Thinking about Extensibility and Scalability in OpenFlow Networking

(CANS 2011)

Zhiliang WANG, Xia YIN, Jun WANG
Tsinghua University

2011-8-9

Outline

- OpenFlow Introduction
- Our ideas on Openflow Research
  - Extensibility
  - Scalability
- Summary

OpenFlow Switching

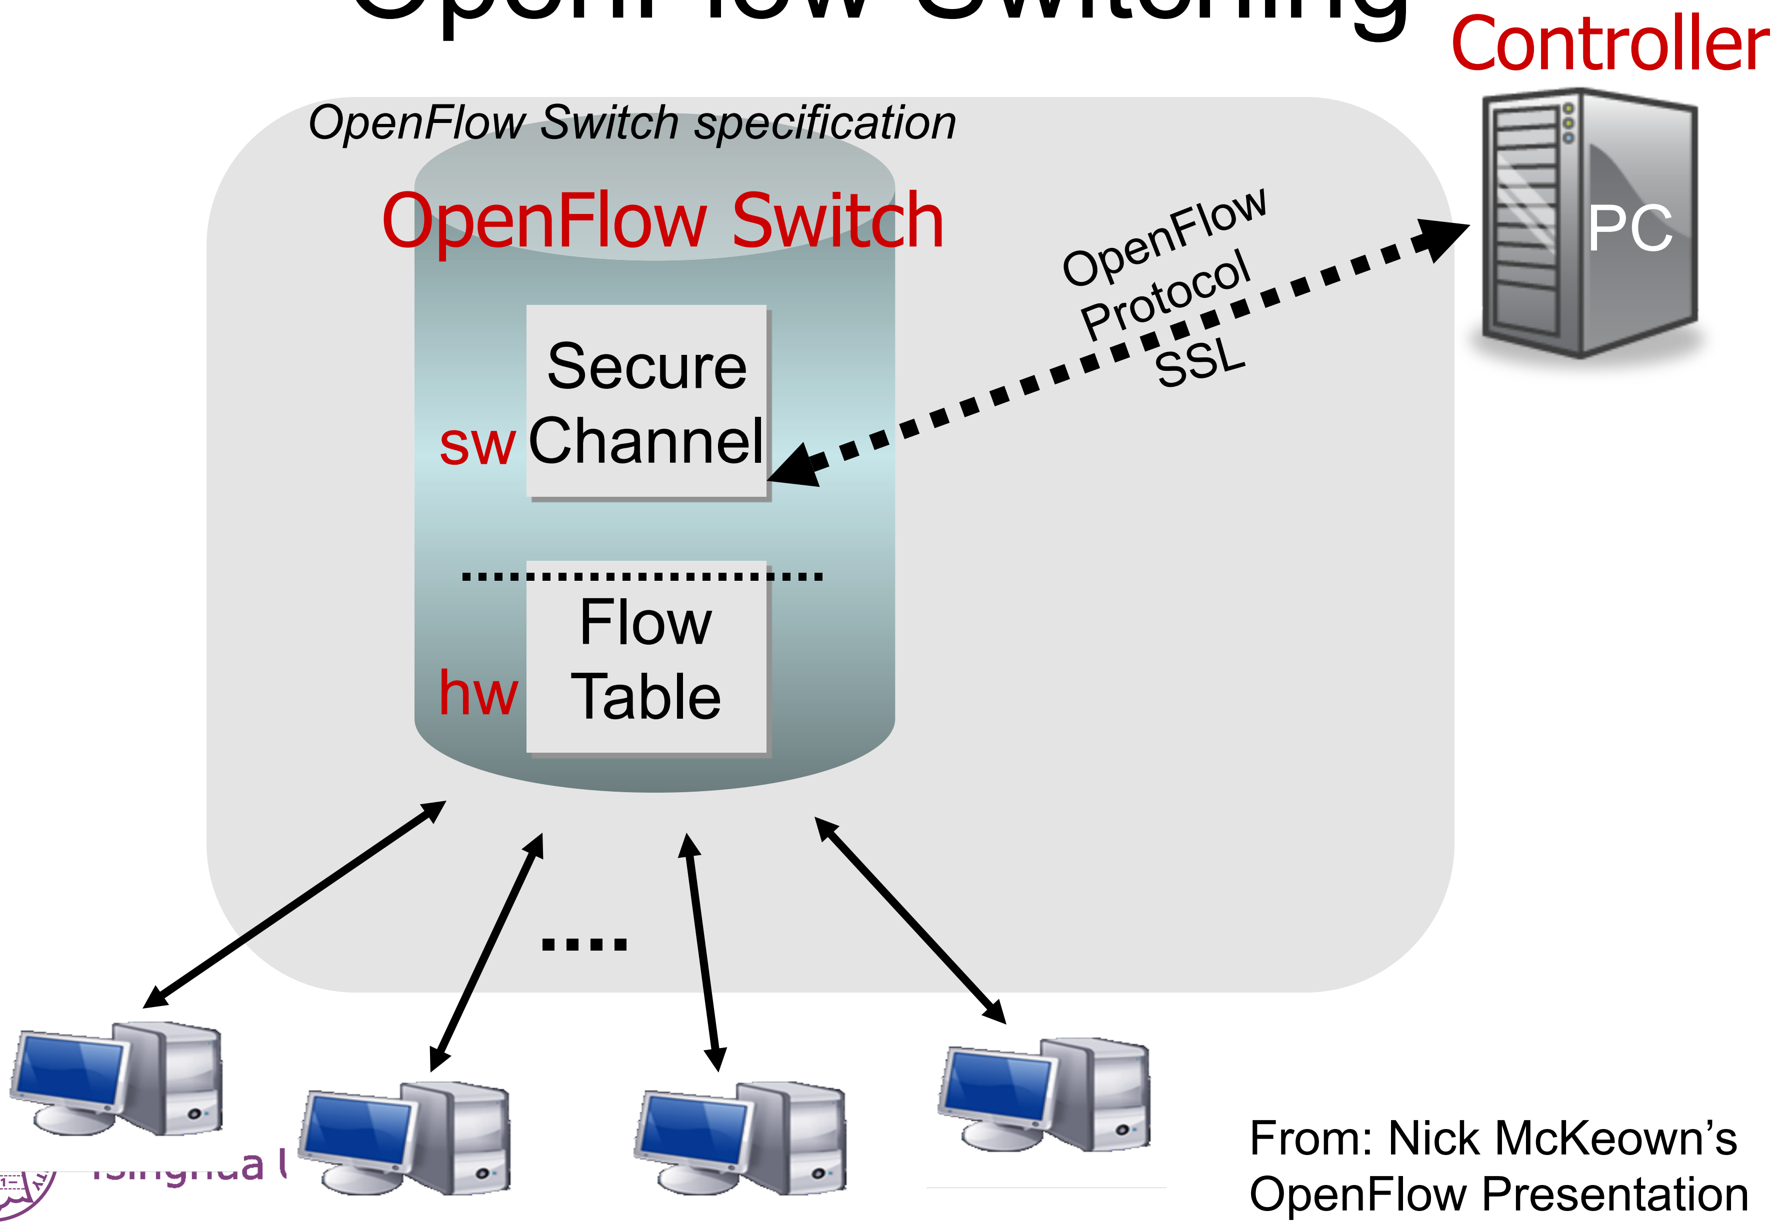

From: Nick McKeown's OpenFlow Presentation

Flow Table(Spec v1.1.0)

Match Fields	Counter	Instructions

Counter: Packets Statistics

Instructions: To change action sets:

- Forward packet to port(s)
- Drop packet
- Send to normal processing pipeline
- Encapsulate and forward to controller

Match Fields:

Switch Port	MAC src	MAC dst	Eth type	VLAN ID	IPv4 Src	IPv4 Dst	IPv4 Prot	TCP sport	TCP dport

OpenFlow is …

- Experimental platform
  - Run experiments in production networks
  - “Enabling innovation on campus”
- New network architecture
  - Open Programmable Networking
  - Software Defined Networking

SDN: Software Defined Networking

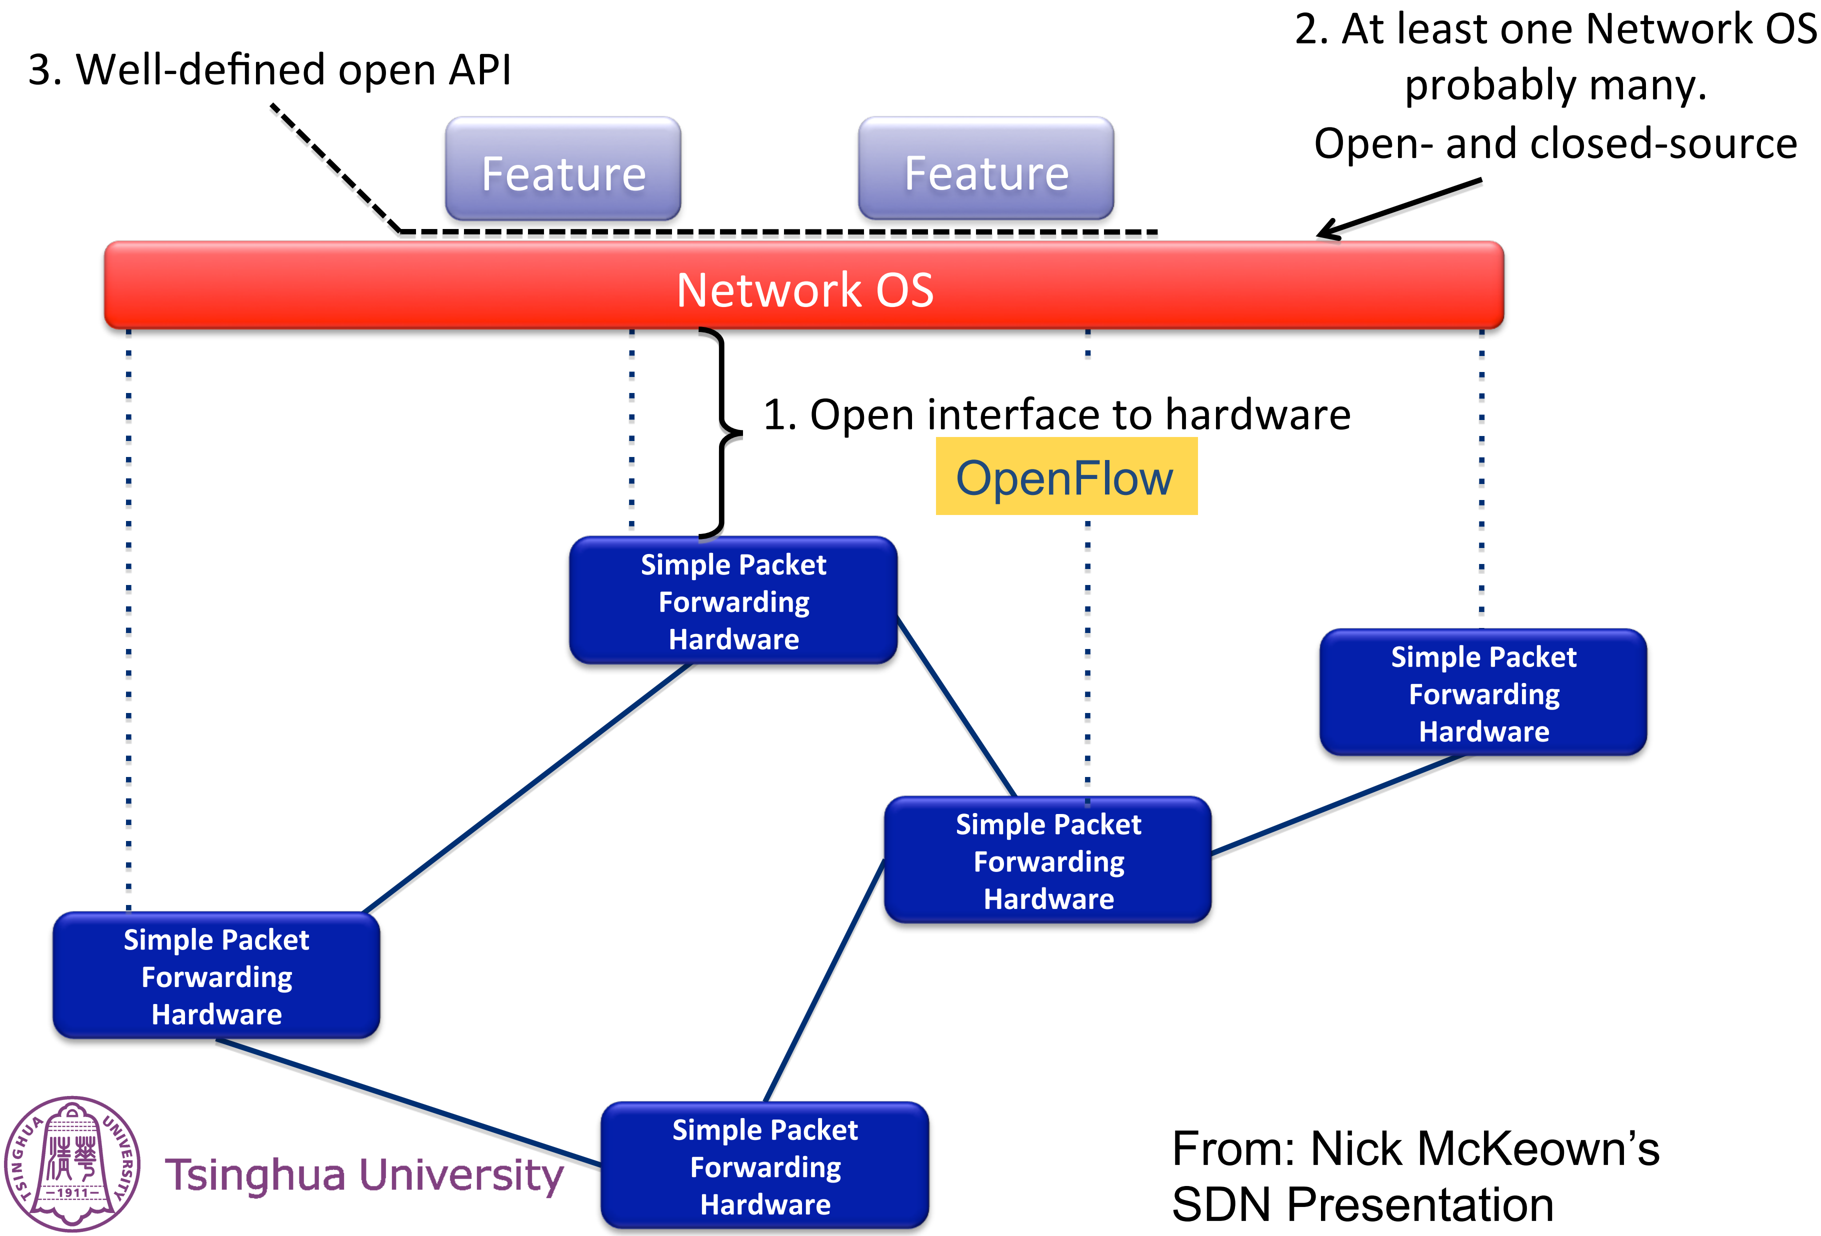

Tsinghua University

From: Nick McKeown's SDN Presentation

ONF: Open Networking Foundation

- a nonprofit organization dedicated to promoting new approach of Software-Defined Networking (SDN)
- Founded in **March, 2011**
- Standardization of OpenFlow Specification
- Member including:
  - ISP, ICP, equipment vendors, networking and virtualization software suppliers, and chip technology providers

Our ideas on OpenFlow Research

Scalability

Extensibility

IPv6 Extensions

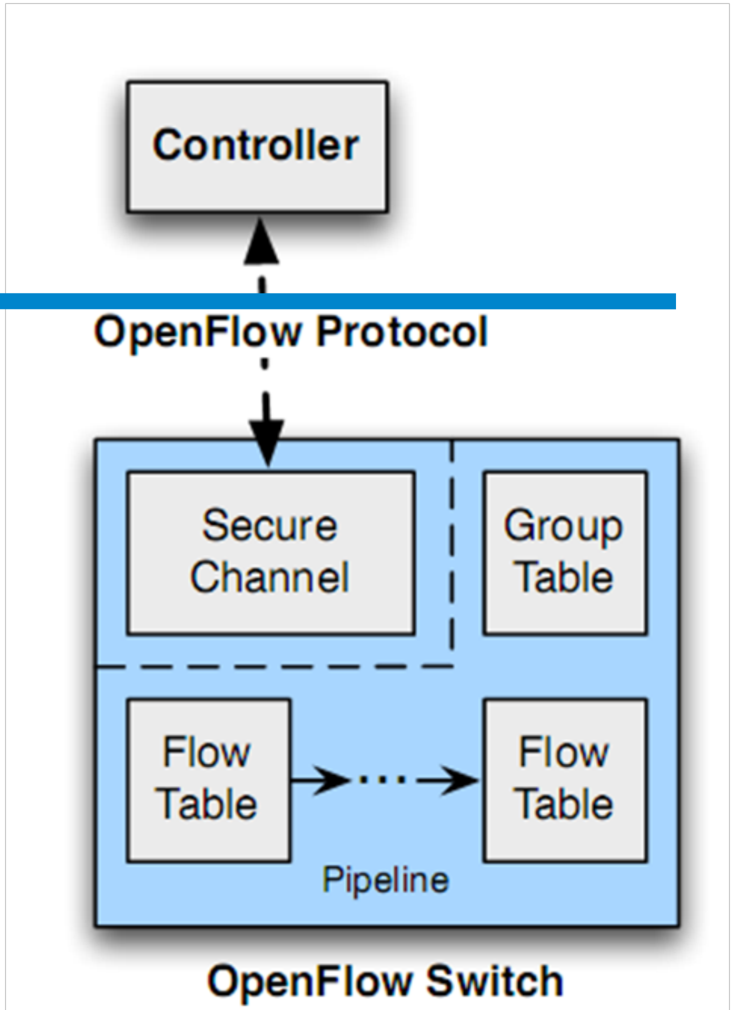

From: OpenFlow Specification v1.1.0

Tsinghua University

清華大學 Tsinghua University

IPv6 Extensions of OpenFlow

IPv6 Extensions of OpenFlow

- Why OpenFlow with IPv6 support?
- The design principle of IPv6 extension
- Next steps

Why OpenFlow with IPv6 support?

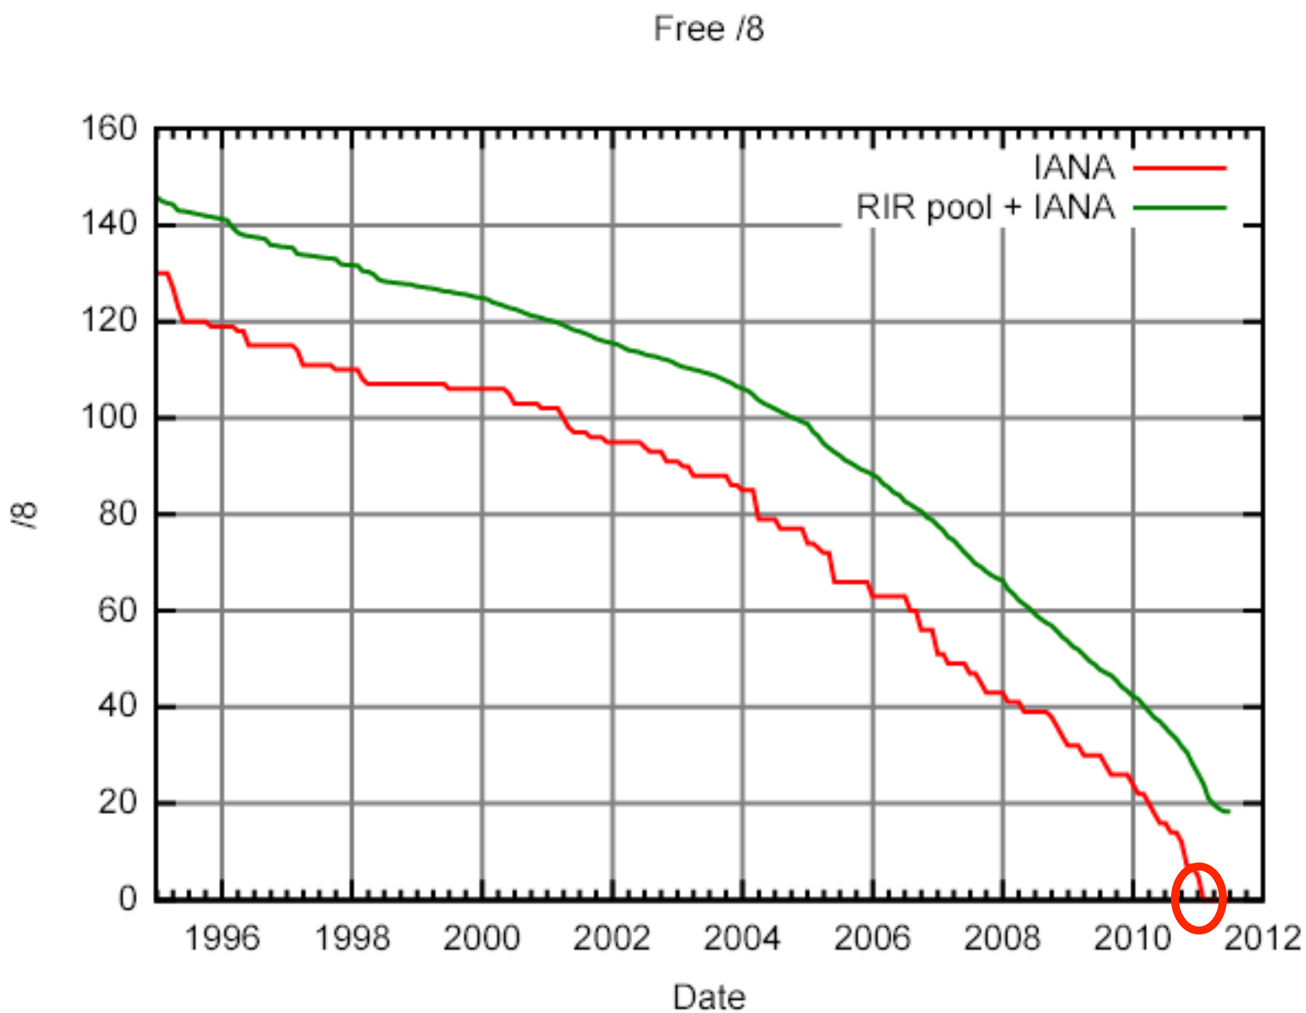

From: wiki

Design Principle

- At least two data structures in OpenFlow needed to extend

Flow Table

Group Table

- Additional instructions or actions needed to design to support the new data structures

Flow Table

- Extending the Flow Table to support IPv6

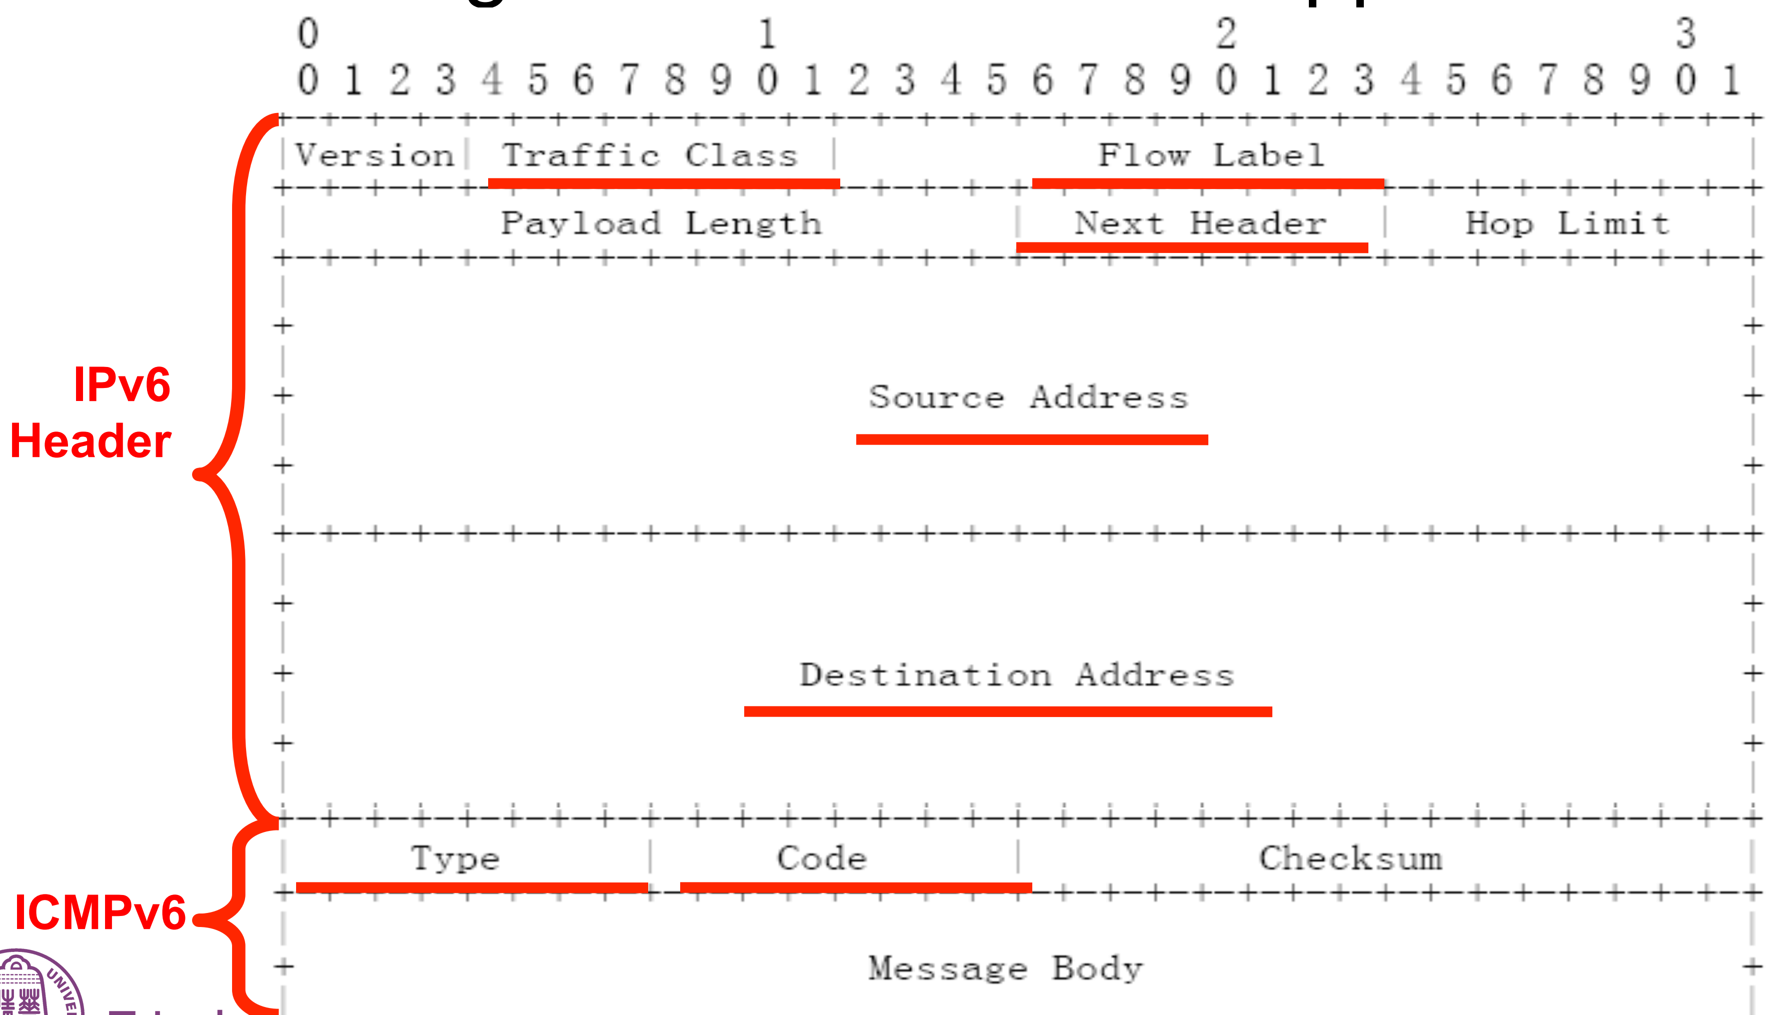

Flow Table

- Flow Table supporting IPv4/IPv6

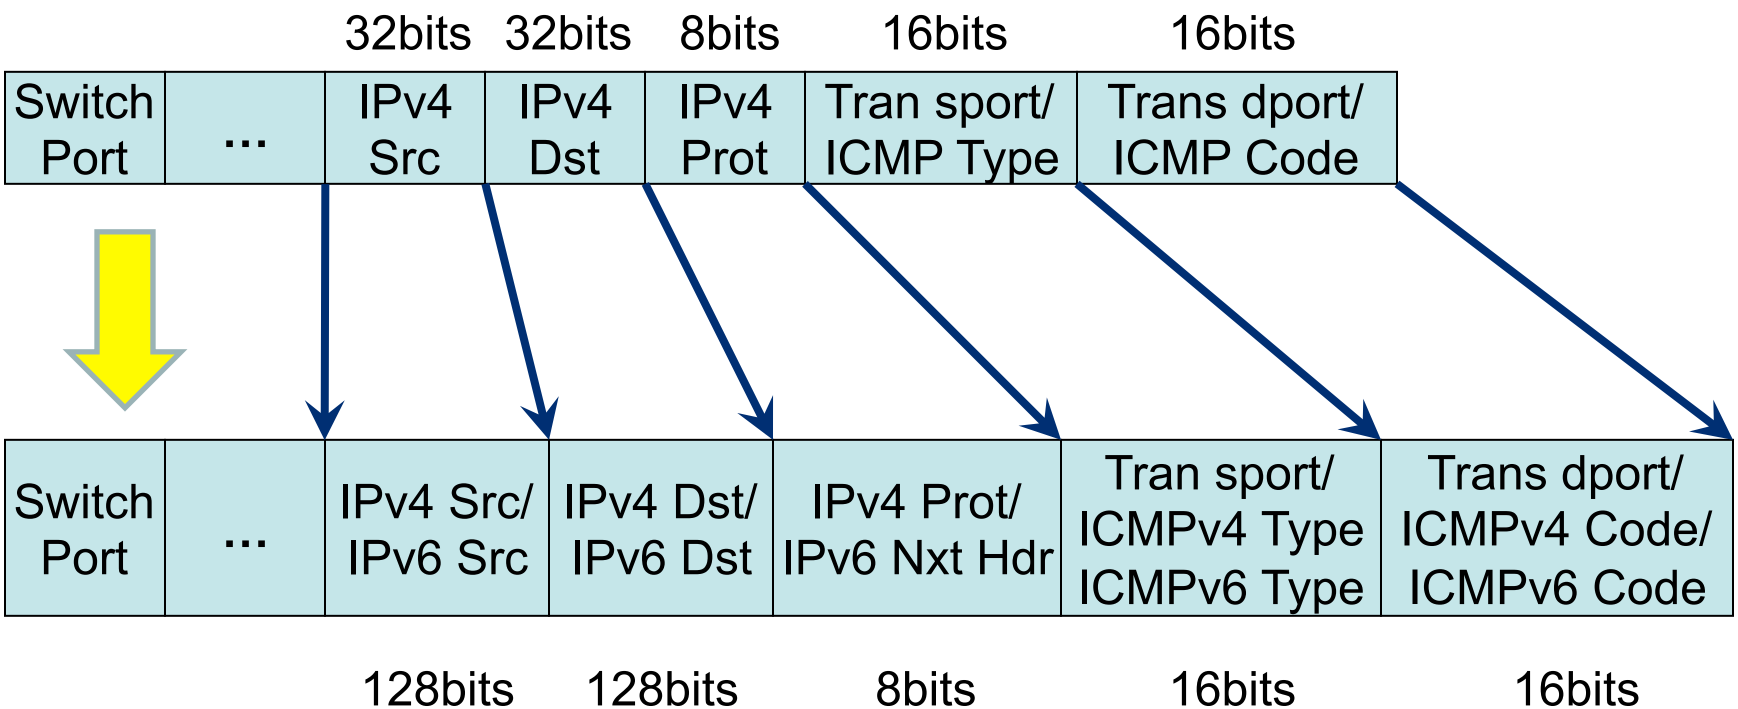

Group Table

- Defined in OpenFlow version 1.1
- Multicast/Broadcast/Anycast can be supported by different group type.

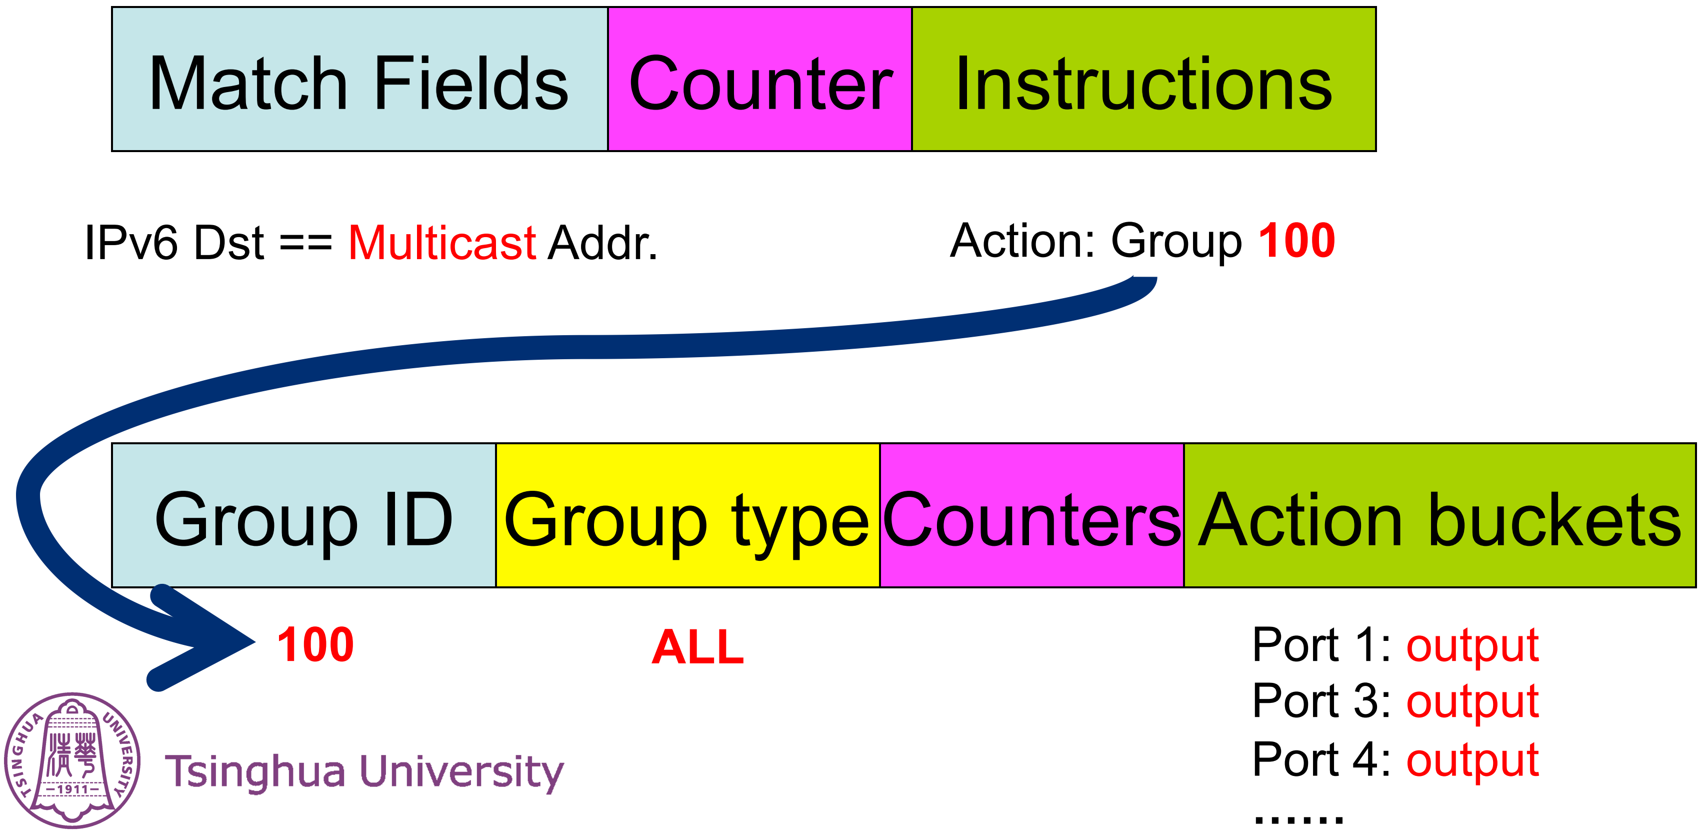

Group Table

- IPv6 has well-known multicast addresses
  - all-nodes multicast address(FF01::1)
  - all-router multicast address(FF01::2)
  - ……
  - These well-known multicast addresses should have predefined group identifier.
  - The corresponding action buckets should also be assigned by the controller as one command.

Next Steps….

- Submit the proposal of OpenFlow IPv6 extensions to ONF
- Prompt standardization of OpenFlow with IPv6 supports
- Demos of OpenFlow IPv6 extensions

Scalability in OpenFlow Networking

Scalability in OpenFlow Networking

- Scalability issues
- Related work
- Our ideas
- More Further…

Scalability Issues

- **Centralized** OpenFlow Control Plane is **lack of scalability**
- Problems
  - There is only one single controller in current deployments
    - Performance issues of controller in large scale network
  - Interconnections of remote OpenFlow Islands
    - Large establishment delay of new flow entry

Related Work to Solve Scalability Issues

- 1. Using optimization techniques
  - – Maestro uses the technique of parallelism
    - multi-threading to handle the flow requests from OpenFlow switches
    - batching sending when the controller sends control configurations back to the switches

A Z. Cai, A. L. Cox, E. Ng. Maestro: A System for Scalable OpenFlow Control. Technical Report of Rice University. Dec, 2010.

Related Work to Solve Scalability Issues

- 2. Devolving some control functions back to the switches
  - DevoFlow (Devolved OpenFlow) [HotNets10][Sigcomm11]
    - Decreasing the interactions between OpenFlow switches and controller
  - DIFANE (DIstributed Flow Architecture for Networked Enterprises)[Sigcomm10]
    - Distributing the rules across "authority switches"

[Sigcomm 11] A. R. Curtis, J. C. Mogul, J. Tourrilhes, et al. DevoFlow: Scaling Flow Management for High-Performance Networks. SIGCOMM 2011.
[Sigcomm 10]M. Yu, J. Rexford, M. J. Freedman, J. Wang. Scalable Flow-Based Networking with DIFANE. SIGCOMM 2010.

Related Work to Solve Scalability Issues

- 3. Designing a distributed control platform
  - HyperFlow [INM/WREN 10]
    - Distributed event-based control plane for OpenFlow,
    - logically centralized but physically distributed
  - Onix[OSDI10]
    - Distributed control platform for large-scale networks
    - providing more general APIs than previous systems and flexible distribution primitives

[INM/WREN 10] A. Tootoonchian, Y. Ganjali. HyperFlow: A Distributed Control Plane for OpenFlow Networks. INM/WREN 2010.
[OSDI10]T. Koponen, M. Casado, N. Gude, et al. Onix: A Distributed Control Platform for Large-scale Production Networks. OSDI 2010.

Our Ideas

- Scalable Distributed Control Platform
  - physically distributed
  - logically distributed
  - Each controller serves for its managed OpenFlow Cloud
  - Control Platform Overlay
  - Control Protocol

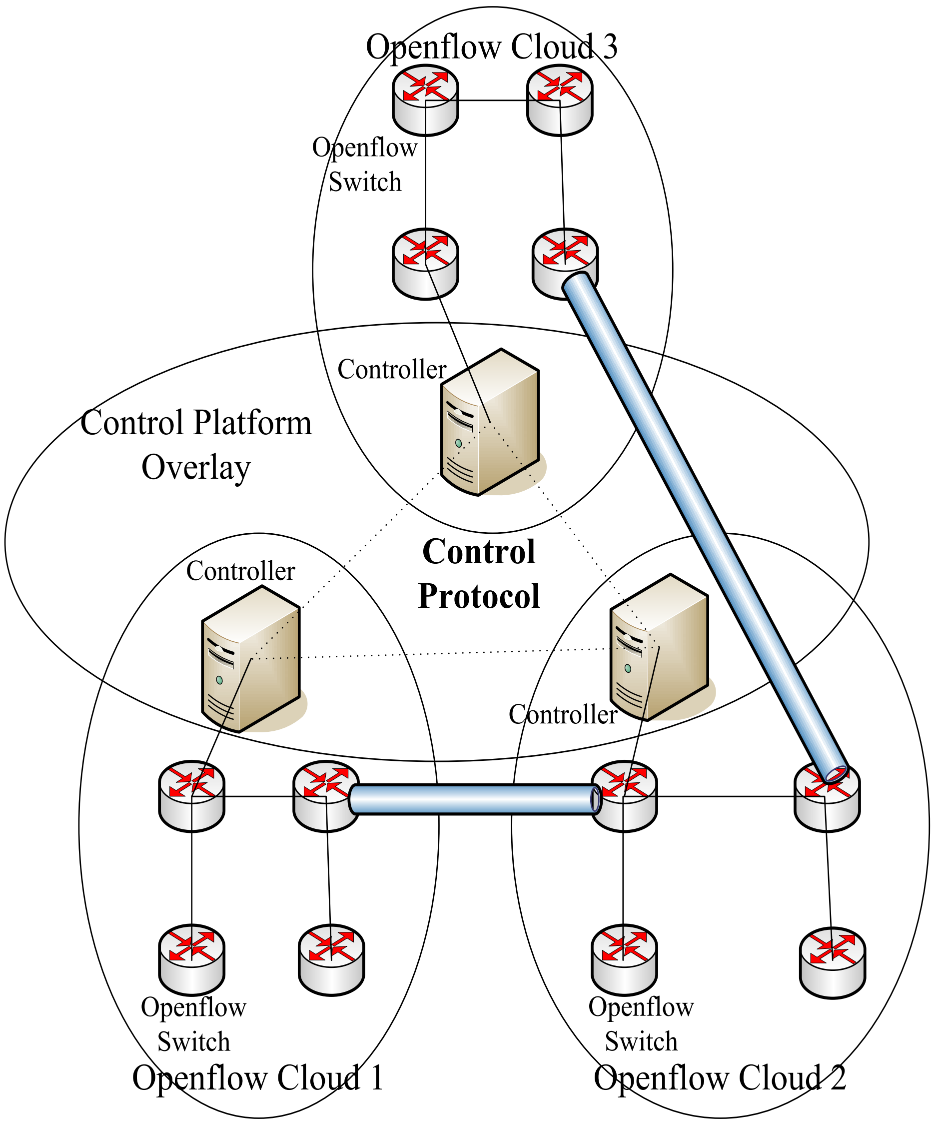

More Further…

- Inter-Domain Scenarios?

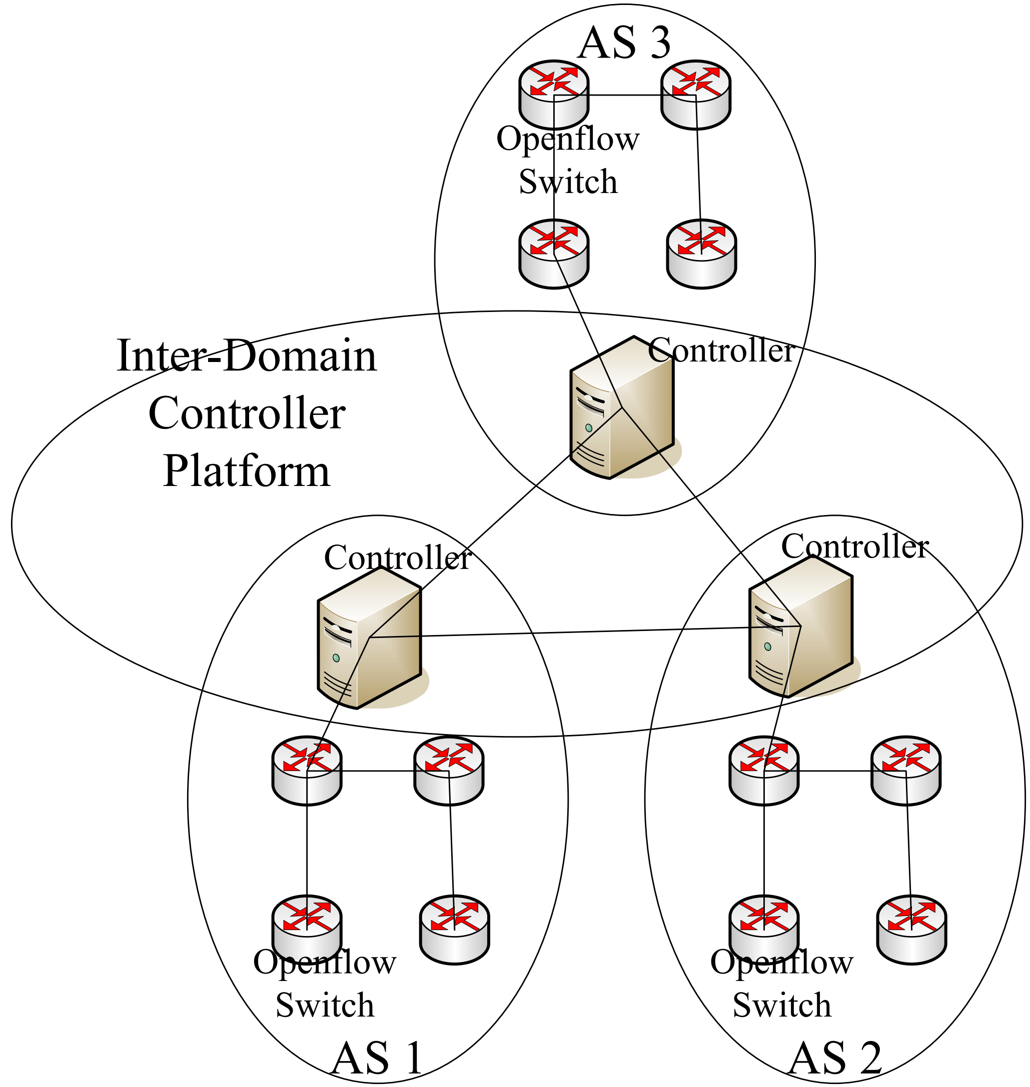

Tsinghua University

Summary

- Extensibility of OpenFlow
  – IPv6 Extensions

- Scalability of OpenFlow
  – Scalable Distributed Control Platform

清華大學 Tsinghua University

Thank you!

Q&A

http://netarchlab.tsinghua.edu.cn/~wzl

wzl@cernet.edu.cn

Tsinghua University