Metron also utilizes the hardware switch for dispatching traffic. Metron compiles NFs into a single process, and runs-to-completion each chain as a thread. We have implemented both Metron's runtime and scaling algorithm. Each Metron runtime is a single-thread process that runs an NF chain. Packets are transmitted between NFs with no isolation mechanisms (as Metron does not consider isolation).

We compare Galleon and Metron in these end-to-end experiments. Before experiments, an NF chain specification
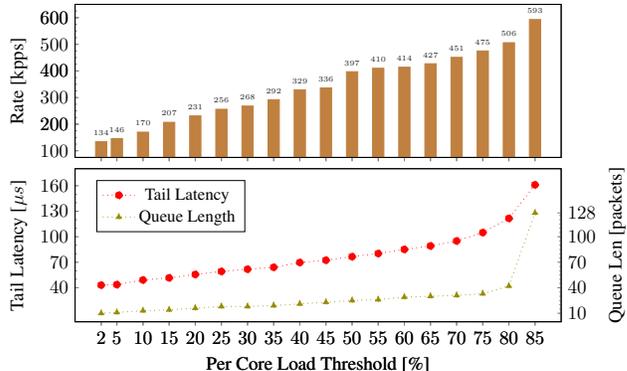
Figure 4: NF Chain Profiling: Effect of per-core load threshold on tail latency (p99), queue length, and rate.

is passed to both system's controllers to deploy NFs in the test cluster. For each experiment, we run a DPDK-based flow generator to generate traffic at 20 flows/s with a median packet size of 1024-byte. We estimate this level of flow rate by analyzing an ISP trace [29]. All incoming flows are forwarded to the system ingress and to be processed by their destined NF chains. The traffic generator gradually increases the number of flows and reaches the maximum throughput after 60 seconds, and results in an average of 1.2K flows at the peak load of 18 Gbps. This number of flows is similar to that used by prior work [8, 28]. Then the traffic generator stays at the steady state to keep generating traffic at the maximum rate until the total number of packets reaches the 100 million. We compare end-to-end performance metrics, including the tail latency, the maximum number of cores, and the time-averaged number of cores for deploying the chain.

Results. Across all experiments, both systems achieve a zero loss rate. Thus, we compare two systems by looking at the tail (p99) latency of the CPU core usage when they serve the test traffic (100 million packets). Figure 5 shows this for different latency targets. For all chains, Galleon can precisely meet the SLO for its tail latency. Metron ignores the latency SLO and targets a deployment without losses.

Take NF chain 1 as an example. Metron fails to meet a low latency SLO target (50 us). Both Galleon and Metron can satisfy the medium latency SLO (100 us). In terms of CPU usage, we notice that Galleon and Metron use the same number of cores at peak. Galleon uses a slightly higher time-averaged number of cores (15.5%). We see this gap as the overhead of achieving isolation among NFs.

7.3 Overhead

Next, we show a breakdown of Galleon overhead. Galleon's overhead comes from spatiotemporal isolation, including 1) the spatial isolation of a virtual NIC to isolate chains and packet copy from the NIC's buffer to the NF chain's buffer, and 2) the temporal isolation provided by transparent cooperative scheduling, with extra context switches among NFs (unlike in a single process).

10
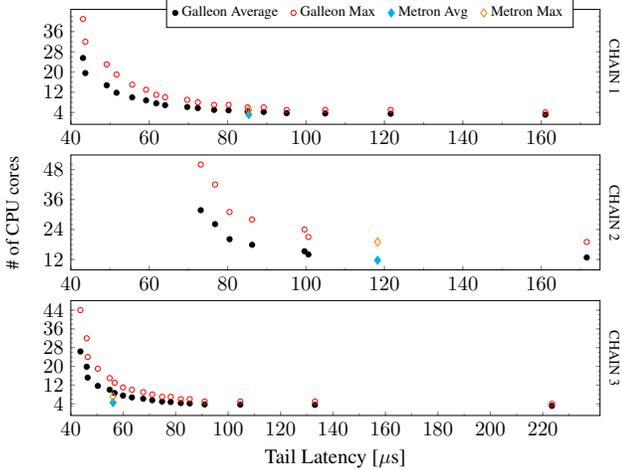
Figure 5: Core usage of NF chains implemented in Galleon and Metron as a function of achieved tail latency. Per-core thresholds are set as in Figure 4.
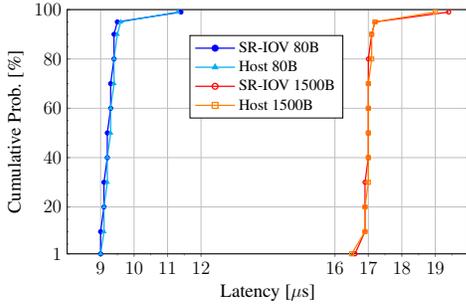
**Figure 6:** End-to-end latency CDF with SR-IOV on and off.

7.3.1 Spatial isolation overhead

To evaluate Galleon's spatial isolation overhead, we quantify the SR-IOV NIC's overhead by comparing performance metrics when running a test NF with and without enabling SR-IOV. Our test NF is a MACSwap module that swaps the dst and src Ethernet addresses of a packet.

Results. Figure 6 shows that running the NF with an SR-IOV enabled NIC only increases the latency by 0.1 us for both 80-byte and 1500-byte packets. We also find that the maximum throughput achieved by a SR-IOV enabled NIC is always greater than or equal to 99.6% of the throughput achieved by a NIC running in a non-virtualized mode.

7.3.2 Temporal isolation overhead

For a chain, the ownership of a packet is transferred between NFs as the chain proceeds. Packet isolation requires that NFs in the same chain should only acquire the ownership of a packet after its predecessor finishes processing it (§4.4). To quantify this temporal isolation overhead, we use a multi-NF chain in which the first NF holds a NIC VF that is dedicated to this chain. The NIC hardware DMAs incoming packets to the NIC packet buffer that resides in this NF's memory space. The packet isolation requirement prevents other NFs from accessing the NIC packet buffer directly, as packets in
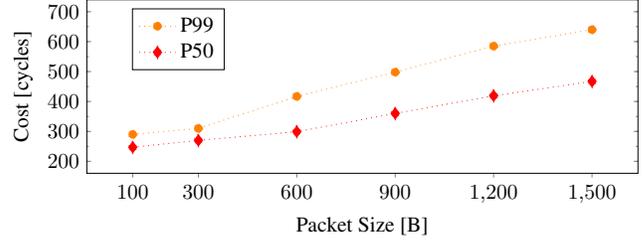
Figure 7: Per-packet cost of copying packets of different sizes

that memory region are only destined to the first NF. For the rest of the chain, NFs share the access of a per-chain packet buffer and wait to be scheduled by the Cooperative Scheduler to process the same batch of packets in the correct sequence.

Results. Figure 7 shows the p50 and p99 CPU cycle cost for copying one packet of different sizes. At median, copying a 100-byte packet costs 247 cycles, while doing so for a 1500-byte packet costs 467 cycles. The small difference of copying a small and a large packet is because the cost of allocating a packet struct takes a significant share in the total cost.

Scheduling NFs cooperatively involves context switches between NF threads that belong to different NF processes. We profile the average cost of context switches between NFs: 2143 cycles per context switch. Note that this context switch cost is amortized among the batch of packets in each execution. For a default 32-packet batch, the amortized cost is only 67 cycles per packet. This cost is 27 | 14% of the cost for copying a 64 | 1500-byte packet respectively. Further, it is only 31% of the cost for forwarding a packet via a vSwitch with packet copying, as in NFVNice. (Galleon has zero packet switching cost because it uses the ToR switch and the NIC's L2 to dispatch packets to different chains.)

Alternative packet ownership transfer. We also consider a traditional mechanism for ensuring isolation for a shared memory region in a multi-process environment by using `unmap` and `map` to explicitly move the ownership of the shared packet buffer. We profile the average cost for these two options: `unmap` is 4083 cycles, while `map` is 8495 cycles. With all packets placed in the same memory page, we need one `unmap` and `map` to transfer the memory page to a different process. We conclude that their costs are significantly larger (5.87×) than the context switch overhead.

Cooperative scheduling effects. For a chain, cooperative scheduling involves context switches between different NF processes. This operation can also flush caches and TLBs; we conduct an experiment to understand these effects.

Here we run the same test chain, with 5 BPF modules, as in 7.1 with four experimental groups: 1) Galleon: the vanilla deployment w/o adaptive batch optimization; 2) Local mem: the vanilla deployment that operates on one dummy packet in the shared memory region; 3) Dummy NF: a chain of dummy NFs that do not process packets, but simulate NF cycle costs; 4) Single thread: a chain that runs in a single thread. None use adaptive batch optimization so that they run with the same

Metric		Galleon	Local mem	Dummy NF	Single thread
Per packet cycles		2846	2746	2730	2592
Chain delay	P50	121536	116232	116008	85980
(in cycles)	P99	128116	122324	117892	87492
Misses	dTLB	72,864,218	68,259,827	61,251,661	1,185,591
	dTLB (%)	0.52%	0.48%	4.07%	0.00%
	iTLB	11,542,564	12,325,256	9,615,381	591,737
	iTLB (%)	478.18%	488.28%	379.31%	182.39%
	Cache	18,923,855	6,710,255	6,699,896	12,963,766
	L1 dcache	508,578,460	417,298,551	333,262,806	327,837,034
	L1 icache	41,272,281	37,127,027	30,856,178	17,568,605

Table 1: Overheads under isolation variants.

batch size. One traffic generator produces excessive traffic (1024B packets) to saturate the chain's NIC queue so that each chain runs at a batch size of 32, the NIC's default batch size. The TLB and cache misses are measured as the average value for a 15-second execution duration for 5 measurements. *Results.* Table 1 shows NF runtime stats. For all multiple-process groups, we see higher iTLB and dTLB misses. As shown, the number of dTLB misses reaches less than 1% of dTLB hits for cases that run a non-dummy NF, though dTLB misses are less important for an NF's performance.

All multi-process groups see higher iTLB misses compared to the single thread case because NF processes do not share code in memory. The 'local mem' and 'dummy NF' cases perform similarly in terms of per-packet cycle cost (and the number of cache misses). This is because the 'local mem' case processes one packet that resides in the chain's local memory and is likely to benefit from the L3 cache. The 'Galleon' case has a slightly higher per-packet cost compared to the other two multi-process cases. Breaking down per-chain cycle cost into the per-NF level, we find that the extra cost only comes from the first NF that copies incoming packets. The 2nd-4th NFs in 1)-3) have the same cycle cost (509 cycles / packet). These NFs benefit from L3 caching as the first NF's runtime loads when copying packets from the NIC's packet buffer to the per-chain packet buffer.

In the above four cases, two major differences explain the per-chain cycle cost: a) iTLB misses when deploying as a multi-process chain and b) L3 cache misses when processing network traffic. The 'Galleon' case has both; the 'Local mem' and 'Dummy NF' cases have a); and 'Single thread' has b). We can calculate the breakdown cycle cost for each via simple algebraic manipulations for this 5-NF chain: a) 254 cycles / packet (or 50.8 cycles / packet / hop); b) 100 cycles / packet. Among these two, a) is the extra overhead of a context switch, which could be reduced by tagged TLBs; b) is not an overhead as a chain must load packets into the CPU cache once when processing network traffic. Thus the amortized TLB overhead is relatively small compared to the context switch itself.

8 Related Work

Serverless platforms. Serverless platforms are available as both commercial cloud services [4, 5, 18, 26], and as open-source offering [7, 20]. They are designed for latency-insensitive applications, and not suited for running NFV workloads for reasons that have been discussed along this paper.

Research on serverless platforms has explored two issues. The first improves aspects of FaaS platforms. Sock [19] and LightVM [17] optimizes sandbox startup time. SAND [1] optimizes IPC performance. E3 [16] accelerates serverless execution with SmartNIC to improve cost-efficiency. The second explores new applications only made possible with serverless, such as real-time big data analysis [10, 23] and scalable video encoding [3]. ServerlessNF [27, 28] executes NFs on serverless platforms. While offering insights on handling different flow sizes, it does not propose changes to the infrastructure, and thus does not leverage techniques for building high-performance NFV platforms. Indeed, its main focus is on deploying a single NF without meeting SLOs for NF chains [31, 34]. ServerlessNF's centralized resource manager (RM) incurs extra overhead of CPU cores and NIC bandwidth, and poses a new challenge to make RM scalable while meeting SLOs.

NFV frameworks. NFV frameworks [9, 11, 21, 22, 34] deploy, execute, and orchestrate NF chains in a highly optimized way. E2 [21] deploys NF chains using commodity servers, but lacks optimizations and isolation mechanisms. Subsequent work explored a number of specific topics. **Isolation:** NetBricks [22] isolates NFs with a safe runtime based on Rust. This assumes that NF vendors offer source code written in a certain language, which is uncommon in practice. In contrast, most commercial NFs are packaged containers or VMs without source. EdgeOS [24] employs an expensive isolation mechanism via packet copying for each NF. **Hardware acceleration:** a number of papers examine accelerating NFV by leveraging specialized hardware such as P4 switches [9, 13, 34]. Galleon sees the NFV's fundamental motivation as reducing infrastructure management cost, and does not use specialized hardware platforms that are not widely available in today's cloud infrastructure. **Performance tuning:** Some [11, 14, 31] examine dirty-slate solutions to optimize NFV's performance. They often work in a restricted setting without considering multi-tenant cluster deployment; that said, this line of work is largely compatible with Galleon and could be integrated for further performance gains.

Stateful NFs. Another line of research [8, 28, 33] considers optimizations for stateful NFs. Galleon supports stateful NFs but does not specifically implement some of these more advanced techniques; these are complementary mechanisms that can be integrated into Galleon.

9 Conclusions

Galleon supports NFV in edge and cloud computing environments using commodity software and hardware components, fulfilling NFV's original ambitions. It minimally extends FaaS abstractions, and eases the deployment of third-party NFs. With its spatiotemporal packet isolation mechanism, it outperforms state-of-the-art NFV platforms that use alternative isolation mechanisms, and has performance comparable to custom NFV systems that do not consider NF isolation.

References

[1] Istemi Ekin Akkus, Ruichuan Chen, Ivica Rimac, Manuel Stein, Klaus Satzke, Andre Beck, Paarijaat Aditya, and Volker Hilt. SAND: Towards high-performance serverless computing. In *2018 USENIX Annual Technical Conference (USENIX ATC 18)*, pages 923–935, Boston, MA, 2018. USENIX Association.

[2] Dmitry Duplyakin, Robert Ricci, Aleksander Maricq, Gary Wong, Jonathon Duerig, Eric Eide, Leigh Stoller, Mike Hibler, David Johnson, Kirk Webb, Aditya Akella, Kuangching Wang, Glenn Ricart, Larry Landweber, Chip Elliott, Michael Zink, Emmanuel Cecchet, Snigdhaswin Kar, and Prabodh Mishra. The design and operation of CloudLab. In *Proceedings of the USENIX Annual Technical Conference (ATC)*, pages 1–14, July 2019.

[3] Sadjad Fouladi, Riad S Wahby, Brennan Shacklett, Karthikeyan Vasuki Balasubramaniam, William Zeng, Rahul Bhalerao, Anirudh Sivaraman, George Porter, and Keith Winstein. Encoding, fast and slow: Low-latency video processing using thousands of tiny threads. In *Proceedings of USENIX NSDI*, pages 363–376, 2017.

[4] The Apache Software Foundation. Apache OpenWhisk. http://openwhisk.org/, 2017. Accessed on May 2020.

[5] Google. Google Cloud Functions. https://cloud.google.com/functions/, 2017. Accessed on May 2020.

[6] Sangjin Han, Keon Jang, Aurojit Panda, Shoumik Palkar, Dongsu Han, and Sylvia Ratnasamy. softnic: A software nic to augment hardware. Technical report.

[7] Scott Hendrickson, Stephen Sturdevant, Tyler Harter, Venkateshwaran Venkataramani, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau. Serverless computation with openlambda. In *8th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 16)*, pages 1–6, Denver, CO, 2016. USENIX Association.

[8] Murad Kablan, Azzam Alsudais, Eric Keller, and Franck Le. Stateless network functions: Breaking the tight coupling of state and processing. In *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*, pages 97–112, Boston, MA, 2017. USENIX Association.

[9] Georgios P. Katsikas, Tom Barbette, Dejan Kostić, Rebecca Steinert, and Gerald Q. Maguire Jr. Metron: NFV service chains at the true speed of the underlying hardware. In *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18)*, pages 171–186, Renton, WA, April 2018. USENIX Association.

[10] Ana Klimovic, Yawen Wang, Patrick Stuedi, Animesh Trivedi, Jonas Pfefferle, and Christos Kozyrakis. Pocket: Elastic ephemeral storage for serverless analytics. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*, pages 427–444, Carlsbad, CA, 2018. USENIX Association.

[11] Sameer G. Kulkarni, Wei Zhang, Jinho Hwang, Shriram Rajagopalan, K. K. Ramakrishnan, Timothy Wood, Mayutan Arumaithurai, and Xiaoming Fu. Nfvnice: Dynamic backpressure and scheduling for nfv service chains. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, SIGCOMM '17, pages 71–84, New York, NY, USA, 2017. ACM.

[12] Surendra Kumar, Mudassir Tufail, Sumandra Majee, Claudiu Captari, and Shunsuke Homma. Service Function Chaining Use Cases In Data Centers. Internet-Draft draft-ietf-sfc-dc-use-cases-06, Internet Engineering Task Force, February 2017. Work in Progress.

[13] Yanfang Le, Hyunseok Chang, Sarit Mukherjee, Limin Wang, Aditya Akella, Michael M. Swift, and T. V. Lakshman. *UNO: Uniflying Host and Smart NIC Offload for Flexible Packet Processing*, page 506–519. Association for Computing Machinery, New York, NY, USA, 2017.

[14] Tamás Lévai, Felicián Németh, Barath Raghavan, and Gabor Retvari. Batchy: Batch-scheduling data flow graphs with service-level objectives. In *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*, pages 633–649, Santa Clara, CA, February 2020. USENIX Association.

[15] Bpf documentation. https://www.kernel.org/doc/html/latest/bpf/index.html.

[16] Ming Liu, Simon Peter, Arvind Krishnamurthy, and Phitchaya Mangpo Phothilimthana. E3: Energy-efficient microservices on smartnic-accelerated servers. In *2019 {USENIX} Annual Technical Conference ({USENIX}{ATC} 19)*, pages 363–378, 2019.

[17] Filipe Manco, Costin Lupu, Florian Schmidt, Jose Mendes, Simon Kuenzer, Sumit Sati, Kenichi Yasukata, Costin Raiciu, and Felipe Huici. My vm is lighter (and safer) than your container. In *Proceedings of the 26th Symposium on Operating Systems Principles*, SOSP '17, pages 218–233, New York, NY, USA, 2017. ACM.

[18] Microsoft. Azure Functions. https://azure.microsoft.com/en-us/services/functions/, 2017. Accessed on May 2020.

[19] Edward Oakes, Leon Yang, Dennis Zhou, Kevin Houck, Tyler Harter, Andrea Arpaci-Dusseau, and Remzi Arpaci-Dusseau. SOCK: Rapid task provisioning with serverless-optimized containers. In *2018 USENIX Annual Technical Conference (USENIX ATC 18)*, pages 57–70, Boston, MA, 2018. USENIX Association.

[20] Openfaas. https://www.openfaas.com/.

[21] Shoumik Palkar, Chang Lan, Sangjin Han, Keon Jang, Aurojit Panda, Sylvia Ratnasamy, Luigi Rizzo, and Scott Shenker. E2: A framework for nfv applications. In *Proceedings of the 25th Symposium on Operating Systems Principles*, SOSP '15, page 121–136, New York, NY, USA, 2015. Association for Computing Machinery.

[22] Aurojit Panda, Sangjin Han, Keon Jang, Melvin Walls, Sylvia Ratnasamy, and Scott Shenker. Netbricks: Taking the v out of NFV. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pages 203–216, Savannah, GA, November 2016. USENIX Association.

[23] Qifan Pu, Shivaram Venkataraman, and Ion Stoica. Shuffling, fast and slow: Scalable analytics on serverless infrastructure. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*, pages 193–206, Boston, MA, 2019. USENIX Association.

[24] Yuxin Ren, Guyue Liu, Vlad Nitu, Wenyuan Shao, Riley Kennedy, Gabriel Parmer, Timothy Wood, and Alain Tchana. Fine-grained isolation for scalable, dynamic, multi-tenant edge clouds. In *2020 USENIX Annual Technical Conference (USENIX ATC 20)*, pages 927–942. USENIX Association, July 2020.

[25] Salvatore Sanfilippo. Redis. https://github.com/redis/redis, May 2009.

[26] Amazon Web Services. AWS Lambda. https://aws.amazon.com/lambda/, 2017. Accessed on May 2020.

[27] Arjun Singhvi, Sujata Banerjee, Yotam Harchol, Aditya Akella, Mark Peek, and Pontus Rydin. Granular computing and network intensive applications: Friends or foes? In *Proceedings of the 16th ACM Workshop on Hot Topics in Networks*, HotNets-XVI, pages 157–163, New York, NY, USA, 2017. ACM.

[28] Arjun Singhvi, Junaid Khalid, Aditya Akella, and Sujata Banerjee. SNF: Serverless Network Functions. In *Proceedings of the 11th ACM Symposium on Cloud Computing*, SoCC '20, page 296–310, New York, NY, USA, 2020. Association for Computing Machinery.

[29] The caida ucsd. https://www.caida.org/data/[dataset-URL.

[30] An introduction to sr-iov technology. http://goo.gl/m7jP3.

[31] Amin Tootoonchian, Aurojit Panda, Chang Lan, Melvin Walls, Katerina Argyraki, Sylvia Ratnasamy, and Scott Shenker. Resq: Enabling slos in network function virtualization. In *Proceedings of USENIX NSDI*, pages 283–297, 2018.

[32] Liang Wang, Mengyuan Li, Yinqian Zhang, Thomas Ristenpart, and Michael Swift. Peeking behind the curtains of serverless platforms. In *2018 {USENIX} Annual Technical Conference ({USENIX}{ATC} 18)*, pages 133–146, 2018.

[33] Shinae Woo, Justine Sherry, Sangjin Han, Sue Moon, Sylvia Ratnasamy, and Scott Shenker. Elastic scaling of stateful network functions. In *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18)*, pages 299–312, Renton, WA, 2018. USENIX Association.

[34] Jane Yen, Jianfeng Wang, Sucha Supittayapornpong, Marcos A. M. Vieira, Ramesh Govindan, and Barath Raghavan. Meeting slos in cross-platform nfv. In *Proceedings of the 16th International Conference on Emerging Networking EXperiments and Technologies*, CoNEXT '20, page 509–523, New York, NY, USA, 2020. Association for Computing Machinery.