

WireCAP: a Novel Packet Capture Engine for Commodity NICs in High-speed Networks

Wenji Wu, Phil DeMar

Fermilab Network Research Group

wenji@fnal.gov, demar@fnal.gov

SC'15 2015

November 15-20, 2015 Austin, TX, USA



Managed by Fermi Research Alliance, LLC for the U.S. Department of Energy Office of Science

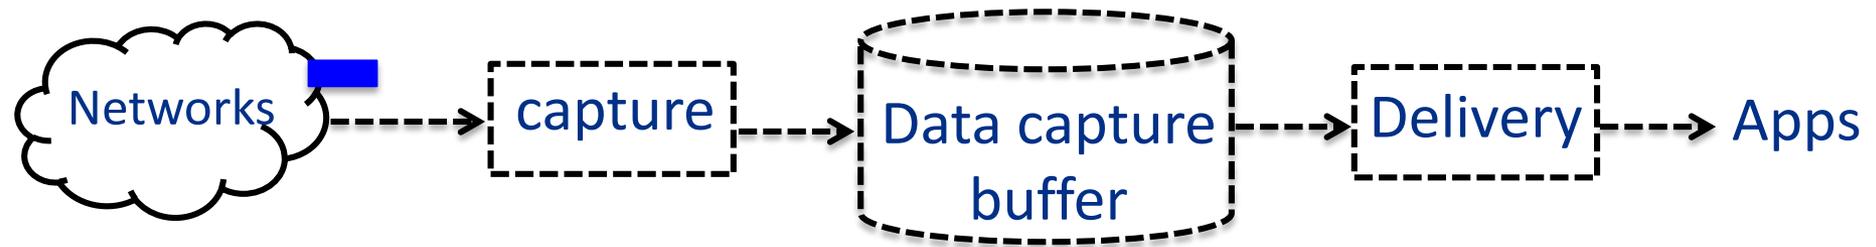
WireCAP: a Novel Packet Capture Engine for Commodity NICs in High-speed Networks

Wenji Wu, Phil DeMar
SC'15 Austin, TX, USA
November 15-20 2015

Packet Capture is an essential function for many network applications

- **Security analysis (e.g., IDS/IPS)**
 - Snort, www.snort.org
 - Suricata, <http://suricata-ids.org>
- **Network and application performance analysis**
 - Riverbed SteelCentral NetProfiler
 - Netscout Sniffer Analysis
 - Wildpackets OmniPeek Network Analyzer
- **Traffic characterization studies**
 - Benson, IMC'10
 - And more

A general packet capture process

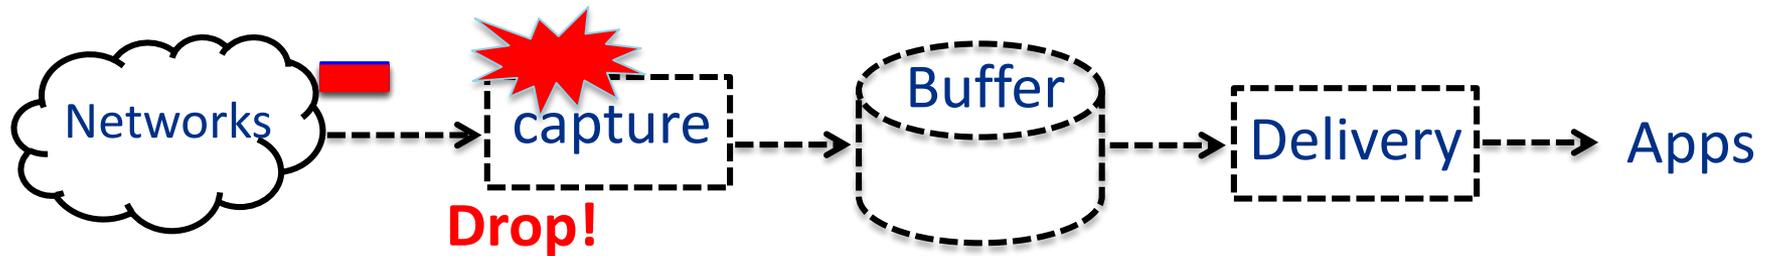


Typically computationally and I/O throughput intensive

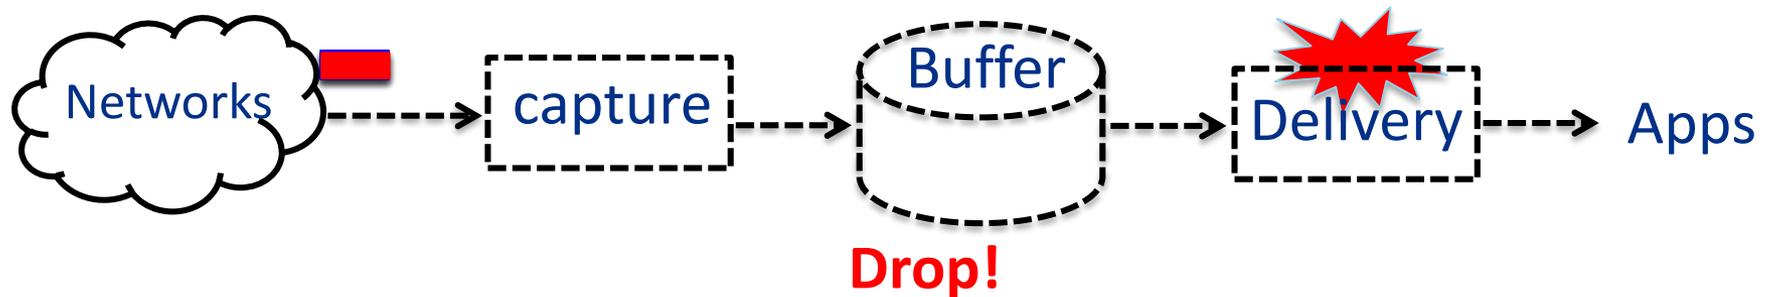
Significant performance challenges in high-speed networks

Packet drop is a major problem with packet capture in high-speed networks

Type I packet drop: packet capture drops



Type II packet drop: packet delivery drops



Packet drops degrade the accuracy & integrity of network monitoring applications!

Fundamental design goal in packet capture tools: Avoid packet drops!

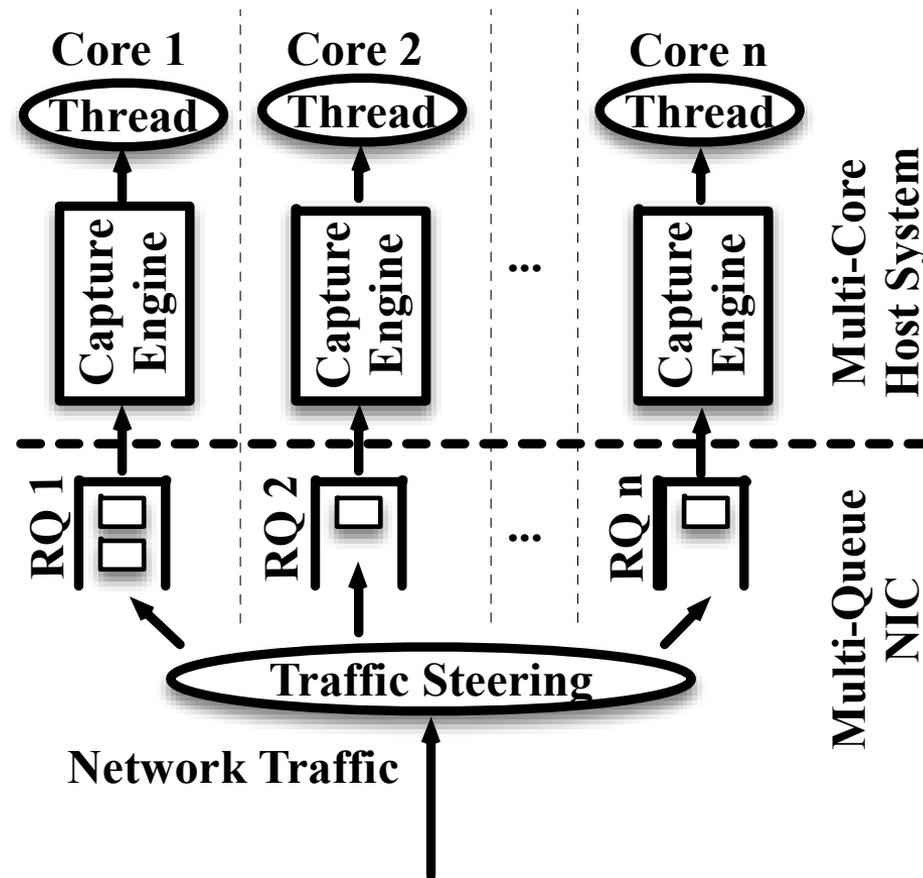
Packet capture approaches

1. **Use dedicated packet capture card**
 - **Pros**
 - The least amount of CPU intervention
 - Lossless packet capture and delivery
 - **Cons: costly, relatively inflexible, and not very scalable**
2. **Use a commodity system with a commodity NIC**
 - A commodity NIC in promiscuous mode to intercept pkts
 - A capture engine provides capture and delivery services
 - **Pros: flexible and cost effective**
 - **Cons: significant system CPU and memory resources**

The 2nd approach becomes more appealing with recent advances in multicore systems and multi-queue NICs:

- **A new paradigm in packet capturing and processing.**
- **This is our research focus.**

A new packet capturing and processing paradigm



Assumption: the hardware-based traffic-steering mechanism is capable of evenly distributing the incoming traffic among cores.

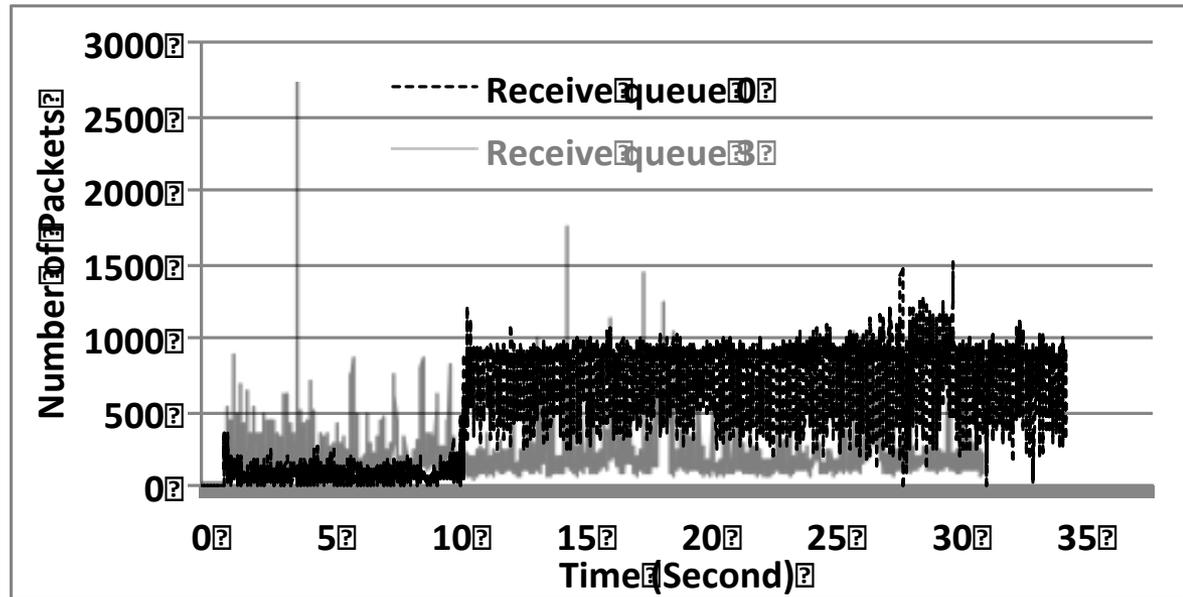
Question 1

Can NIC hardware-based traffic-steering mechanisms evenly distribute the incoming traffic among cores?

Question 2

Can existing packet capture engines support this new packet capture and processing paradigm?

Observation 1: load imbalance occurs frequently in a multicore system



Two types of load imbalance

- Short-term load imbalance (short-term burst of packets)
- Long-term load imbalance (queue 0 receives much more traffic than queue 3)

Existing NICs distribute traffic on a per-flow basis!

Observation 2: existing packet capture engines suffer significant packet drops

	<u>NETMAP</u>	<u>DNA</u>	<u>PF_RING</u>
<u>Receive Queue 0</u>			
Packet Capture Drops	46.5%	50.1%	0%
Packet Delivery Drops	0%	0%	56.8%
<u>Receive Queue 3</u>			
Packet Capture Drops	33.4%	9.3%	0.8%
Packet Delivery Drops	0%	0%	0%

Packet drop rates with a heavy packet-processing load

Existing packet capturing engines can suffer significant packet drops with load imbalance of either type on a multicore system!

Our Solution

WireCAP: a Novel Packet Capture Engine for Commodity NICs in High-speed Networks

What is WireCAP?

- **A packet capture engine for commodity network interface cards (NICs) in high-speed network**
- **Designed and developed by Fermilab network research group**

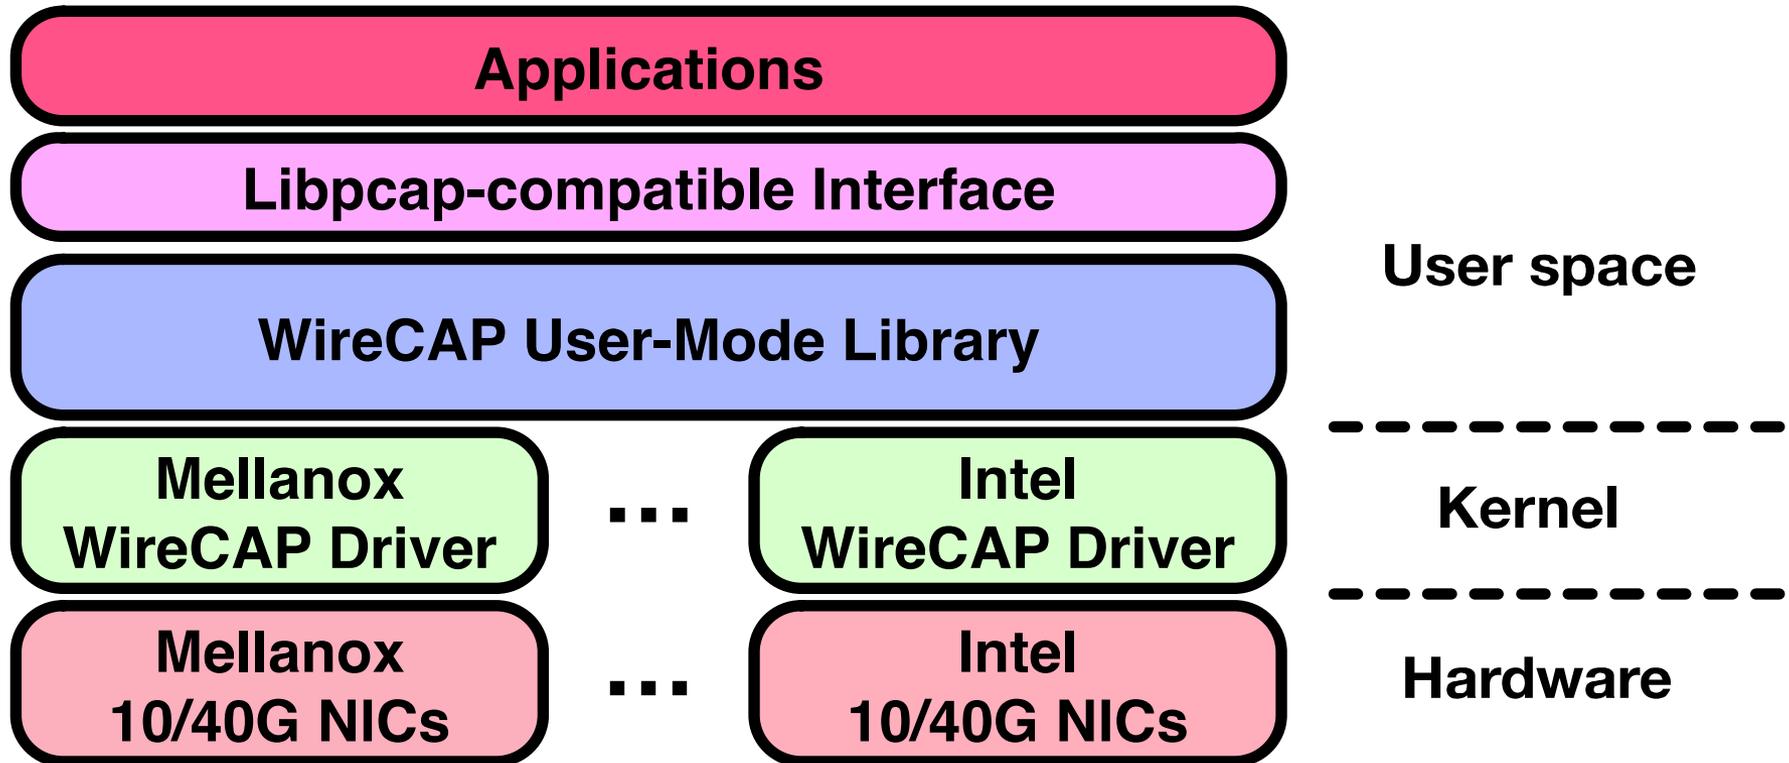
WireCAP Key Features

- **Lossless zero-copy packet capture and delivery**
- **Zero-copy packet forwarding**
 - To support middlebox-type applications
- **A Libpcap-compatible interface for low-level network access**
 - Wide applicability

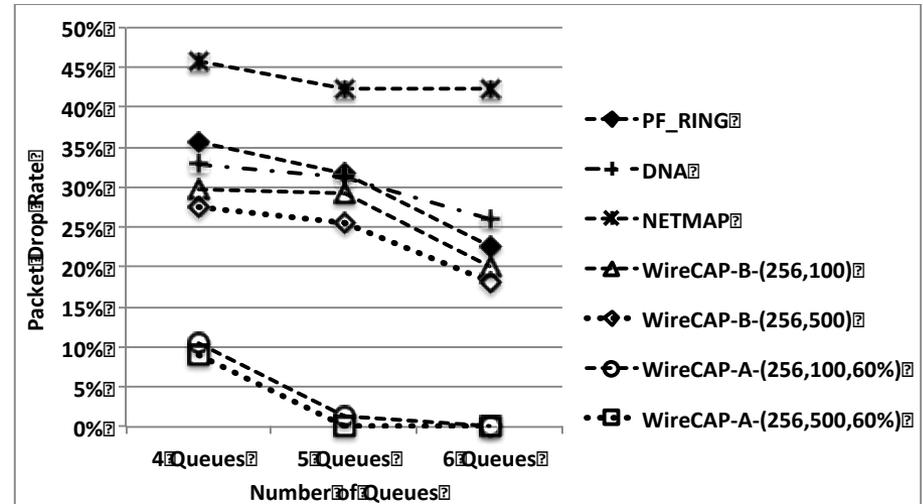
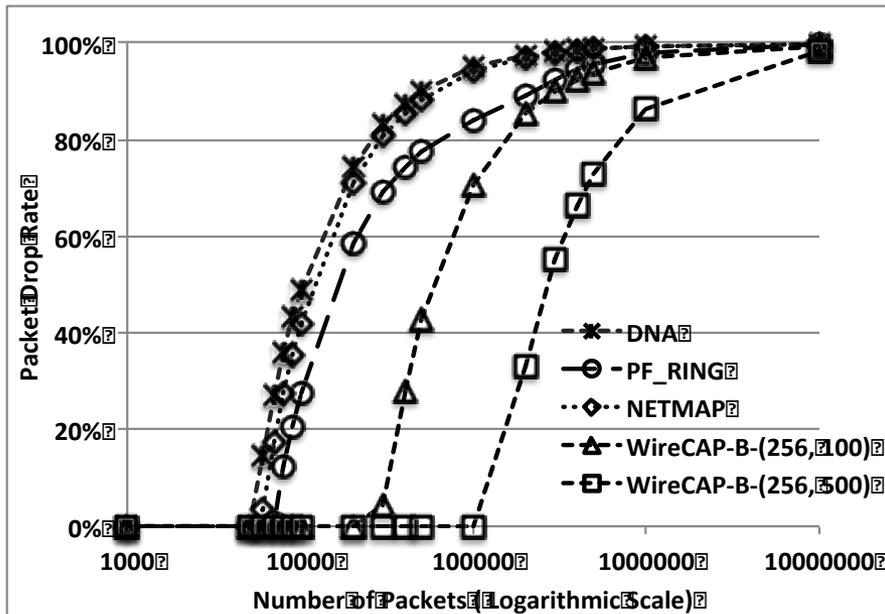
WireCAP Key Techniques

- **A ring-buffer-pool mechanism**
 - To handle short-term load imbalance
- **A buddy-group-based offloading mechanism**
 - To handle long-term load imbalance
- **Optimization techniques**
 - Pre-allocated large packet buffers
 - Zero-copy
 - Packet-level batching processing

WireCAP Architecture



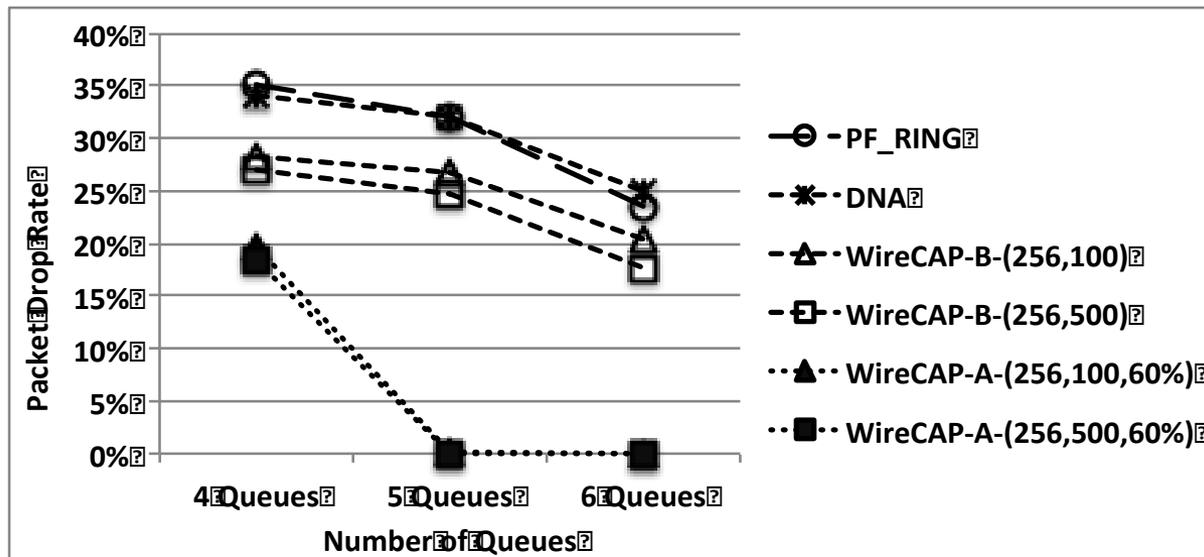
WireCAP Evaluation – Packet Capturing



WireCAP vs. existing packet capture engines

- **WireCAP demonstrates superior buffering capabilities for short-term bursts of packets**
- **The buddy-group-based offloading mechanism achieved significant improved performance**

WireCAP Evaluation - Packet Forwarding



WireCAP vs. existing packet capture engines

- **WireCAP's packet forwarding function is capable of supporting middlebox applications**
- **The *buddy-group-based offloading* mechanism can achieve a significant improved performance**

WireCAP vs. existing packet capture engines

Engines	Efficiency	Handle Short-term load imbalance ?	Handle Long-term load imbalance?
PF_RING	✗	✗	✗
DNA	✓	✗	✗
NETMAP	✓	✗	✗
WireCAP	✓	✓	✓

WireCAP Status

- **Publication**

- Wenji Wu, Phil DeMar, “WireCAP: a Novel Packet Capture Engine for Commodity NICs in High-speed Networks,” IMC’14, November 5 – 7 2014, Vancouver, BC, Canada.

- **Patent pending**

- **WireCAP website:**

- <http://wirecap.fnal.gov>

- **WireCAP source code is available:**

- Please contact Fermilab’s Office of Partnerships and Technology Transfer (OPTT) <optt@fnal.gov> to obtain a copy of WireCAP

WireCAP Summary

- **WireCAP is a packet capture engine for commodity network interface cards (NICs) in high-speed network**
- **WireCAP can be used to support middle-box-type applications**
- **WireCAP provides a new packet I/O framework for commodity NICs in high-speed networks**

Questions?

WireCAP website:

<http://wirecap.fnal.gov>