



中科院计算所  
INSTITUTE OF COMPUTING TECHNOLOGY, CAS

# PEARL

## **P**rogrammable **V**irtual Router **P**latform **E**nabling Future Internet **I**nnovation

**Hongtao Guan**

Ph.D., Assistant Professor

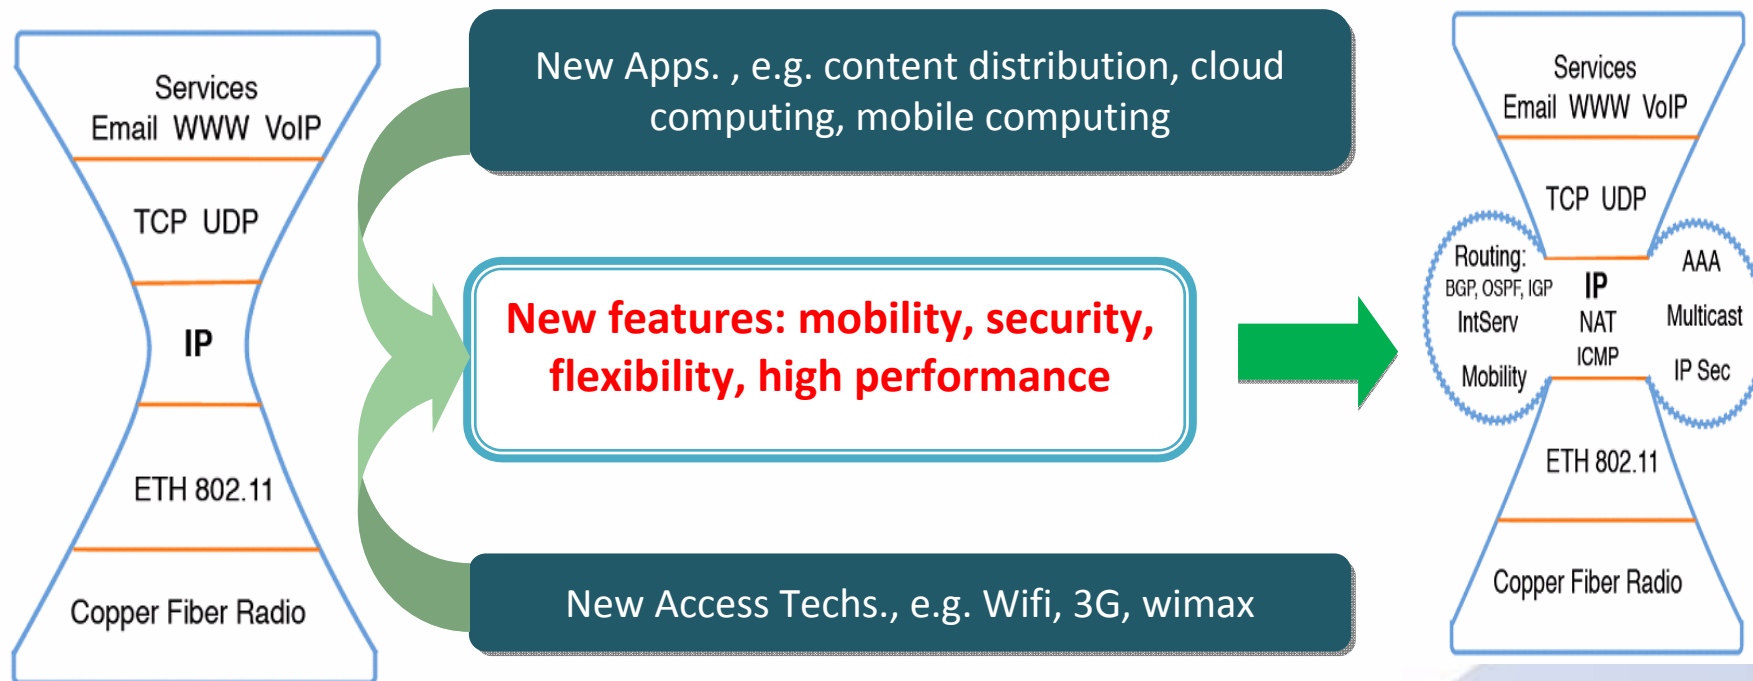
Network Technology Research Center

Institute of Computing Technology, Chinese Academy of Sciences

email: [guanhongtao@ict.ac.cn](mailto:guanhongtao@ict.ac.cn)

<http://fi.ict.ac.cn>

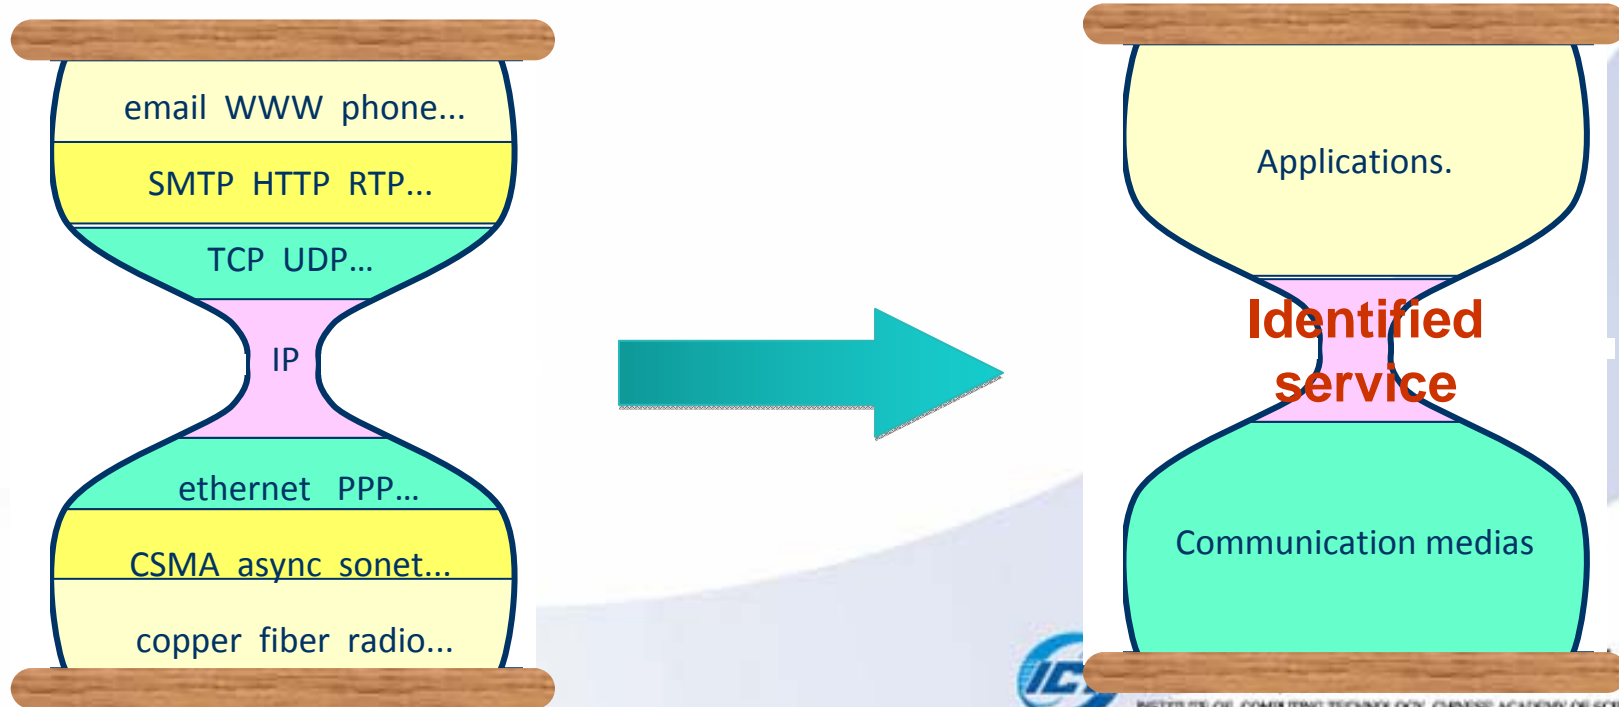
# Future Internet



- Internet is not designed for what it's doing now
- We need new Internet architecture, techniques and algorithms

# Future Internet Architecture Research

- Content Orientated Future Internet (also called NDN)
  - Users are interested in *what* content – not *where* it is
- Mobility First, XIA, Nebula ...
- SOFIA (Service-Oriented Future Internet Architecture)
  - IEEE Infocomm 2011 poster



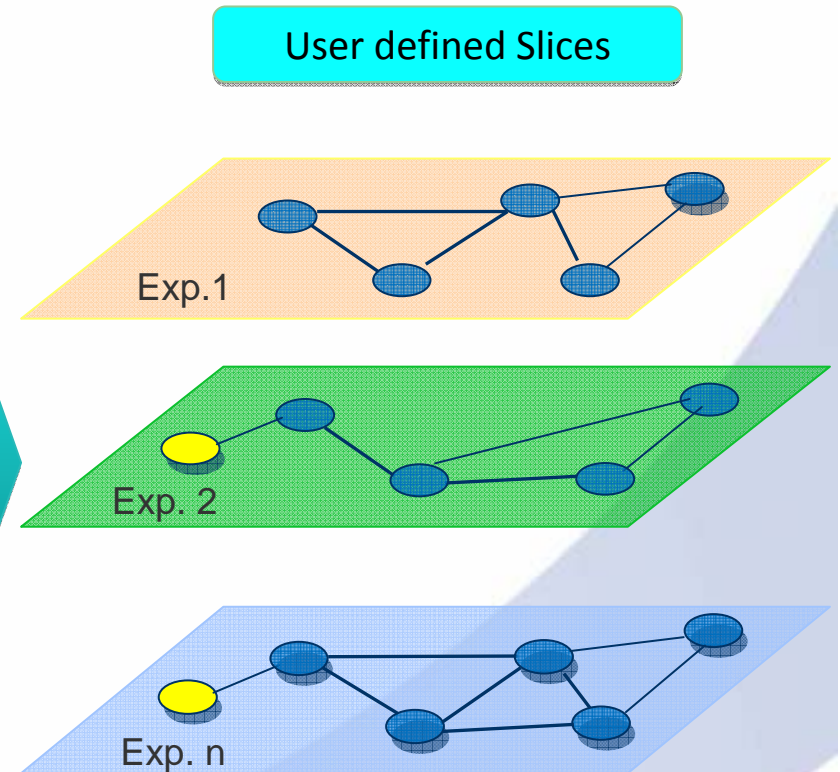
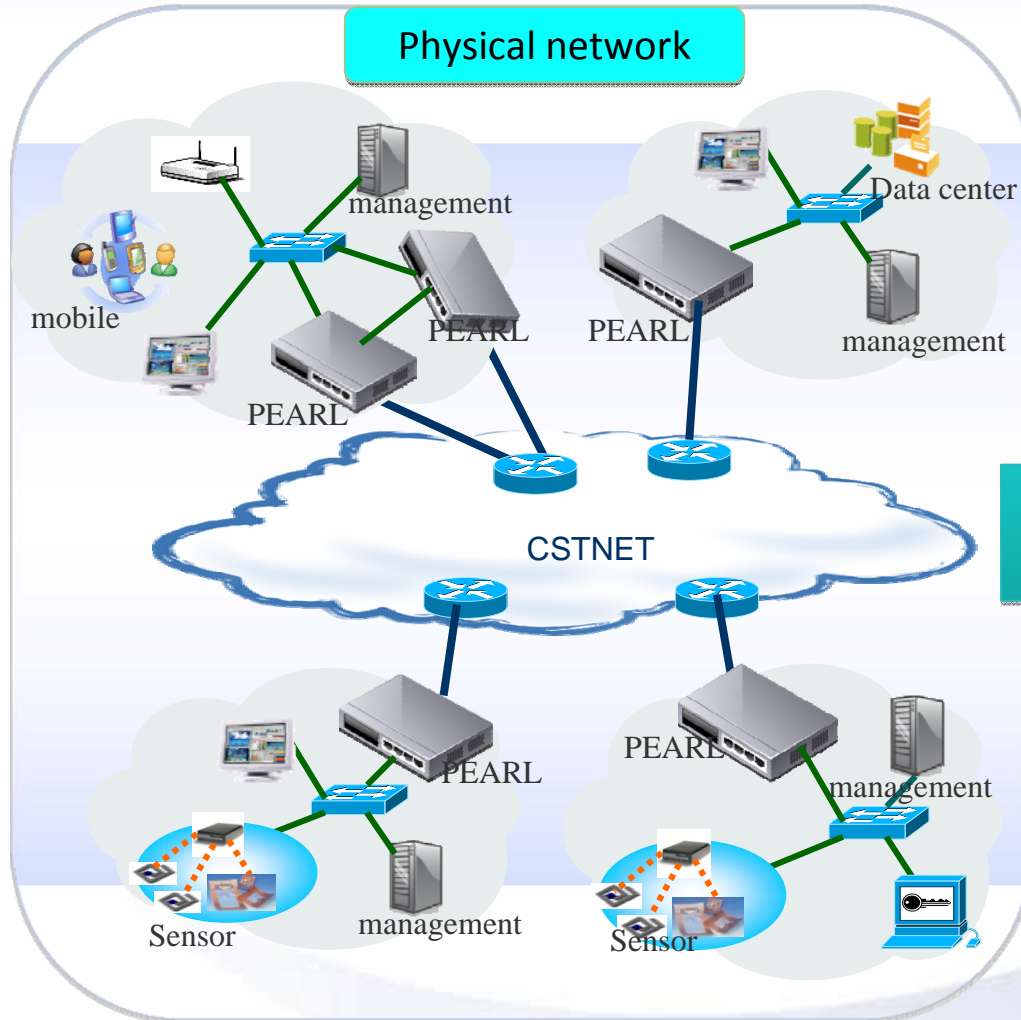
# Motivation of Testbed

- **Deployment** on real networks is an inevitable step for new architecture
- Internet is a good testbed, but it has many limitations.
- An example about router buffer size problem.
- A testbed dedicated to experiments is necessary.

# Requirements for testbed

- Data/control planes need to **be easily changed**
  - SOFIA is not mature
- Running multiple protocols in parallel, **but independently**
  - SOFIA protocols, current Internet protocols, others
- **High performance** in terms of computation and I/O
  - Routing, content cache, flow-balance, security...

# Planning Testbed



# Goals of PEARL

- Programmable virtual Router Platform with following features
  - **Programmability**: open to change data/ctrl planes
  - **High performance**: computation capacity, storage and packet forwarding rate
  - **Virtualization**: several virtual data/ctrl planes are separately run in parallel


# Technical Problems

- Trade-off between flexibility and high performance
  - Flexibility → software
  - High performance → hardware



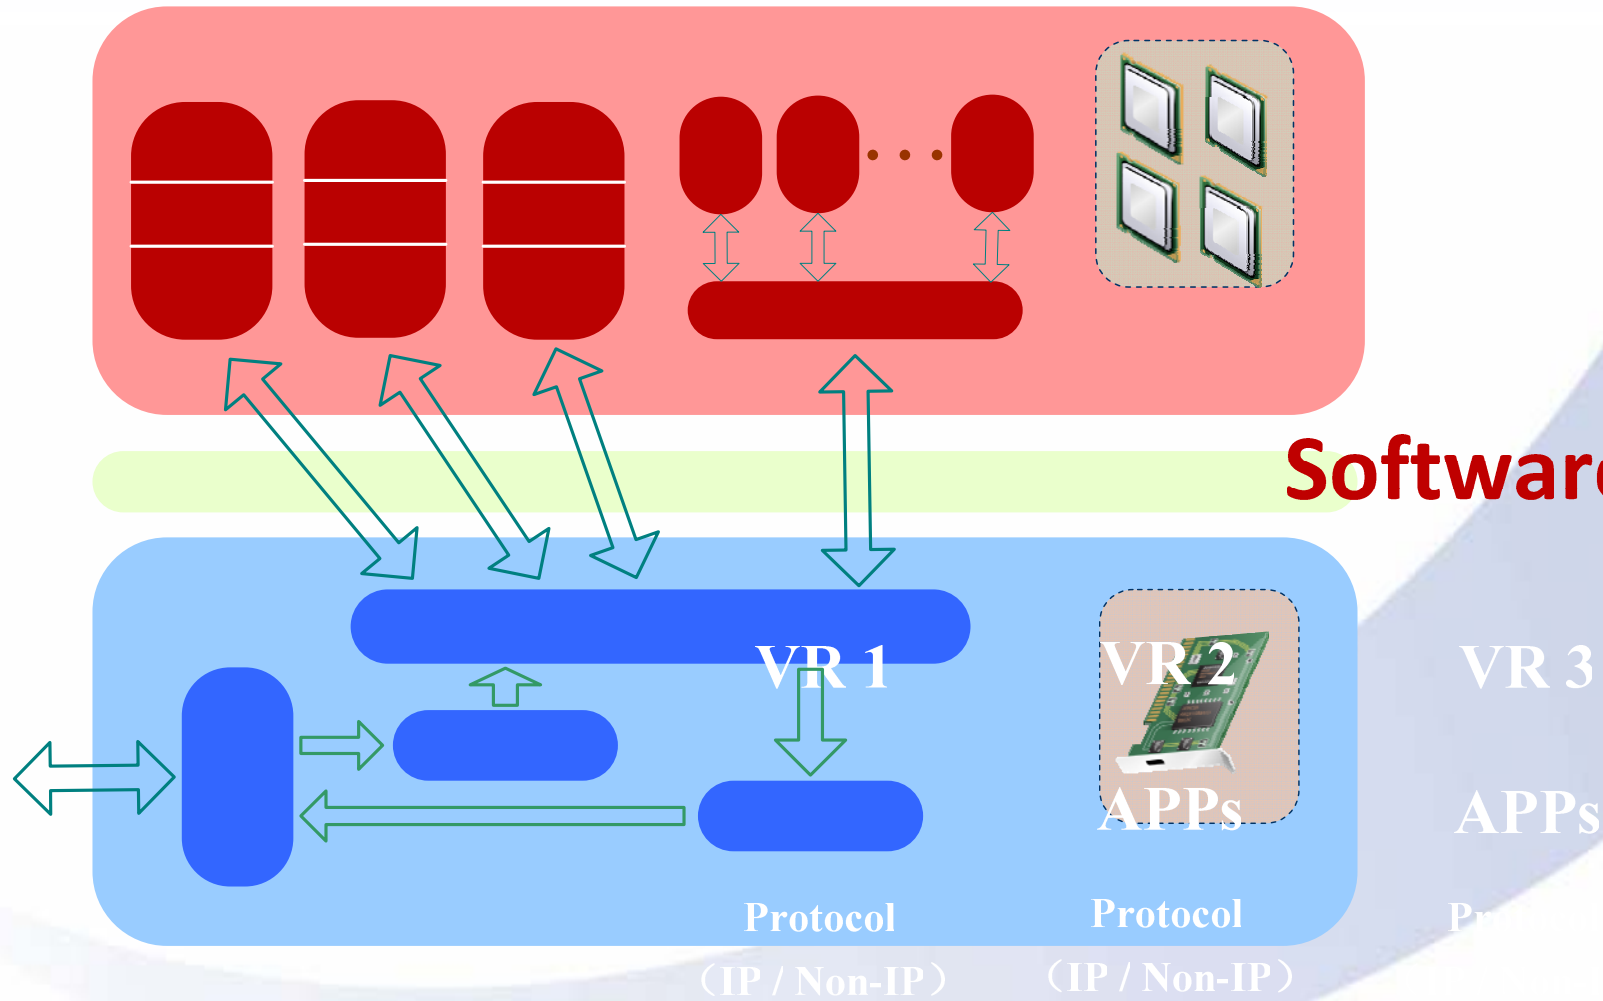
- The impact of virtualization on performance
  - Virtualization techniques: full/para, OS-level virtualization
- Isolation: Independently running virtual routers in parallel

# PEARL v1



Hardware	Description
CPUs	2 CPUs 8 cores Intel E5506 2.13Ghz
Memory	16G DDR3
PCIe bus	Version 1.1
Processing Cards	At most 2 cards 4 GE/2 OC-192 POS ports

# PEARL v1



Software Platform

VR 3


APPS

Protocol

# Enabling Techniques

## Software Platform

Multi-core  
generalized CPUs



Forwarding Information Base  
(FIB) APIs for programmability

OS level virtualization with  
resource guarantees



Packet  
processing card

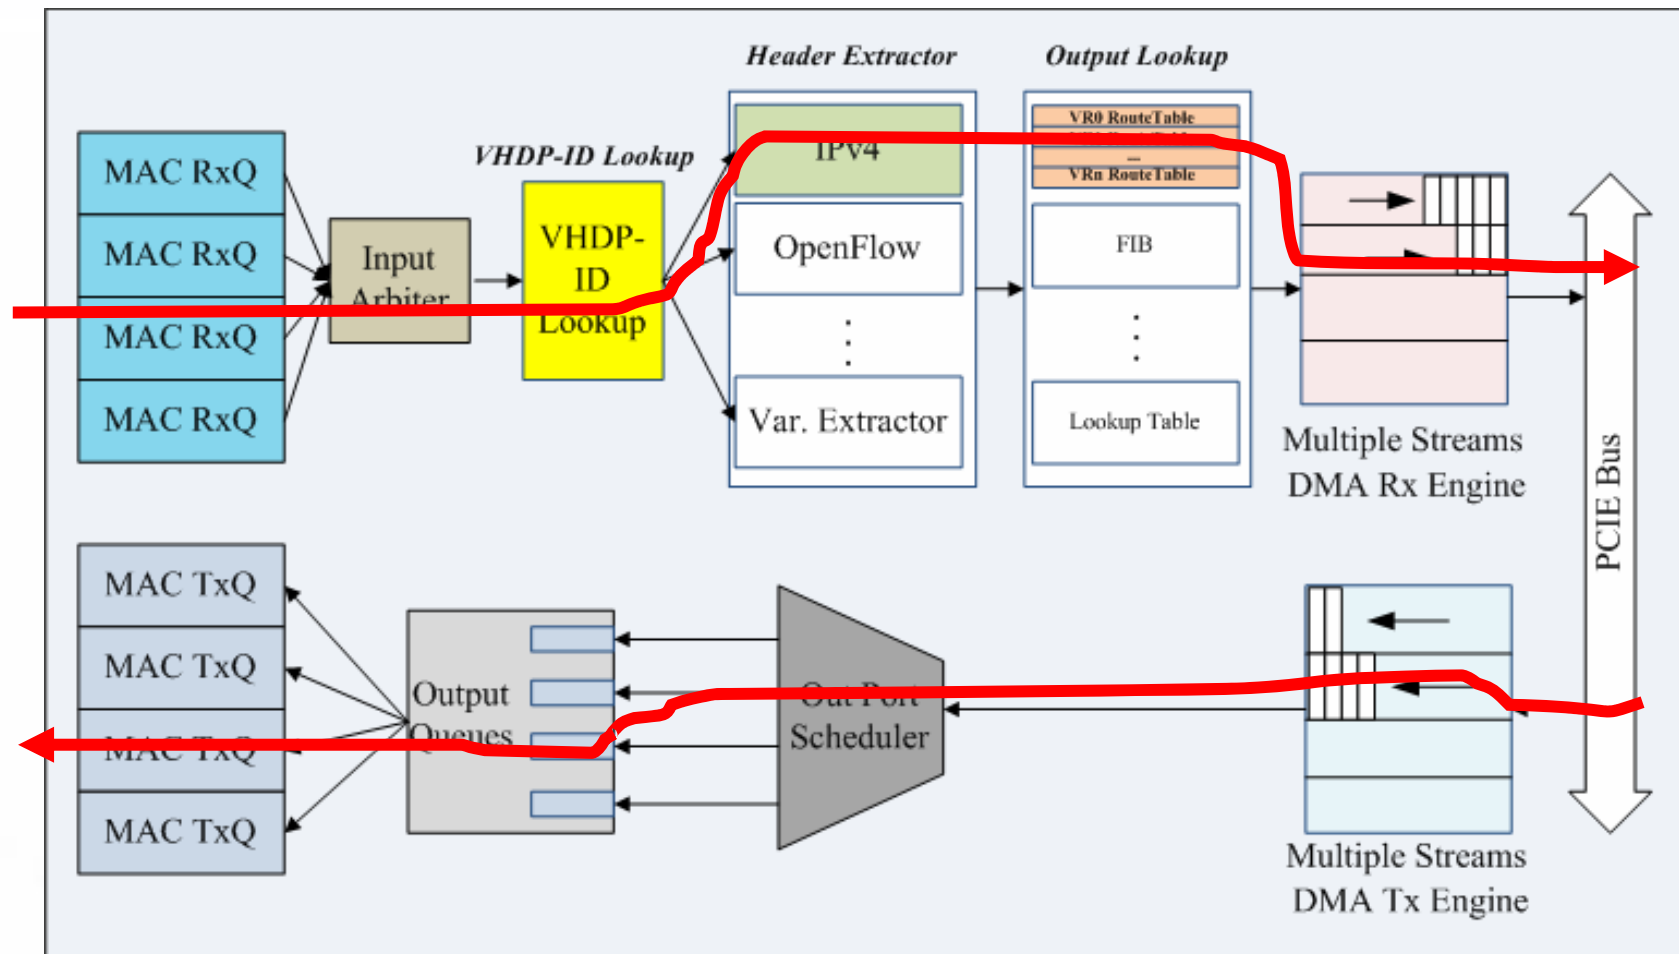


## Hardware Platform

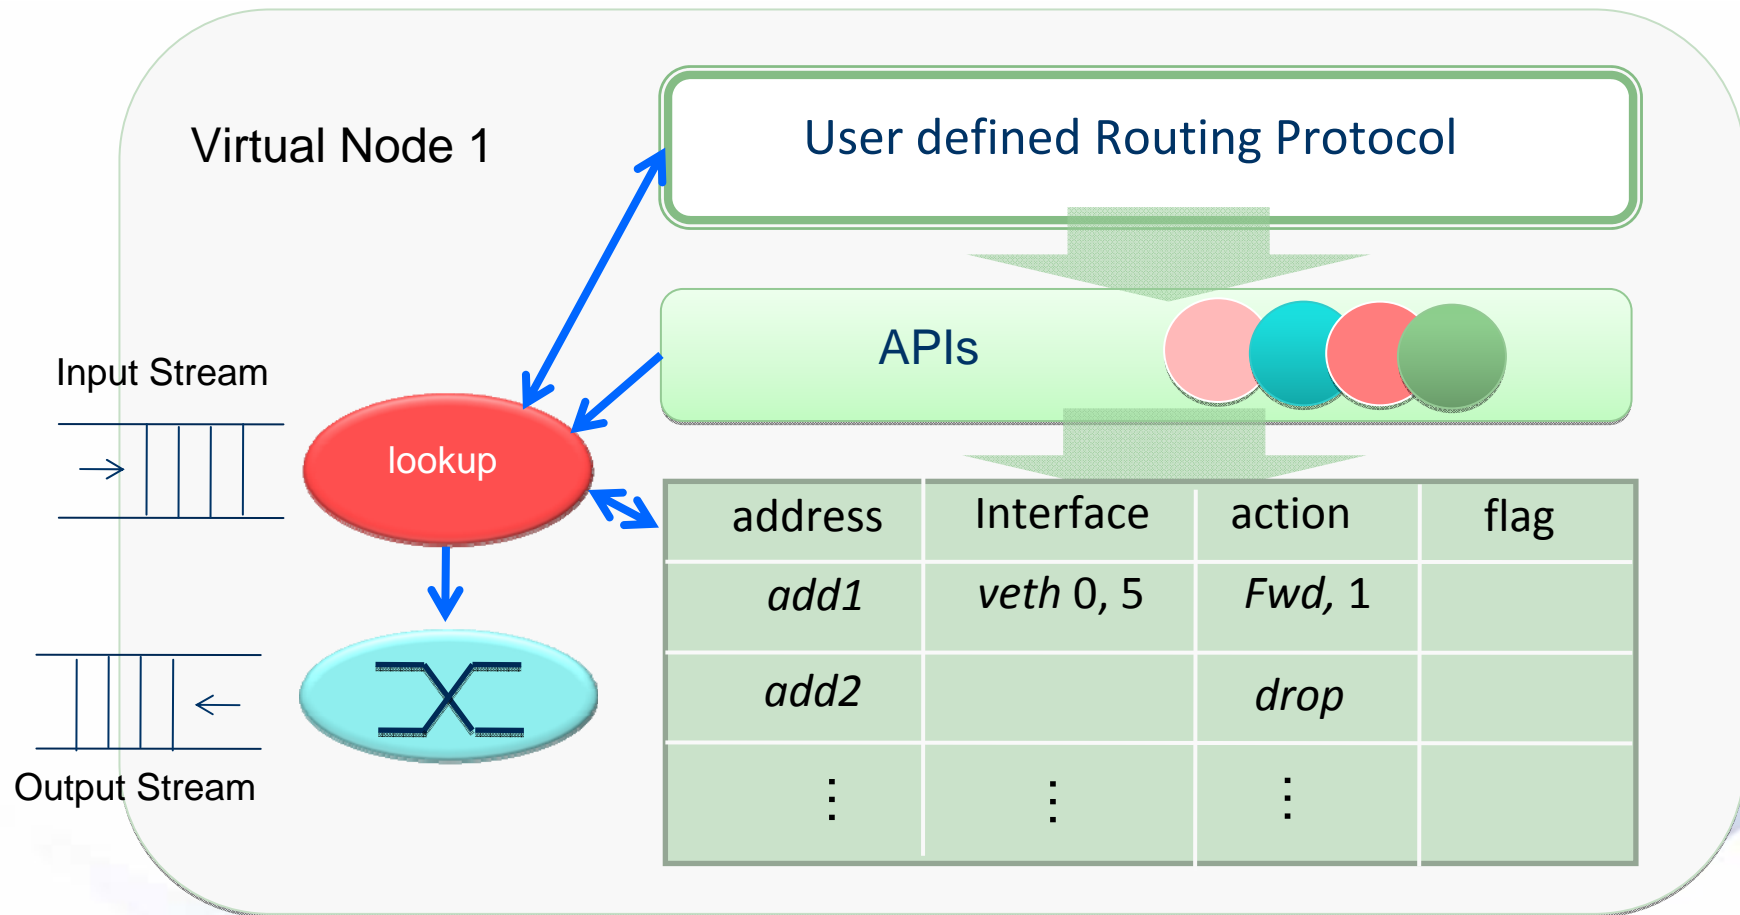
Configurable multiple virtual  
data planes (VDP) with own  
forwarding tables

I/O virtualization for isolation

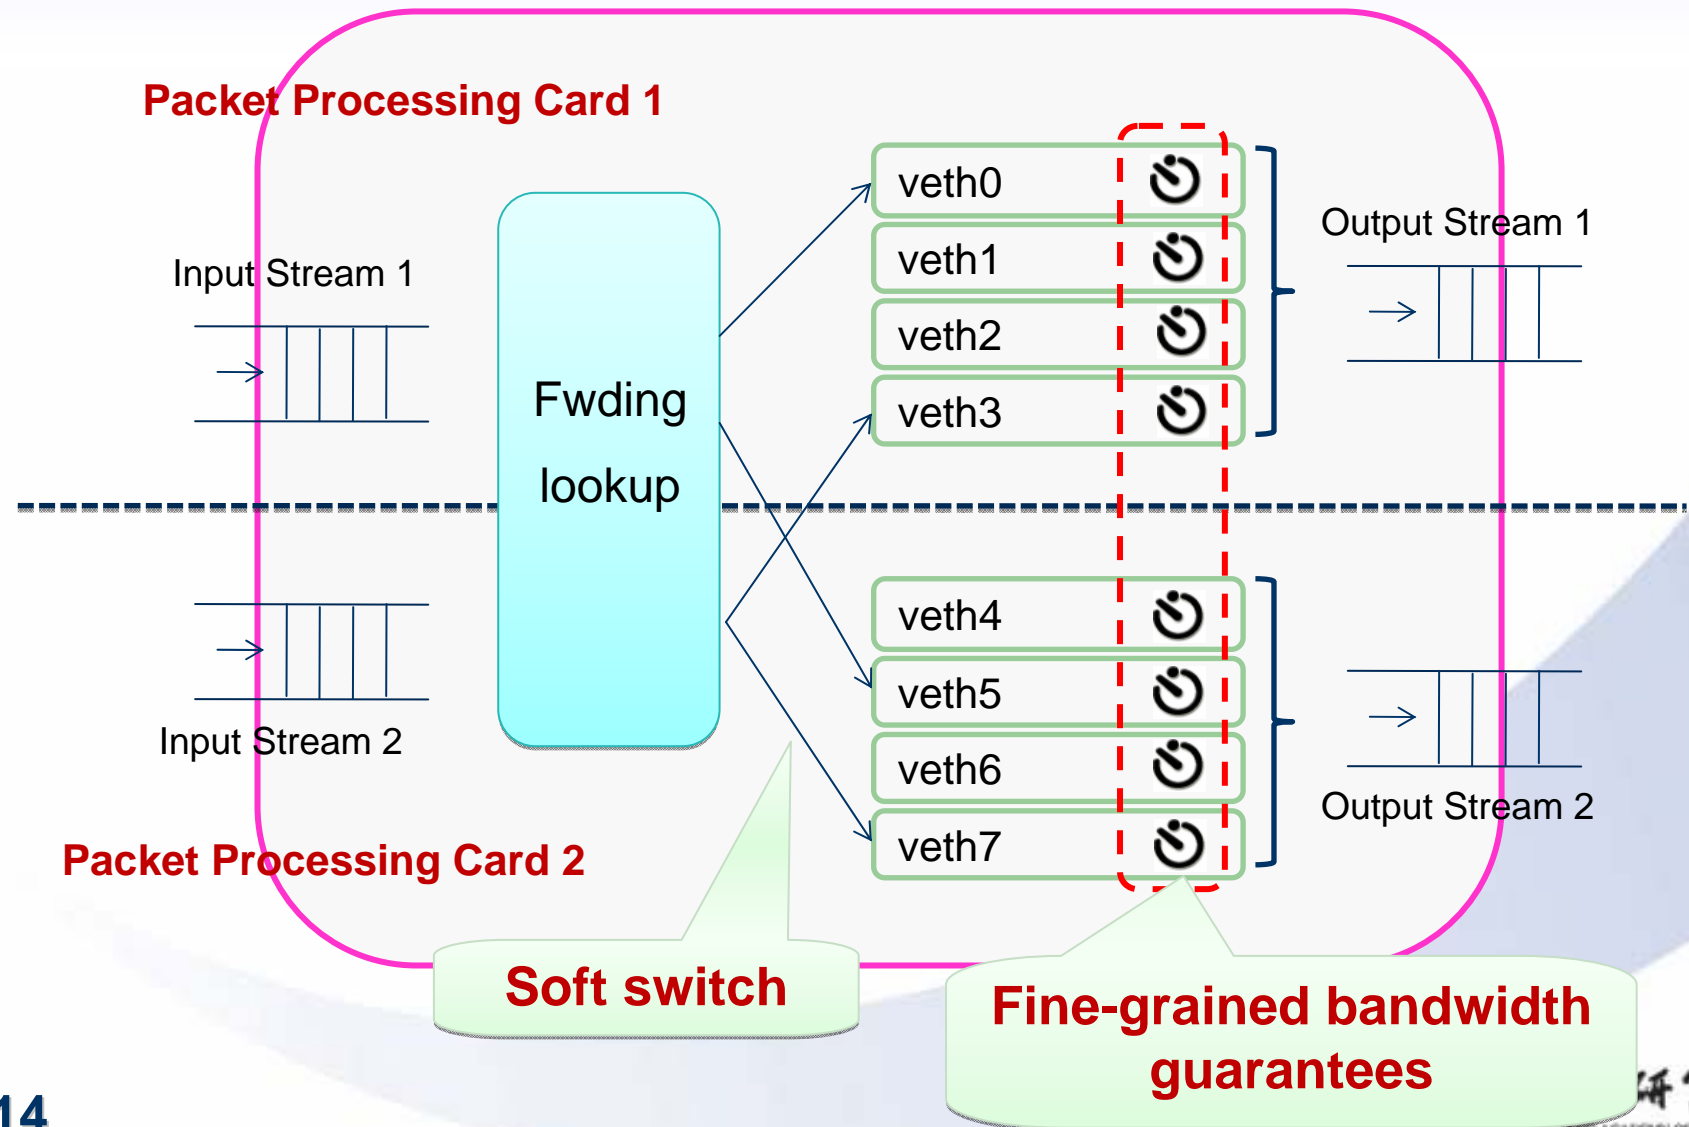
# Packet Processing Card



# Software Platform



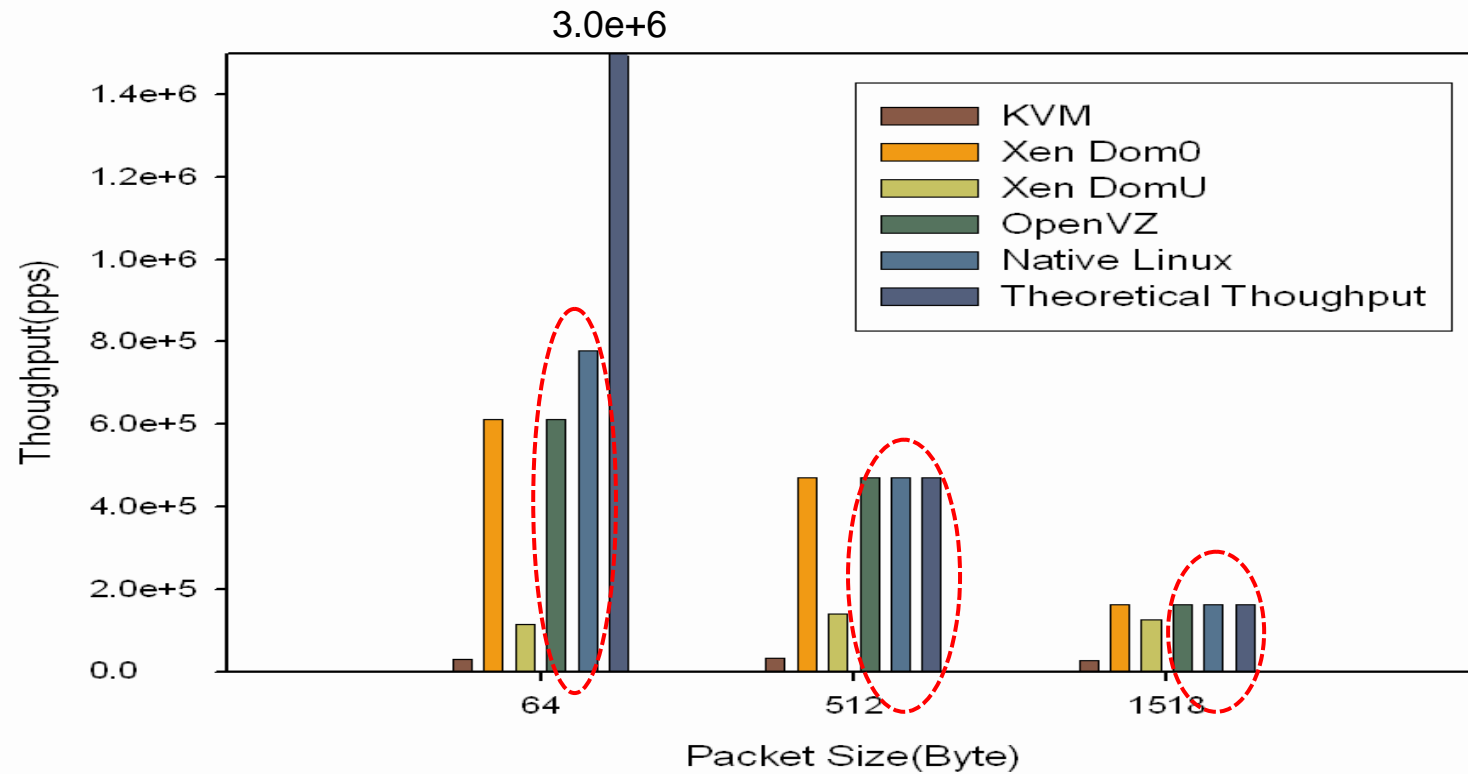
# Switch Component



# Virtualization Techniques

- Full/Para virtualization
  - **Pro**: hypervisor between OS and hardware, resulting in robust
  - **Con**: low I/O performance
  - Examples: KVM, Xen
- OS-level virtualization
  - **Pro**: containers of processes, small overhead
  - **Con**: sharing a common Kernel
  - Examples: OpenVZ, LXC, linux Vserver (Planetlab)


# Impacts on performance



- OS-level virtualization has a comparable performance to native linux

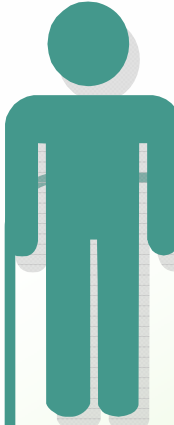
# Two roles of the virtual router

- Manager



Virtual router  
**manager** setup  
Physical router  
and virtual  
routers

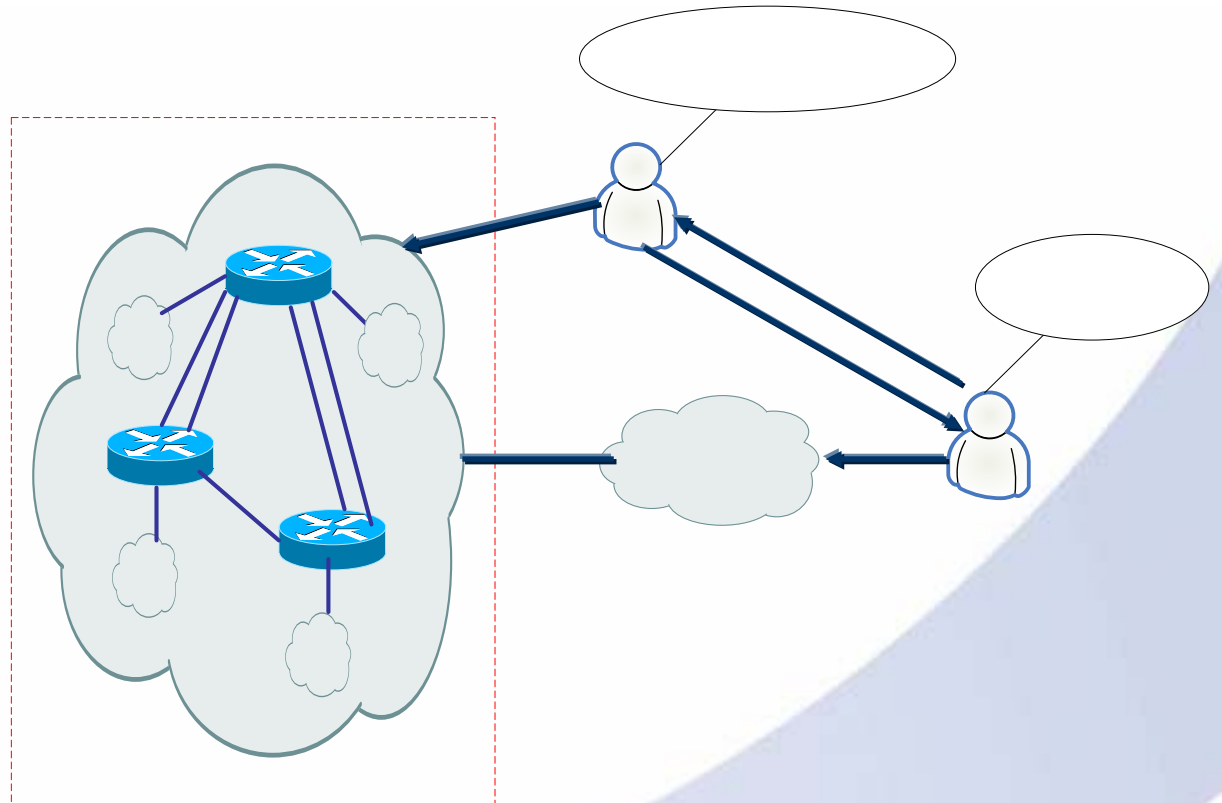
- Experimenter



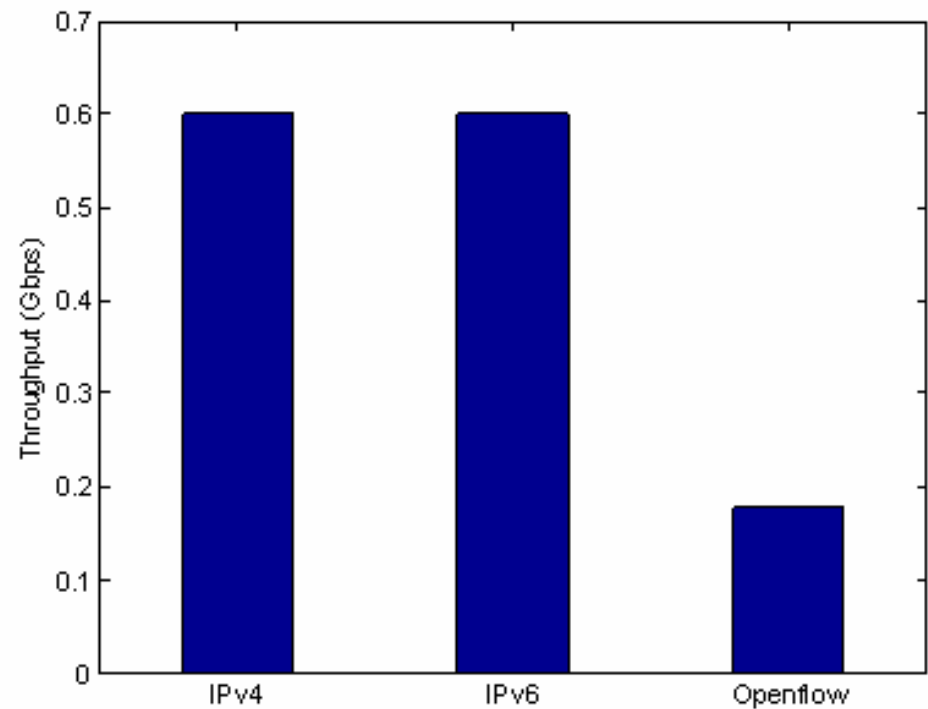

Virtual router  
**Experimenter**  
configure virtual  
routers

# Deployment

- Step1: experimenter requests to manager
- Step2: manager setup virtual routers
- Step3: manager tell experimenter each router's accounts and passwords
- Step4: experimenter login to the virtual router



# Evaluation

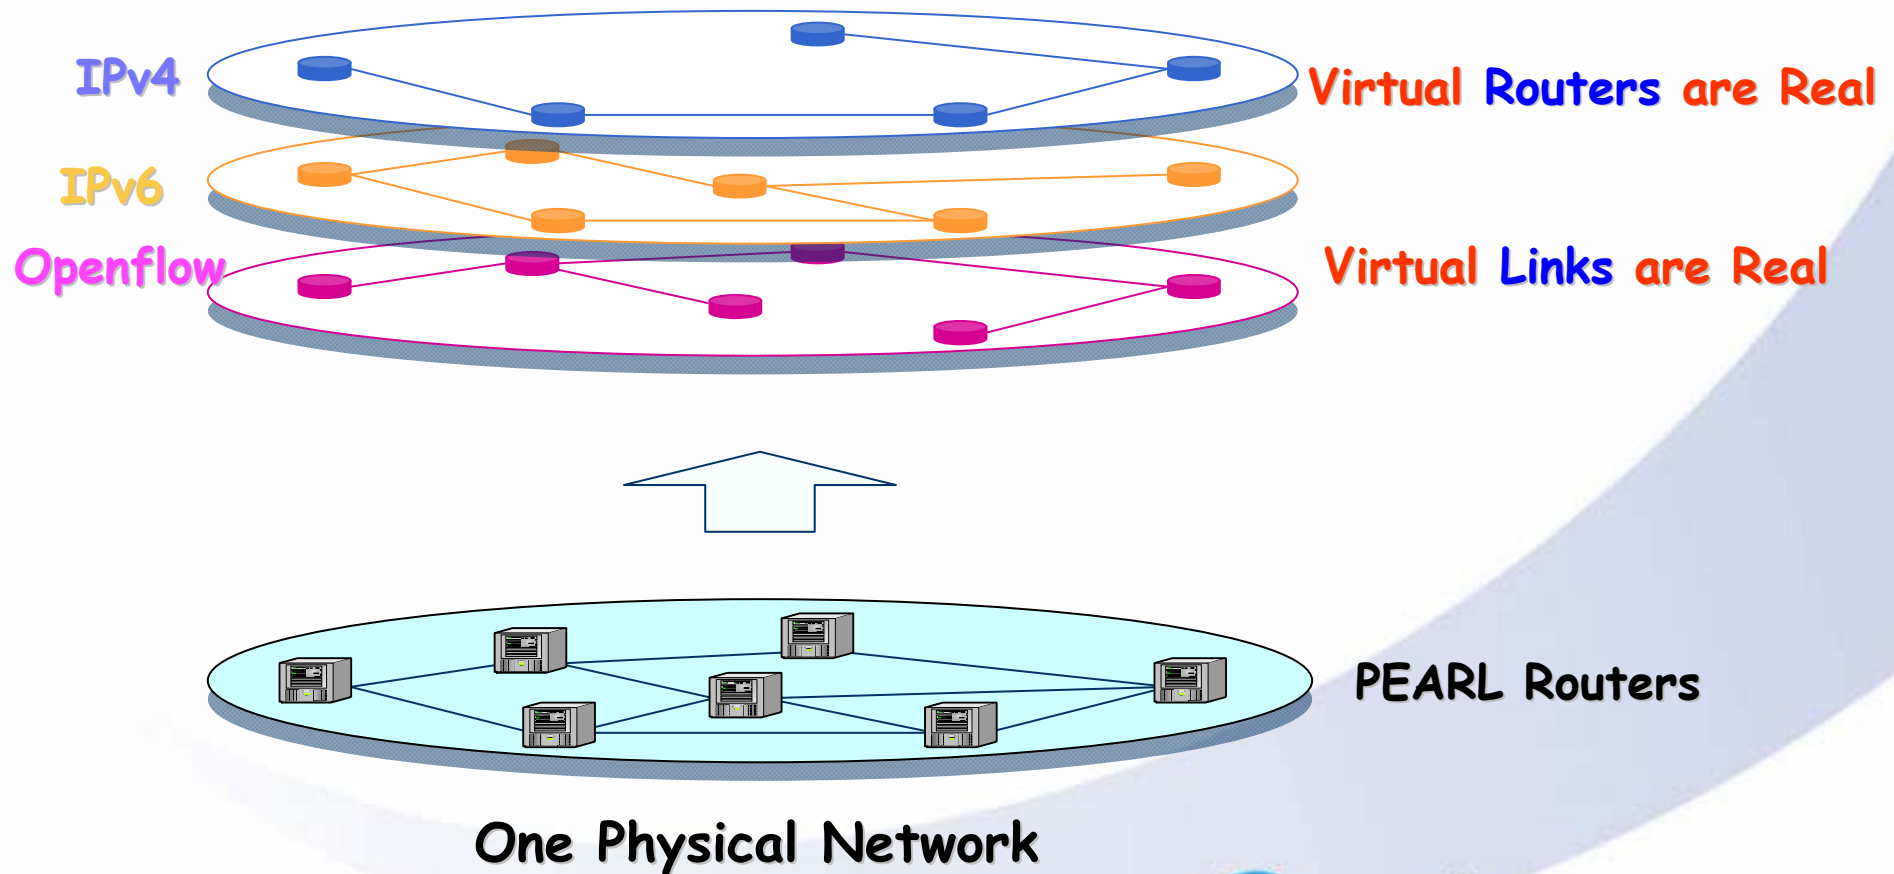


# Planning Experiments

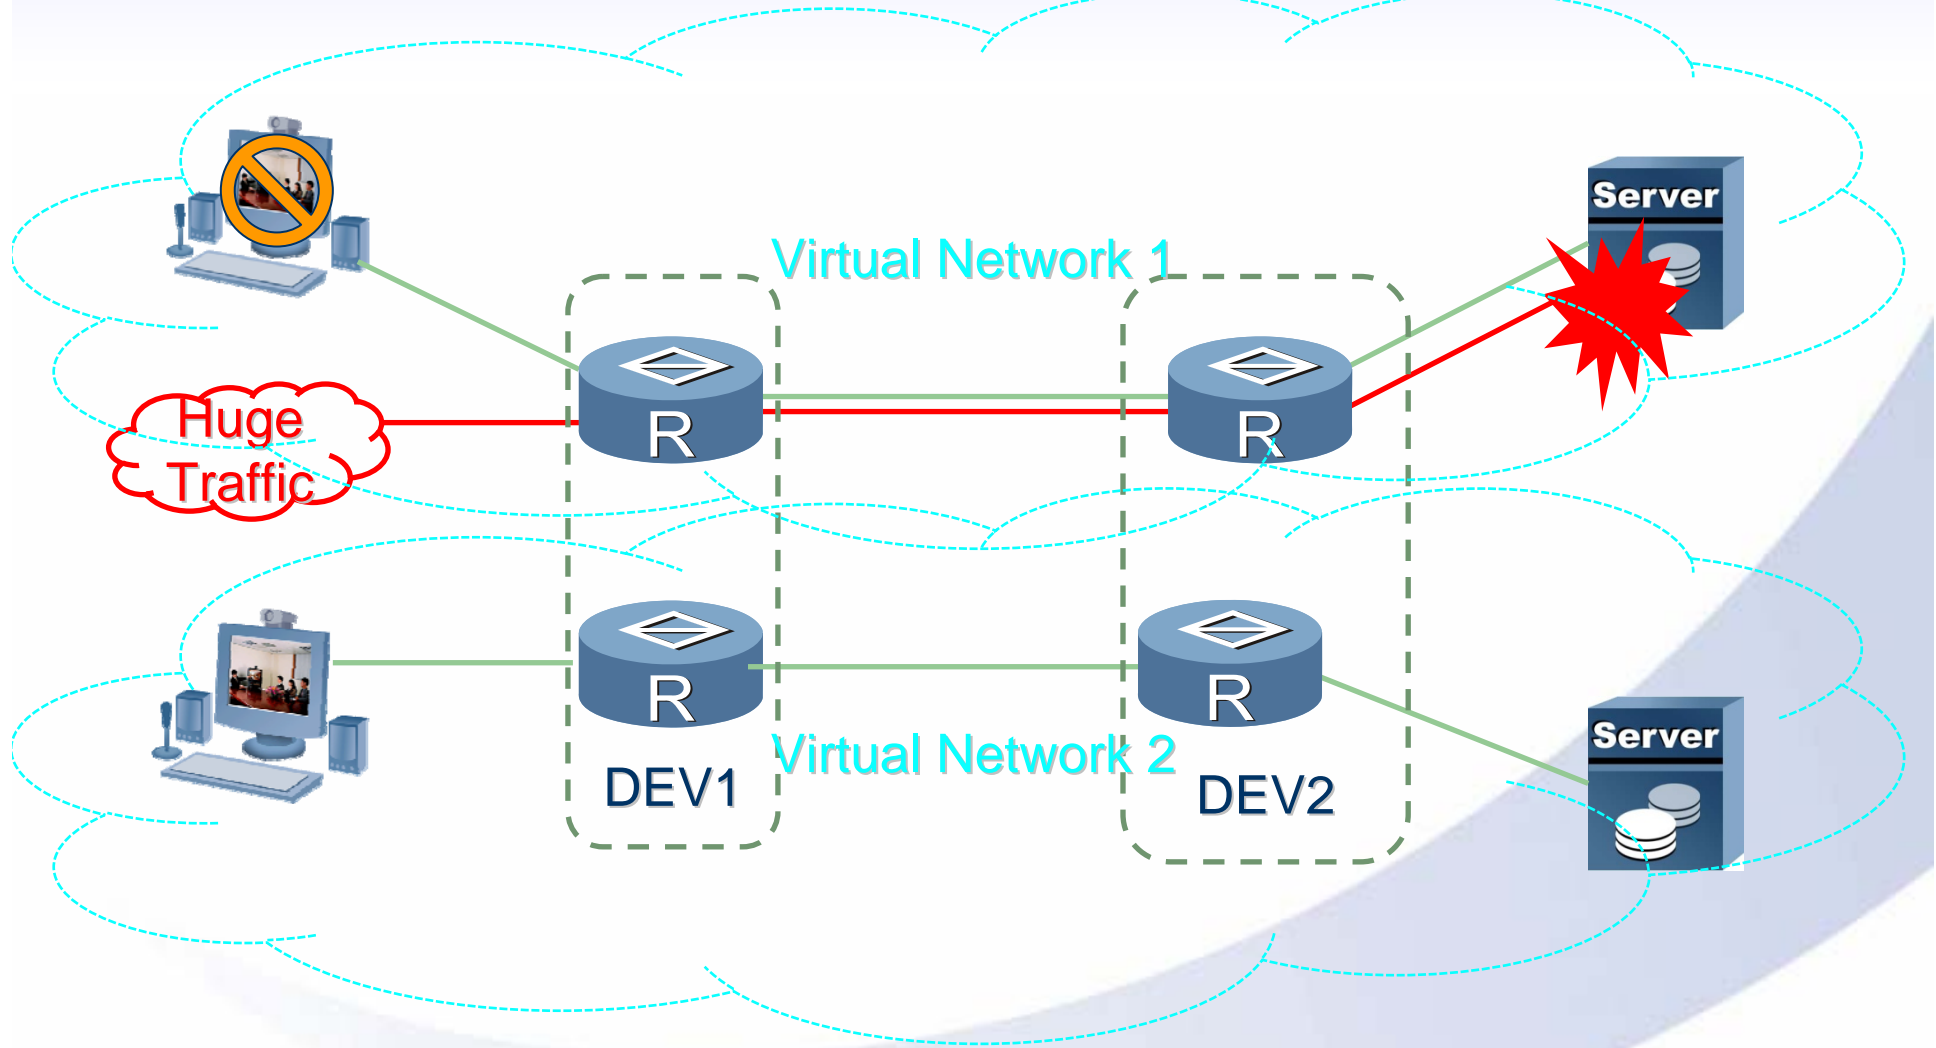
- Multiple experiments running on isolated network slices concurrently
  - Named Data Networking
    - Routing and content distribution performance
  - Openflow
    - Not only switch, but routing
  - ...

# Multiple Virtual Networks with Different Protocols

## Multiple Virtual Network




# Isolation between Virtual Networks



# PEARL v2

- Chassis platform
  - 14 cards on a physical node
  - Card with 24 GE ports
  - Card with 4 OC-192 POS
- High Performance
  - 320 Gbps



**Thank you for your attention**

**Hongtao Guan**  
**guanhongtao@ict.ac.cn**

**<http://fi.ict.ac.cn>**

