

# NetFPGA Summer Course



**Presented by:**

**Andrew W Moore, Noa Zilberman, Gianni Antichi  
Stephen Ibanez, Marcin Wojcik, Jong Hun Han,  
Salvator Galea, Murali Ramanujam, Jingyun Zhang,  
Yuta Tokusashi**

**University of Cambridge  
July 24 – July 28, 2017**

**<http://NetFPGA.org>**

# The power of OpenSource: build your own proof of concept!

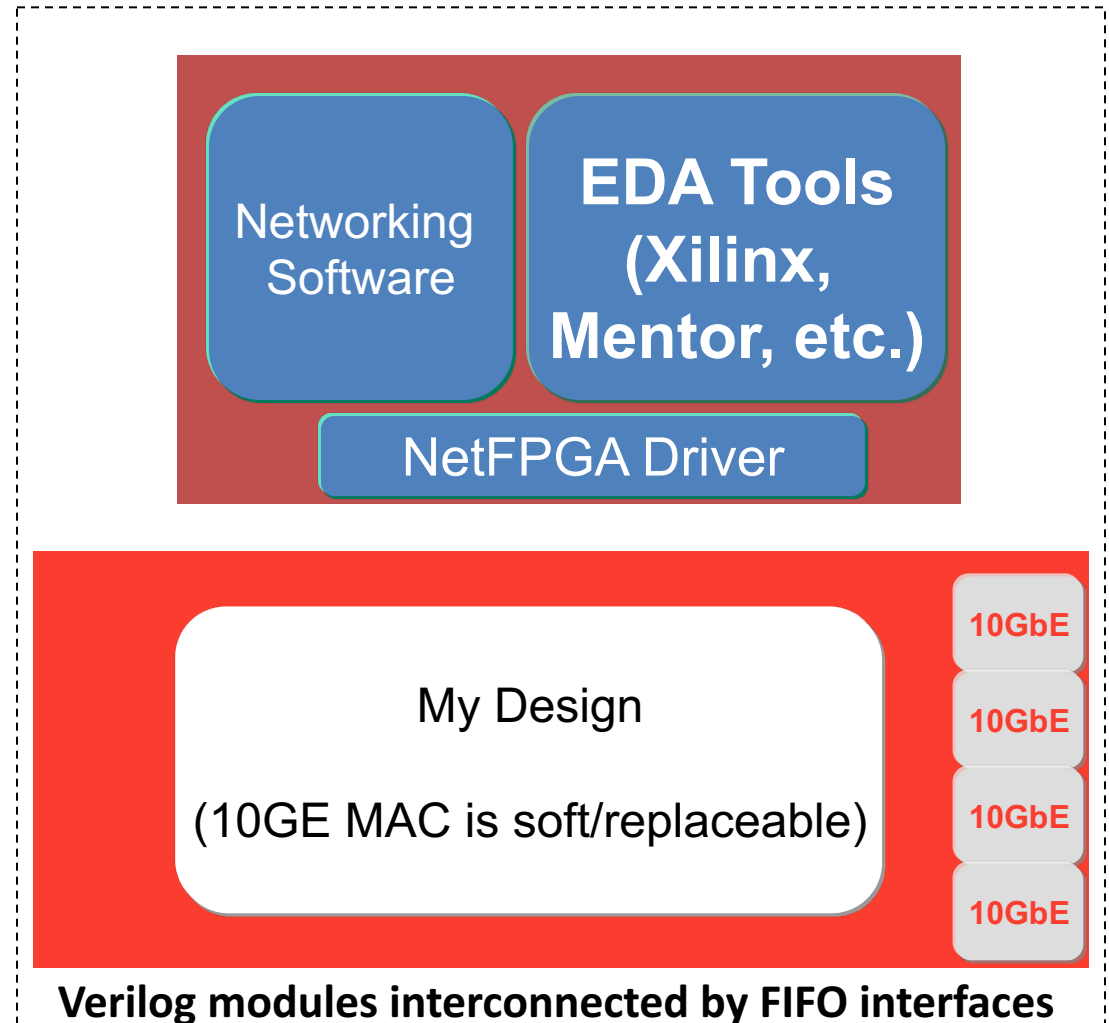
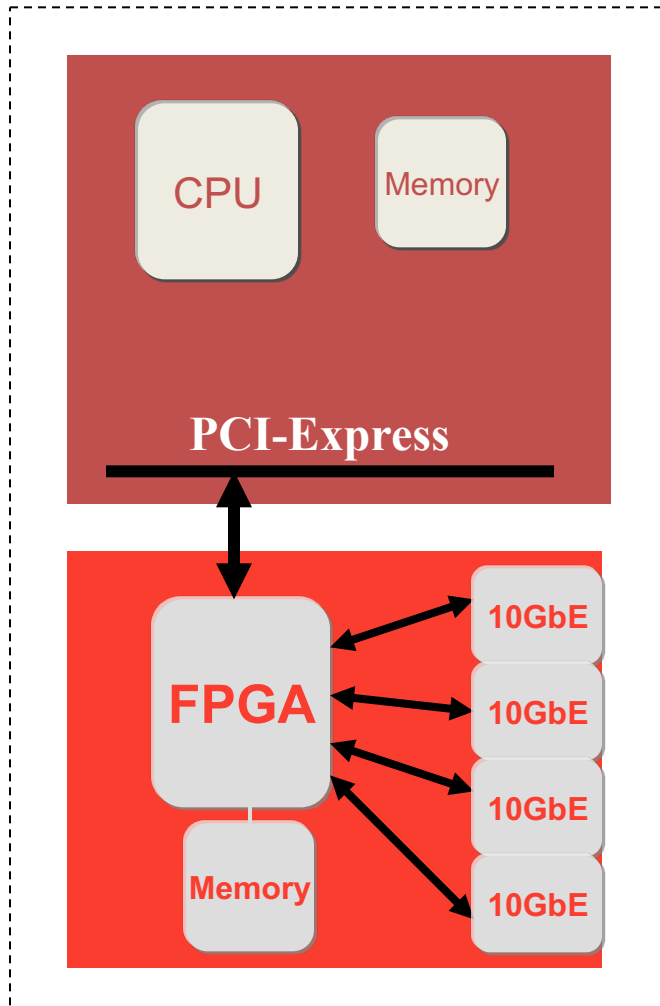
Gianni Antichi

Computer Laboratory  
University of Cambridge

[gianni.antichi@cl.cam.ac.uk](mailto:gianni.antichi@cl.cam.ac.uk)  
<https://www.cl.cam.ac.uk/~ga288>

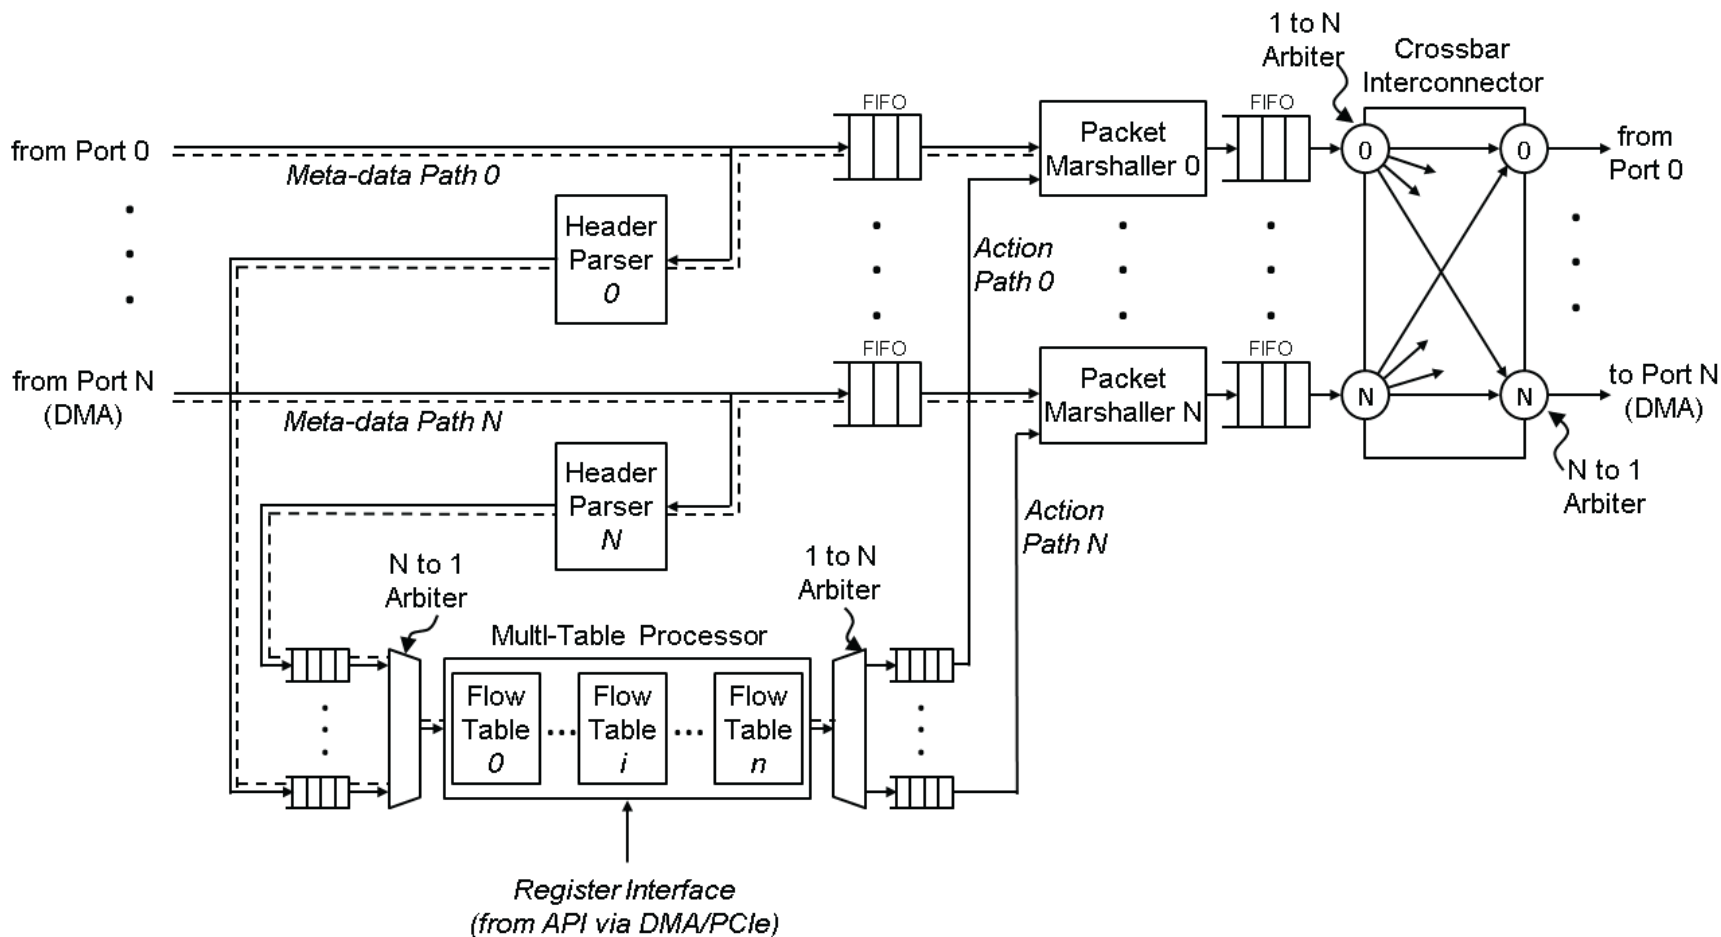


# NetFPGA: Networked FPGA



# BlueSwitch: A Multi Table OpenFlow Switch

**Your design can look completely different!**

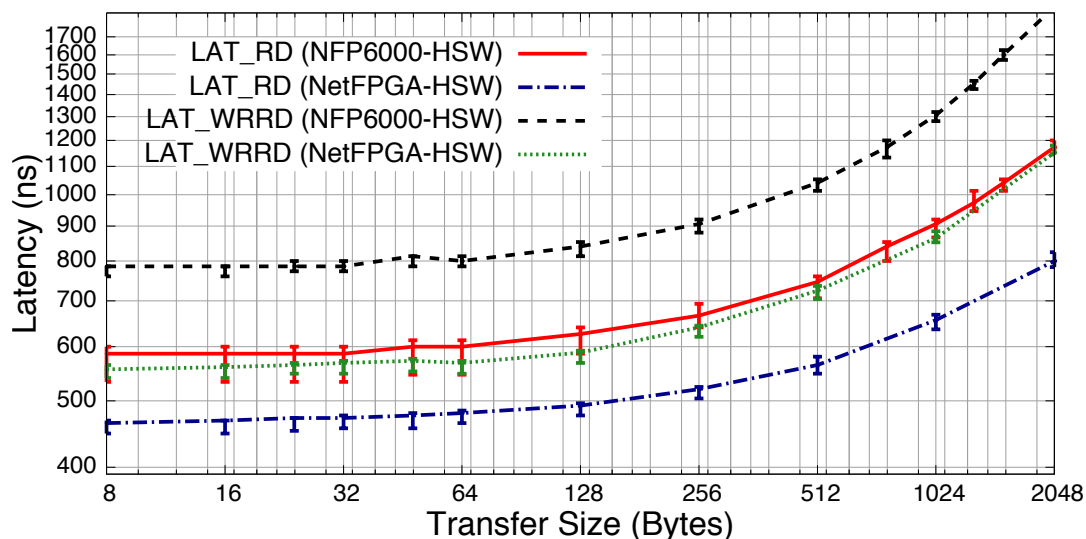




# pciebench: an open source tool for benchmarking PCI Express


there is a limited understanding of PCIe functionality, nor the trade-offs that must be made to get best-performance from PCIe systems

- pciebench tool open source available
- It builds on NetFPGA and Netronome boards



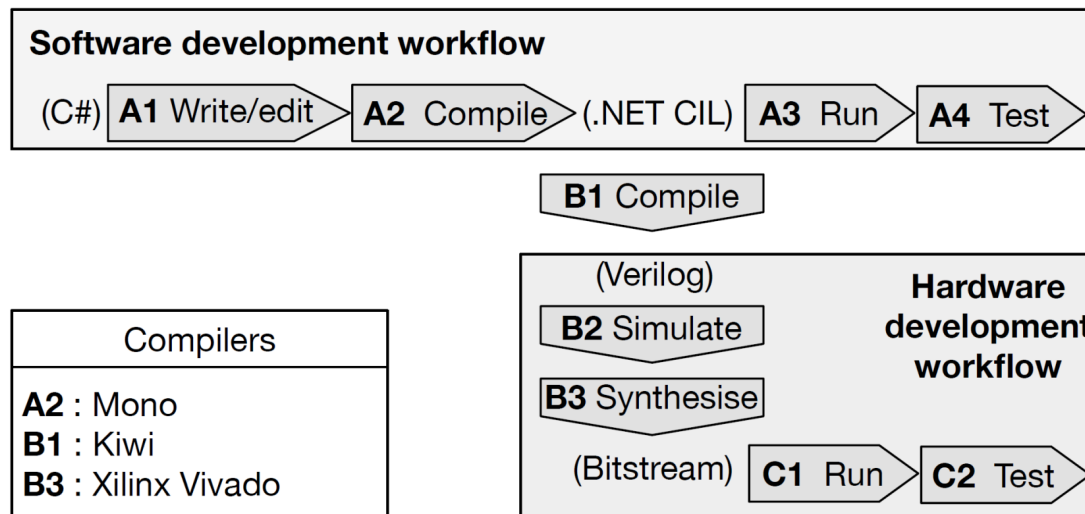

# Power Efficient MAC

- A Platform for 100Gb/s power-saving MAC design (e.g. lights-out MAC)
- Porting MAC design to SUME permits:
  - Power measurements
  - Testing protocol's response
  - Reconsideration of power-saving mechanisms
  - Evaluating suitability for complex architectures and systems




# Emu: Rapid FPGA Prototyping of Networking Services in C#

- Accelerating networking services
- Compiling .Net programs
  - To x86
  - To simulation environment
  - To multiple FPGA targets






# Open Source Network Tester



- NetFPGA platform enabled the first prototype of OSNT.
- The open nature of NetFPGA ecosystem represents the best starting point for open HW/SW community-oriented projects.
- OSNT aims to build a community as NetFPGA did.

# OSNT: Open Source Network Tester
















Open source hardware and software platform for network monitoring and testing.

<https://osnt.org>

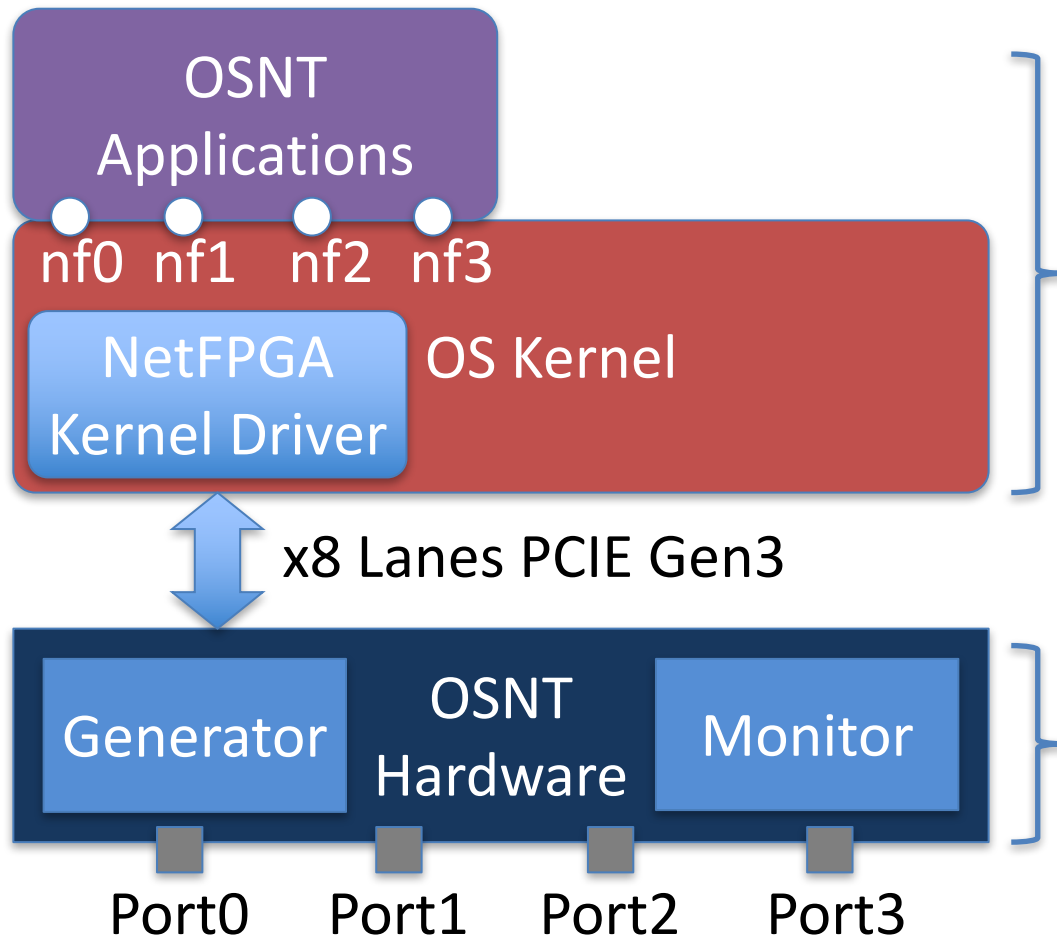
Low cost, flexible to update, scale-out, no CPU usage, nanosecond resolution measurements

# Network Tester Comparison

	Cost	Flexibility	Precision	Line Rate
   endace power to see all	\$\$\$ \$\$\$			
DPDK, Moongen	(\$)			(  )
	\$			

# Open Source Network Tester

- OSNT is an open source HW/SW platform for network testing



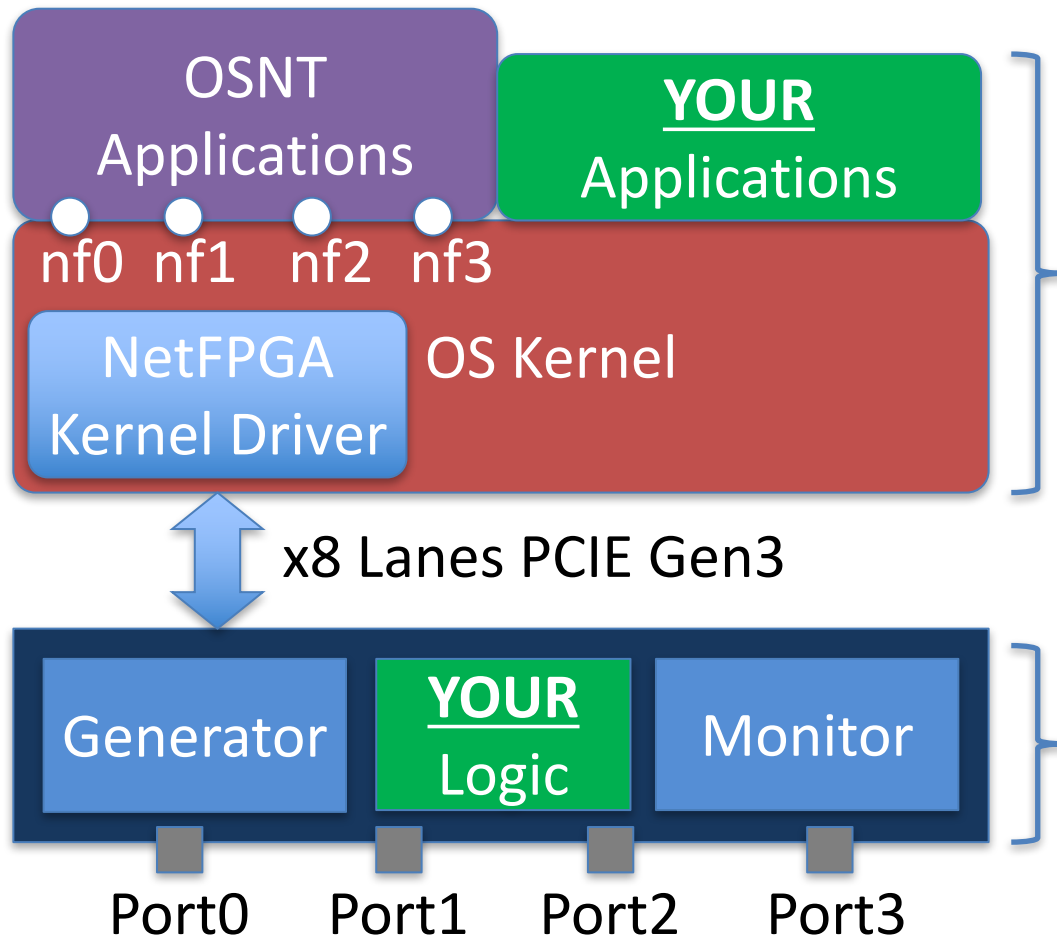
- Written in C, Python
- Open API and registers

- Written in Verilog HDL using standard Xilinx protocols

- Users can add and modify the modules

# Open Source Network Tester

- OSNT is an open source HW/SW platform for network testing



- Written in C, Python
- Open API and registers

- Written in Verilog HDL using standard Xilinx protocols
- Users can add and modify the modules

# Open Source Network Tester

OSNT **currently** is:

- 4x10Gbps traffic generator.
- Capture card with high resolution timestamp (6.4nsec).
- GPS-ready synchronized measurement kit.

# Open Source Network Tester

OSNT **currently** is:

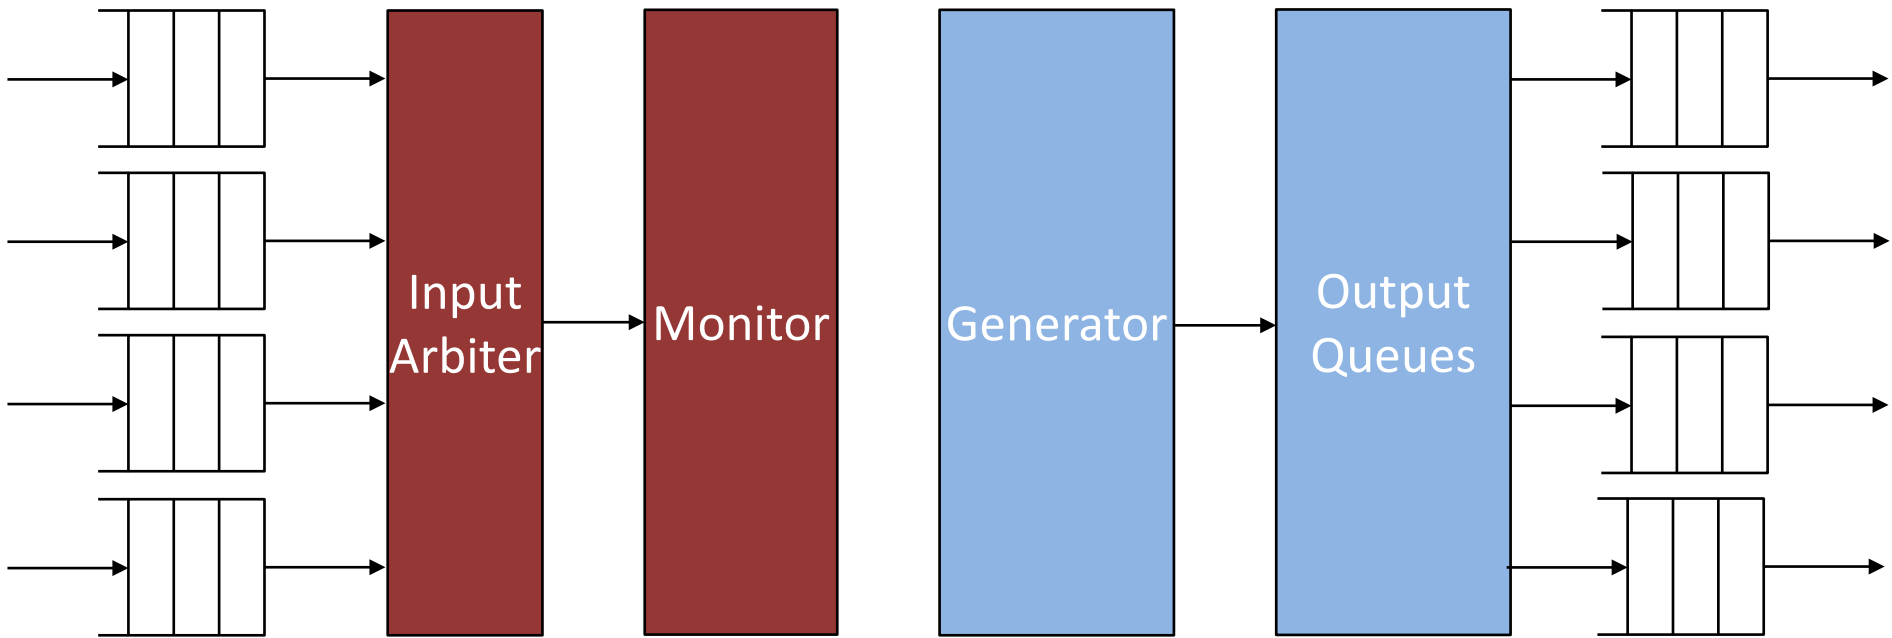
- 4x10Gbps traffic generator.
- Capture card with high resolution timestamp (6.4nsec).
- GPS-ready synchronized measurement kit.

**a starting point**

# OSNT Architecture

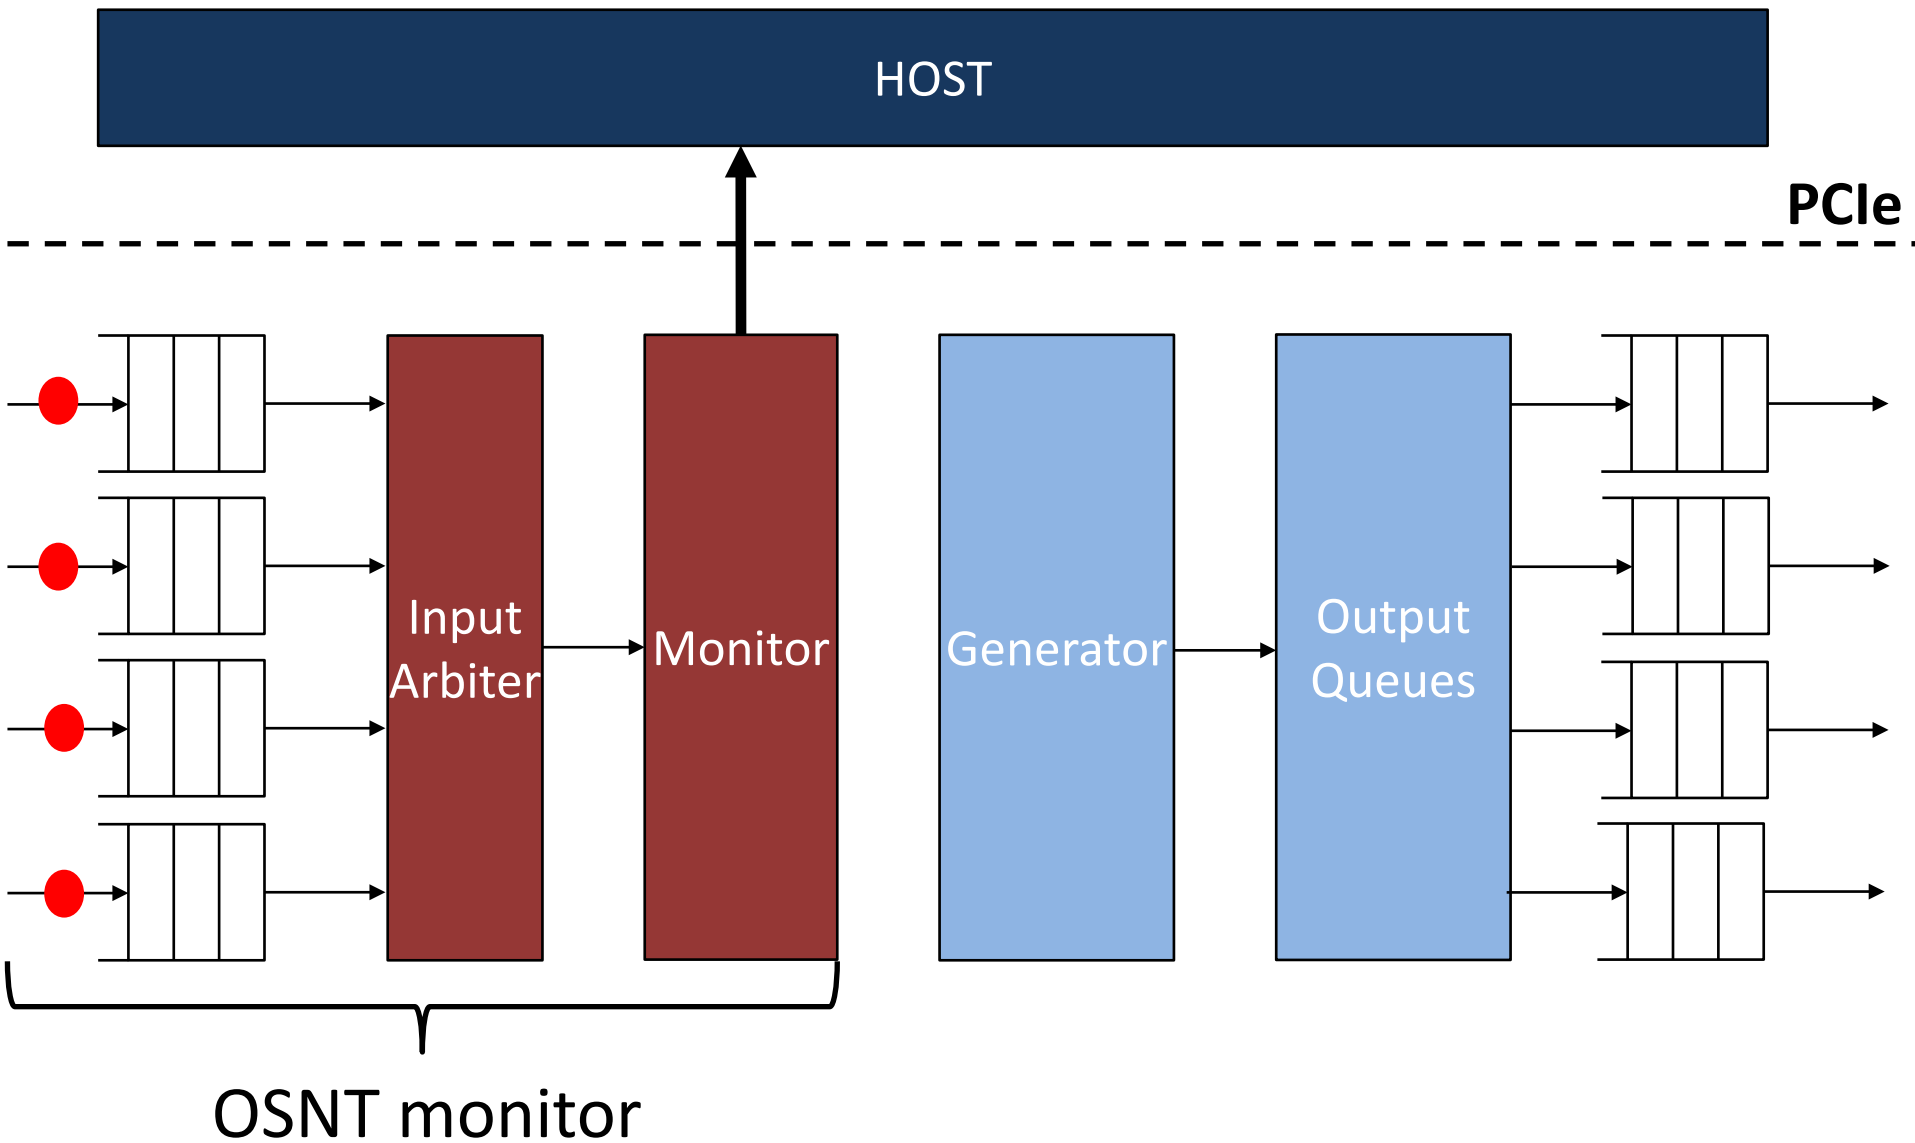
HOST

PCIe

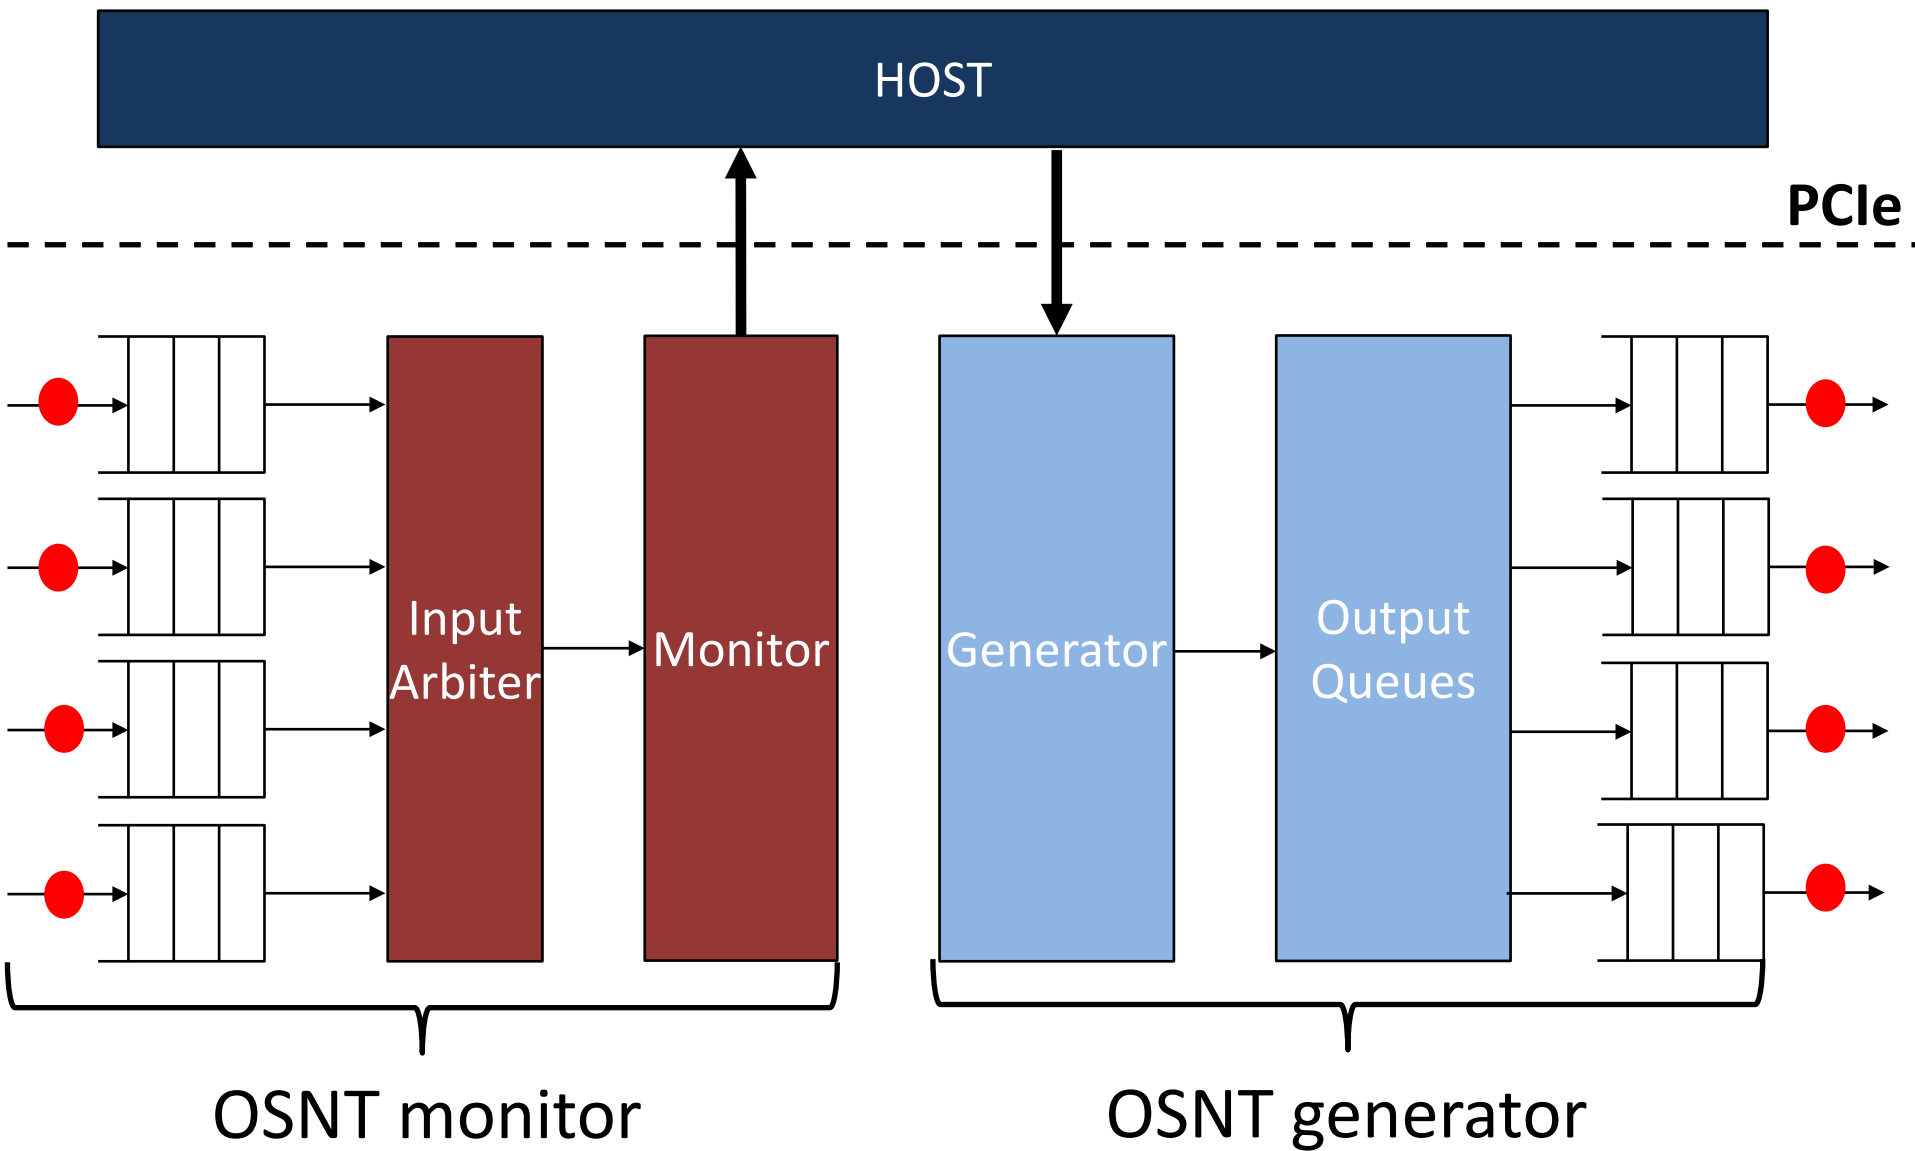




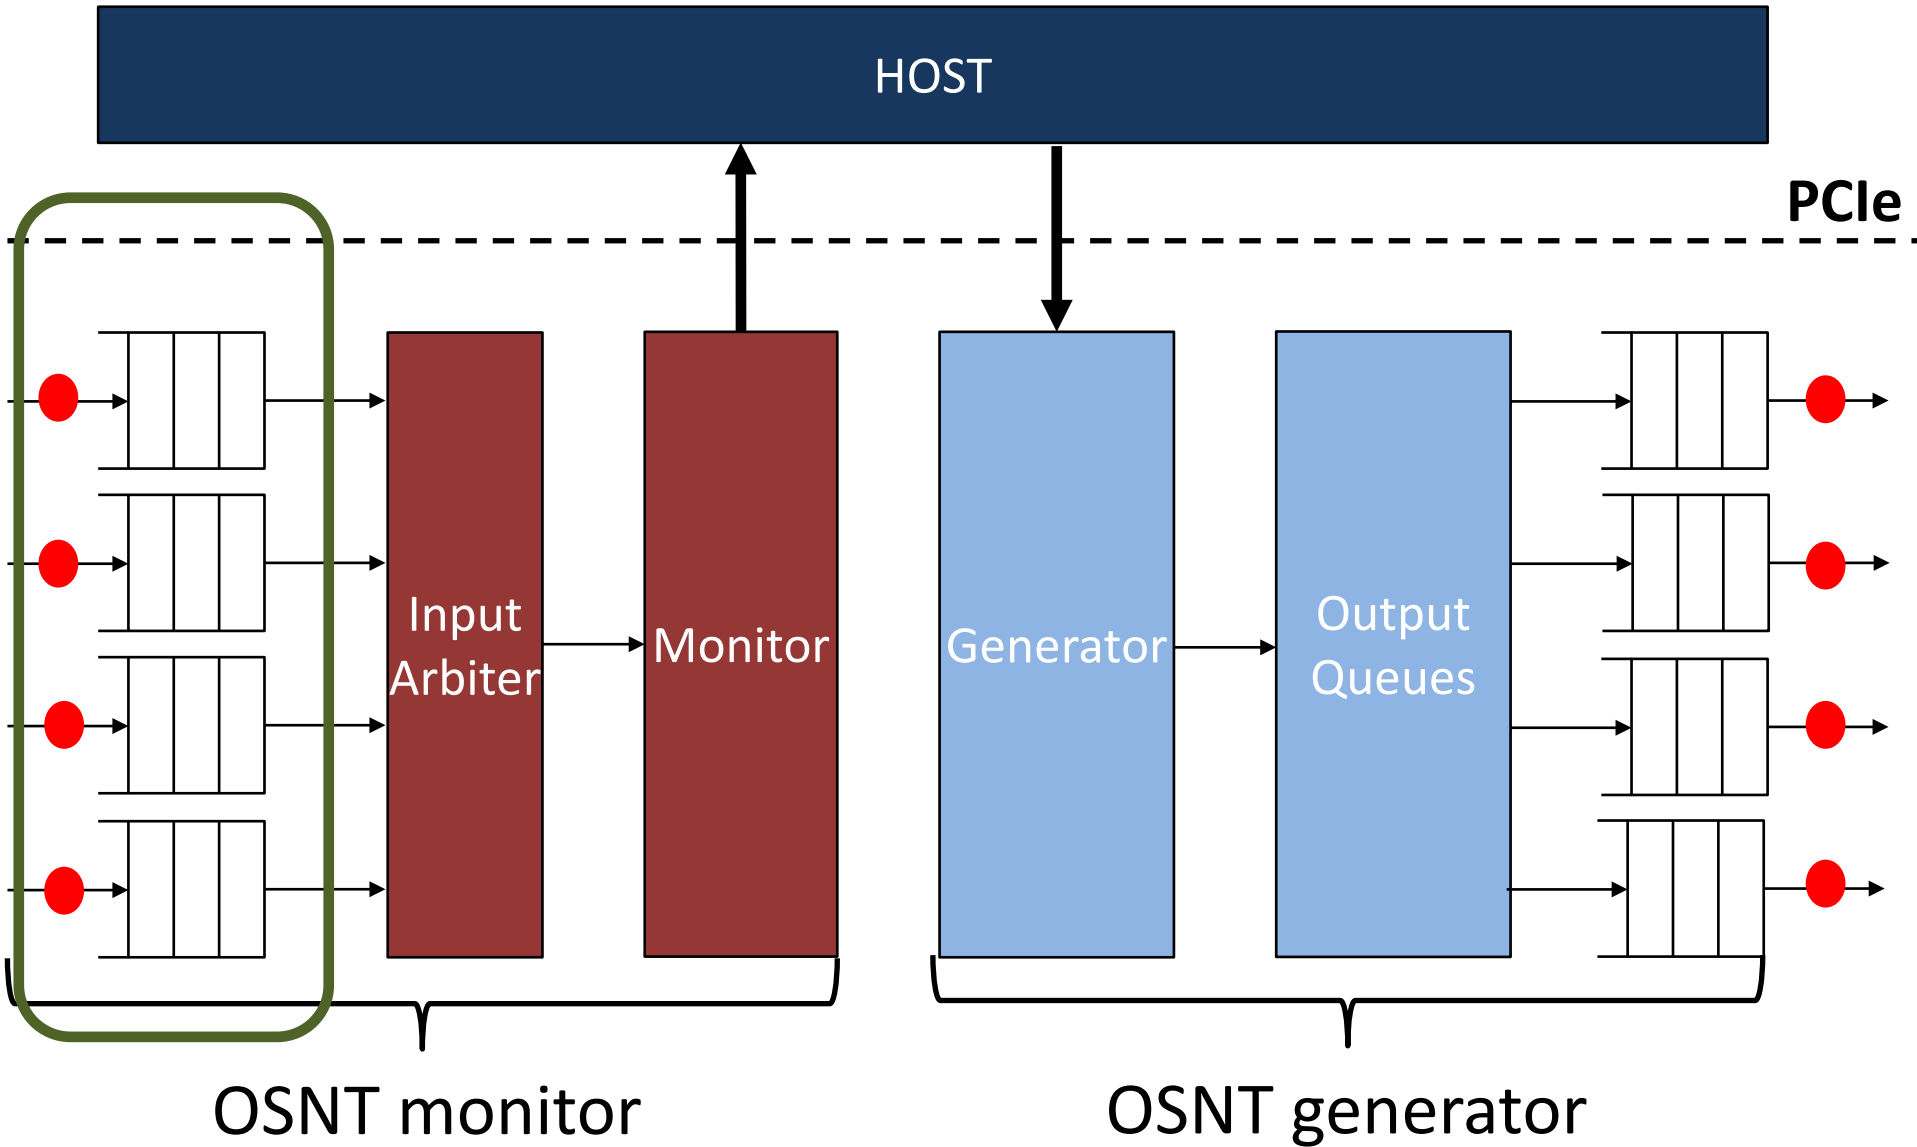
# OSNT Architecture



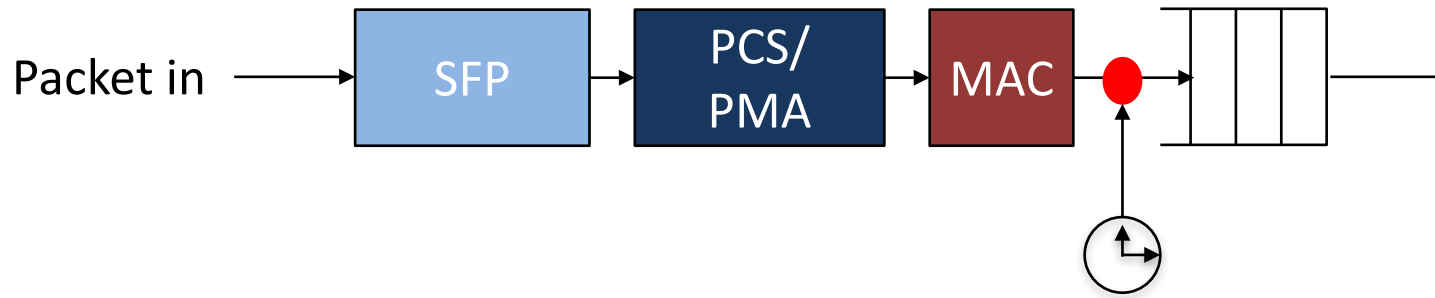
# OSNT Architecture



# OSNT Architecture

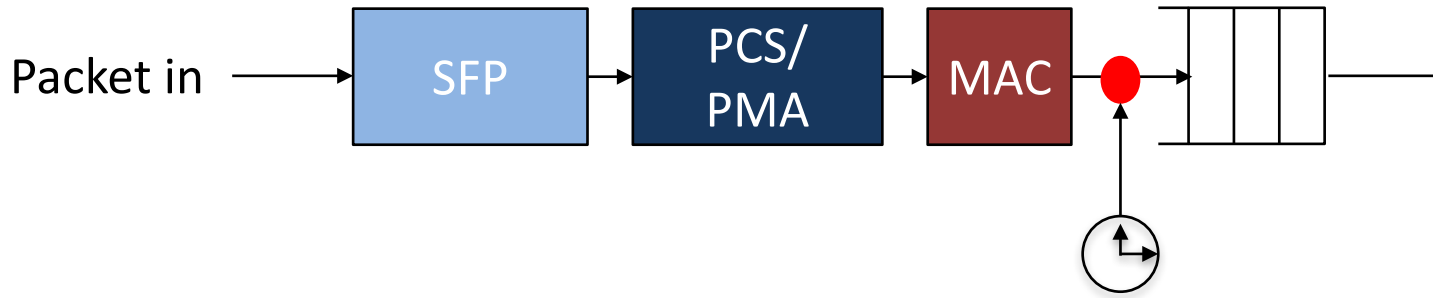


# OSNT RX interface



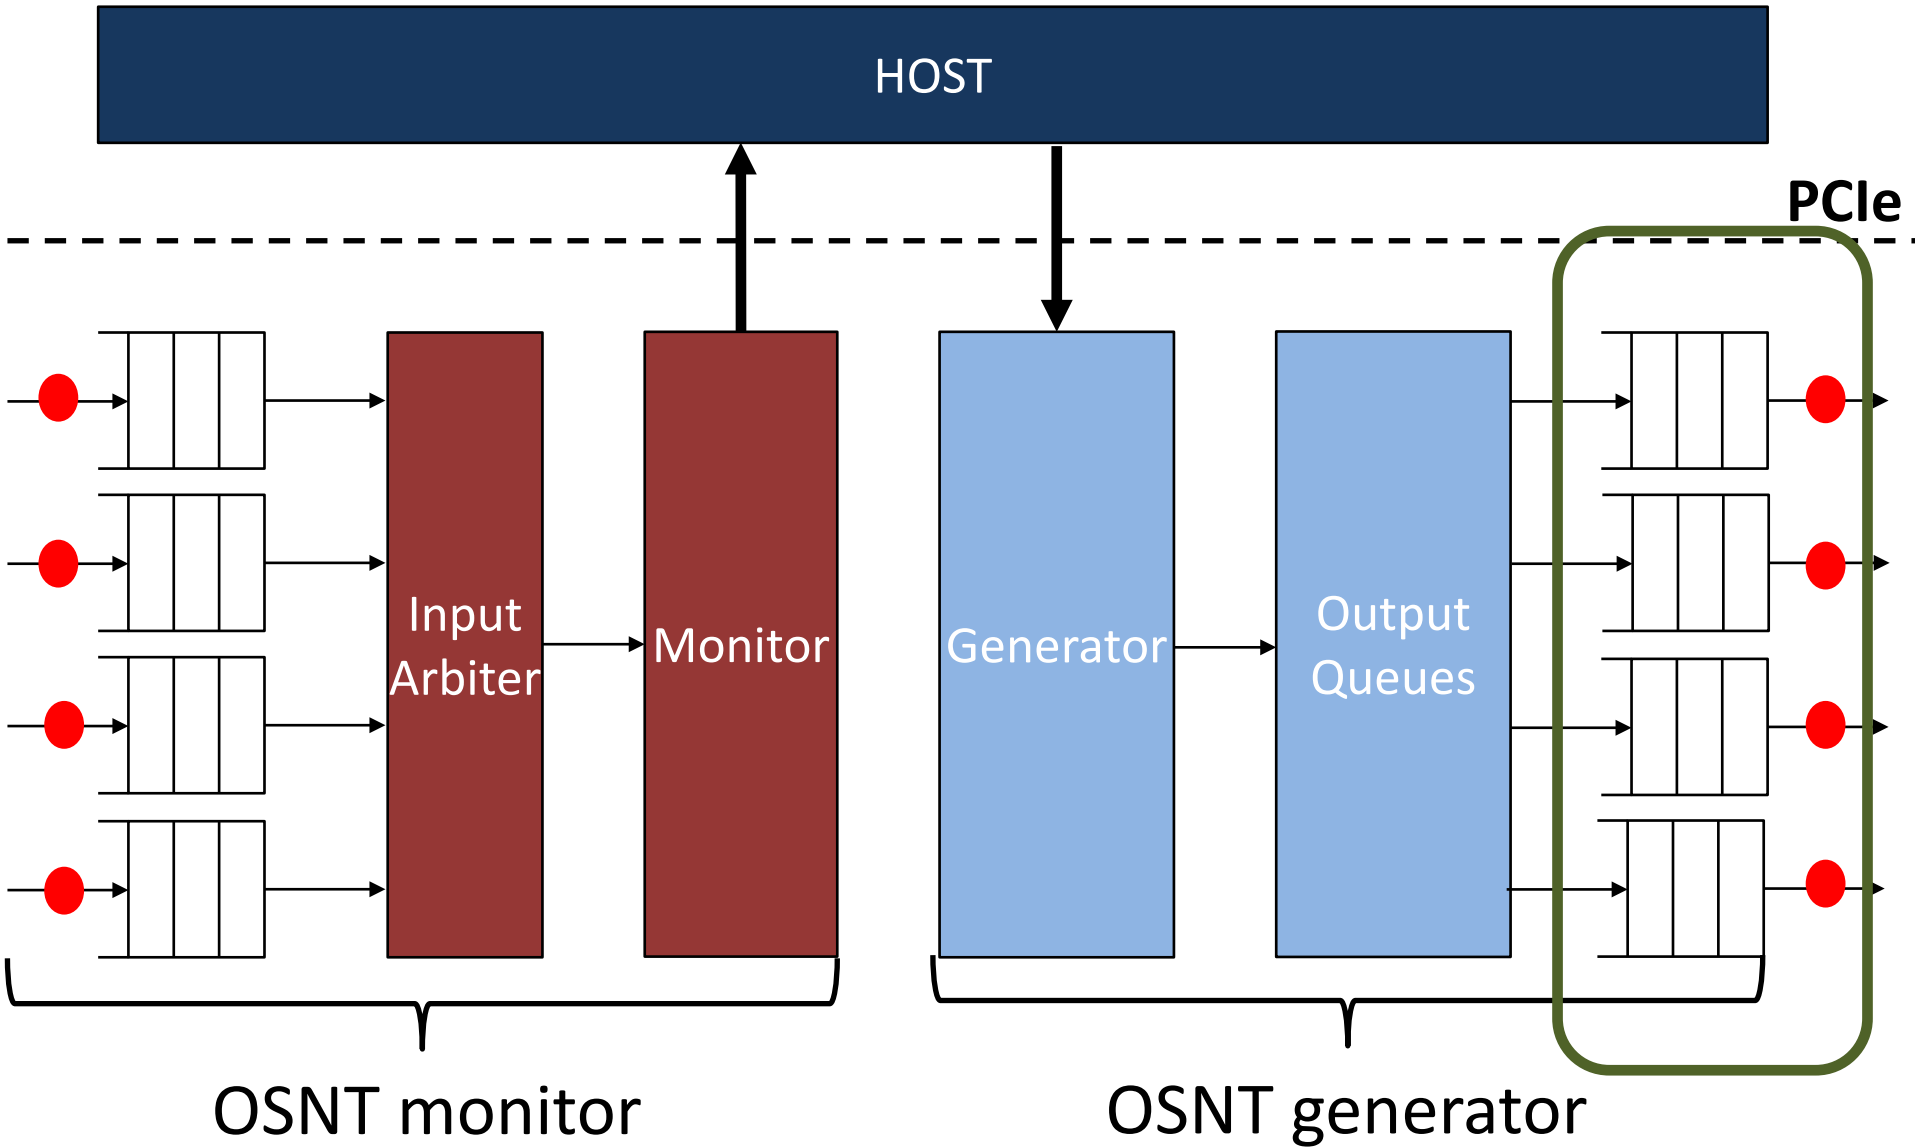
- Timestamp taken before RX queues to reduce FIFO-induced jitter

# OSNT RX interface

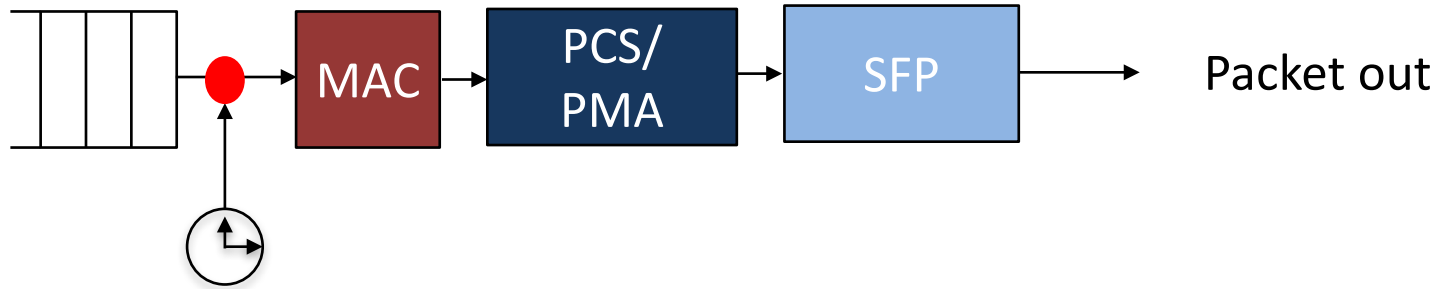


- Timestamp taken before RX queues to reduce FIFO-induced jitter
- Timestamp overwrites packet data at a configurable offset

# OSNT Architecture




# OSNT TX interface



- Timestamp taken after TX queues to reduce FIFO-induced jitter

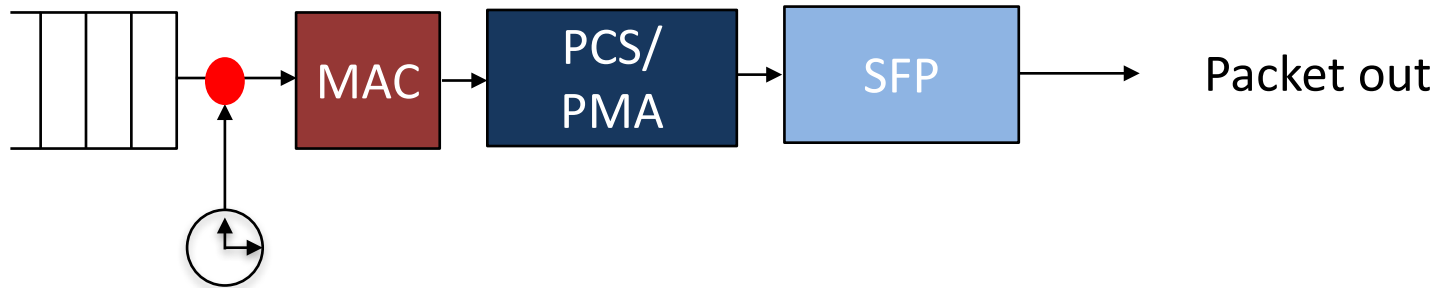
# OSNT TX interface



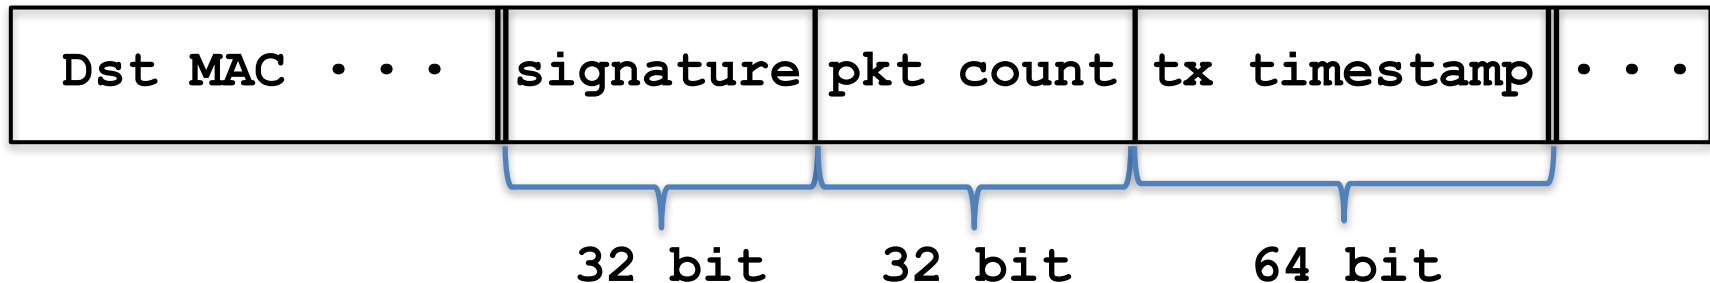
- Timestamp taken after TX queues to reduce FIFO-induced jitter
- Timestamp overwrites packet data at a configurable offset



# OSNT TX interface



- Timestamp taken after TX queues to reduce FIFO-induced jitter
- Timestamp overwrites packet data at a configurable offset
- If enabled, it will overwrite 128bit data:



# OSNT Timestamp

- Free running counter?
- We could use a 64-bit counter driven by the 156.25MHz system clock (naïve solution)



# OSNT Timestamp

- Free running counter?
- We could use a 64-bit counter driven by the 156.25MHz system clock (naïve solution)
  - provides no means by which to correct oscillator frequency drift



# OSNT Timestamp

- Free running counter?
- We could use a 64-bit counter driven by the 156.25MHz system clock (naïve solution)
  - provides no means by which to correct oscillator frequency drift
  - produces timestamps expressed in unit of 6.4 ns



# OSNT Timestamp

- Free running counter?
- We could use a 64-bit counter driven by the 156.25MHz system clock (naïve solution)
  - provides no means by which to correct oscillator frequency drift
  - produces timestamps expressed in unit of 6.4 ns
  - fixed-point representation of time in seconds more useful to host



# OSNT Timestamp

- Direct Digital Synthesis (DDS) is the **solution**!!
- DDS is a technique by which arbitrary variable frequencies can be generated



# OSNT Timestamp

- Direct Digital Synthesis (DDS) is the **solution**!!
- DDS is a technique by which arbitrary variable frequencies can be generated
  - need a time reference to correct DDS rate (the GPS provides long-term stability)



# OSNT Timestamp

- Direct Digital Synthesis (DDS) is the **solution**!!
- DDS is a technique by which arbitrary variable frequencies can be generated
  - need a time reference to correct DDS rate (the GPS provides long-term stability)
  - allow 64 bit value in fixed-point representation



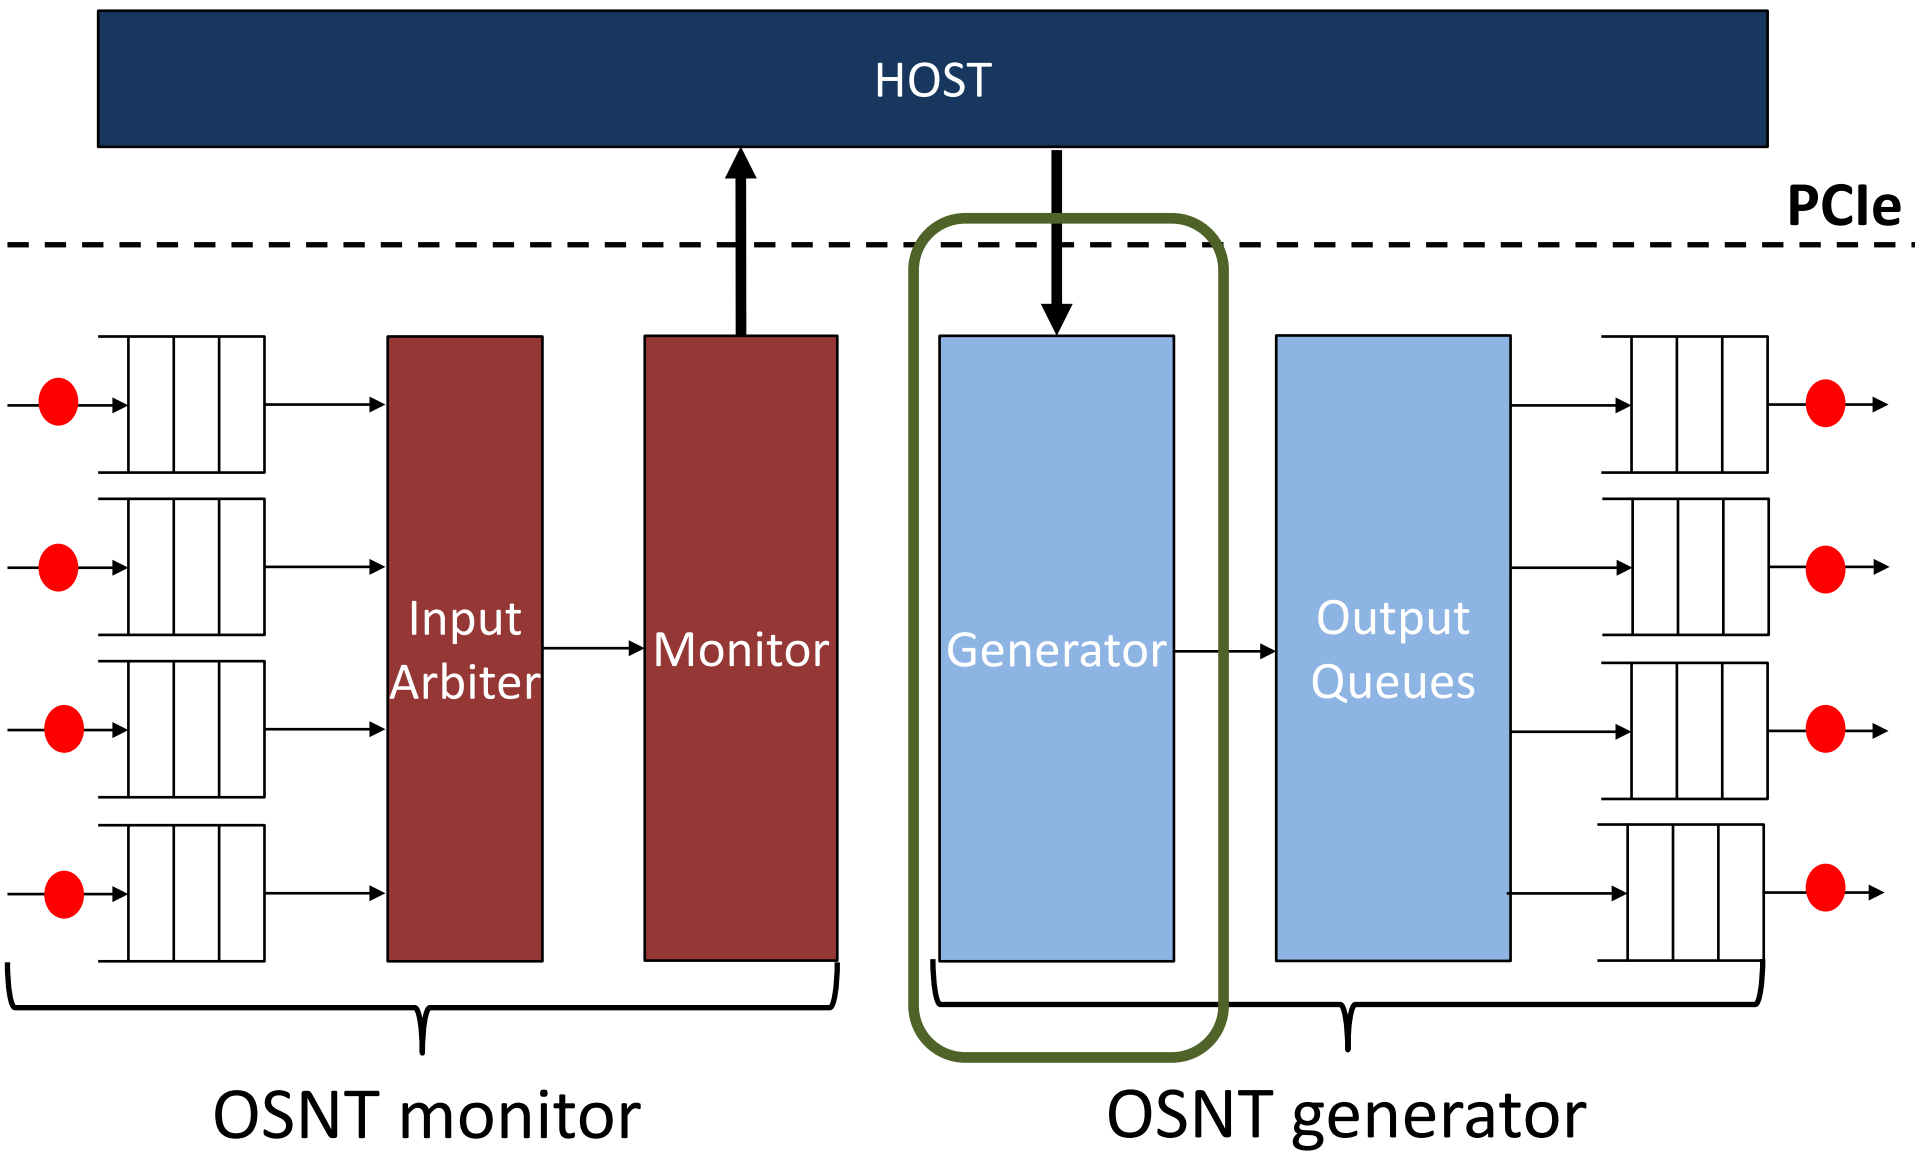


# OSNT Timestamp

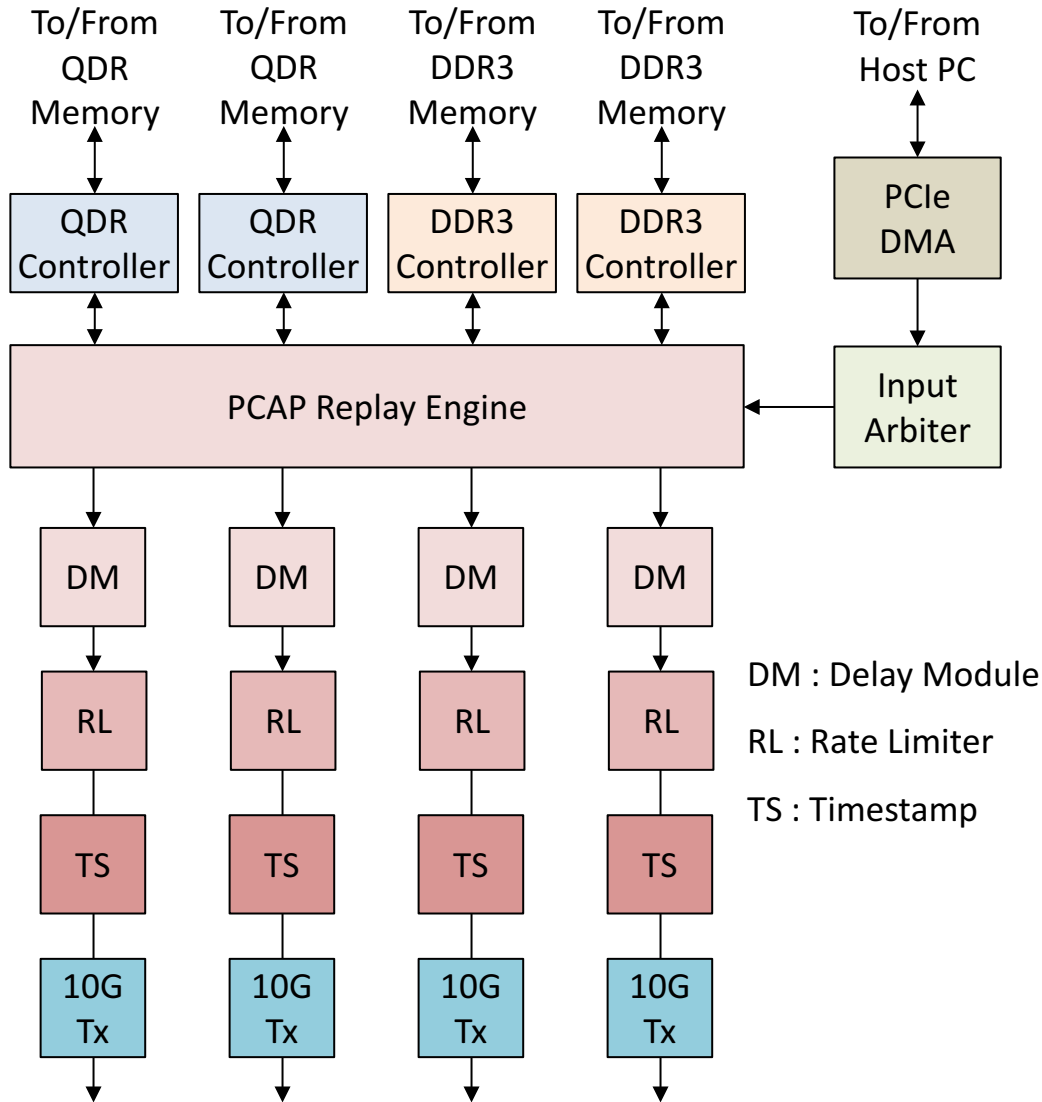
- Direct Digital Synthesis (DDS) is the **solution**!!
- DDS is a technique by which arbitrary variable frequencies can be generated
  - need a time reference to correct DDS rate (the GPS provides long-term stability)
  - allow 64 bit value in fixed-point representation
  - how Endace DAG card works!



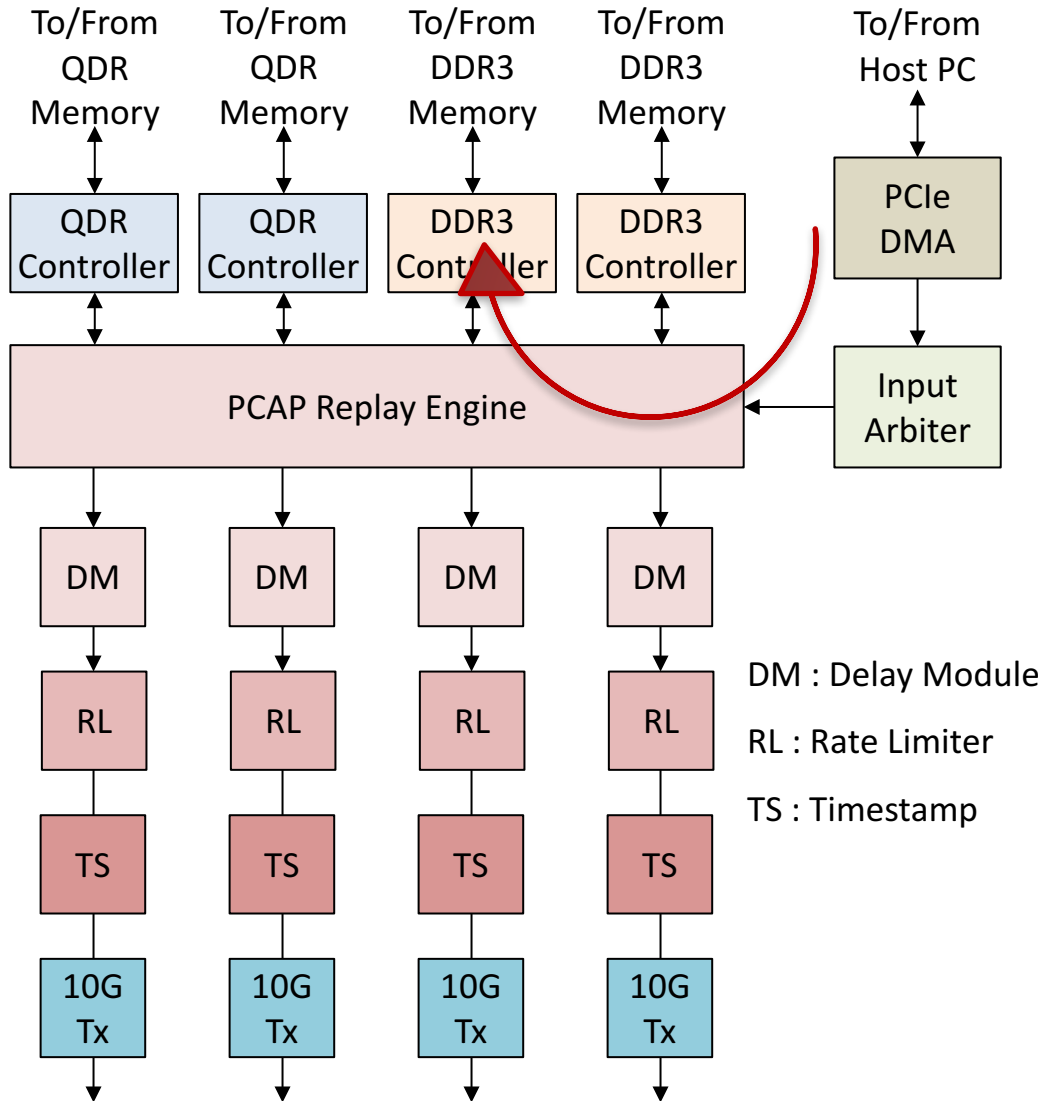
# OSNT Architecture



# OSNT Generator

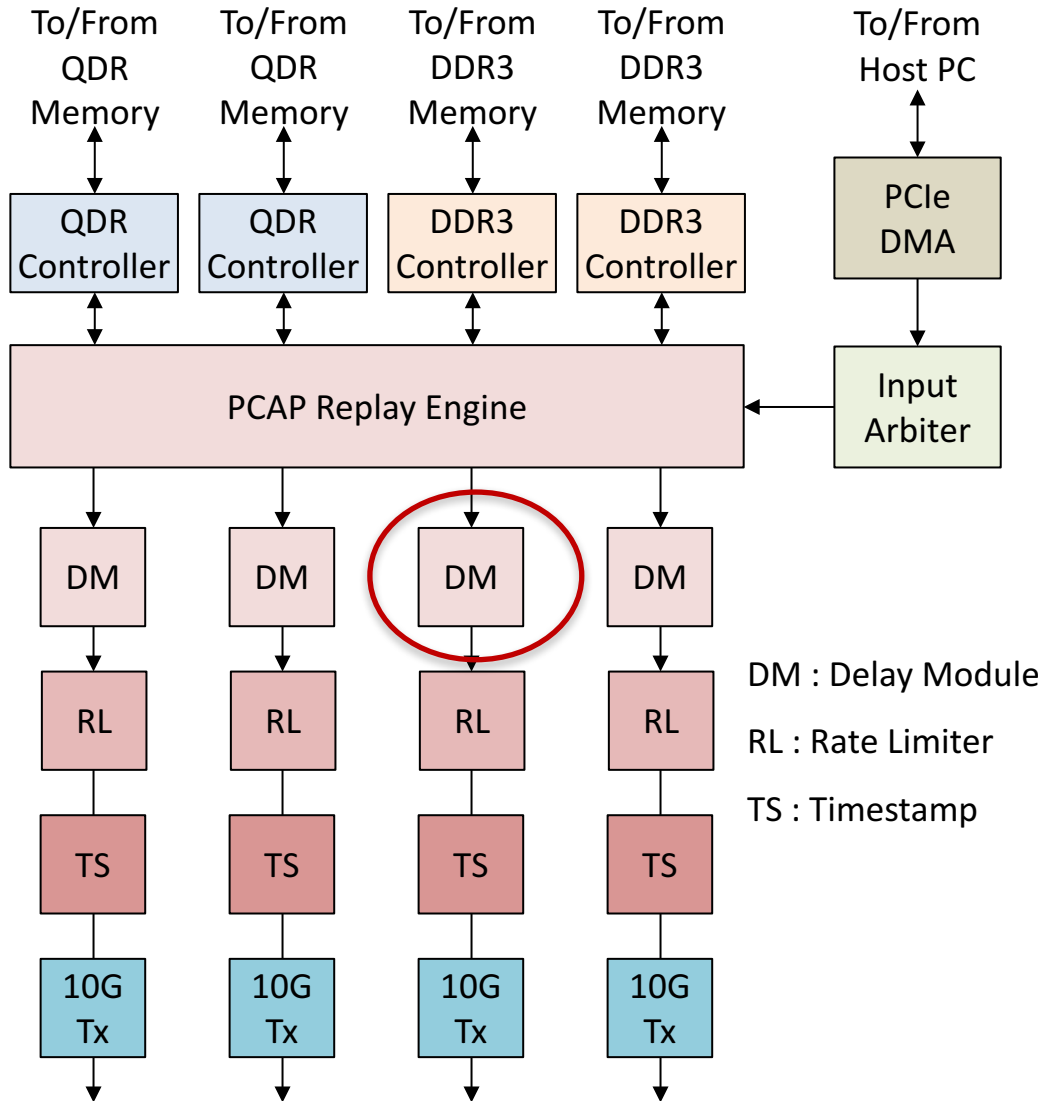


# OSNT Generator



- 4x10G PCAP replay engine
- SRAM: 27MB
- DRAM: 8GB

# OSNT Generator



- 4x10G PCAP replay engine
- SRAM: 27MB
- DRAM: 8GB

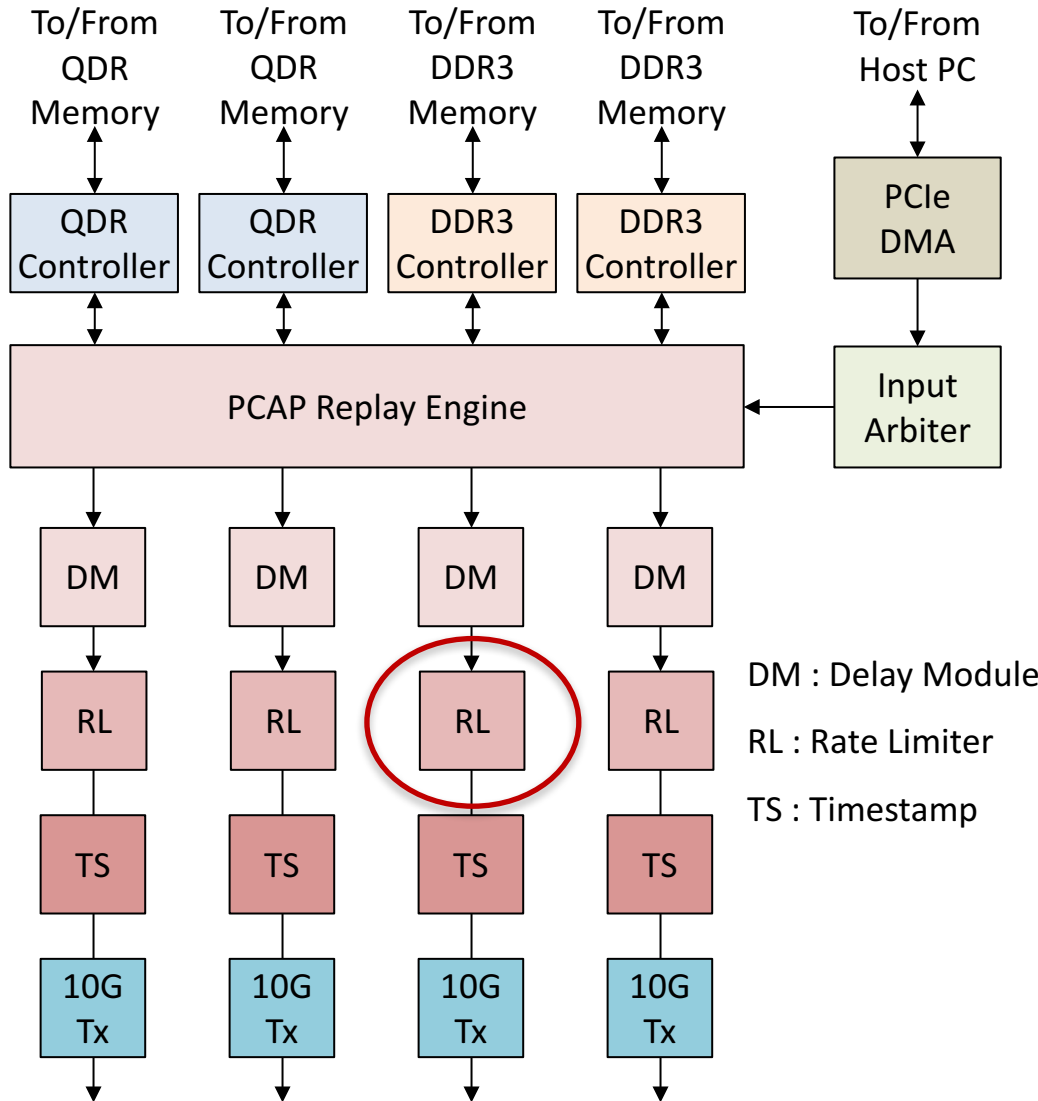
- Delay module

DM : Delay Module

RL : Rate Limiter

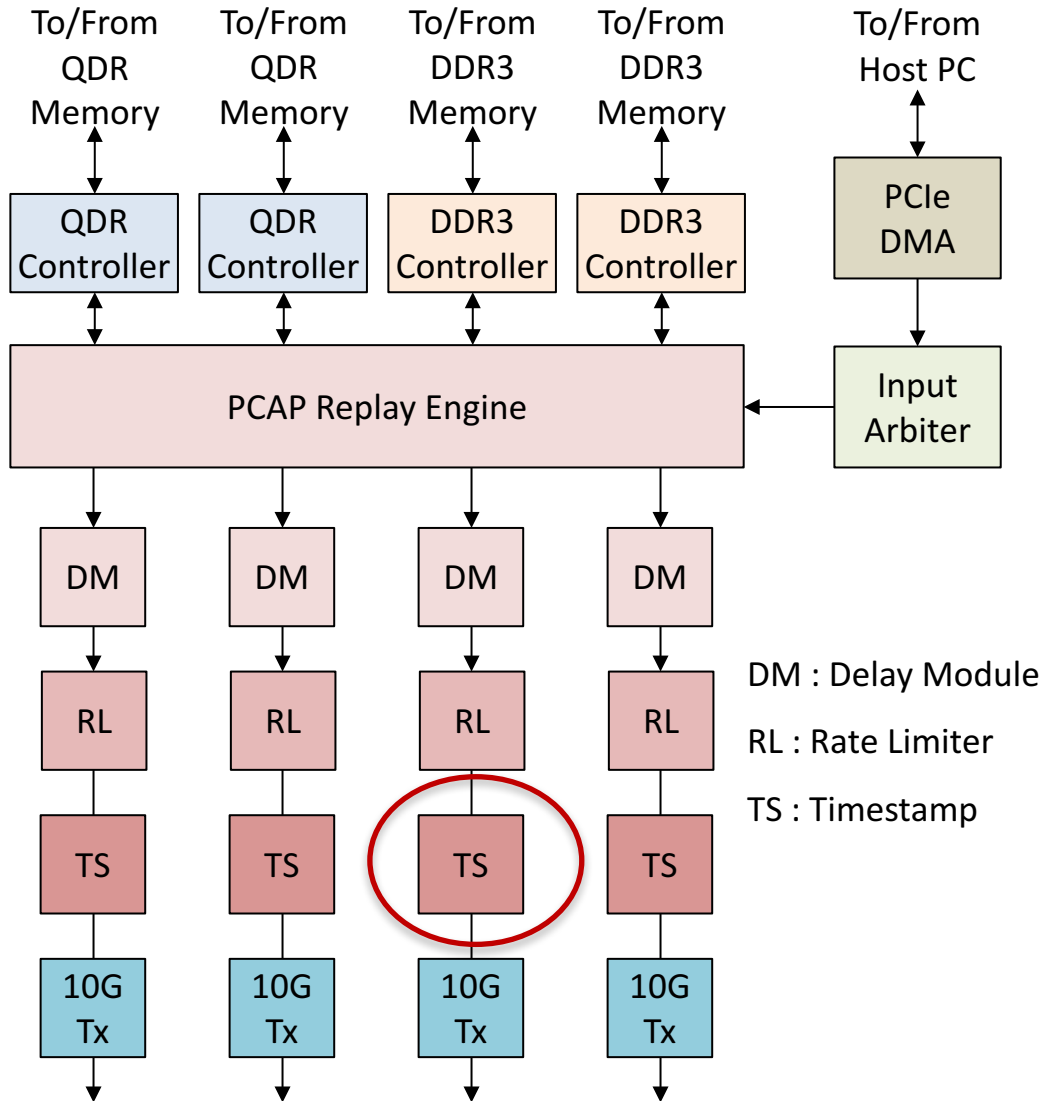
TS : Timestamp

# OSNT Generator



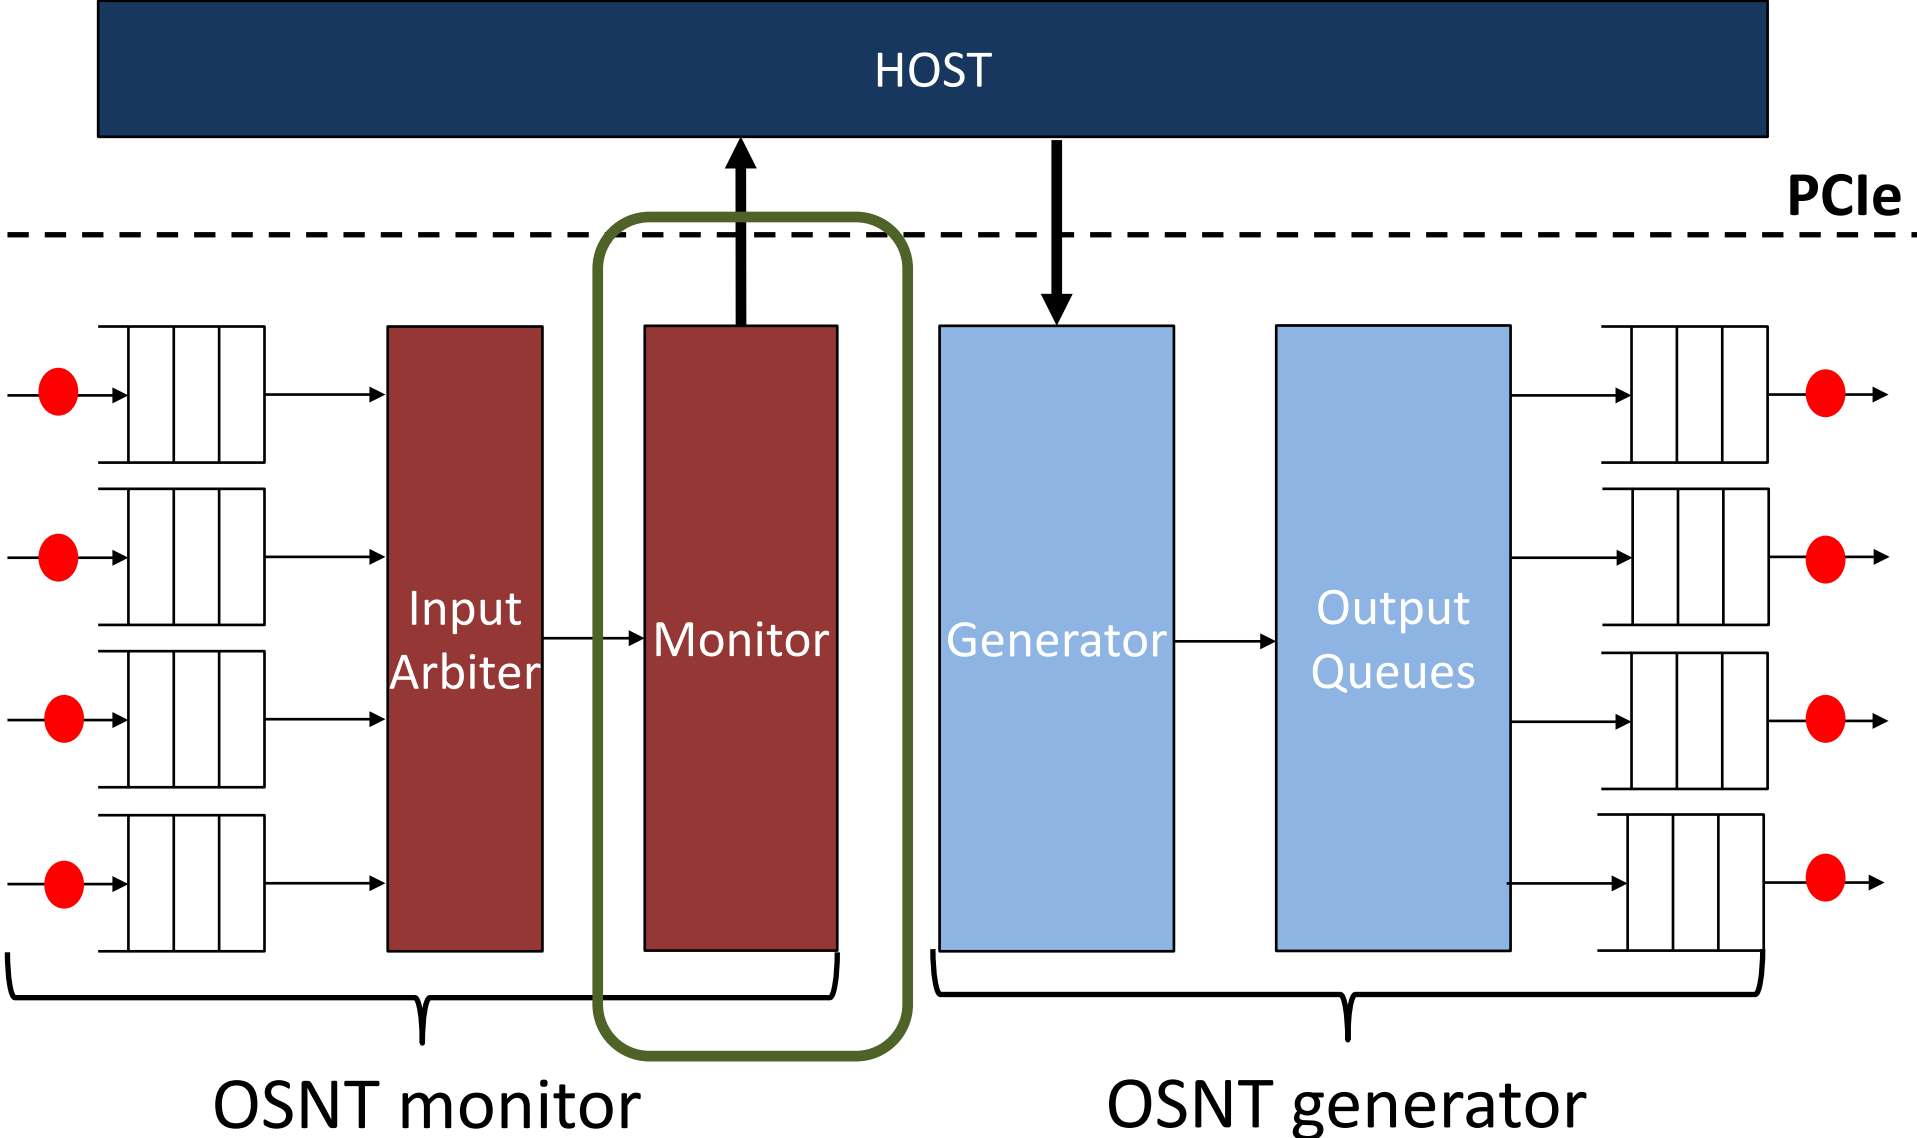
- 4x10G PCAP replay engine
- SRAM: 27MB
- DRAM: 8GB
  
- Delay module
- Rate limiter

# OSNT Generator



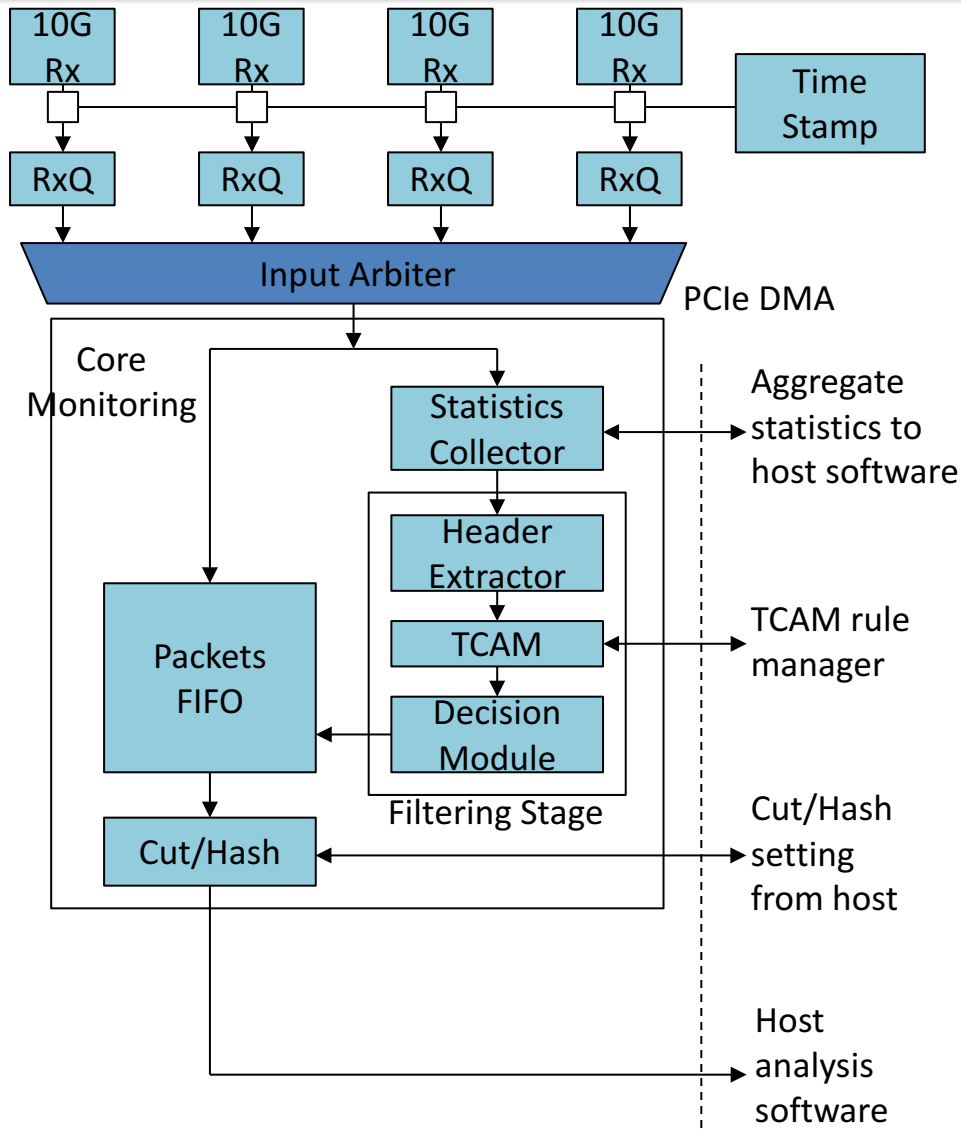
- 4x10G PCAP replay engine
- SRAM: 27MB
- DRAM: 8GB
  
- Delay module
- Rate limiter
- TX timestamping

# OSNT Architecture

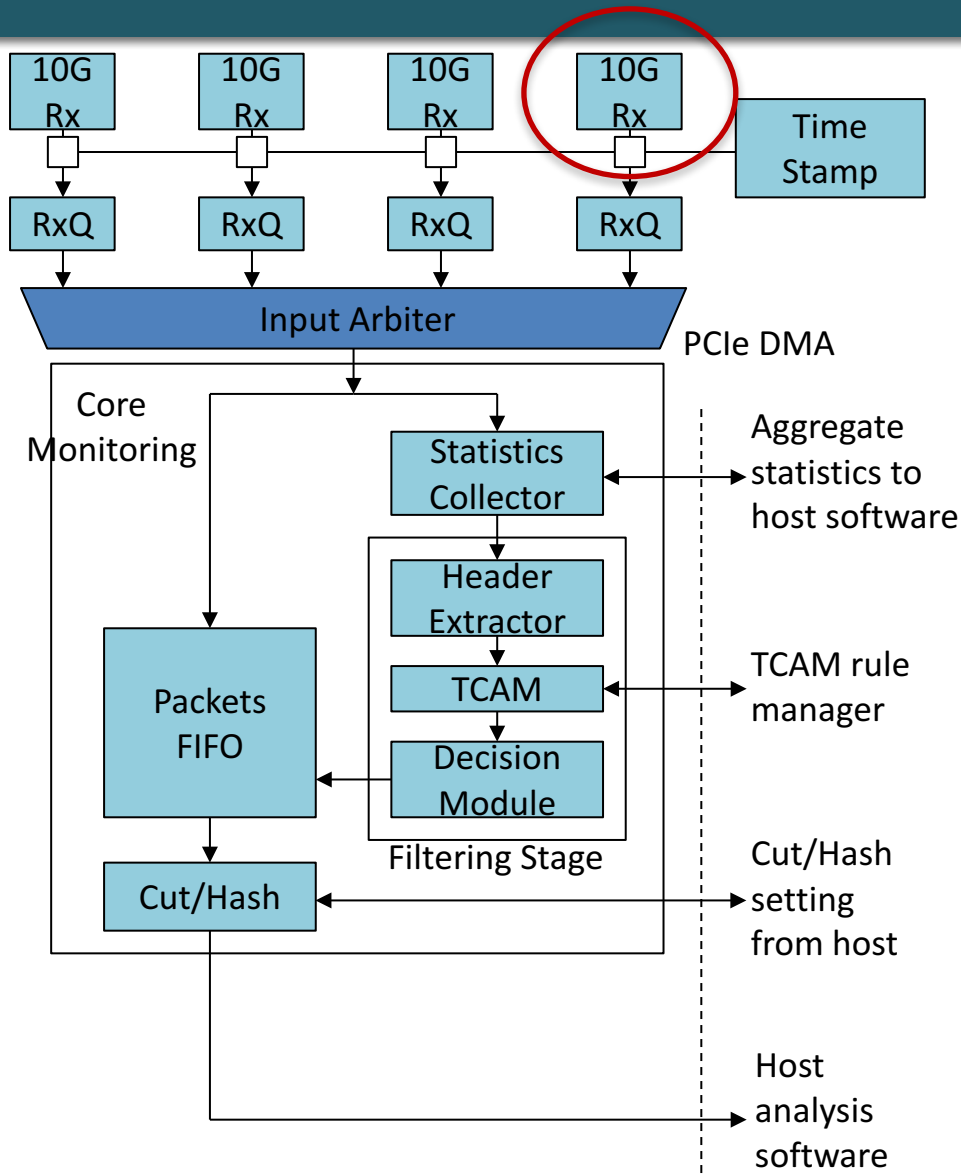




# OSNT Monitor

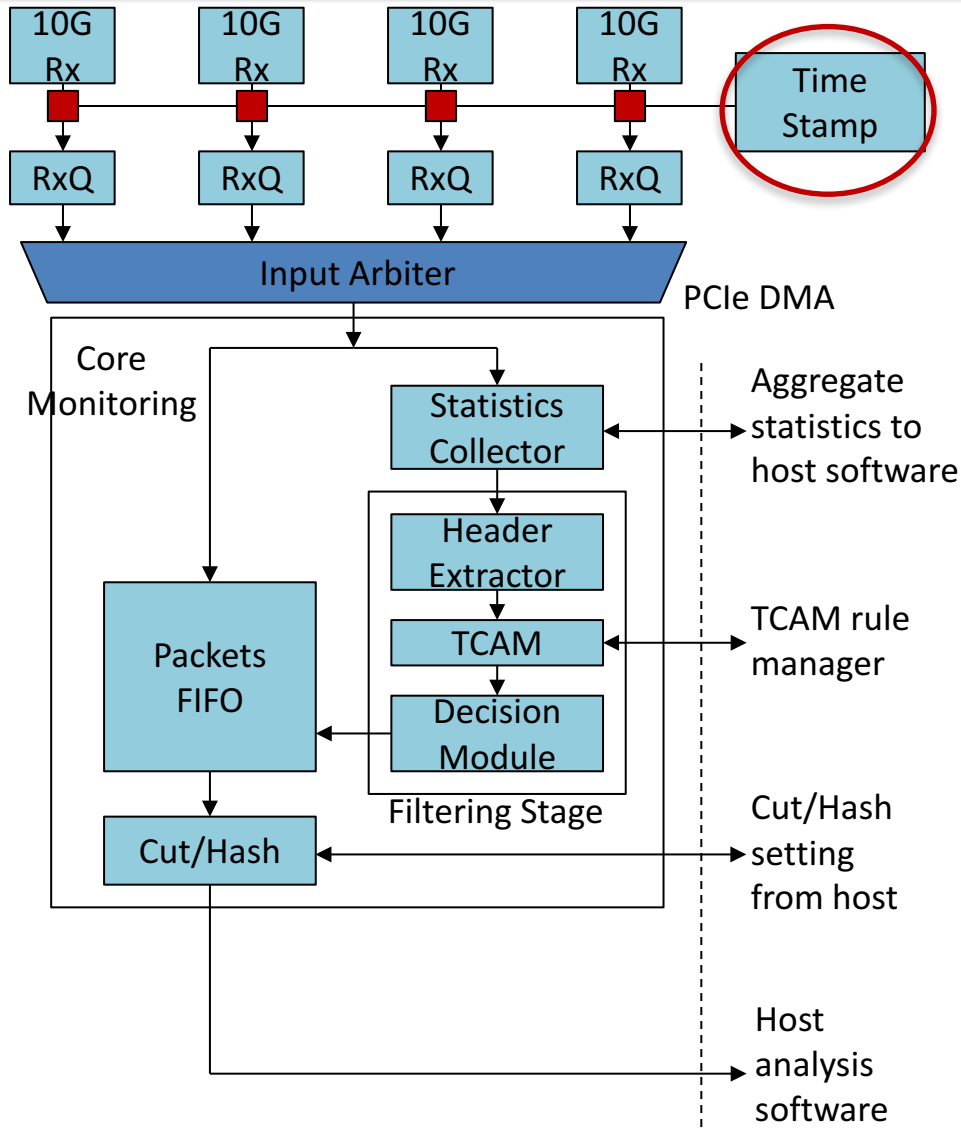


# OSNT Monitor



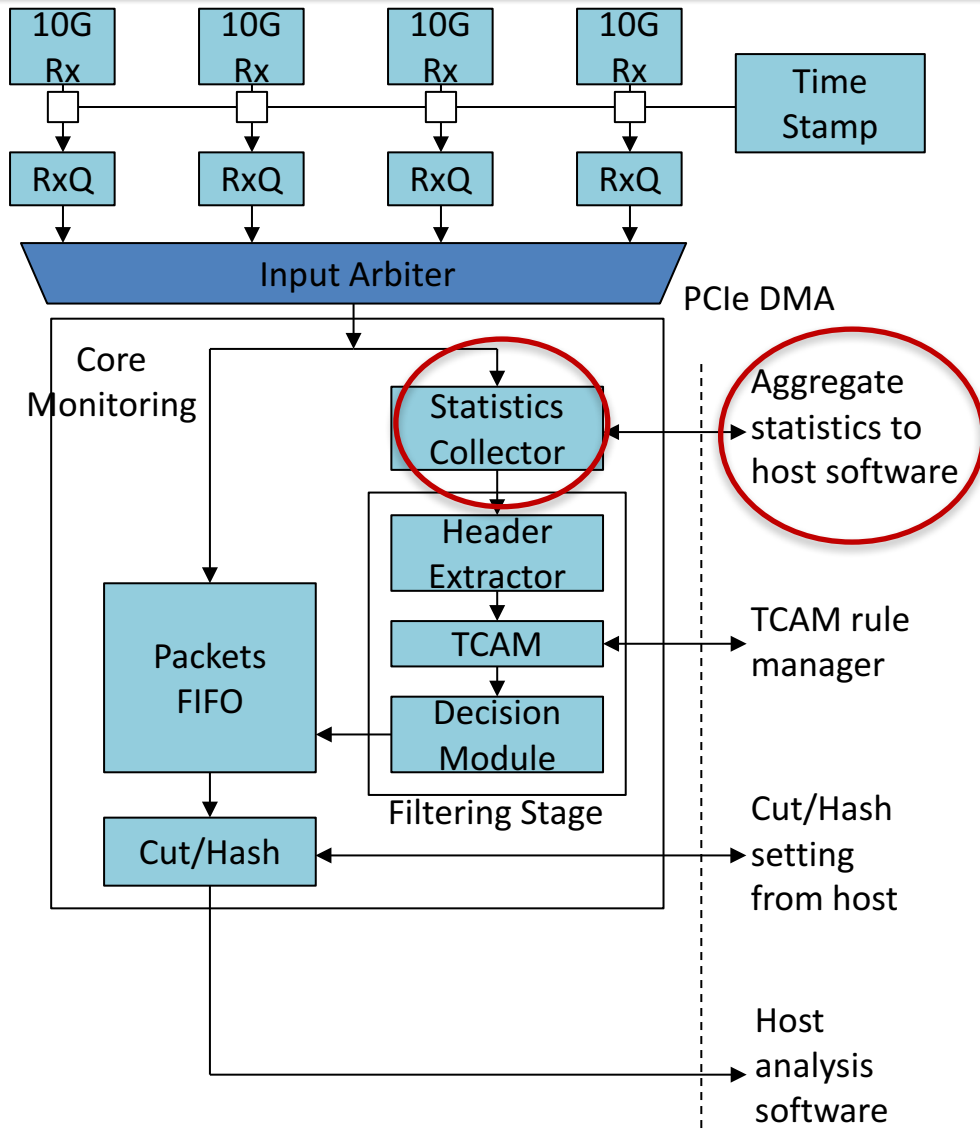
packet entering the board

# OSNT Monitor



- RX timestamp

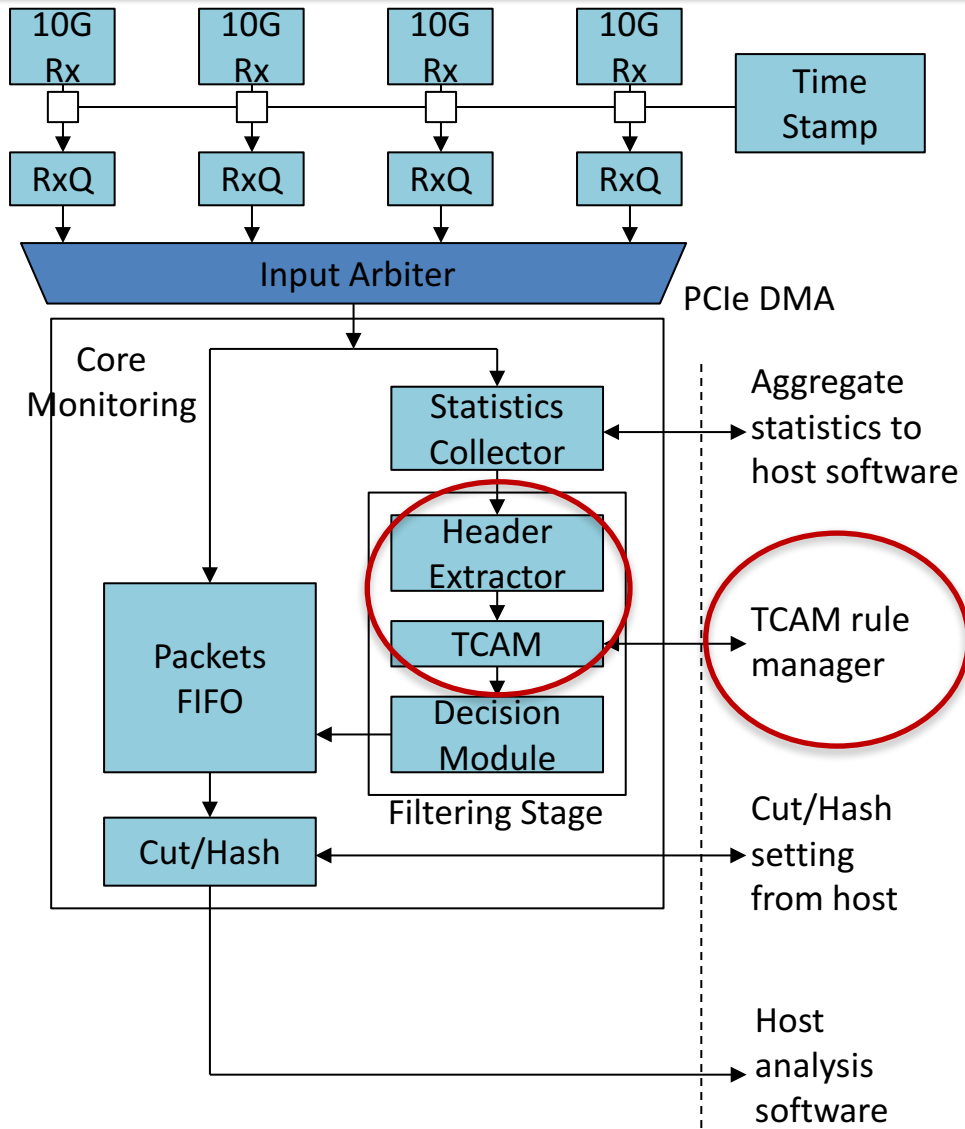
# OSNT Monitor



- RX timestamp

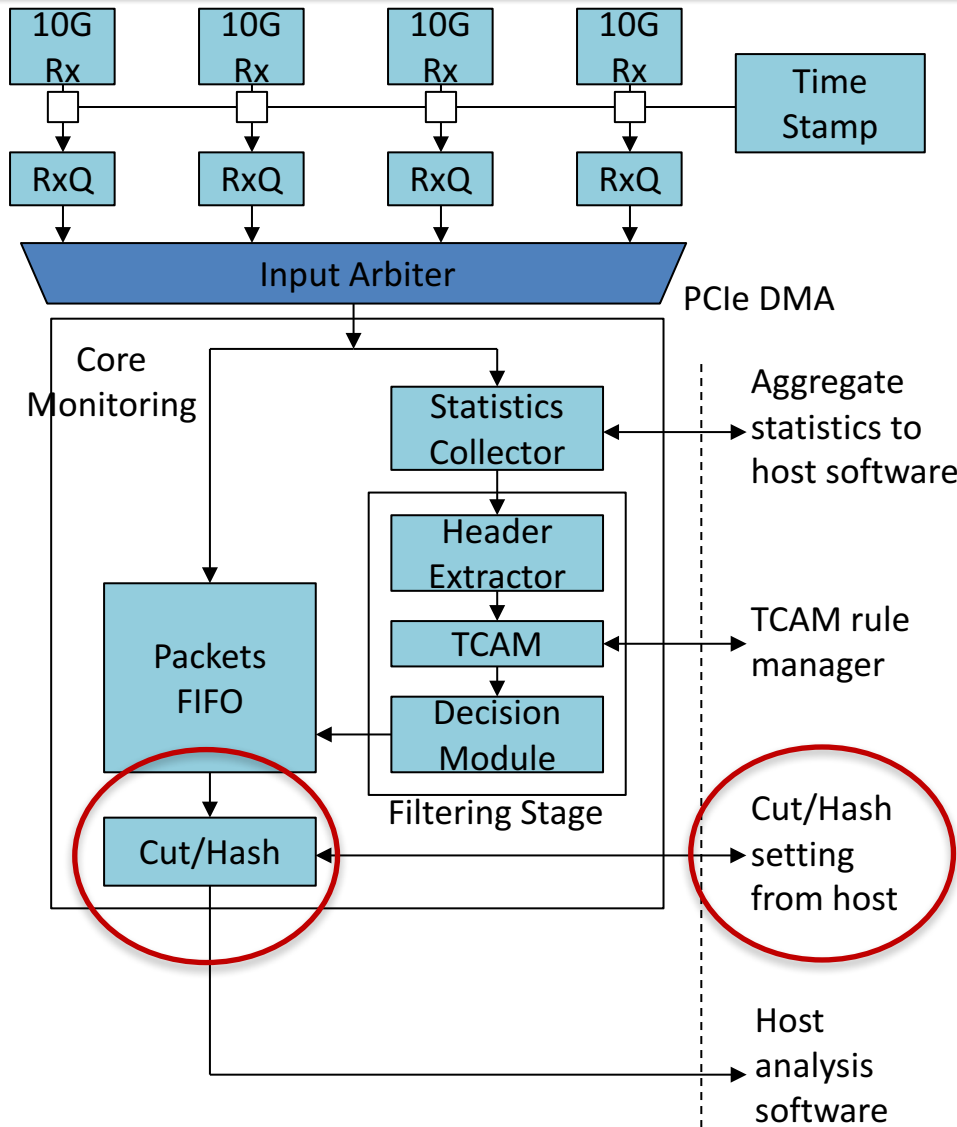
- Stats collector

# OSNT Monitor



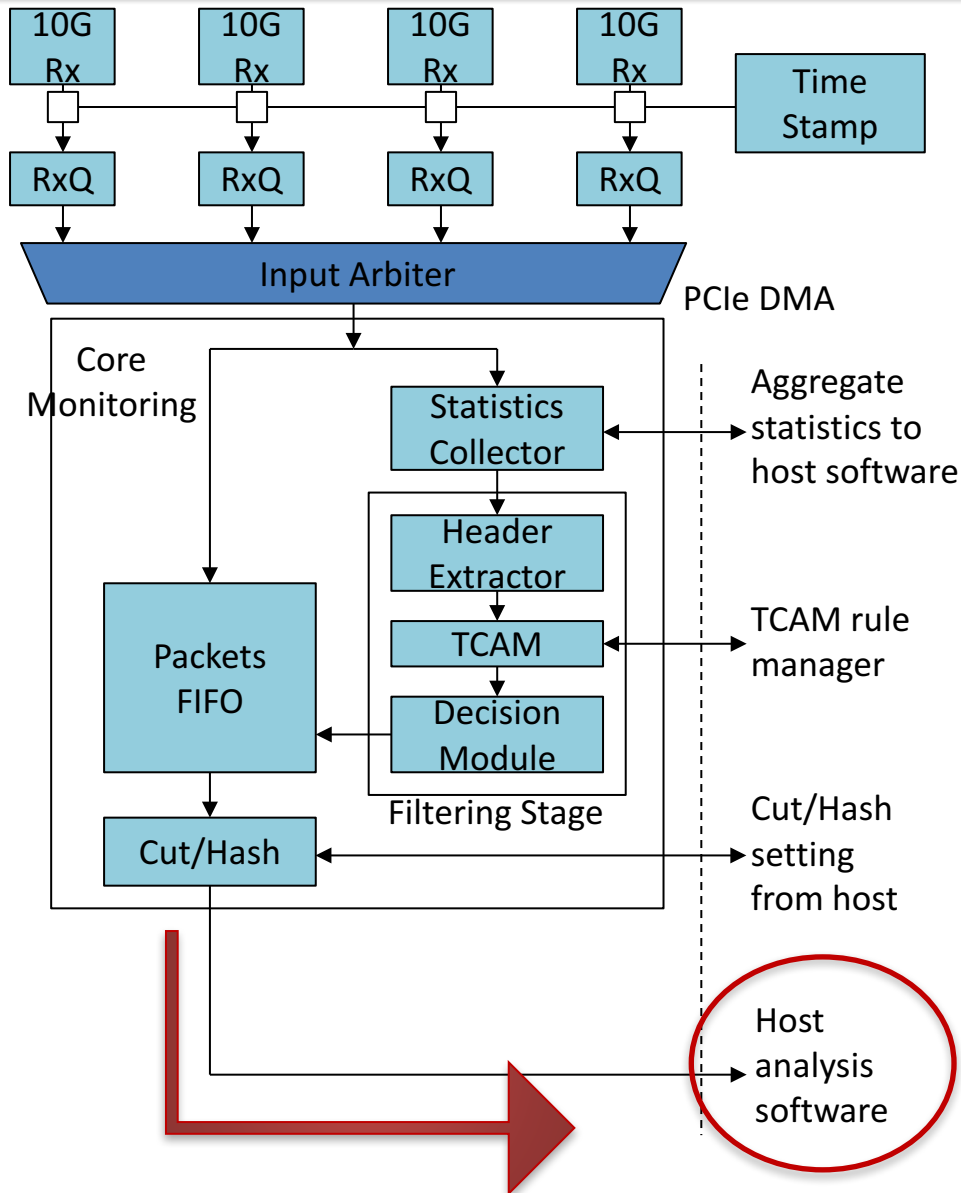
- RX timestamp
- Stats collector
- TCAM-based Packet Filter (5-tuple)

# OSNT Monitor



- RX timestamp
- Stats collector
- TCAM-based Packet Filter (5-tuple)
- Cut-Hash function

# OSNT Monitor



- RX timestamp
- Stats collector
- TCAM-based Packet Filter (5-tuple)
- Cut-Hash function

# OSNT GUI

- OSNT GUI – Extensible Generator and Monitor GUI in Python.
- Command-Line-Interface is also available.

**Generator GUI**

Console

**PCAP ENGINE**

Interface	Pcap File	Replay Cnt	Replay Cnt Display	Mem_addr_low	Mem_addr_high
0	1500.cap	100000000	100000000	0x0	0x30
1	Select Pcap File	0	0	0x30	0x30
2	Select Pcap File	0	0	0x30	0x30
3	Select Pcap File	0	0	0x30	0x30

**RATE LIMITER**

Interface	Rate Input	Rate Display	Enable	Reset
0	9.87Gbps	100.0000%	Enable	Reset
1	9.87Gbps	100.0000%	Enable	Reset
2	9.87Gbps	100.0000%	Enable	Reset
3	9.87Gbps	100.0000%	Enable	Reset

**INTER PACKET DELAY**

Interface	Delay Source	Delay Reg Input	Delay Reg Display	Enable	Reset
0	Set IPG	0	0ns	Enable	Reset
1	Set IPG	0	0ns	Enable	Reset
2	Set IPG	0	0ns	Enable	Reset
3	Set IPG	0	0ns	Enable	Reset

**Timestamp Rx and Tx Position**

Interface	RX TS Pos	TX TS Pos
0	0	0
1	0	0
2	0	0
3	0	0

**Monitor GUI**

Console

**STATS**

Port	Pkt Cnt	Vlan Cnt	IP Cnt	UDP Cnt	TCP Cnt	Pkts/s	Bits/s
0	12061355	0	12061355	12061355	0	822.368K	9.868G
1	0	0	0	0	0	0.0	0.0
2	0	0	0	0	0	0.0	0.0
3	0	0	0	0	0	0.0	0.0

**FILTER RULES**

Entry	SRC IP	SRC IP MASK	DST IP	DST IP MASK	L4 PORT	L4 PORT MASK	PROTO	PROTO MASK
0	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
1	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
2	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
3	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
4	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
5	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
6	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
7	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
8	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
9	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
10	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A

**CUTTER and TIMER**

Cut to Length: N/A      FPGA Timer: 183.13438784      Hash



# OSNT command line

- Command-Line-Interface is available to create a script automating the test process.

```
/root/osnt-sume-master/OSNT-SUME-live/projects/osnt/sw/host/app/tools
/root/osnt-sume-master/OSNT-SUME-live/projects/osnt/sw/host/app/tools 113x40
total 6256
drwxr-xr-x 2 root root 4096 Apr 11 13:31 .
drwxr-xr-x 7 root root 4096 Mar 25 10:53 ..
-rw-r--r-- 1 root root 2778 Apr 4 14:55 ext_mem_access.py
-rwxr-xr-x 1 root root 4624 Apr 11 13:30 gen_pcap_pkts.py
-rw-r--r-- 1 root root 0 Mar 25 10:53 __init__.py
-rw-r--r-- 1 root root 4749847 Apr 10 15:11 latency_dump_conv.pcap
-rw-r--r-- 1 root root 1512024 Apr 10 15:11 latency_dump.pcap
-rwxr-xr-x 1 root root 2772 Apr 10 14:48 osnt-test-template-1-2.sh
-rw-r--r-- 1 root root 2520 Apr 4 15:26 osnt-test-template-3.sh
-rw-r--r-- 1 root root 2696 Apr 4 15:25 osnt-test-template-5.sh
-rw-r--r-- 1 root root 2612 Apr 4 15:22 osnt-test-template-6.sh
-rwxr-xr-x 1 root root 2311 Mar 25 10:53 osnt-test-template-7.sh
-rw-r--r-- 1 root root 4 Apr 10 15:19 rd_reg32.dat
-rwxr-xr-x 1 root root 1665 Apr 11 13:31 run_pcap_gen.sh
-rw-r--r-- 1 root root 80364 Apr 10 15:11 tcpdump_latency_data.dat
-rw-r--r-- 1 root root 964 Apr 11 13:31 test_pcap_01.ccap
-rw-r--r-- 1 root root 3192 Mar 25 10:53 timestamp_capture_cli.py
root@nf-test111:tools$
root@nf-test111:tools$
root@nf-test111:tools$ python ../cli/osnt-tool-cmd.py -ifp0 ../sample_traces/1500.ccap -flt ../g
0 10000 -rpn0 1000 -txs0 6 -rxs0 7 -lpn 1000 -lty0 -rnm
```

```
/root/osnt-sume-master/OSNT-SUME-live/projects/osnt/sw/host/app/cli
/root/osnt-sume-master/OSNT-SUME-live/projects/osnt/sw/host/app/cli 113x40
OSNT Monitor Stats (SUME-NetFPGA)
nf0 =>
Packet No : 0      Byte No : 0
VLAN No  : 0      IP No  : 0      UDP No : 0      TCP No : 0
=====
Pkt/Sec  :0.0      Byte/Sec :0.0

nf1 =>
Packet No : 0      Byte No : 0
VLAN No  : 0      IP No  : 0      UDP No : 0      TCP No : 0
=====
Pkt/Sec  :0.0      Byte/Sec :0.0

nf2 =>
Packet No : 0      Byte No : 0
VLAN No  : 0      IP No  : 0      UDP No : 0      TCP No : 0
=====
Pkt/Sec  :0.0      Byte/Sec :0.0

nf3 =>
Packet No : 0      Byte No : 0
VLAN No  : 0      IP No  : 0      UDP No : 0      TCP No : 0
=====
Pkt/Sec  :0.0      Byte/Sec :0.0

OSNT TimsStamp Counter: 34.381888 sec.  Cutter size : Disabled
Press Ctrl-C to exit...
```


# OSNT in action



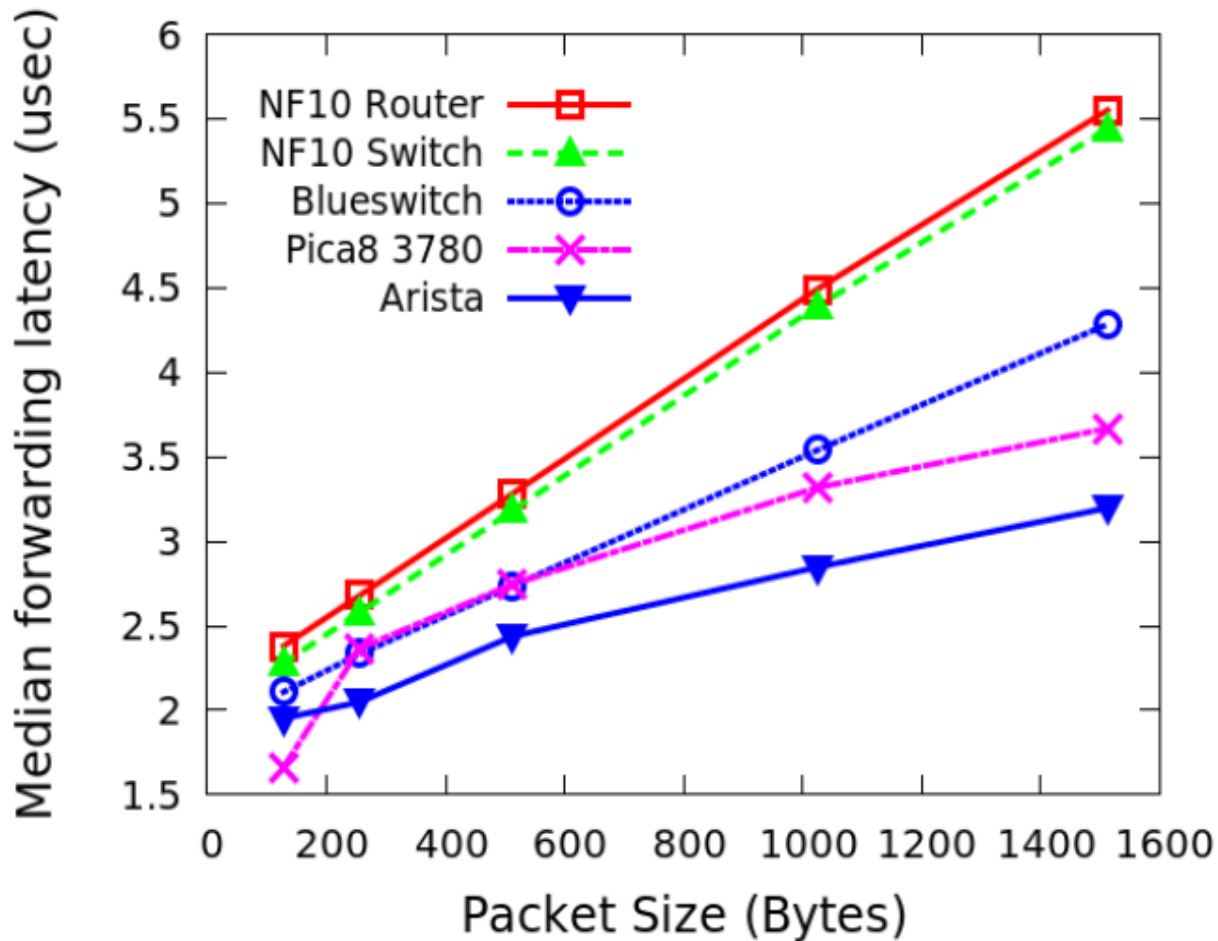
*Enabling network innovation with accurate networking systems  
characterization*

# Forwarding latency measurement

- Unloaded switches baseline latency no cross traffic



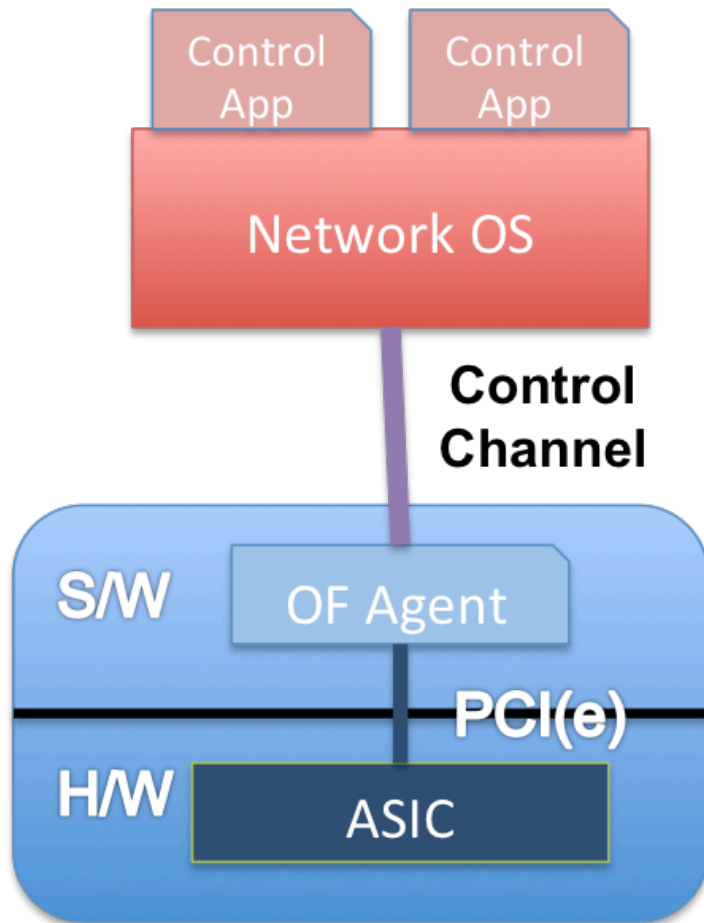
# Forwarding latency measurement



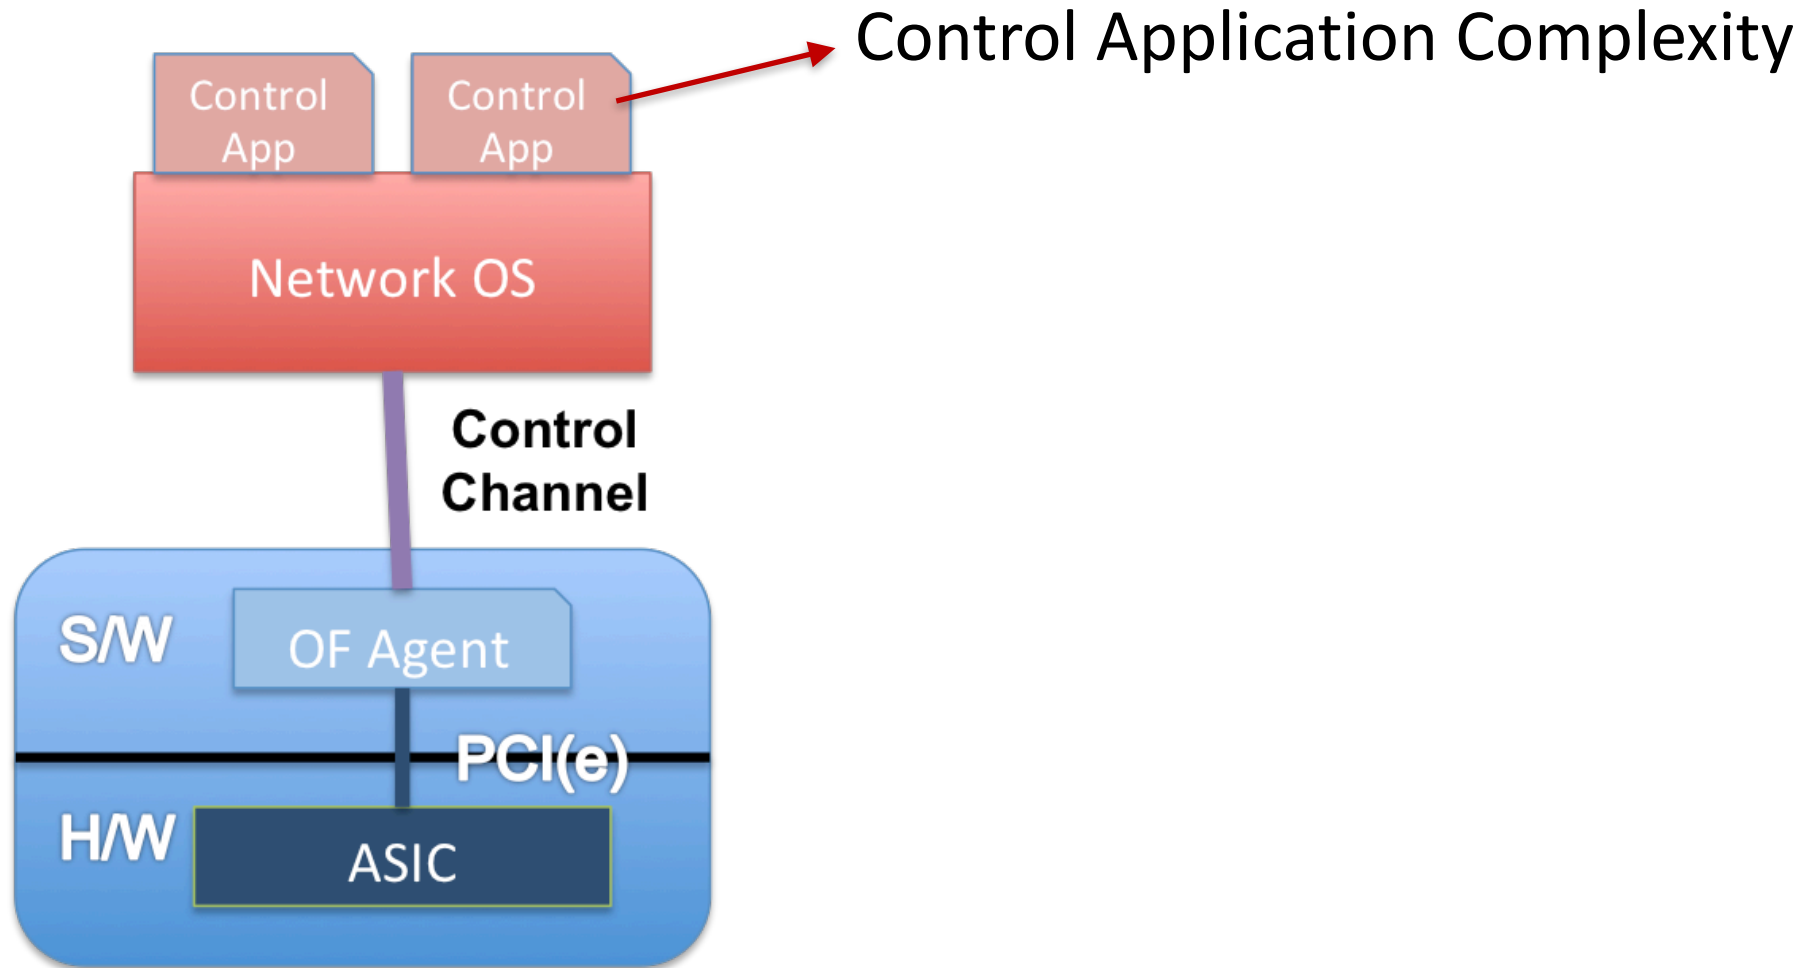
# SDN Testing Suite

- SDN enables unprecedented flexible and extensible network control
- OpenFlow specifications lack performance semantics
  - What does a barrier reply signifies?
- OpenFlow performance aspects are yet to be explored
  - How do you compare two OpenFlow switches?
- OpenFlow flexibility is not always portable on switch ASIC

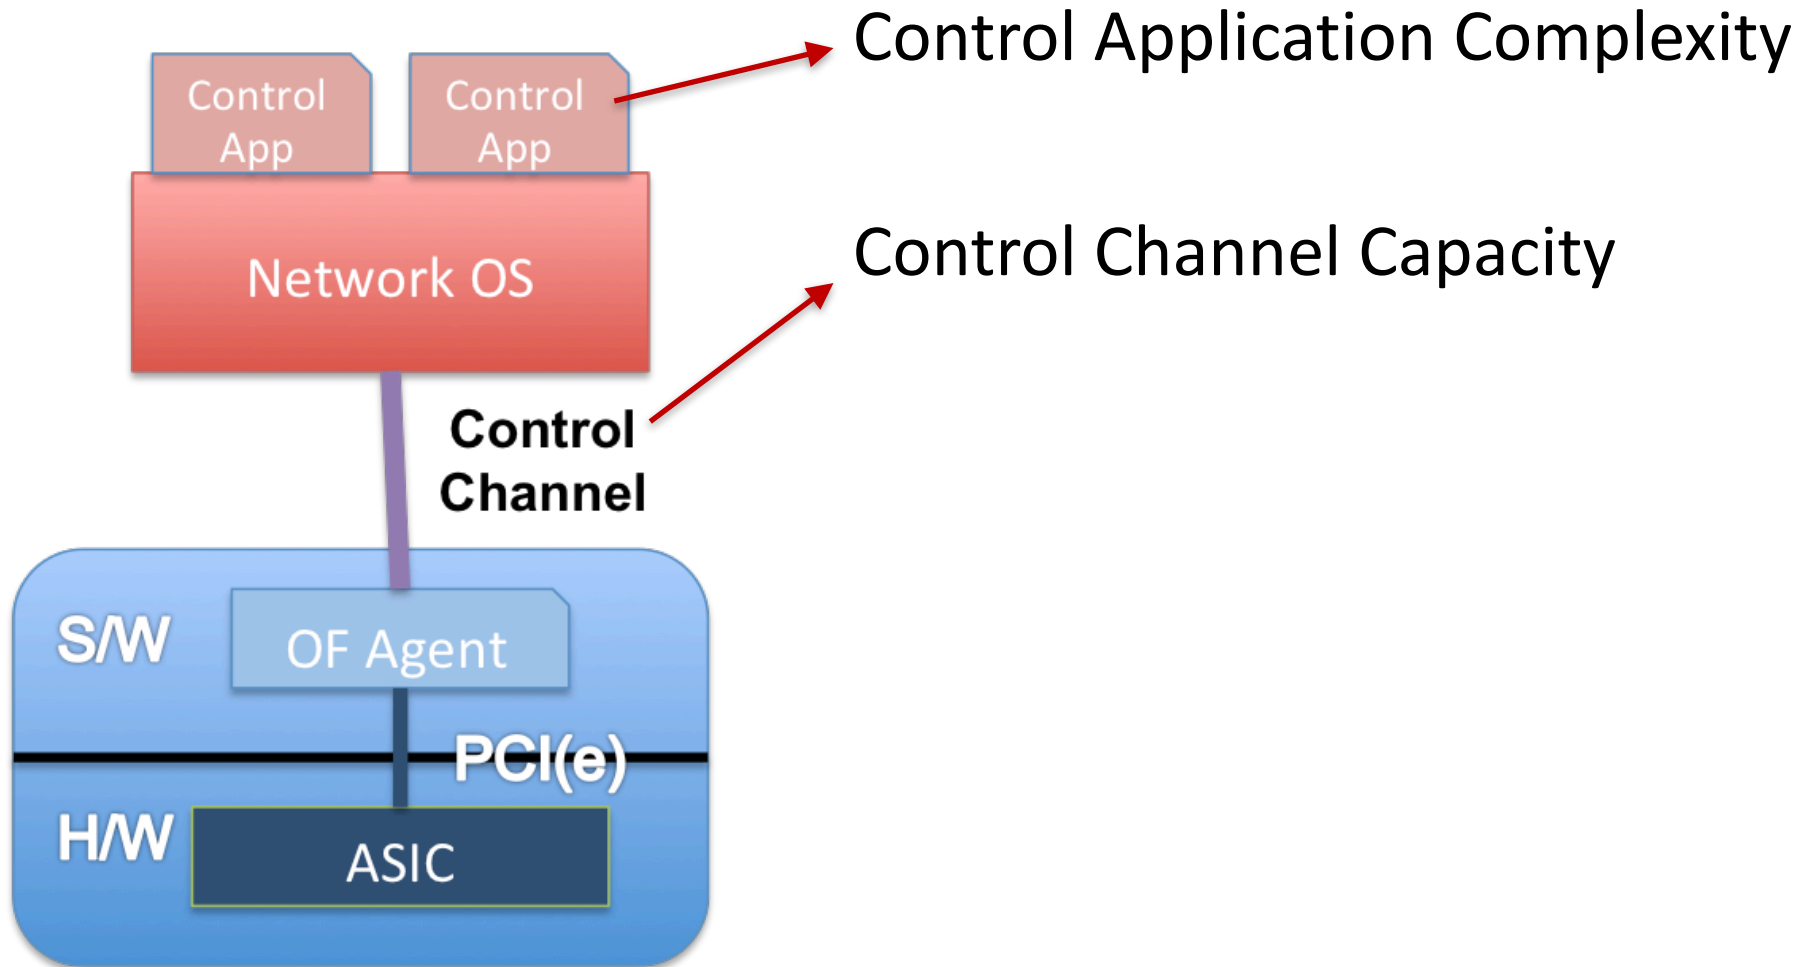
# OpenFlow toolstack X-Ray



# OpenFlow toolstack X-Ray

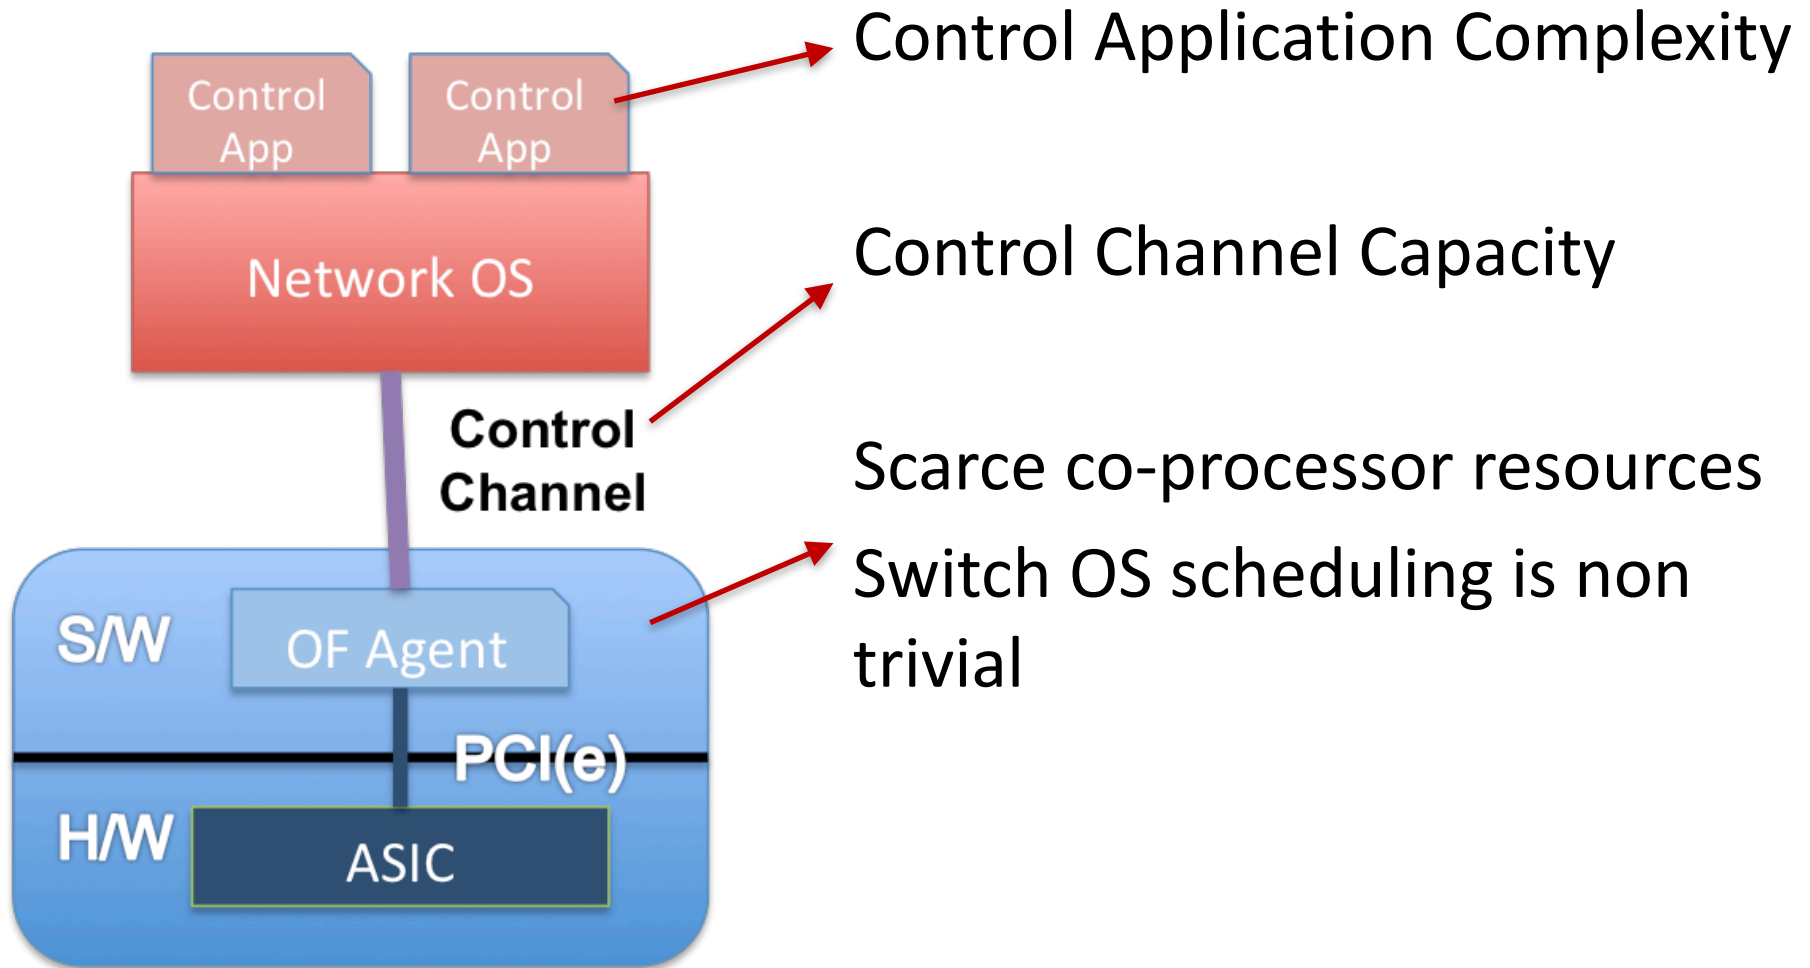


# OpenFlow toolstack X-Ray

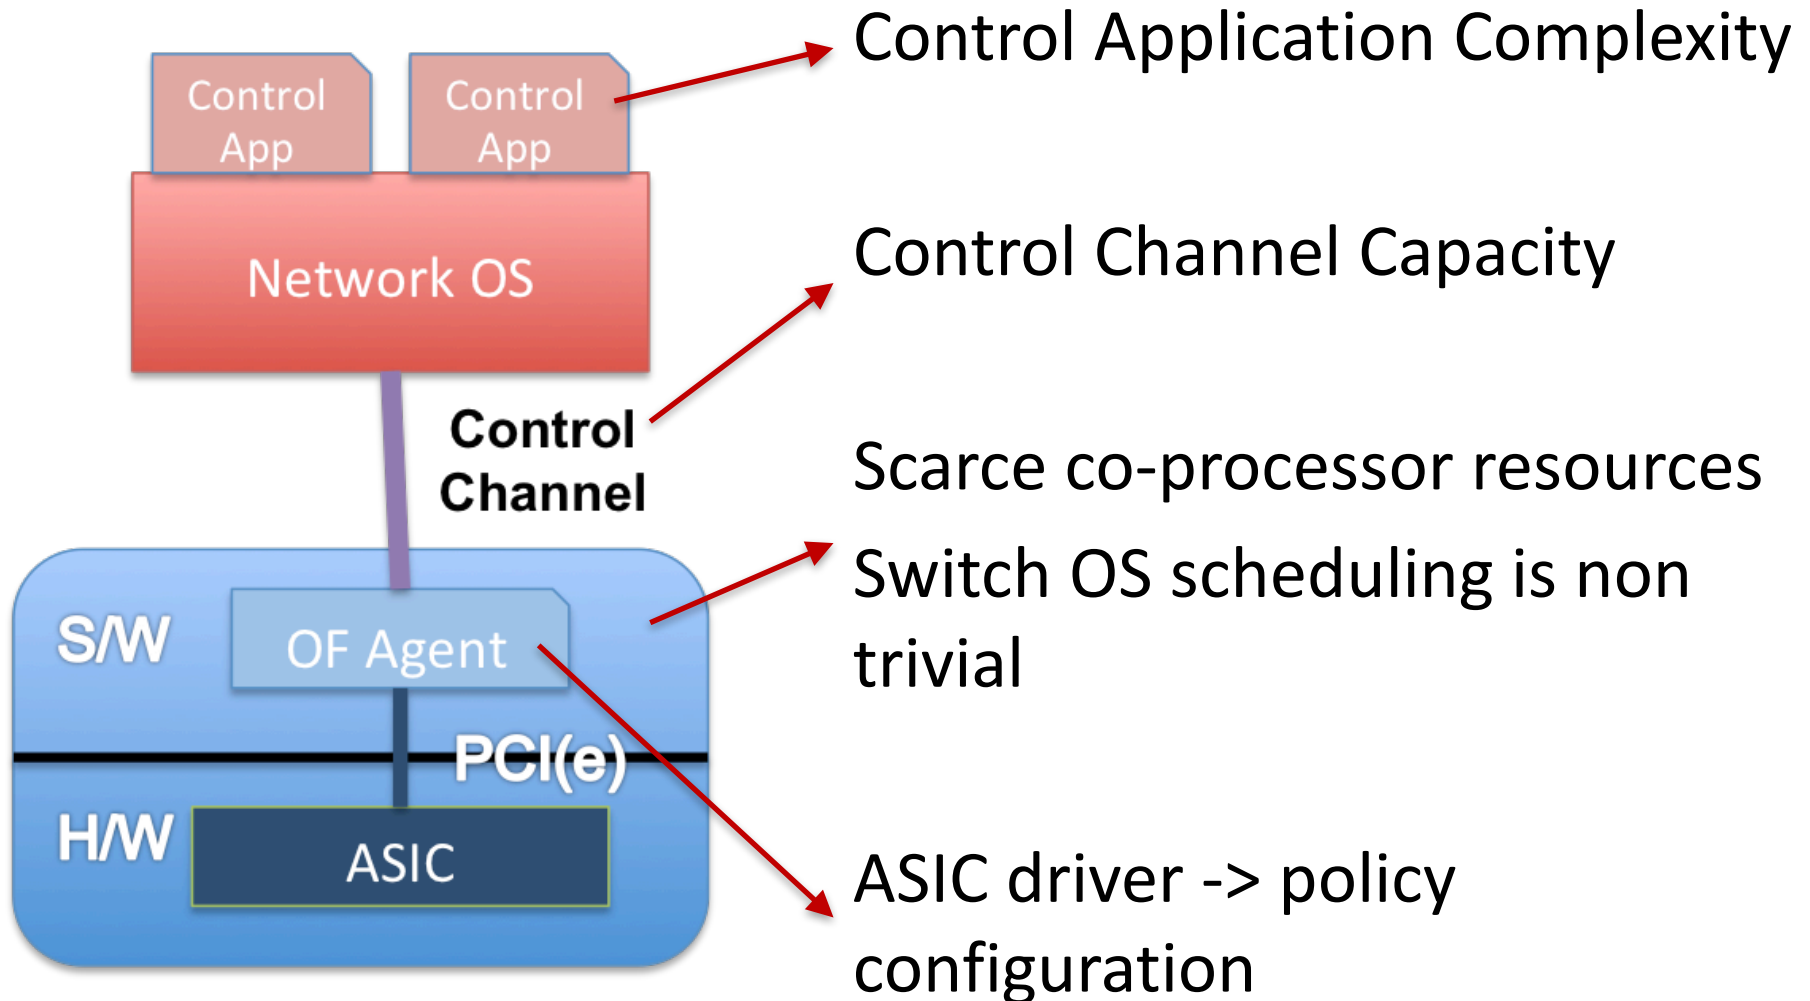




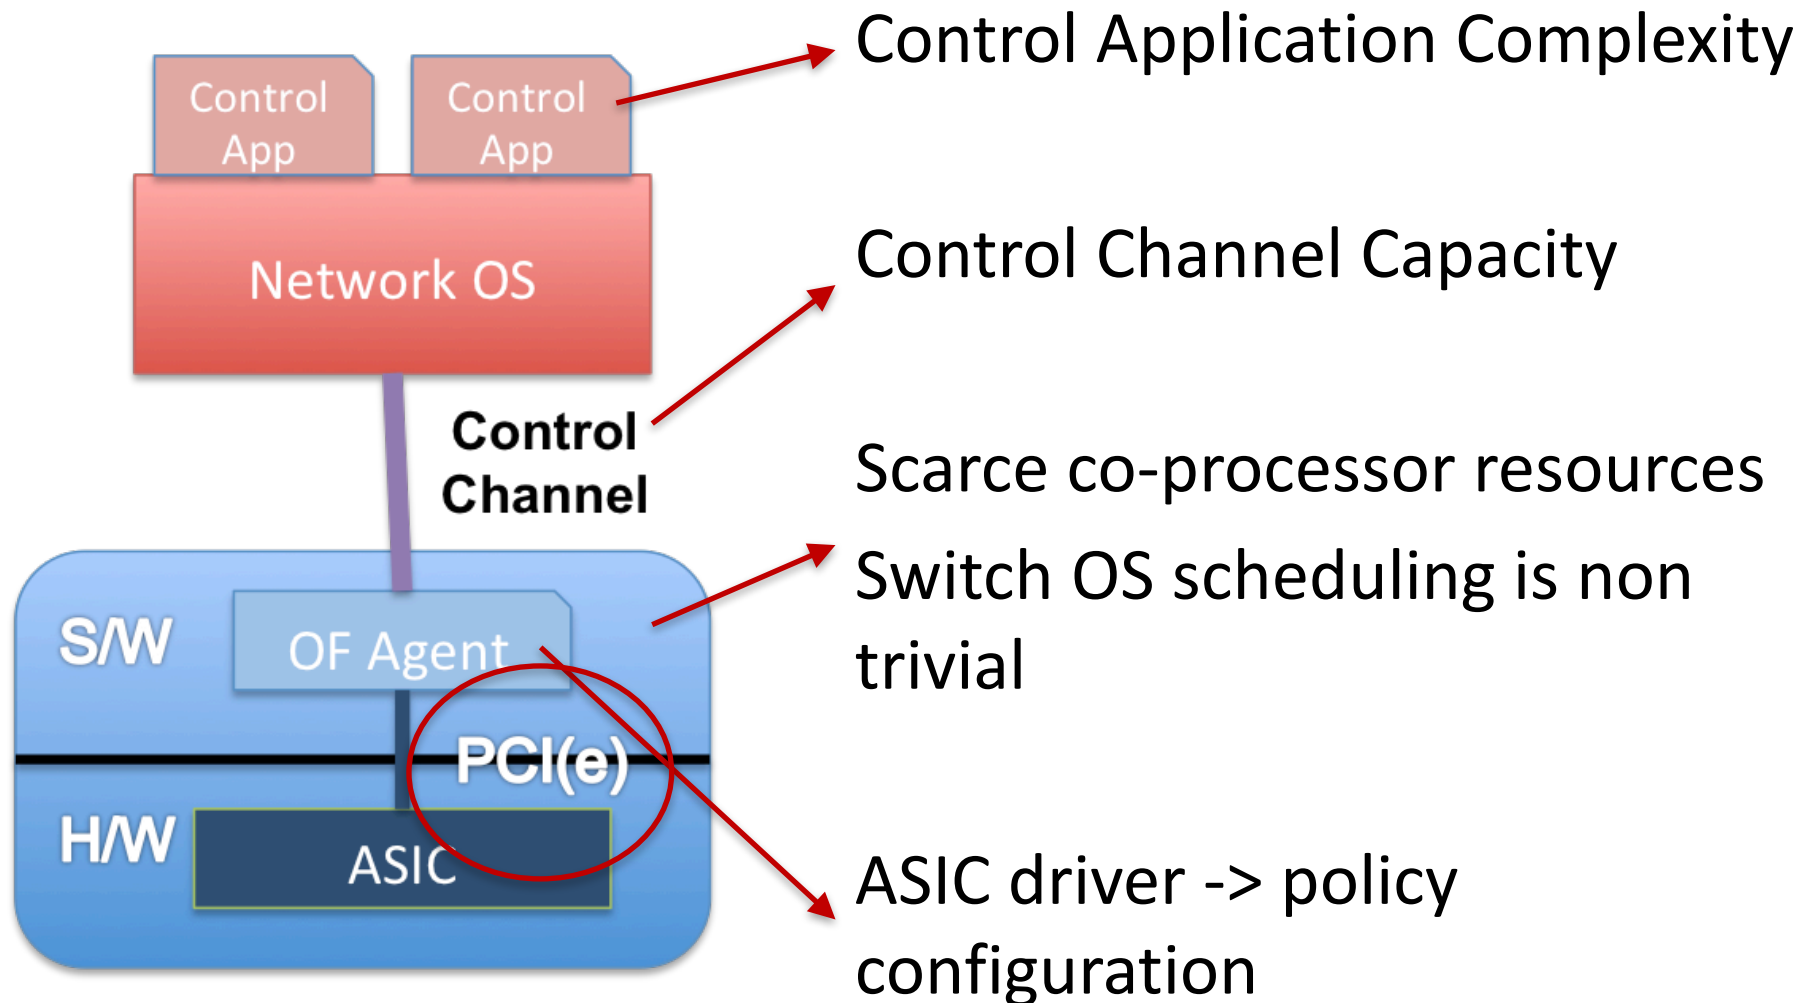
# OpenFlow toolstack X-Ray



# OpenFlow toolstack X-Ray

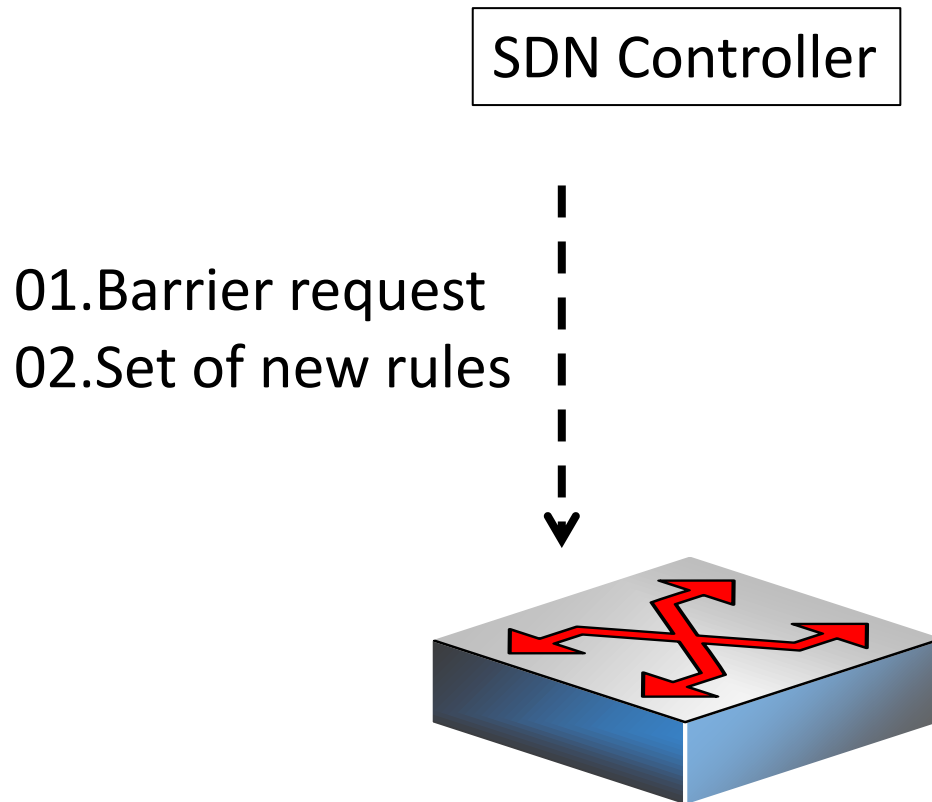


# OpenFlow toolstack X-Ray




# SDN networks performances

## THE PROACTIVE CASE



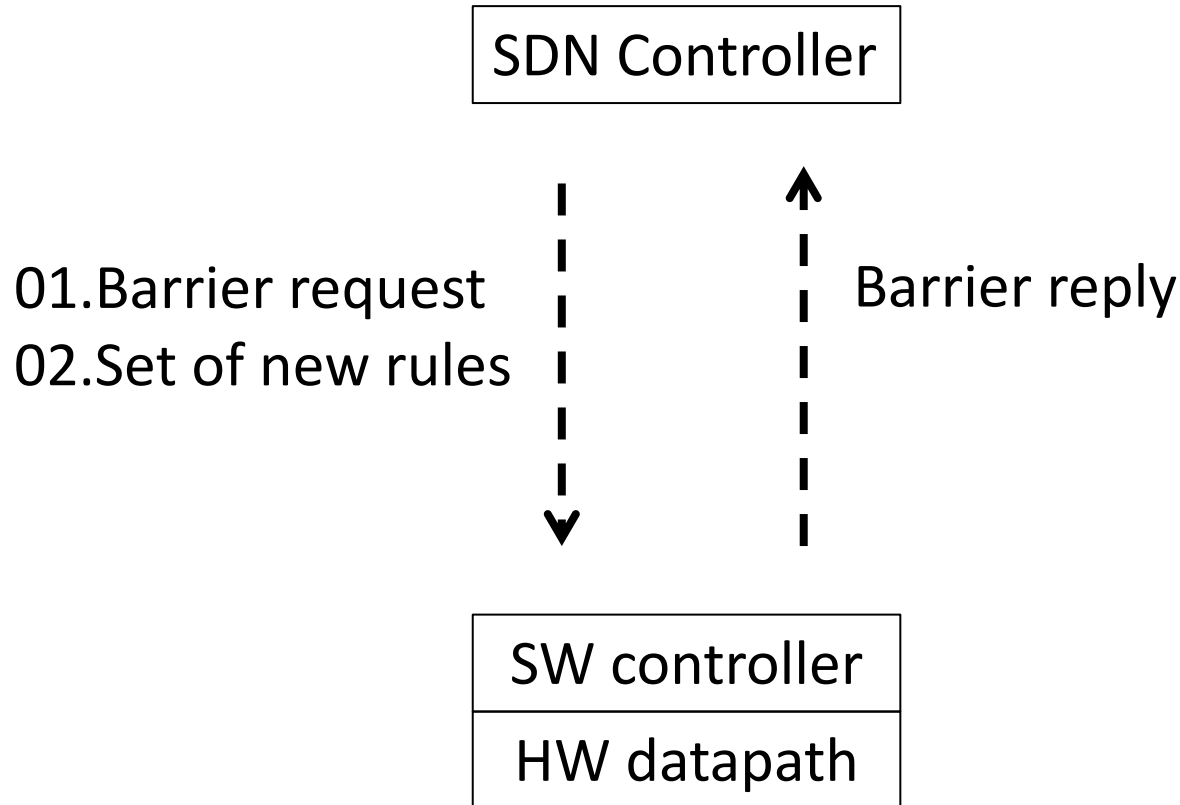
# SDN networks performances

## THE PROACTIVE CASE



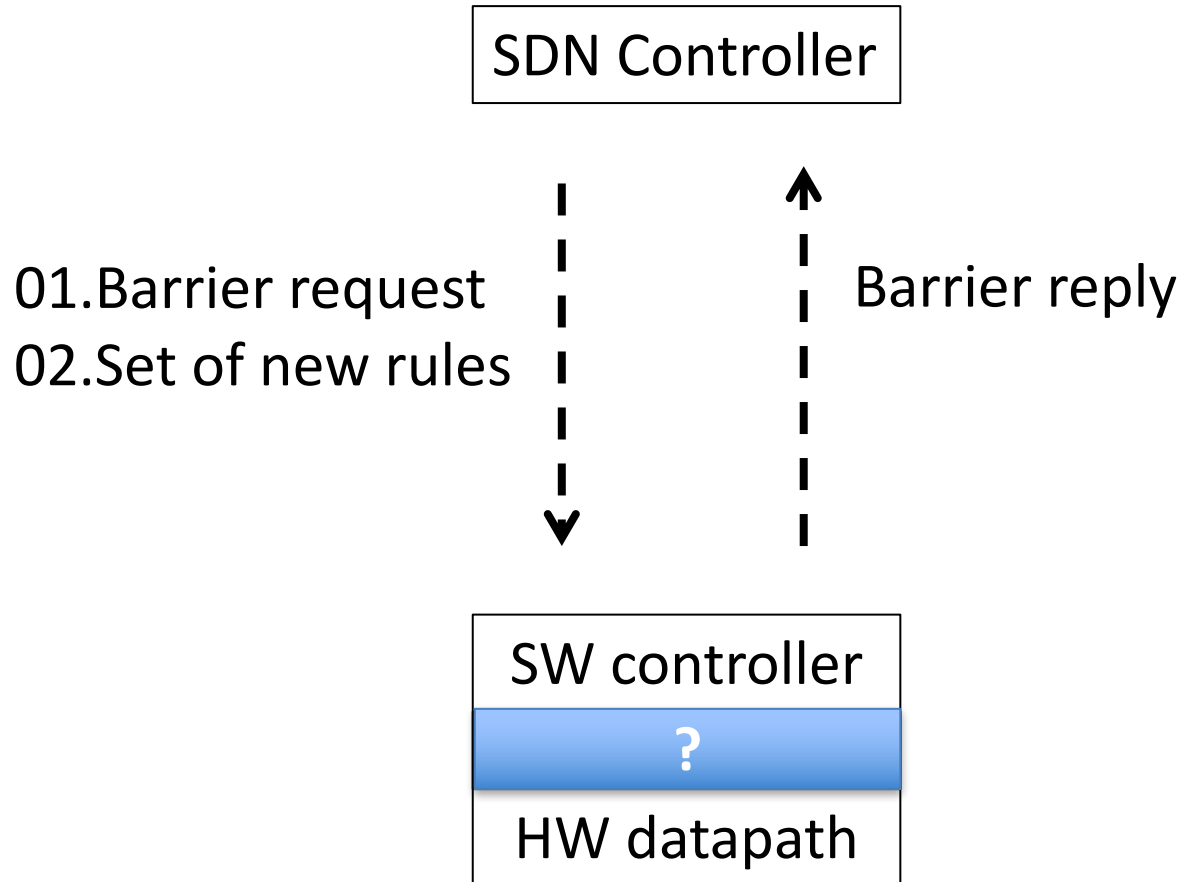
# SDN networks performances

## THE PROACTIVE CASE



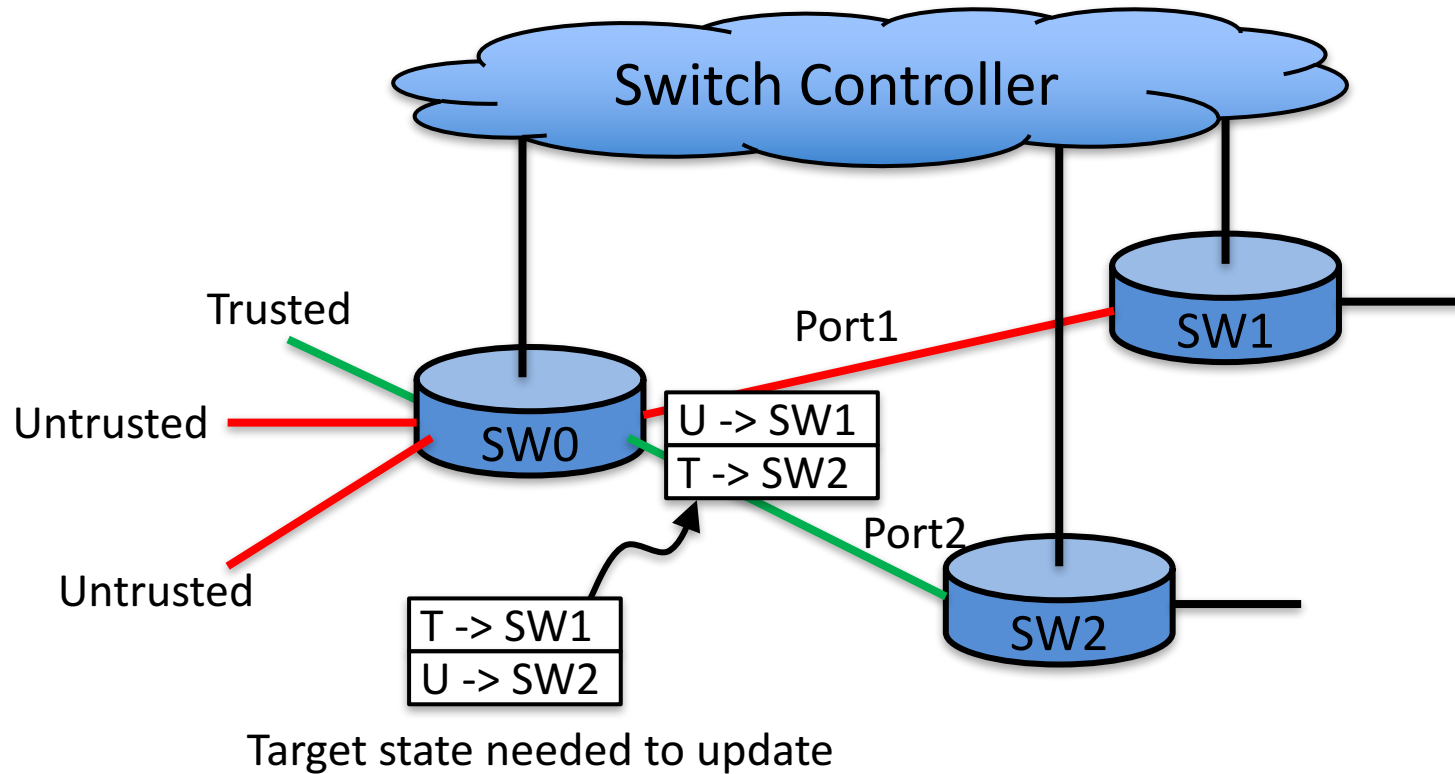
# SDN networks performances

## THE PROACTIVE CASE



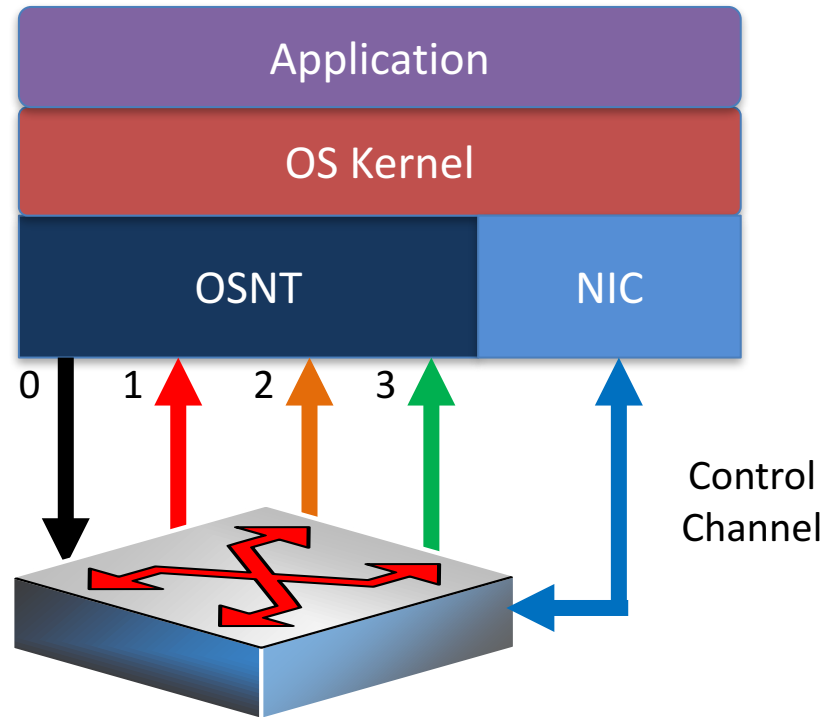
# Control/Data plane consistency

Consistent policy update affects security in SDN.

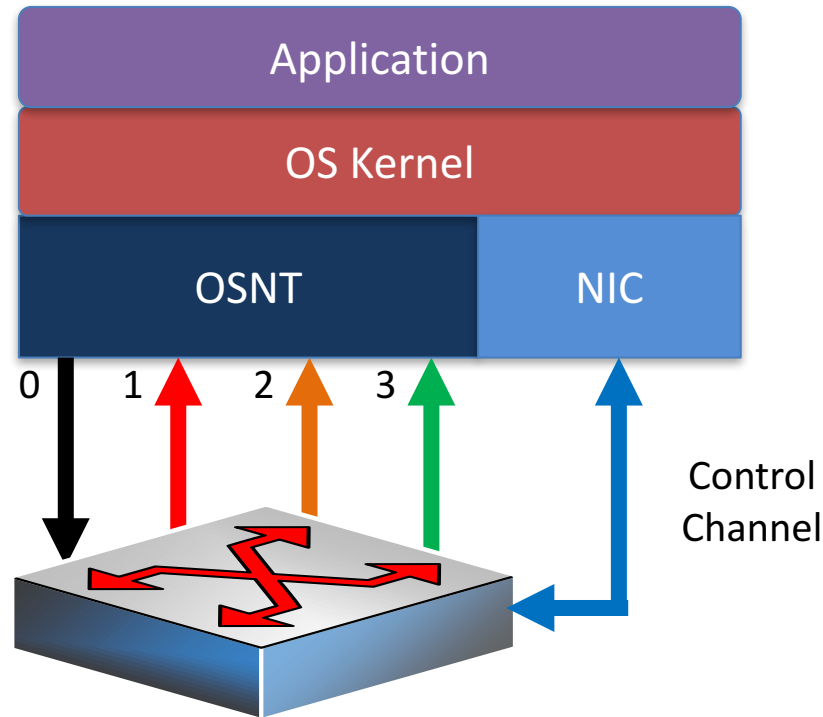




# Control/Data plane consistency




# Control/Data plane consistency



1. Initial rule :  $0 \rightarrow 1$  (as a set of different IP rules)
2. Rule update :  $0 \rightarrow 2$


# Control/Data plane consistency

We generate an aggregate 2Gbps with 150B packets

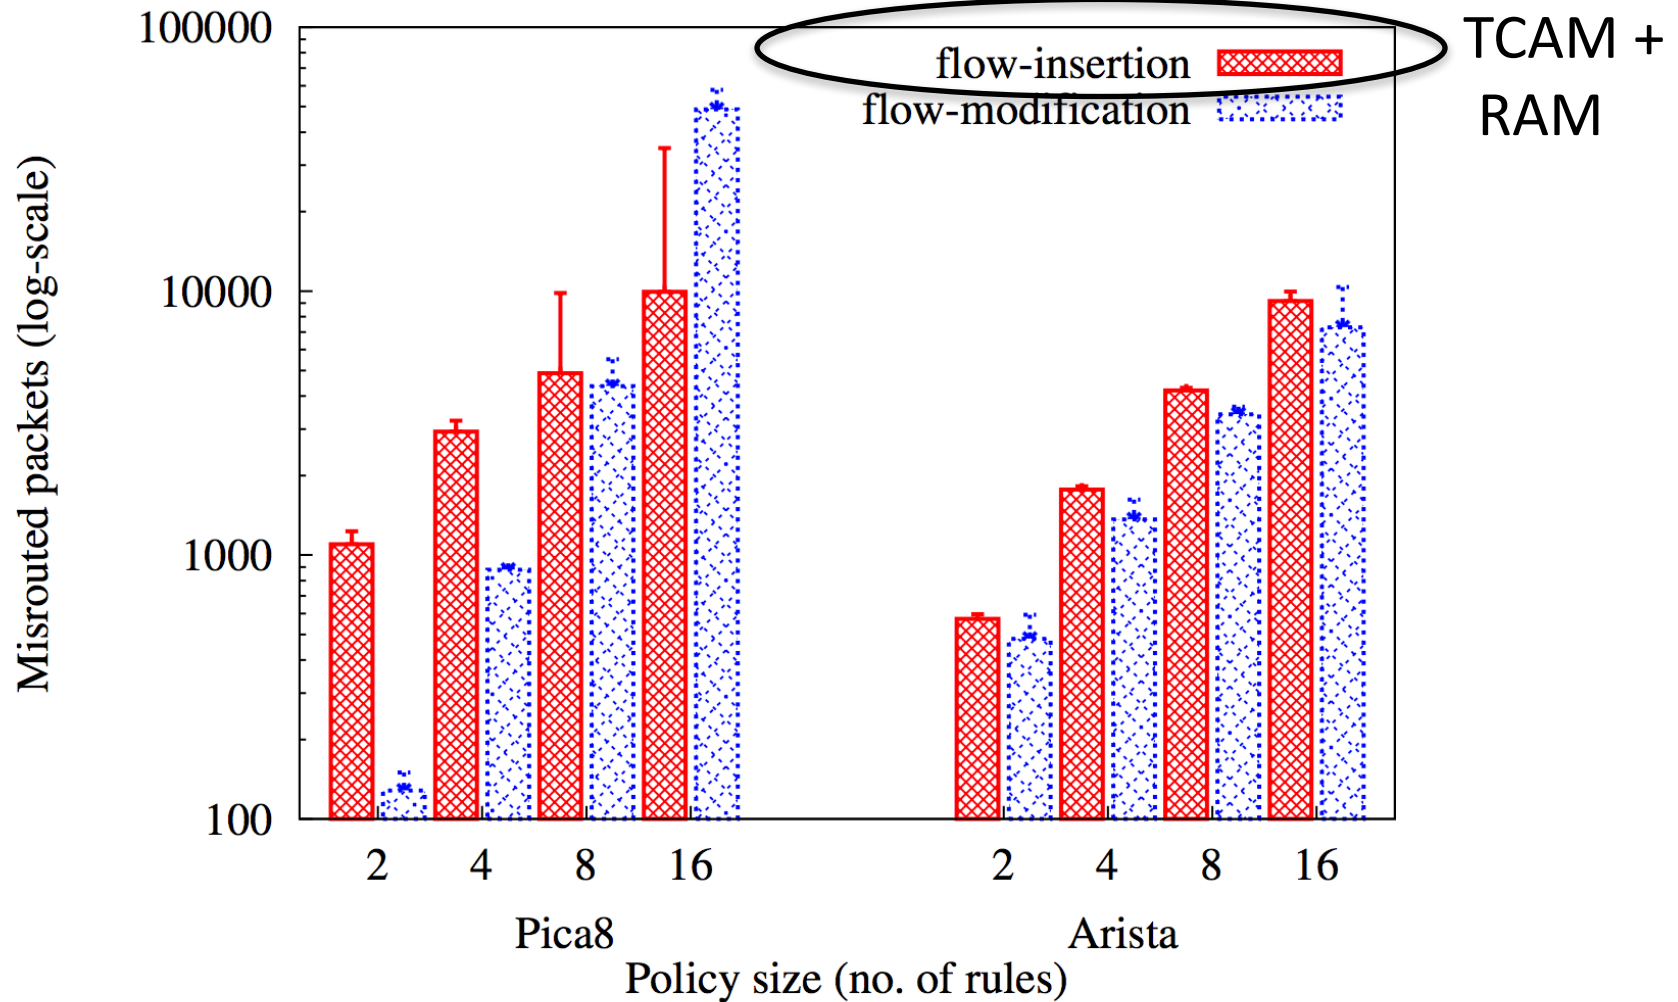


1. Initial rule :  $0 \rightarrow 1$  (as a set of different IP rules)
2. Rule update :  $0 \rightarrow 2$

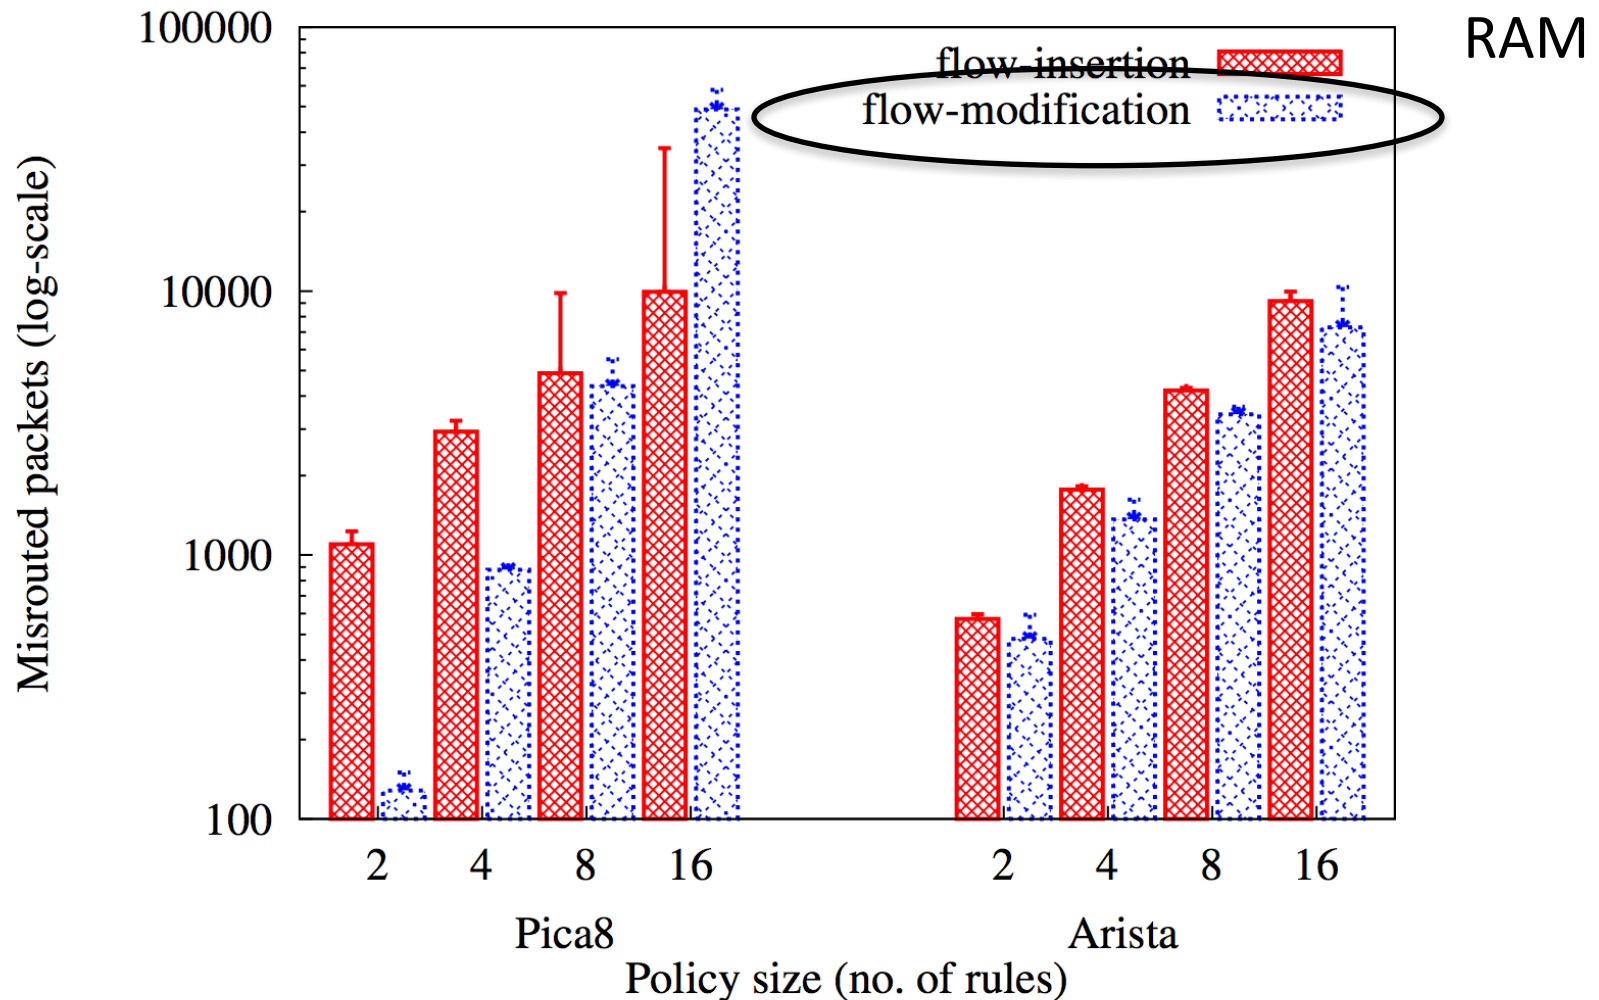
# Control/Data plane consistency



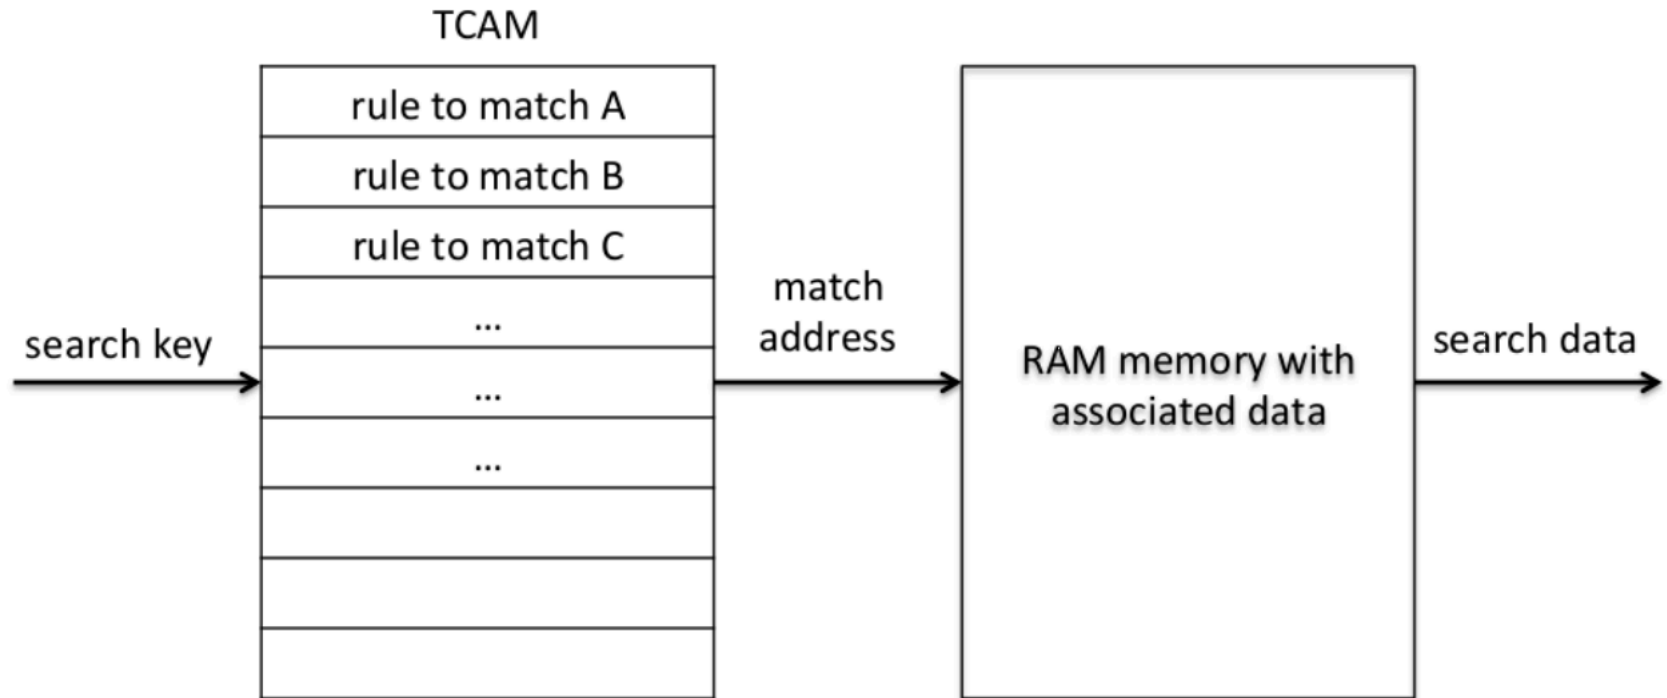
# Control/Data plane consistency



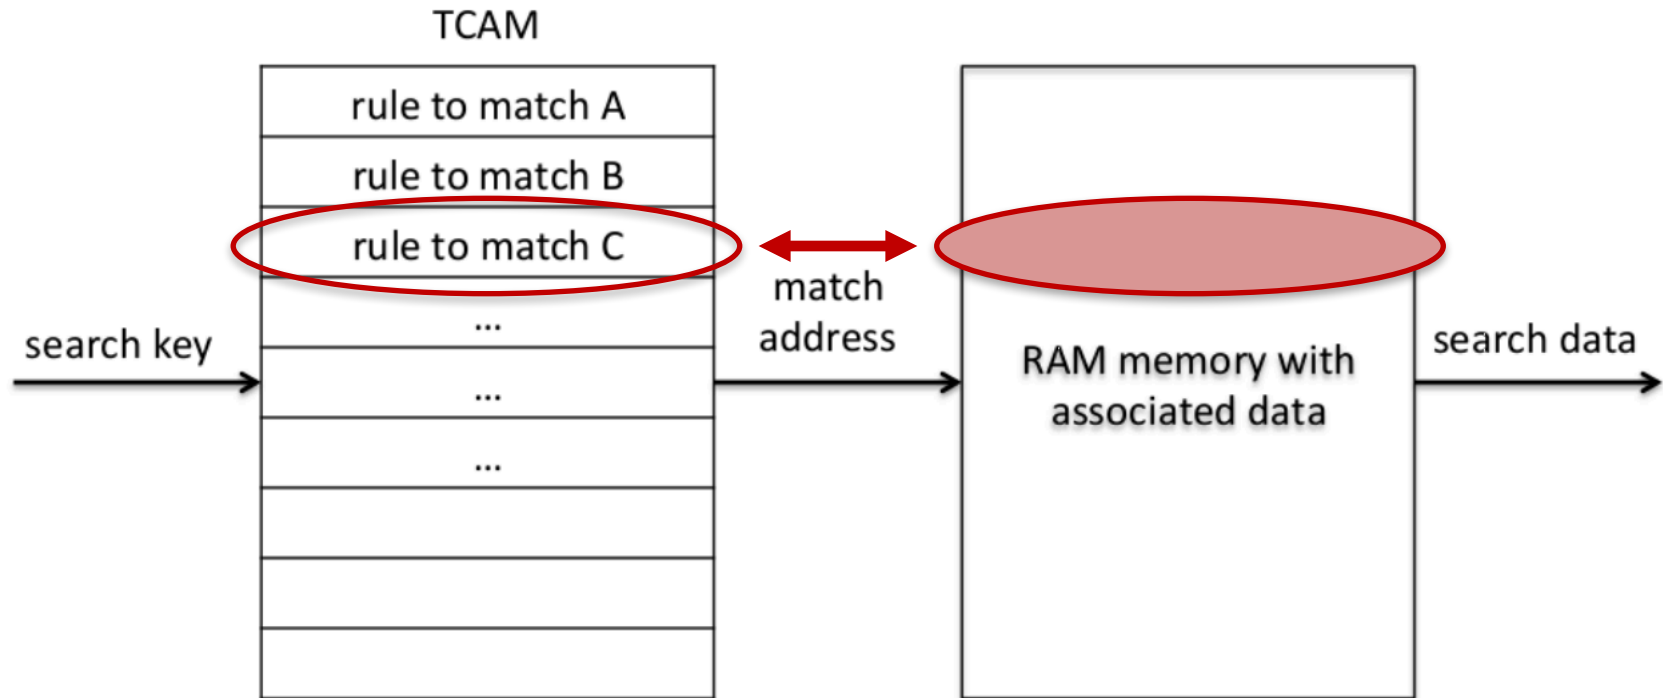
# Control/Data plane consistency



# Control/Data plane consistency



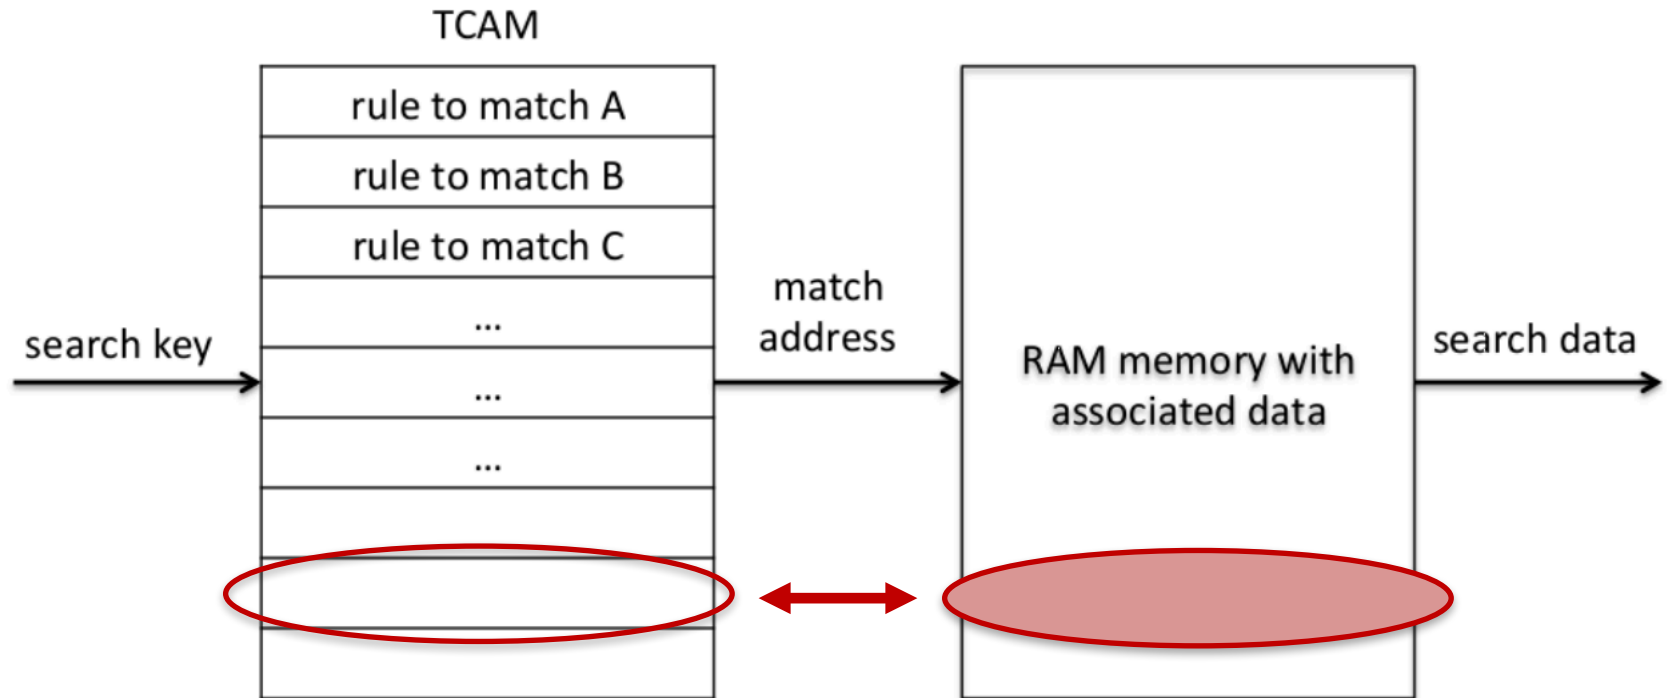
# Control/Data plane consistency



A flow modification requires only to change the “action” in the RAM. The flow is already present in the TCAM.

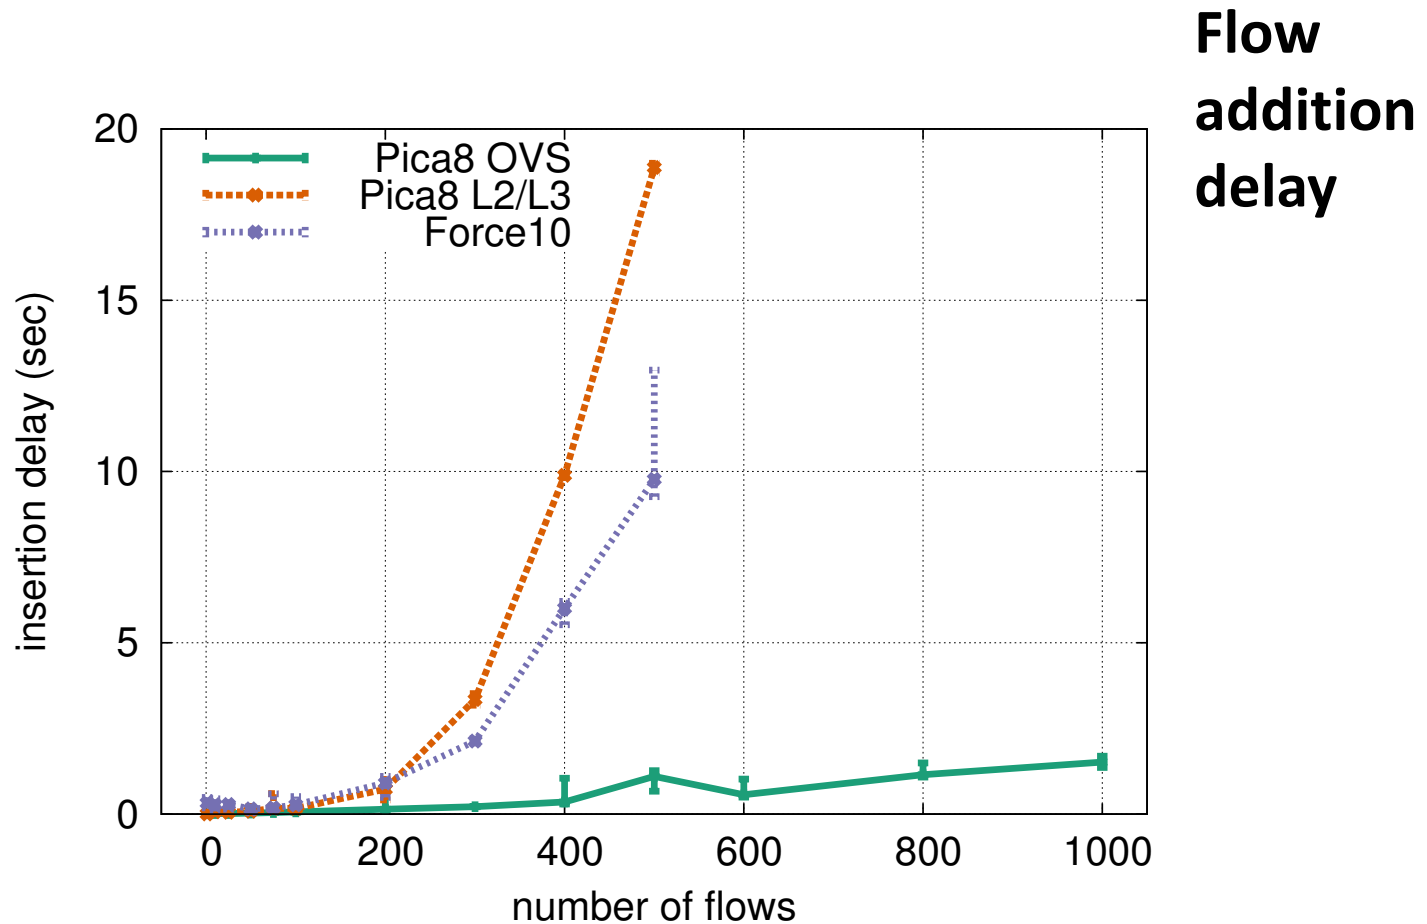


# Control/Data plane consistency




A flow insertion requires also to write the matching fields in the TCAM.

# Dataplane performances

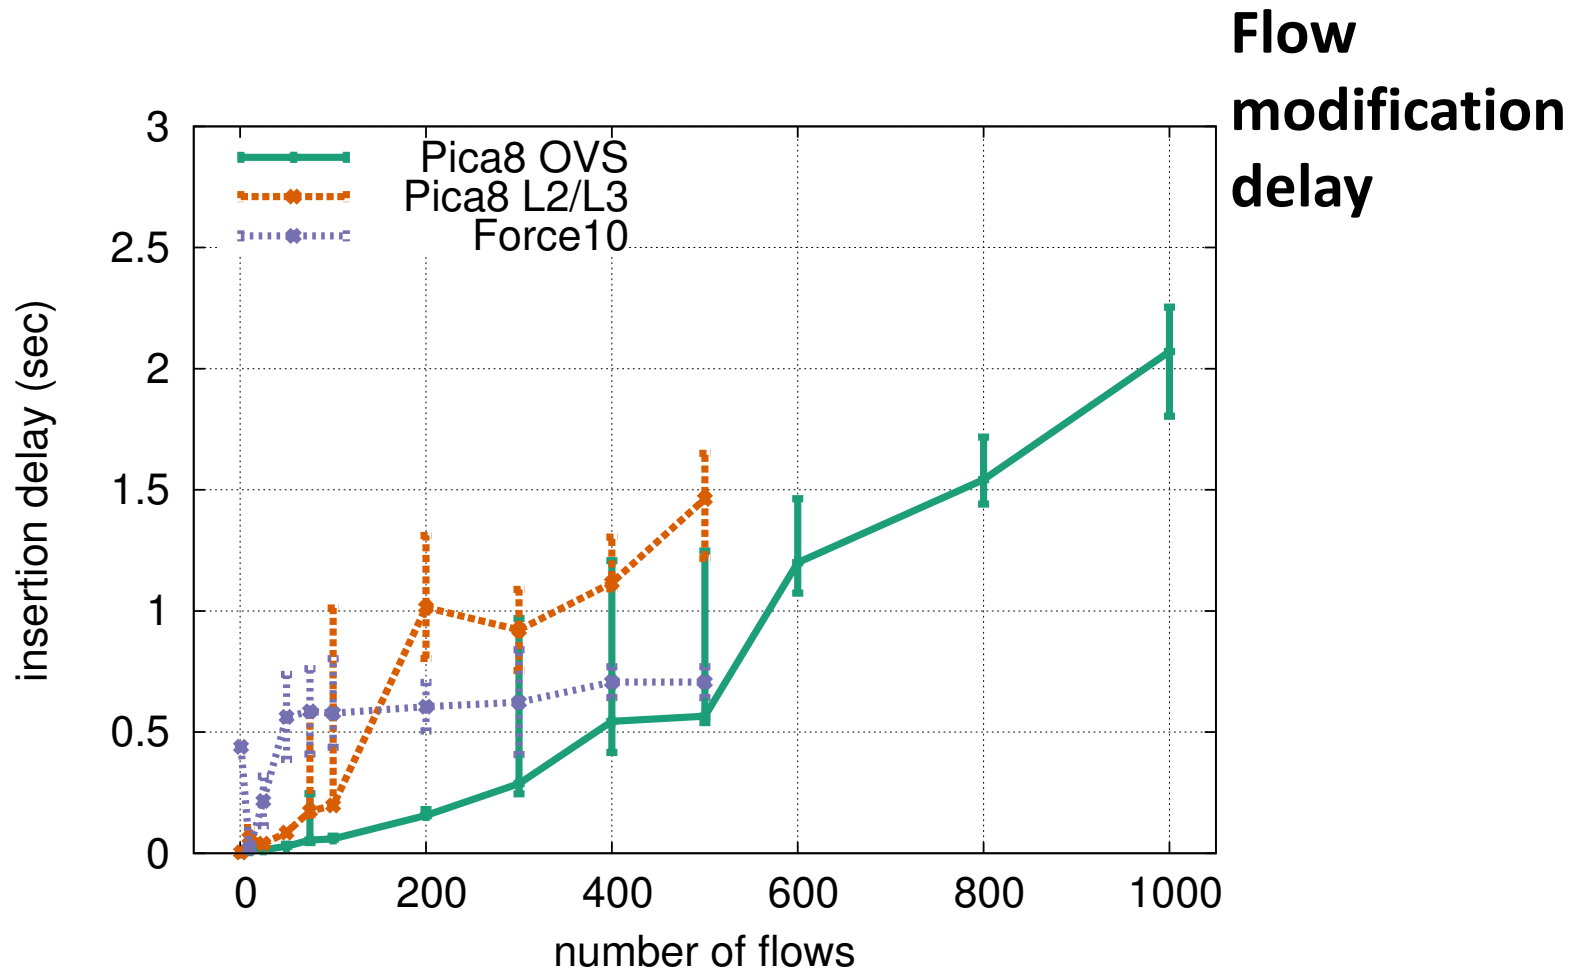


# Dataplane performances




**Flow  
addition  
delay**

# Dataplane performances

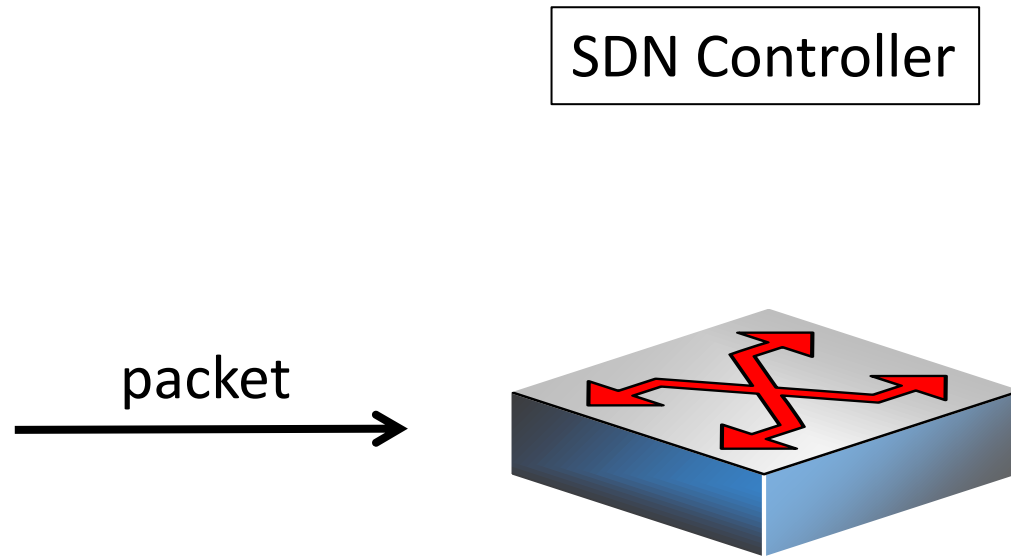


# Dataplane performances




# SDN networks performances

## THE REACTIVE CASE



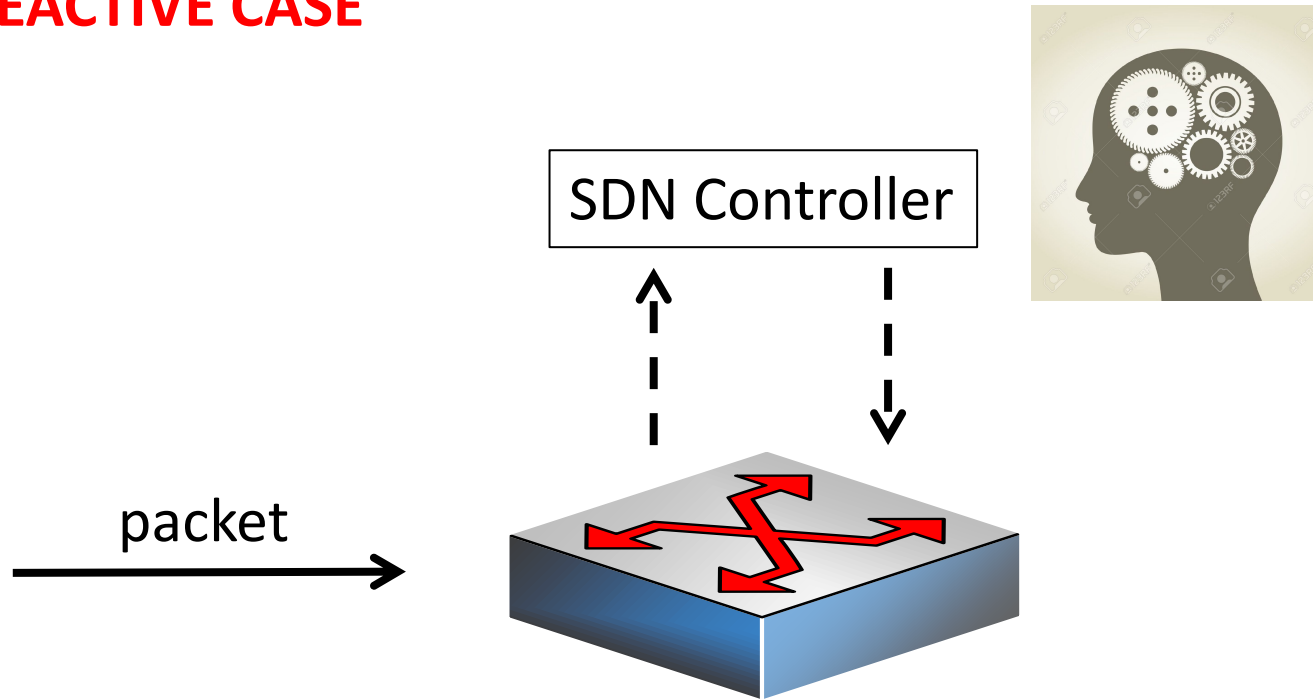
# SDN networks performances

## THE REACTIVE CASE



# SDN networks performances


## THE REACTIVE CASE



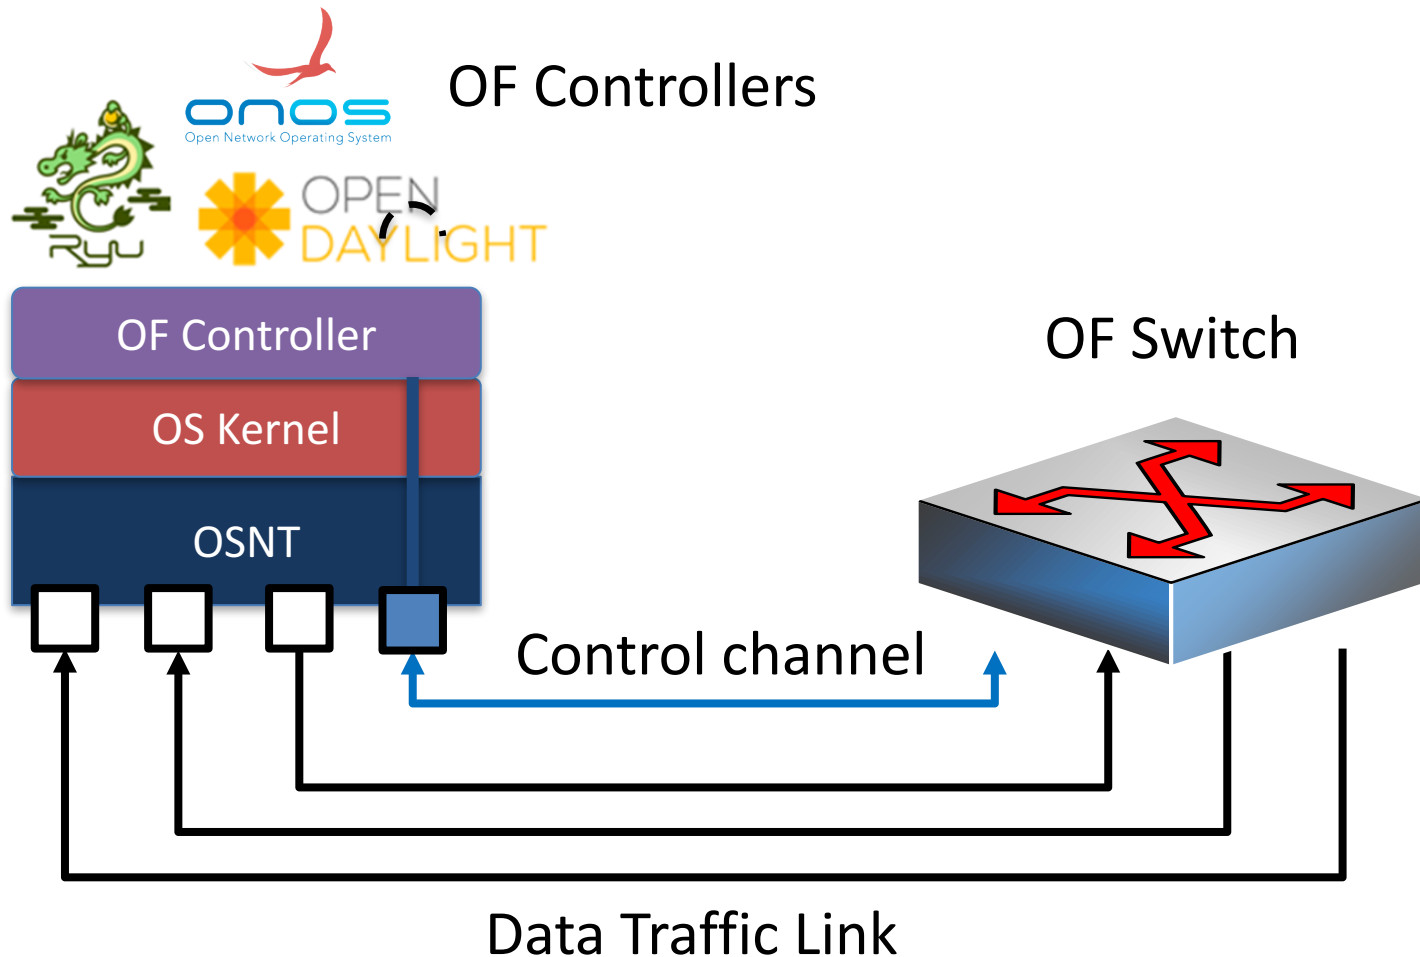


# SDN networks performances


## THE REACTIVE CASE



# Controller performances in SDN networks



# Controller performances in SDN networks



# OSNT-SUME-live Github

- OSNT-SUME-live is publicly available.

NetFPGA / OSNT-SUME-live Private

Unwatch 20 Star 0 Fork 2

Code Issues 0 Pull requests 0 Projects 0 Wiki Pulse Graphs Settings

OSNT for NetFPGA-SUME board Edit

New Add topics

47 commits 1 branch 4 releases 3 contributors


Branch: master New pull request Create new file Upload files Find file Clone or download

jhhan Merge branch 'master' of https://github.com/NetFPGA/OSNT-SUME-live Latest commit 1a058a2 7 days ago

contrib/challenge2017	NetFPGA design challenge 2017, added the first test	12 days ago
lib	Update the extmem packet replay controller and add the qdr i/f contro...	7 days ago
projects	Update minor in 10g rx and tx cores. Add comments to the timestamp mo...	a month ago
scripts	Add the first clean repo.	8 months ago
util	update pcap_gen to support configurable source and dest MAC addresses	12 days ago
Makefile	Add Makefile for ip core and sw driver generation.	
README.md	Update the release not link.	

README.md

OSNT

A photograph of a NetFPGA-10G-SUME board, a high-speed networking board featuring a Virtex-7 FPGA, a 10GbE SFP+ port, and various peripheral components like RAM and connectors.

# Acknowledgments (I)

## ***NetFPGA Team at University of Cambridge (Past and Present):***

Andrew Moore, David Miller, Muhammad Shahbaz, Martin Zadnik, Matthew Grosvenor, Yury Audzevich, Neelakandan Manihatty-Bojan, Georgina Kalogeridou, Jong Hun Han, Noa Zilberman, Gianni Antichi, Charalampos Rotsos, Hwanju Kim, Marco Forconesi, Jinyun Zhang, Bjoern Zeeb, Robert Watson, Salvator Galea, Marcin Wojcik, Diana Andreea Popescu, Murali Ramanujam

## ***NetFPGA Team at Stanford University (Past and Present):***

Nick McKeown, Glen Gibb, Jad Naous, David Erickson, G. Adam Covington, John W. Lockwood, Jianying Luo, Brandon Heller, Paul Hartke, Neda Beheshti, Sara Bolouki, James Zeng, Jonathan Ellithorpe, Sachidanandan Sambandan, Eric Lo, Stephen Gabriel Ibanez

## ***All Community members (including but not limited to):***

Paul Rodman, Kumar Sanghvi, Wojciech A. Koszek, Yahsar Ganjali, Martin Labrecque, Jeff Shafer, Eric Keller, Tatsuya Yabe, Bilal Anwer, Yashar Ganjali, Martin Labrecque, Lisa Donatini, Sergio Lopez-Buedo, Andreas Fiessler, Robert Soule, Pietro Bressana, Yuta Tokusashi  
Steve Wang, Erik Cengar, Michael Alexander, Sam Bobrowicz, Garrett Aufdemberg,  
Patrick Kane, Tom Weldon  
Patrick Lysaght, Kees Vissers, Michaela Blott, Shep Siegel, Cathal McCabe

# Acknowledgements (II)



UNIVERSITY OF  
CAMBRIDGE

**EPSRC**

Pioneering research  
and skills



The Leverhulme Trust



**ALGO-LOGIC**

