Future Internet Architecture

Jun Bi
Tsinghua University
junbi@tsinghua.edu.cn
Sep. 14, 2010

Outline

- Internet Architecture
- NSF FIA Projects
- Named Data Networking （NDN）
- Openflow
- Evolvable Internet Architecture (EIA)
- Brief Introduction to other new architectures
- David Clark’s Video
- Q&A

Internet Architecture

Network Architecture

- Network architecture is the design of a communications network. It is a framework for the specification of a network's physical components and their functional organization and configuration, its operational principles and procedures, as well as data formats used in its operation.
- http://en.wikipedia.org/wiki/Network_architecture

Clark88

- David Clark, Design Philosophy of the DARPA Internet Protocols, ACM SIGCOMM 1988
- Internet design goals and priority
- Different design goals and priority leads to different network architectures design

Internet Design Philosophy (Clark'88)

Basic Goal： Connect existing networks *

2nd Level Goal:

- Survivability *
- Support multiple types of services *
- Must accommodate a variety of networks *
- Allow distributed management
- Be cost effective
- Allow host attachment with a low level of effort
- Allow resource accountability

End-to-End Arguments (Saltzer'84)

- [J. Saltzer, et.al., *end-to-end arguments in system design,* ACM Trans. Comput. Syst. 2(4): 277-288, 1984]:
- If the application can do it, don't do it at a lower layer -- anyway the application knows the best what it needs
  - add functionality in lower layers iff it is (1) used and improves performances of a large number of applications, and (2) does not hurt other applications
- Internet is one of the successful example

Example: Reliable File Transfer

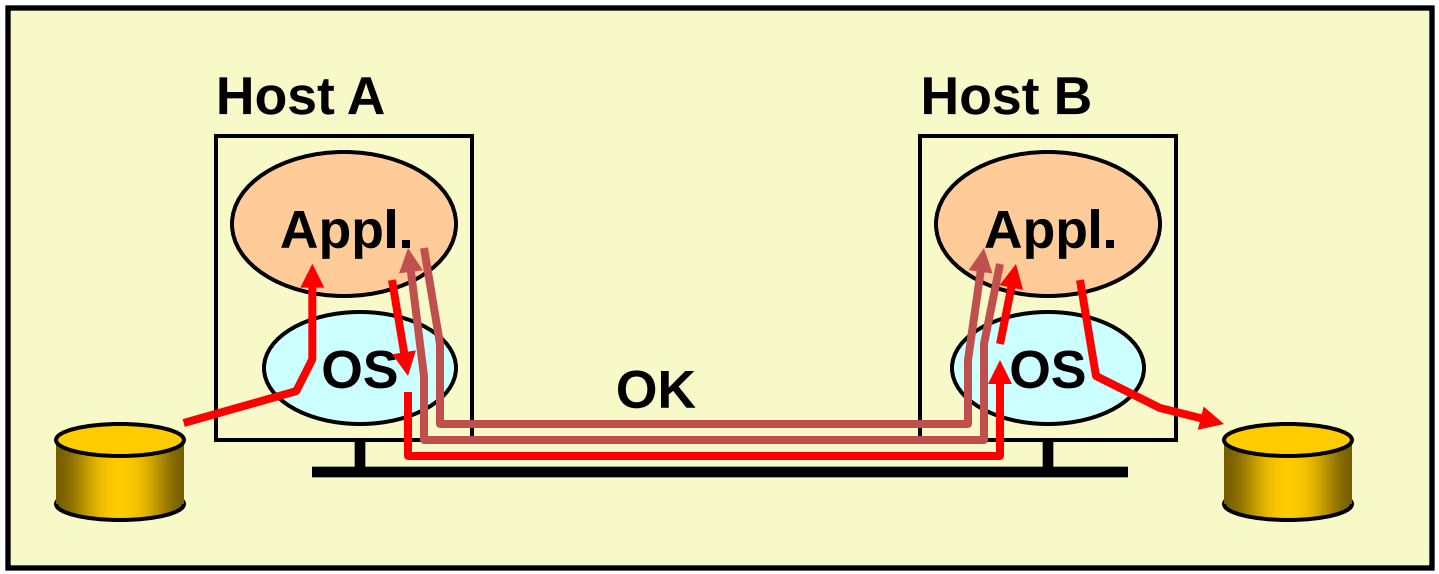

- Solution 1: make each step reliable, and then concatenate them
- Solution 2: each step unreliable: end-to-end check and retry

Internet Architecture (Hourglass model)

- IP over Everything
- Everything over IP
- The dummy network provides minimal functions while enabling the upper application on END SYSTEM evolution and competition. （end system matches PC age）

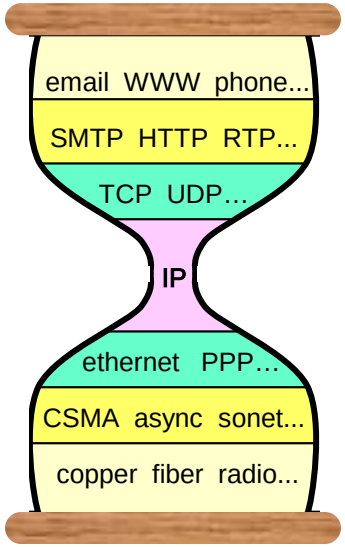

Applications

Technology

- Shortcoming： IP layer is toooo important to evolve（ ***"Narrow waist")***

Internet Architecture

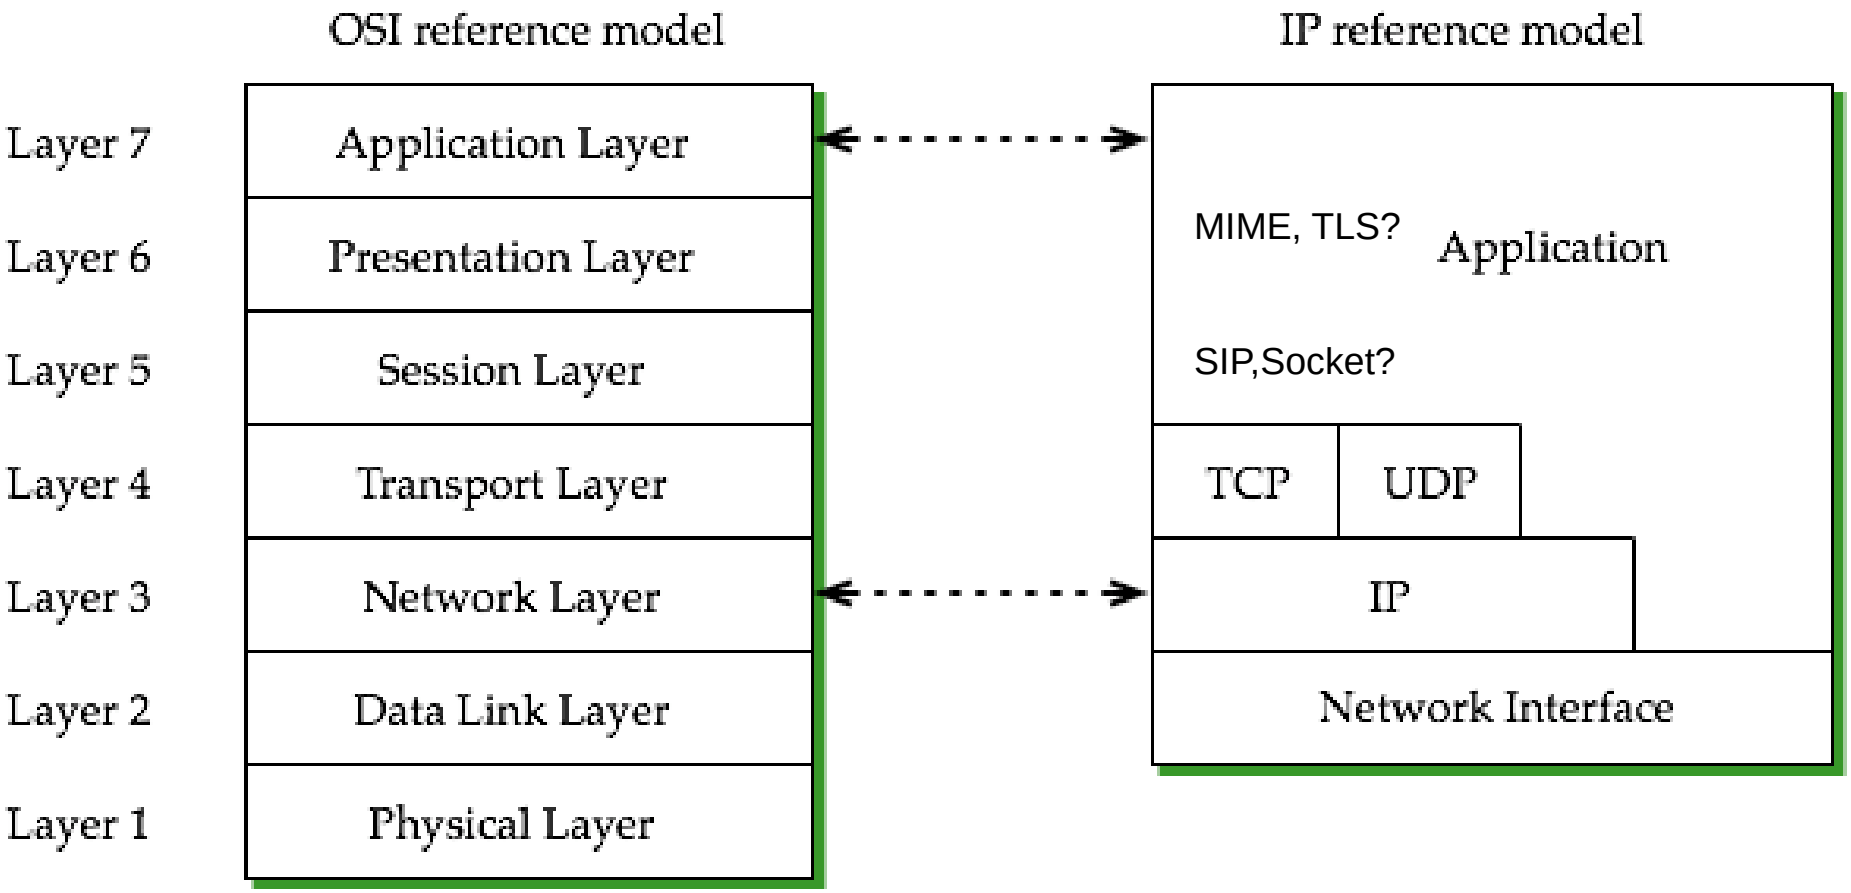

Major Challenges to Future Internet

- “Successful Disaster”
- Evolvable
- Scalability
- Mobility/Wireless
- Security/Trust
- Quality of Service
- “the other billions”
- The green Internet
- Network Management
- Economics

Future Internet Architecture Research

- Approaches：
  - Evolutionary
  - Clean Slate

- Projects
  - NewArch (-2003), FIND (2005-), FIA(2010)
  - EuroNF
  - Asia
  - …….

NSF FIA Projects

Basic Information

- **NSF Future Internet Architectures (FIA) :**
  - http://www.nsf.gov/pubs/2010/nsf10528/nsf10528.htm
- **Program Title:**
  - Future Internet Architectures (FIA)
- **Full Proposal Deadline :**
  - April 22, 2010
- **Result Announcement:**
  - August 27, 2010

INTRODUCTION

- Success of Internet:
- Crisis of Internet:
  - The continued success of the Internet, however, is increasingly threatened by the ever-mounting sophistication of security attacks and by the lack of performance reliability of Internet services.
  - As our reliance on a secure and highly dependable information technology infrastructure continues to increase, it is no longer clear that emerging and future needs of our society can be met by the current trajectory of incremental changes to the current Internet.

INTRODUCTION

- Objective of this program:
  - Our objective is to engage the research community in collaborative, long-range, transformative thinking - unfettered by the constraints of today's networks yet inspired by lessons learned and promising new research ideas - to design and experiment with new network architectures and networking concepts that take into consideration the larger social, economic and legal issues that arise from the interplay between the Internet and society.

PROGRAM DESCRIPTION

- Background：
  - Driven by technological innovations and the requirements of emerging and yet to be discovered applications, the Internet of the future is likely to be different from that of today. Responsive proposals should not focus on making the existing Internet better through incremental changes, but rather should focus on designing comprehensive architectures that can meet the challenges and opportunities of the 21st century.

PROGRAM DESCRIPTION

- Requirements to Proposing teams:
  - To capitalize on lessons learned from the past;
  - To explore emerging concepts and promising research outcomes recently reported by the research community;
  - To engage in transformative research that has the potential to enable the creation of comprehensive architectures that reach beyond current core networking components, mechanisms and application requirements.

PROGRAM DESCRIPTION

- Requirements to Proposals submitted:
  - Other design requirements such as, but not limited to, scalability, openness, ubquitous access, innovation-enabling, manageability, evolvability and economic viability, may also be considered.
  - To prototype and evaluate the proposed architectures; this may require the construction of new artifacts or the use of research infrastructure like GENI (Global Environment for Network Innovations) or the NCR (National Cyber Range).

NSF Future Internet Architecture Awards (**August 27, 2010**)

- Four new projects, each worth up to $8 million over three years, as part of the Future Internet Architecture (FIA) program.
  - **Named Data Networking**
  - **MobilityFirst**
  - **NEBULA**
  - **eXpressive Internet Architecture**

Named Data Networking

- **Principal Investigator**:
  - Lixia Zhang, UCLA
- **Collaborating Institutions**:
  - Colorado State University,
  - PARC, University of Arizona,
  - University of Illinois/Urbana-Champaign,
  - UC Irvine,
  - University of Memphis,
  - UC San Diego,
  - Washington University,
  - Yale University

MobilityFirst

- **Principal Investigator**:
  - Dipankar Raychaudhuri, Rutgers University/New Brunswick
- **Collaborating Institutions**:
  - Duke University,
  - Massachusetts Institute of Technology,
  - University of Massachusetts/Amherst,
  - University of Massachusetts/Lowell,
  - University of Michigan,
  - University of Nebraska/Lincoln,
  - University of North Carolina/Chapel Hill

MobilityFirst

- Current Internet, originally designed to support communications between fixed end-points
- MobilityFirst project proposes an architecture centered on mobility as the norm, rather than the exception
- Generalized Delay-tolerant networking (GDTN) to provide robustness even in presence of link/network disconnections
- Tradeoffs between mobility and scalability
- Tradeoffs on opportunistic use of network resources to achieve effective communications among mobile endpoints.

NEBULA

- **Principal Investigator**:
  - Jonathan Smith, University of Pennsylvania
- **Collaborating Institutions**:
  - Cornell University,
  - Massachusetts Institute of Technology,
  - Princeton University,
  - Purdue University,
  - Stanford University,
  - Stevens Institute of Technology,
  - University of California/Berkley,
  - University of Delaware,
  - University of Illinois/Urbana-Champaign,
  - University of Texas,
  - University of Washington

NEBULA

- “Cloud” creating opportunities for global-scale, network-centric computing infrastructure
  - Fast resource provisioning
  - Fast utility pricing
  - Fast consistent
  - easy management
- NEBULA is an architecture in which cloud computing data centers are
  - primary repositories of data
  - primary locus of computation
- The project focuses
  - developing new trustworthy data, control
  - core networking approaches to support the emerging cloud computing model of always-available network services.
- This project addresses the technical challenges in creating a cloud-computing-centric architecture.

eXpressive Internet Architecture

- **Principal Investigator**:
  - Peter Steenkiste, Carnegie Mellon University
- **Collaborating Institutions**:
  - Boston University,
  - University of Wisconsin/Madison

eXpressive Internet Architecture

- XIA addresses following issues:
  - the growing diversity of network use models,
  - the need for trustworthy communication,
  - the growing set of stakeholders who coordinate their activities to provide Internet services.
- XIA solves these problems by:
  - offers inherent support for communication between current communicating principals—including: hosts, content, and services, accommodating unknown future entities.
  - defining a narrow waist For each type of principal
  - enables flexible context-dependent mechanisms for establishing trust between the communicating principals
  - By bridging the gap between human and intrinsically secure identifiers
- This project evaluates and refines the interface between the network and users, and studies
- This project analyzes the relationship between technical design decisions, and economic incentives and public policy.

Named Data Networking

What's the problem

- Morgan Stanley View:
  - Video streaming is the main contributor for global IP traffic growth in the following years, accouning for 50% of the total traffic.
- A hot YouTube video downloaded 1,000,000 times from the same servers
  - YouTube request is looking for the data,
  - but the network only knows how to find that data from specific server
- Need new protocols for moving data around for every new application□
  - today's data dissemination is application specific
- Tolerating delay/disruption in data delivery
  - Especially for mobile networking

How Naming Data Solves the Problem

Just thinking about routing on data names

- YouTube request reaches a nearby copy of the Data
  - Rather than going to the specific server
- If all requests directed to data nearby:
  - Network natually provides scalable data dissemination
  - Disk technology made massive caching feasible

The Issues of Today's Internet

- **Availability**: Fast, reliable content access requires awkward, pre-planned, application-specific mechanisms like CDNs and P2P networks, and/or imposes excessive bandwidth costs.
- **Security**: Trust in content is easily misplaced, relying on untrustworthy location and connection information.
- **Location-dependence**: Mapping content to host locations complicates configuration as well as implementation of network services.

Solution to The Issues

- To replace *where* with *what*.
- To focus *named data* rather than *named hosts*.

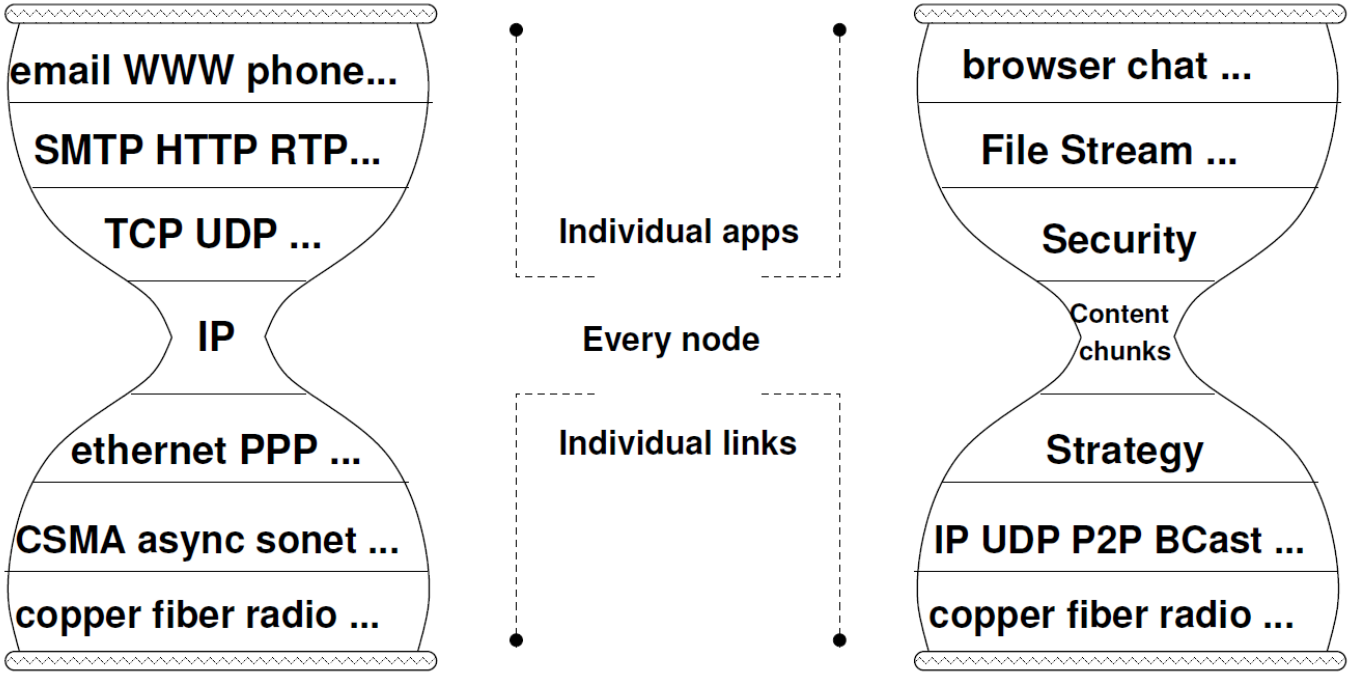

NDN moves the universal component of the network stack from IP to chunks of named content

Advantages

- Strategy
  - NDN can take maximum advantage of multiple simultaneous connectivities.
  - The strategy makes the fine-grained, dynamic optimization choices needed to best exploit multiple connectivities under changing conditions.
- Security
  - NDN secures content itself (Section 5), rather than the connections over which it travels, thereby avoiding many of the host-based vulnerabilities that plague IP networking.

NDN packet types

- Interest

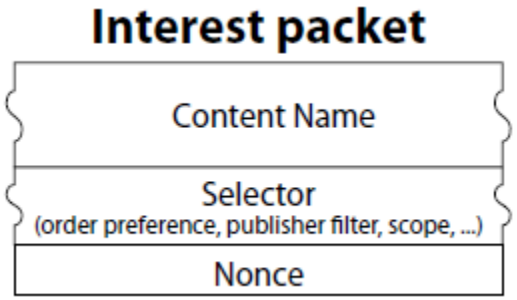

- Data

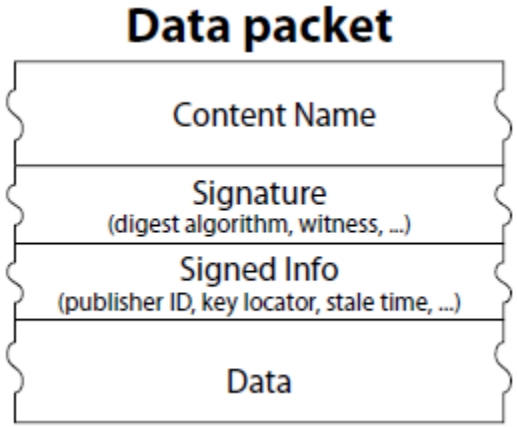

1. A consumer asks for content by broadcasting its interest over all available connectivity.
2. Any node hearing the interest and having data that satisfies it can respond with a Data packet.
3. Data is transmitted only in response to an Interest and consumes that Interest.

NDN forwarding engine model

A packet arrives on a face, a longest-match look-up is done on its name, and then an action is performed based on the result of that lookup

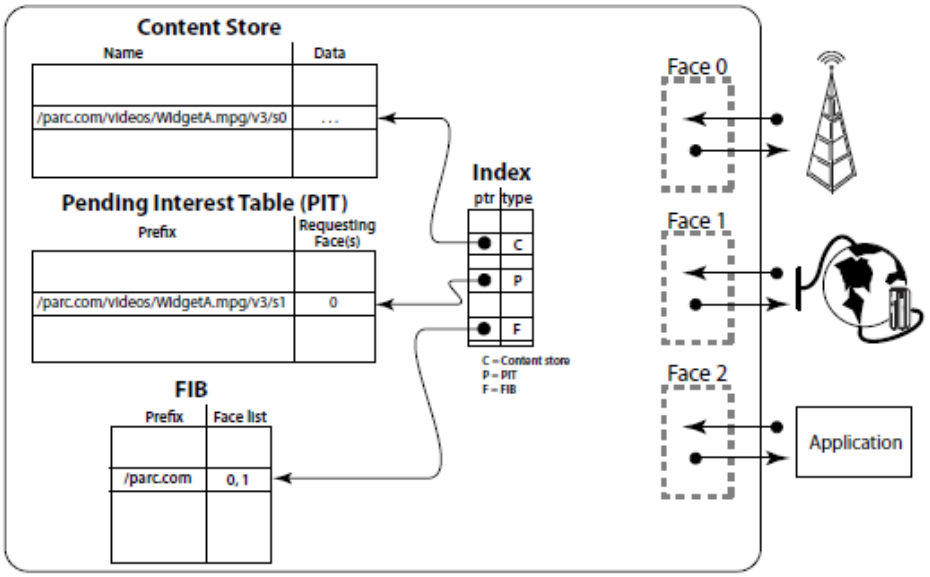

NDN forwarding engine model

FIB

- Be similar with IP FIB
- Forward Interest packets toward potential source(s) of matching Data
- A list of outgoing faces

- **NDN route Lookup is more efficient than ip, because NDN names used for lookup is structural.**

NDN forwarding engine model (Con'd)

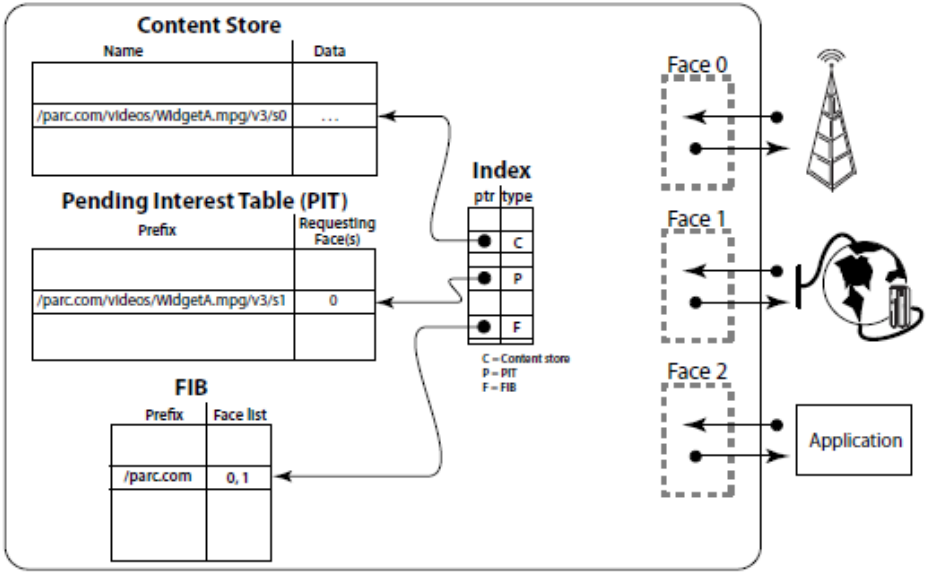

NDN forwarding engine model

Content Store

- Like the buffer memory of an IP router
- Remember arriving Data packets as long as possible

- **Because NDN is Content-based, it can deal with caching contents better.**

NDN forwarding engine model(Con'd)

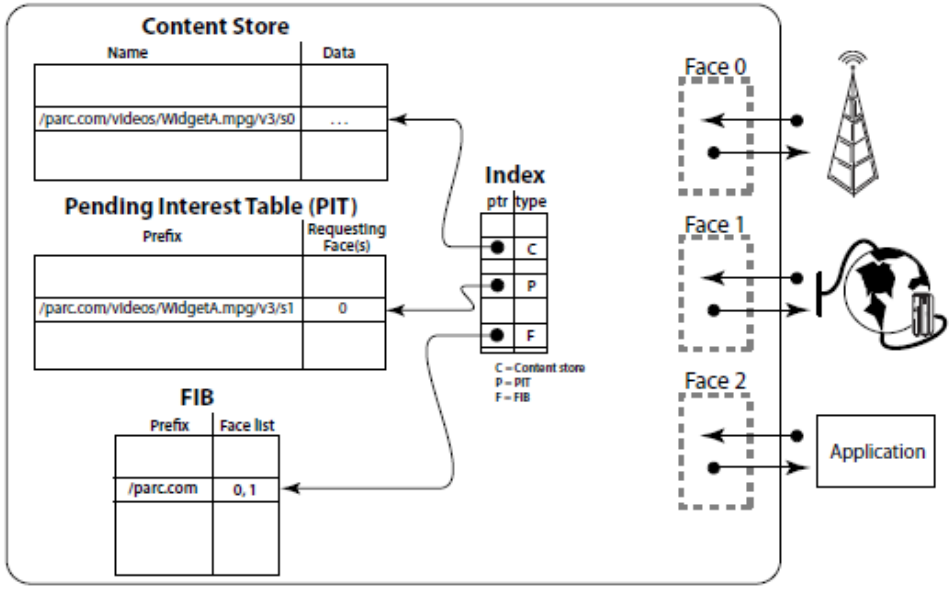

NDN forwarding engine model

PIT

•keeps track of Interests forwarded upstream toward content source(s) so that returned Data can be sent downstream to its requester(s)

When a interest arrives

- There is already a Data packet in the ContentStore
  - it will be sent out the face the Interest arrived on
  - the Interest will be discarded
- Otherwise, if there is an exact-match PIT entry
  - the Interest's arrival face will be added to the PIT entry's RequestingFaces list.
  - the Interest will be discarded.
- Otherwise, if there is a matching FIB entry
  - the Interest needs to be sent upstream towards the data.
  - The arrival face is removed from the face list of the FIB entry
  - if the resulting list is not empty, the Interest is sent out all the faces that remain
  - a new PIT entry is created from the Interest and its arrival face.
- If there is no match for the Interest
  - it is discarded

Sequencing

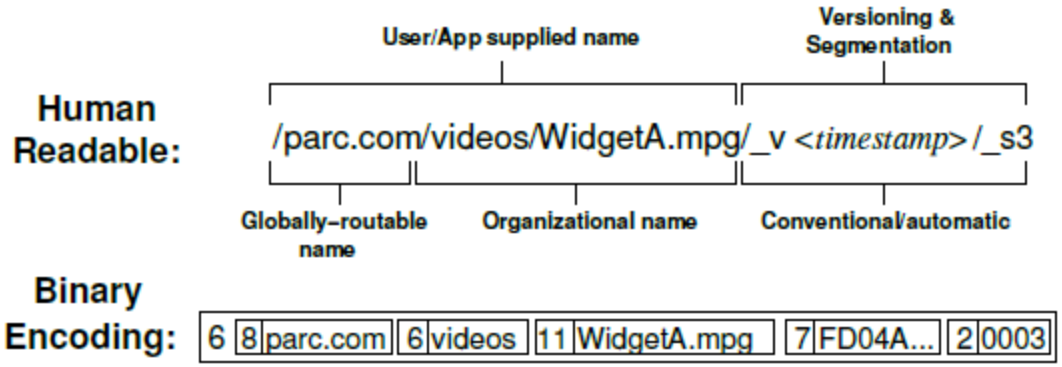

Example Data name

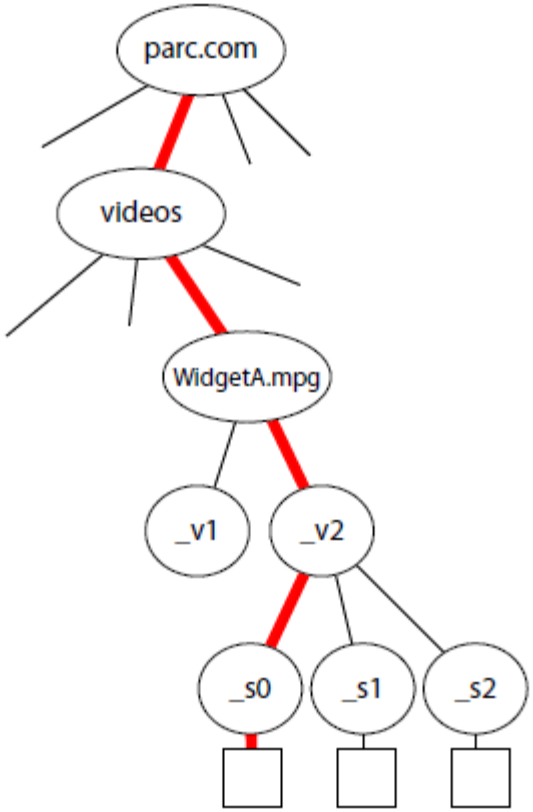

Name tree traversal

Rich Connectivity, Mobility and Strategy

- **Rich Connectivity, Mobility**

  NDN talks about data, not to nodes, so it does not need to obtain or bind a layer 3 identity (IP address) to a layer 2 identity such as a MAC address
- **Strategy**
  - sendToAll,
  - sendToBest,
  - markAsBest,
  - triggers(interestSatisfied, interestTimedOut, faceDown)

OpenFlow

Nick McKeown
Stanford University

Why OpenFlow

RESEARCH

Many great ideas

- *Mobility Management*
- *Network Security*
- *Energy Management*
- *New routing*
- *Network Measurement*
- *Things we haven't thought of…*

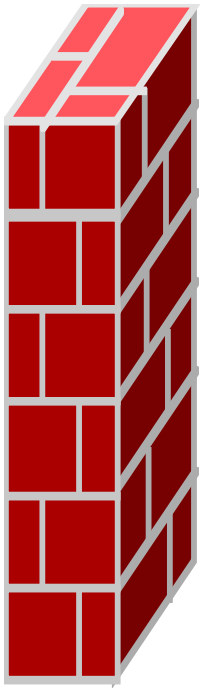

PRODUCTION

none transferred to Production networks

NEEDED: Realistic ways to try out ideas at scale

Glacial process of innovation made worse by captive standards process

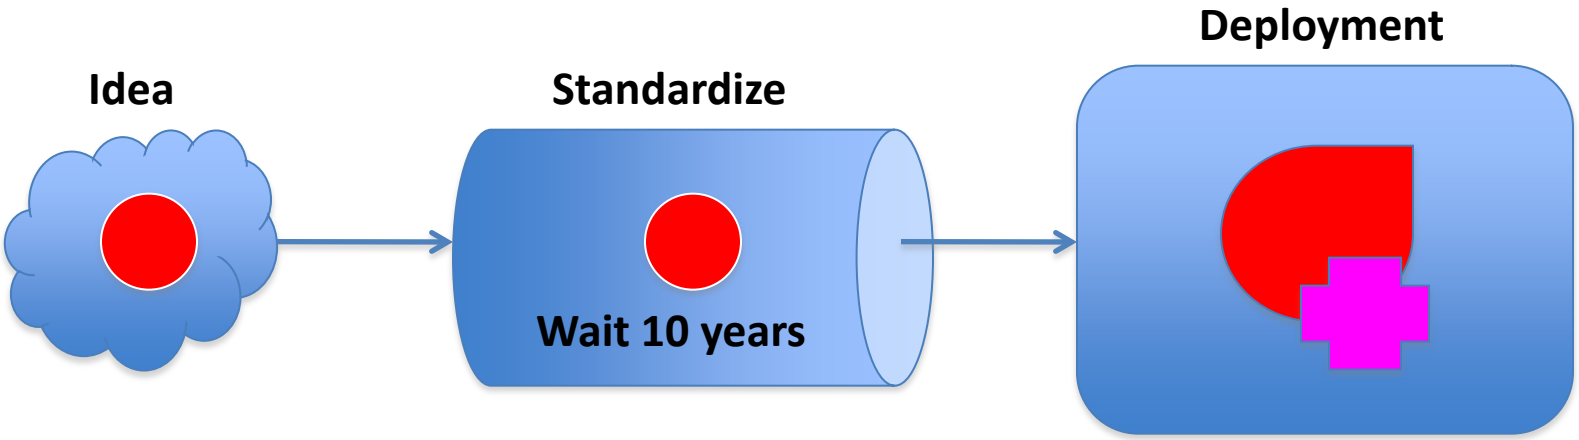

Restructuring Networks

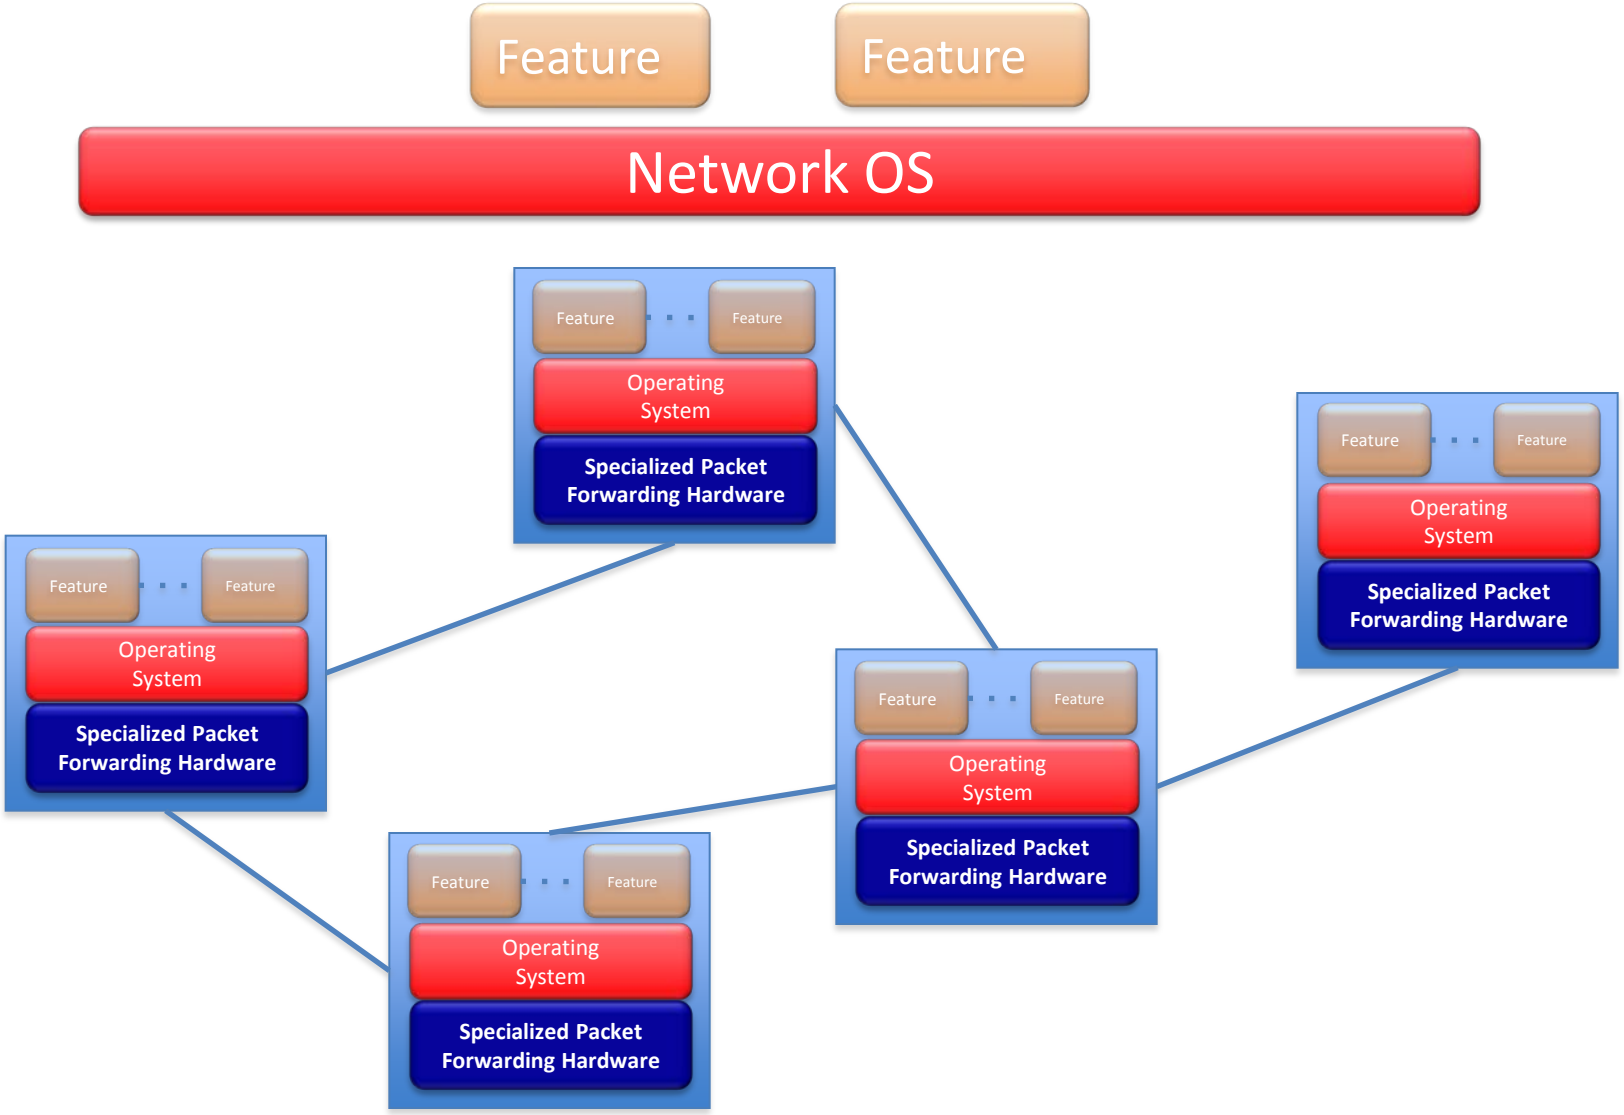

The “Software-defined Network”

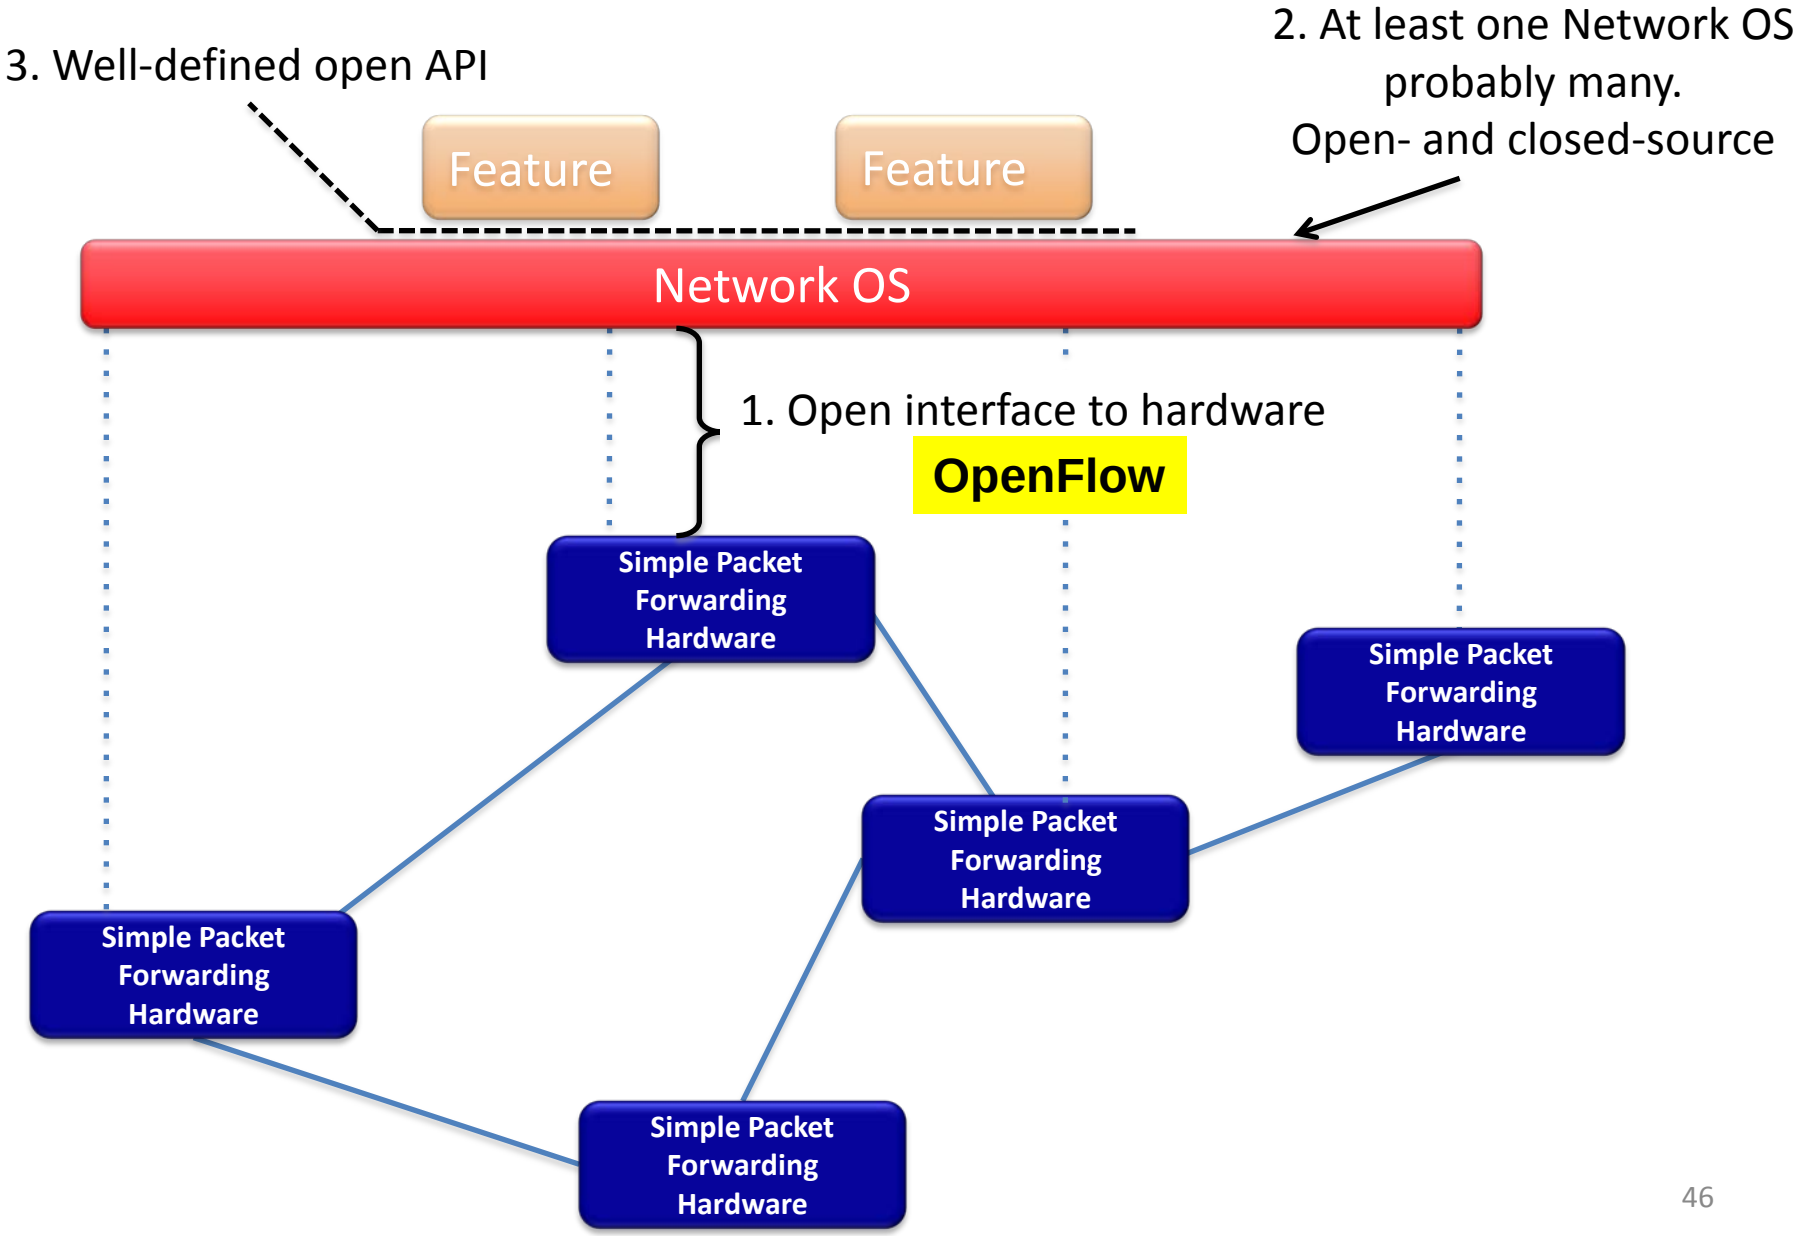

What's OpenFlow

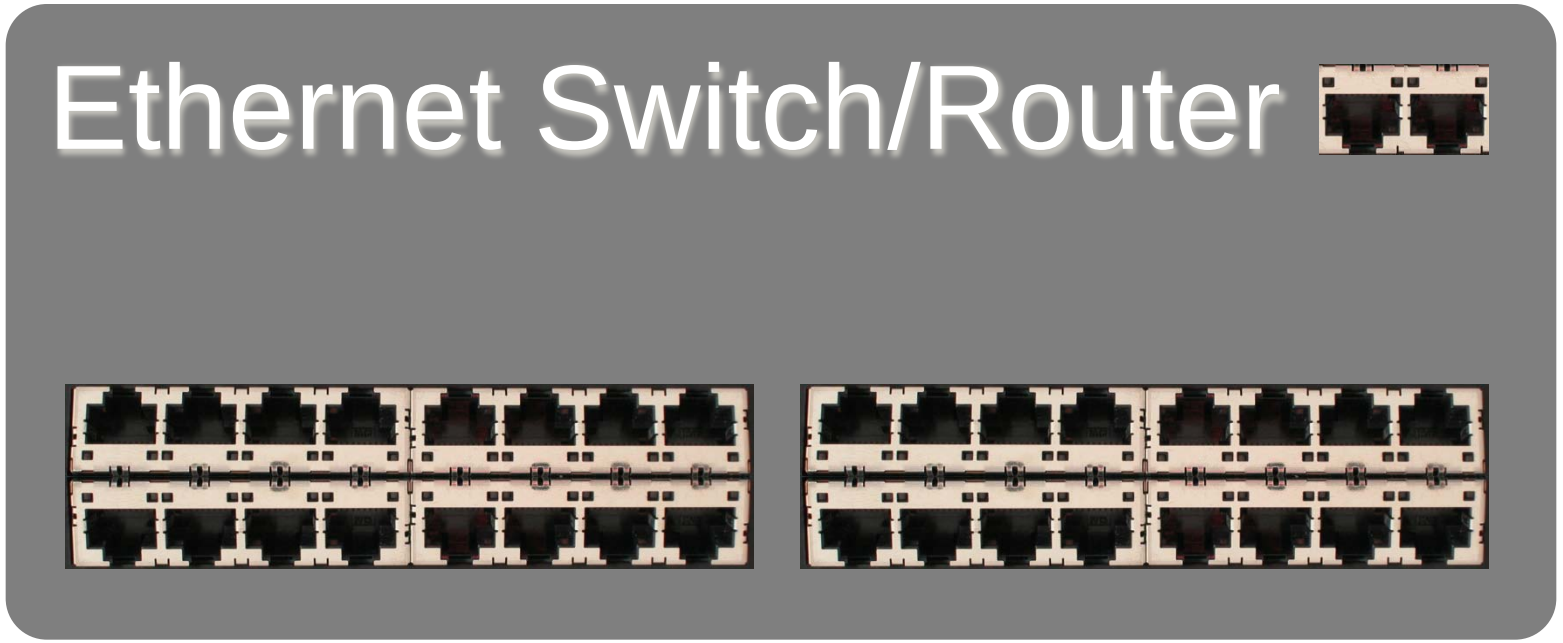

What's OpenFlow

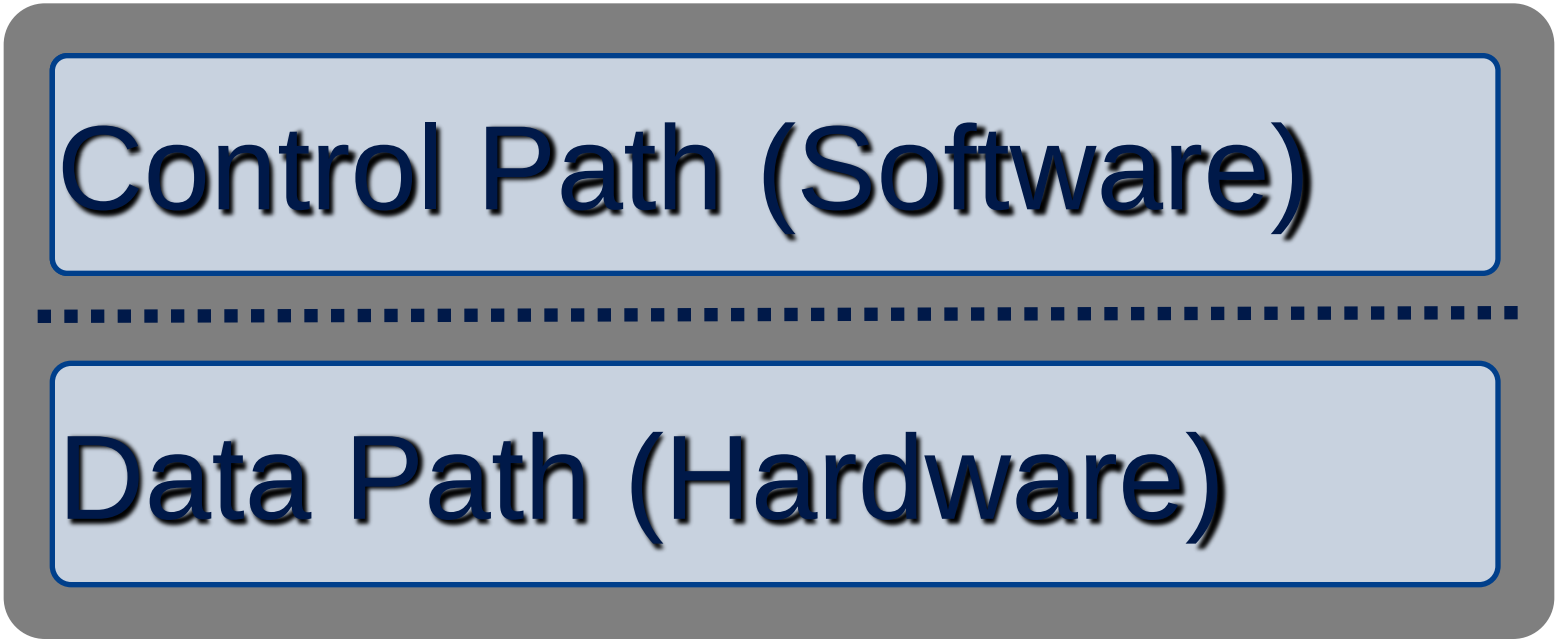

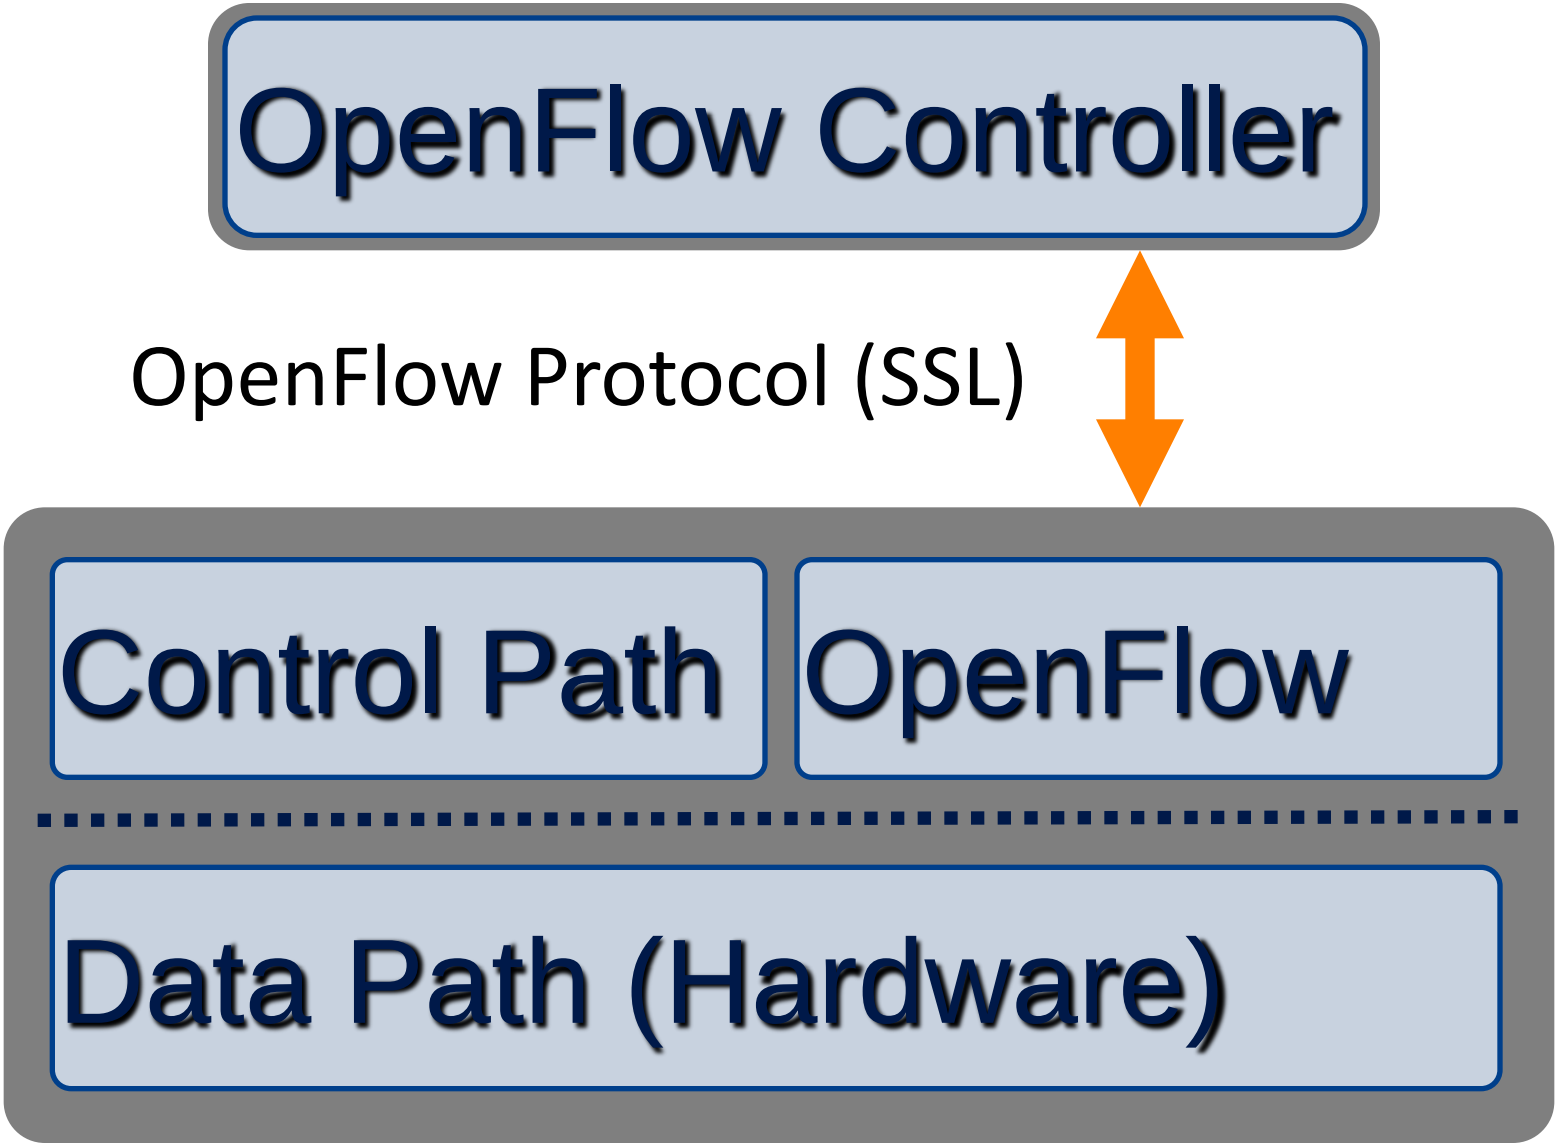
OpenFlow Controller
OpenFlow Protocol (SSL)
Control Path
OpenFlow
Data Path (Hardware)

De-coupling Control & Data Planes

Control

Routing

Signaling

Data

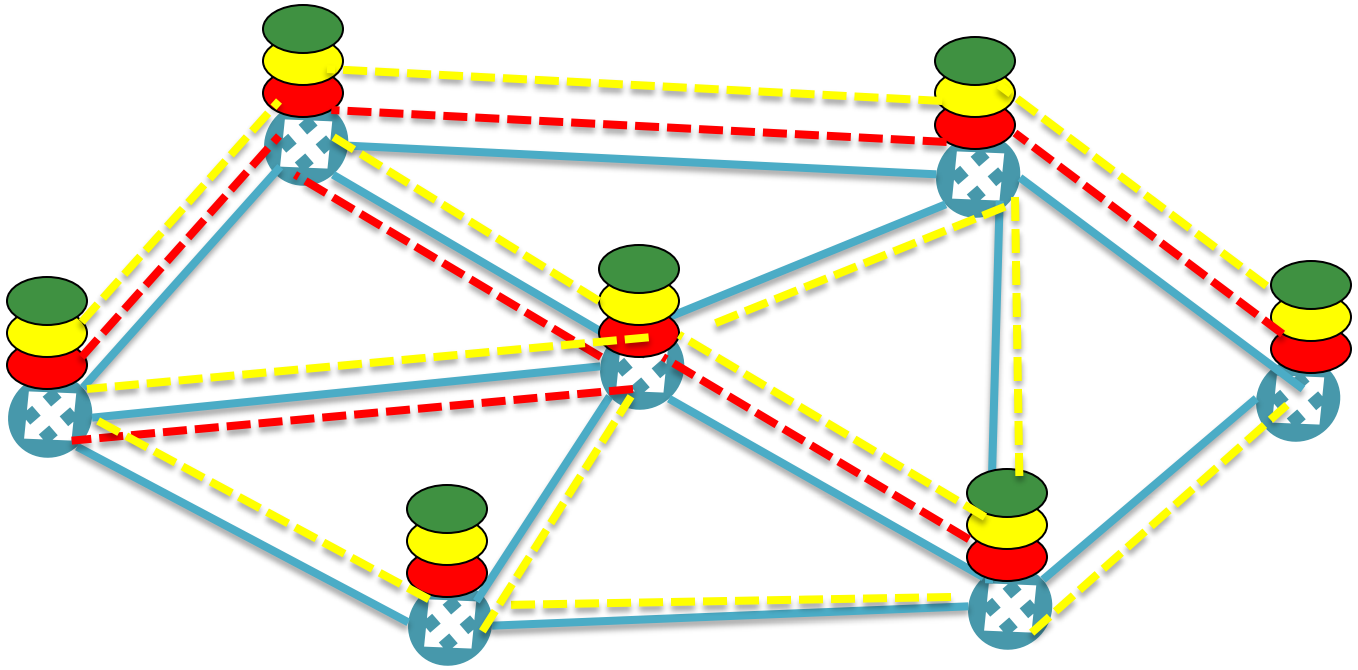

De-coupling Control & Data Planes

Controller

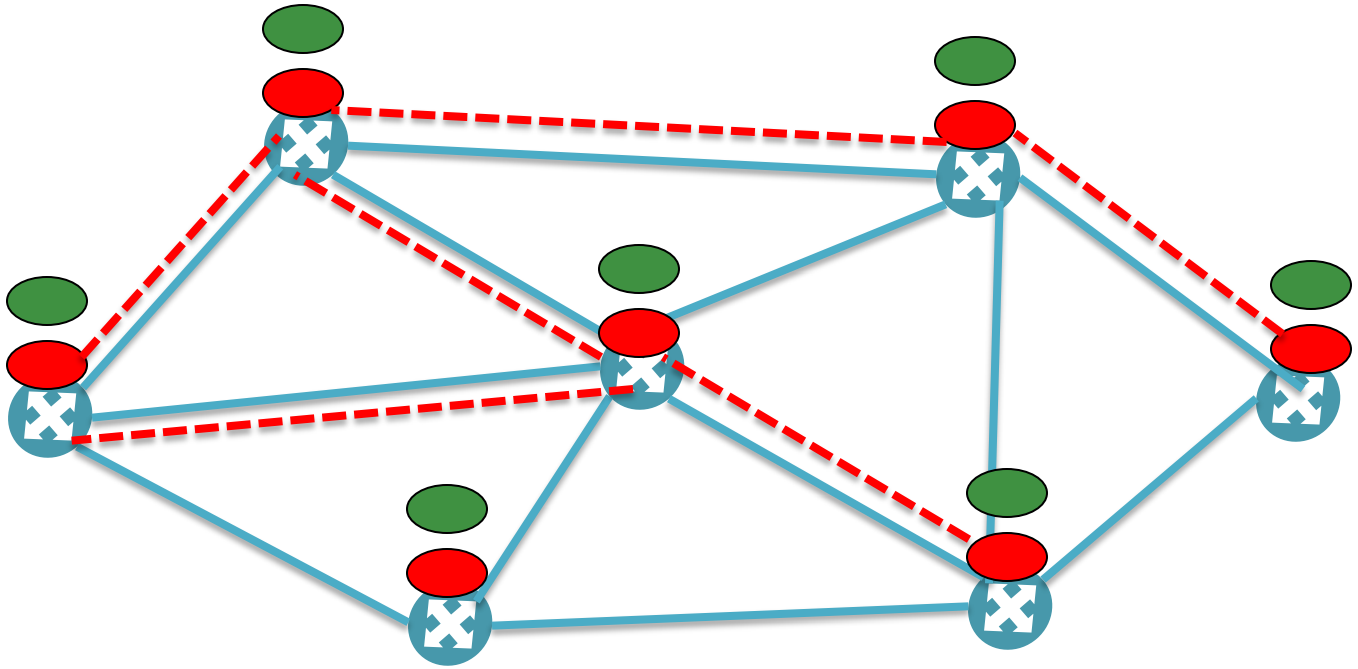

Signaling

Data

De-coupling Control & Data Planes

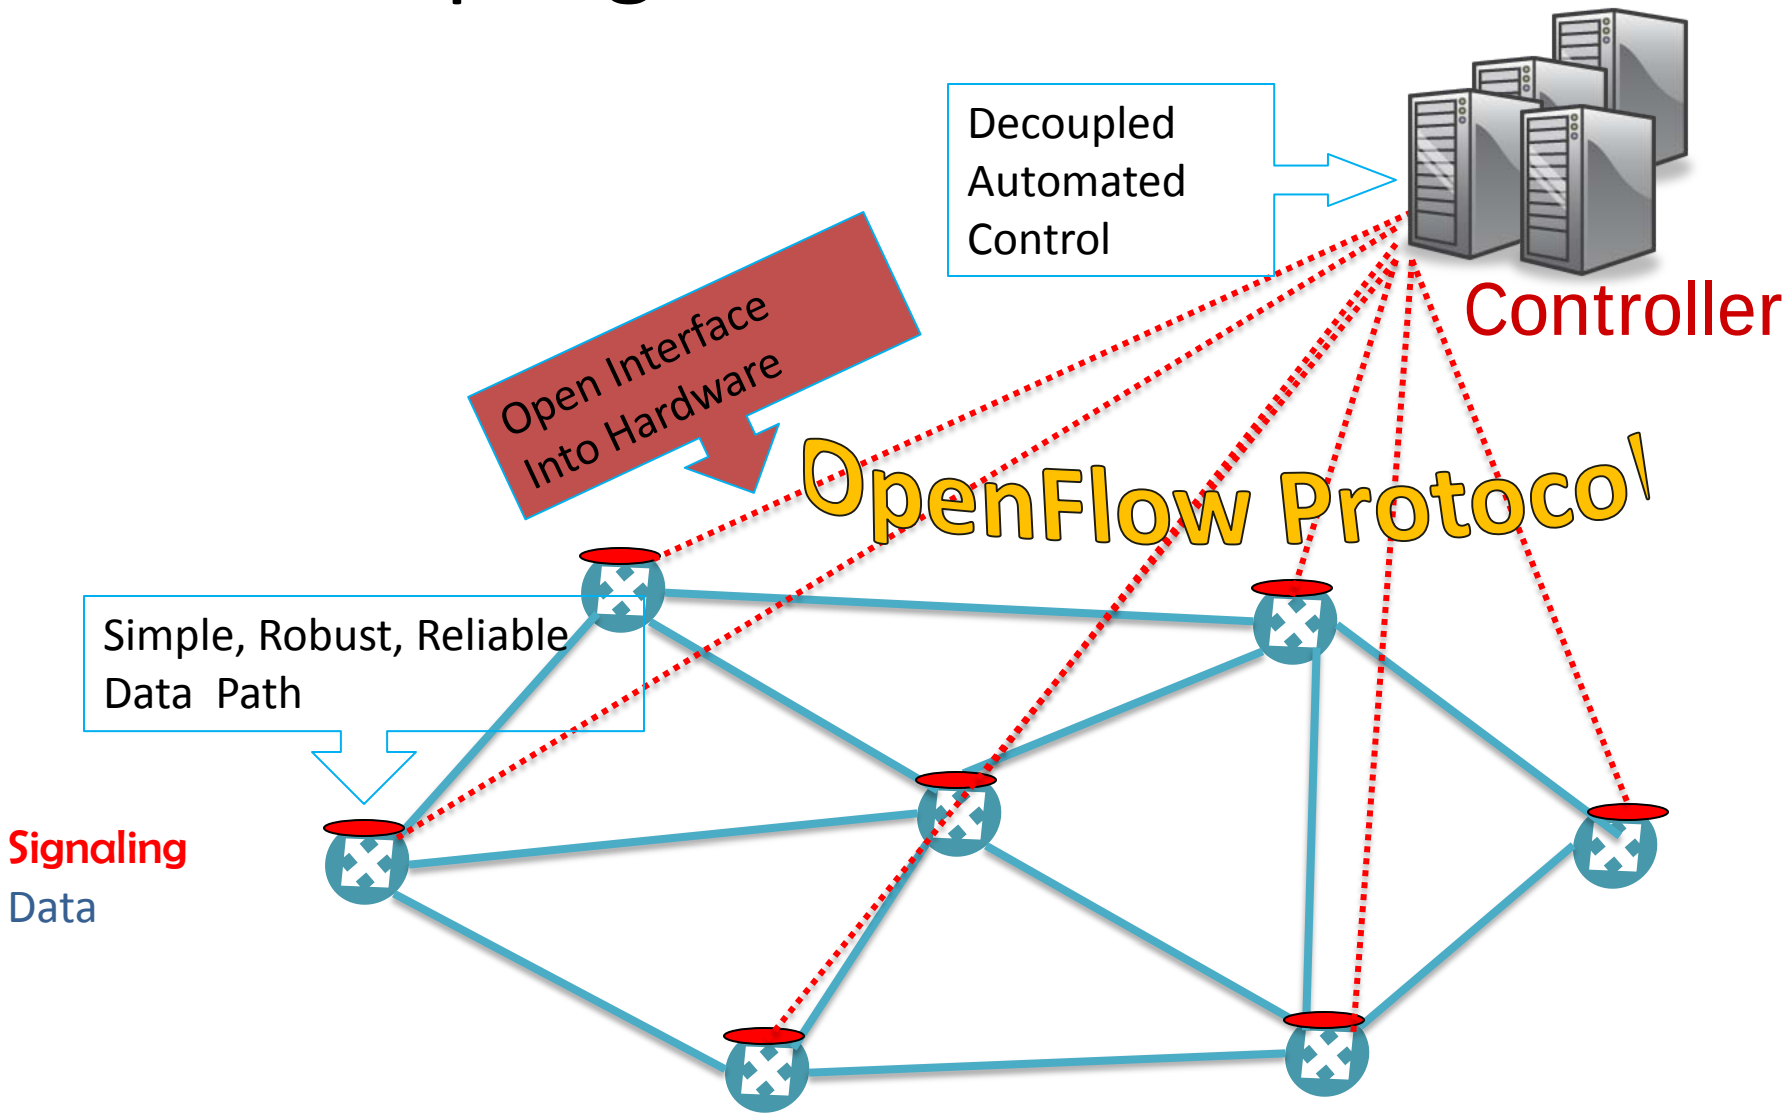

The Flow Abstraction

Exploit the flow table in switches, routers, and chipsets

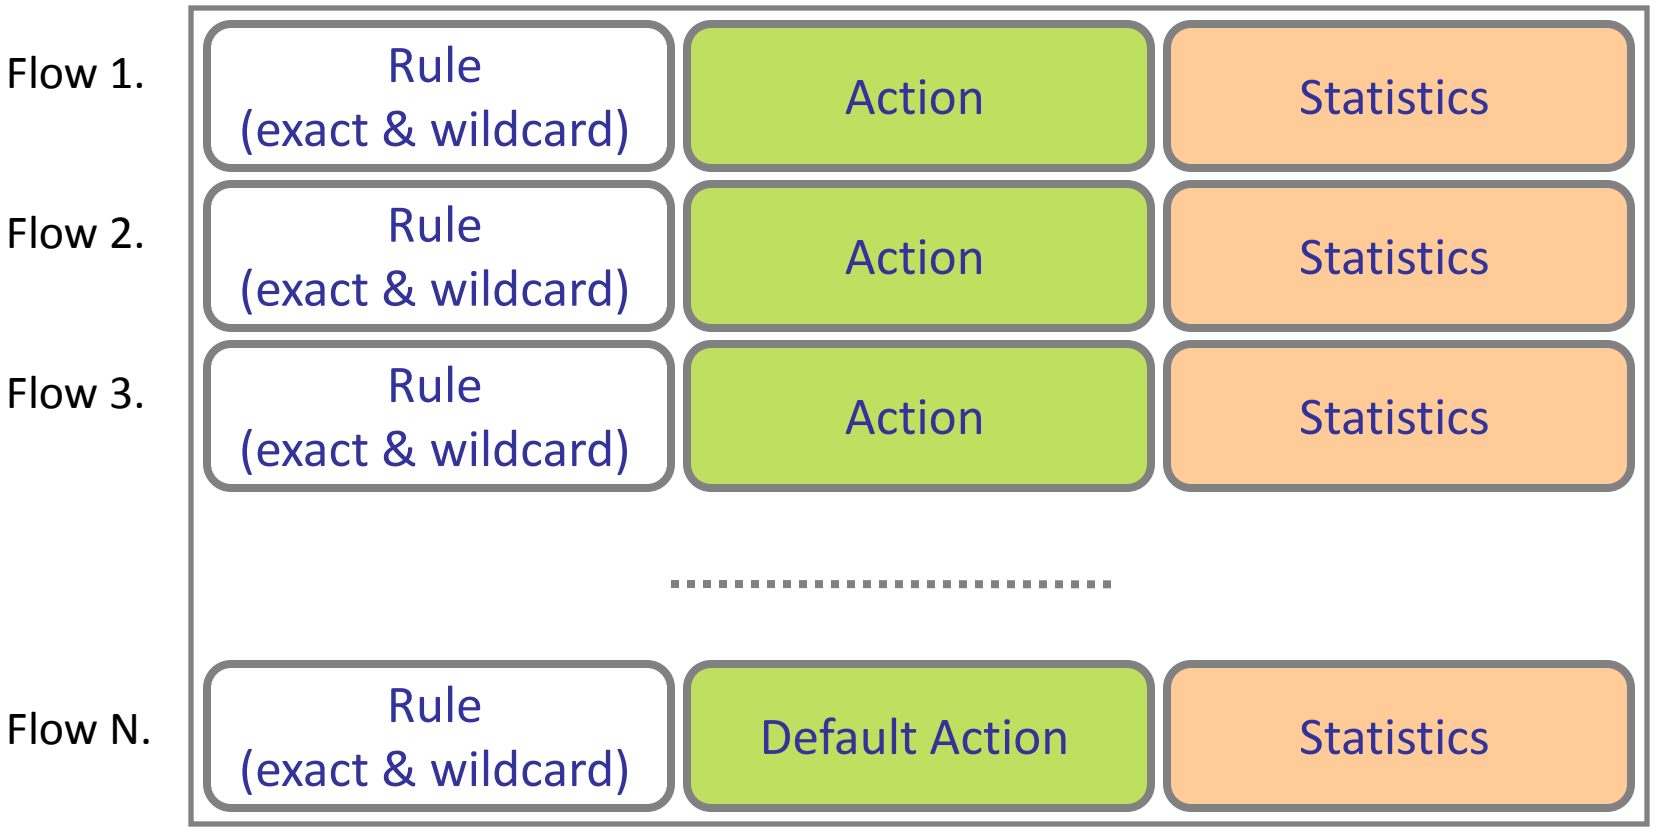

OpenFlow Switching

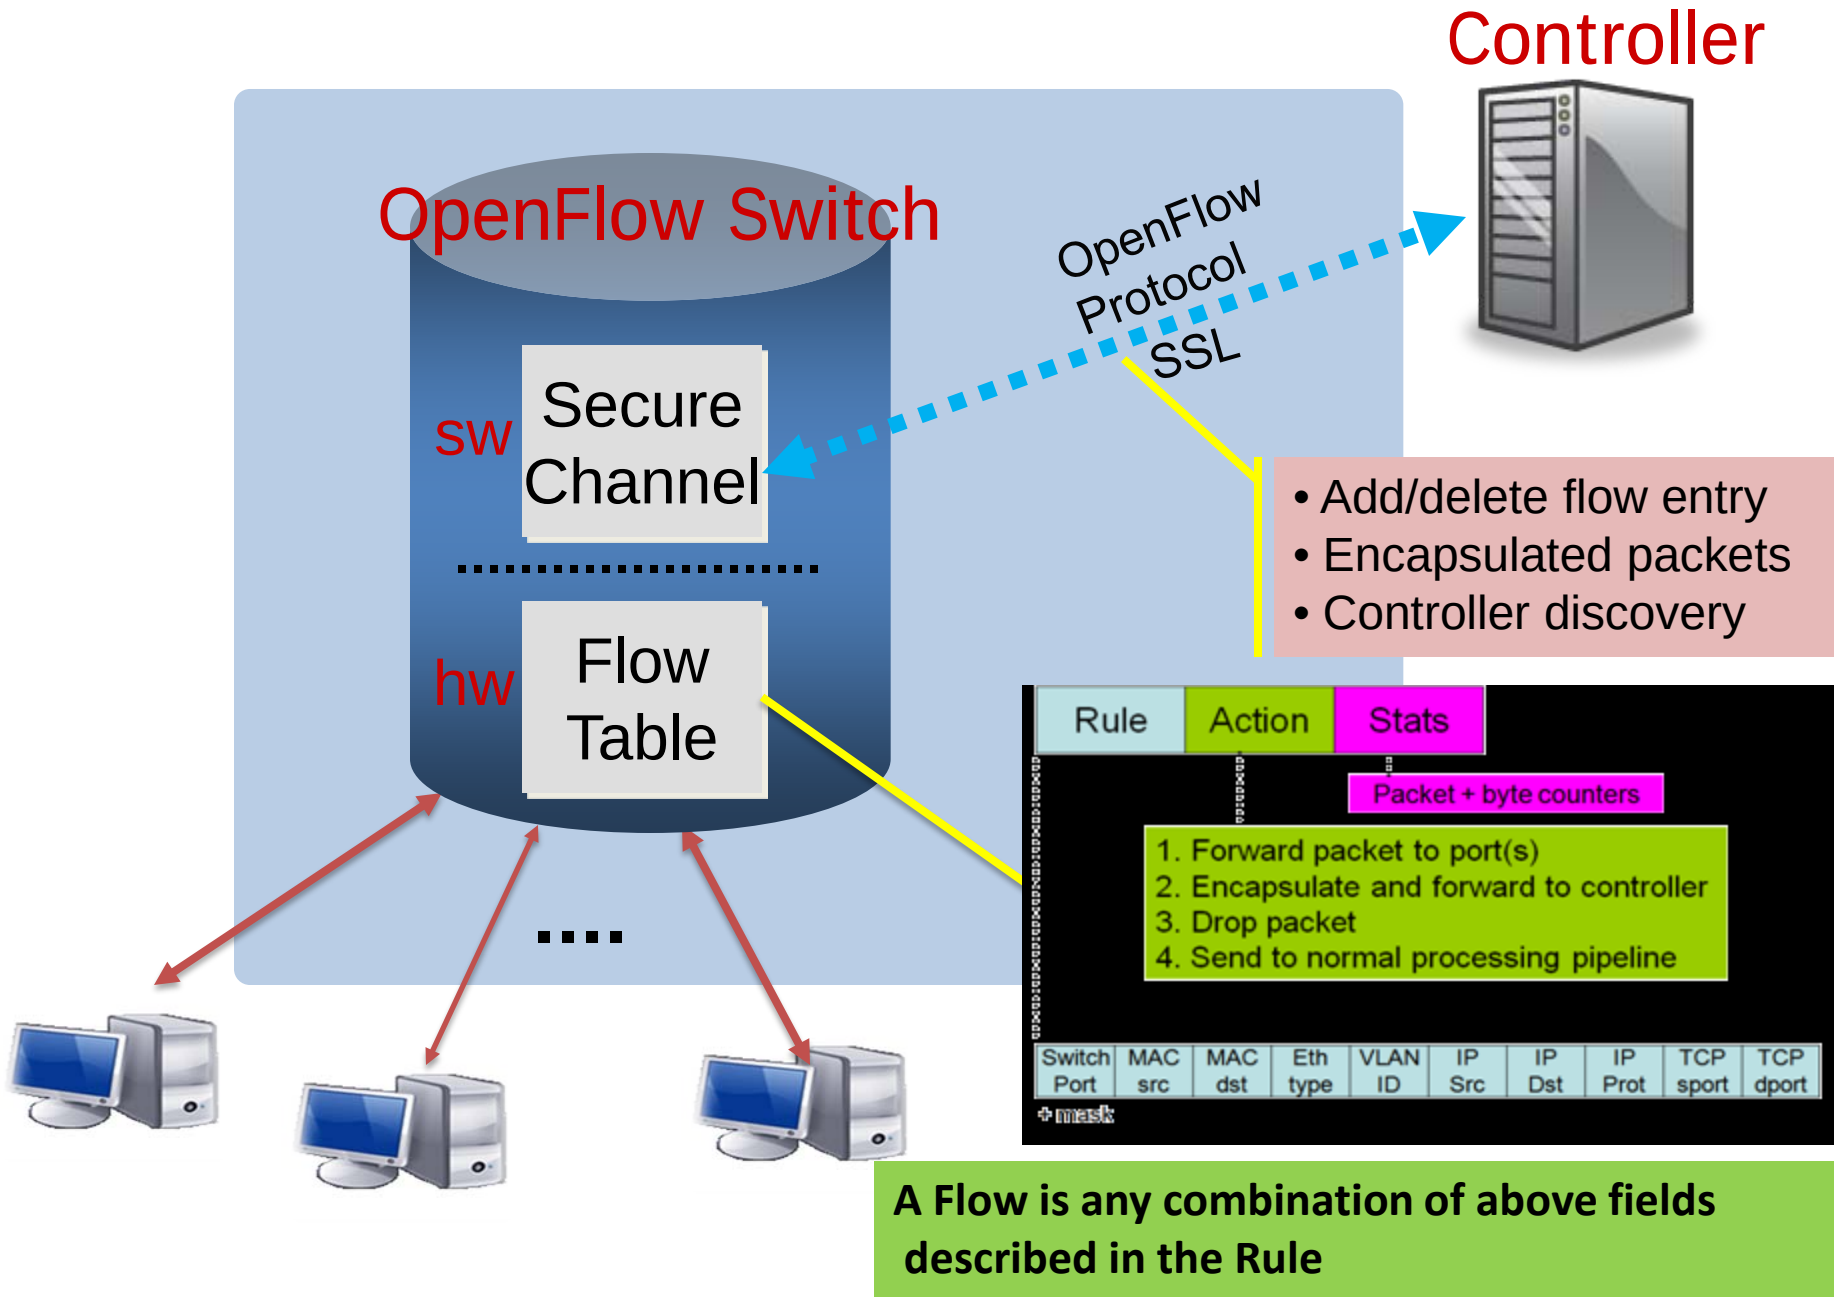

Flow Example

A Flow is the fundamental
unit of manipulation within a switch

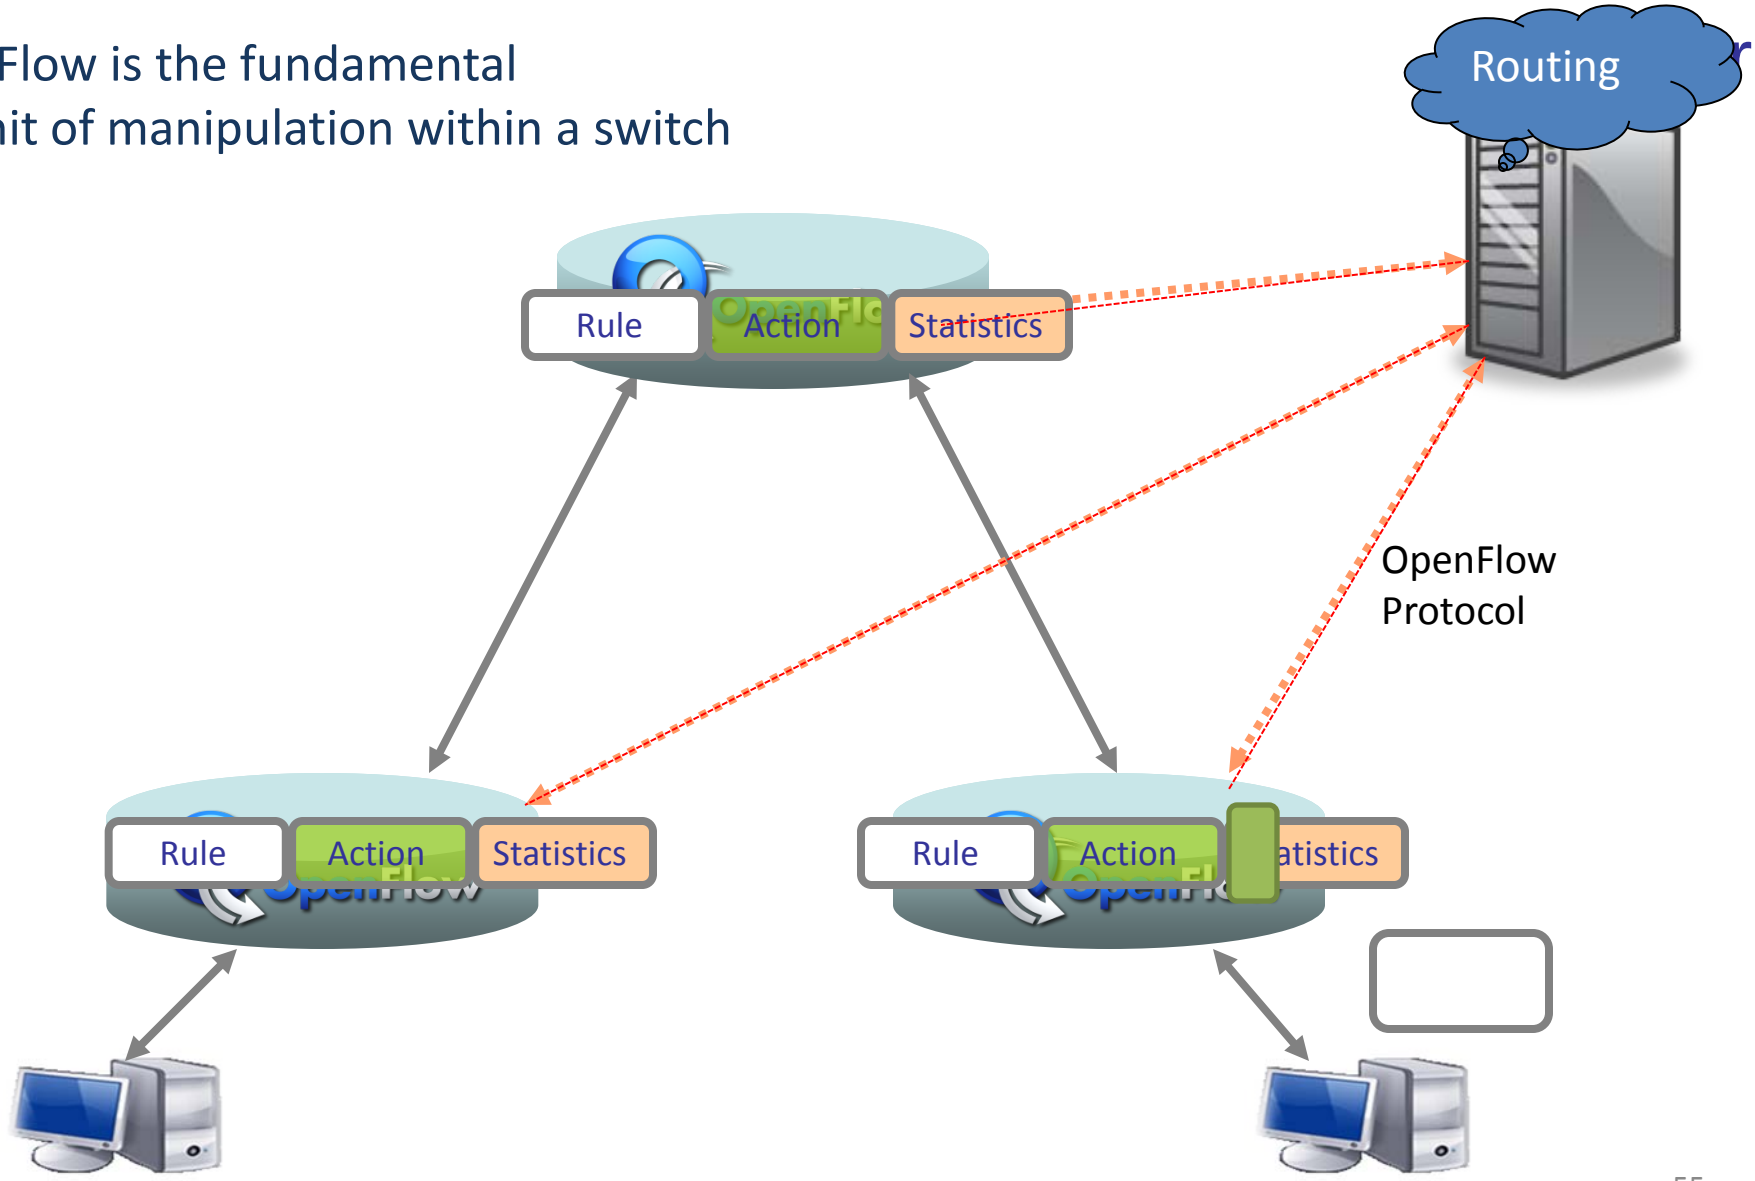

What is happening now?

OpenFlow

- Maturing and being added to existing hardware
- New hardware will emerge

Controllers

- Research: NOX http://noxrepo.org
- Commercial controllers under development
- Production deployments scheduled late 2010

Software Defined Networking

- Companies signing up to support trend

Companies involved

Network operators

- Google, Facebook, Rackspace
- DT, BT, Level-3, DoCoMo

Switch vendors

- HP, Dell, NEC, Juniper, Arista, Quanta, Extreme

Chip vendors

- Broadcom, Marvell, Fulcrum

EIA: An Evolvable Internet Architecture

Jun Bi

Tsinghua University

INTRODUCTION

- New protocols related to core network or network forwarding equipments are hard to implement and deploy, which hinders the evolution of the core network technology.

- EIA is a programmable platform, abstracts and open up a set of public package level instructions or interfaces for developers to operate the hardware resources in network forwarding device.

INTRODUCTION

- Various network architectures or protocols can plug themselves as new modules into EIA network forwarding equipments, and run simultaneously for experiment and actual development.
- At the same time, users can select one or several different network architectures on EIA to use.
- Architectures upon EIA coexist and compete with each other, and some of them maybe victory or eliminated under natural selection (competition in the market).
- EIA uses such competition to promote the evolution of the core level network, and to achieve the prosperity of the future Internet.

RELATED WORK

- OpenFlow is a solution providing real data flows for researchers to carry out their Internet innovations, by an OpenFlow controller and OpenFlow switch structure.
  - Currently, it works in the local area network, and mainly deals with IPv4 protocol.
- PlanetLab is an open, shared platform for developing, deploying, and accessing planetary scale applications.
  - However it is mainly made up by end hosts but not packet forwarding device.
- VINI is also a test bed on which simultaneous experiments or slices run software router (by Click) over virtual topologies.
  - It does not provide the hardware router support which can be provided in EIA.
- ……

THE EIA NETWORK ARCHITECTURE

- EIA aims to be an “architecture of architectures”
- EIA make architecture more evovlable

THE EIA NETWORK ARCHITECTURE

EIA targets:

- 1) Researchers can program and test their new architectures or protocols in real network device.
  - Everyone can participate in the development of the switch
  - Experimental flow and normal flow can coexist and don't affect each other.
- 2) EIA is able to accommodate varieties of other architectures, and works as a platform to enable the competition among the upper architectures.
  - Architectures upon EIA can compete to solve a same problem, or form a complementary relationship to solve different problems.
  - Different architectures share the same underlying hardware resources.

THE EIA NETWORK ARCHITECTURE

EIA composition:

- **1) EIA host**
  - Normal hosts with EIA software to support multi-architecture;
- **2) EIA network equipment:**
  - EIA supported network devices (switches or routers).

THE EIA NETWORK ARCHITECTURE

EIA host:

- EIA module is a number of function interfaces located in the link layer.
  - users just need to download the EIA software and install it
- EIA module passes packets to different architectures above it.
- The existing TCP/IP communication protocol stacks (IPv4, IPv6) run as two special architectures and coexist with other new architectures.

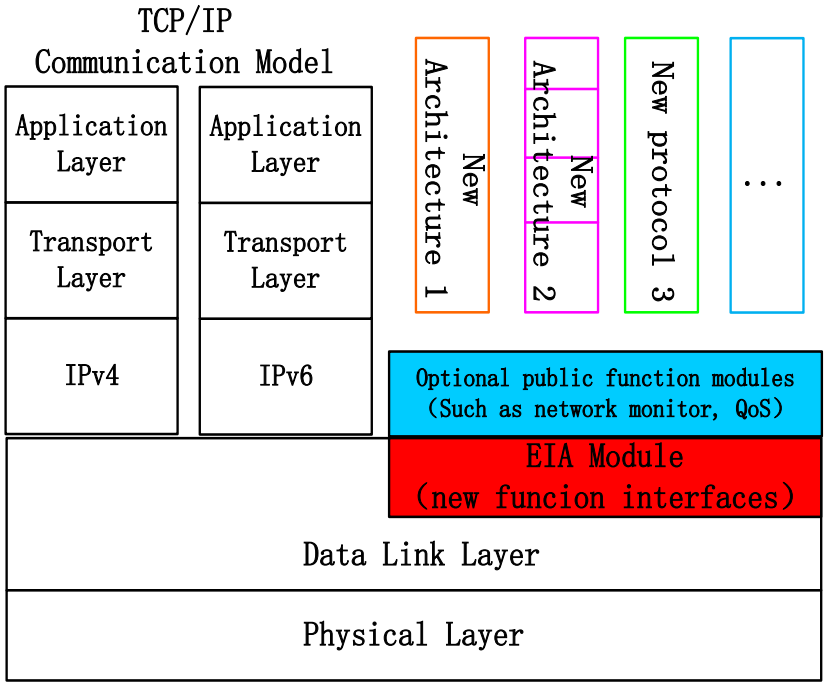

Basic model of EIA host to support multiple architectures

THE EIA NETWORK ARCHITECTURE

EIA network equipment:

- The instruction set in EIA network device will be provided as primitive operation to directly operate the hardware resources
  - similar to the x86 instruction in computer.

- Package action interface layer encapsulates the lower layer instructions to a packet level interface for the upper architectures
  - such as packet forwarding, discarding, and rewriting functions and so on.

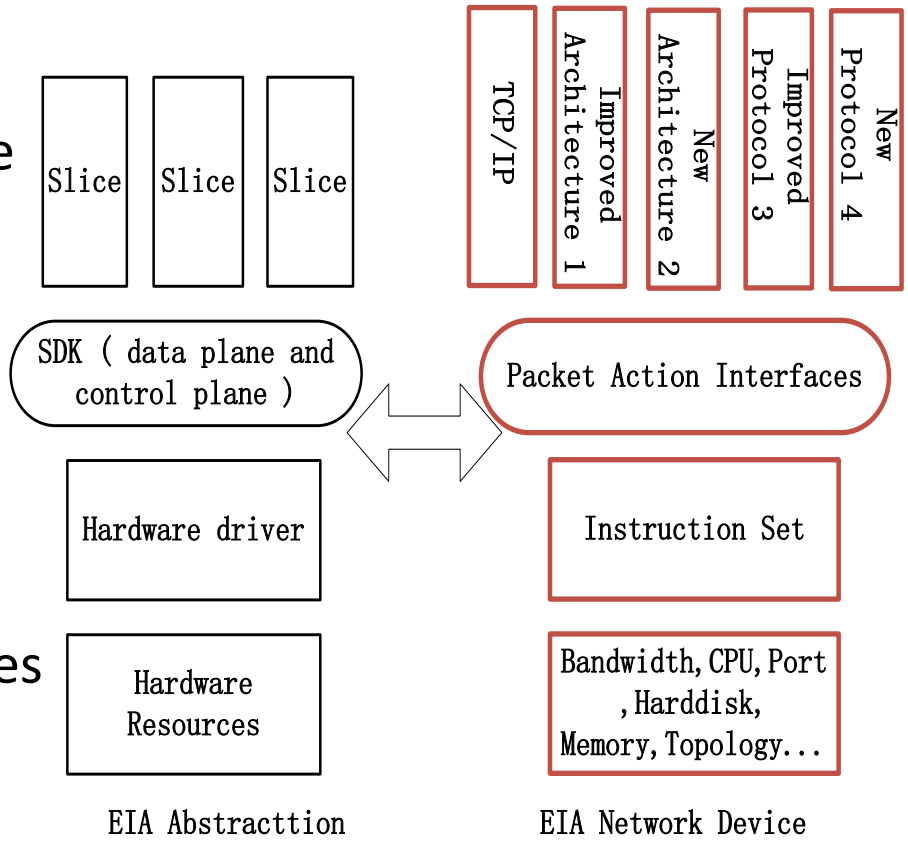

EIA network equipment

THE EIA NETWORK ARCHITECTURE

How to support EIA for commercial network equipment?

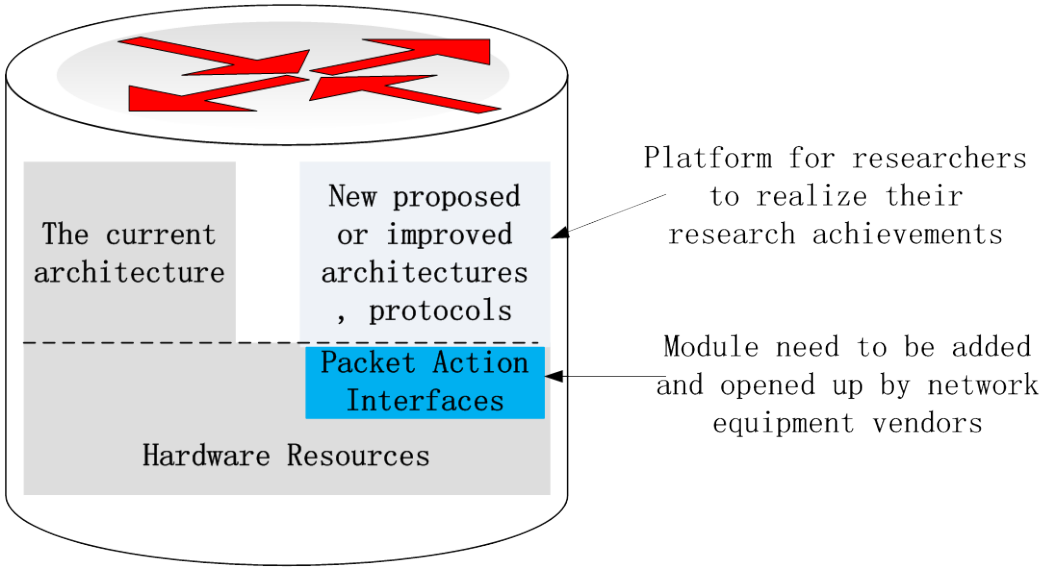

Commercial switches to support EIA

THE EIA NETWORK ARCHITECTURE

Current network substrate with EIA routers:

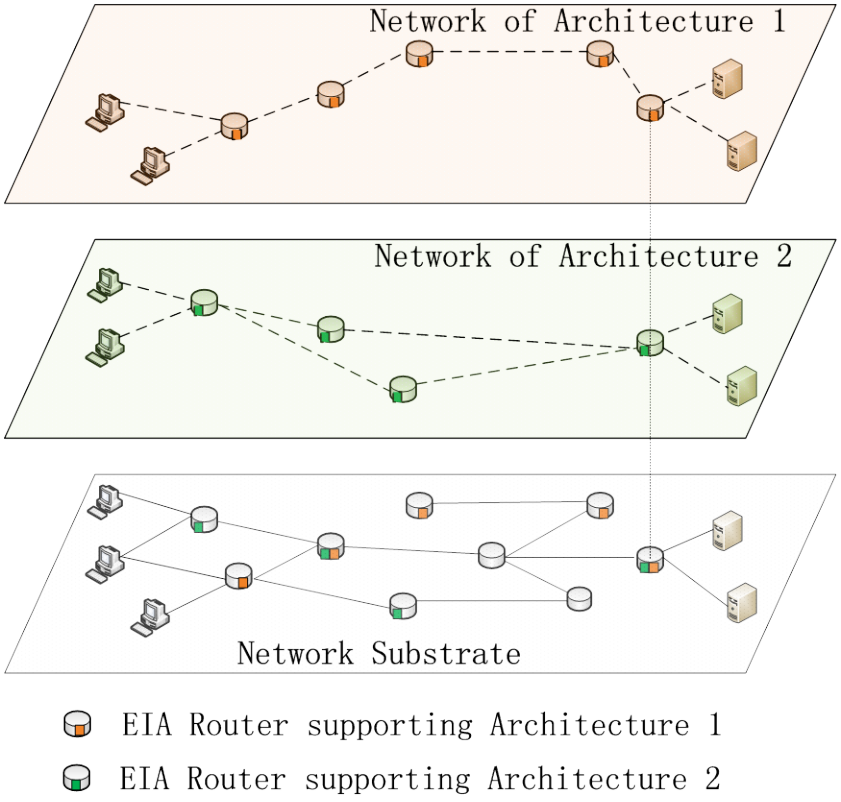

EVALUATION

The cycle time for a new protocol from the conception design to the commercial use is greatly shortened:

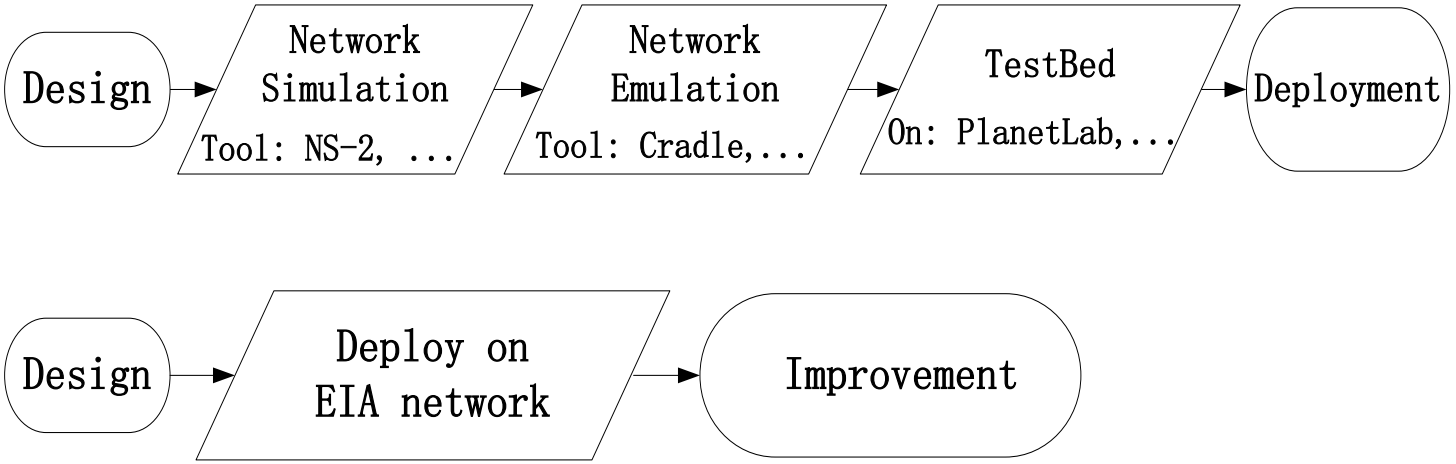

Traditional way to deploy a new protocol to the Internet (above), the deploy process on EIA network (blow)

EVALUATION

The benefit to normal Internet users:

- For users, if they want to enjoy a good service from a specific network on EIA. They just need to download an EIA software.
  - For example, they can access network A which provides high-definition multicast IPTV service by downloading software A, or access network B on which it is surprisingly simple and convenient for them to do the online payment by downloading software B.

EVALUATION

Business model or Incentives for ISP:

- 1) Through government support.
  - Currently academic achievements are always difficult to be deployed to the real network.
  - EIA can increase the conversion rate of the scientific research achievements.

- 2) If service in the new architecture or protocol on EIA can gain profits, this service or protocol can pay for his partner ISP.

CONCLUSION

- EIA proposed runs as a carrier or socket for other network architectures.
- EIA acts as a special programmable platform for researchers to realize their ideal future Internet architectures.
- EIA breaks the situation that one network architecture dominates the Internet.
- Architectures upon EIA can coexist and compete to achieve the continual evolution to the core level of the network.
- EIA is a open platform, with EIA switches, everyone can participate in the development instead of relying on vendors to implement.

Recursive Networks

Joe Touch
USC/ISI

Introduction

- Virtualization requires recursion
- Recursion supports layering
- Recursion is a native network property
  - Integrates and virtualization, forwarding and layering **in a single mechanism**

RNA Metaprotocol

- Template of basic protocol service:
  - Establish / refresh state
  - Encrypt / decrypt message
  - Apply filtering
  - Pace output via flow control
  - Pace input to allow reordering
  - Multiplex/demultiplex
    - includes switching/forwarding

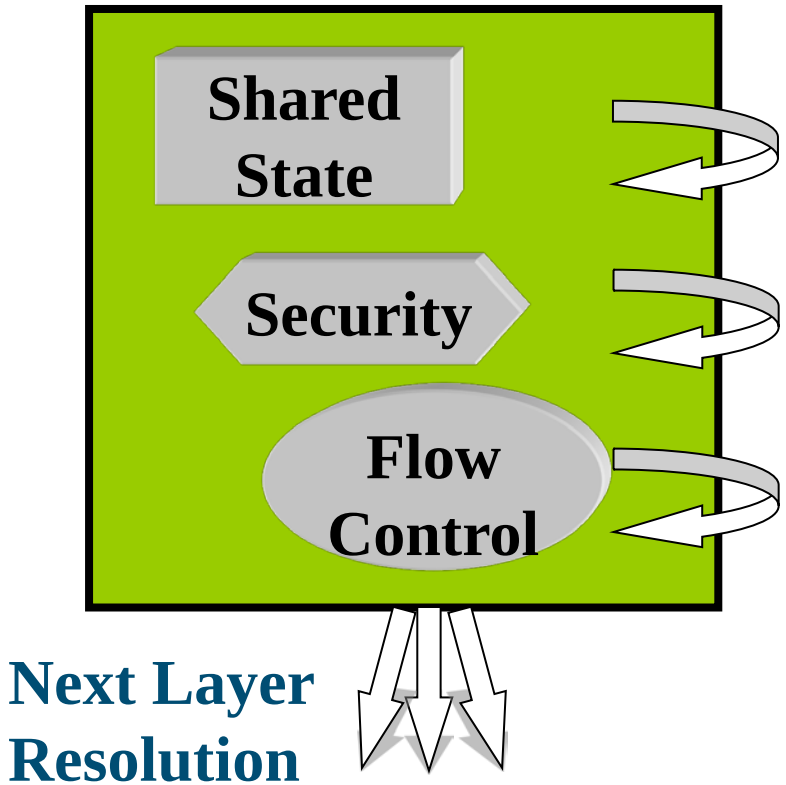

Recursion requires new layers – where? Why?

- Wedge between (IPsec, left)
  or replicate (virtualization, right)

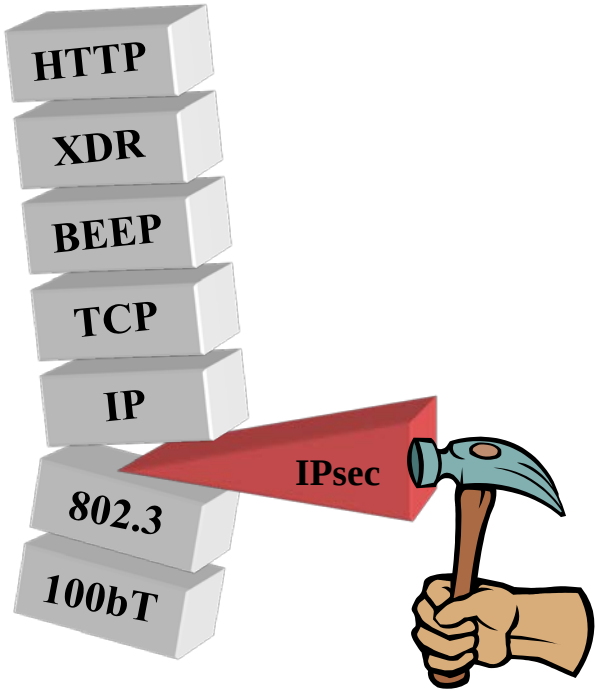

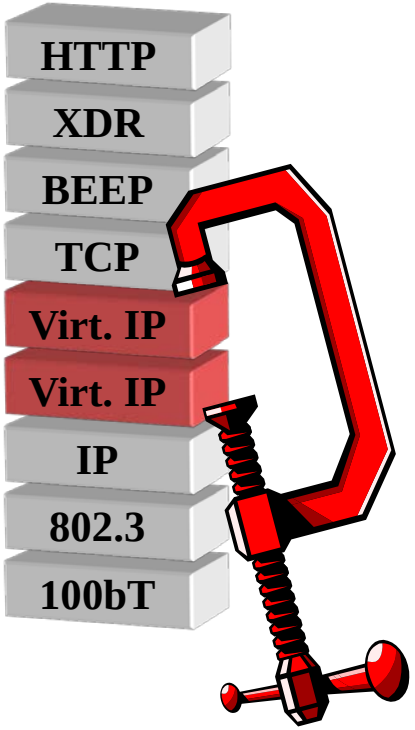

Trustworthy Internet

Tsinghua University

Trustworthy New Generation Internet

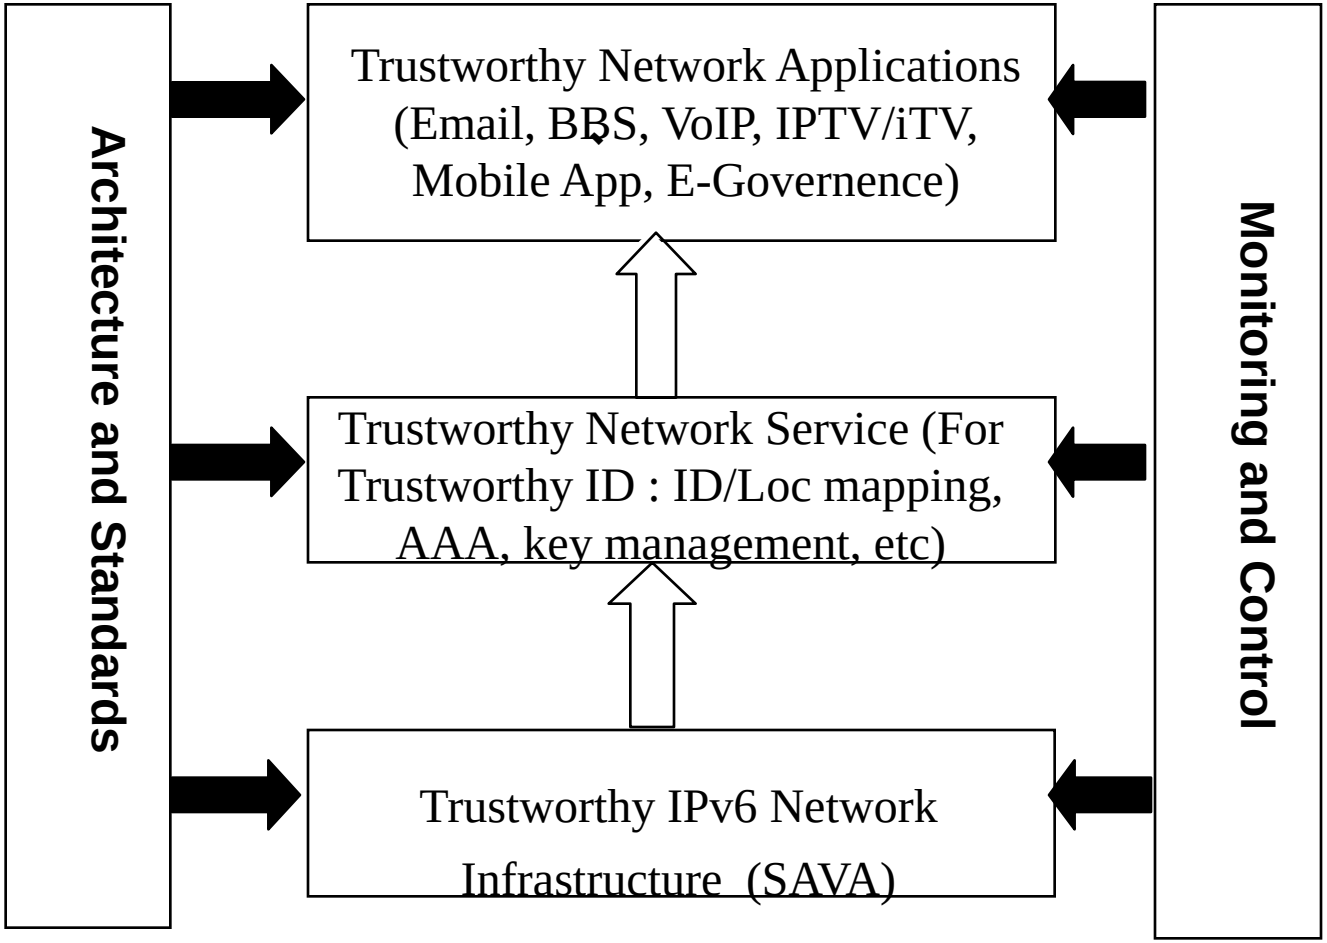

Video：David Clark，”Mapping the Future of the Internet”, AsiaFI Future Internet Architecture Workshop, Aug 27, 2010.

Q&A