Software-Define Networking

Soheil Abbasloo
Department of Computer Science
University of Toronto

Fall 2022

Outline

- Innovation: Computers vs. Networks
- SDN
- OpenFlow

Innovation
Computers vs. Networks

- How difficult is it to create/modify a computer application?
- How difficult is it to create/modify a network feature?
- What is the difference?
- What are the tools available for each?

Innovation in Applications

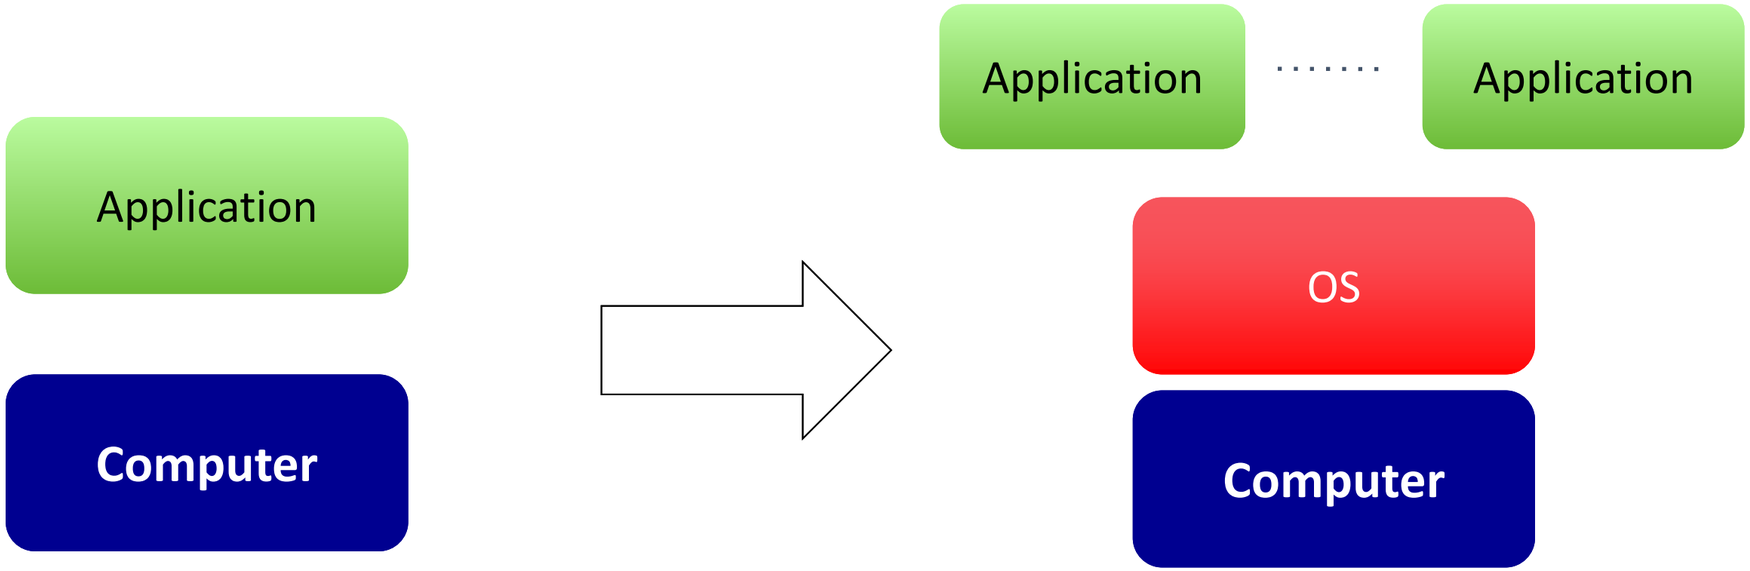

OS abstracts hardware substrate
→ Innovation in applications

Innovation in OS and Applications

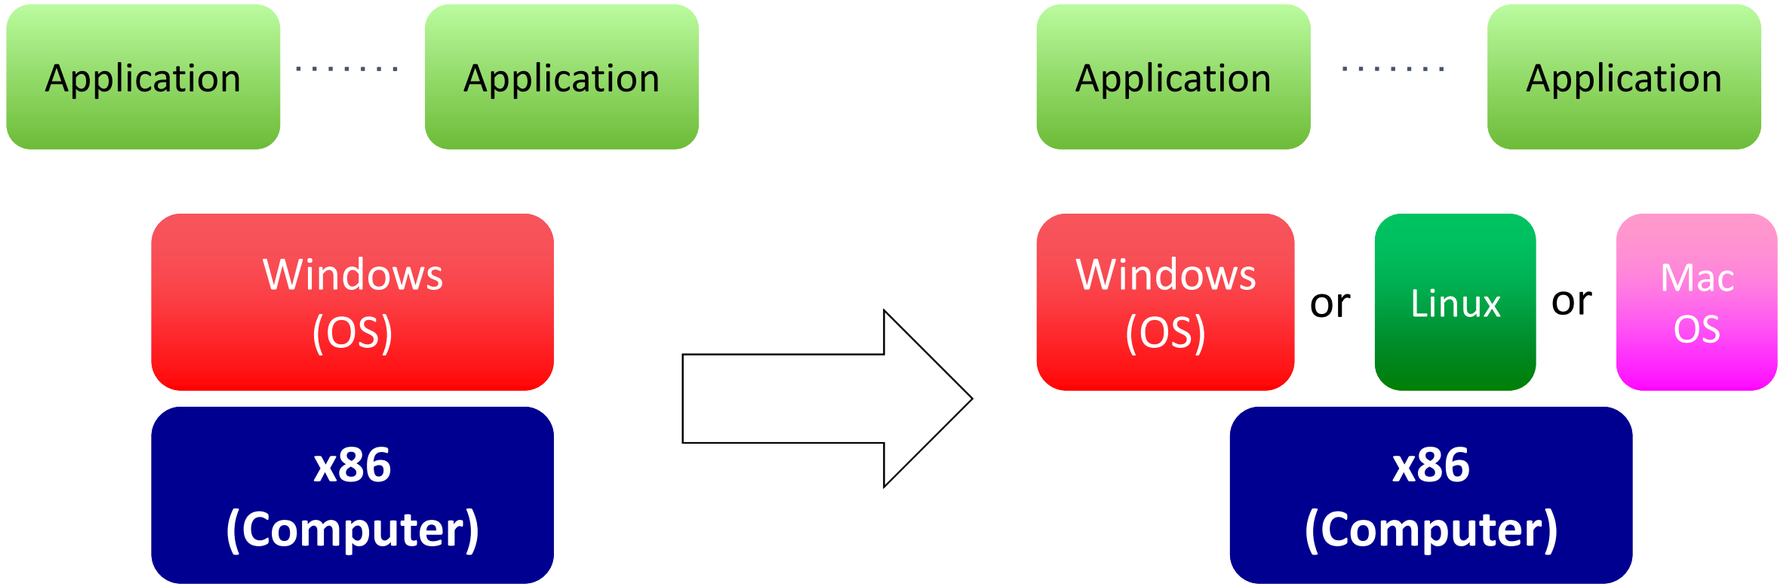

Simple, common, stable, hardware substrate below
+ Programmability
+ Competition
→ Innovation in OS and applications

Innovation in Infrastructure

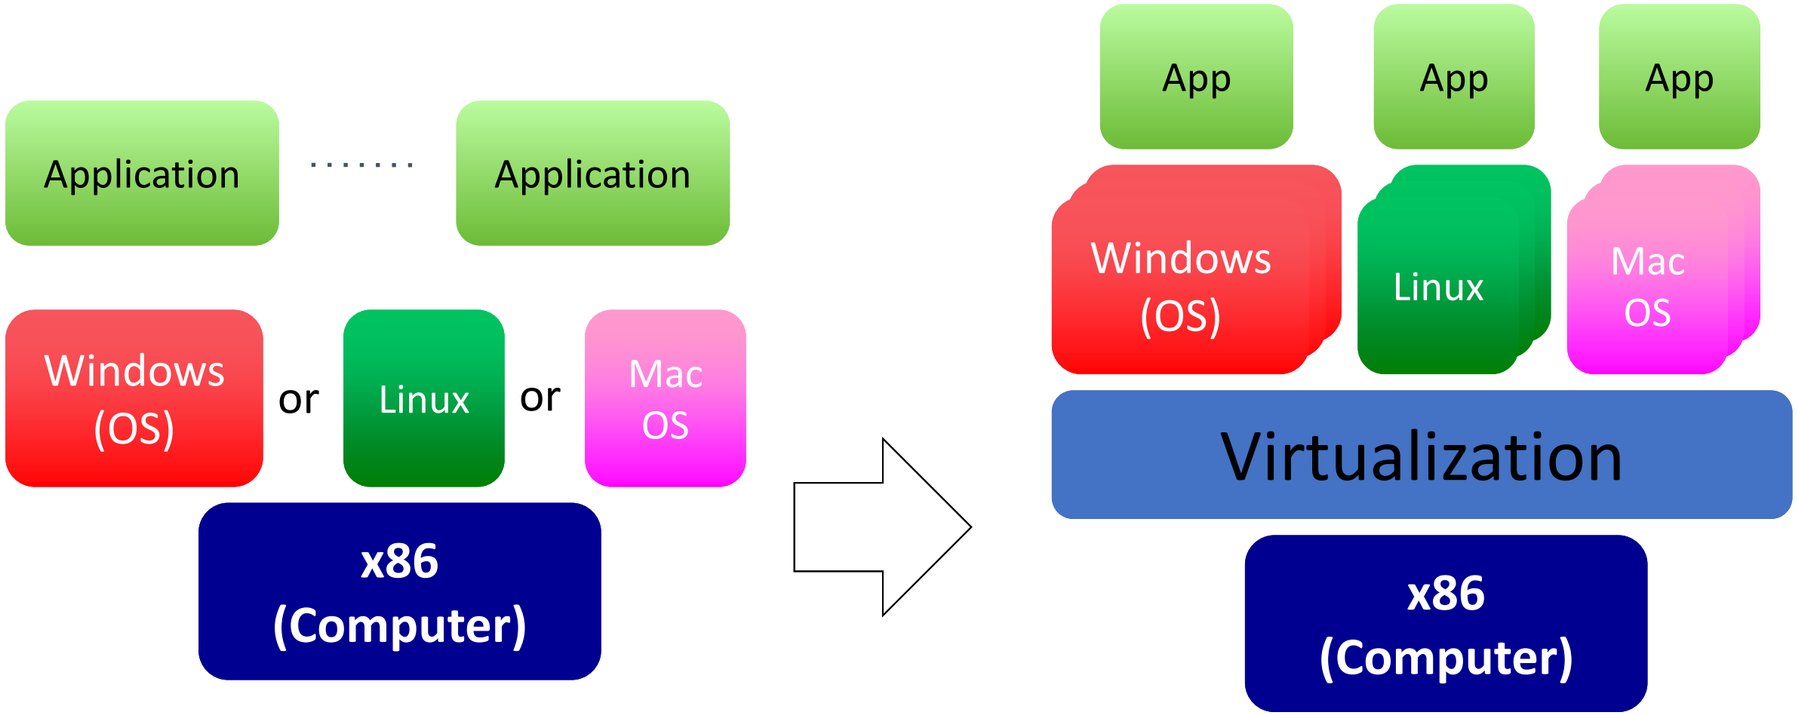

Simple, common, stable, hardware substrate below
+ Programmability
+ Strong isolation model
+ Competition above
→ Innovation in infrastructure

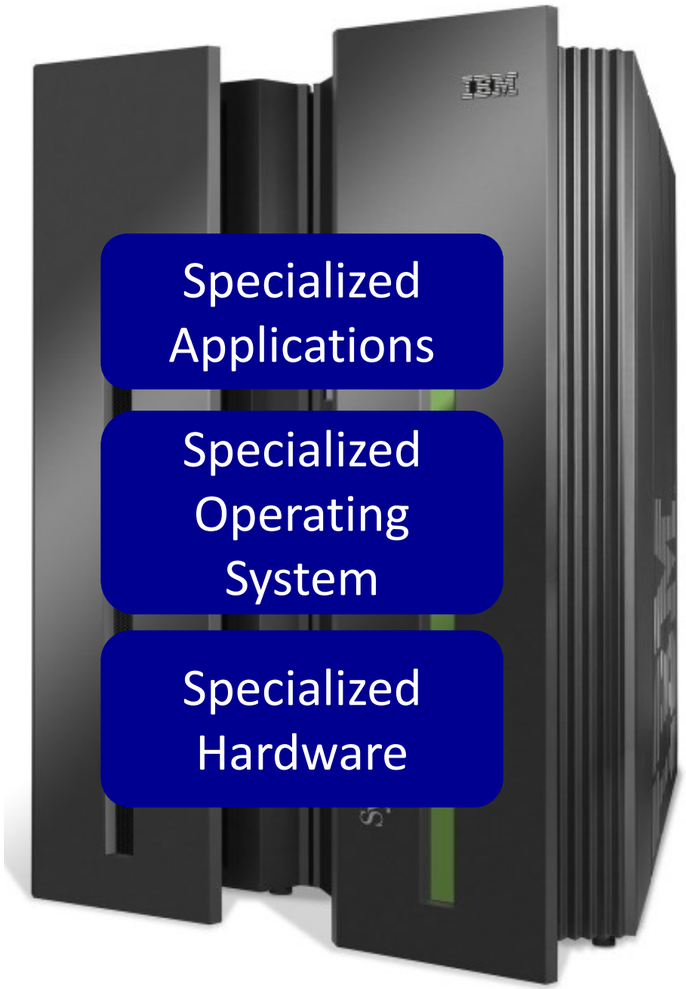

Vertically integrated
Closed, proprietary
Slow innovation
Small industry

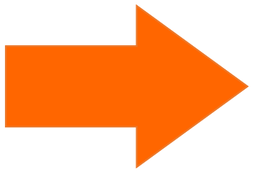

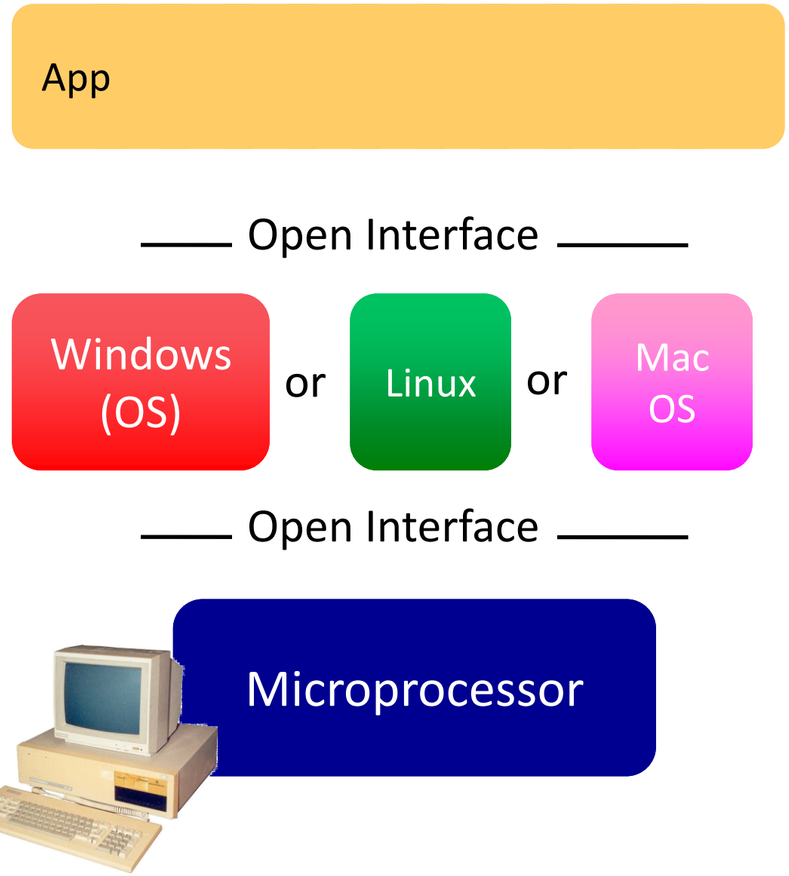

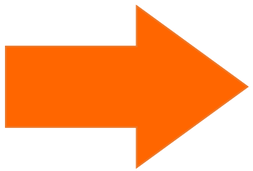

Horizontal
Open interfaces
Rapid innovation
Huge industry

We Have Lost Our Way

Routing, management, mobility management, access control, VPNs, …

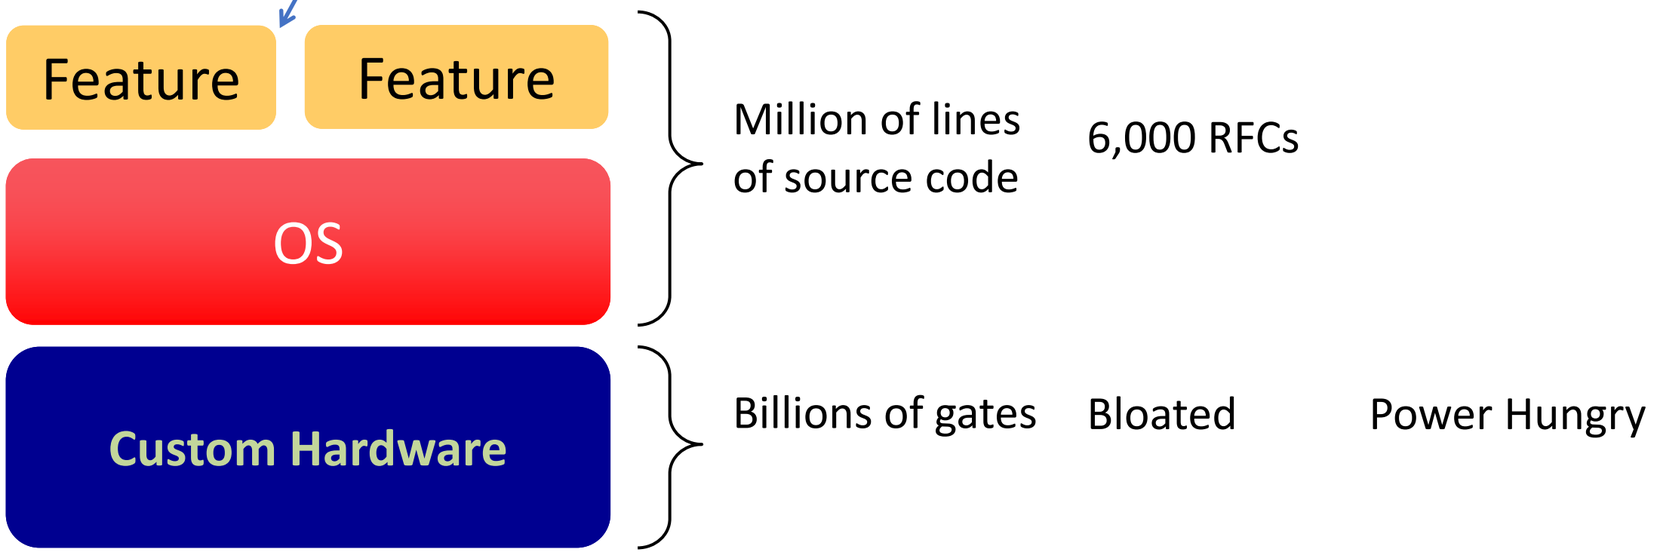

- Vertically integrated, complex, closed, proprietary
- Networking industry with “mainframe” mind-set

Reality is Even Worse

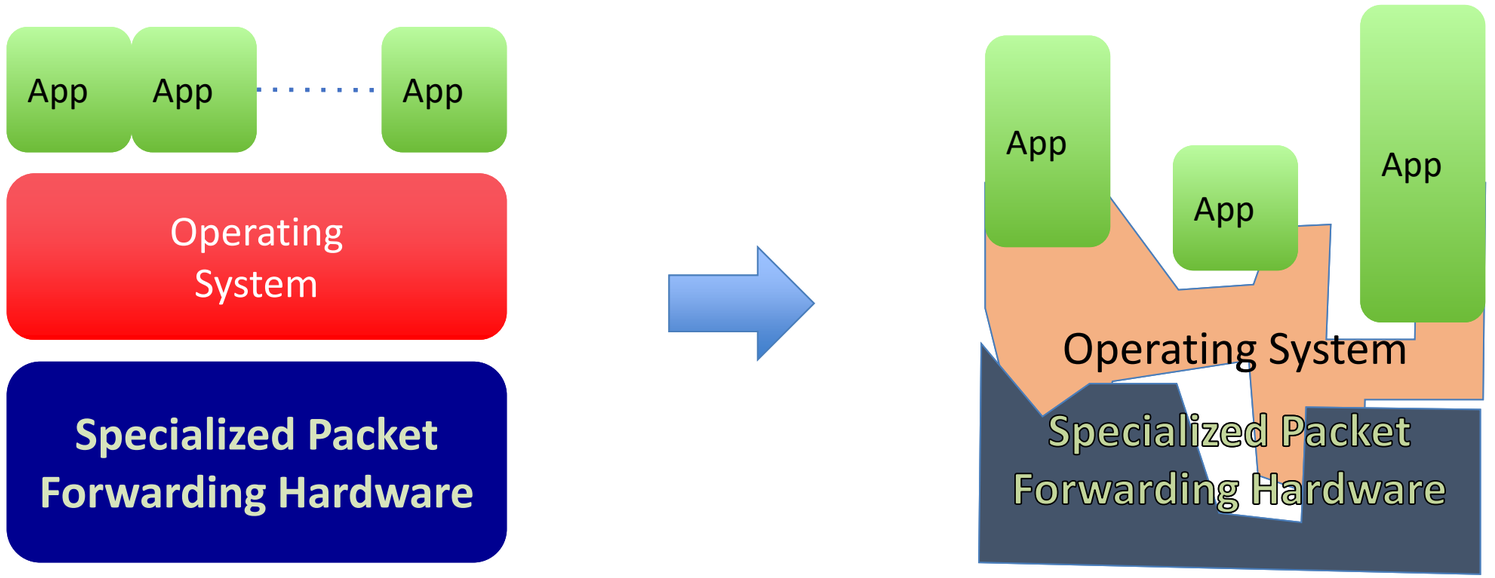

- Lack of competition means glacial innovation
- Closed architecture means blurry, closed interfaces

App

Open Interface

Control Plane or Control Plane or Control Plane

Open Interface

Merchant Switching Chips

Vertically integrated
Closed, proprietary
Slow innovation

Horizontal
Open interfaces
Rapid innovation

What we need …

1) Separate Intelligence from Datapath

Operators, users, 3rd party developers, researchers, …

New function!

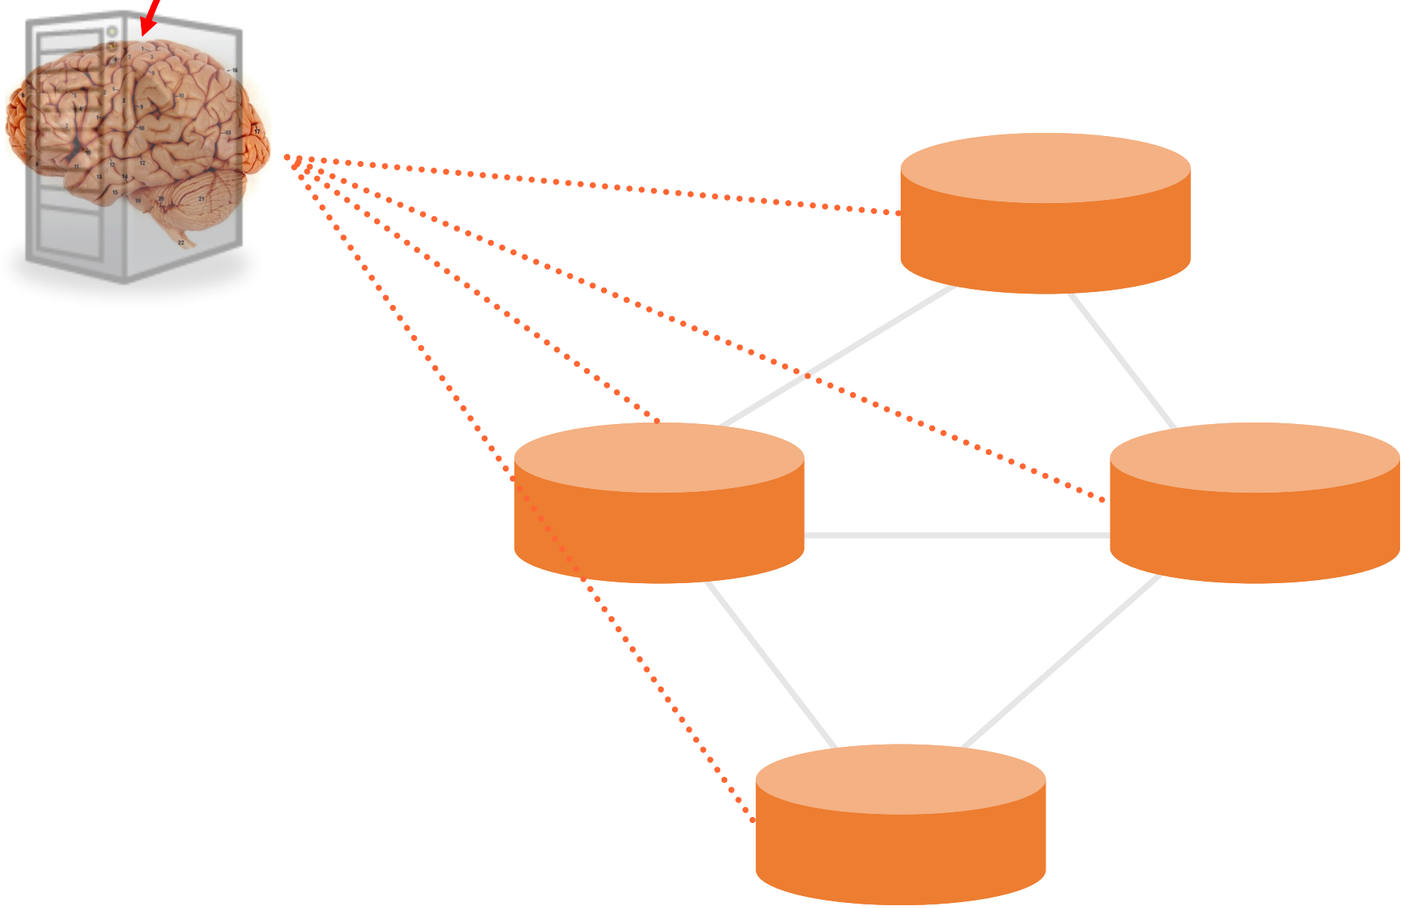

2) Cache Decisions

- In minimal flow-based datapath

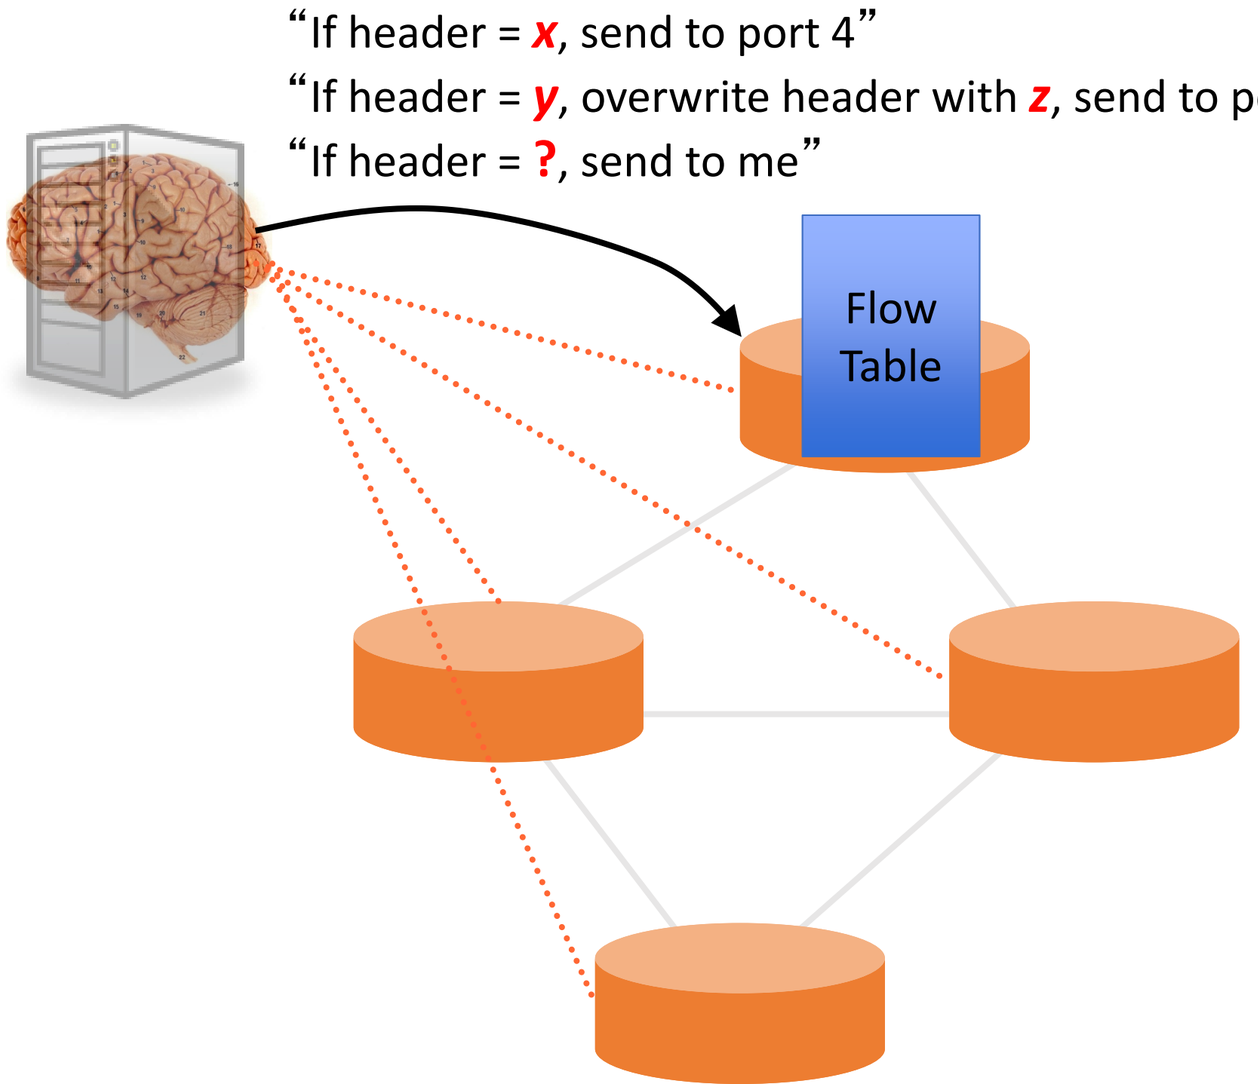

How Can We Do This?

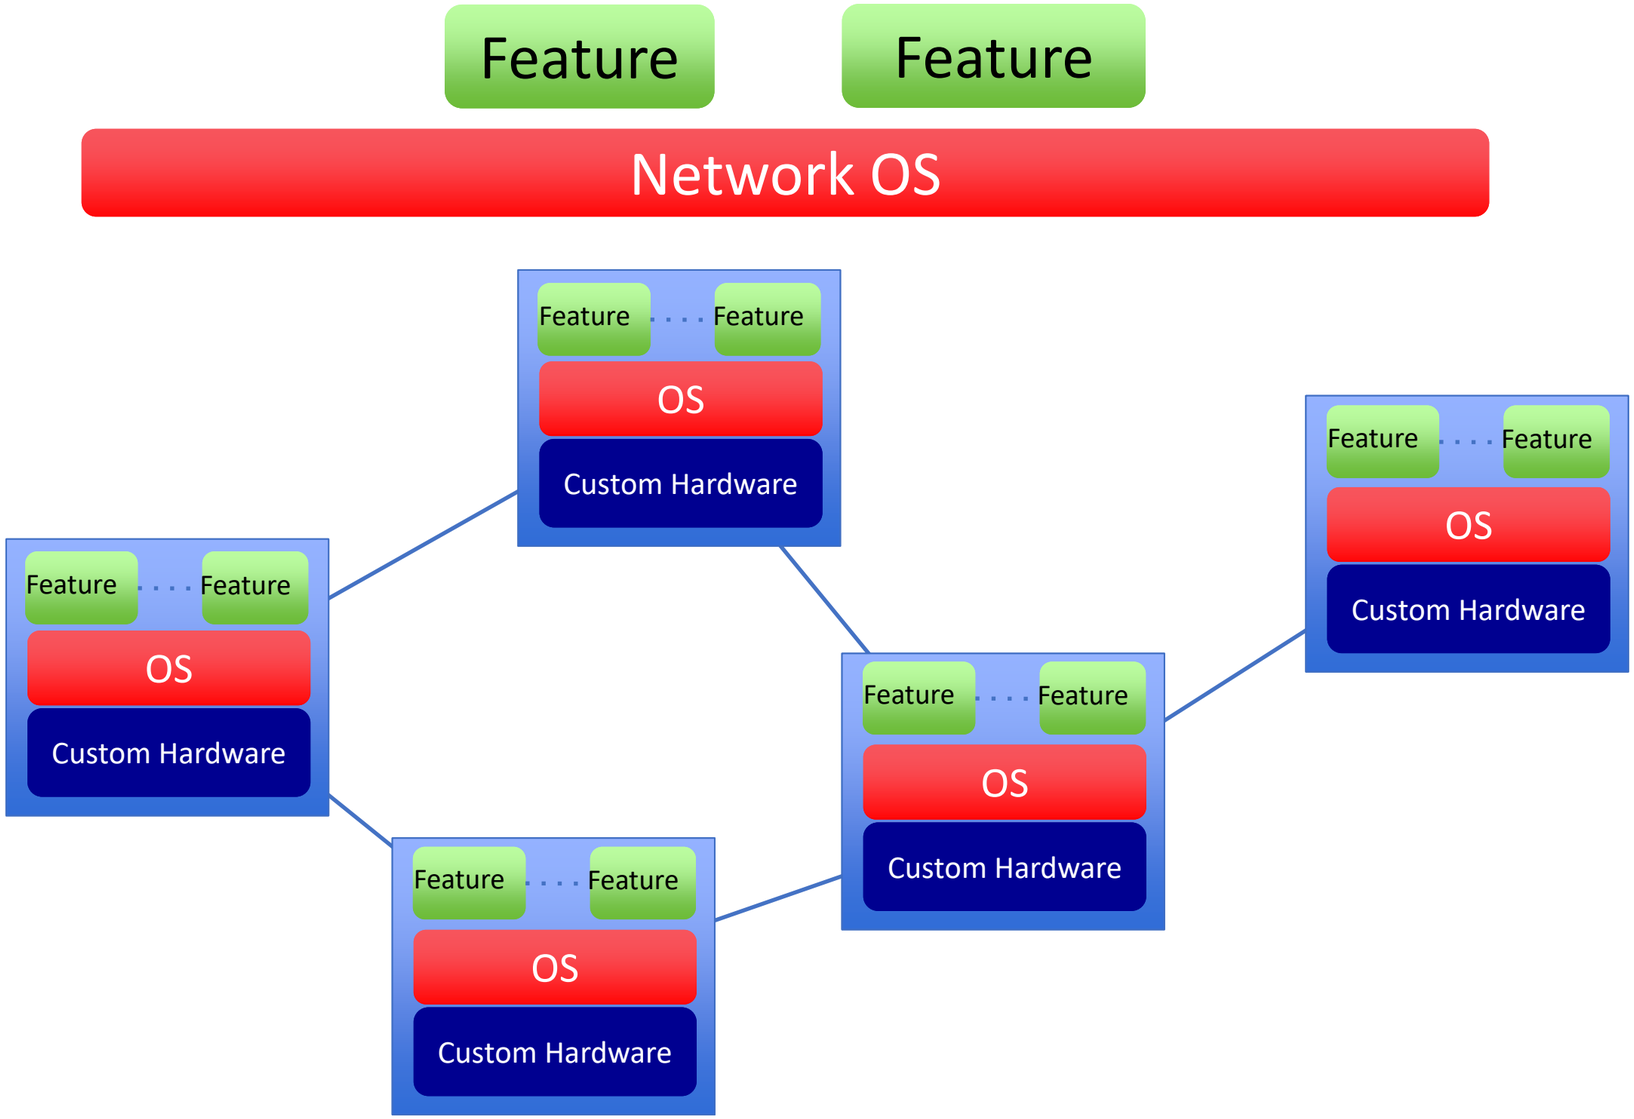

Software Defined Network (SDN)

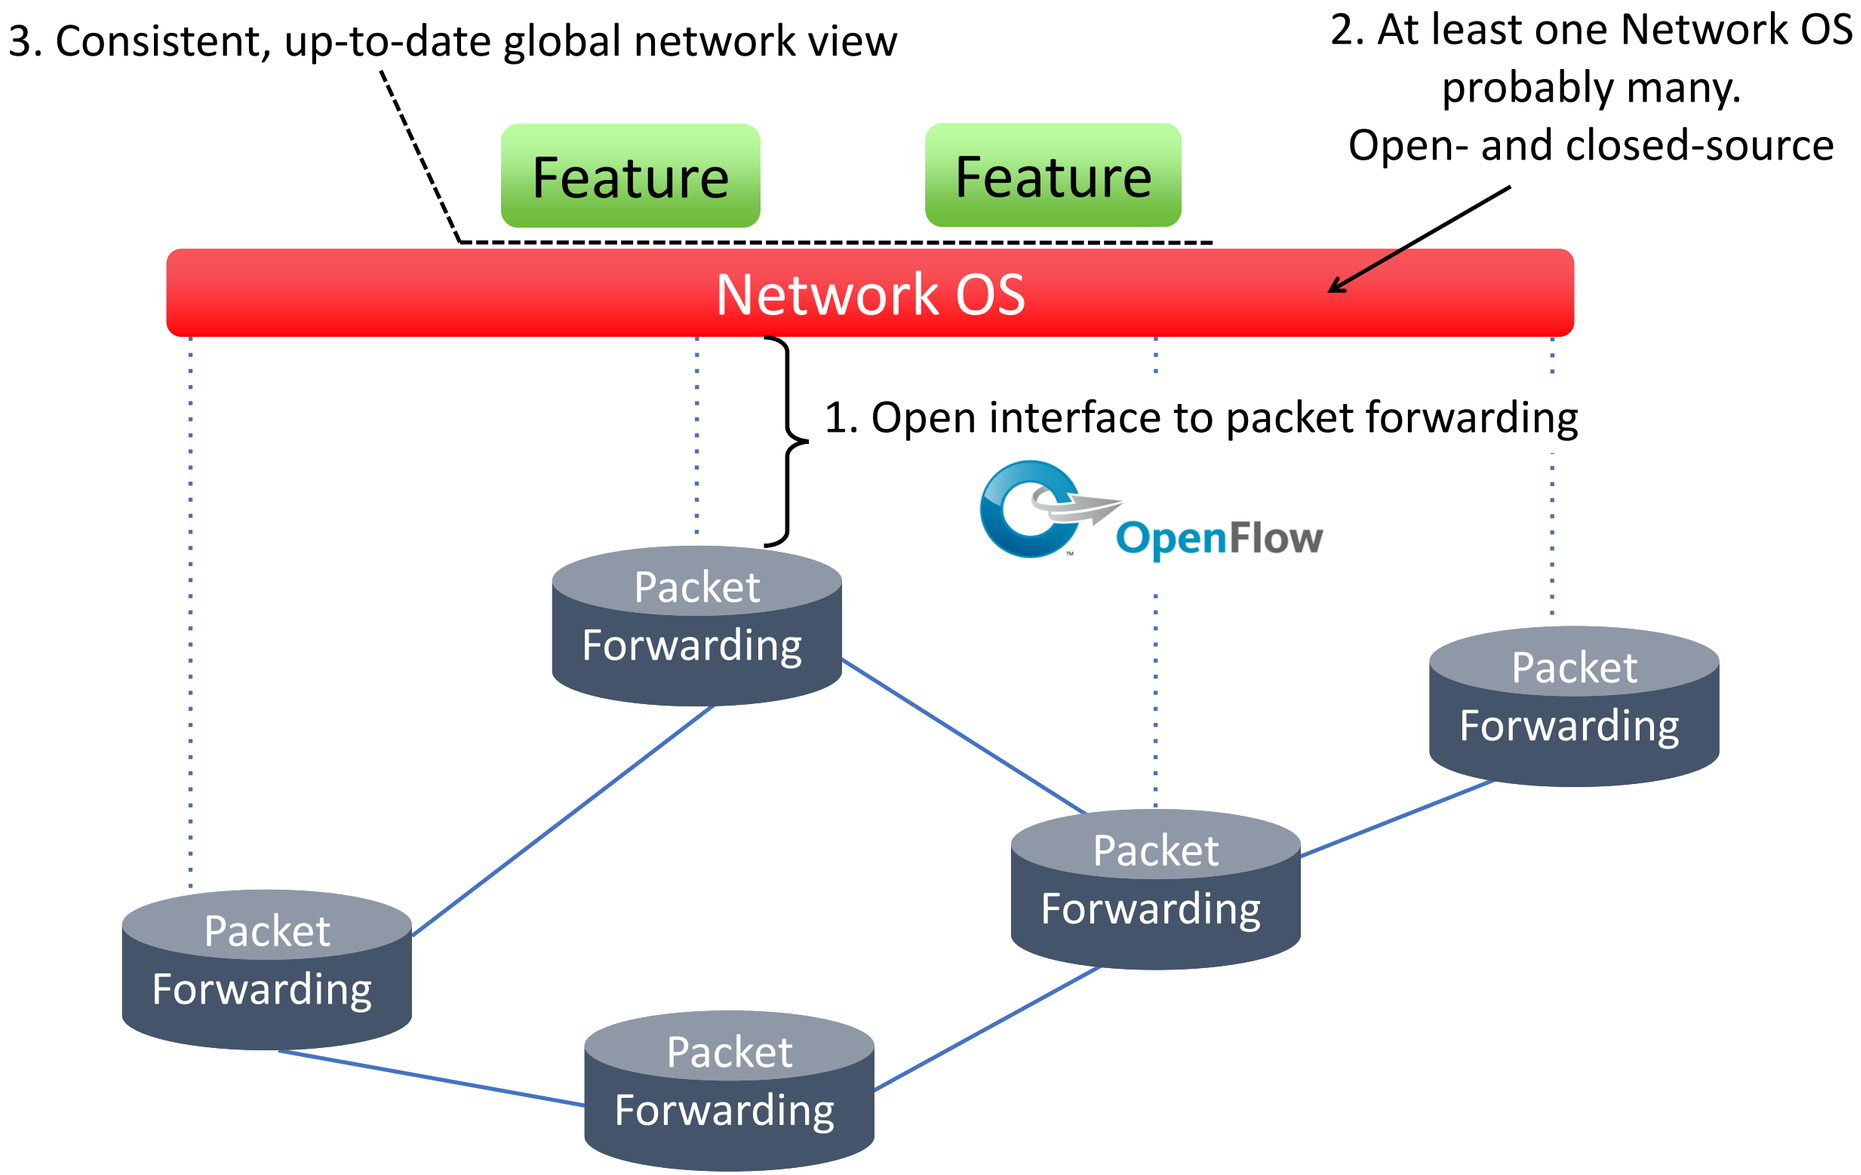

Consequences

- More innovation in network services
  - Owners, operators, 3rd party developers, researchers can improve the network
  - E.g. energy management, data center management, policy routing, access control, denial of service, mobility
- Lower barrier to entry for competition
  - Healthier market place, new players
- Lower cost
  - Infrastructure
  - Management

Example: Routing

- OSPF
  - RFC 2328: 245 pages
- Distributed System
  - Builds consistent, up-to-date map of the network: 101 pages
- Dijkstra’s Algorithm
  - Operates on map: 4 pages

Example: Routing

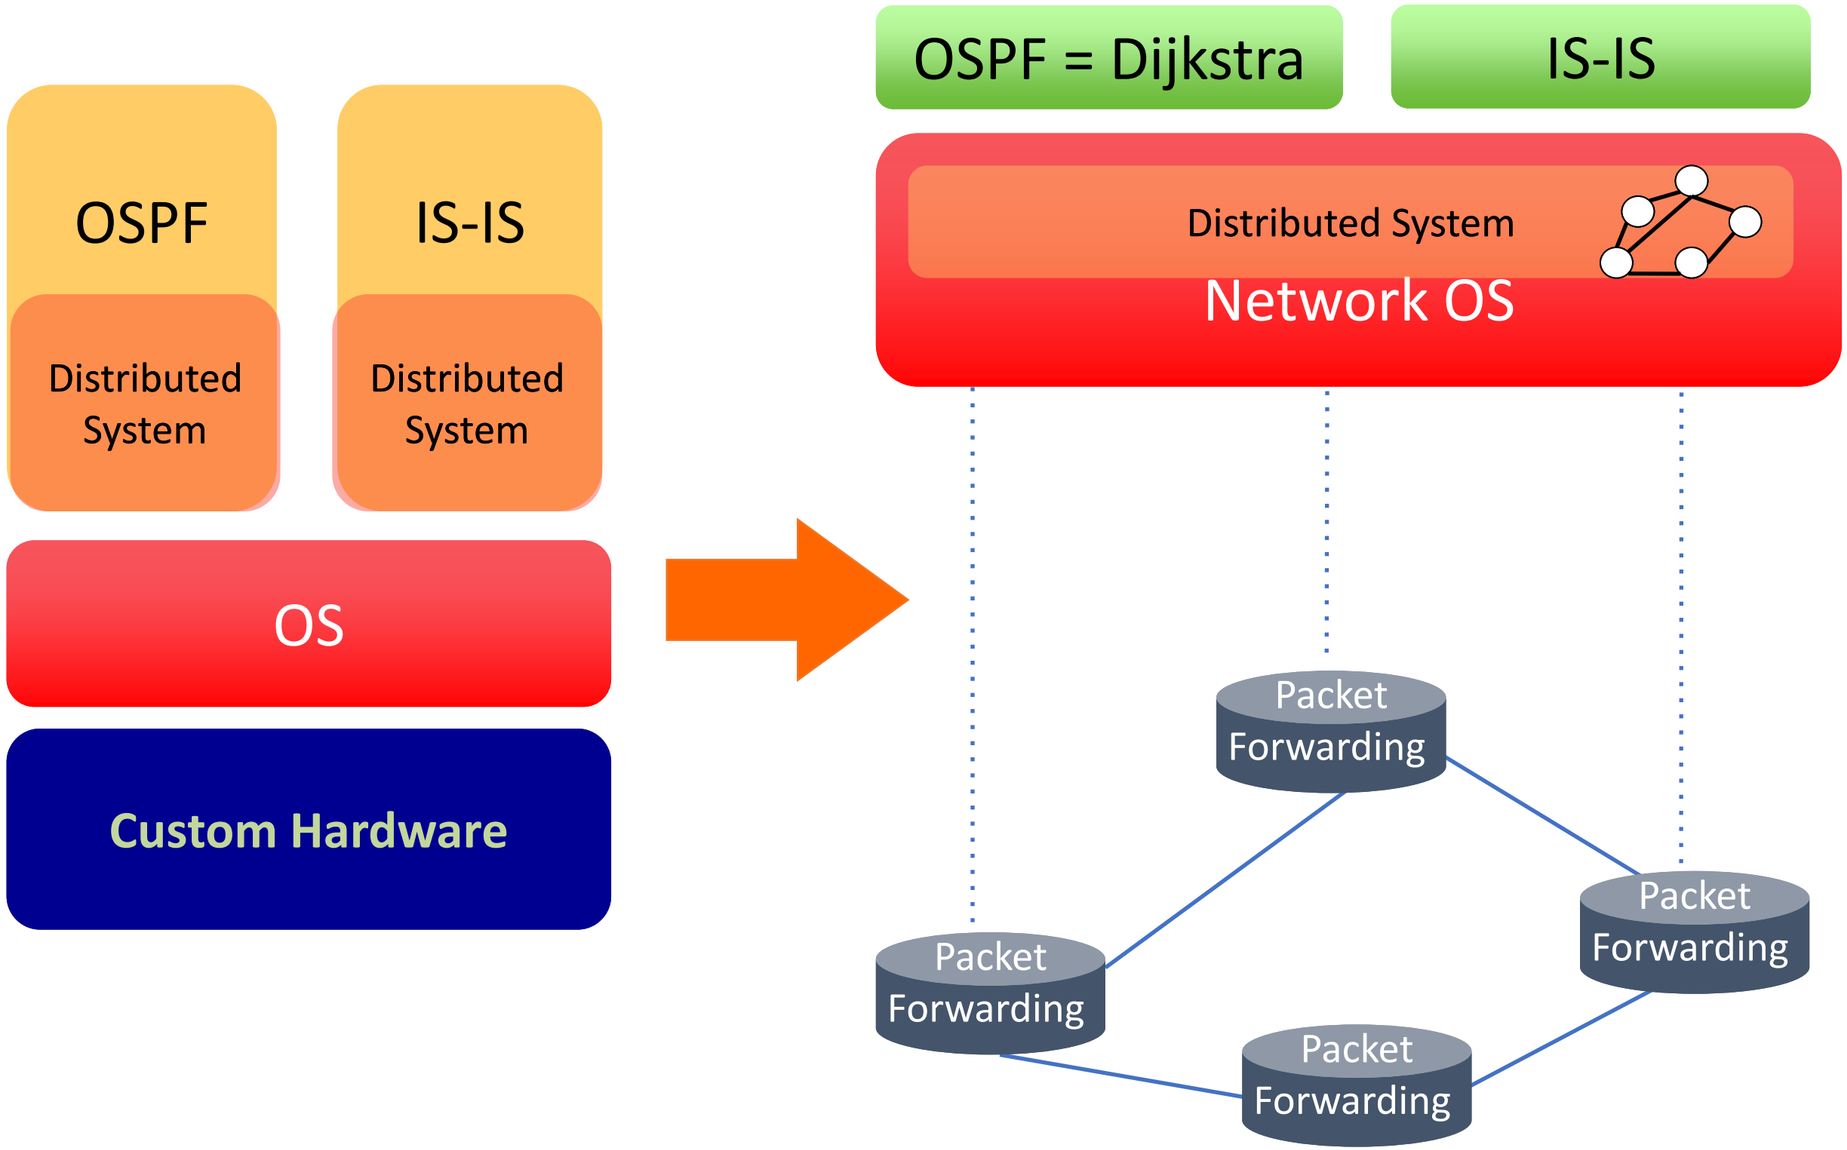

Back to the main storyline …

Software Defined Network (SDN)

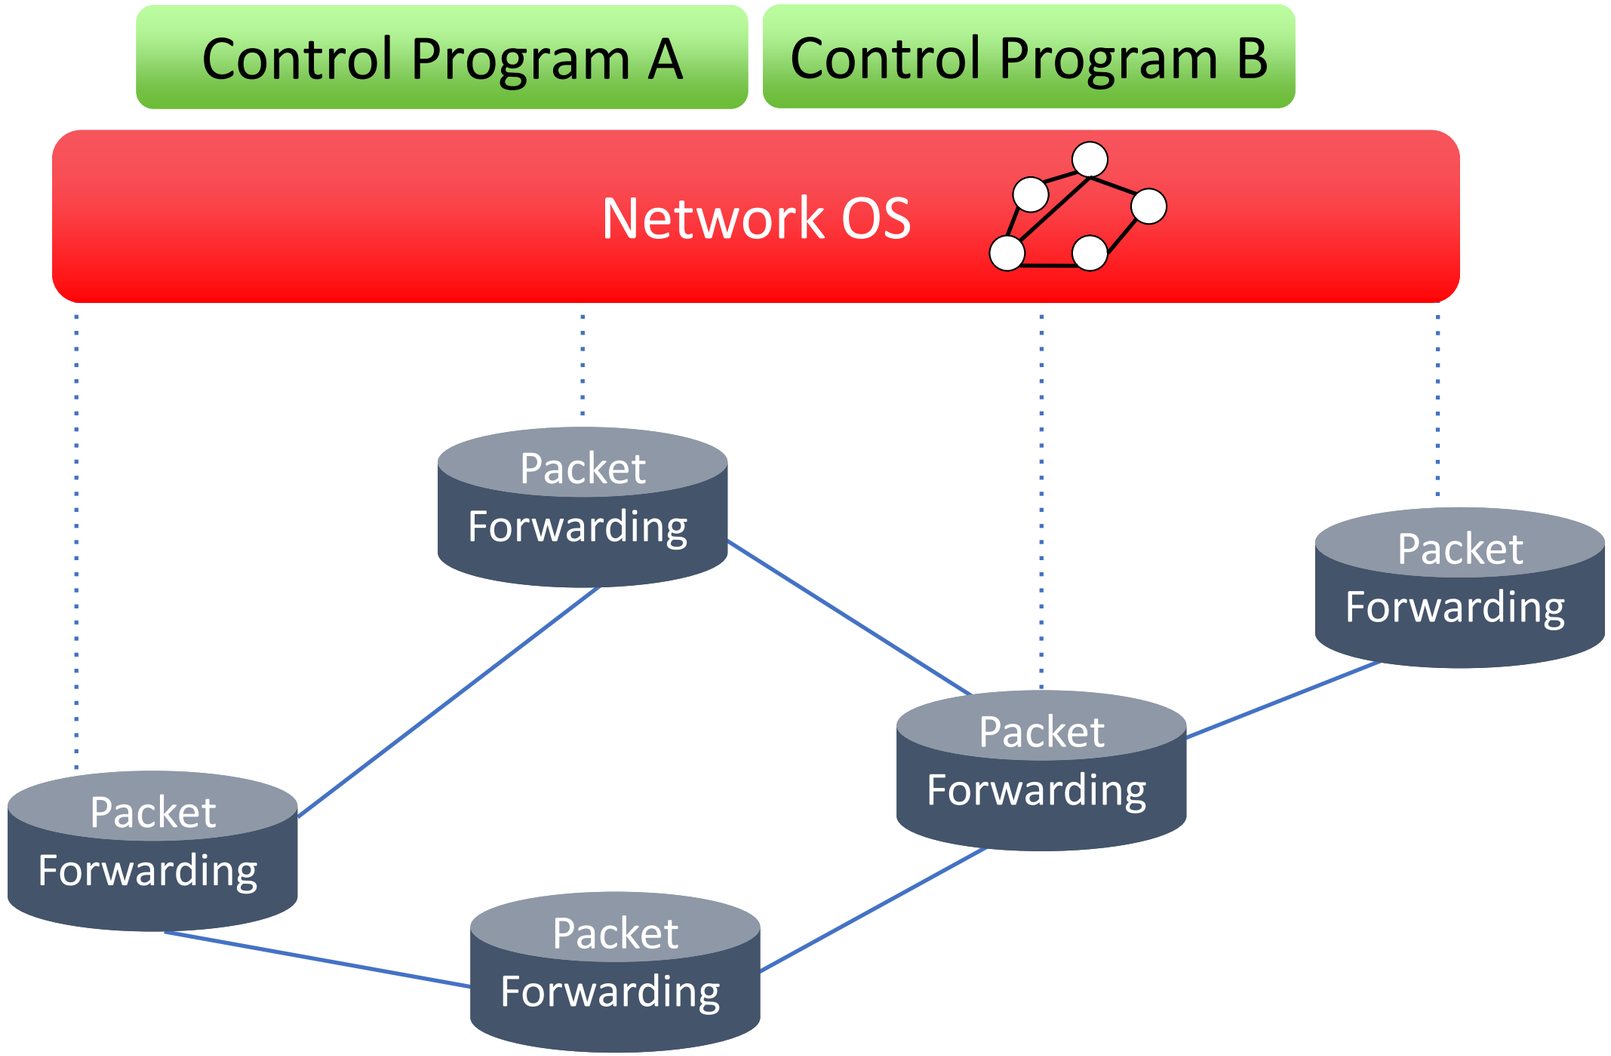

1) Network OS

Network OS: distributed system that creates a consistent, up-to-date network view

- Runs on servers (controllers) in the network
- NOX, ONIX, HyperFlow, Kandoo, Floodlight, Trema, Beacon, Maestro, Beehive, OpenDayLight, … + more

Uses forwarding abstraction to:

- Get state information **from** forwarding elements
- Give control directives **to** forwarding elements

Software Defined Network (SDN)

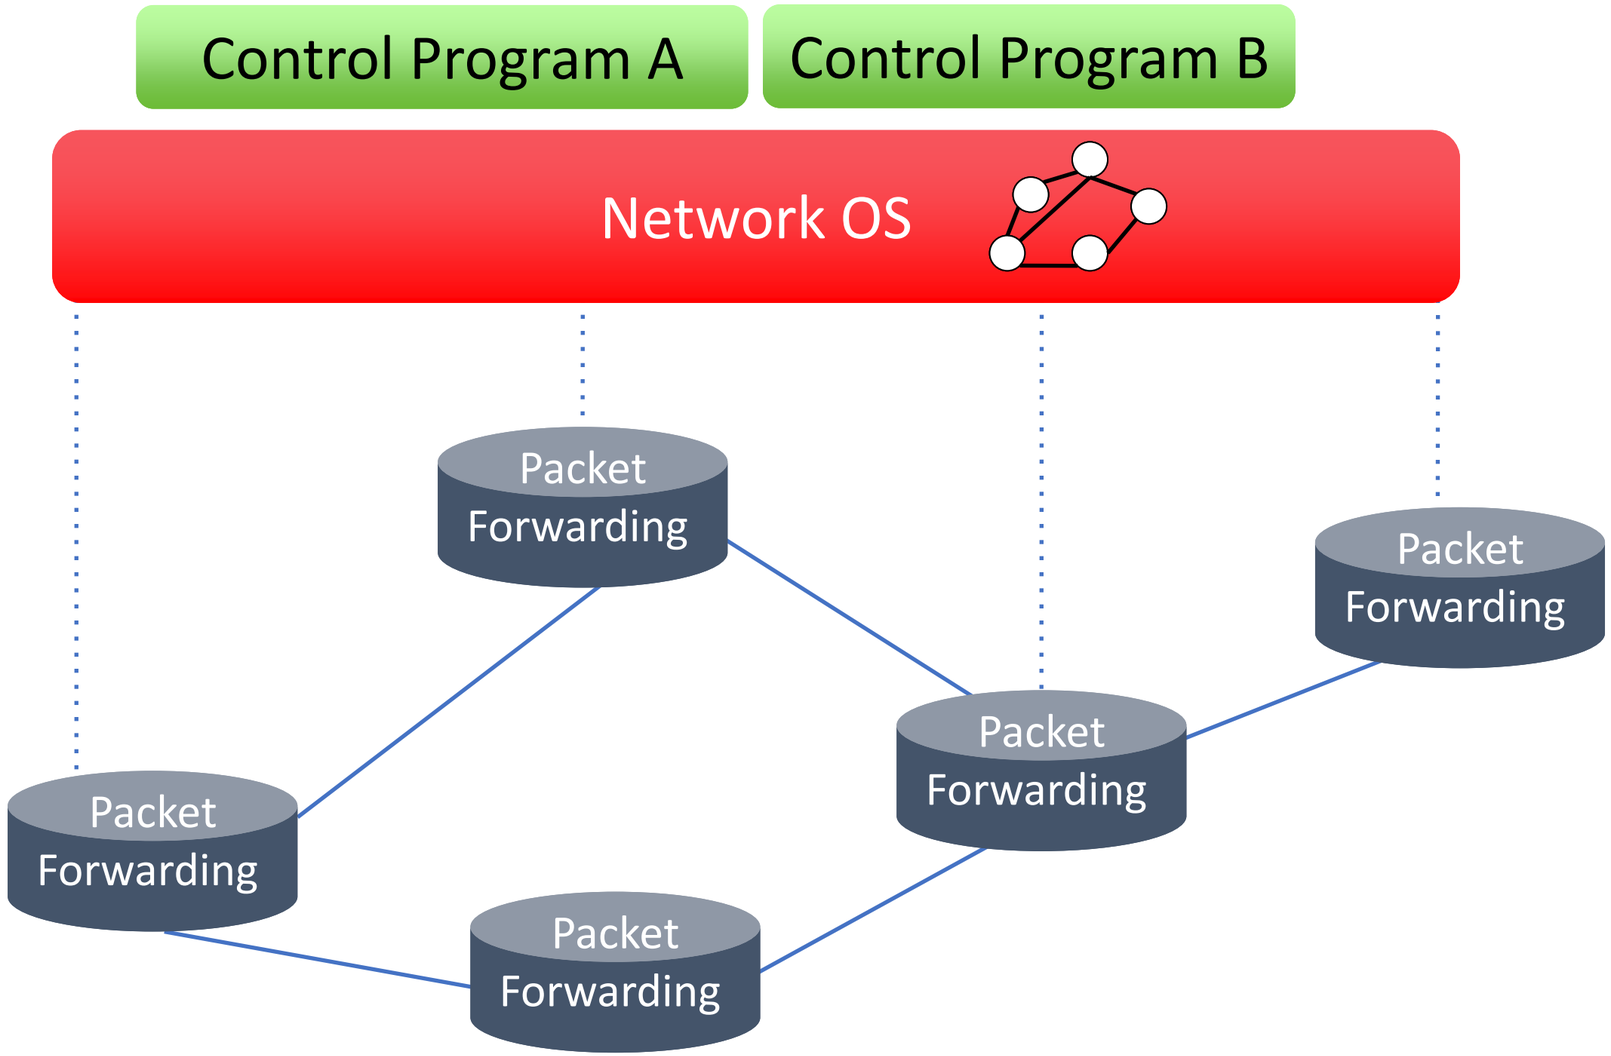

2) Control Program

Control program operates on view of network

- **Input**: global network view (graph/database)
- **Output**: configuration of each network device

Control program is not a distributed system

- Abstraction hides details of distributed state

Software Defined Network (SDN)

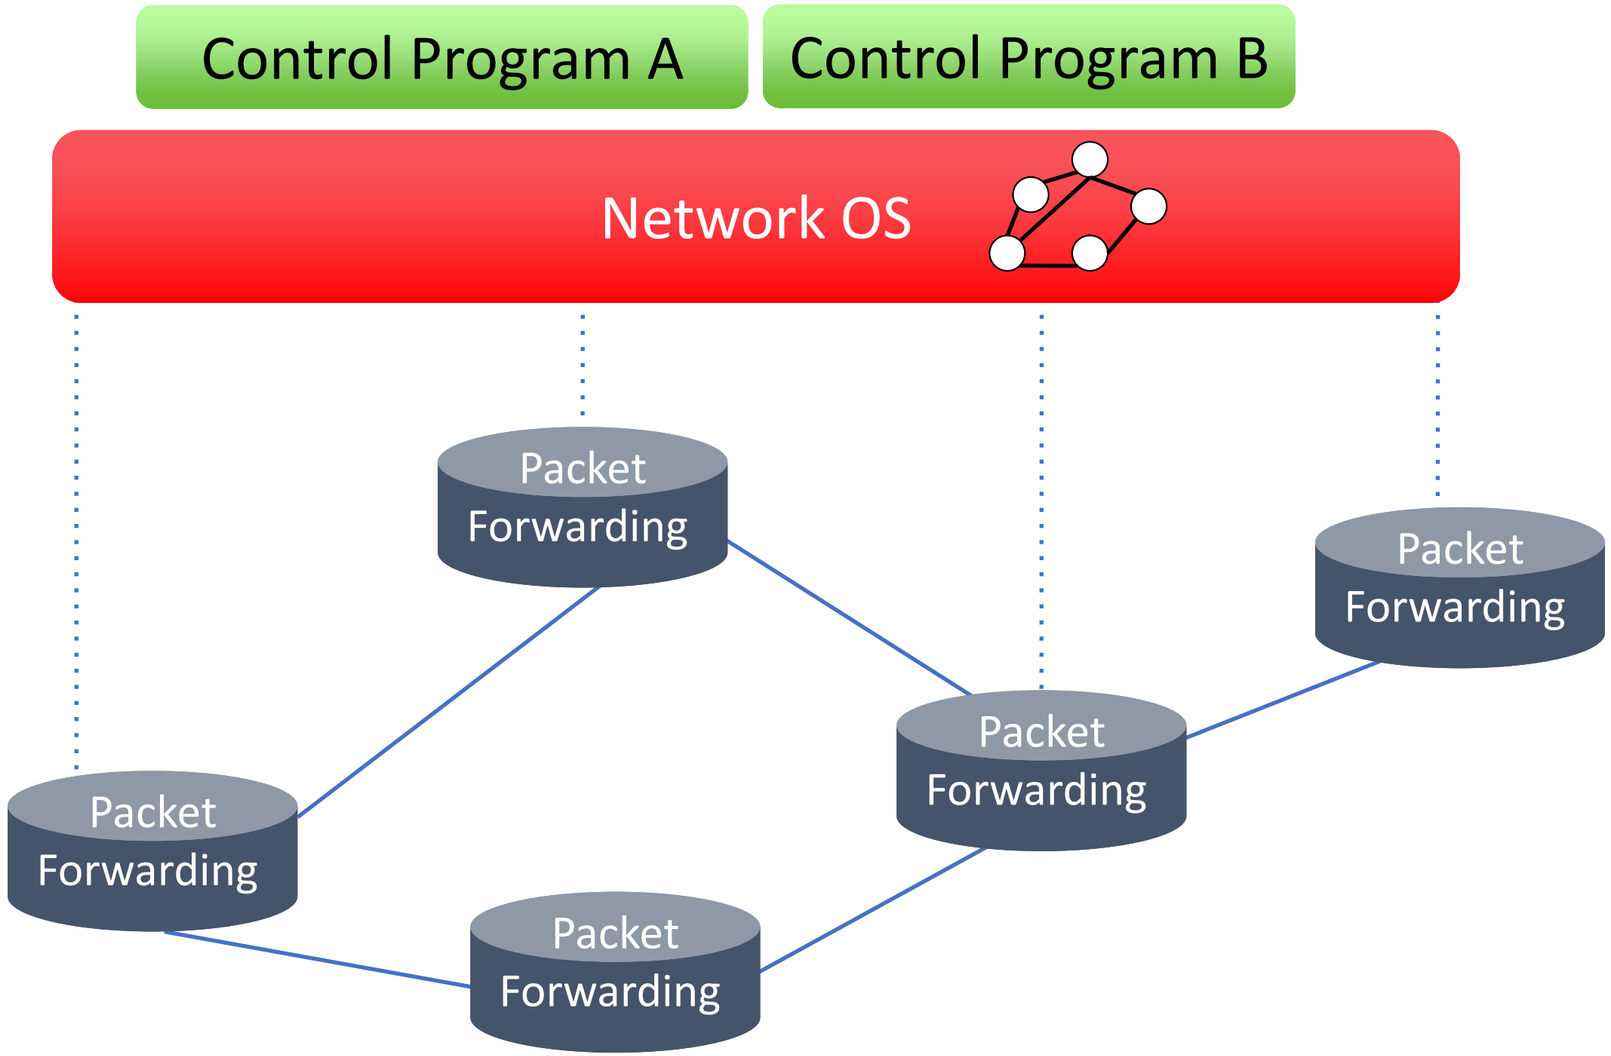

3) Forwarding Abstraction

Purpose: Abstract away forwarding hardware

- Flexible
  - Behavior specified by control plane
  - Built from basic set of forwarding primitives
- Minimal
  - Streamlined for speed and low-power
  - Control program not vendor-specific

- OpenFlow is an example of such an abstraction

4) Forwarding Substrate

- Flow-based (next slides)

- Small number of actions for each flow
  - Plumbing: Forward to port(s)
  - Control: Forward to controller
  - Routing between flow-spaces: Rewrite header
  - Bandwidth isolation: Min/max rate

- External open API to flow-table

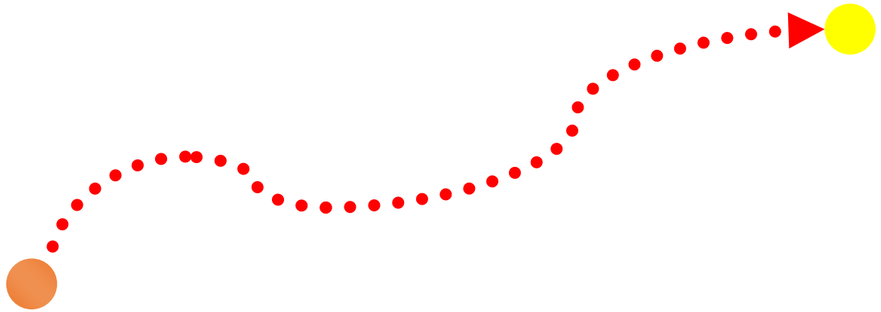

What is a flow?

- Application flow
- All http
- Jim’s traffic
- All packets to Canada
- …

Types of action

- Allow/deny flow
- Route & re-route flow
- Isolate flow
- Make flow private
- Remove flow

Flowspace

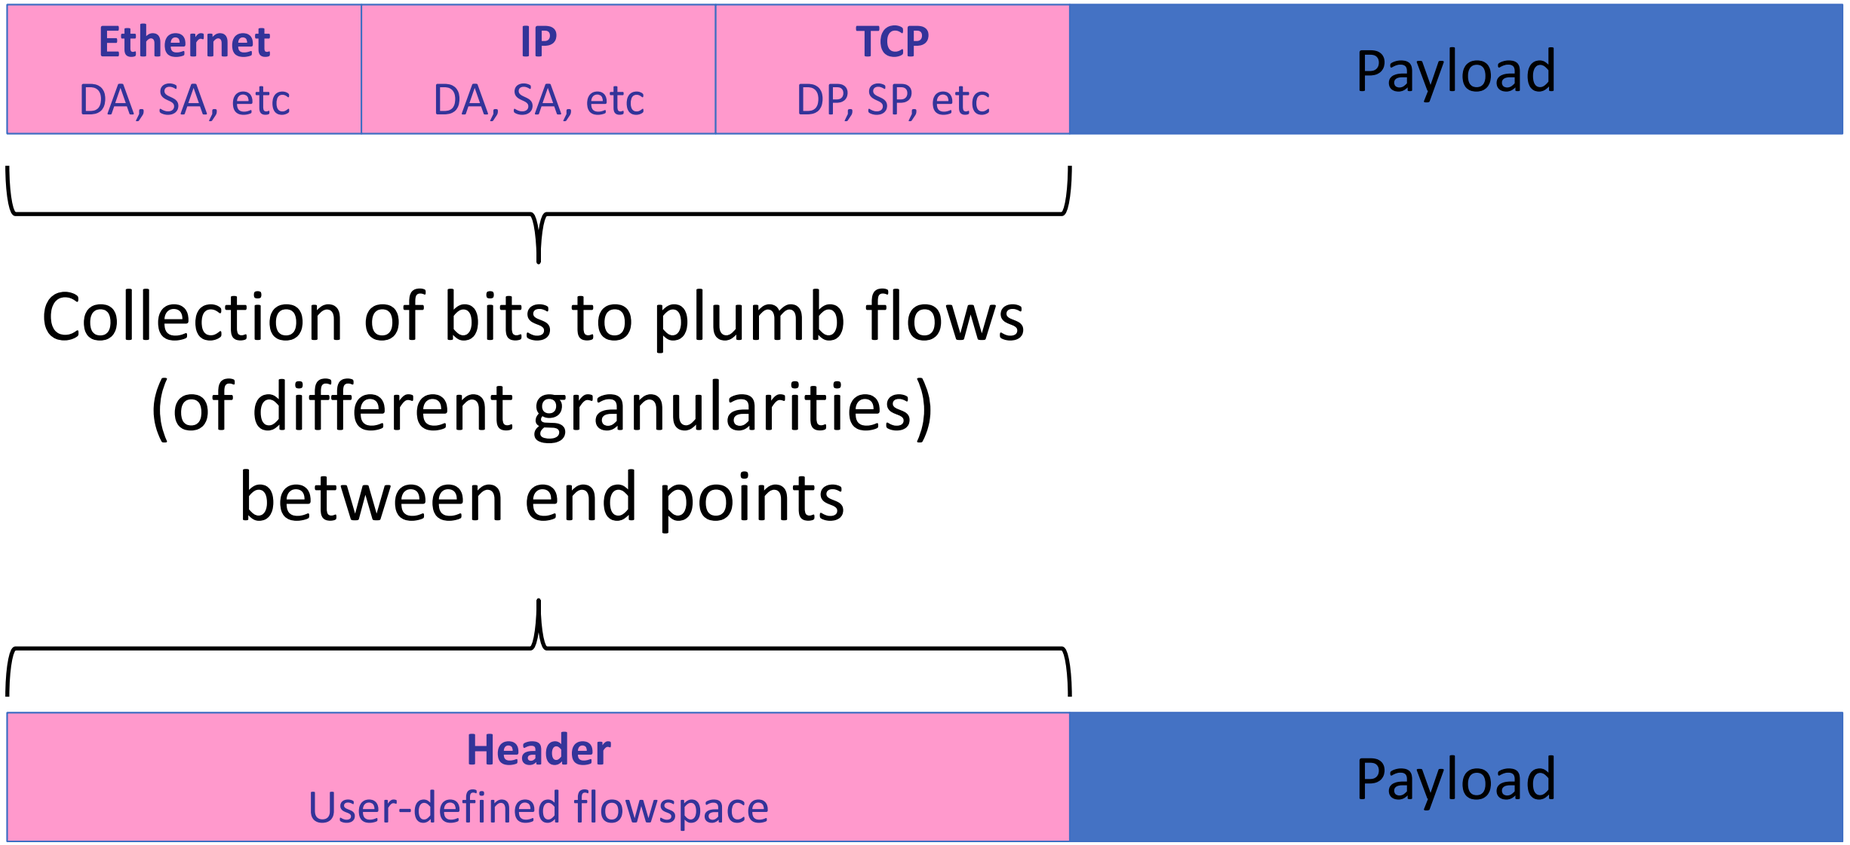
Ethernet
DA, SA, etc
IP
DA, SA, etc
TCP
DP, SP, etc
Payload
Collection of bits to plumb flows
(of different granularities)
between end points
Header
User-defined flowspace
Payload

OpenFlow

- Open standard to run experimental protocols in production networks
  - API between the forwarding elements and the network OS
- Based in Stanford, supported by various companies (Cisco, Juniper, HP, NEC, …)
- Used by universities to deploy innovative networking technology

Traditional Switch

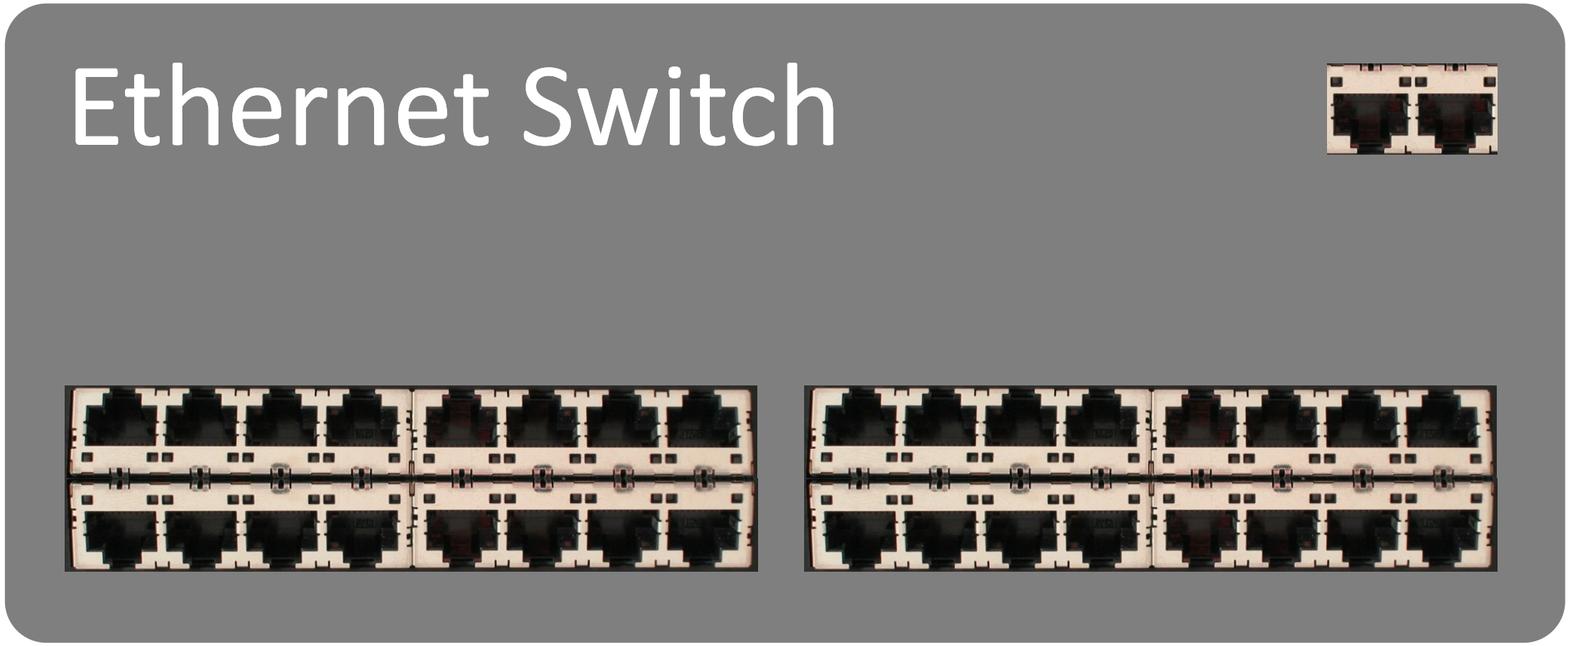

Traditional Switch

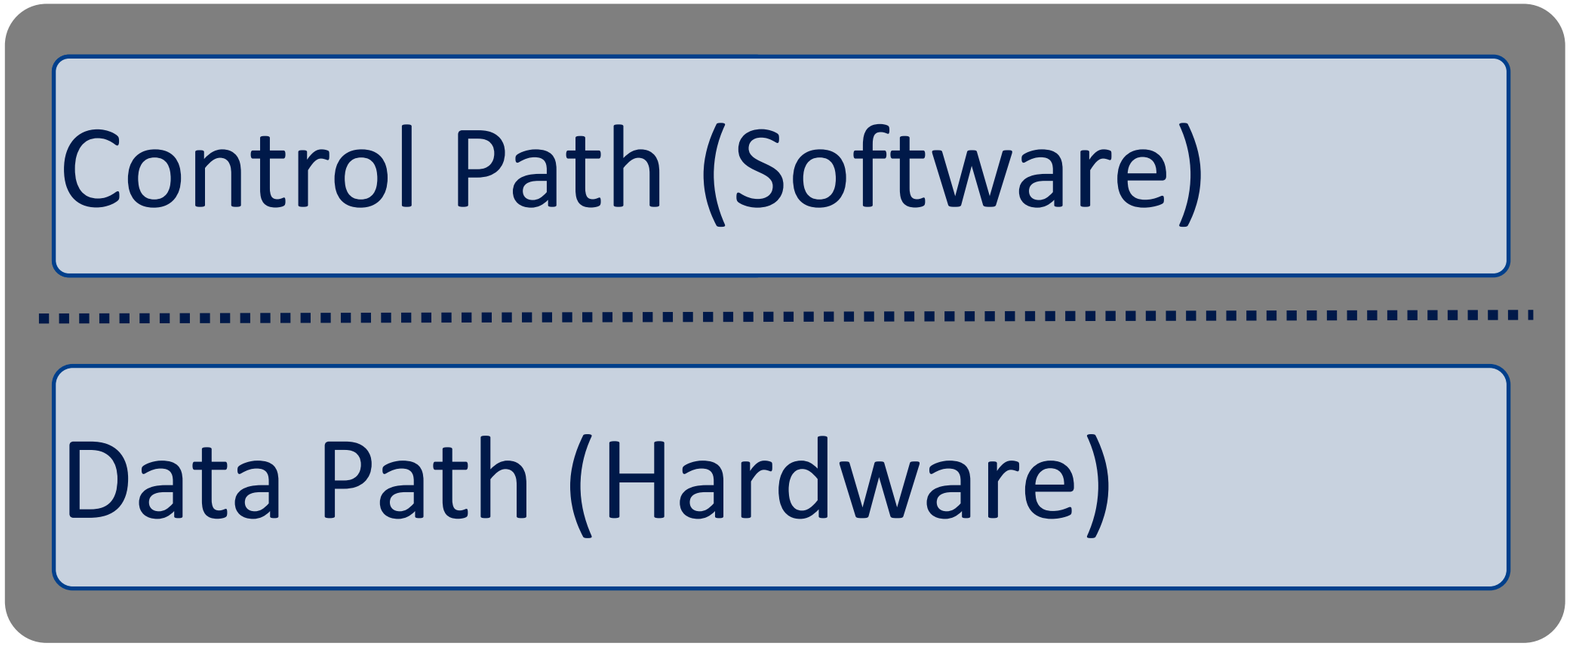

OpenFlow Switch

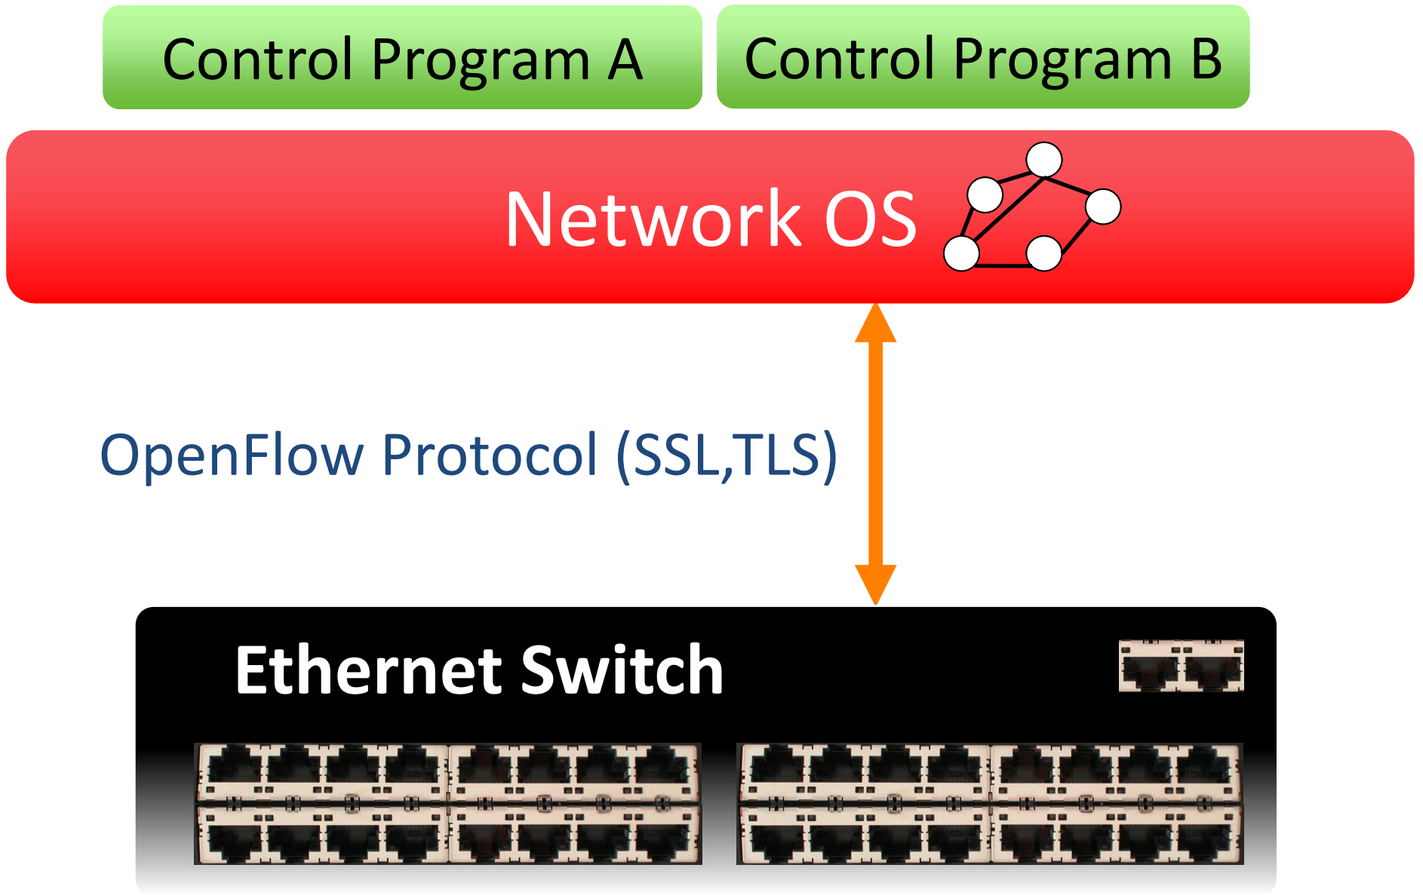

OpenFlow Rules

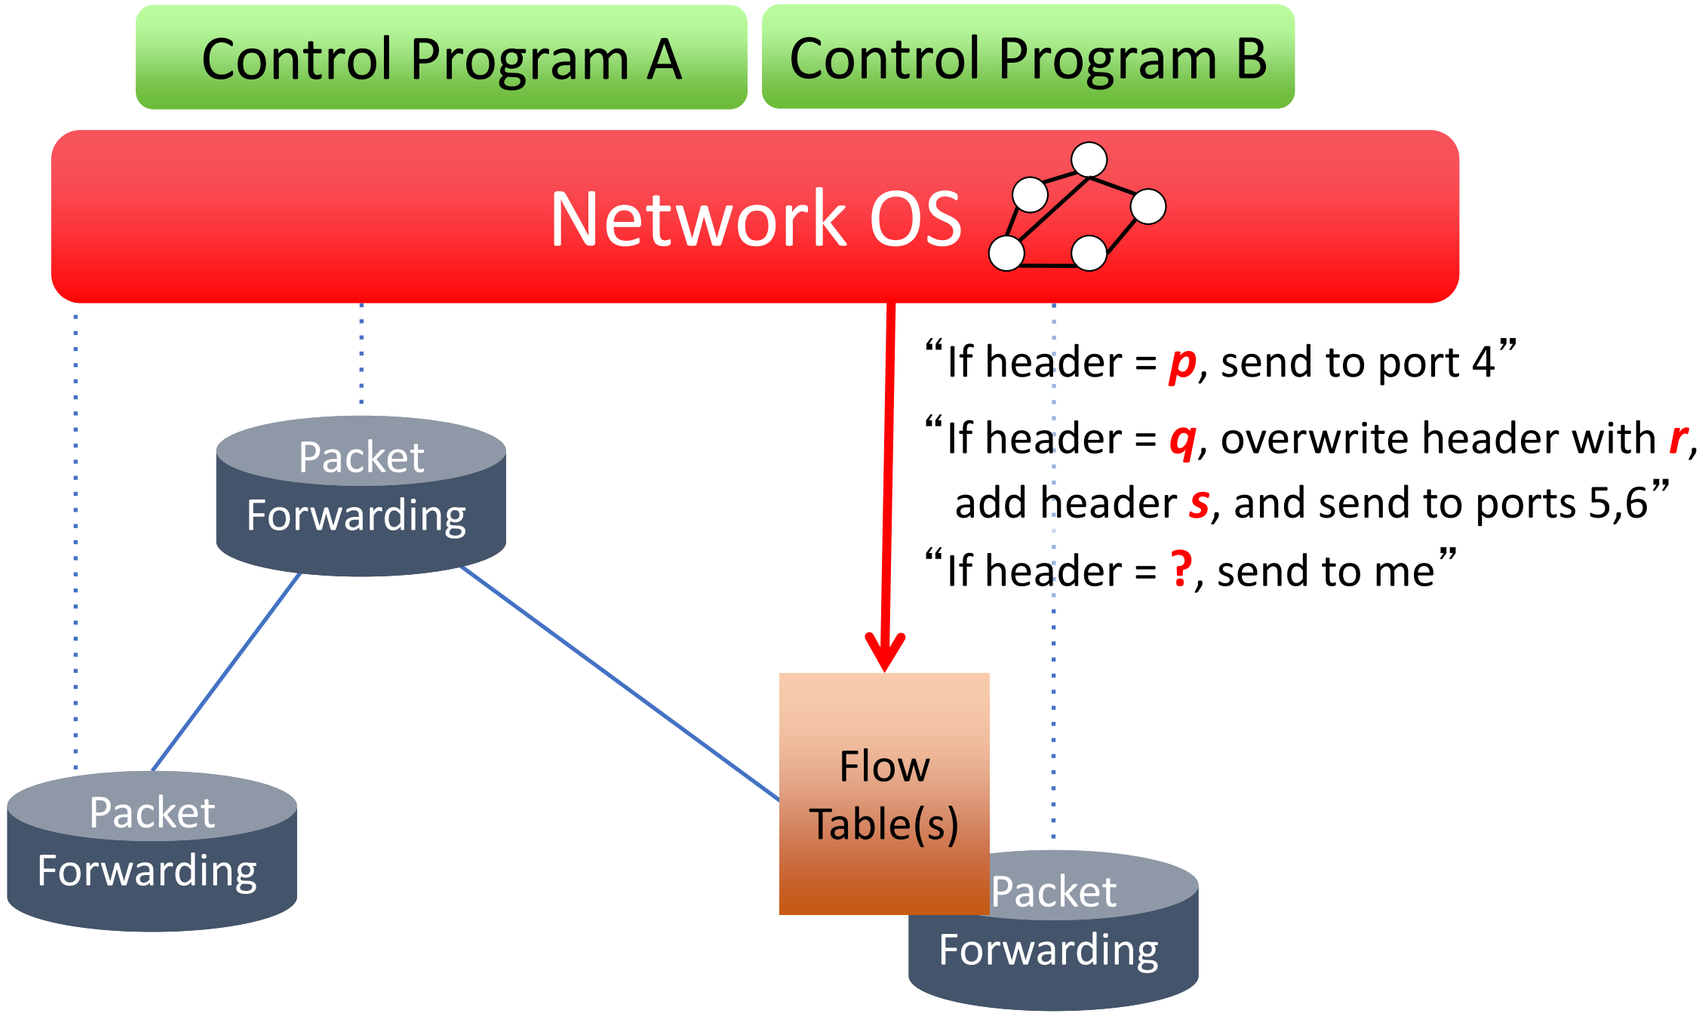

Plumbing Primitives

<Match, Action>

Match arbitrary bits in headers:

Match: 1000x01xx0101001x

- Match on any header, or new header
- Allows any flow granularity

Action

- Forward to port(s), drop, send to controller
- Overwrite header with mask, push or pop
- Forward at specific bit-rate

OpenFlow Rules – Cont'd

- Exploit the flow table in switches, routers, and chipsets

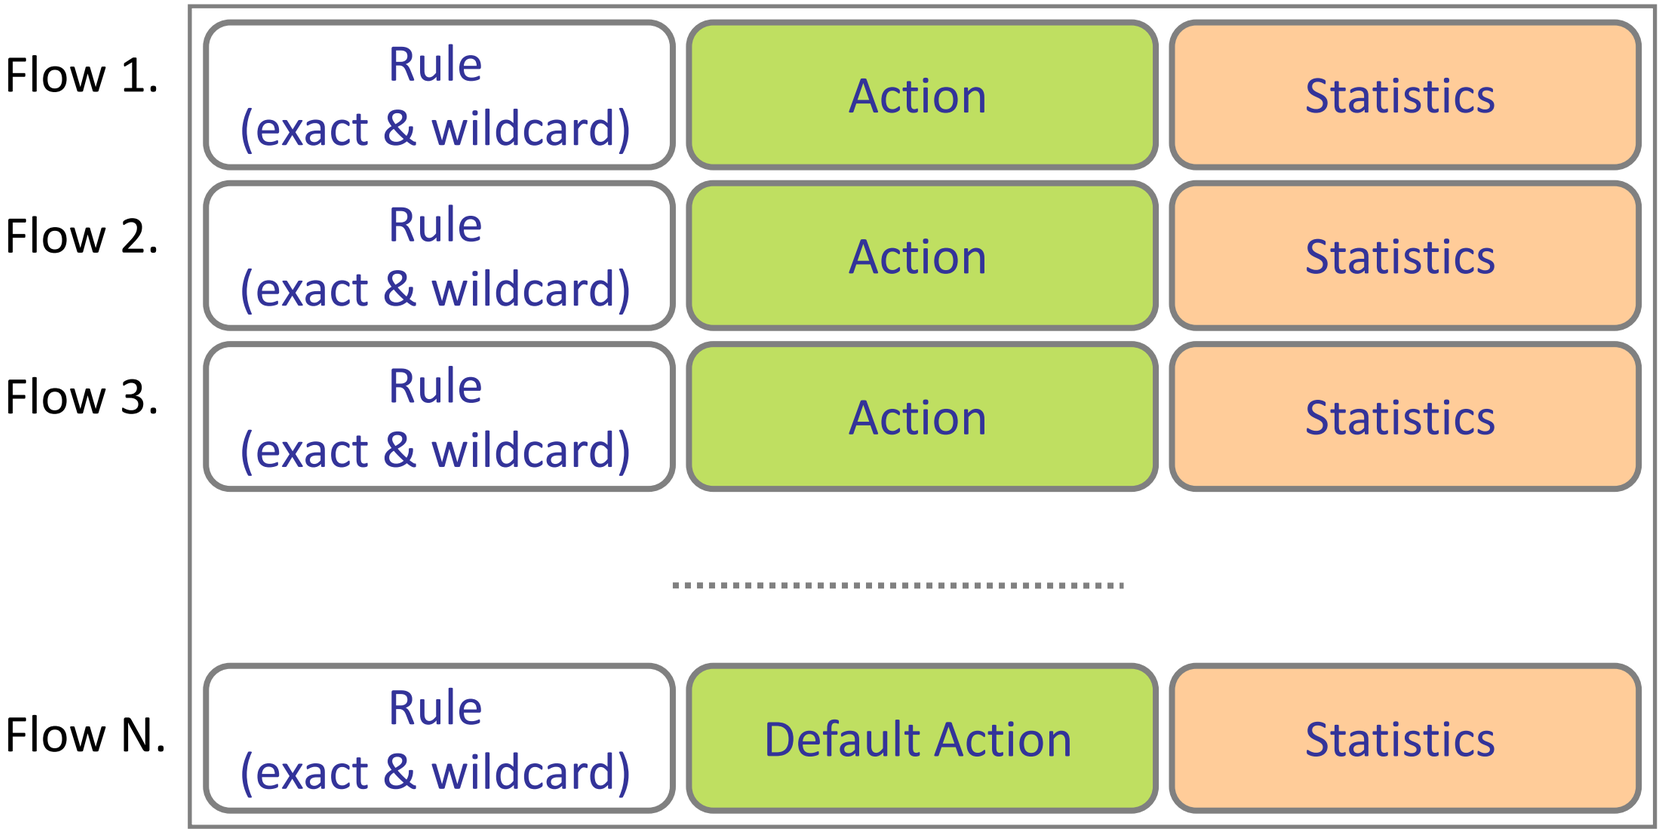

Flow Table Entry

OpenFlow Protocol Version 1.0

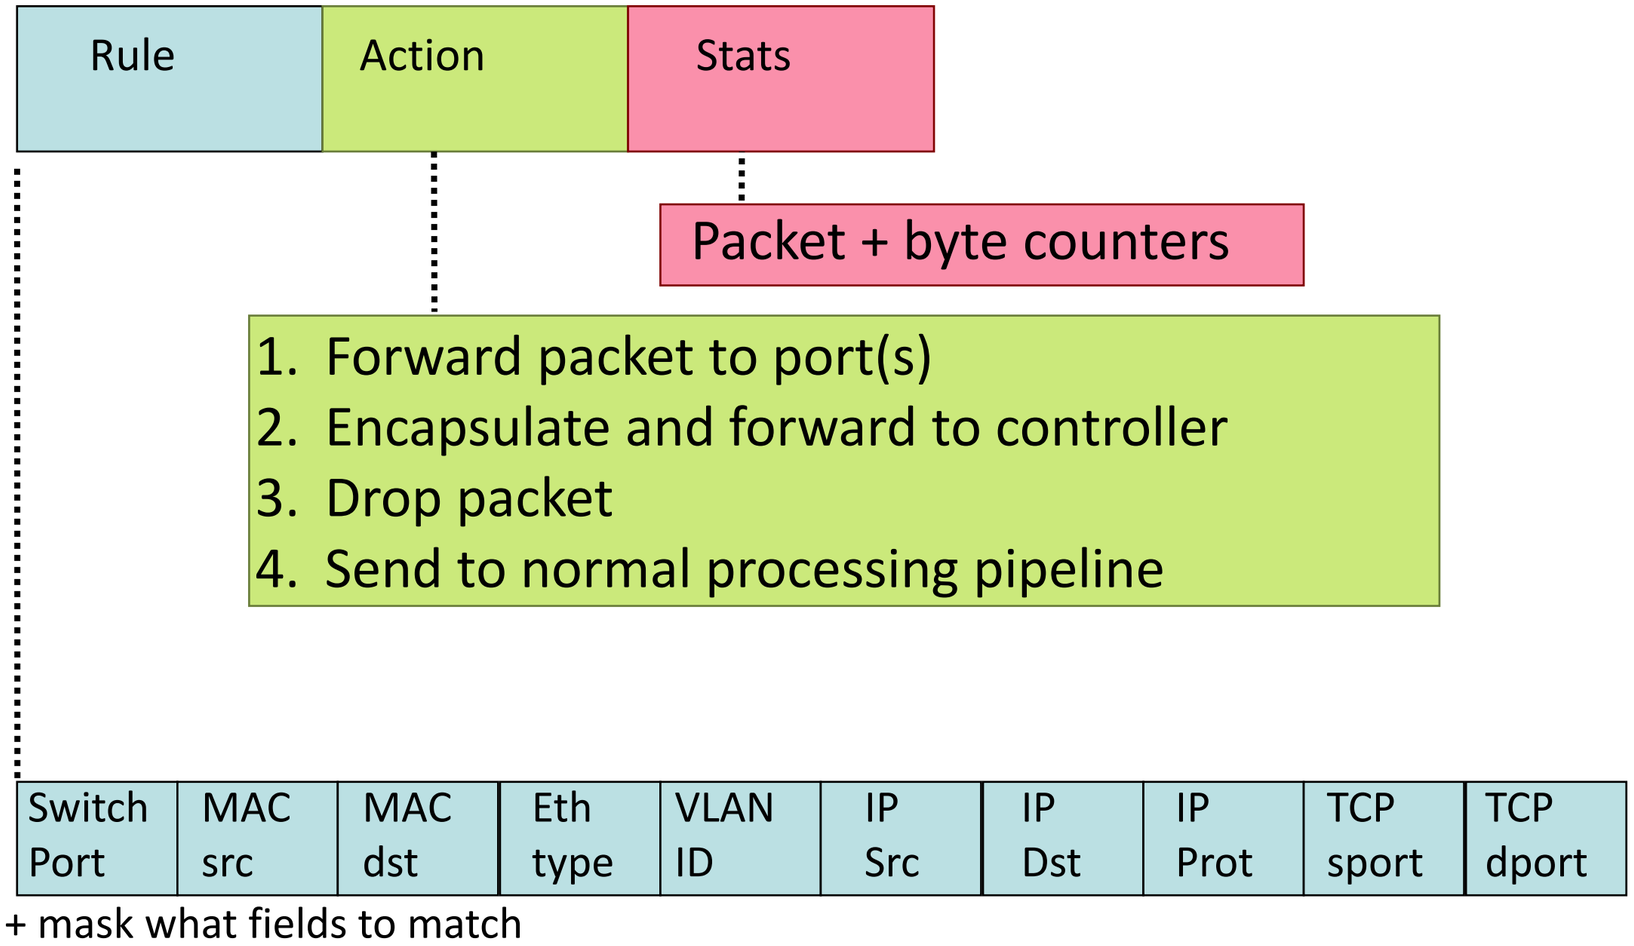

Examples

Switching

Switch Port	MAC src	MAC dst	Eth type	VLAN ID	IP Src	IP Dst	IP Prot	TCP sport	TCP dport	Action
*	*	00:1f:..	*	*	*	*	*	*	*	port6

Flow Switching

Switch Port	MAC src	MAC dst	Eth type	VLAN ID	IP Src	IP Dst	IP Prot	TCP sport	TCP dport	Action
port3	00:2e..	00:1f..	0800	vlan1	1.2.3.4	5.6.7.8	4	17264	80	port6

Firewall

Switch Port	MAC src	MAC dst	Eth type	VLAN ID	IP Src	IP Dst	IP Prot	TCP sport	TCP dport	Forward
*	*	*	*	*	*	*	*	*	22	drop

Examples

Routing

Switch Port	MAC src	MAC dst	Eth type	VLAN ID	IP Src	IP Dst	IP Prot	TCP sport	TCP dport	Action
*	*	*	*	*	*	5.6.7.8	*	*	*	port6

VLAN

Switch Port	MAC src	MAC dst	Eth type	VLAN ID	IP Src	IP Dst	IP Prot	TCP sport	TCP dport	Action
*	*	*	*	vlan1	*	*	*	*	*	port6, port7,port9

OpenFlow Hardware

Juniper MX-series

NEC IP8800

WiMax (NEC)

HP Procurve 5400

Cisco Catalyst 6k

PC Engines

Quanta LB4G

More ...

Some fun facts …

- VMware purchased Nicira for $1.26 billion in 2012
  - VMware Purchased VeloCloud Networks in 2017 (amount unknown!)
- Orion (2022)
  - Google's 2nd generation SDN control plane
  - Responsible for configuration, management, real-time network control
  - In all their data center, campus, and private Wide Area (B4) networks
  - In production for over 4 years

OpenFlow Usage Example

Dedicated OpenFlow Network

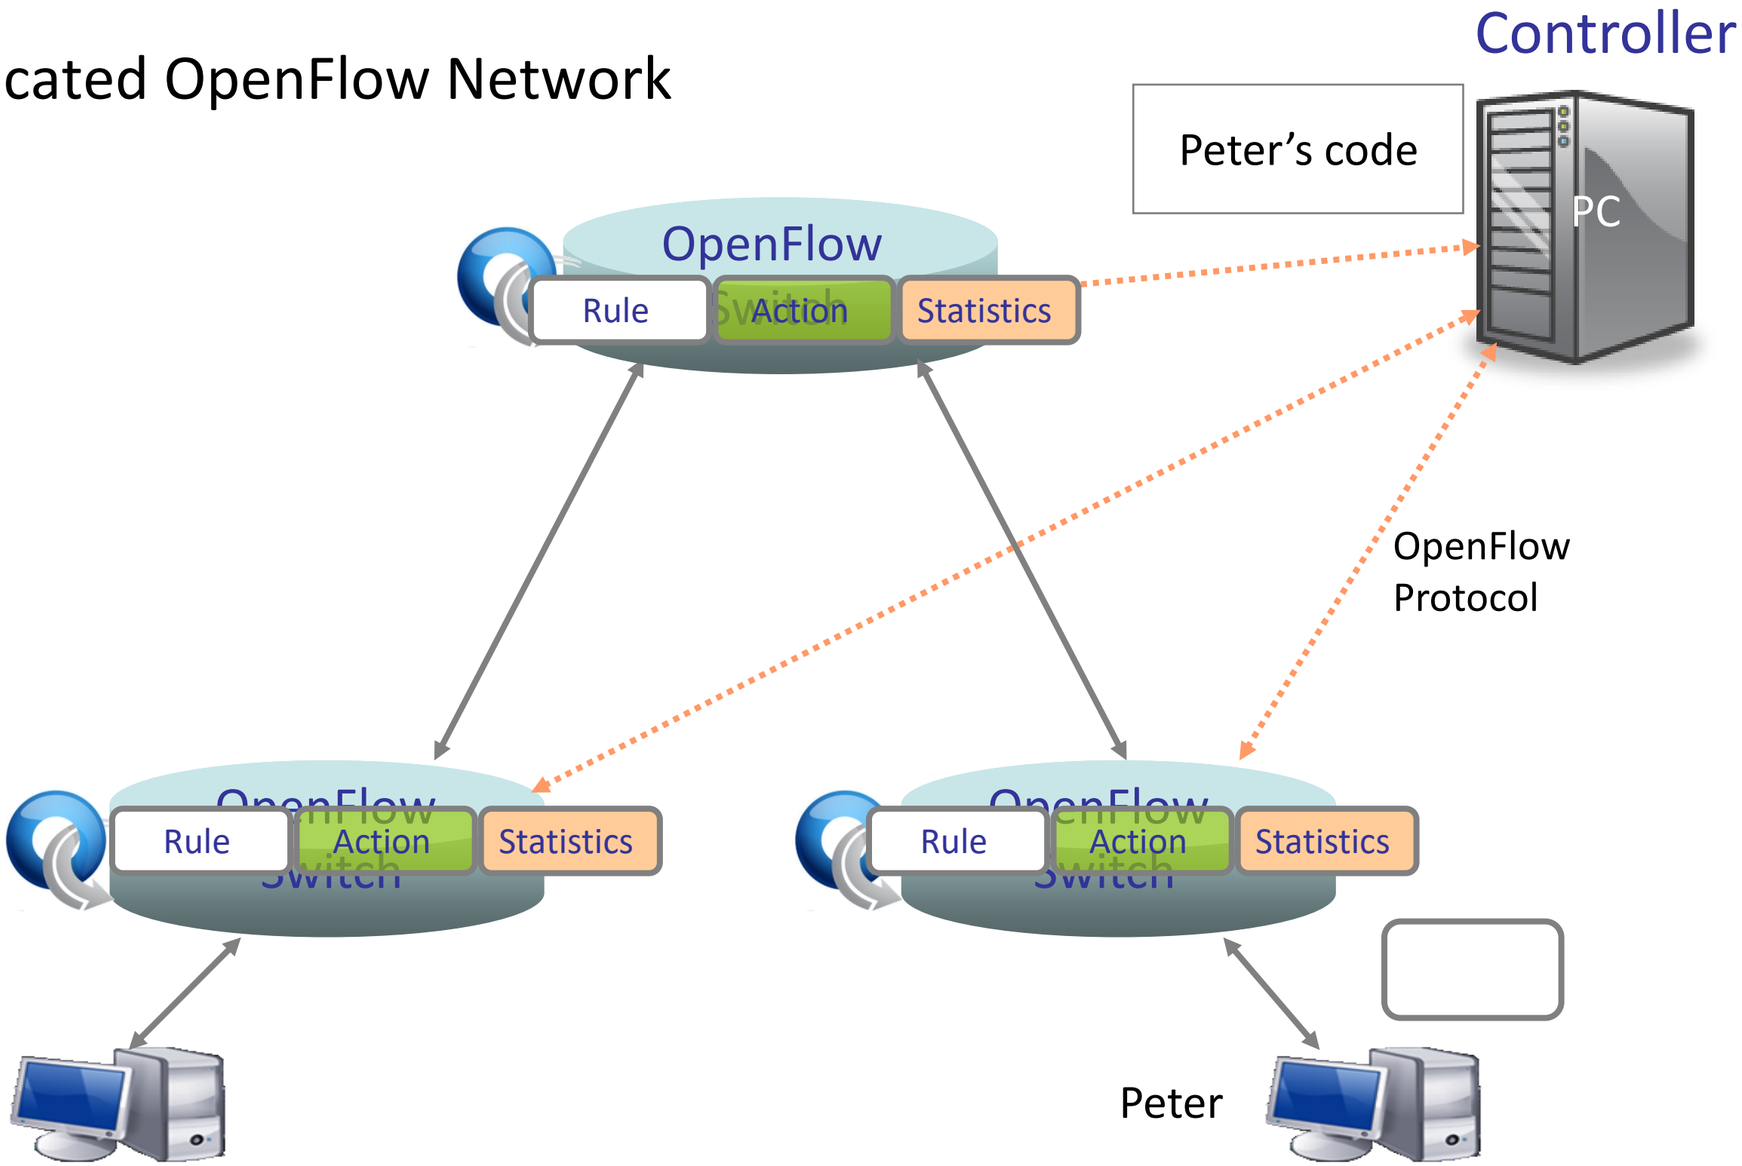

Usage examples

- Peter’s code:
  - Static “VLANs”
  - His own new routing protocol: unicast, multicast, multipath, load-balancing
  - Network access control
  - Home network manager
  - Mobility manager
  - Energy manager
  - Packet processor (in controller)
  - IPvPeter
  - Network measurement and visualization
  - …

Another Example:

Production networks as Research Test-beds!

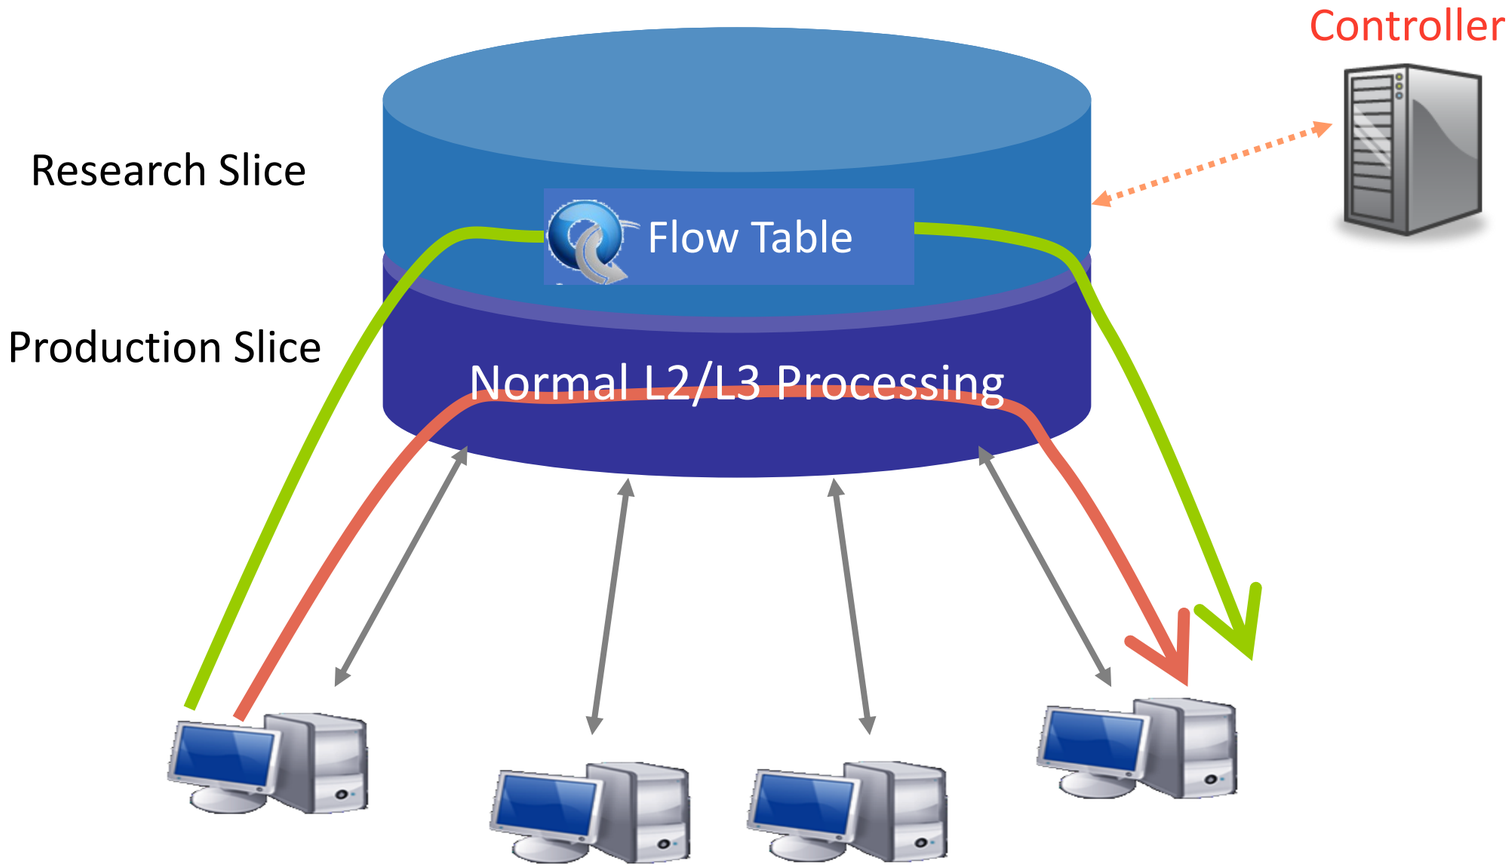

Virtualize OpenFlow Switch

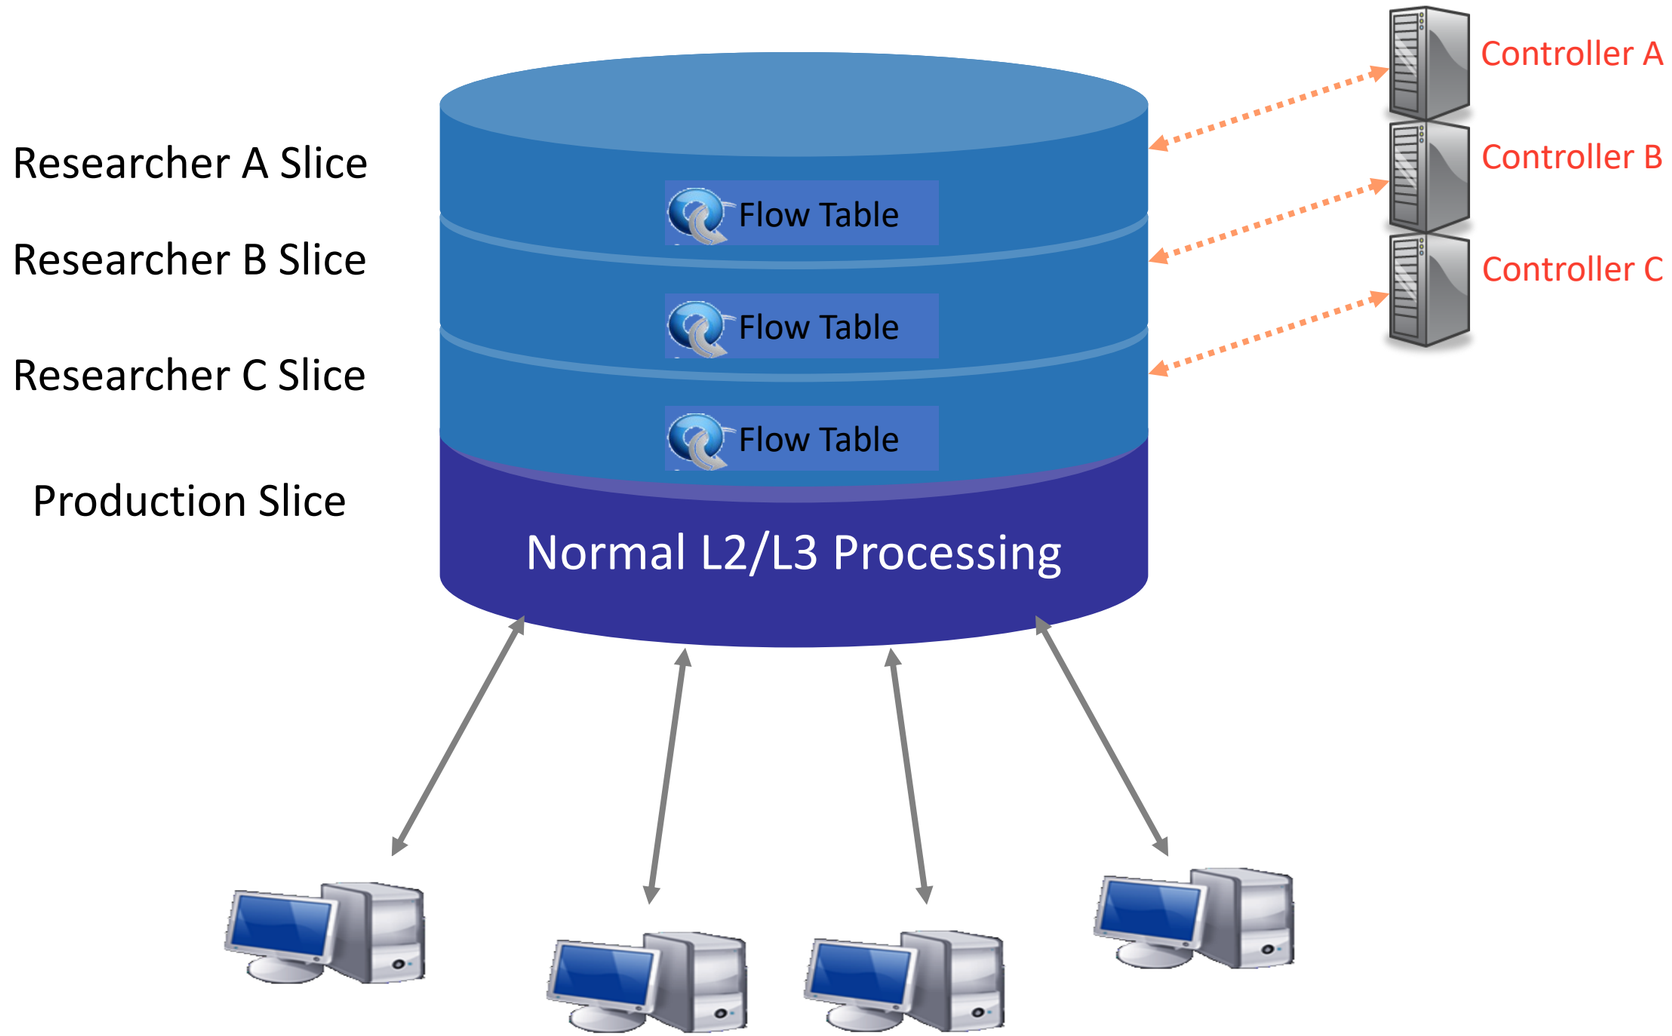

Virtualizing OpenFlow

FlowVisor OSDI'10

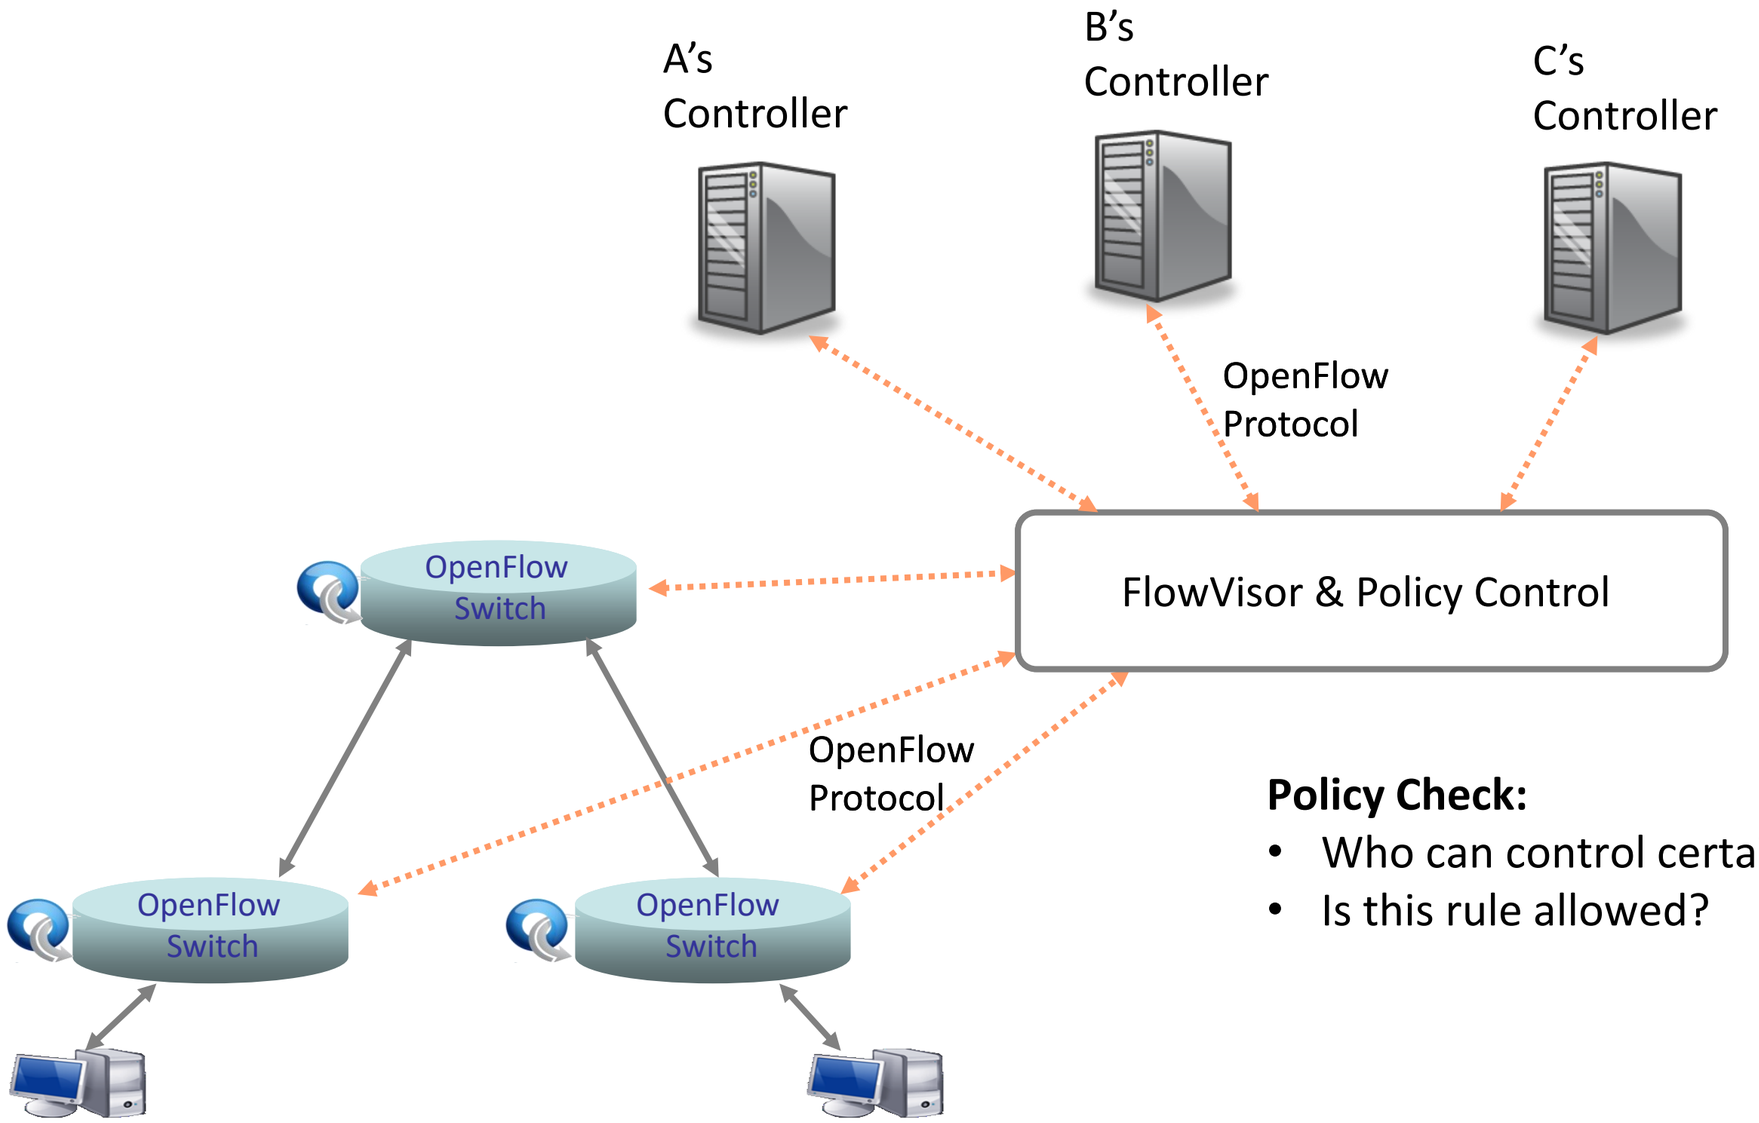

Policy Check:

- Who can control certain packets?
- Is this rule allowed?

Virtualizing OpenFlow

FlowVisor OSDI'10

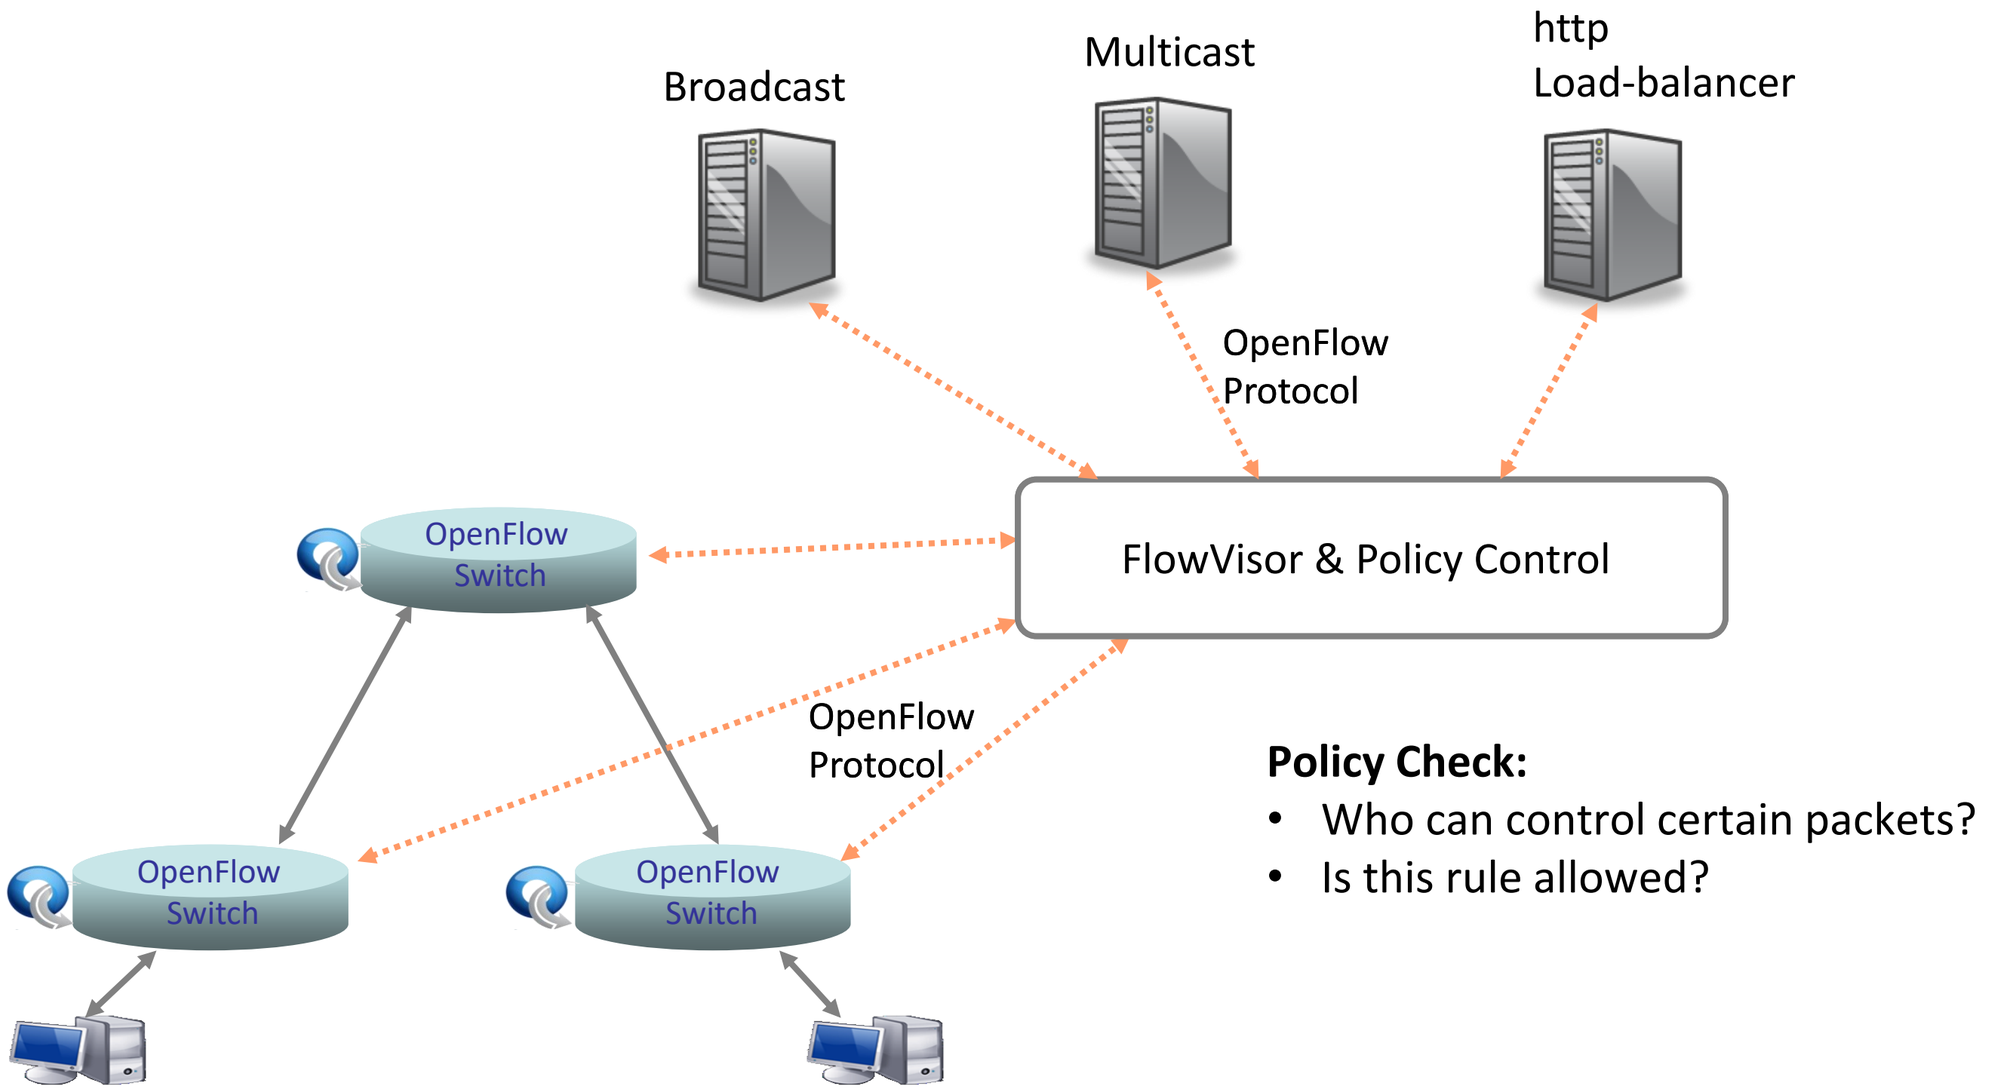

Policy Check:

- Who can control certain packets?
- Is this rule allowed?

Food for Thought

- What are the challenges in switching from traditional networks to OpenFlow networks?

- What are the opportunities?