



SDN@NIT
The Power of Softwarization

NetCloud Hong Kong
2017/12/11

PA-Flow:

Gradual Packet Aggregation at Virtual Network I/O for Efficient Service Chaining

Yuki Taguchi[†], Ryota Kawashima[†], Hiroki Nakayama^{††},
Tsunemasa Hayashi^{††}, and Hiroshi Matsuo[†]

[†] Nagoya Institute of Technology, Japan

^{††} BOSCO Technologies Inc.

Boosting per-flow performance of service chains

Forwarding performance of NFV-nodes is poor.

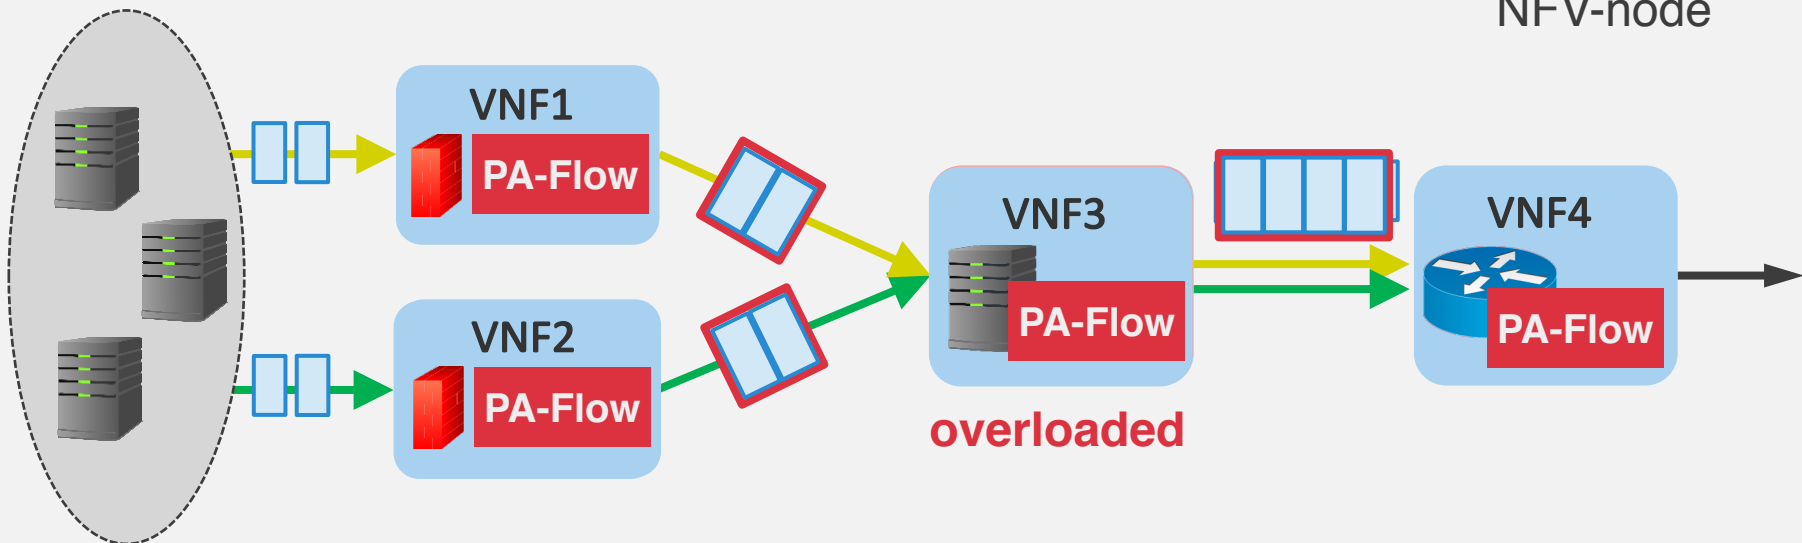
Virtual Network Function

VM or Container

VNF



NFV-node



1 Background

- Problems of Service Chaining

2 Related Work

- Scaling-out
- Scaling-up

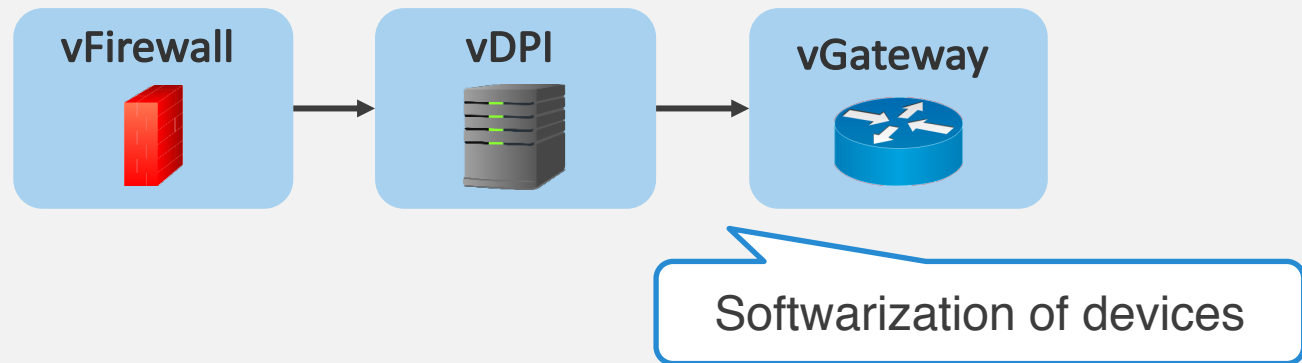
3 Proposal

4 Evaluation

5 Conclusion

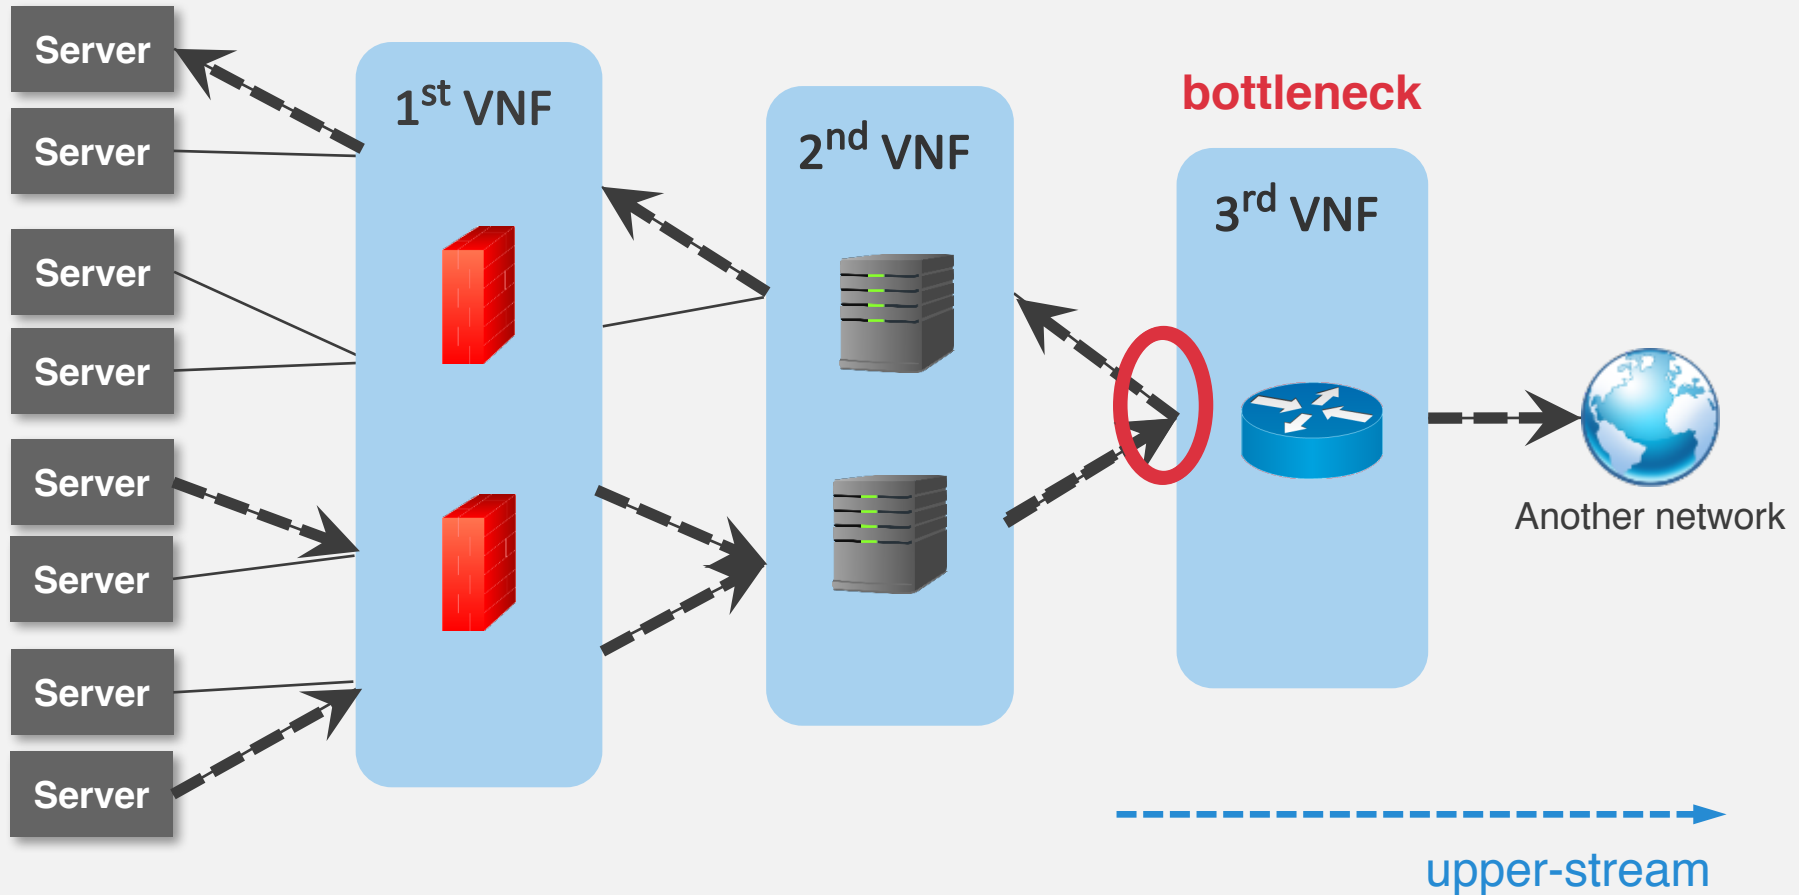
- NFV services have been launched
- Service chaining is a key concept of NFV

Toy-blocking style service chain

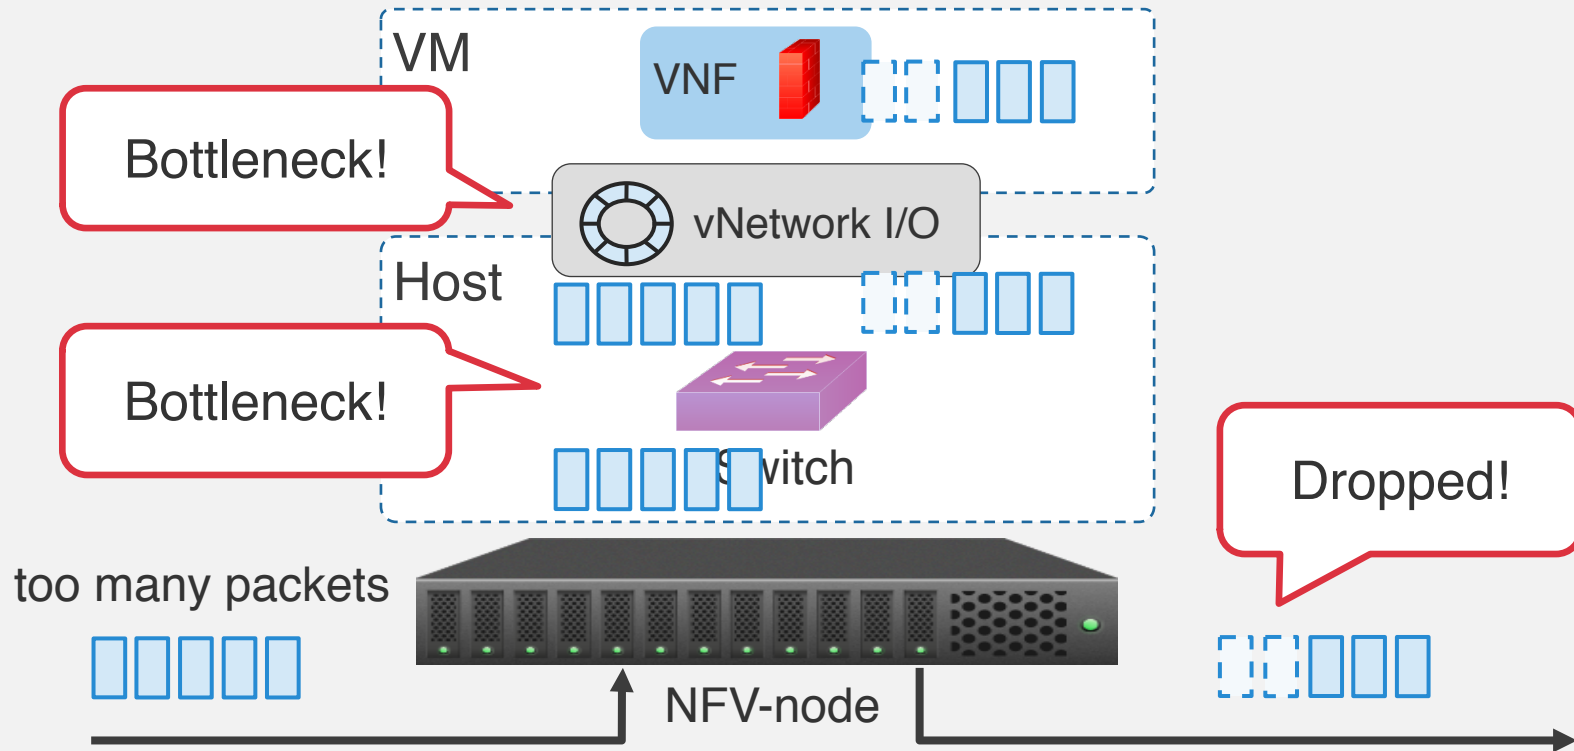


Softwarized devices degrade the performance

- Network-Level problems
- Node-Level problems

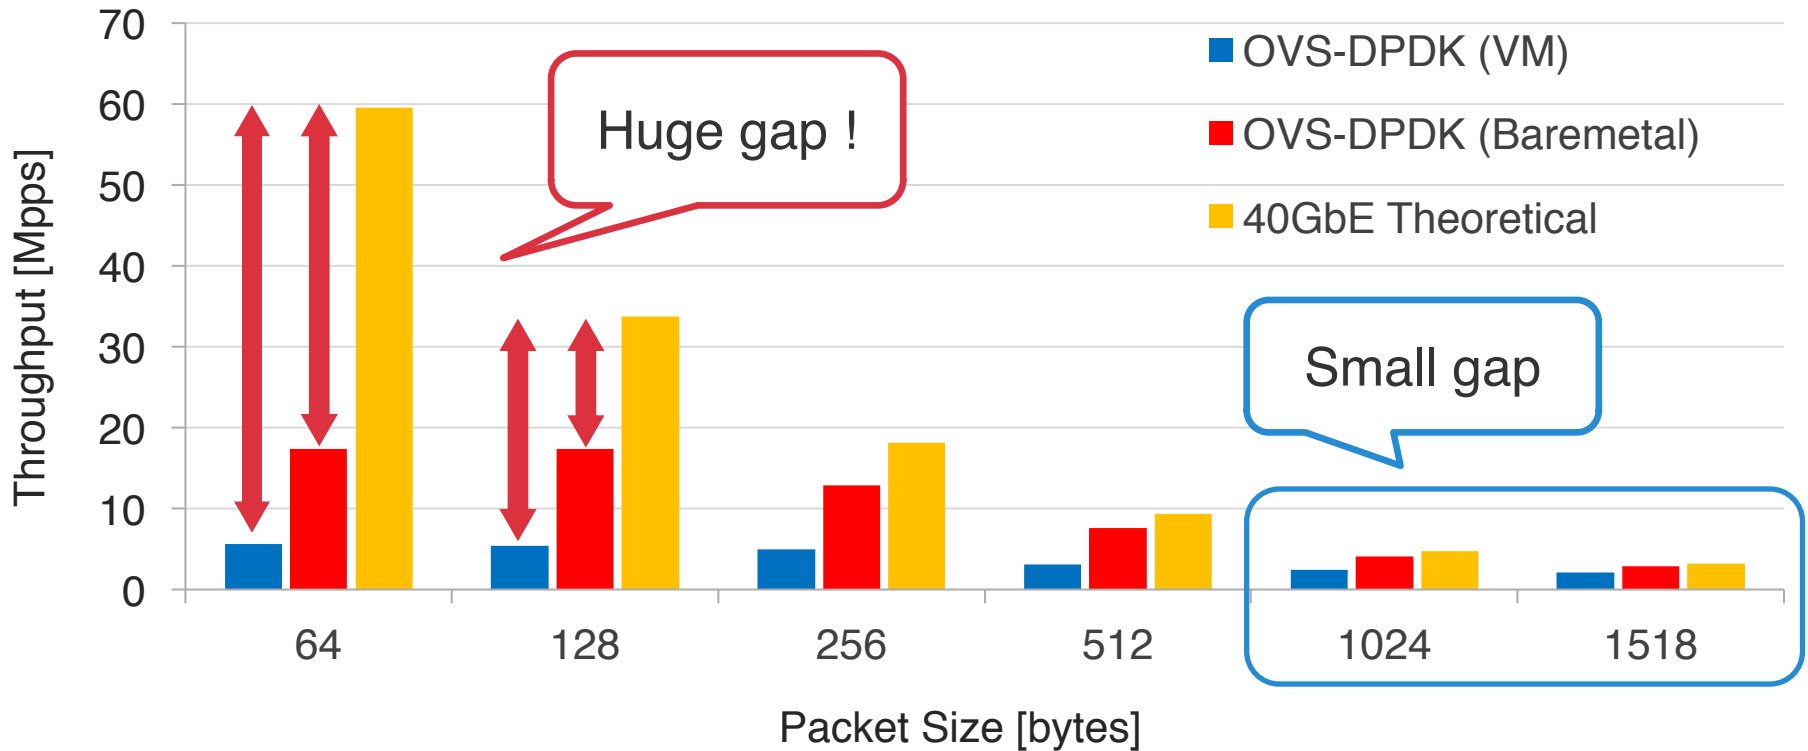


The performance of the upper VNFs must be ensured



Performance degradation factors

- Flow matching by virtual switches
- Packet copy and queueing at vNetwork I/O



“Enlarging packet size” is the key!

† R. Kawashima et al., “Evaluation of Forwarding Efficiency in NFV-nodes toward Predictable Service Chain Performance,” IEEE Trans. on Network and Service Management, 2017.

1 **Background**

- Problems of Service Chaining

2 **Related Work**

- Scaling-out
- Scaling-up

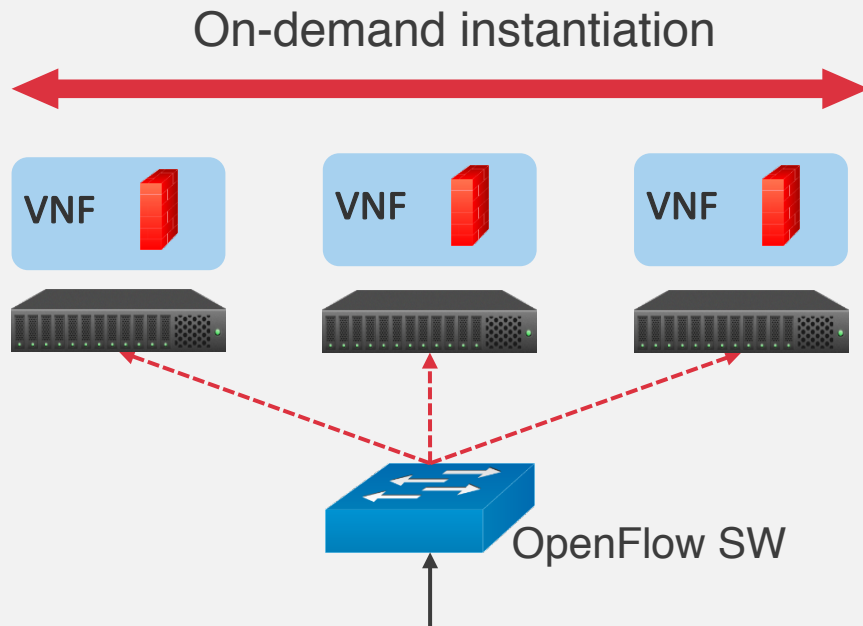
3 **Proposal**

4 **Evaluation**

5 **Conclusion**


● Scaling-Out

1. Auto Scaling





● Scaling-Up

2. NIC Offloading



3. Packet Aggregation

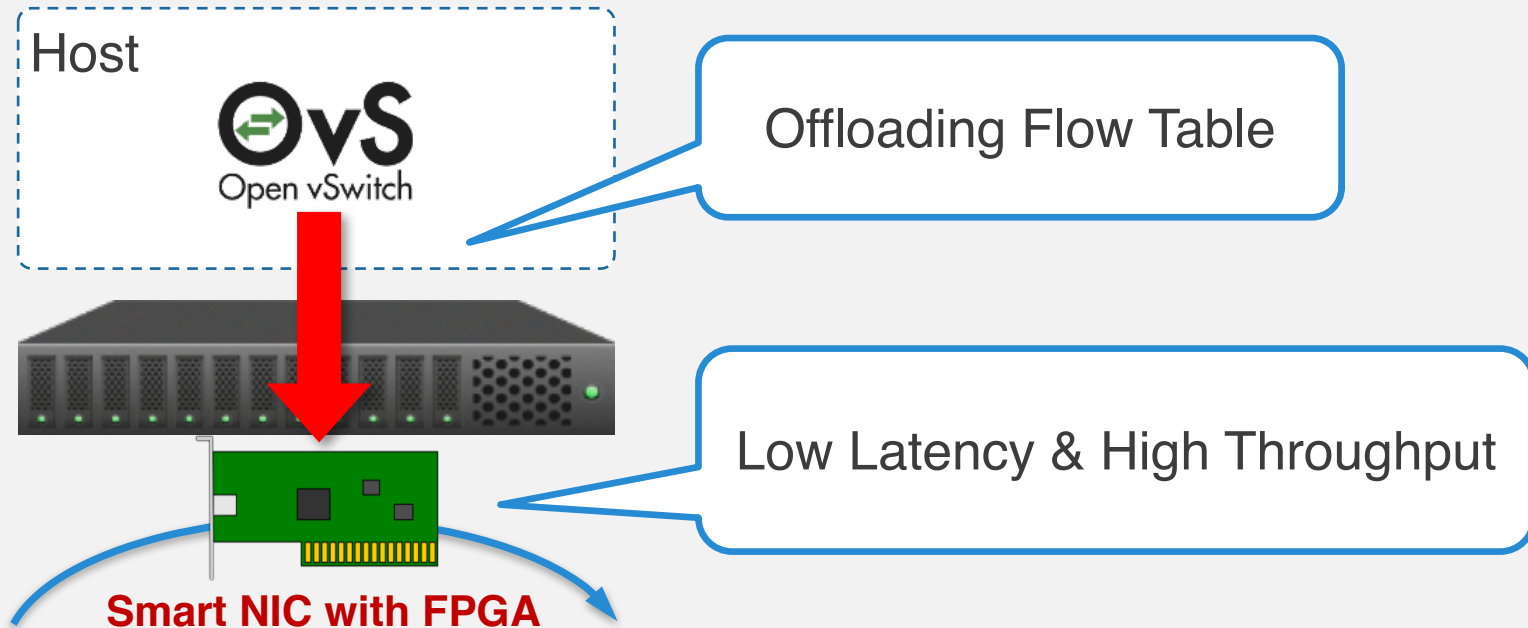




- Physical resources are limited
- There is no de-facto framework for monitoring

† Shoumik Palkar et al., "E2: A Framework for NFV Applications"
ACM Symposium on Operating Systems Principles 2015 (SOSP' 15)

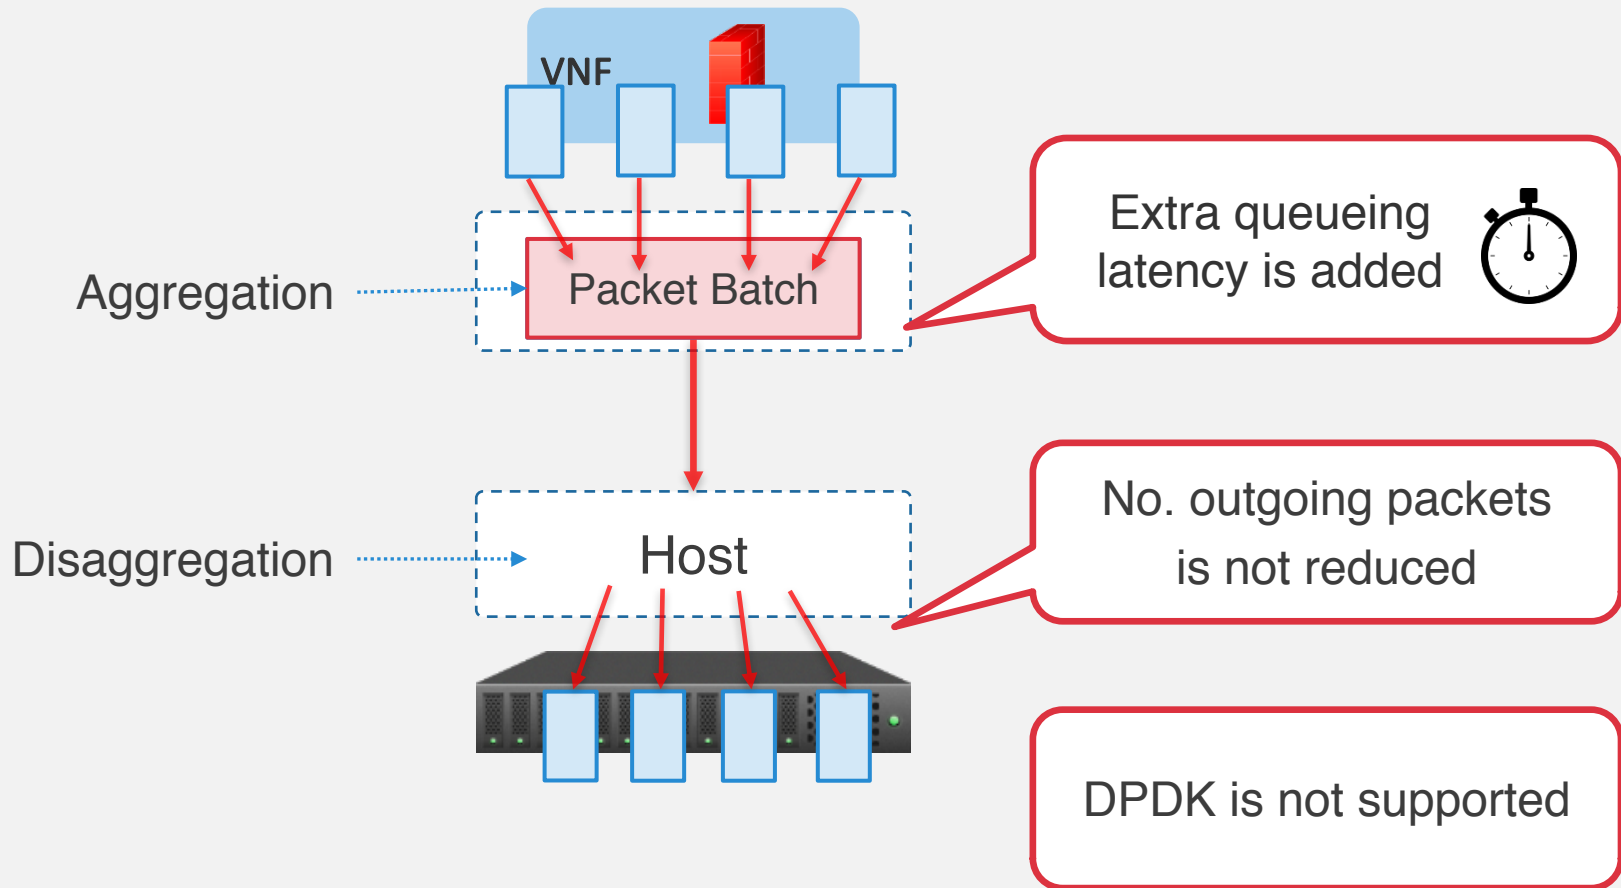
- Smart NIC[†]



- It only works with a specific vSwitch (e.g., Open vSwitch)
- Performance gain is limited

[†] Agilio OVS Software, combined with Agilio SmartNICs
<https://www.netronome.com/products/agilio-software/agilio-ovs-software/>

- Aggregation between VMs and the Host[†]



[†] M. Bourguiba et al., "Improving Network I/O Virtualization for Cloud Computing" IEEE Trans. on Parallel and Distributed Systems, 2014

1 Background

- Problems of Service Chaining

2 Related Work

- Scaling-out
- Scaling-up

3 Proposal

4 Evaluation

5 Conclusion

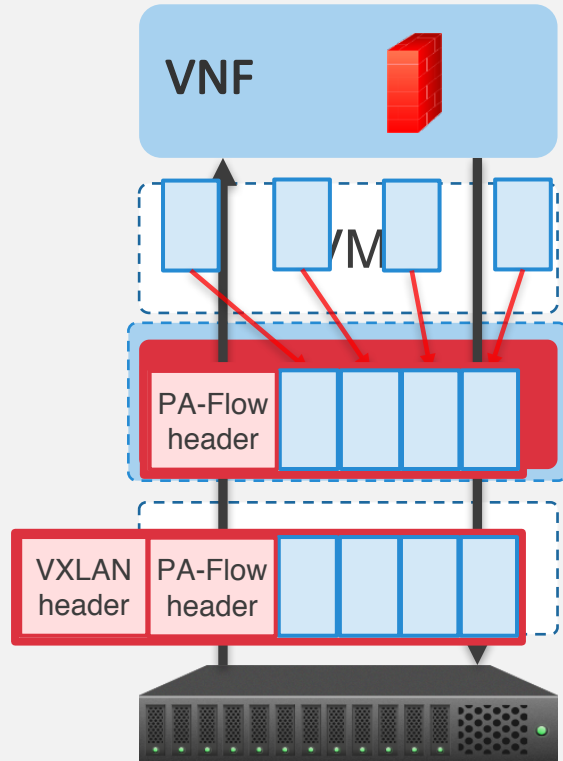
PA-Flow (Packet Aggregation Flow)

- **Key points**

1. Fast and Network-aware Aggregation

2. Gradual Packet Aggregation

3. Next-hop-aware Aggregation



No extra queueing latency!


→ explained later

Outgoing packets are reduced

DPDK is supported

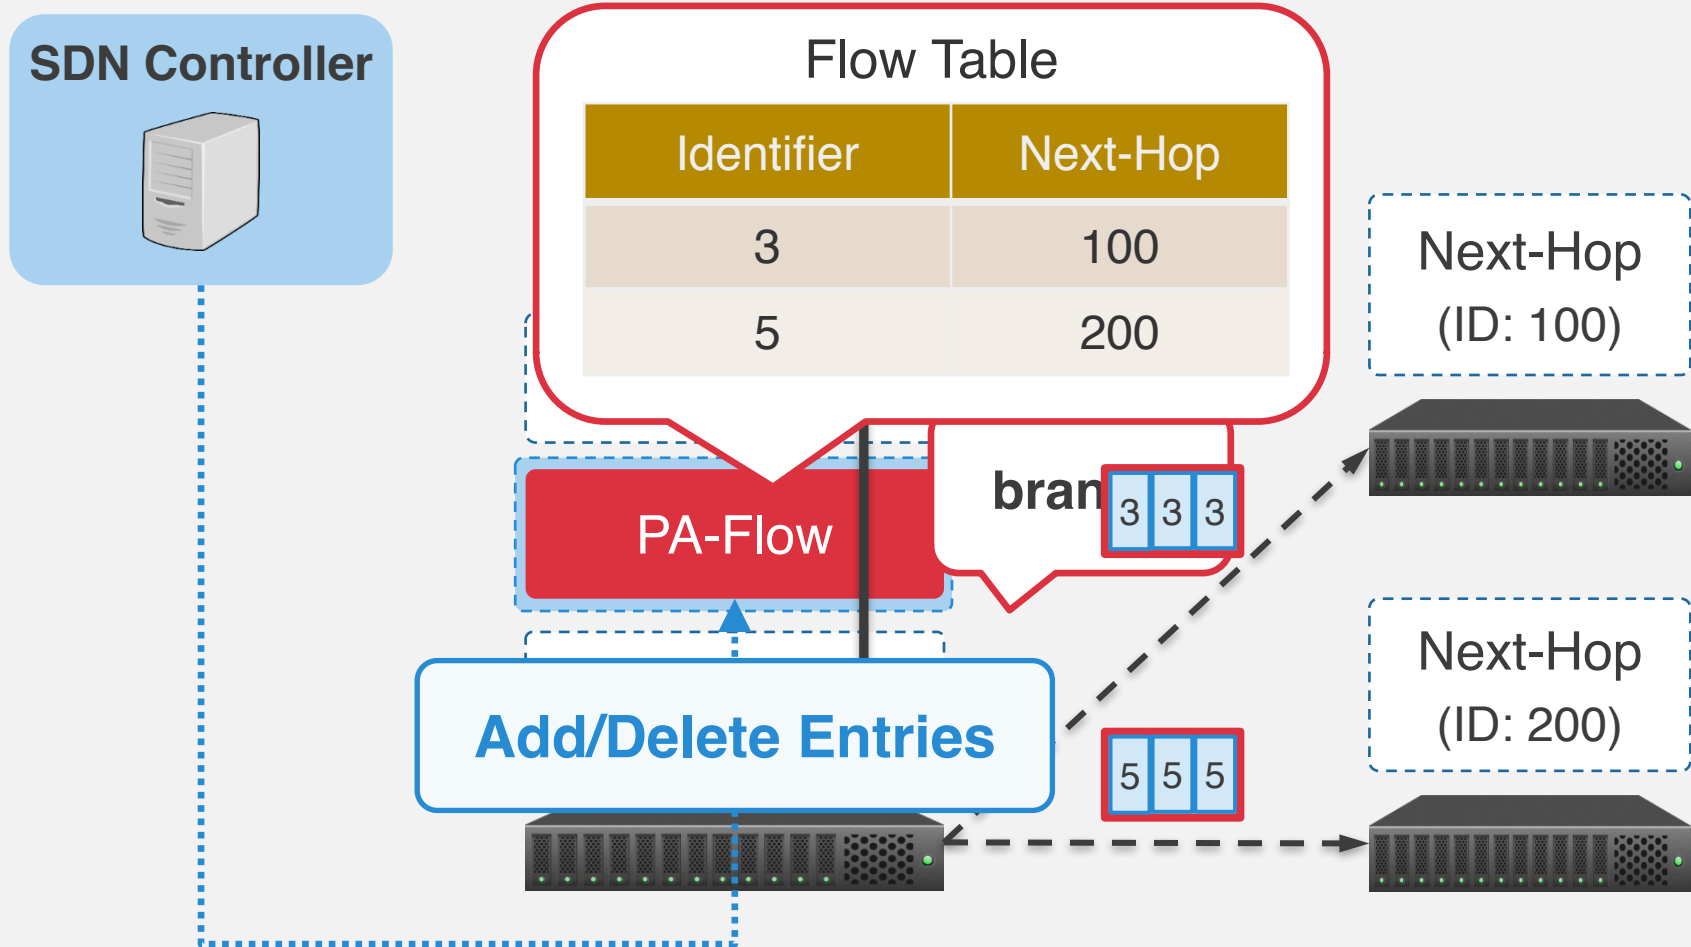
All the existing problems are solved!

Gradually aggregated



The average packet size is enlarged!

3. Next-hop Aware Aggregation

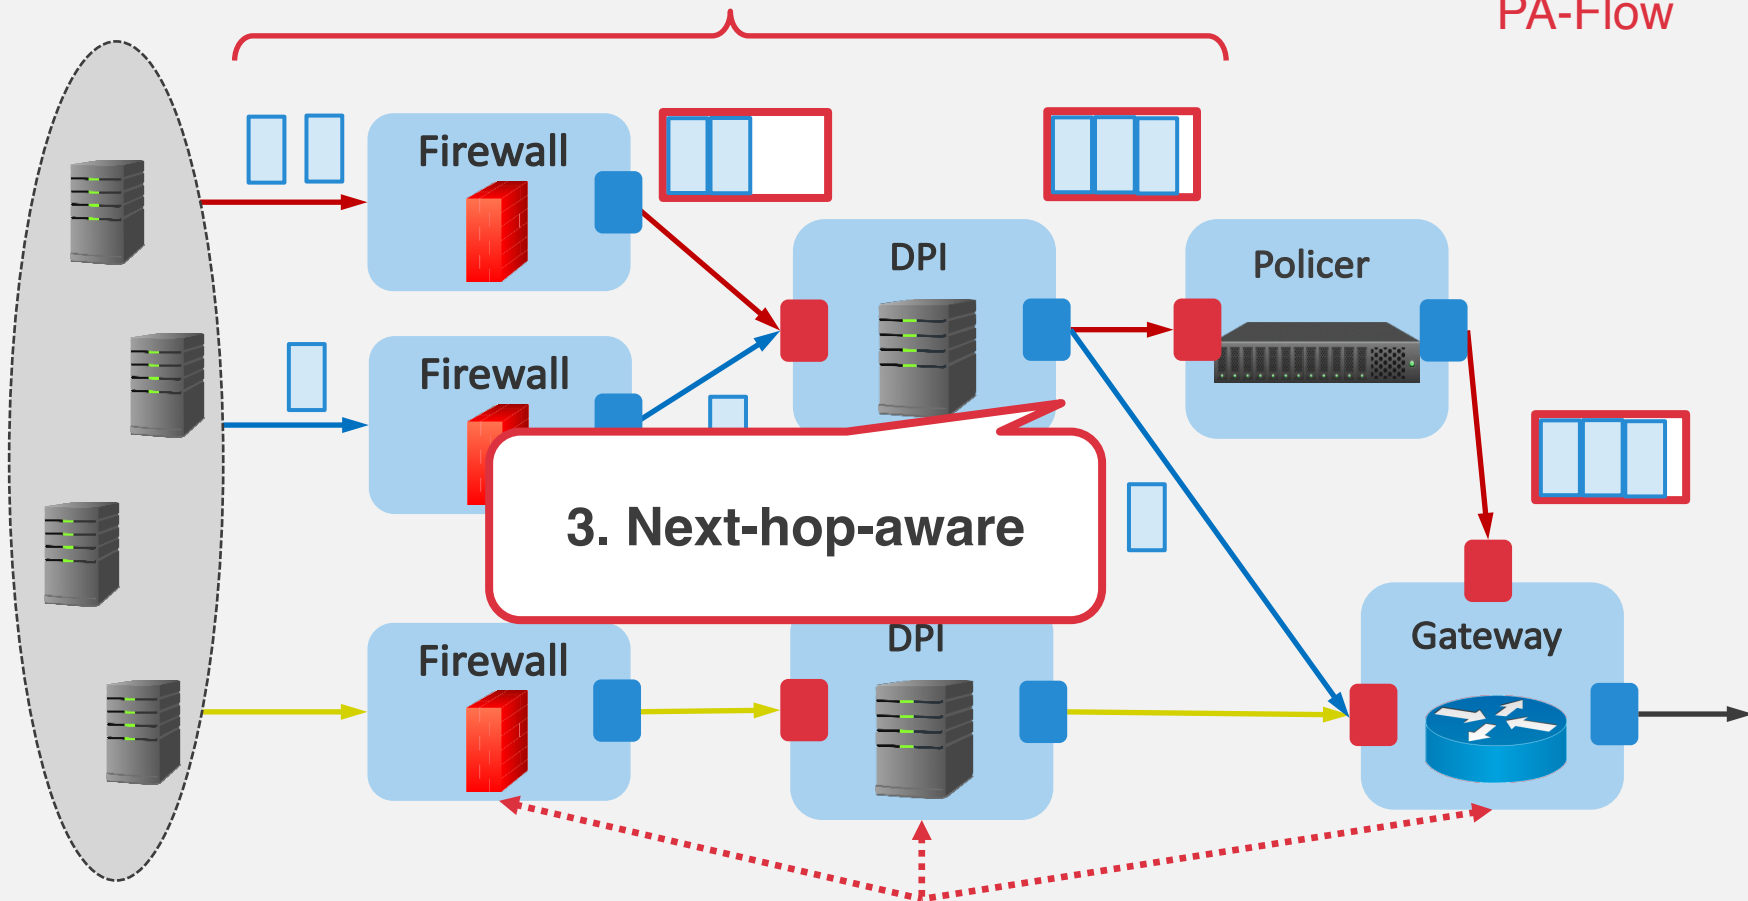


PA-Flow supports branches of service chain

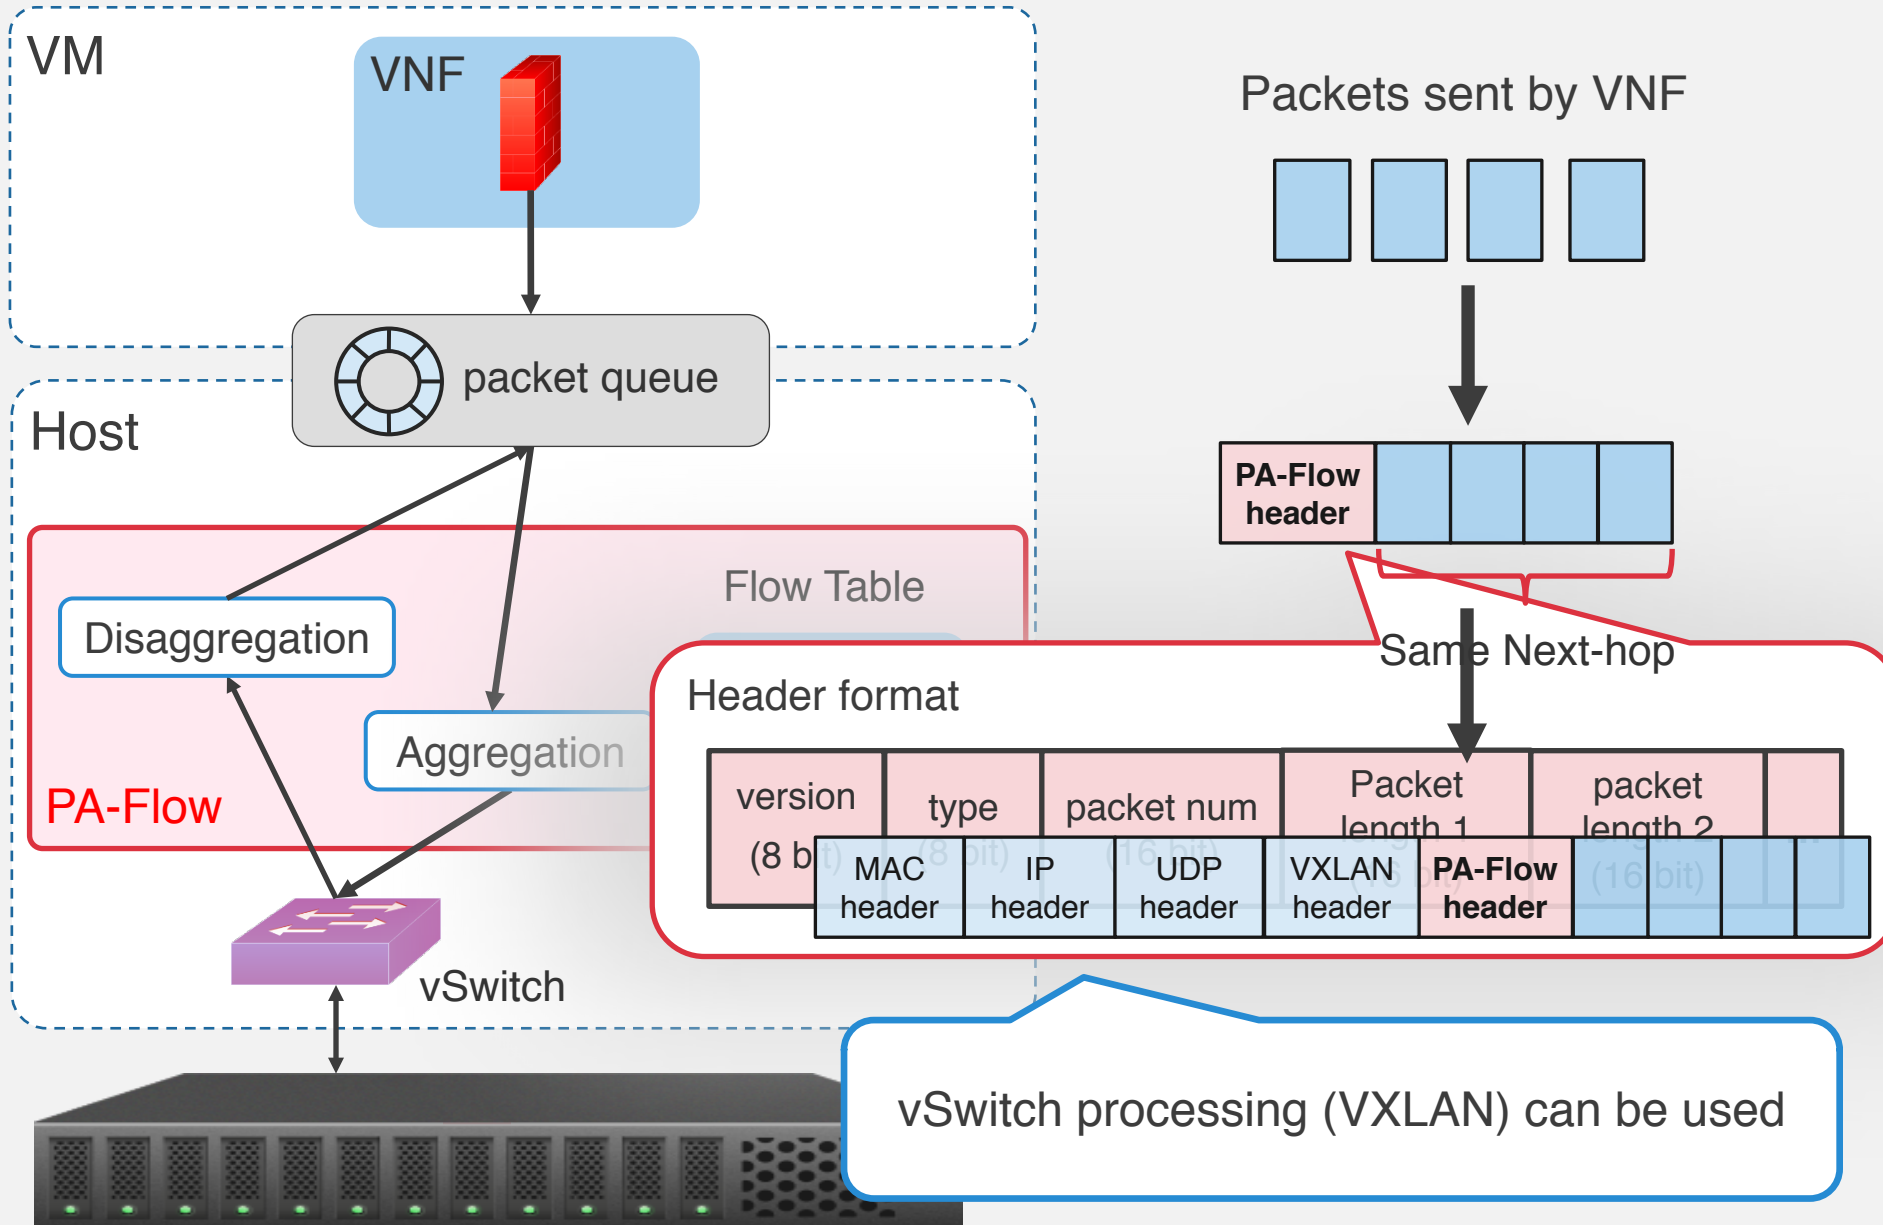
2. Gradual aggregation

■ : aggregation
■ : disaggregation

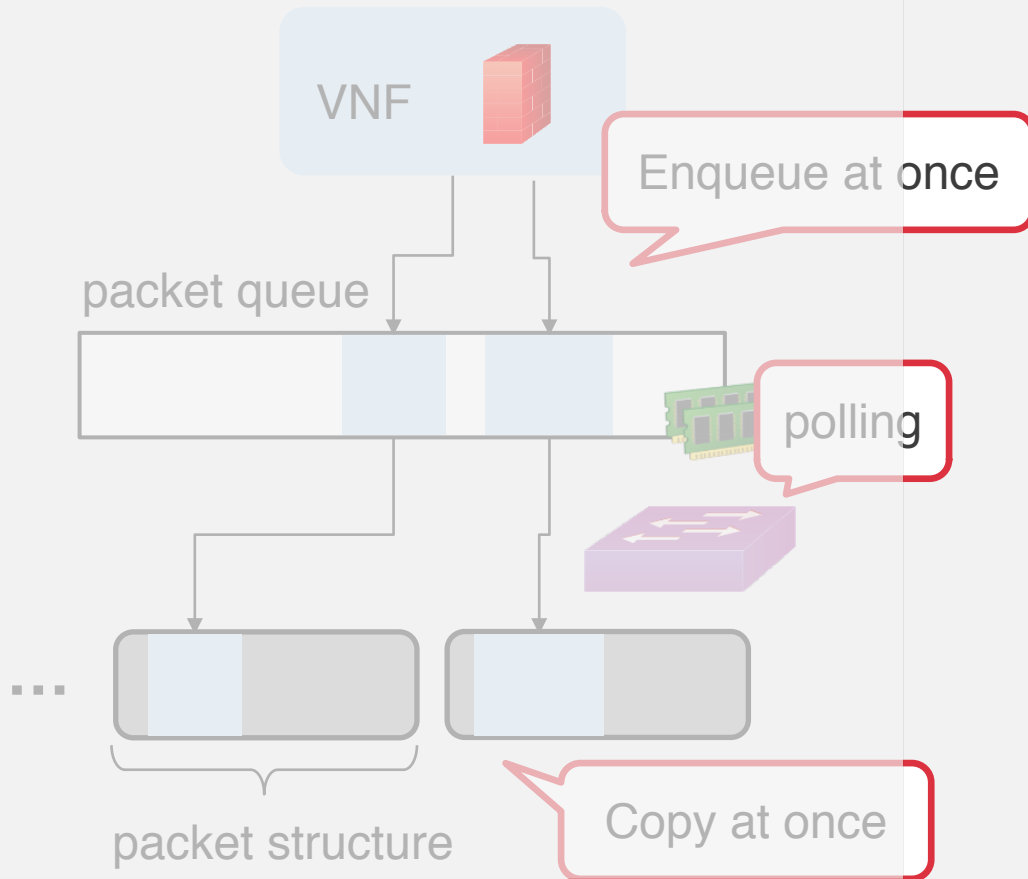
PA-Flow



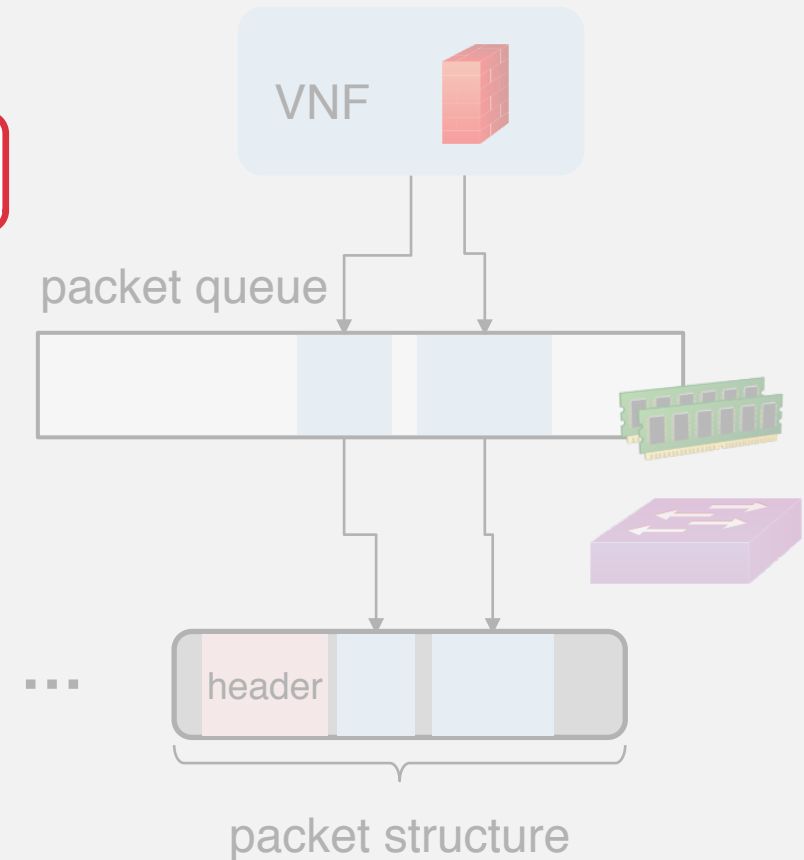
1. Fast and Network-aware aggregation



(a) Default vhost-user



(b) PA-Flow



There is no extra waiting time for aggregation

1 Background

- Problems of Service Chaining

2 Related Work


- Scaling-out
- Scaling-up

3 Proposal

4 Evaluation

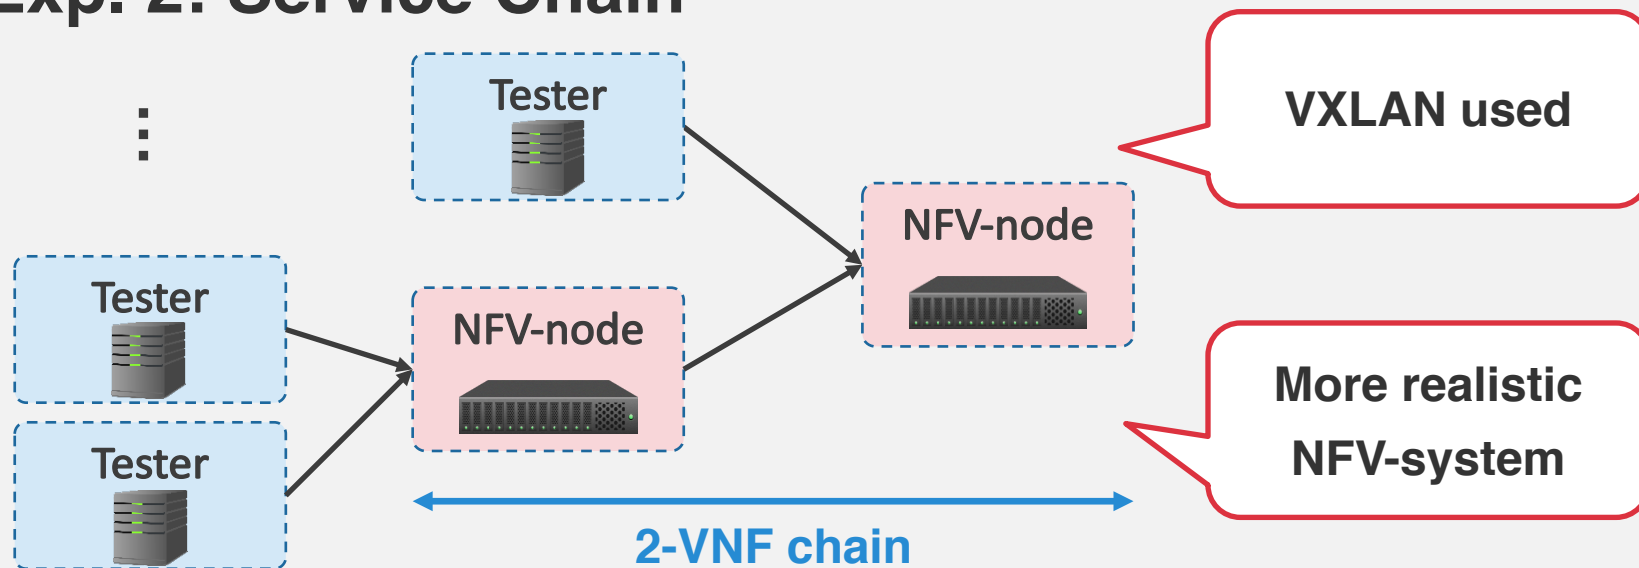
5 Conclusion

● Exp. 1: Baseline



CPU	i7-6900K (8 cores, 3.2 GHz)
NIC	10 GbE (Intel X540)
OS	CentOS 7.3

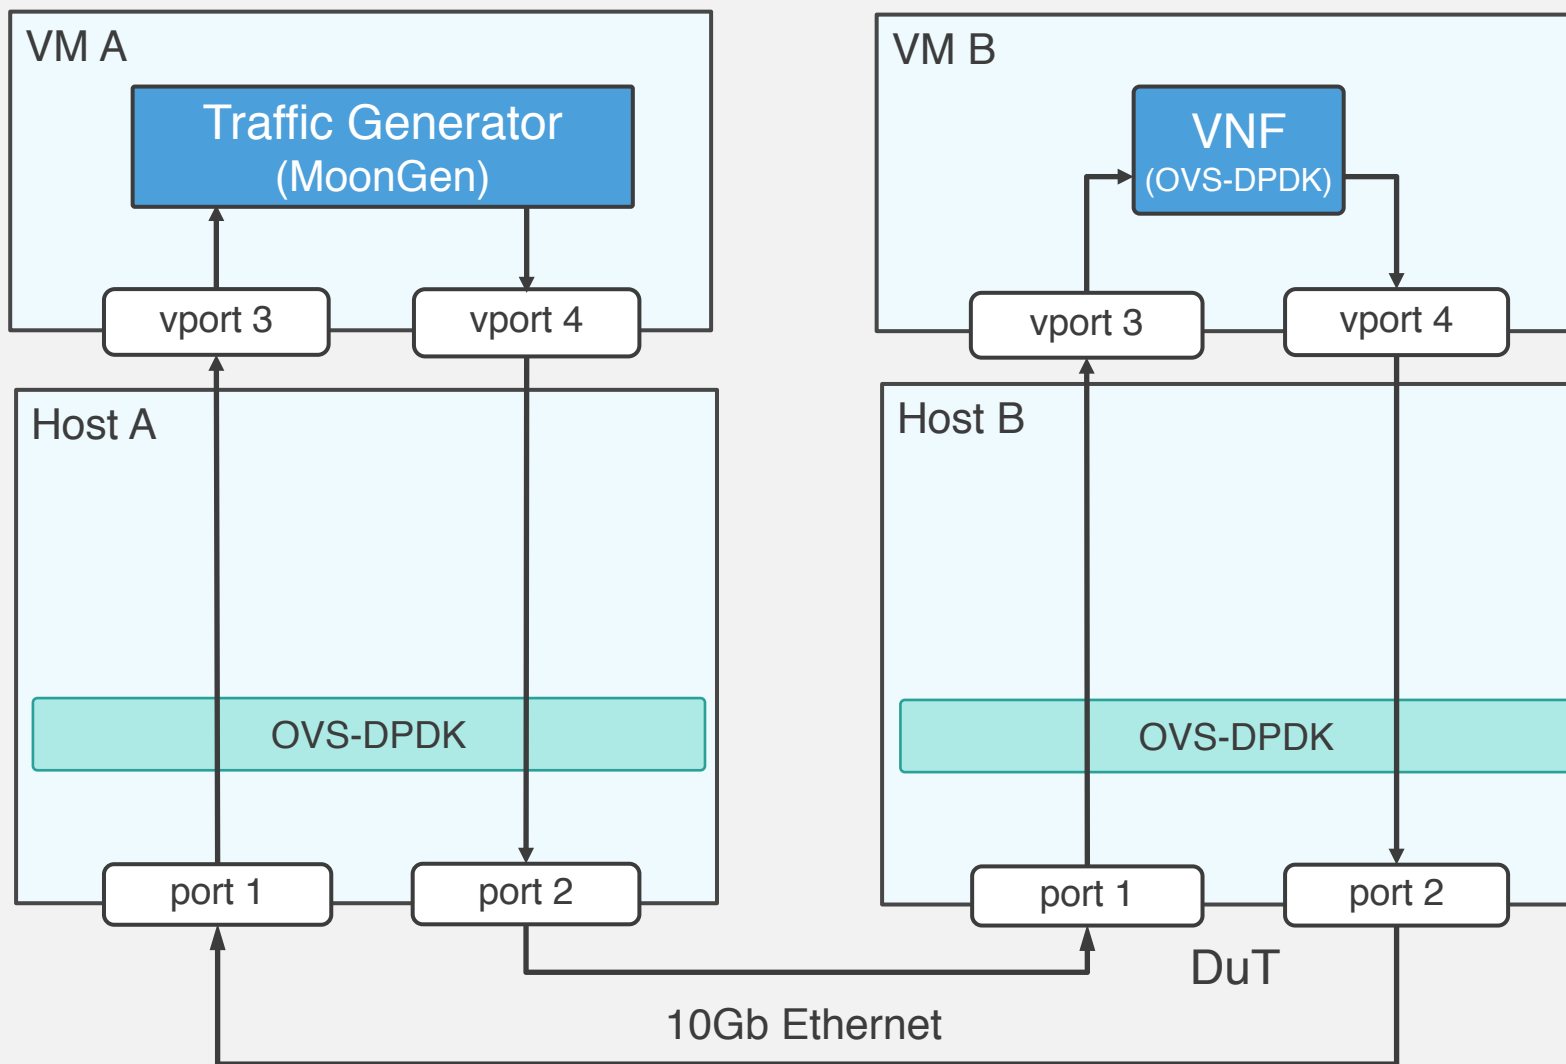
● Exp. 2: Service Chain



(a) Default

(b) PA-Flow

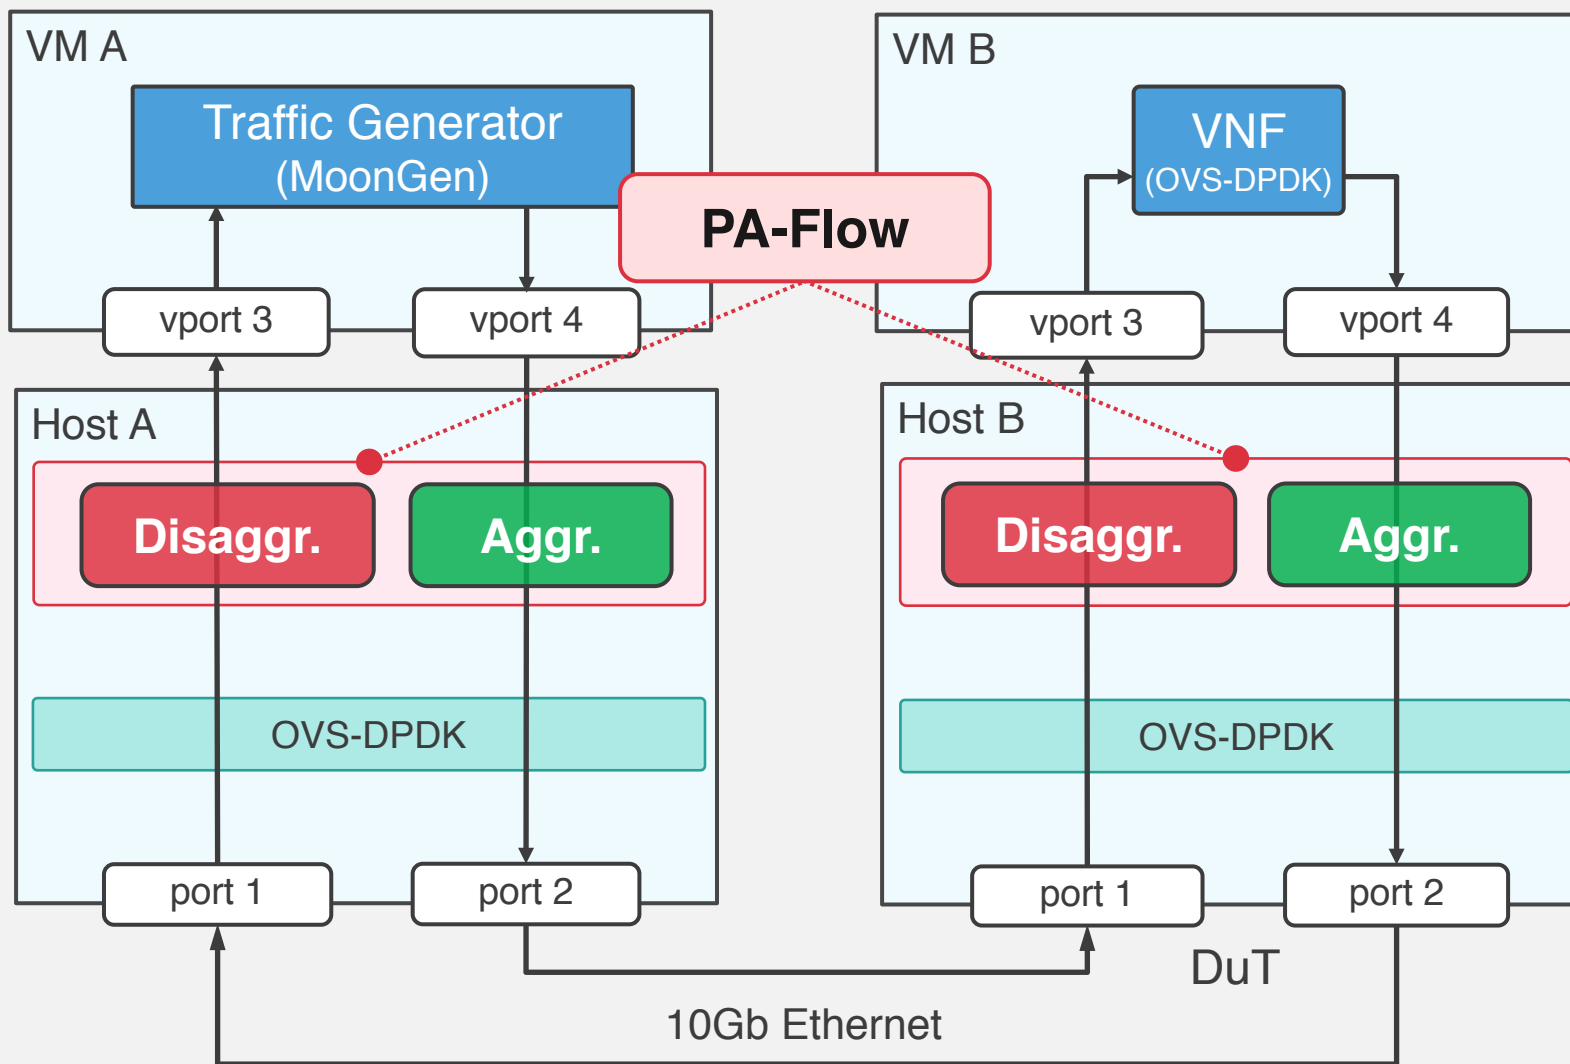
(c) PA-Flow with VNF support



(a) Default

(b) PA-Flow

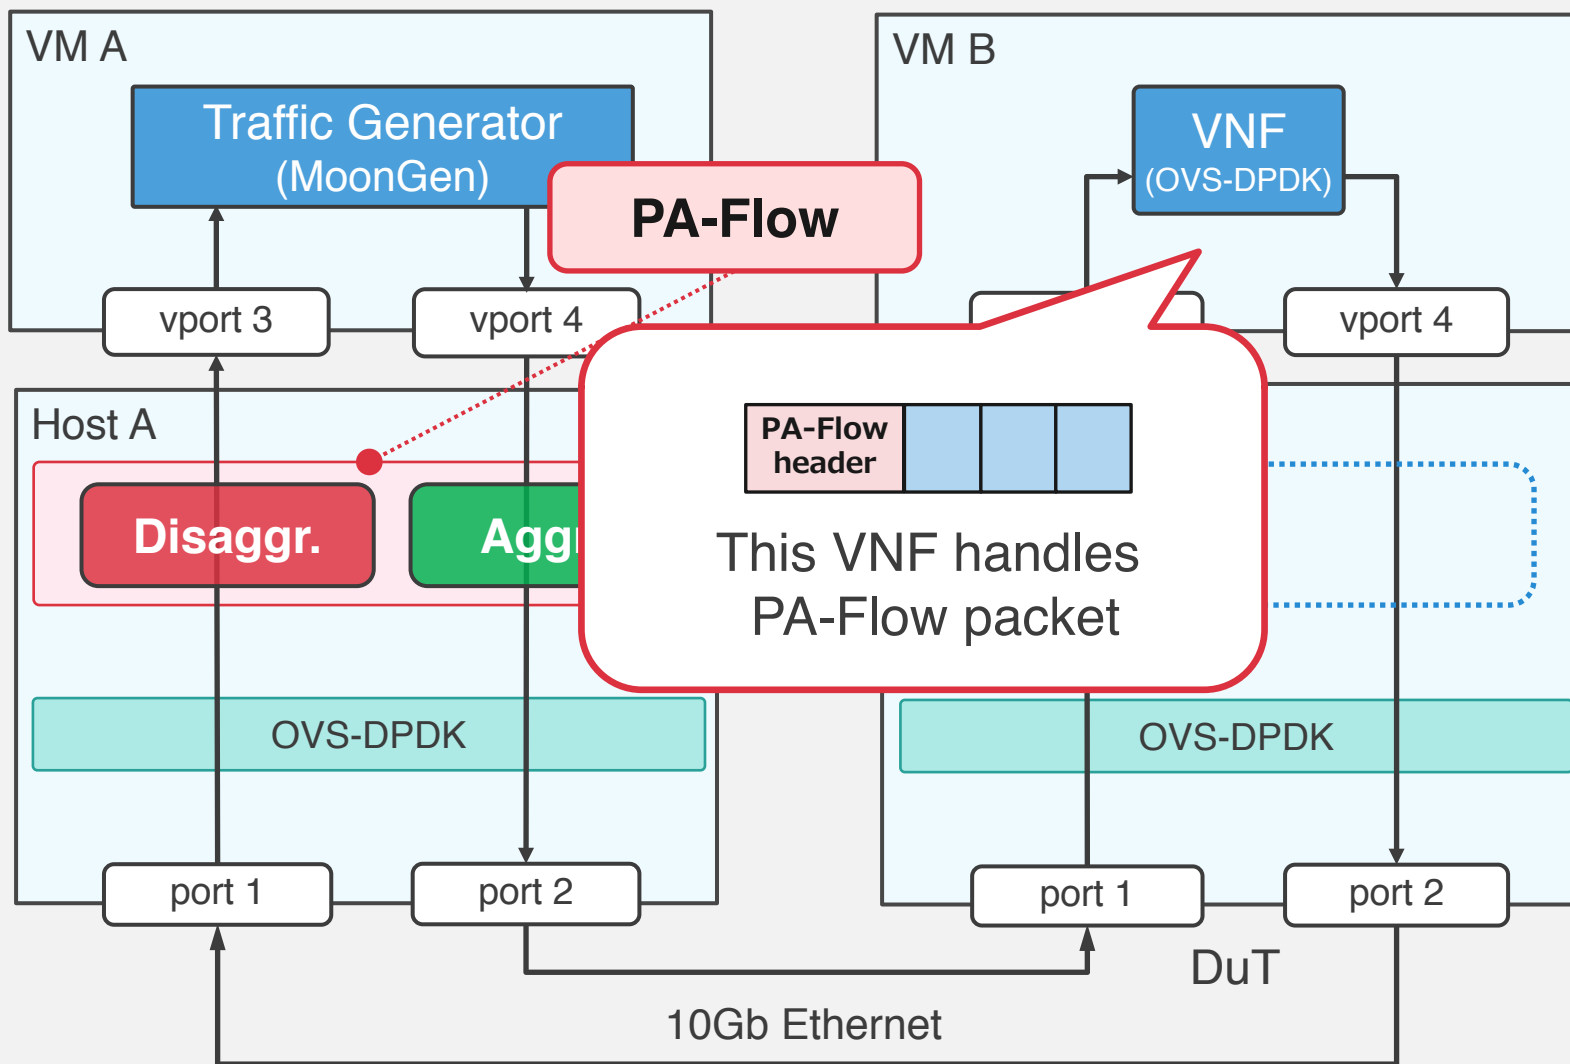
(c) PA-Flow with VNF support

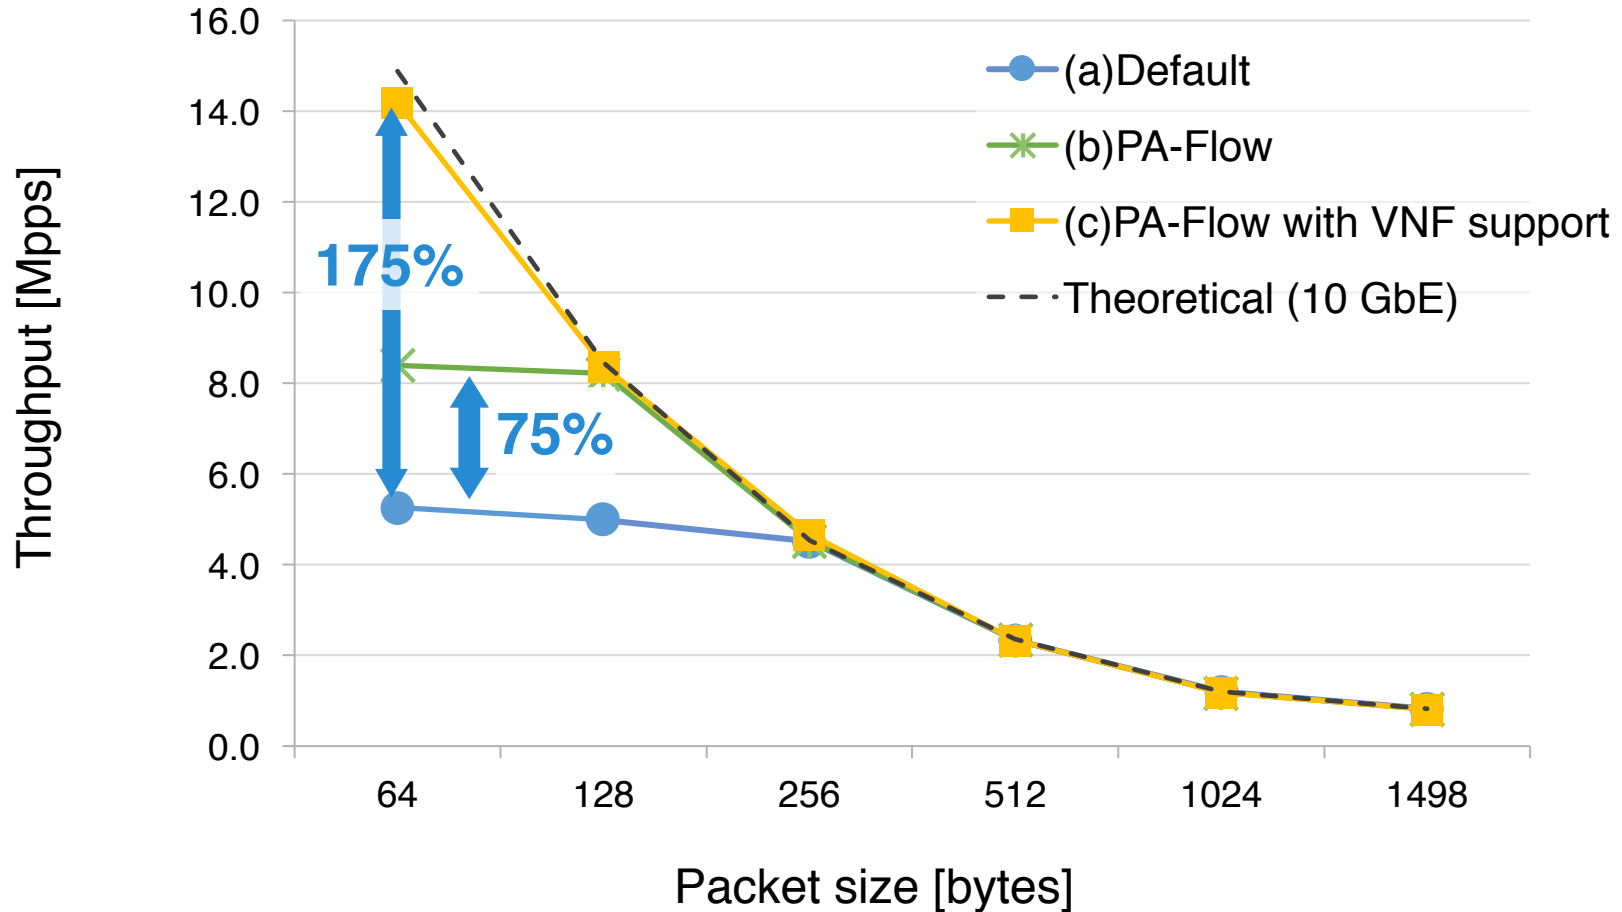


(a) Default

(b) PA-Flow

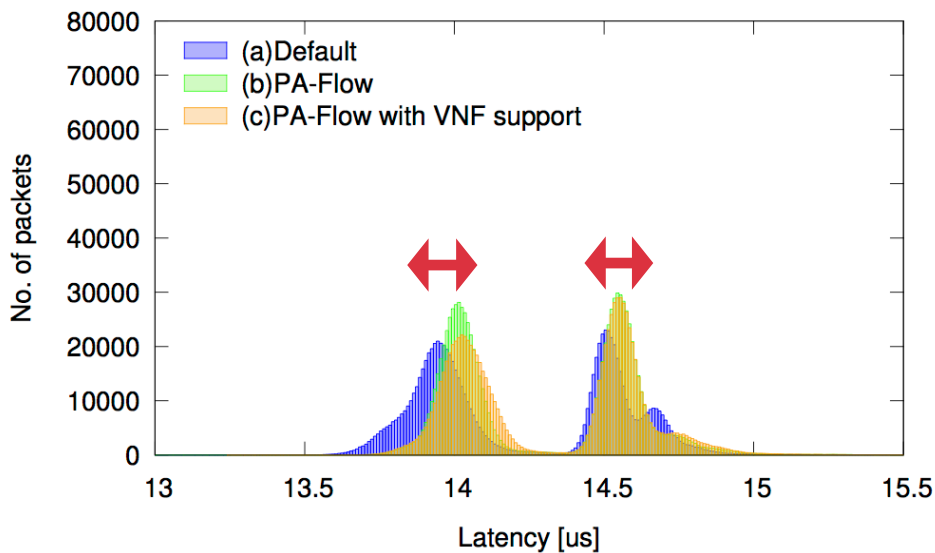
(c) PA-Flow with VNF support






The performance is improved to nearly 10 Gbps

Tx rate: 1 Mpps
(1% aggregated)



200 ns increased

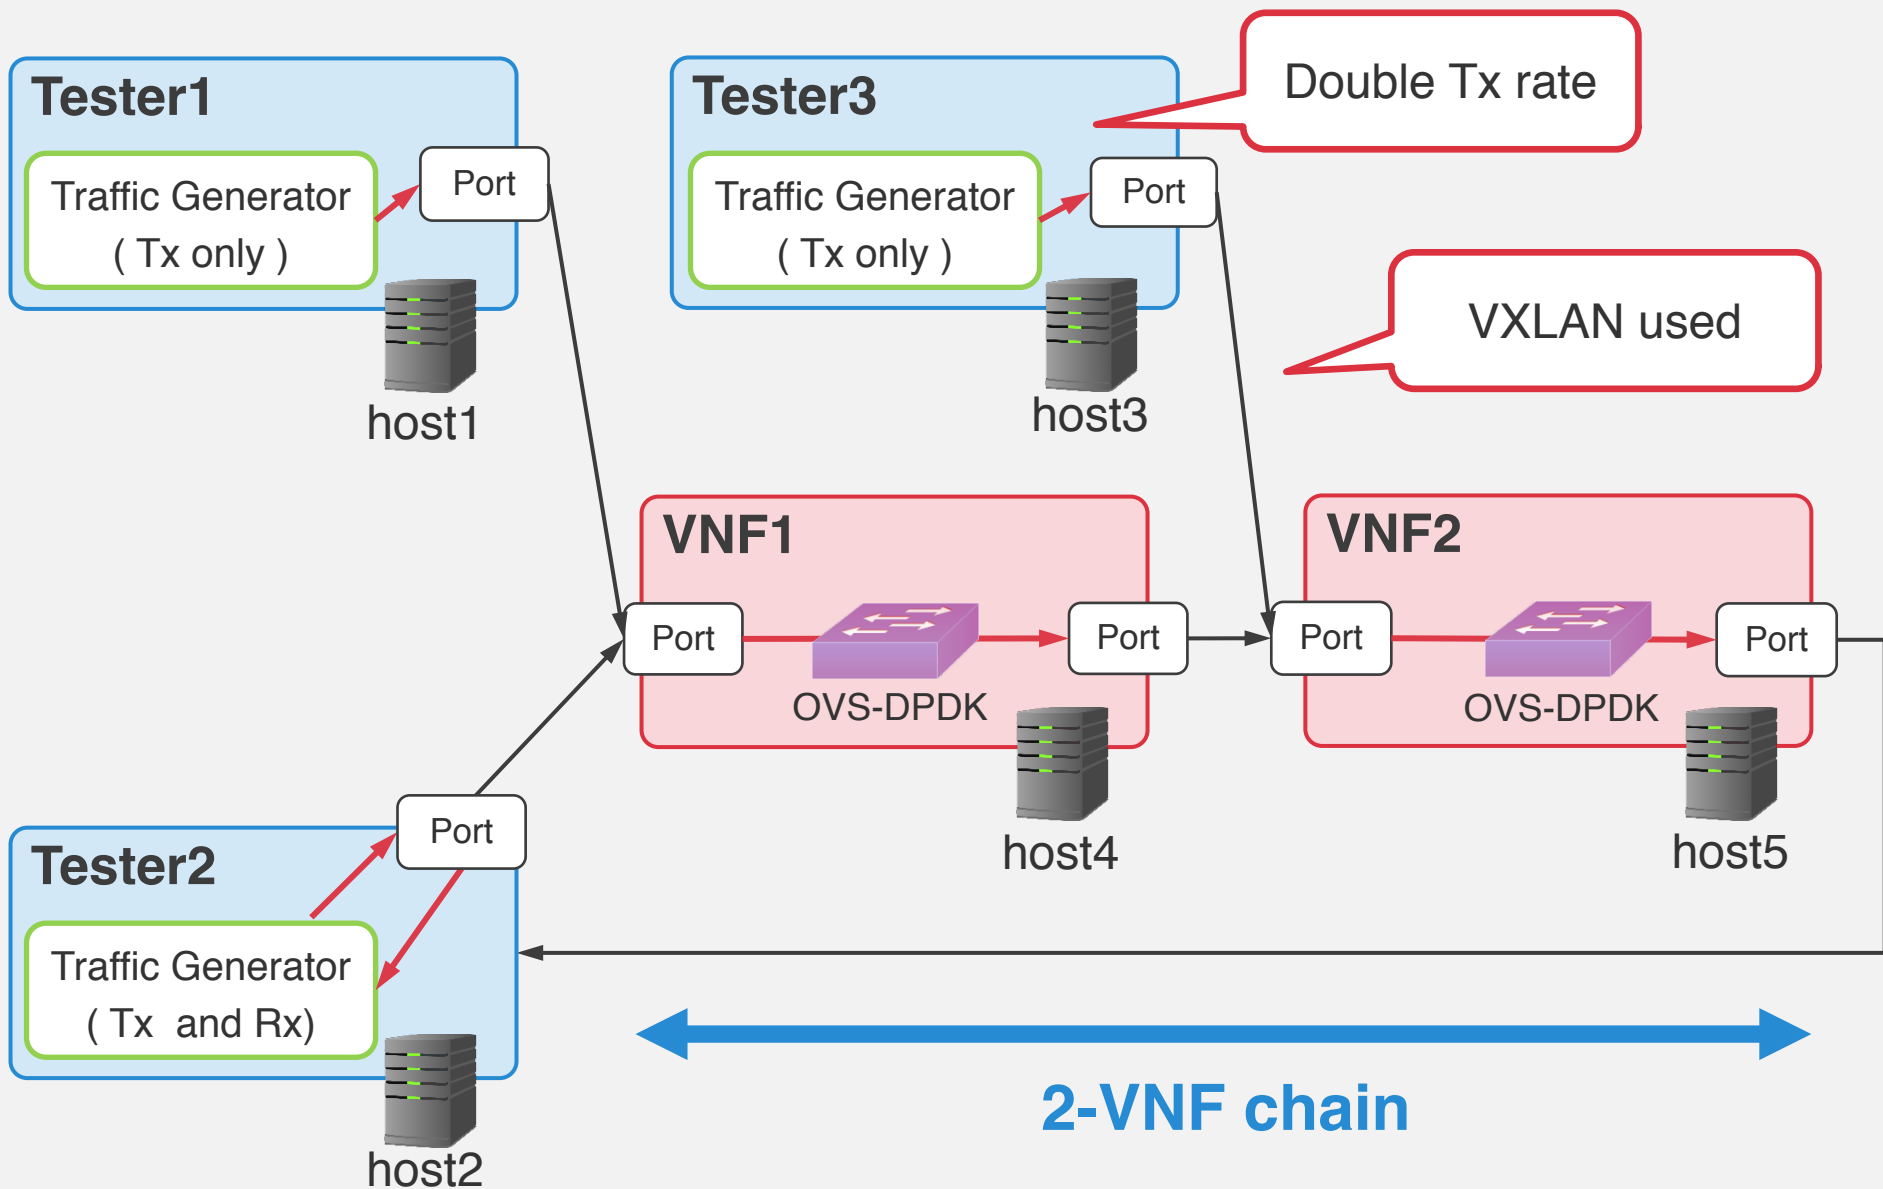
Tx rate: 3 Mpps
(25% aggregated)



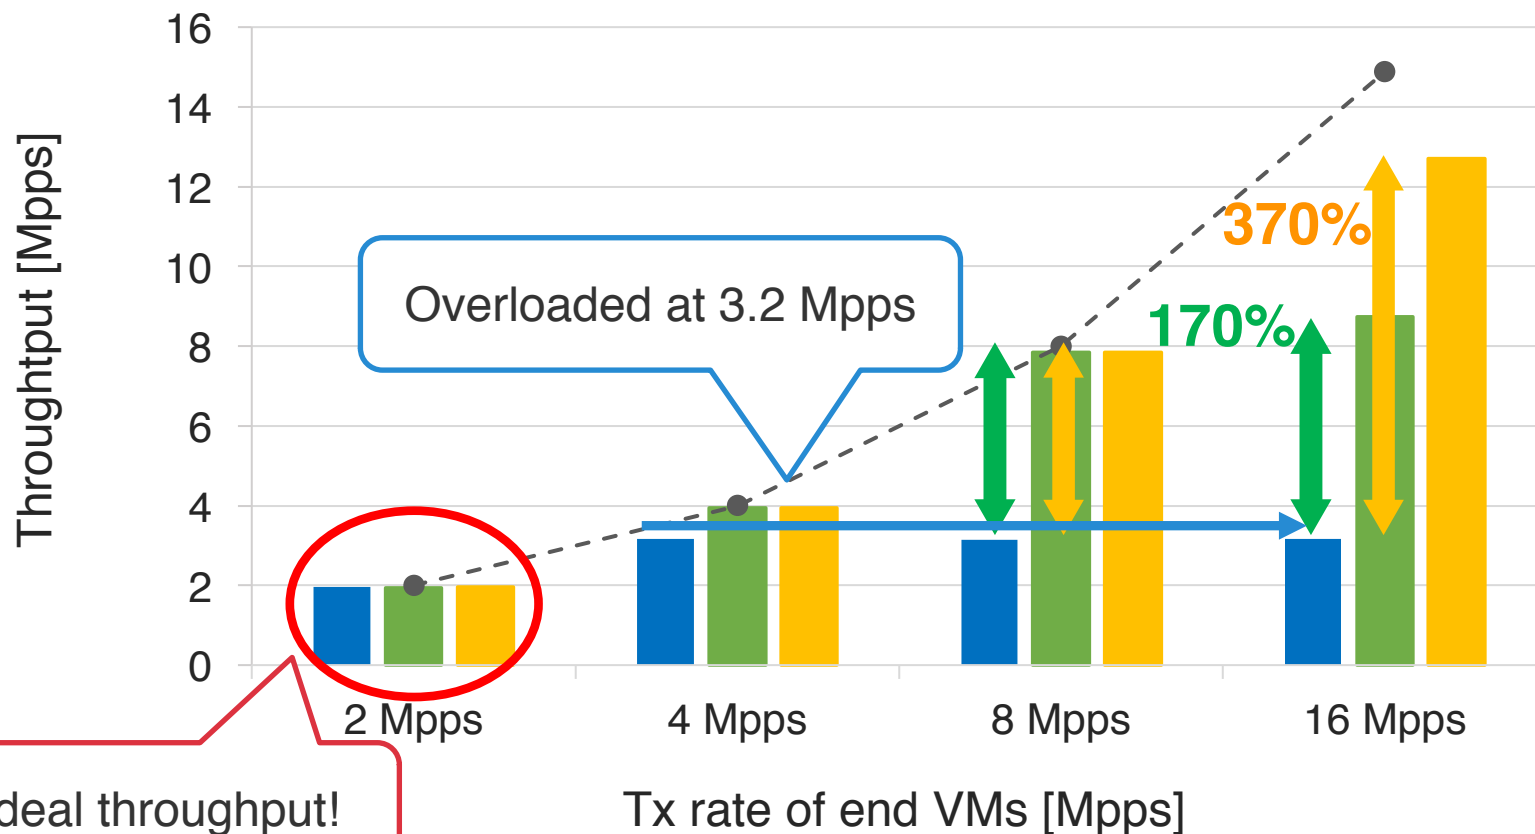
2 μ s reduced

The pure overhead is negligible

The latency reduced under high-rate communications



(a)Default (b)PA-Flow (c)PA-Flow with VNF support -●-ideal throughput



The performance of service chain is drastically boosted

1 Background

- Problems of Service Chaining

2 Related Work

- Scaling-out
- Scaling-up

3 Proposal

4 Evaluation

5 Conclusion

- **Performance limitation of Service Chain**

- Software-based scale-up approach is needed.

- **Proposal: PA-Flow (Packet Aggregation Flow)**

- Gradual aggregation reduces forwarding cost of upper VNFs.
- Flexible operation is realized with next-hop aware aggregation.
- Performance improved by 170% on a 2-VNF service chain.

- **Future work**

- Further performance improvement
 - Porting PA-Flow to a VM-side component
- Evaluation of container