

# Lecture 7: Network Function Virtualization


CS 234 / NetSys 210: Advanced Computer Networks

Sangeetha Abdu Jyothi




This lecture uses material from Radhika Mittal (ECE/CS598HPN ) and Nick McKeown

# Last Class: Network Virtualization




Abstraction



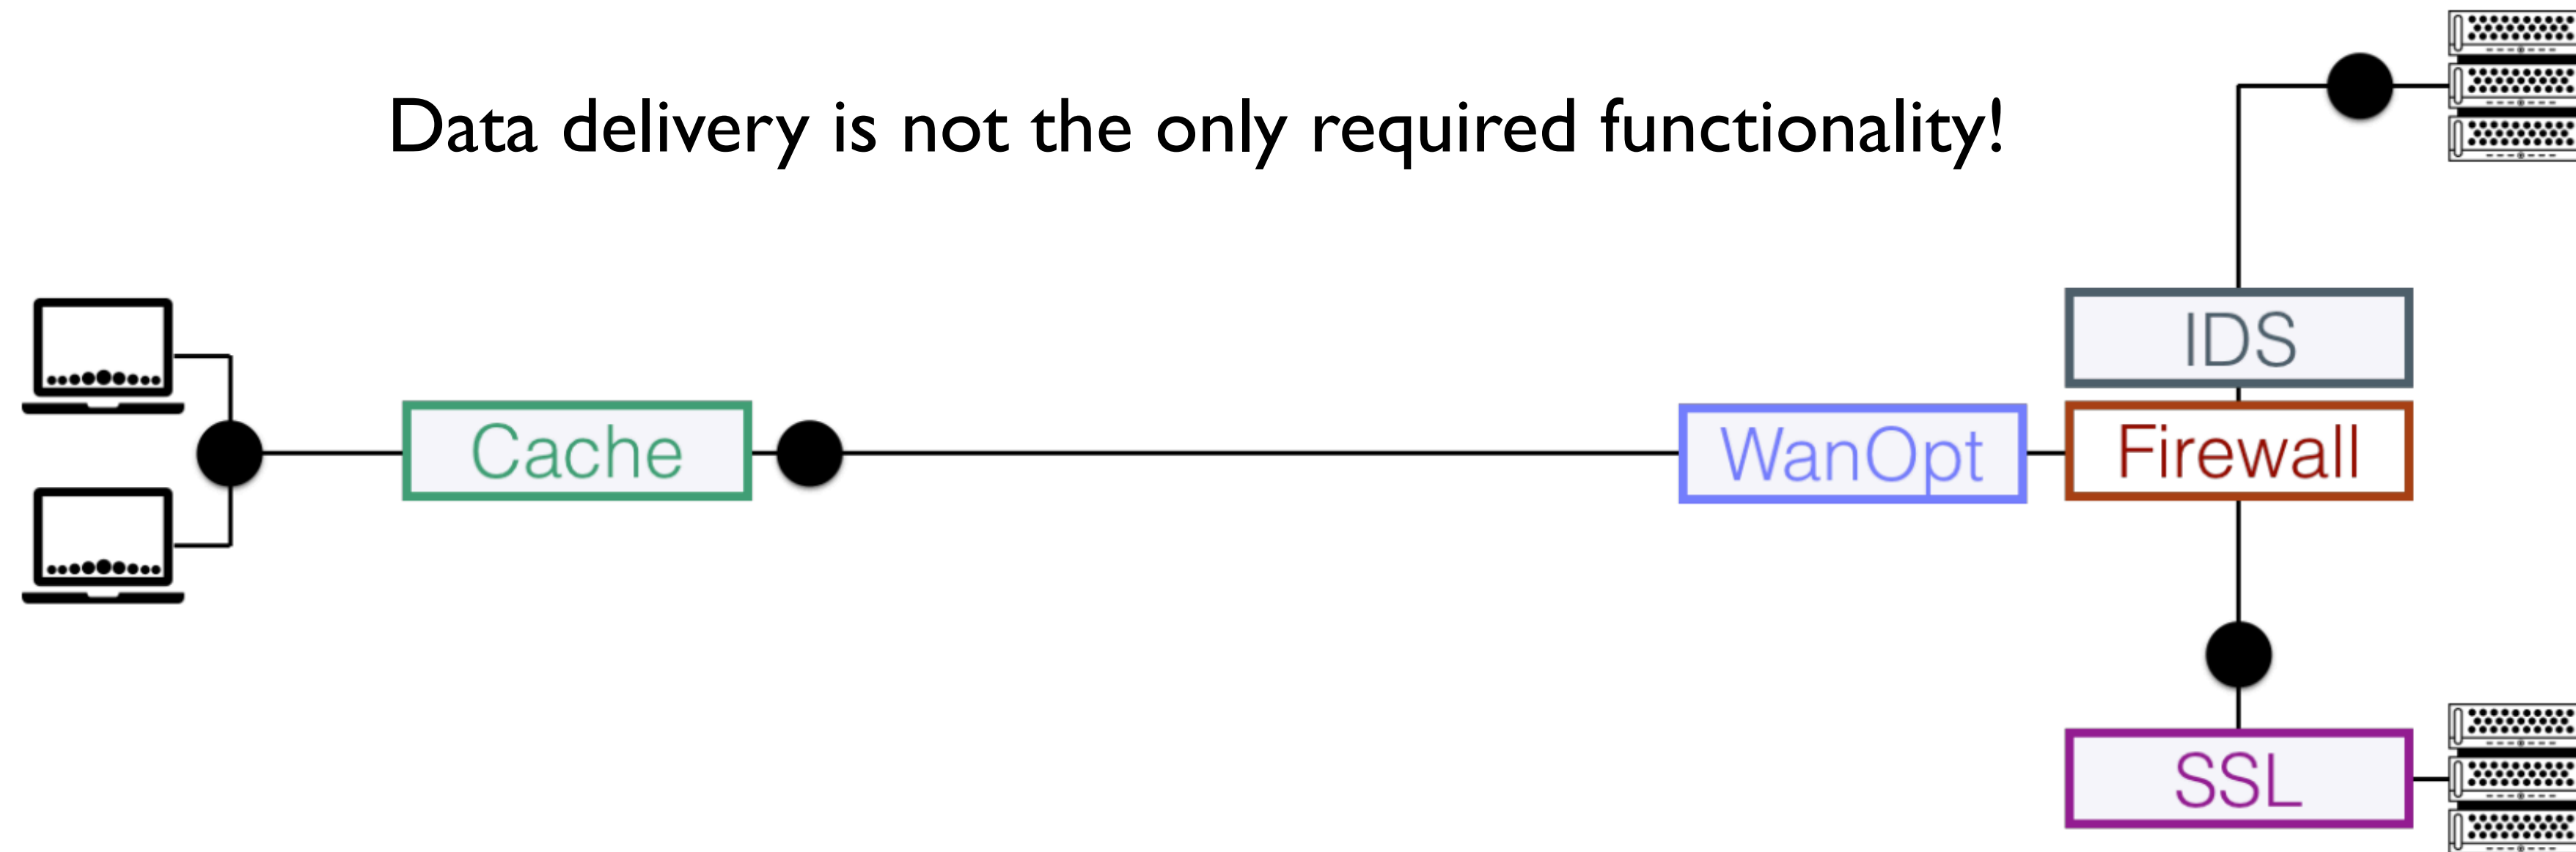
Physical Topology

# Conventional View of Networks



Data delivery is the only functionality provided by such a network

# Rise of Middleboxes



Security (IDS, Firewall): identify and block unwanted traffic

Performance (Cache) : Load content faster

Performance (WanOpt): reduce bandwidth usage

Application support (SSL): protocol for legacy application.

# Middlebox Prevalence

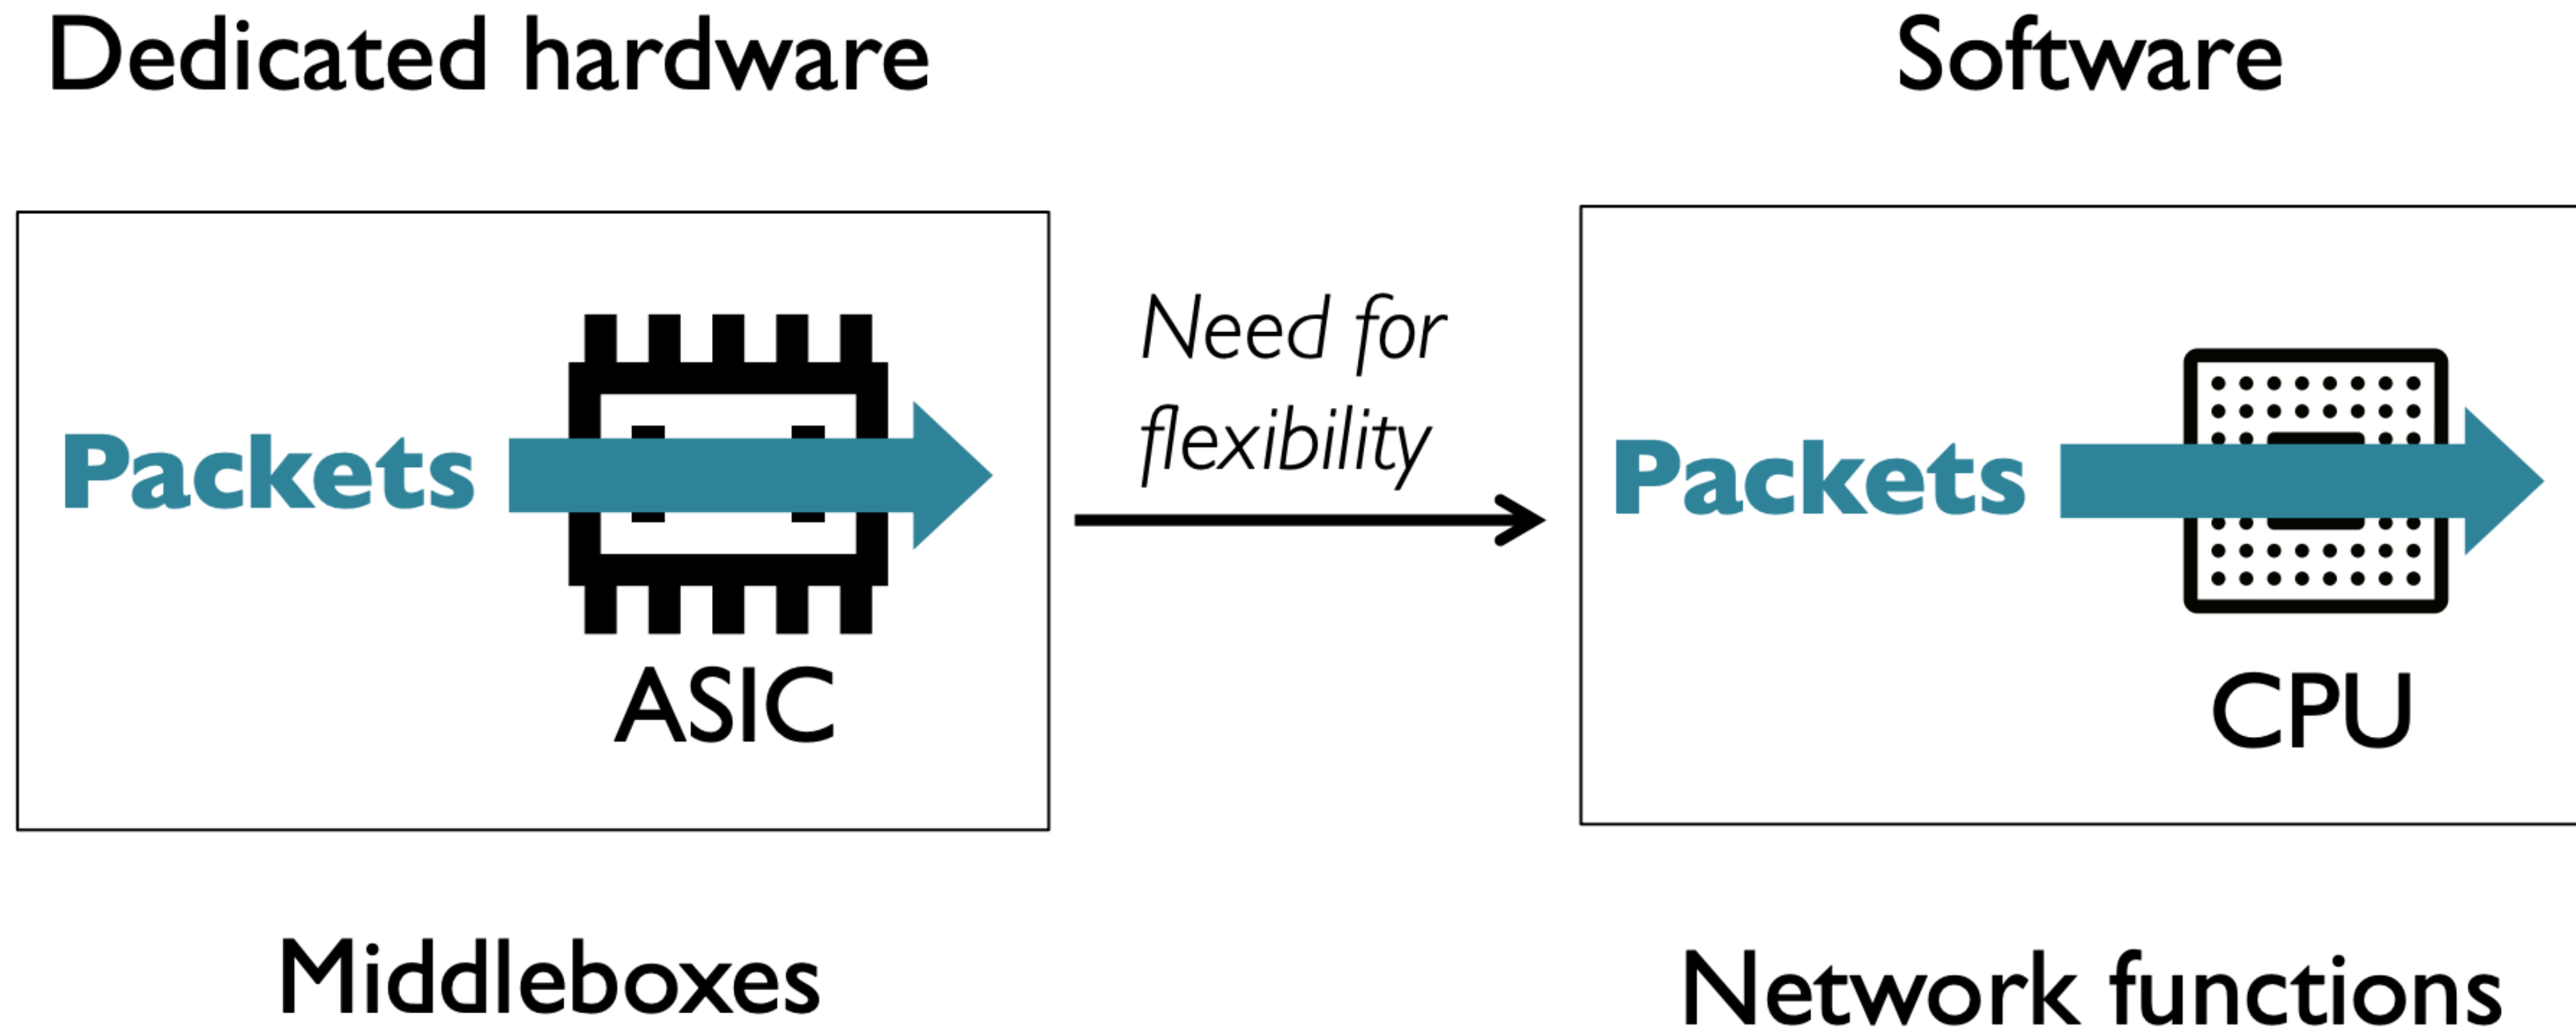
**One-third of all network devices in enterprises are middleboxes!**

[Making middleboxes someone else's problem, SIGCOMM'12]


# Problems with Hardware Middleboxes

- Dedicated
- Fixed function with little/no programmability
- Specialized hardware/software
- Custom Management APIs

# Evolution of Middleboxes: Network Function Virtualization

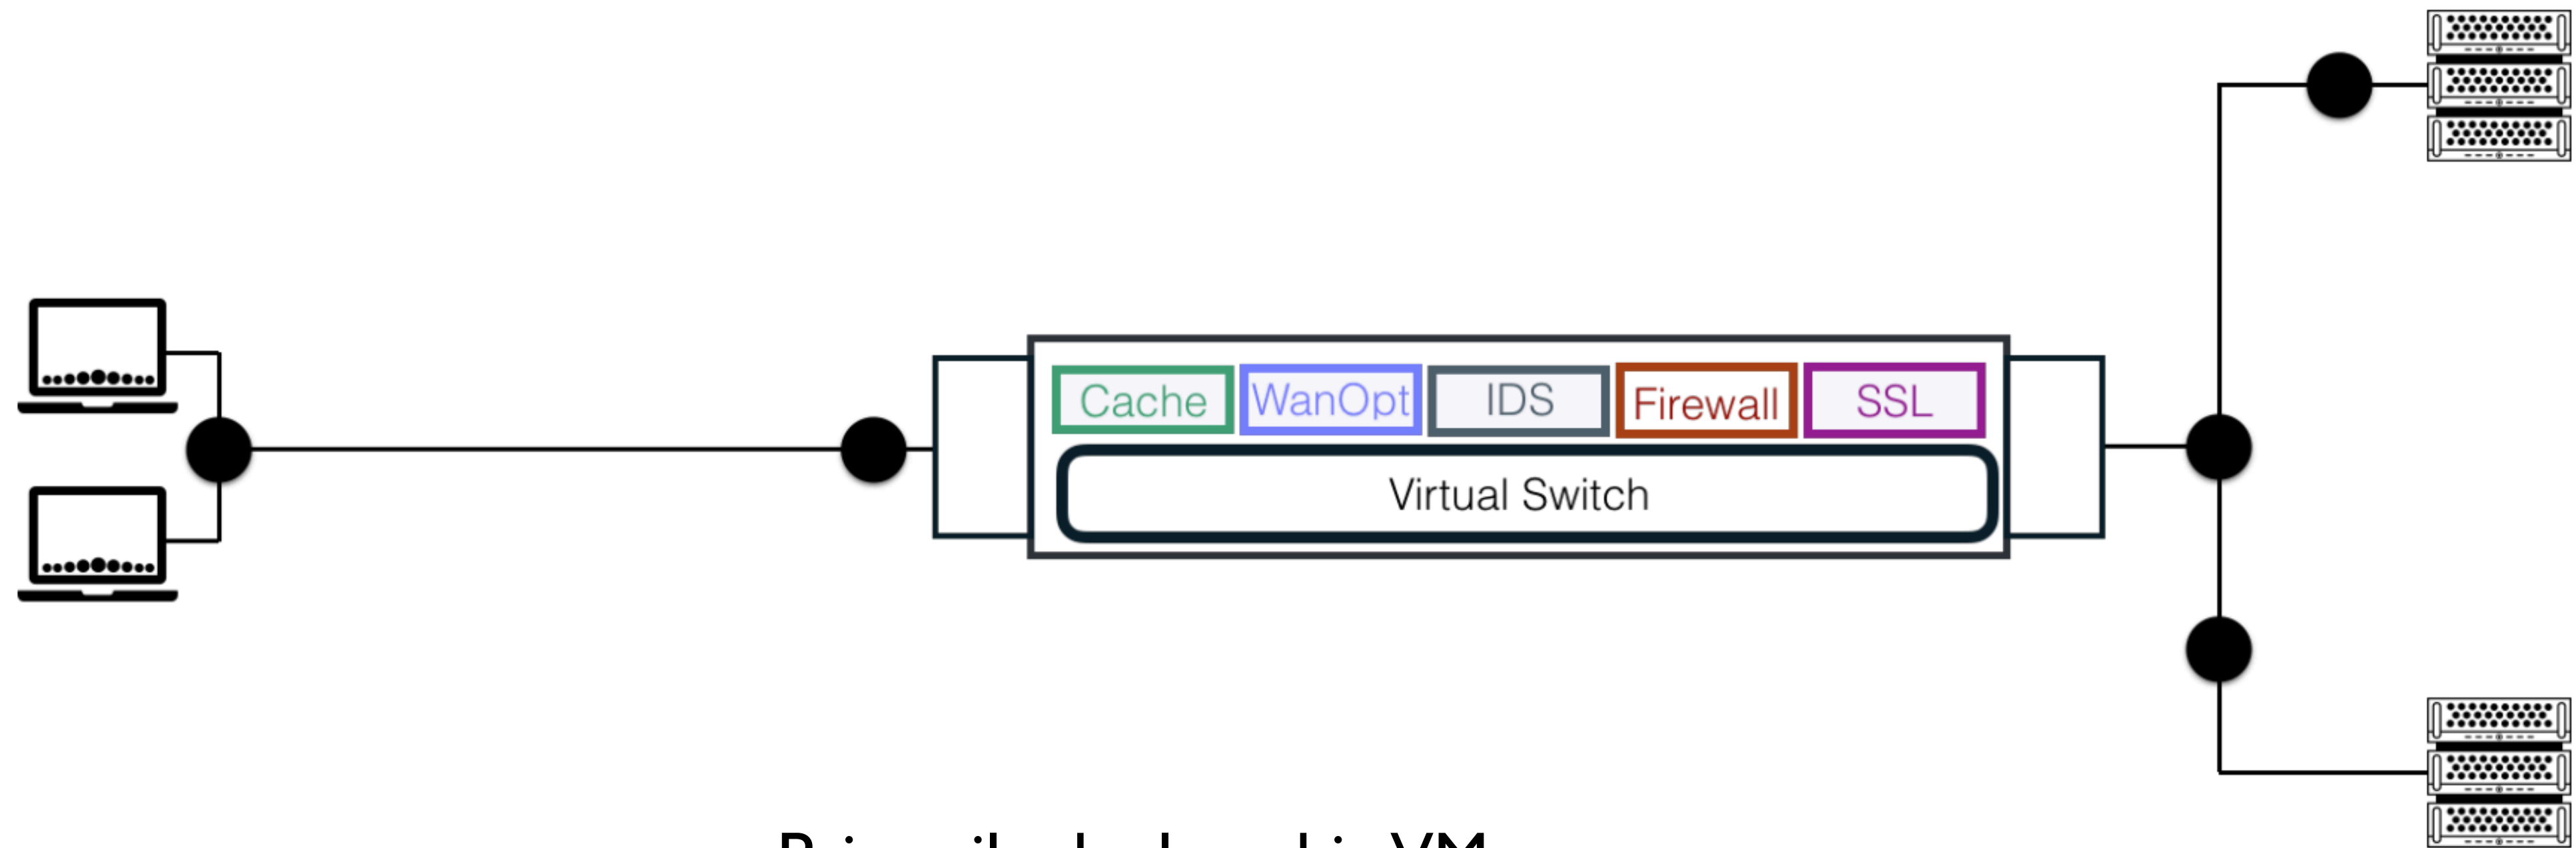


# From Hardware Middleboxes...



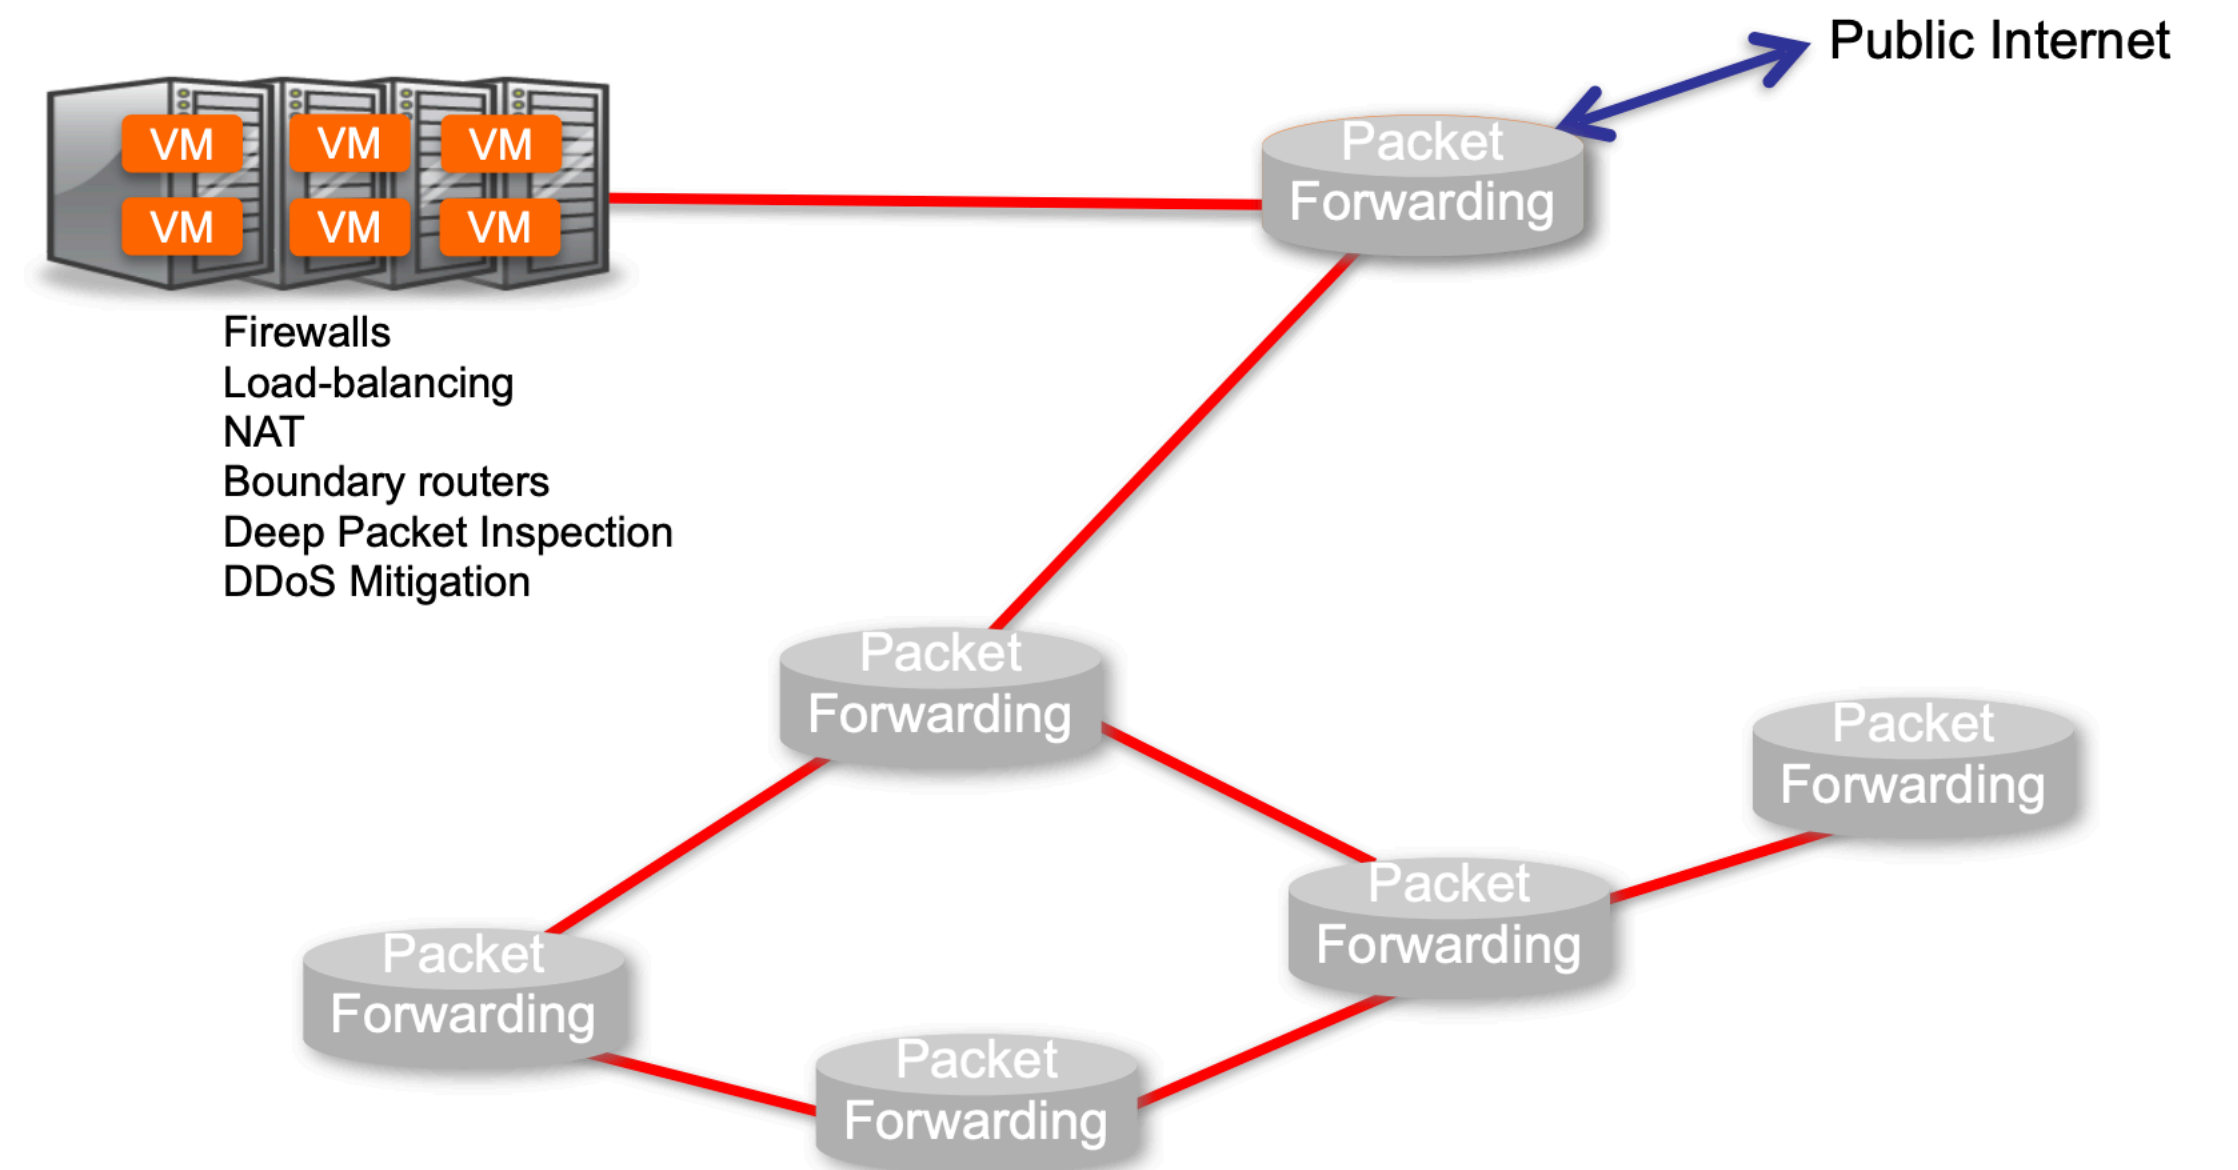
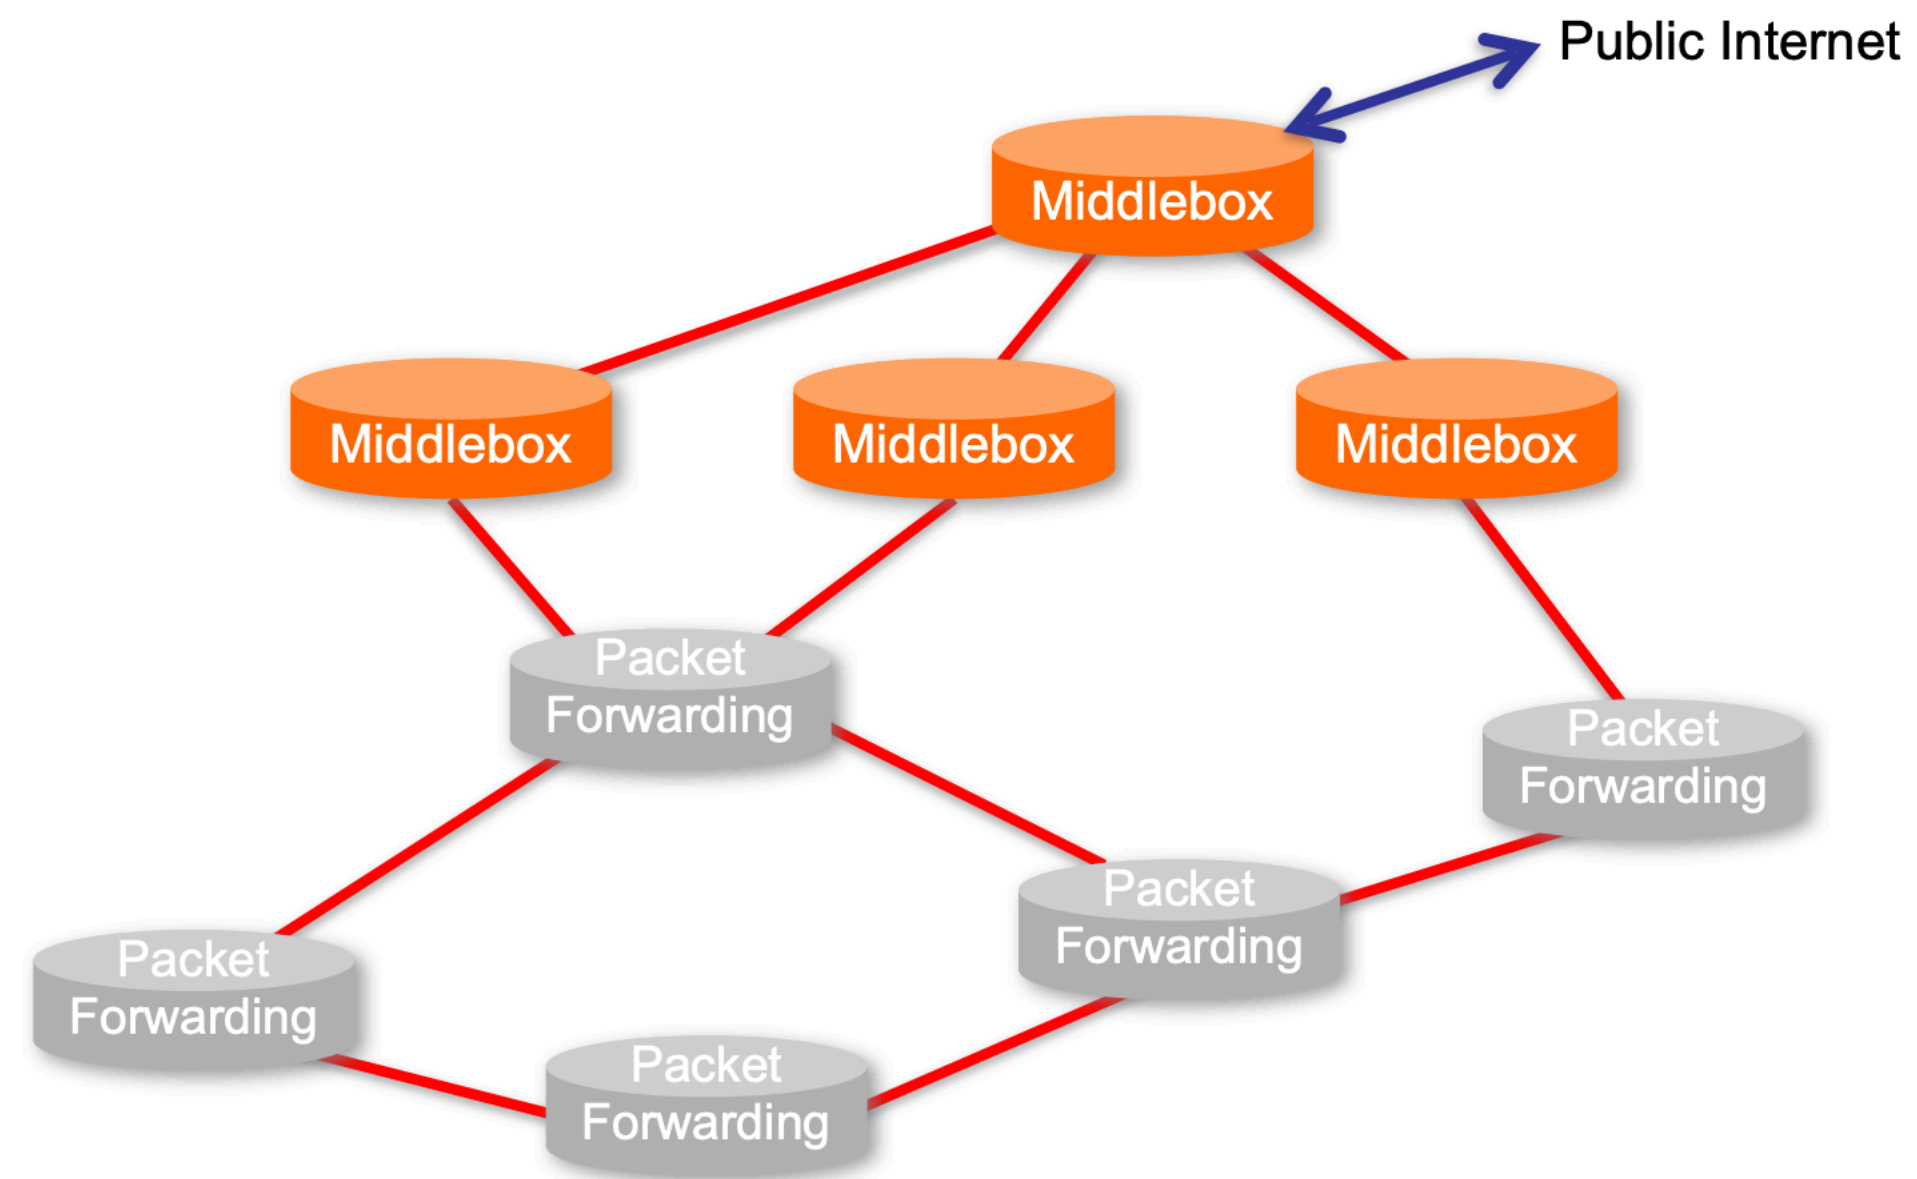


# ... to software Network Functions (NFs)





Primarily deployed in VMs

# Network View



# Key Benefits of Software Network Functions

- Programmability
  - ability to update and create new NFs
- Cost benefits of commodity solutions
- Efficiency of statistical multiplexing
- Ease of deployment, configuration, and management



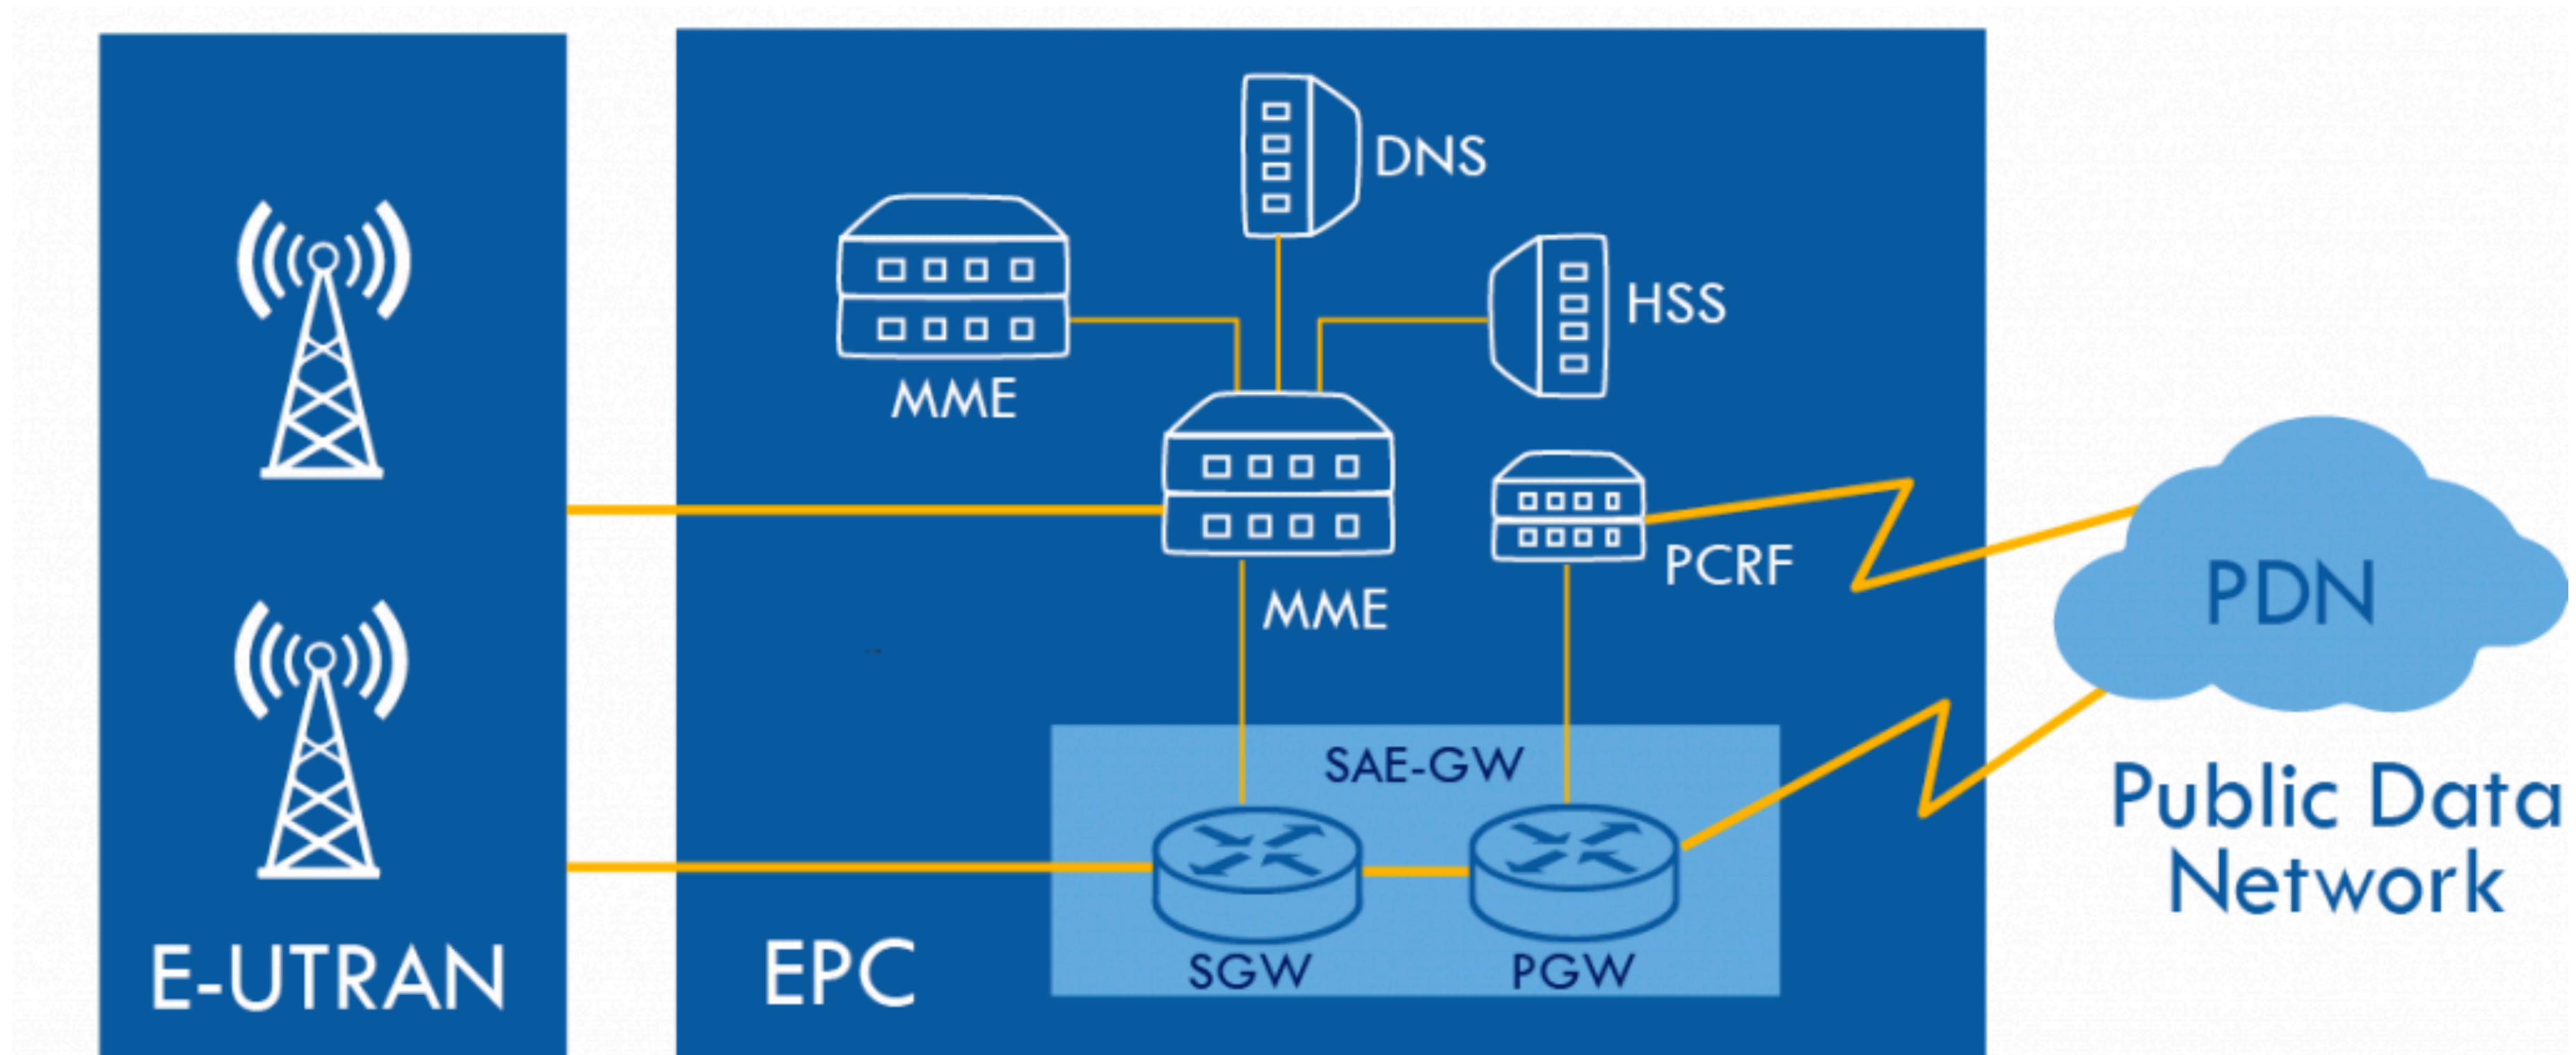
# Challenges with Network Function Virtualization

- Complex and costly state management
- Custom per-app management APIs
- Unpredictable performance
- Performance degradation

# E2: A Framework for NFV Applications

- End-to-end management of Network Functions
- Provide general solutions for common tasks
- Benefits
  - Frees NF developers to focus on NF-specific logic
  - Automates/consolidates management for operators

# Cellular Backend



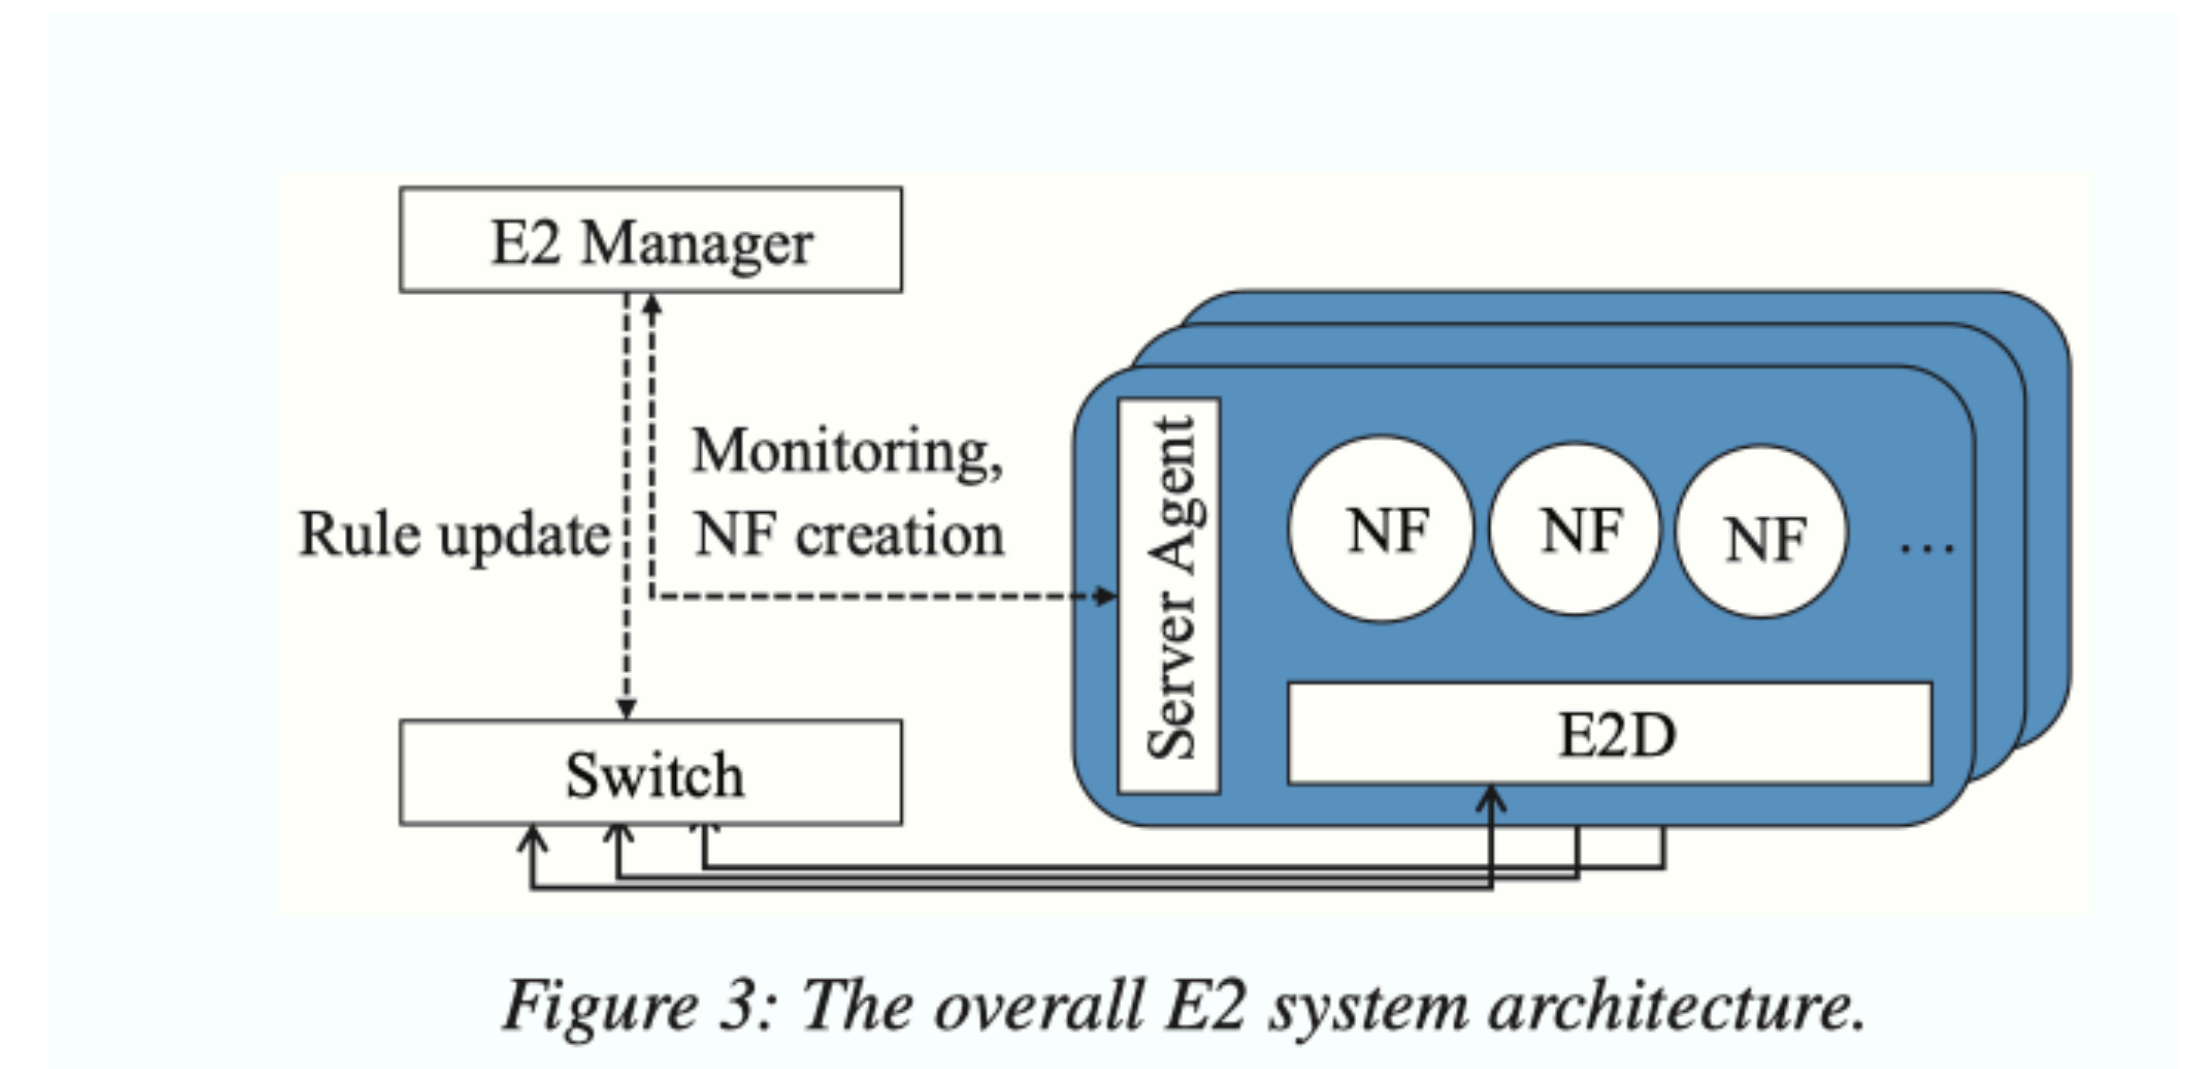
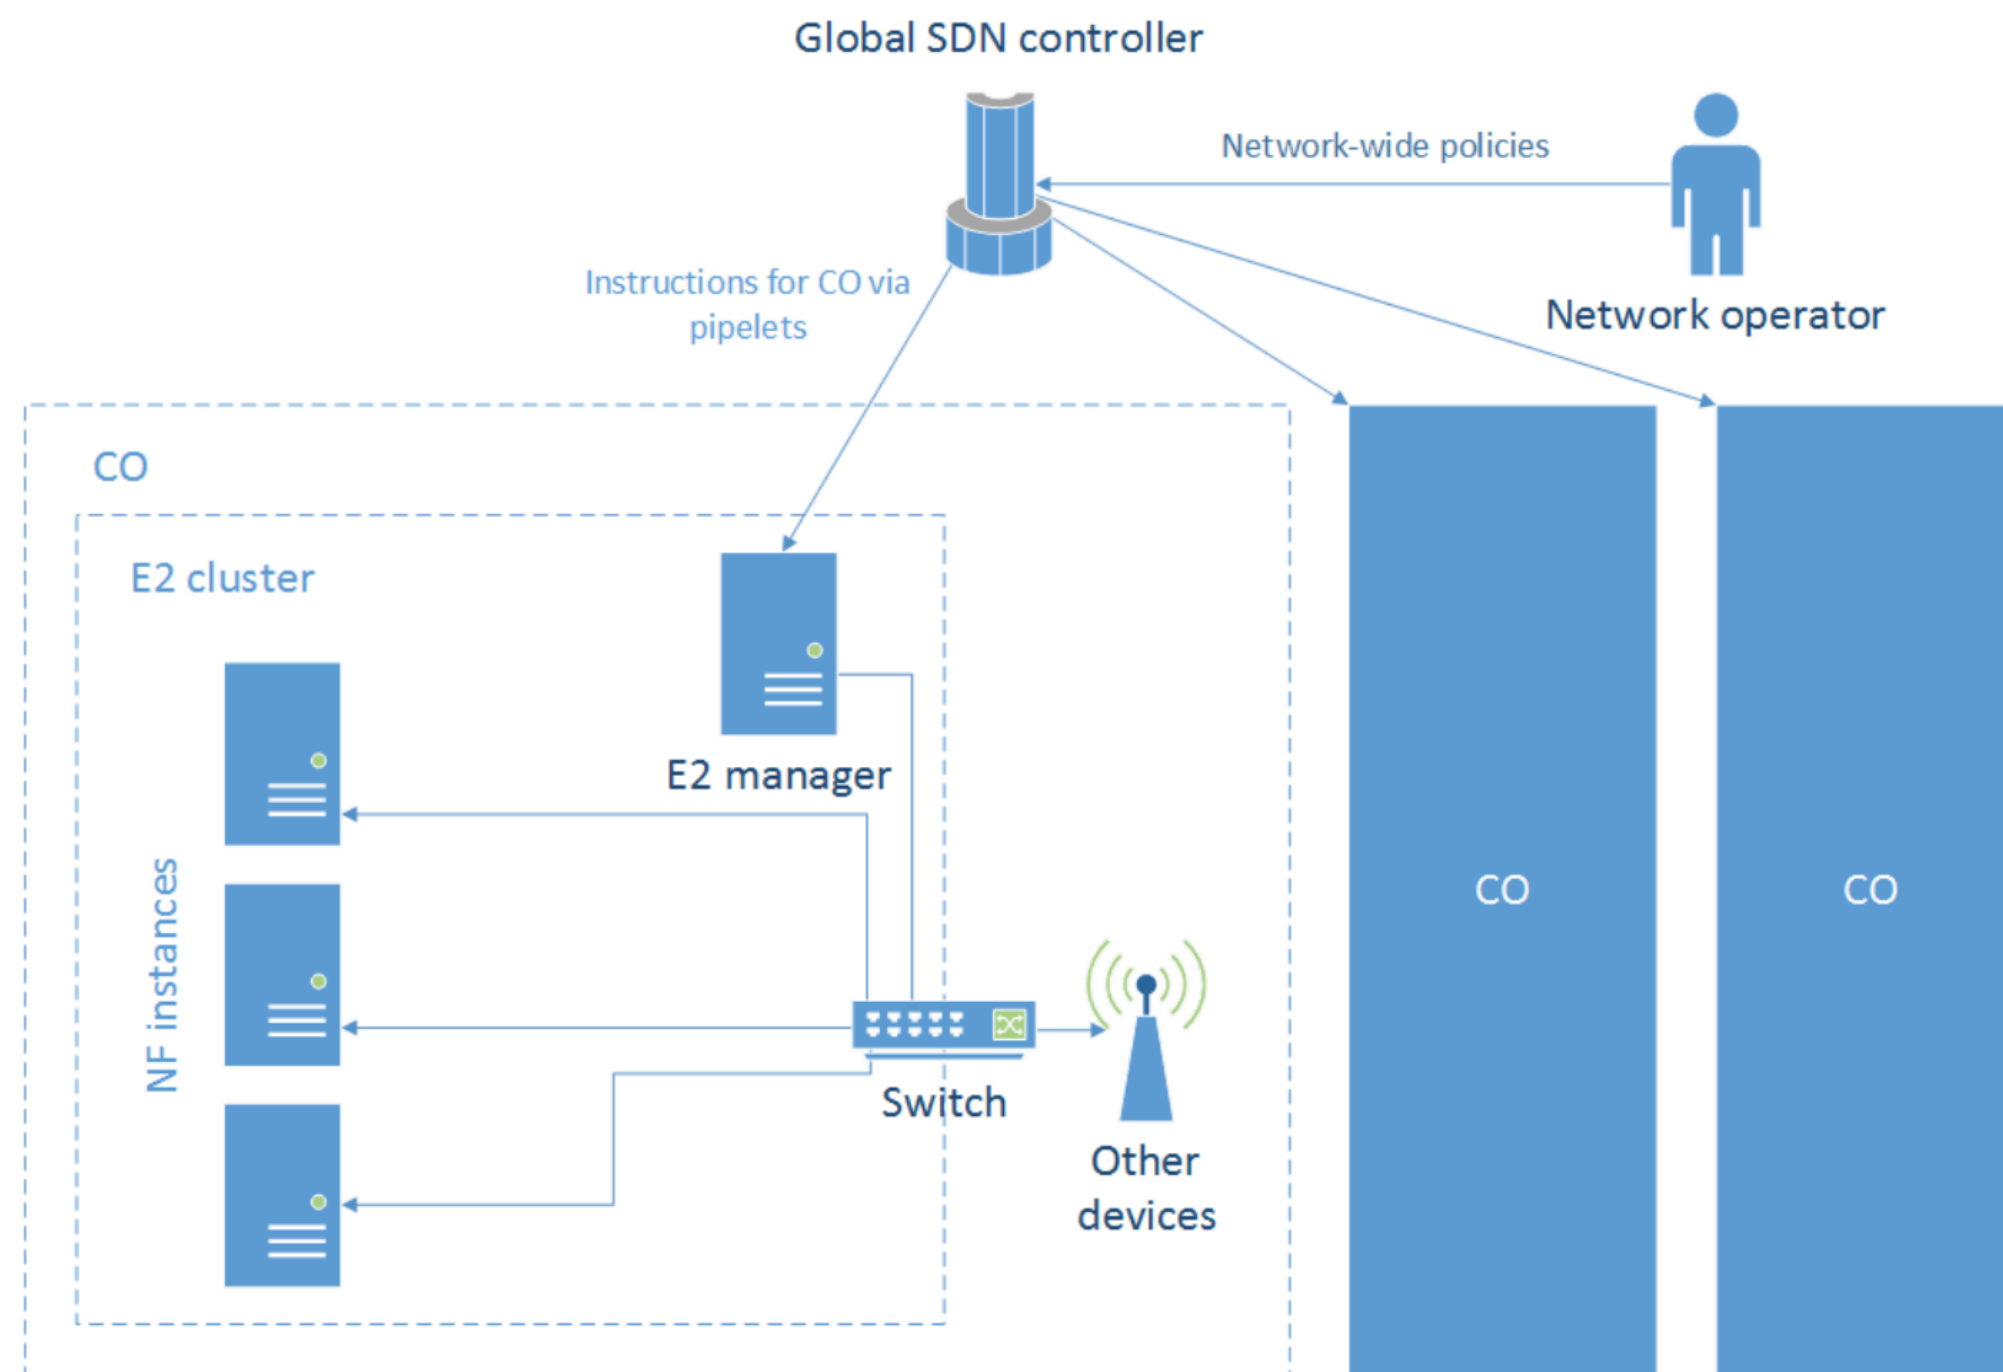
# NF Placement Options

- Thread-Based
  - Lightweight
  - No resource isolation
- Virtual Machine-Based
  - Additional overheads
  - Resource Isolation

E2 is VM-based



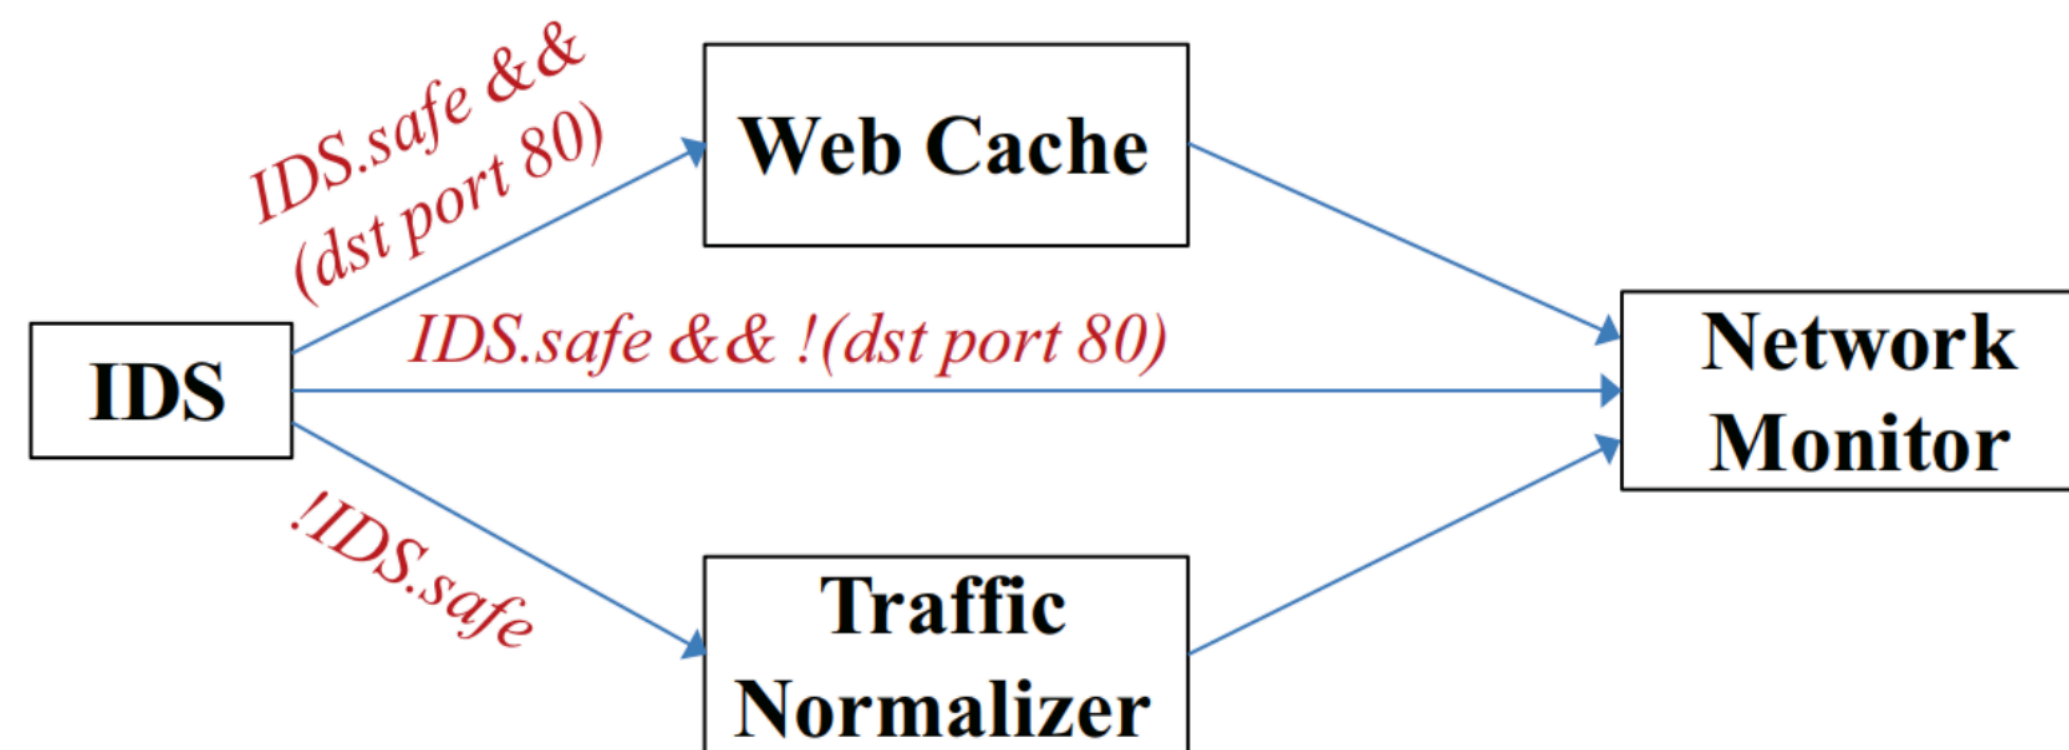
# Design Overview





# Pipelets

- NFV jobs represented as ‘pipelets’
  - a traffic class and a DAG that captures how this traffic class should be processed by NFs



# E2 Dataplane

- Modular architecture based on SoftNIC
- Highly efficient (uses Intel DPDK)
- Why OVS is not suitable?
  - expressiveness and functionality are limited by the flow-table semantics
  - performance optimizations that improve the efficiency of NFs more important in this context

# E2 Control Plane

- Executing Pipelets
  - Sizing: How many NF instances?
  - Placement: Where to place NF instances?
  - Composition: How to steer traffic between NFs?
  - Dynamic scaling: Adapting to traffic changes
  - Ensuring affinity constraints of NFs

# NF placement example

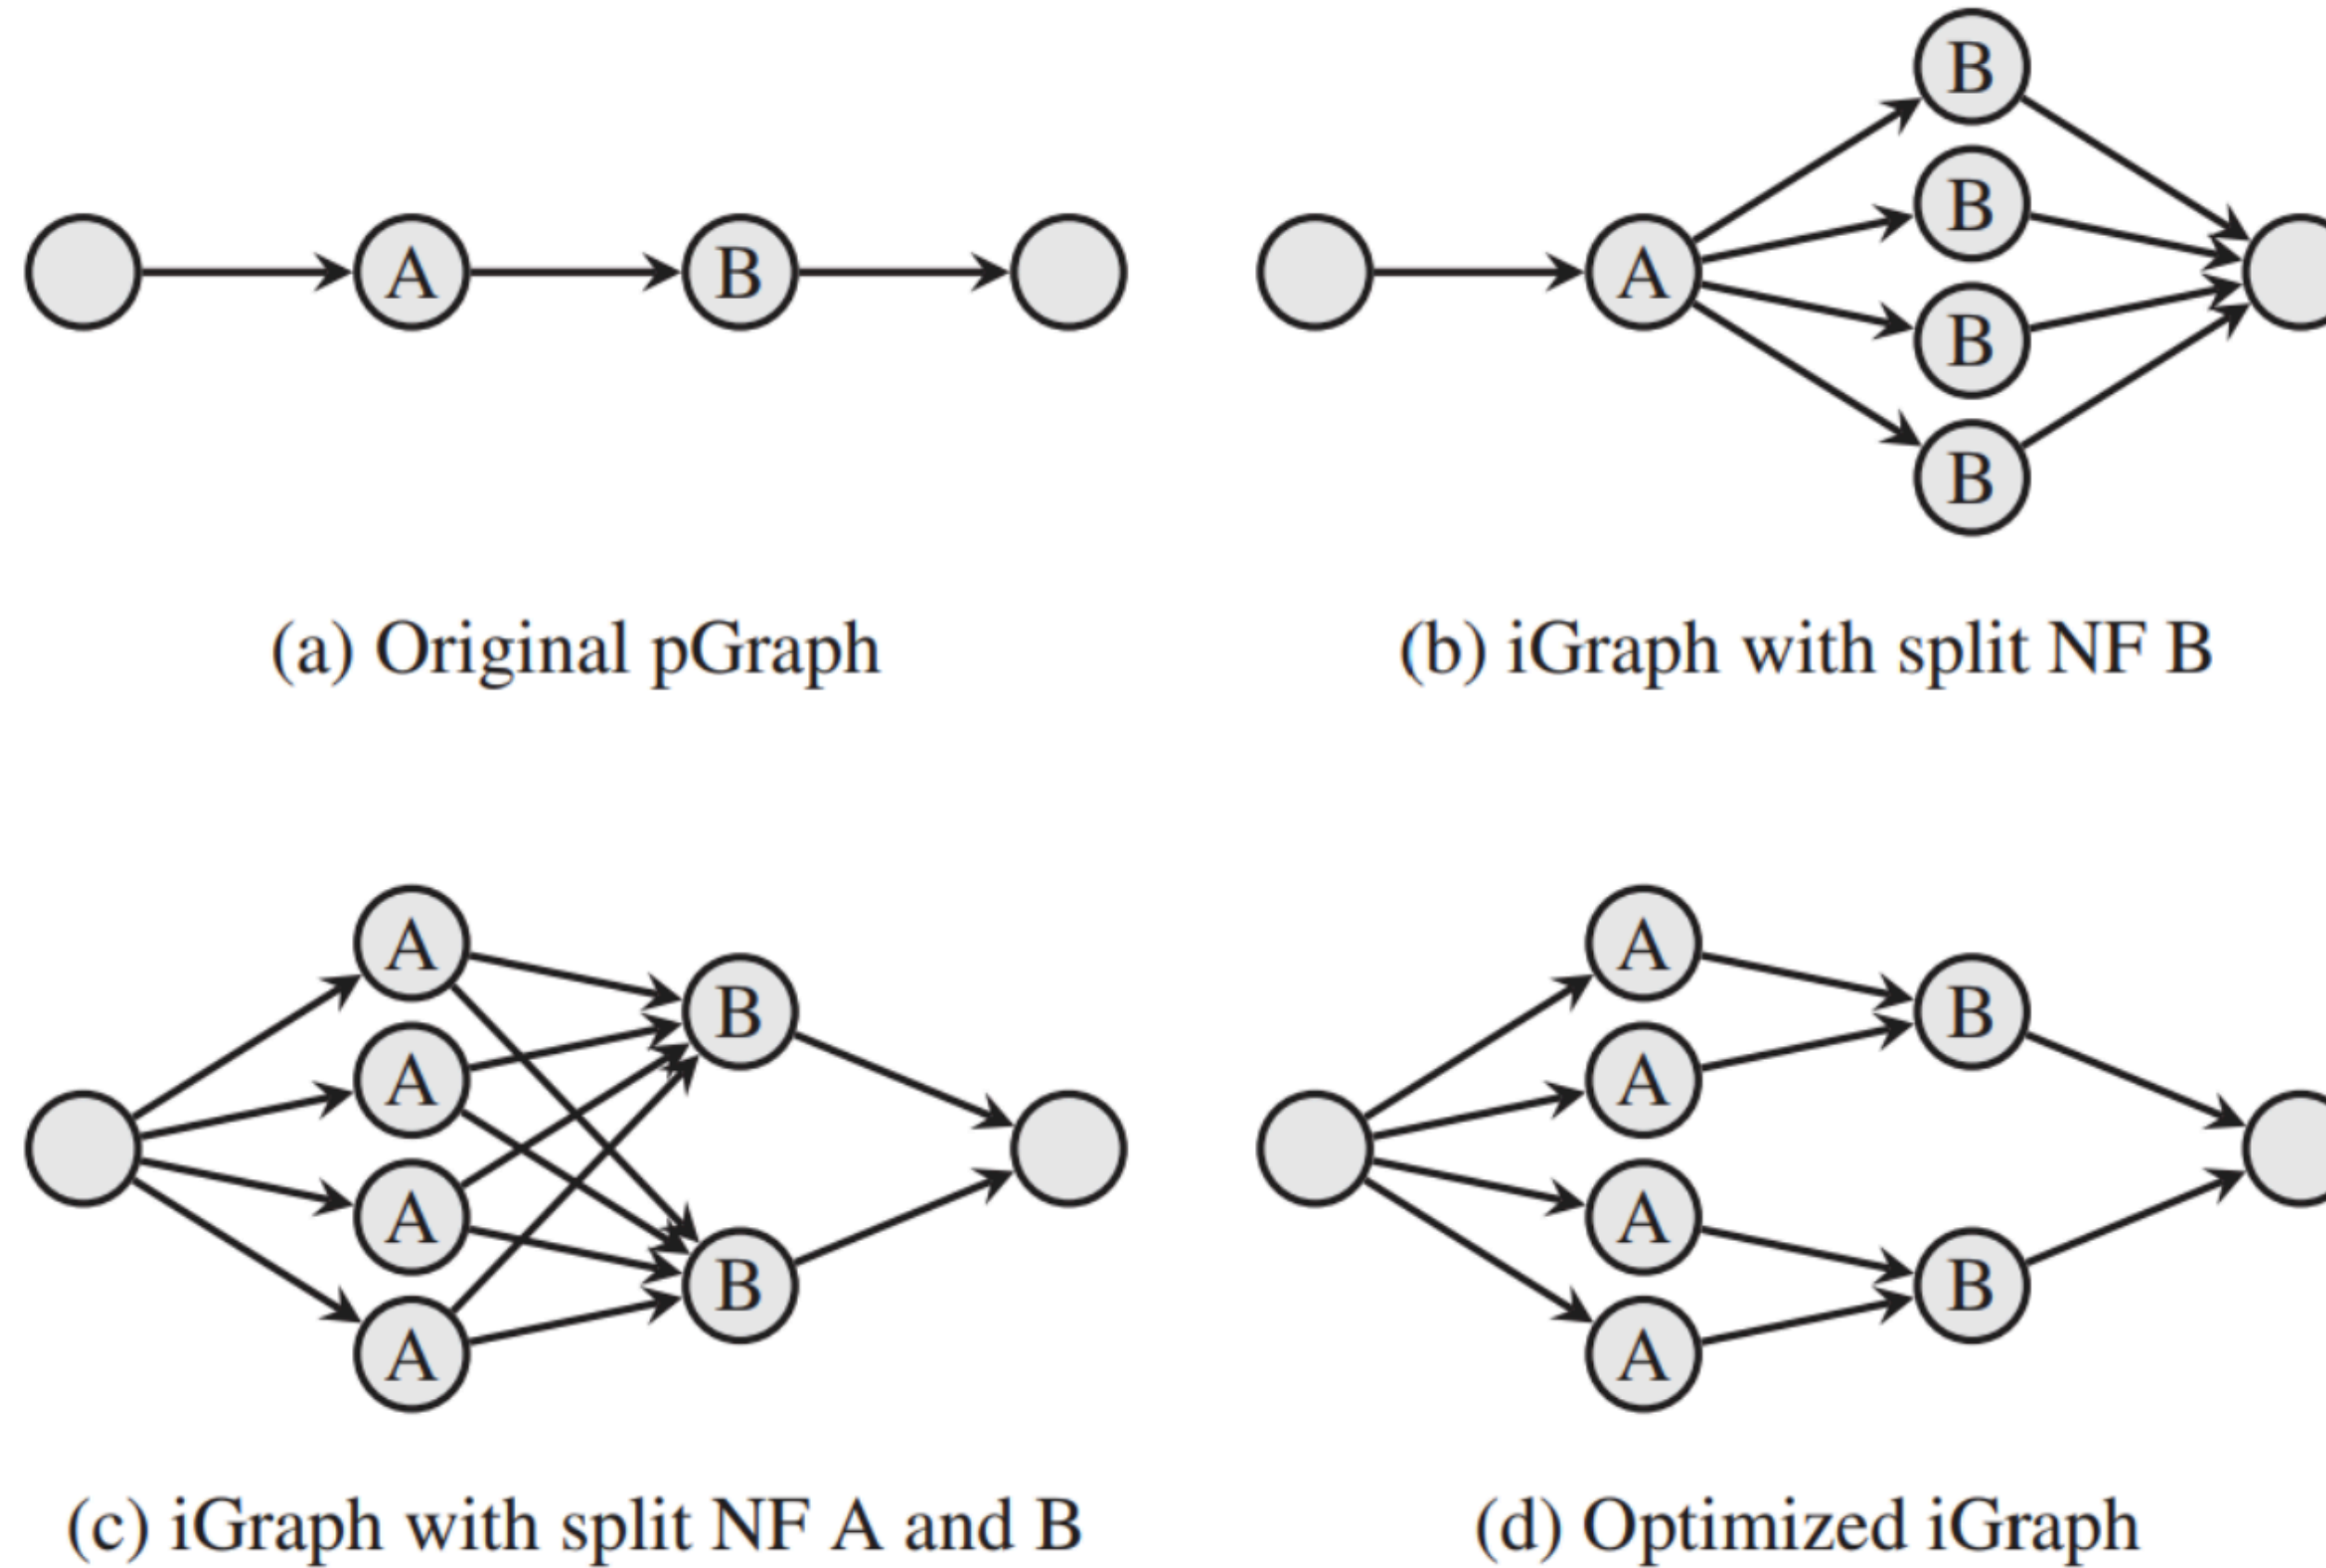


Figure 4: Transformations of a pGraph (a) into an iGraph (b, c, d).

# Comments from students

- Move E2 to container-based implementation - multiple students
- Single point of failure - multiple students
- “There are certain hardware constraints that E2 takes into account. More work is needed to figure out how to exploit richer resources like CPU cache, GPUs, programmable switches, specialized accelerators, etc.” - Rakshit Mehra
- “The paper does not address consistency issues that arise when global or aggregate state is spread across multiple NF instances, which could be a significant challenge for managing cross-NF state in a dynamic scaling scenario.” - Sagar Krishna
- “Future work on the E2 framework should focus on exploring fault-tolerance and energy-efficient management and monitoring to enhance its robustness and sustainability in real-world deployments.” - Yurun Song



# High Performance NF Implementations

## Microboxes: High Performance NFV with Customizable, Asynchronous TCP Stacks and Dynamic Subscriptions

Guyue Liu\*, Yuxin Ren\*, Mykola Yurchenko\*,  
K.K. Ramakrishnan<sup>†</sup>, Timothy Wood\*

\*George Washington University, †University of California, Riverside

## FlowBlaze: Stateful Packet Processing in Hardware

Salvatore Pontarelli<sup>1,2</sup>, Roberto Bifulco<sup>3</sup>, Marco Bonola<sup>1,2</sup>, Carmelo Cascone<sup>4</sup>,  
Marco Spaziani<sup>2,5</sup>, Valerio Bruschi<sup>2,5</sup>, Davide Sanvito<sup>6</sup>, Giuseppe Siracusano<sup>3</sup>,  
Antonio Capone<sup>6</sup>, Michio Honda<sup>3</sup>, Felipe Huici<sup>3</sup> and Giuseppe Bianchi<sup>2,5</sup>

<sup>1</sup>Axbryd, <sup>2</sup>CNIT, <sup>3</sup>NEC Laboratories Europe, <sup>4</sup>Open Networking Foundation, <sup>5</sup>University of Rome Tor Vergata, <sup>6</sup>Politecnico di Milano

# mOS: A Reusable Networking Stack for Flow Monitoring Middleboxes

Muhammad Jamshed, YoungGyoun Moon, Donghwi Kim, Dongsu Han, and KyoungSoo Park  
School of Electrical Engineering, KAIST

## ClickNP: Highly Flexible and High Performance Network Processing with Reconfigurable Hardware

Bojie Li<sup>§†</sup>   Kun Tan<sup>†</sup>   Layong (Larry) Luo<sup>‡</sup>   Yanqing Peng<sup>•†</sup>   Renqian Luo<sup>§†</sup>

Ningyi Xu<sup>†</sup> Yongqiang Xiong<sup>†</sup> Peng Cheng<sup>†</sup> Enhong Chen<sup>§</sup>

<sup>†</sup>Microsoft Research    <sup>§</sup>USTC    <sup>‡</sup>Microsoft    •SJTU

## NetBricks: Taking the V out of NFV

Aurojit Panda<sup>†</sup> Sangjin Han<sup>†</sup> Keon Jang<sup>‡</sup> Melvin Walls<sup>†</sup> Sylvia Ratnasamy<sup>†</sup> Scott Shenker<sup>†\*</sup>

† UC Berkeley ‡ Google \* ICSI

## Abstract

The move from hardware middleboxes to software network functions, as advocated by NFV, has proven more challenging than expected. Developing new NFs remains a tedious process, requiring that developers repeatedly rediscover and reapply the same set of optimizations, while current techniques for providing isolation between NFs (using VMs or containers) incur high performance overheads. In this paper we describe NetBricks, a new NFV framework that tackles both these problems. For building NFs we take inspiration from modern data analytics frameworks (*e.g.*, Spark and Dryad) and build a small set of customizable network processing elements. We also embrace type checking and safe runtimes to provide isolation in software, rather than rely on hardware isolation. NetBricks provides the same memory isolation as containers and VMs, without

standard tools for managing VMs; (c) faster development, which now requires writing software that runs on commodity hardware; and (d) reduced costs by consolidating several NFs on a single machine. However, despite these promised advances, there has been little progress towards large-scale NF deployments. Our discussions with three major carriers revealed that they are only just beginning small scale test deployments (with 10-100s of customers) using simple NFs *e.g.*, firewalls and NATs.

The move from hardware middleboxes to software NFs was supposed to speed innovation, so why has progress been so slow? We believe this delay is because traditional approaches for both *building* and *running* NFs are a poor match for carrier networks, which have the following requirements: *performance*, NF deployments should be able to provide per-packet latencies on the order of 10s of  $\mu$ s, and *throughput* on the order of 10s of Gbps; *efficiency*,

## ABSTRACT

flexible software network functions (NFs) are crucial to enable multi-tenancy in the clouds. Hardware packet processing on a commodity server has capacity and induces high latency. While software could scale out using more servers, doing so adds significant cost. This paper focuses on accelerating NFs with programmable hardware, *i.e.*, FPGA, which is now a mainstream hardware and inexpensive for datacenters. However, it is predominately programmed using low-level hardware description languages (HDLs), which are hard to code and difficult to debug. More importantly, HDLs are almost unusable for most software programmers. This paper presents ClickNP, a FPGA-accelerated platform for highly flexible high-performance NFs with commodity servers. ClickNP is as flexible as it is completely programmable using high-level C-like languages, and exposes a modular programming abstraction that resembles Click Modular Router. ClickNP achieves high performance. Our prototype NFs show that they process high traffic at up to 200 million packets per second at a low latency ( $< 2\mu s$ ). Compared to existing software implementations, with FPGA, ClickNP improves throughput by 10x, while reducing latency by 10x. To the best of our knowledge, ClickNP is the first FPGA-accelerated platform for NFs, written completely in high-level language and achieving 40 Gbps line rate at any packet size.

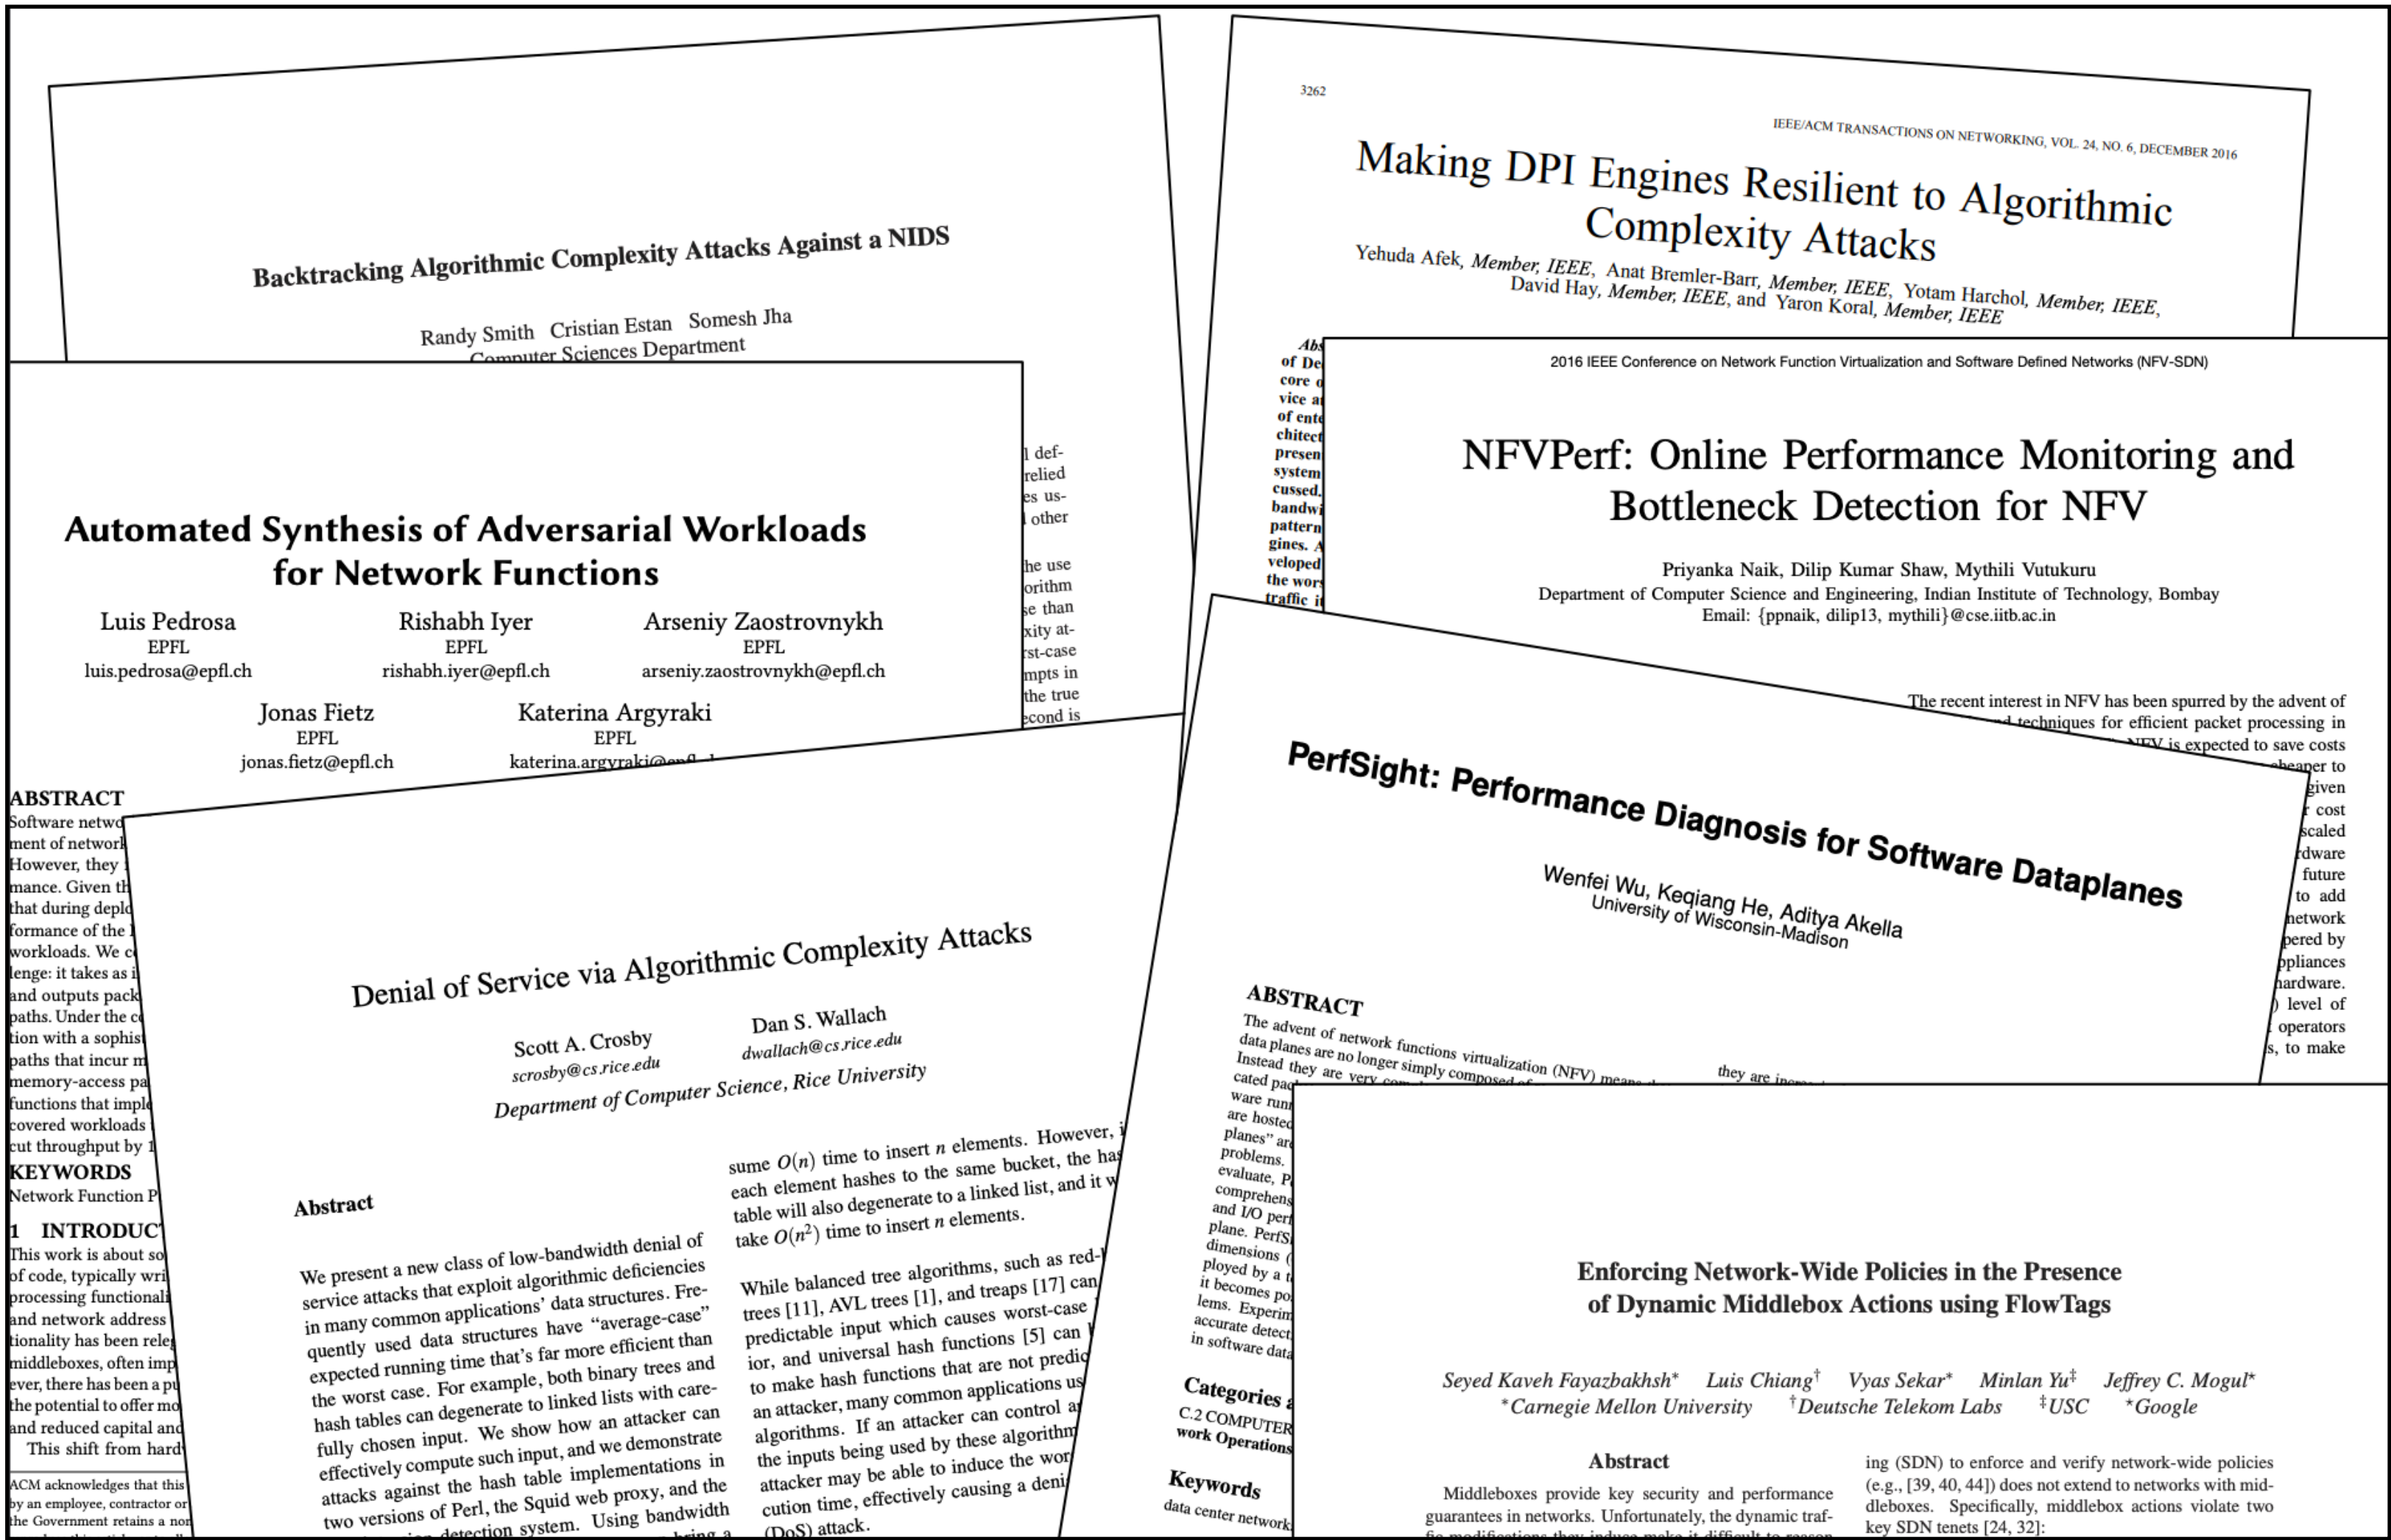
## 1. INTRODUCTION

Modern multi-tenant datacenters provide shared infrastructure for hosting many different types of services from different customers (*i.e.*, tenants) at a low cost. To ensure security and performance isolation, each tenant is deployed in a *virtualized network* environment. Flexible network functions (NFs) are required for datacenter operators to enforce isolation while simultaneously guaranteeing Service Level Agreements (SLAs).

Conventional hardware-based network appliances are not flexible, and almost all existing cloud providers, *e.g.*, Microsoft, Amazon and VMware, have been deploying software based NFs on servers to maximize the flexibility [23, 30]. However, software NFs have two fundamental limitations – both stem from the nature of software packet processing. First, processing packets in software has limited capacity. Existing software NFs usually require multiple cores to achieve 10 Gbps rate [33, 43]. But the latest network links have scaled up to 40~100 Gbps [11]. Although one could add more cores in a server, doing so adds significant cost, not only in terms of capital expense, but also more operational expense as they are burning significantly more energy. Second, processing packets in software incurs large, and highly variable latency. This latency may range from tens of microsecond to milliseconds [22, 33, 39]. For many low latency applications (*e.g.*, stock trading), this inflated latency is un-



# Understanding NF Performance





# Providing Guarantees about NF Behavior





# State Management

**Split/Merge: System Support for Elastic Execution in Virtual Middleboxes**

Shriram Rajagopalan<sup>†‡</sup>, Dan Williams<sup>†</sup>, Hani Jamjoom<sup>†</sup>, and Andrew Warfield<sup>‡</sup>

<sup>†</sup>IBM T. J. Watson Research Center, Yorktown Heights, NY  
<sup>‡</sup>University of British Columbia, Vancouver, Canada

**Elastic Scaling of Stateful Network Functions**

Shinae Woo<sup>\*†</sup>, Justine Sherry<sup>‡</sup>, Sangjin Han<sup>\*</sup>, Sue Moon<sup>†</sup>, Sylvia Ratnasamy<sup>\*</sup>, and Scott Shenker<sup>\*§</sup>

<sup>\*</sup>University of California, Berkeley <sup>†</sup>KAIST <sup>‡</sup>CMU <sup>§</sup>ICSI

**Stateless Network Functions: Breaking the Tight Coupling of State and Processing**

Murad Kablan, Azzam Alsudais, Eric Keller  
 University of Colorado, Boulder

Franck Le  
 IBM Research

**Rollback-Recovery for Middleboxes**

Justine Sherry<sup>\*</sup>, Peter Xiang Gao<sup>\*</sup>, Soumya Basu<sup>\*</sup>, Aurojit Panda<sup>\*</sup>,  
 Arvind Krishnamurthy<sup>\*</sup>, Christian Maciocco<sup>†</sup>, Maziar Manesh<sup>†</sup>, João Martins<sup>‡</sup>,  
 Sylvia Ratnasamy<sup>\*</sup>, Luigi Rizzo<sup>†</sup>, Scott Shenker<sup>\*‡</sup>

<sup>\*</sup>UC Berkeley <sup>•</sup>University of Washington <sup>†</sup>Intel Research <sup>‡</sup>IBM Research

**OpenBox: A Software-Defined Framework for Developing, Deploying, and Managing Network Functions**

Anat Bremler-Barr<sup>\*</sup>, Yotam Harchol<sup>†</sup>, David Hay<sup>†</sup>  
 bremler@idc.ac.il yotamhc@cs.huji.ac.il dhay@cs.huji.ac.il

<sup>\*</sup>School of Computer Science <sup>†</sup>School of Computer Science  
<sup>‡</sup>The Interdisciplinary Center, Herzliya, Israel

**Pico Replication: A High Availability Framework for Middleboxes**

Shriram Rajagopalan<sup>†‡</sup>, Dan Williams<sup>†</sup>, Hani Jamjoom<sup>†</sup>

<sup>†</sup>IBM T. J. Watson Research Center, Yorktown Heights, NY  
<sup>‡</sup>University of British Columbia, Vancouver, Canada

**Paving the Way for NFV: Simplifying Middlebox Modifications using StateAlyzr**

Junaid Khalid, Aaron Gember-Jacobson, Roney Michael,  
 Anubhavnidhi Abhashkumar, Aditya Akella  
 University of Wisconsin-Madison

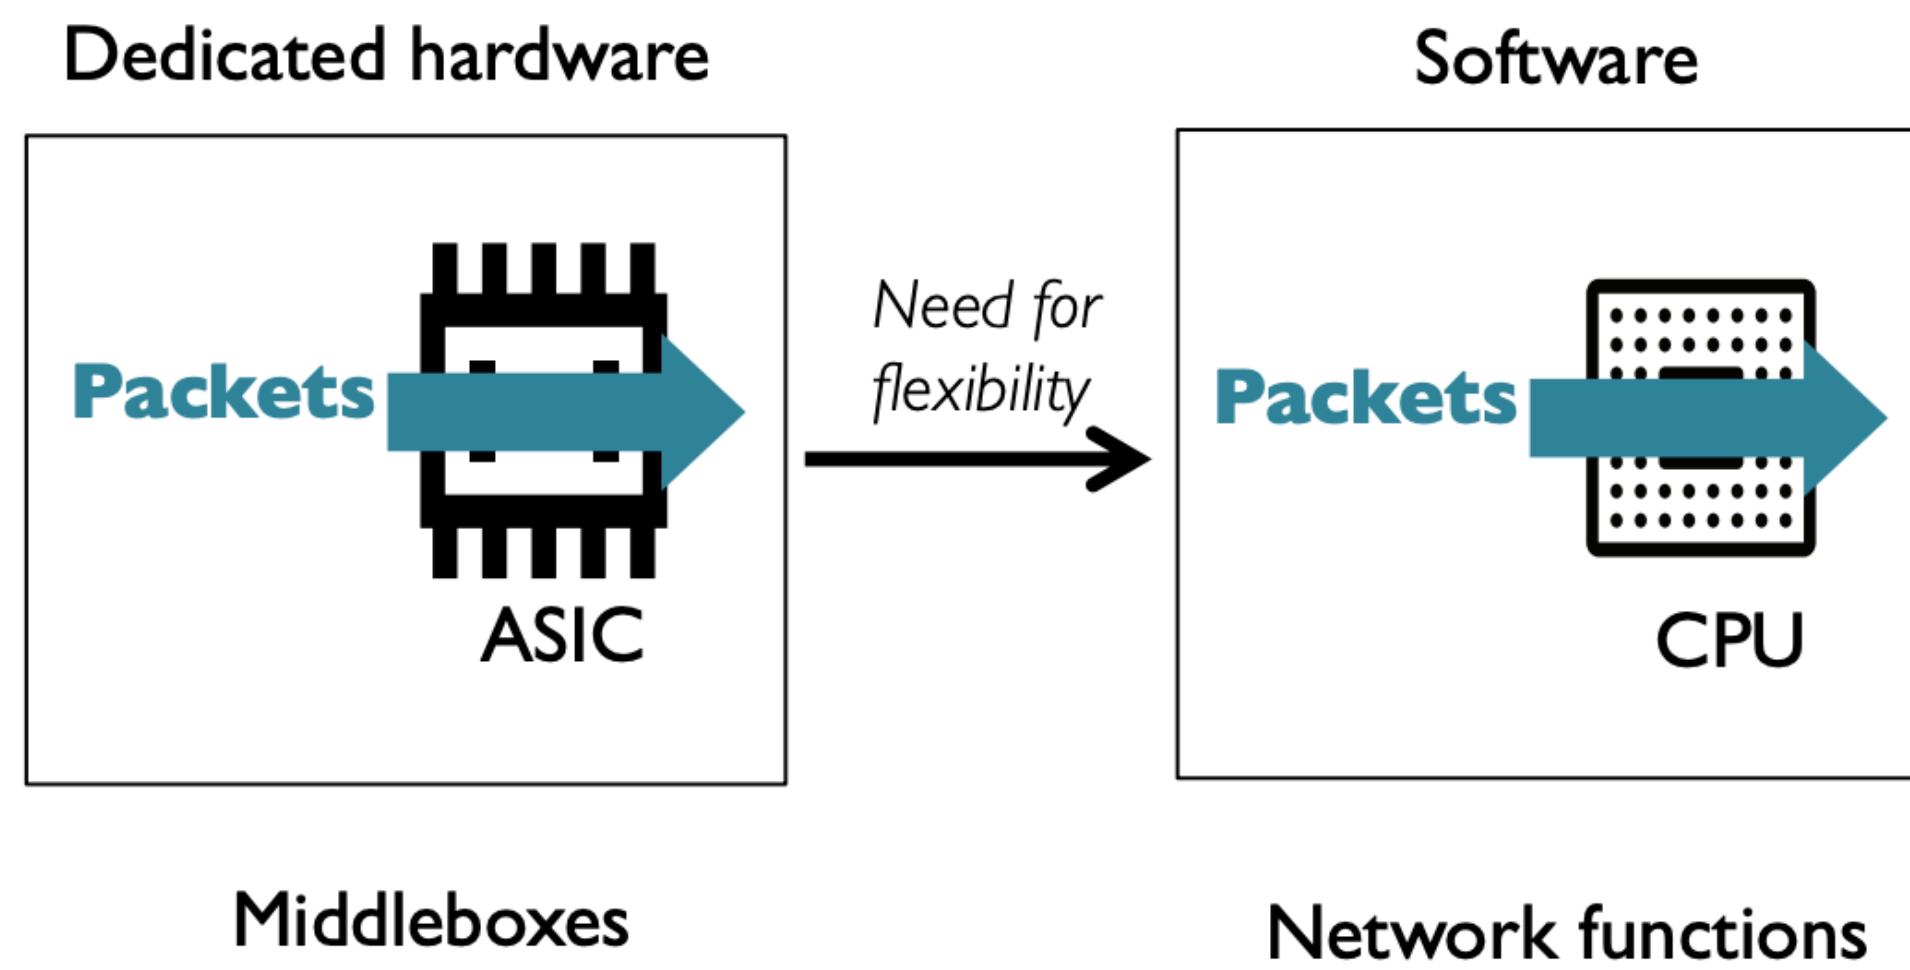
**E2: A Framework for NFV Applications**

Shoumik Palkar<sup>\*</sup>, Chang Lan<sup>\*</sup>, Sangjin Han  
 UC Berkeley UC Berkeley UC Berkeley  
 sppalkar@berkeley.edu clan@eecs.berkeley.edu sangjin@eecs.berkeley.edu

Keon Jang, Aurojit Panda, Sylvia Ratnasamy  
 Intel Labs UC Berkeley UC Berkeley  
 keon.jang@intel.com apanda@cs.berkeley.edu sylvia@eecs.berkeley.edu

Luigi Rizzo, Scott Shenker  
 Università di Pisa UC Berkeley and ICSI  
 rizzo@iet.unipi.it shenker@icsi.berkeley.edu

# Evolution of Middleboxes



Thanks!