Software-defined networks (SDN)

PA160: Net-Centric Computing II.

Tomáš Rebok

Institute of Computer Science Masaryk University

rebok@ics.muni.cz

Motivation

Traditional Computing vs Modern Computing

Vertically integrated
Closed, proprietary
Slow innovation
Small industry

Traditional Computing vs Modern Computing

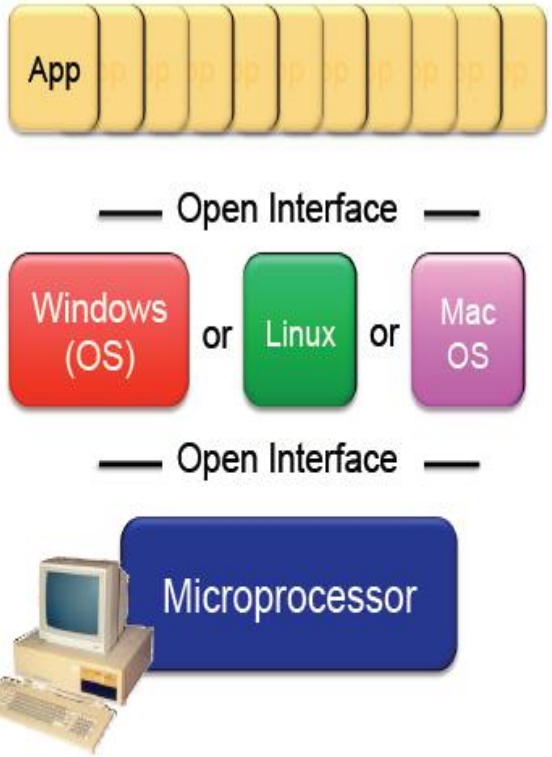

Vertically integrated
Closed, proprietary
Slow innovation
Small industry

Horizontal
Open interfaces
Rapid innovation
Huge industry

Traditional Computing vs Modern Computing

Vertically integrated
Closed, proprietary
Slow innovation
Small industry

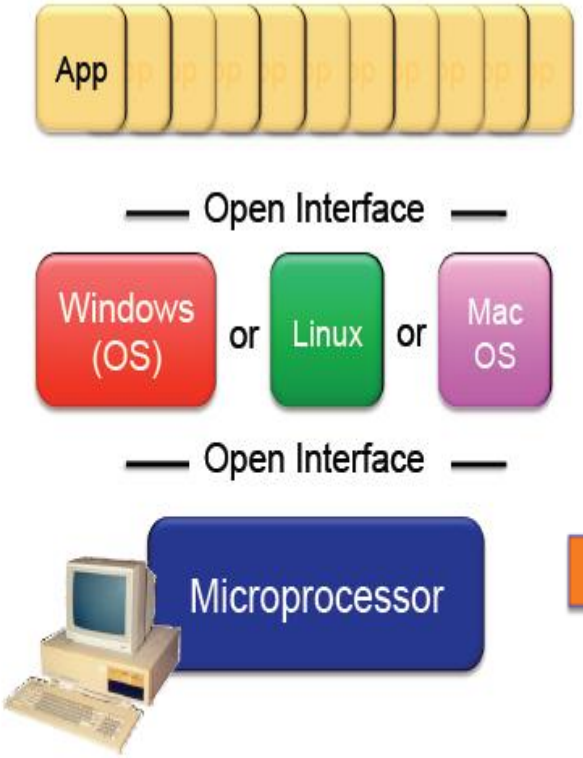

Horizontal
Open interfaces
Rapid innovation
Huge industry

Traditional vs Modern Computing Provisioning Methods

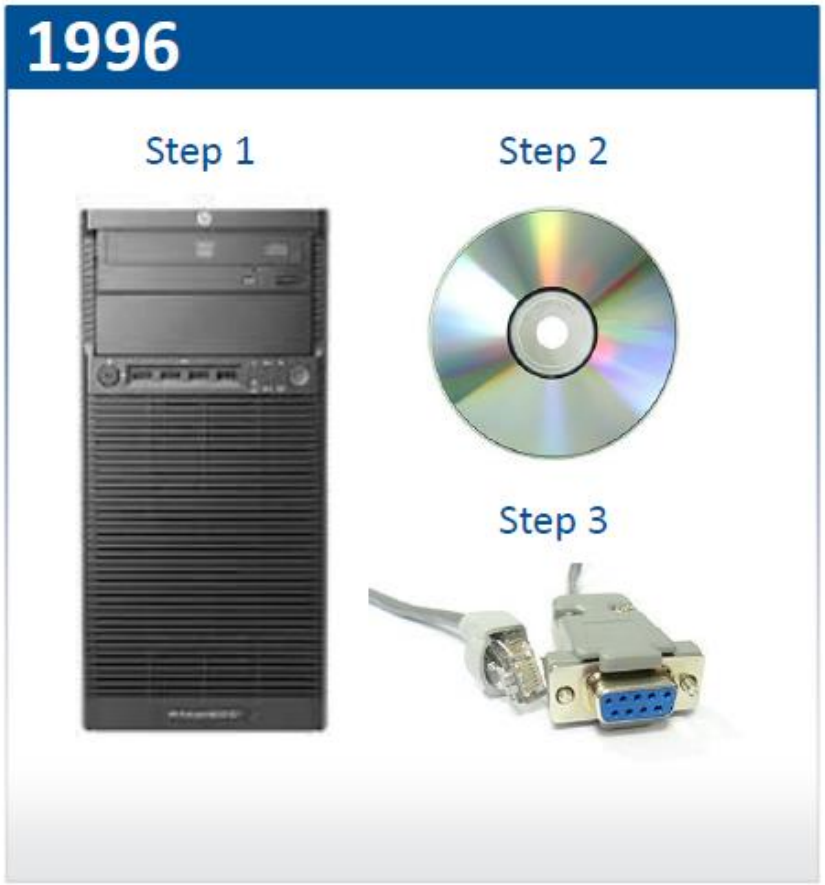

Source: Adopted from Transforming the Network With Open SDN by Big Switch Network

Modern Networking Complexity

Vertically integrated
Closed, proprietary
Slow innovation

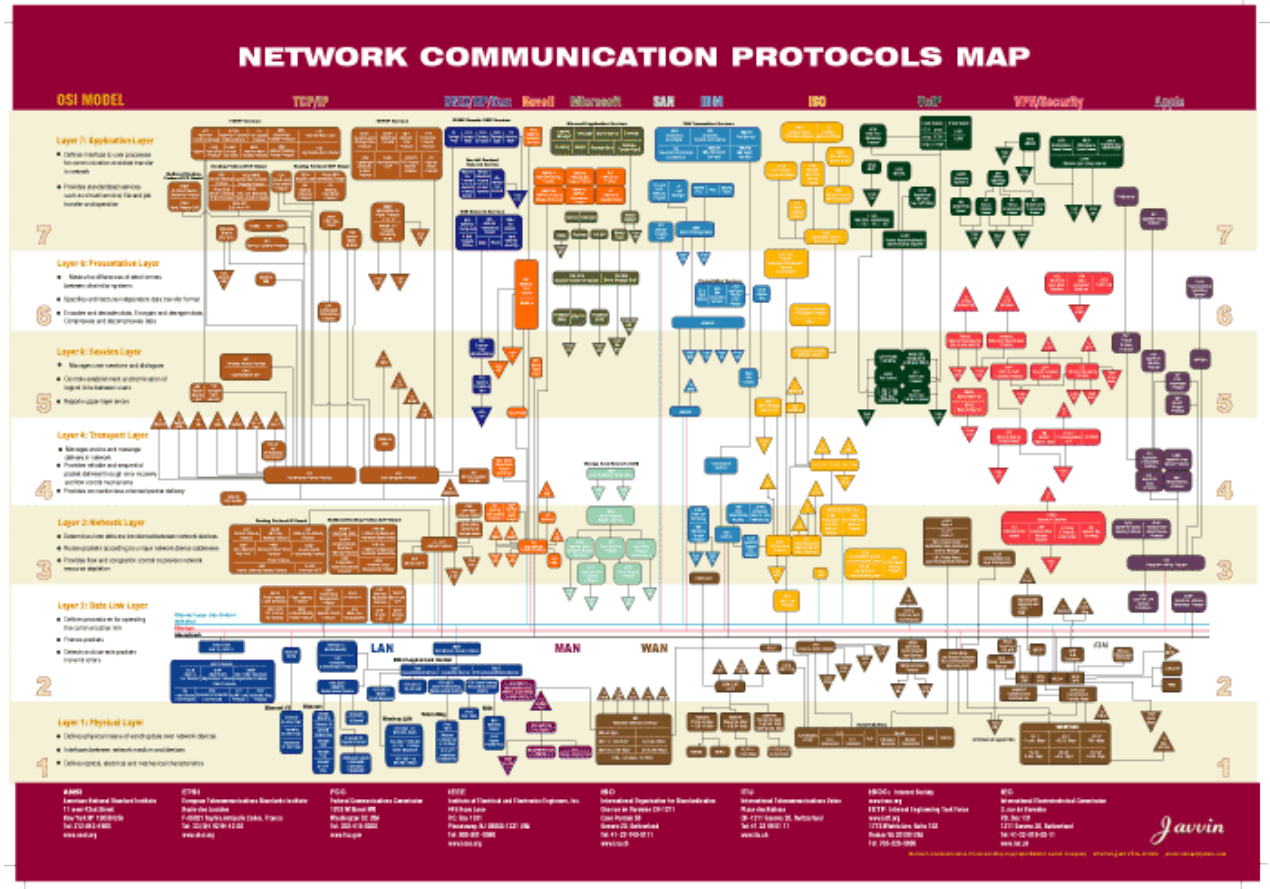

Ref: Javvin

The Ossified Network

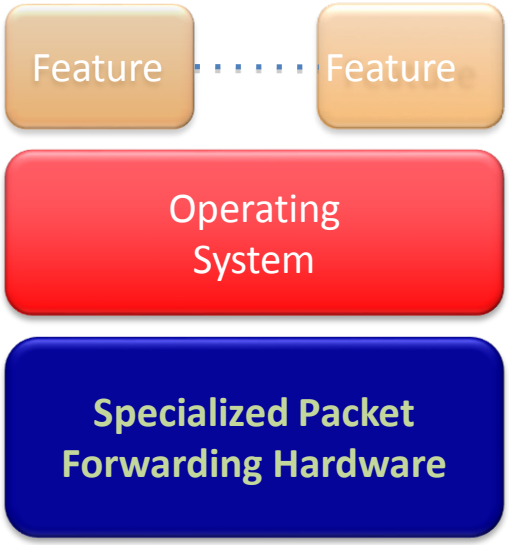

The Ossified Network

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

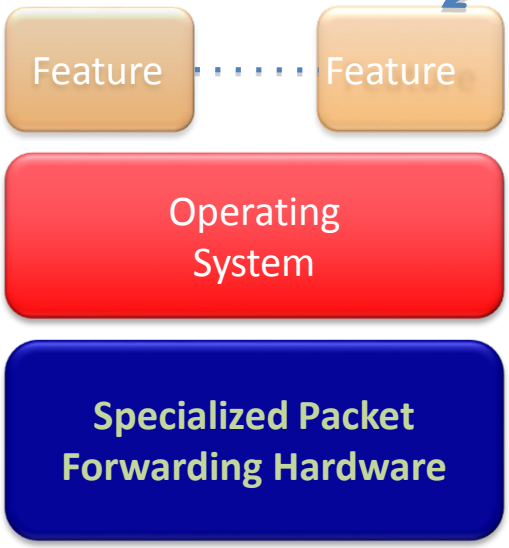

The Ossified Network

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

Feature ..... Feature

Operating System

Million of lines of source code

6000+ RFCs

Barrier to entry

Specialized Packet Forwarding Hardware

Billions of gates

Bloated

Power Hungry

The Ossified Network

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

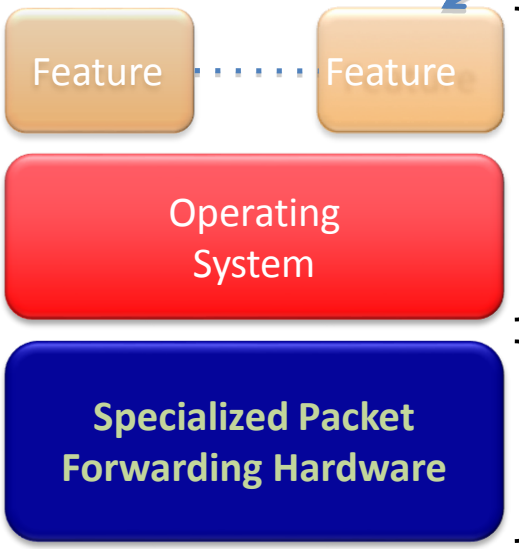

Million of lines of source code

6000+ RFCs

Barrier to entry

Billions of gates

Bloated

Power Hungry

Many complex functions baked into the infrastructure

OSPF, BGP, multicast, differentiated services,
Traffic Engineering, NAT, firewalls, MPLS, redundant layers, …

An industry with a “mainframe-mentality”, reluctant to change

Traditional vs Modern Networking Provisioning Methods

1996

```
Router> enable
Router# configure terminal
Router(config)# enable secret cisco
Router(config)# ip route 0.0.0.0 0.0.0.0 20.2.2.3
Router(config)# interface ethernet0
Router(config-if)# ip address 10.1.1.1 255.0.0.0
Router(config-if)# no shutdown
Router(config-if)# exit
Router(config)# interface serial0
Router(config-if)# ip address 20.2.2.2 255.0.0.0
Router(config-if)# no shutdown
Router(config-if)# exit
Router(config)# router rip
Router(config-router)# network 10.0.0.0
Router(config-router)# network 20.0.0.0
Router(config-router)# exit
Router(config)# exit
Router# copy running-config startup-config
Router# disable
Router>
```

Terminal Protocol: **Telnet**

2013

```
Router> enable
Router# configure terminal
Router(config)# enable secret cisco
Router(config)# ip route 0.0.0.0 0.0.0.0 20.2.2.3
Router(config)# interface ethernet0
Router(config-if)# ip address 10.1.1.1 255.0.0.0
Router(config-if)# no shutdown
Router(config-if)# exit
Router(config)# interface serial0
Router(config-if)# ip address 20.2.2.2 255.0.0.0
Router(config-if)# no shutdown
Router(config-if)# exit
Router(config)# router rip
Router(config-router)# network 10.0.0.0
Router(config-router)# network 20.0.0.0
Router(config-router)# exit
Router(config)# exit
Router# copy running-config startup-config
Router# disable
Router>
```

Terminal Protocol: **SSH**

Source: Adopted from Transforming the Network With Open SDN by Big Switch Network

Computing vs Networking

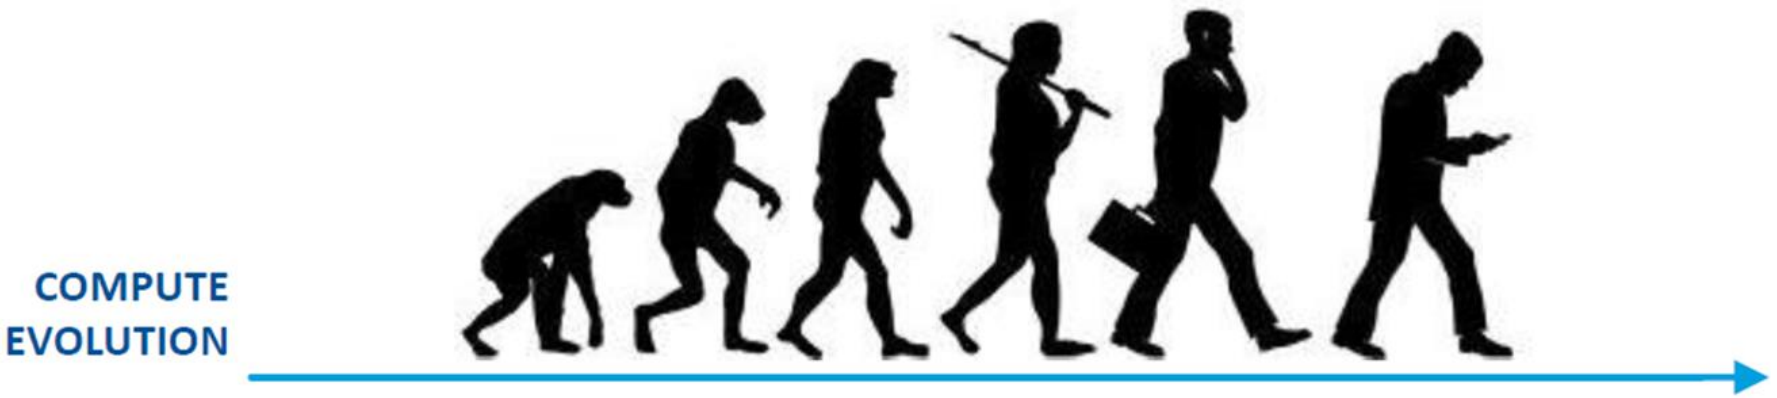

Source: Adopted from Transforming the Network With Open SDN by Big Switch Network

Computing vs Networking

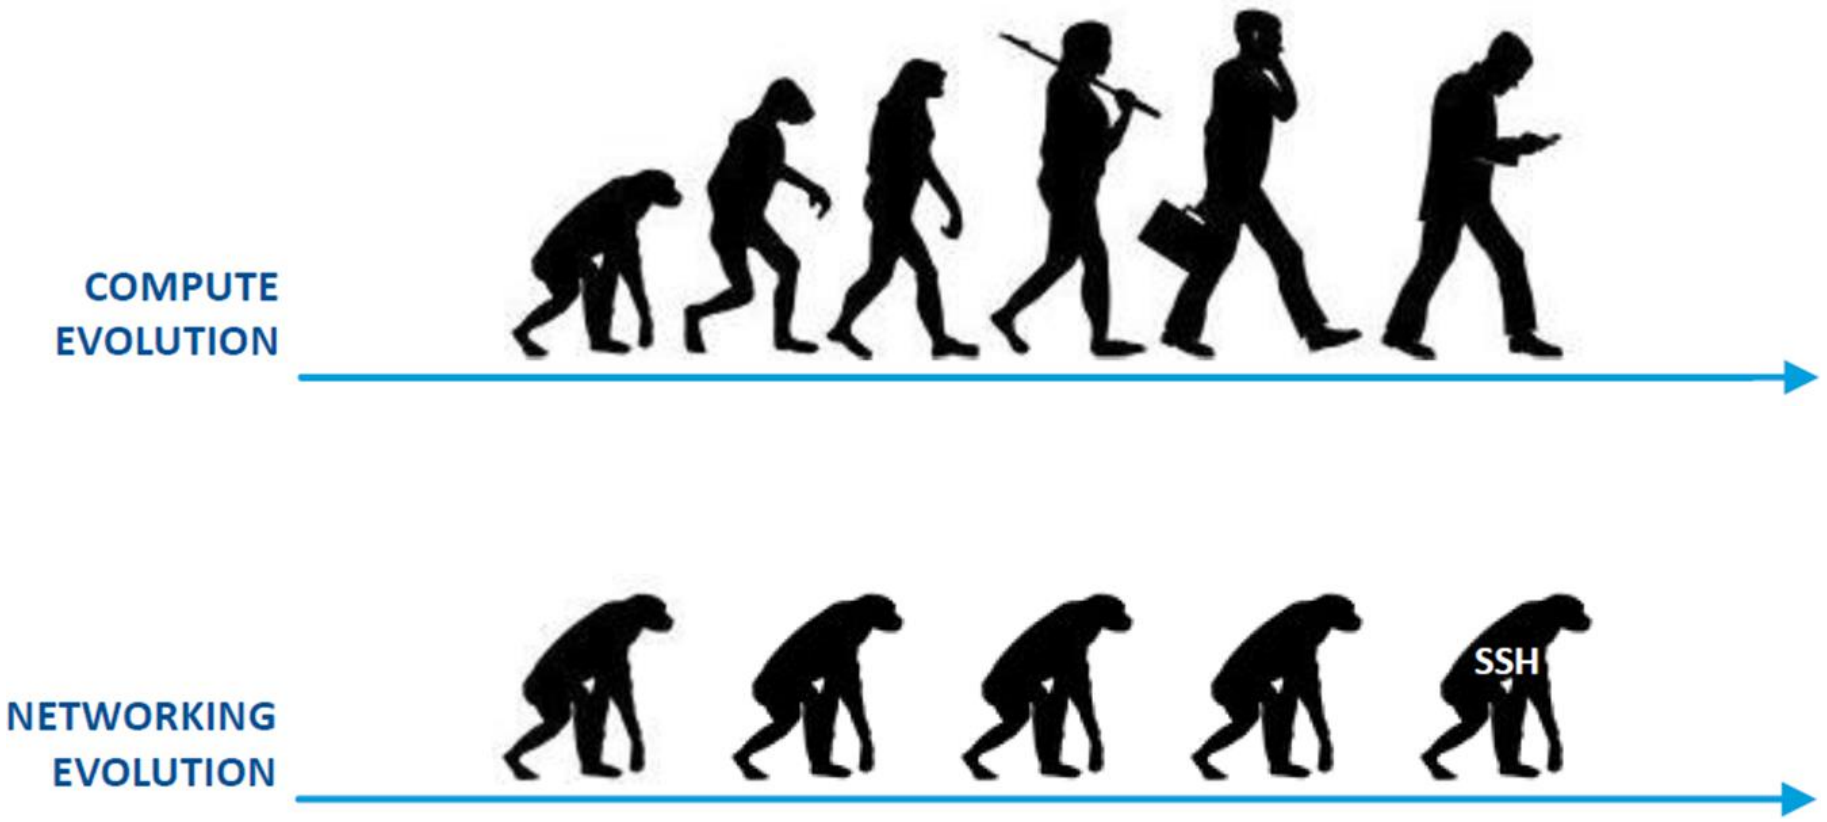

Source: Adopted from Transforming the Network With Open SDN by Big Switch Network

Problems in Networking

- Networks must keep up with exponential increases in traffic and more and more individually managed networked devices
- The result is more networking devices and strain on operations teams (who struggle to provide business value)

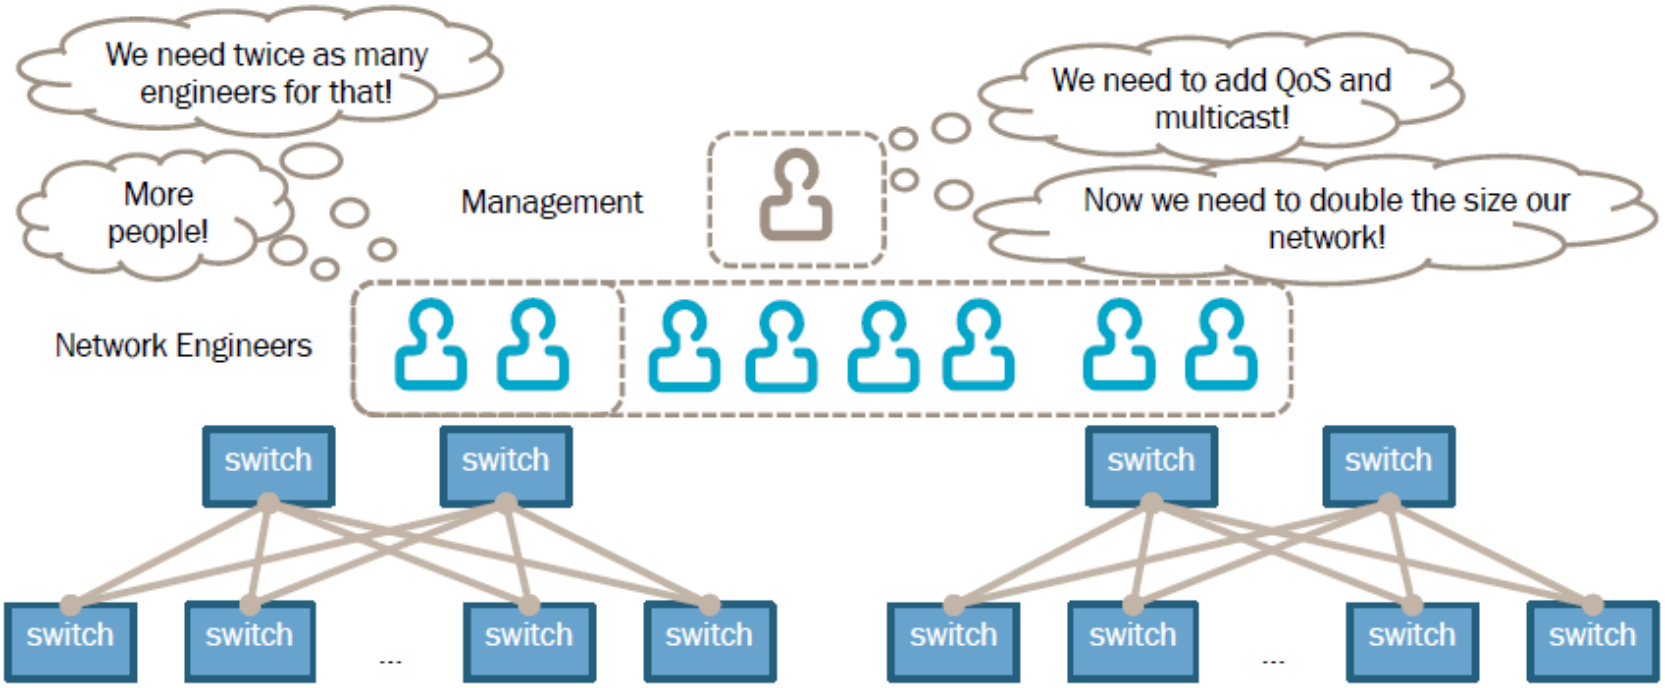

Problems in Networking

SDN Essentials

- Networking is highly prescriptive yet networks are consumed in intents
- There are few (if any) abstractions in traditional networking to hide prescriptive details
- Network details must be exposed to and understood by consumers

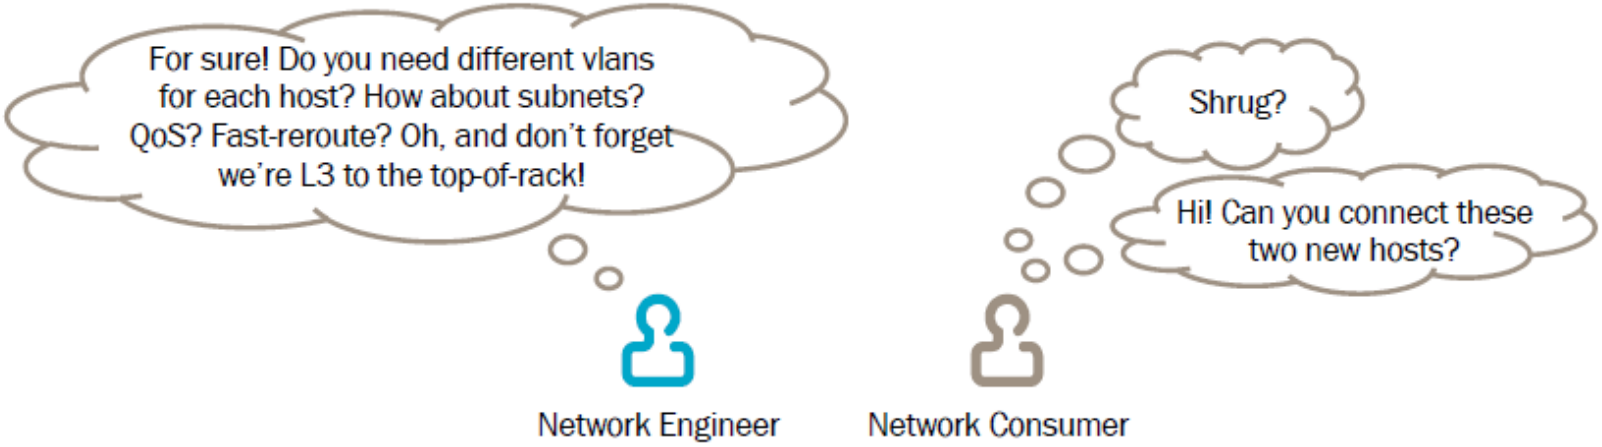

Problems in Networking

- All elements of the traditional networking stack are tightly coupled (read glued together)
- Customers have little choice in selecting elements/ hardware/software for their specific use cases

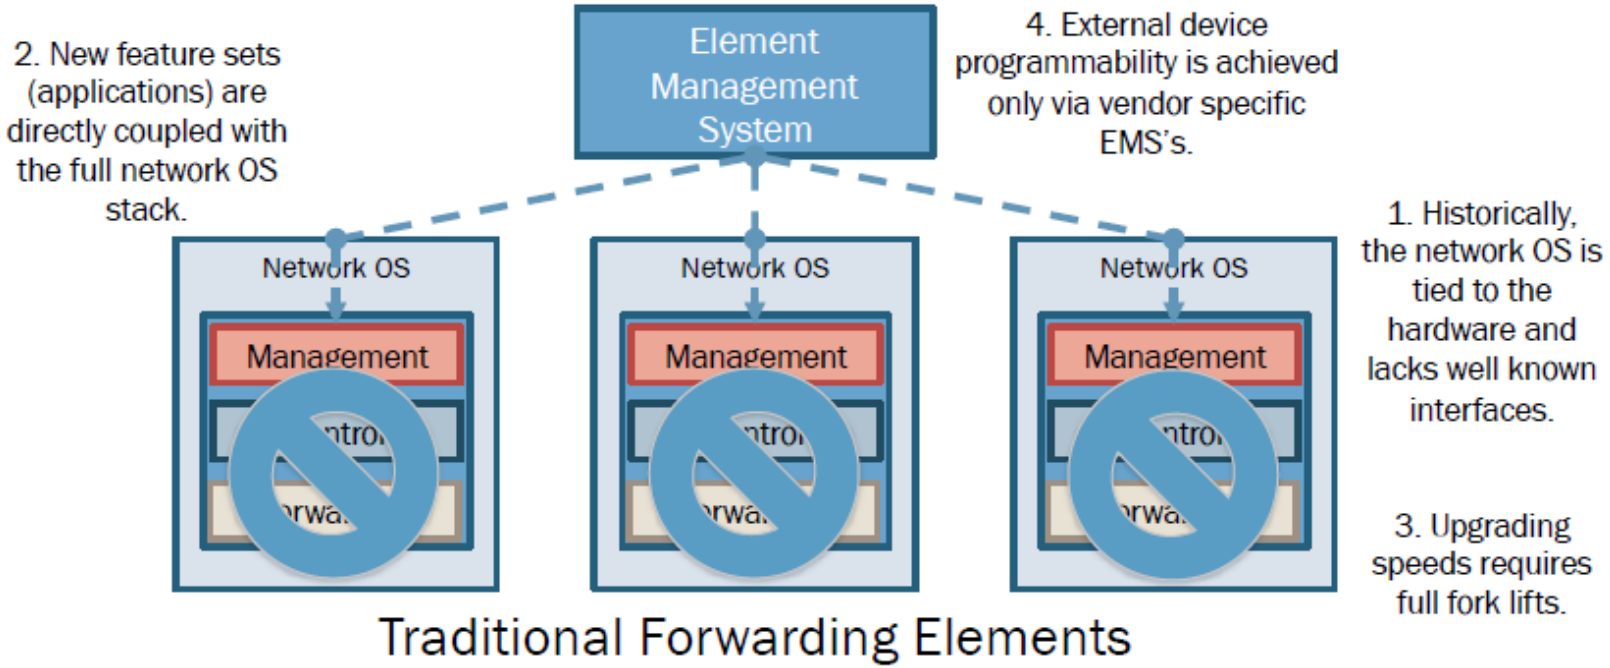

Traditional Forwarding Elements

Problems in Networking

- Network Virtualization is nothing new
  - VRF's
  - Vlans
- However, today network engineers must manually translate high level intents into low level implementations on a per box basis

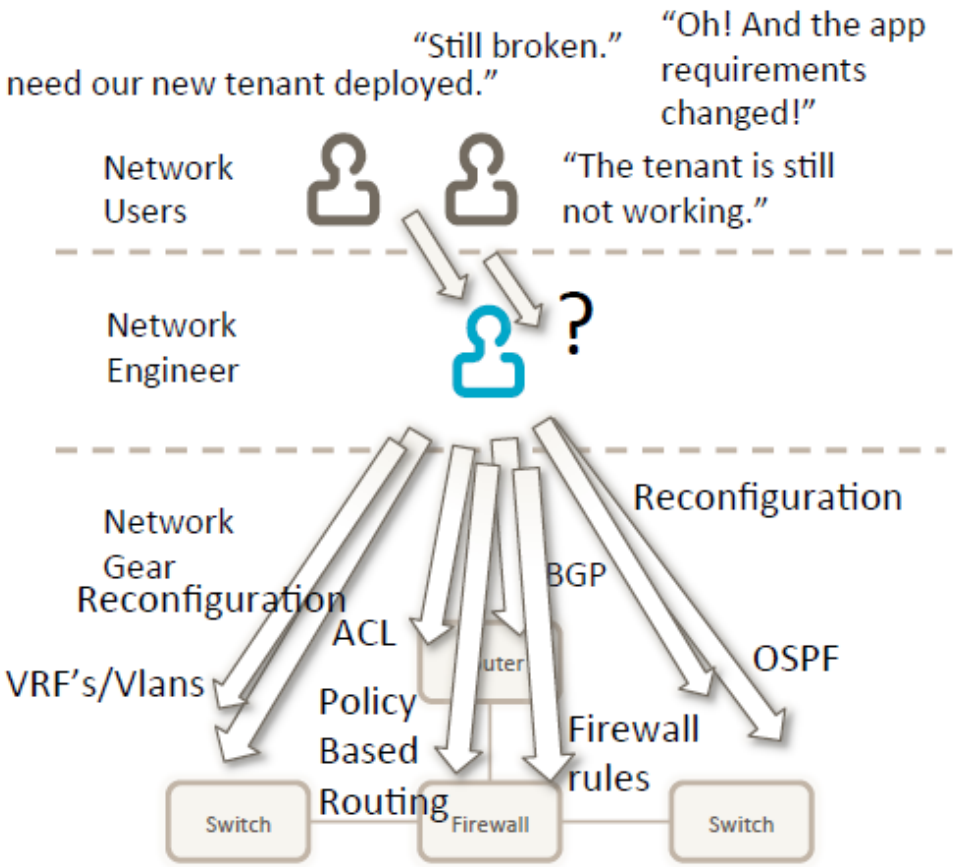

Problems in Networking

SDN Essentials

- Optimal resource utilization is a challenge in networking which typically leads to overprovisioning
  - QoS – Difficult to manage across disparate devices
  - Traffic Engineering – Requires MPLS/RSVP-TE or BGP and static configuration
  - Non-Best Path Forwarding – Requires either RSVP-TE or policy based routing both of which require static configuration which is difficult to scale

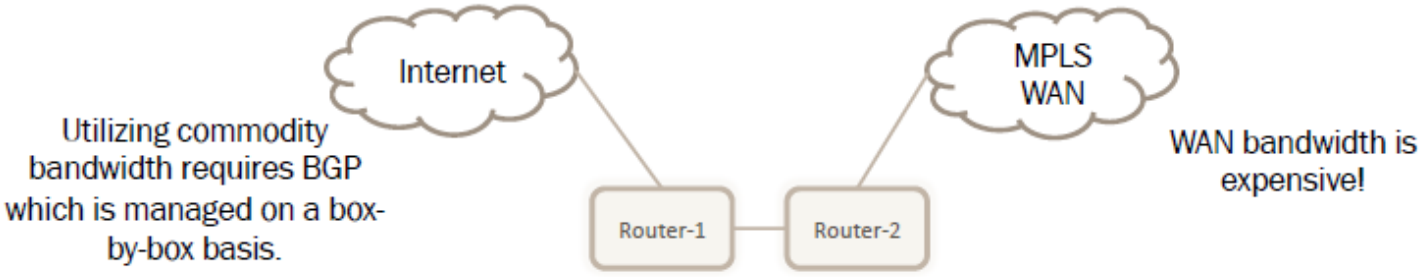

Current Internet

Closed to Innovations in the Infrastructure

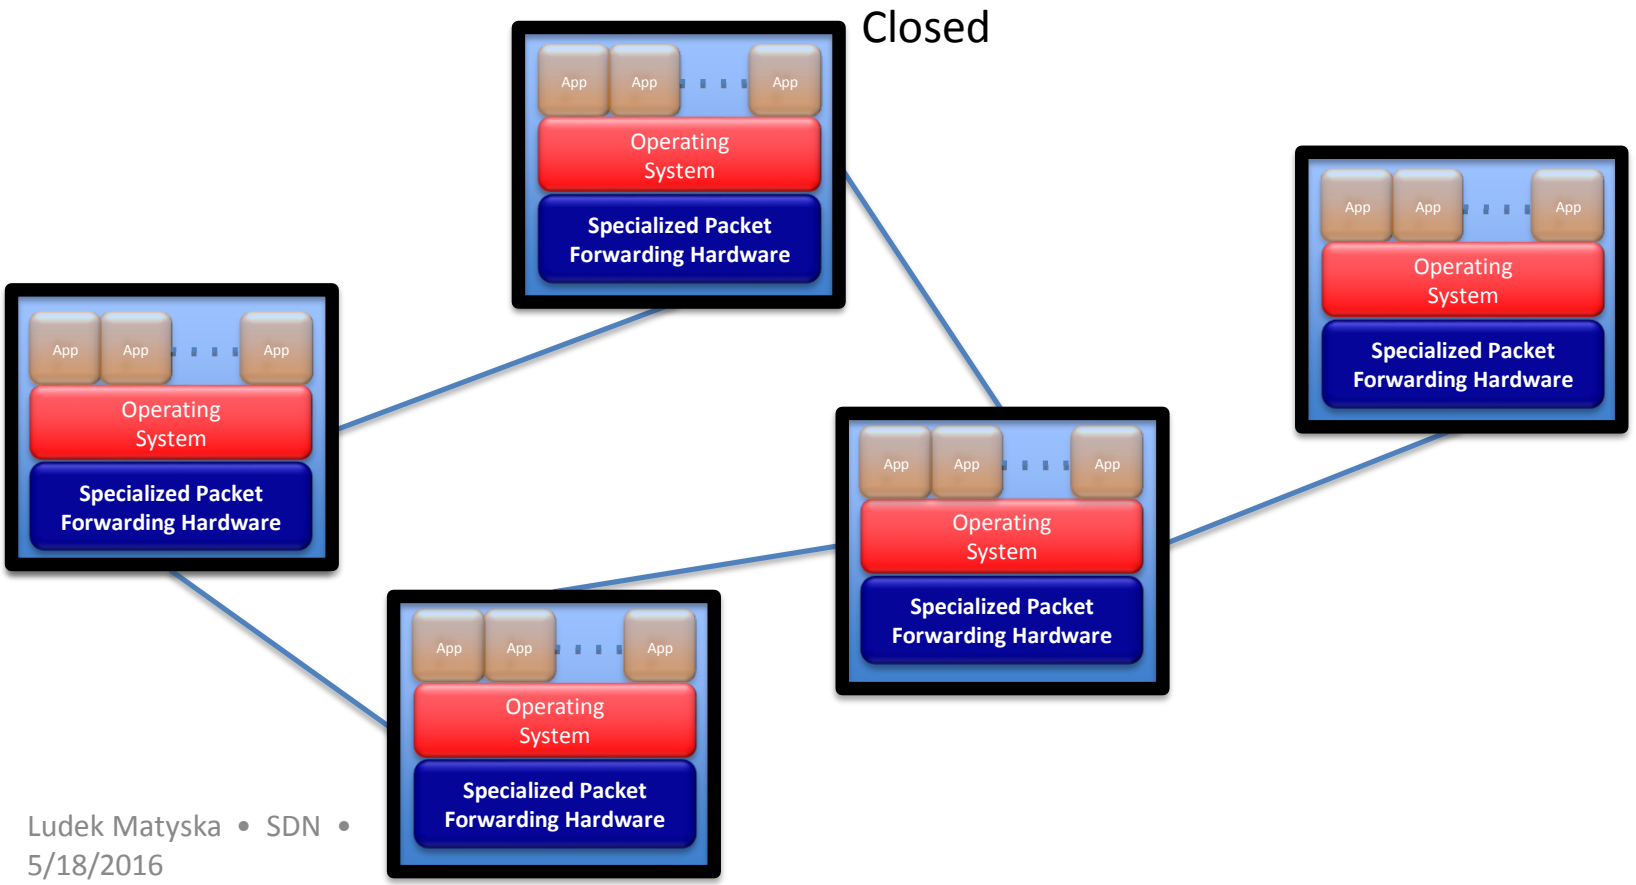

“Software Defined Networking” approach to open it

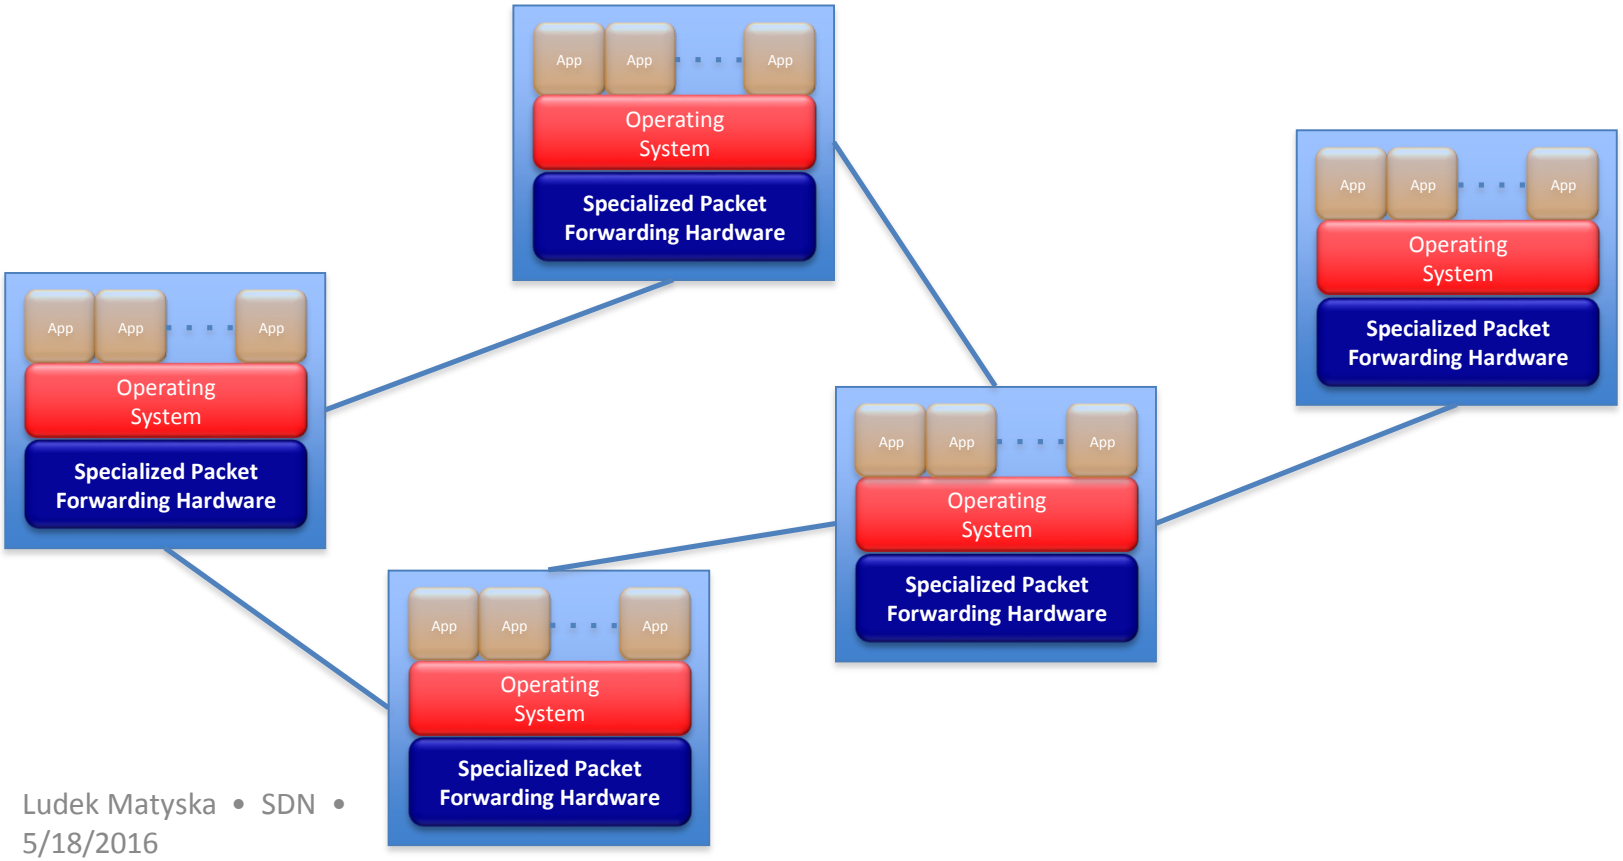

“Software Defined Networking” approach to open it

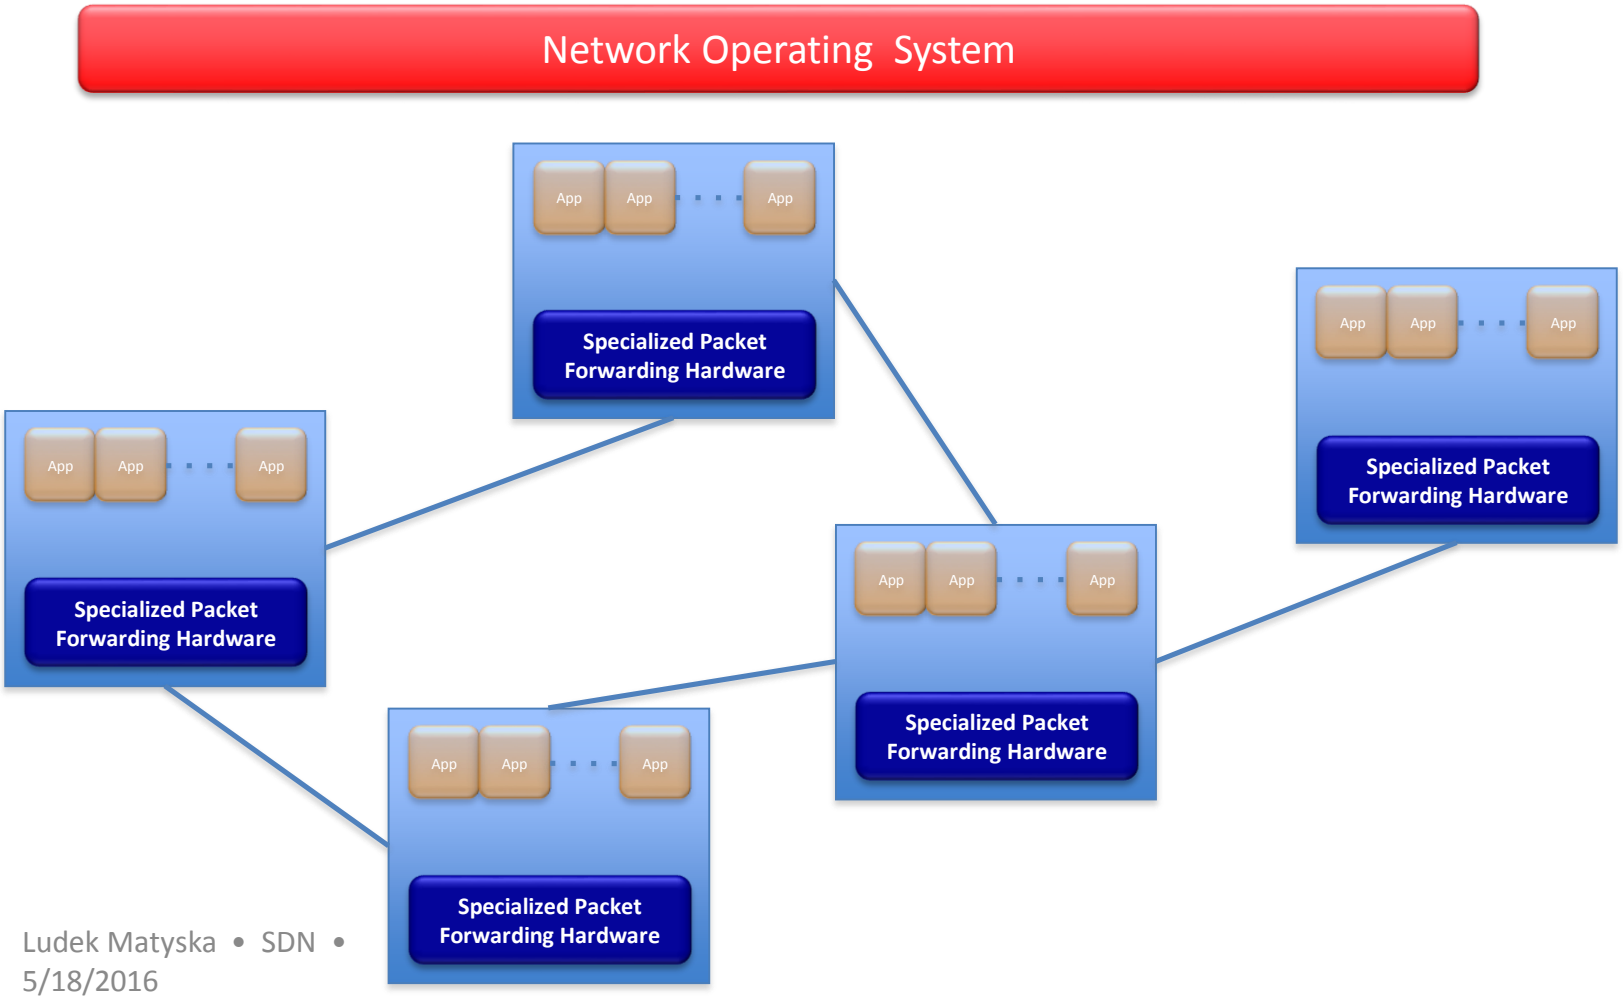

“Software Defined Networking” approach to open it

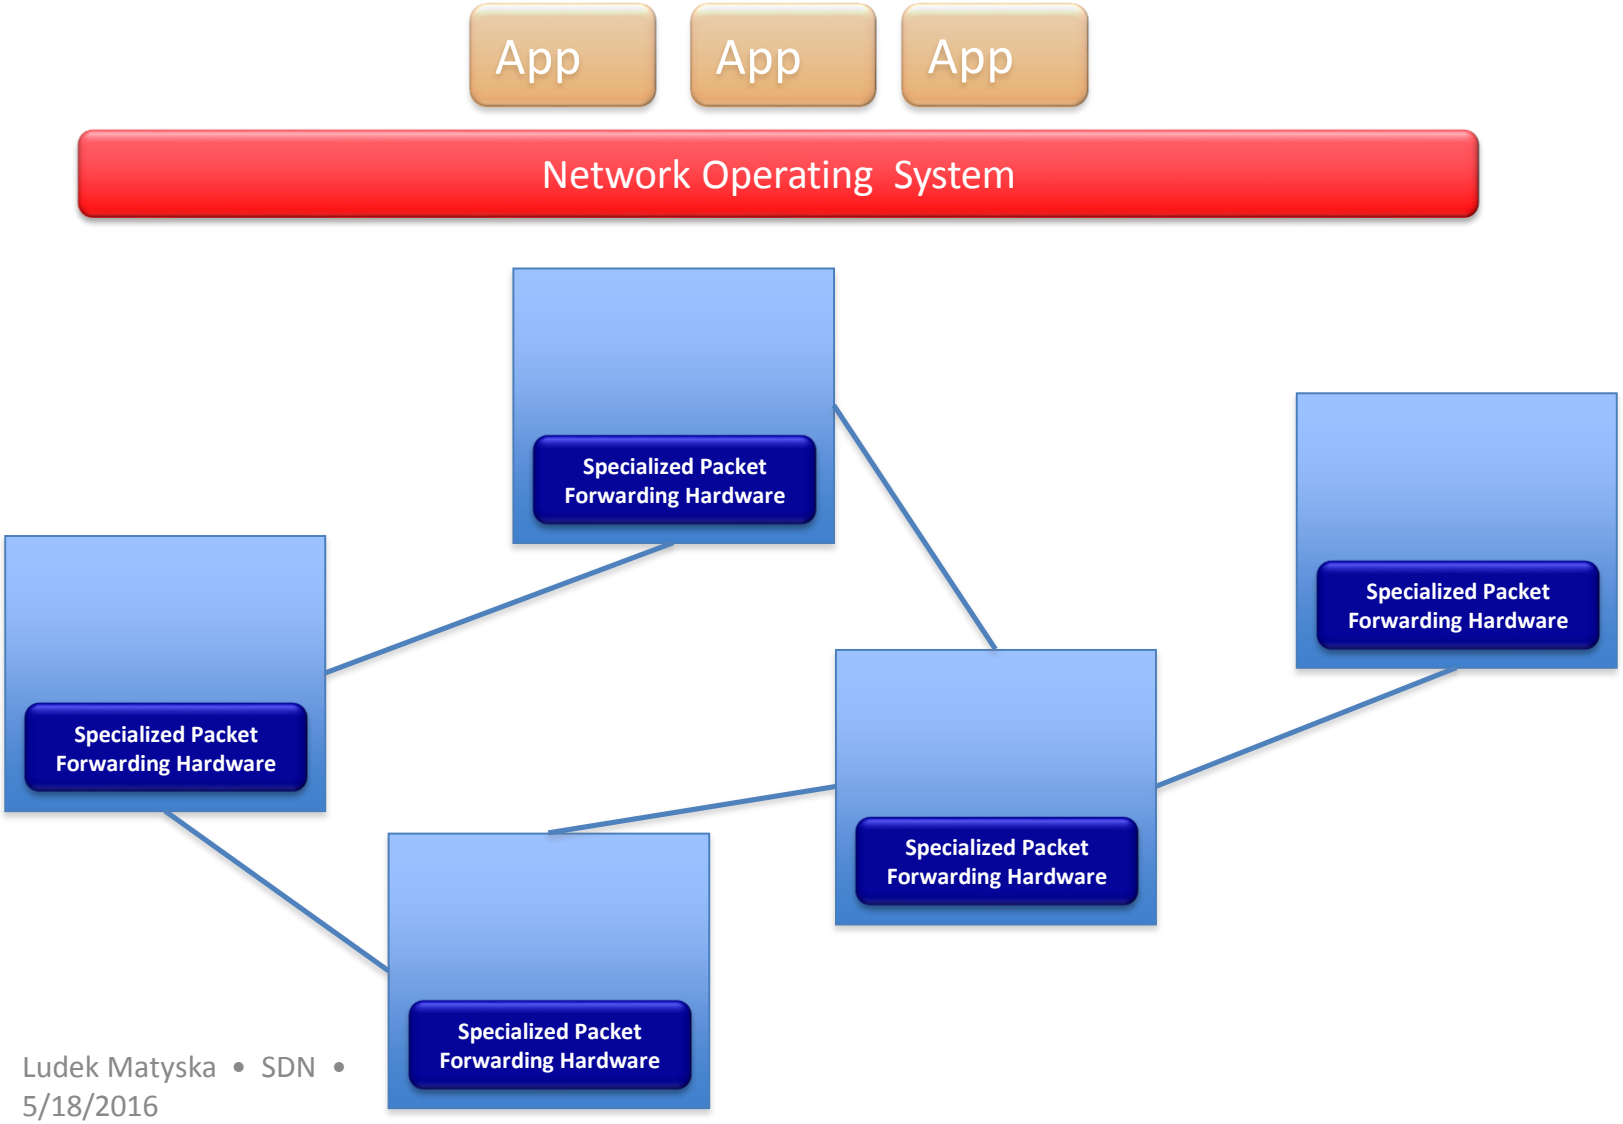

The “Software-defined Network”

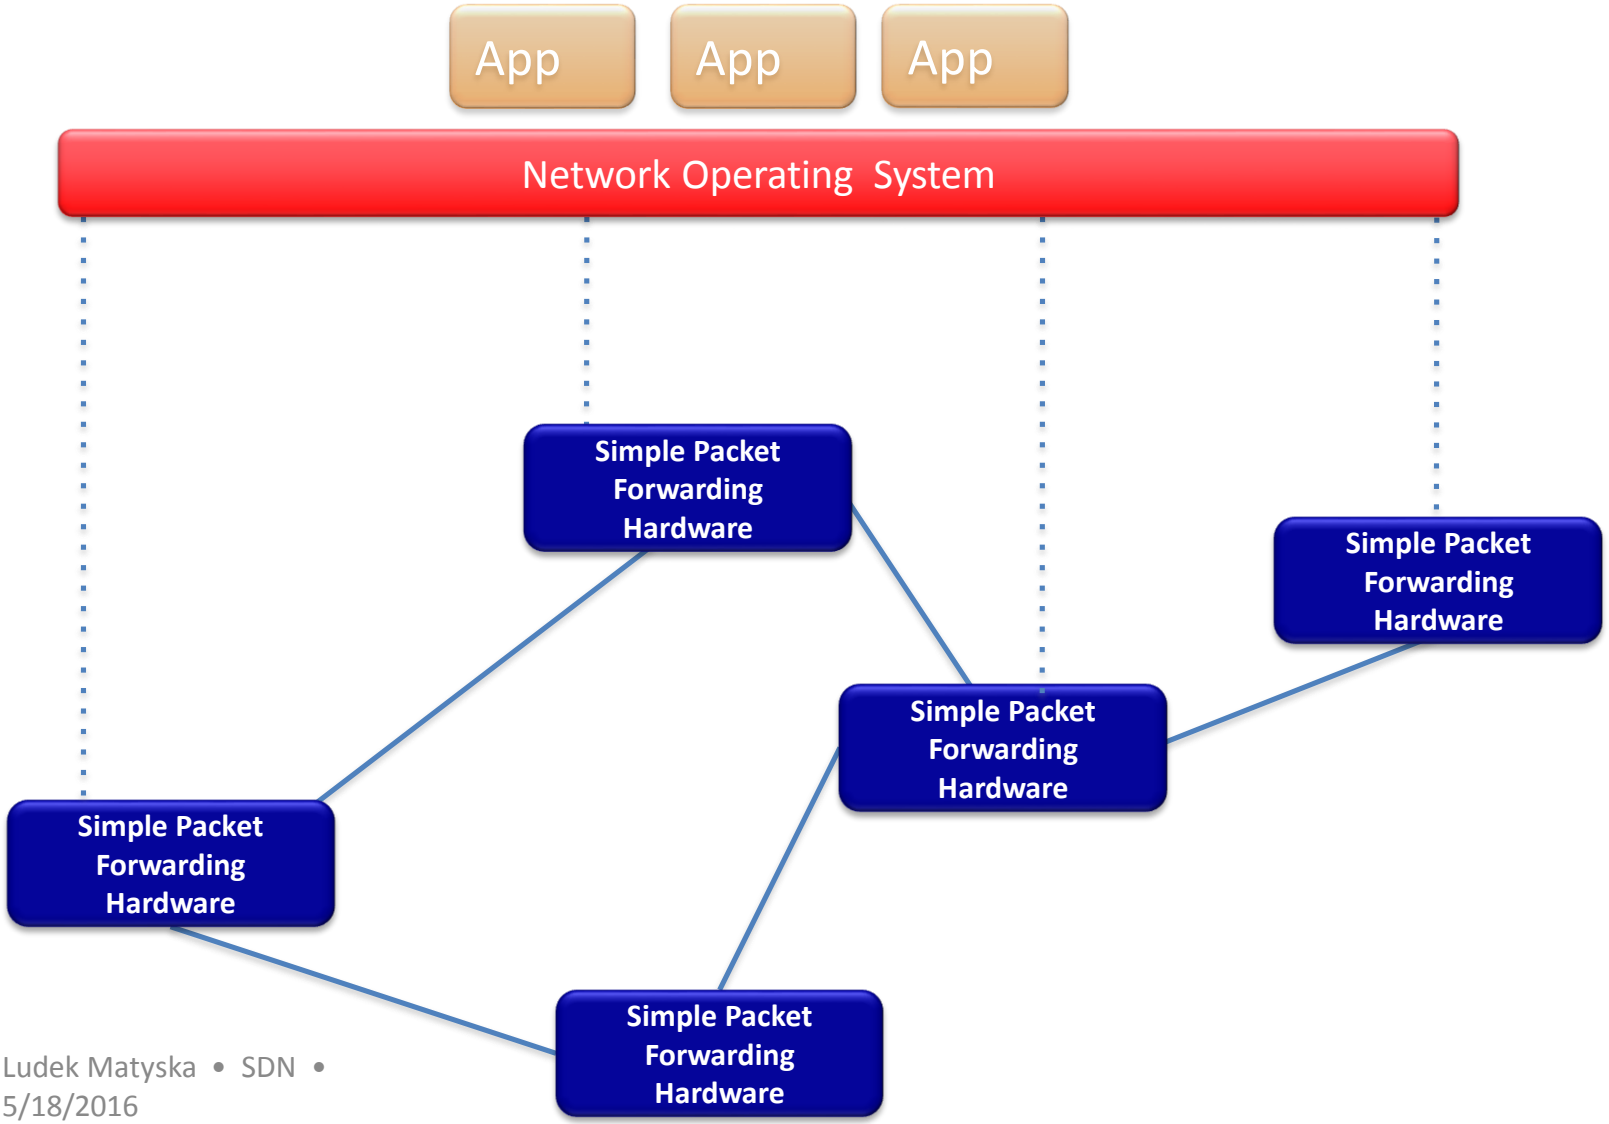

The “Software-defined Network”

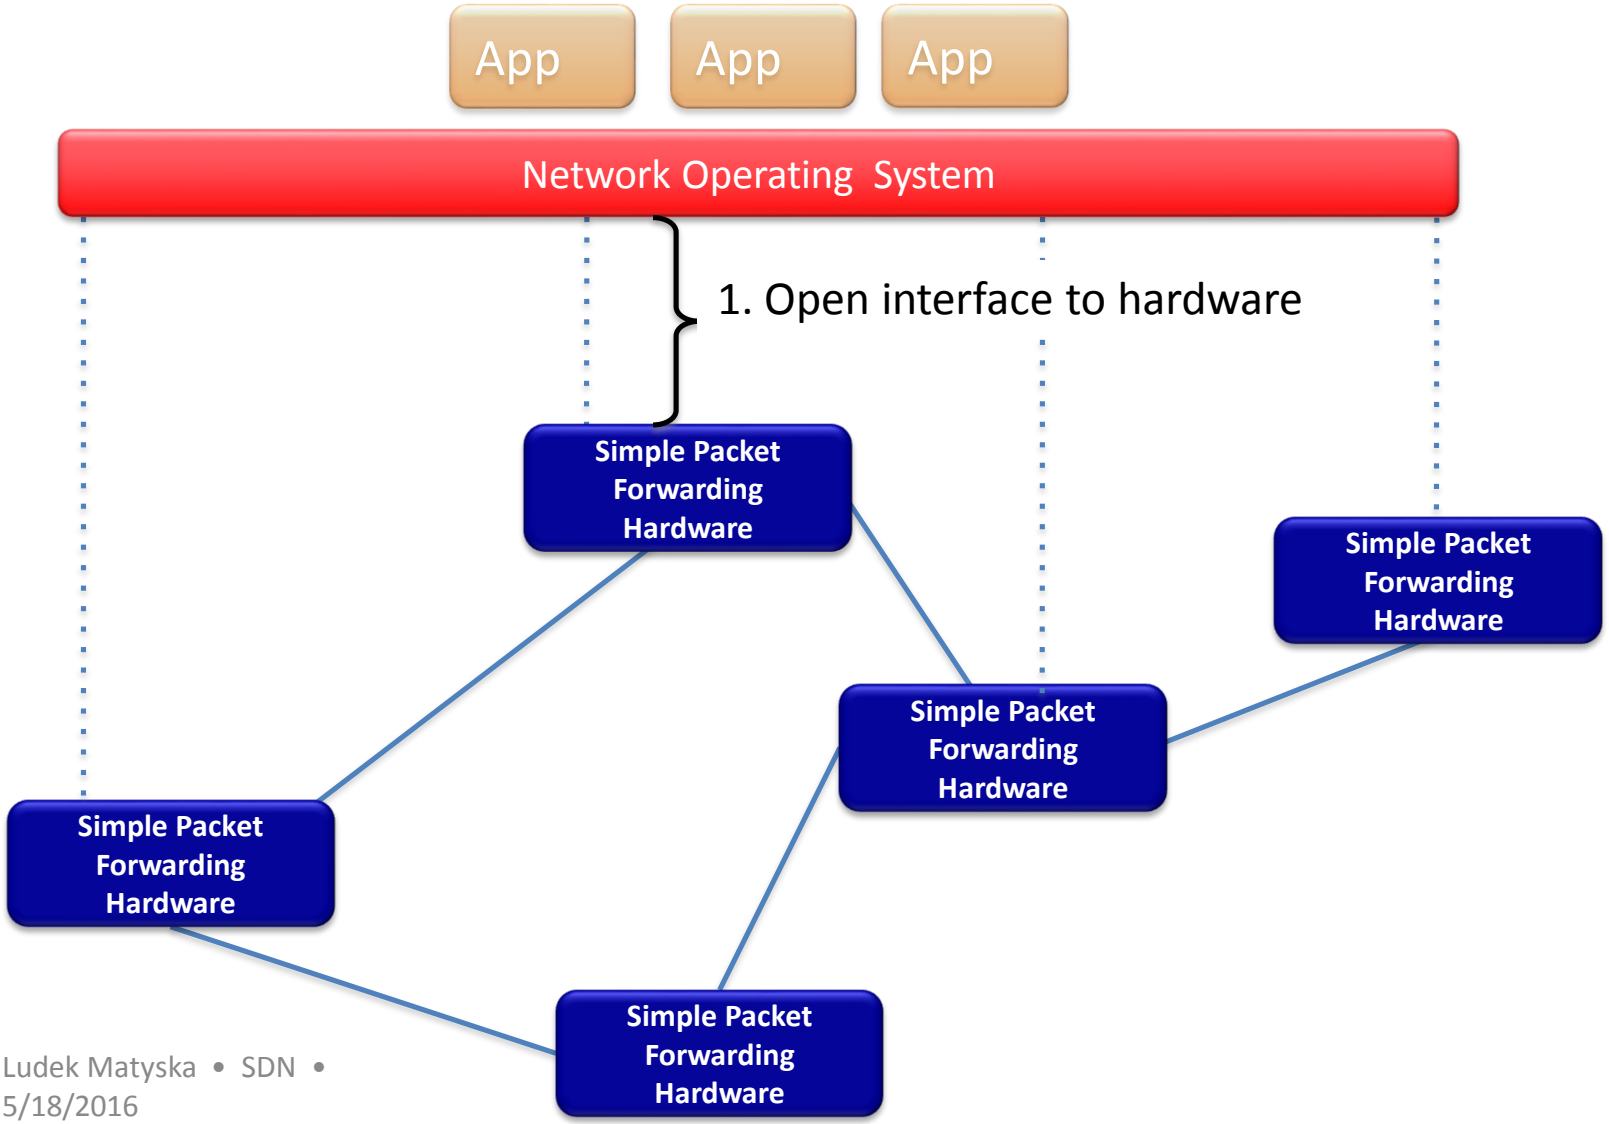

The “Software-defined Network”

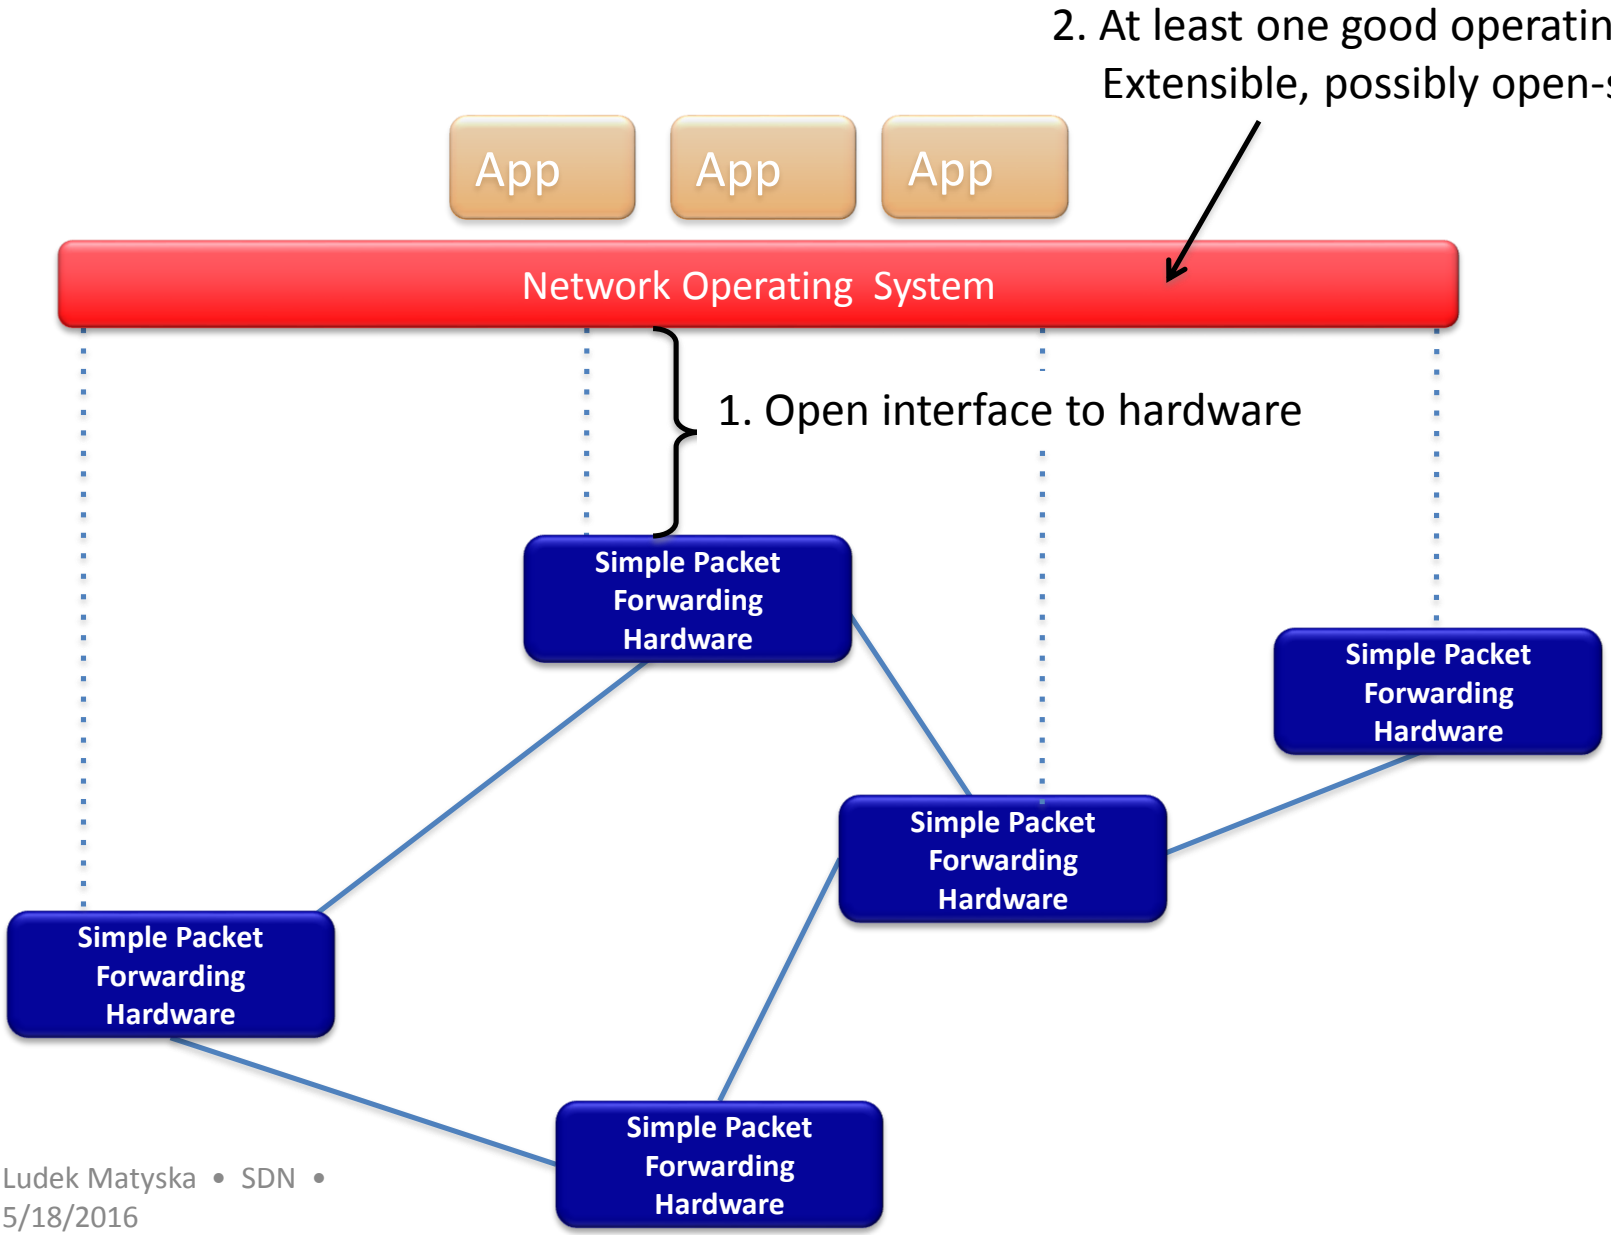

The “Software-defined Network”

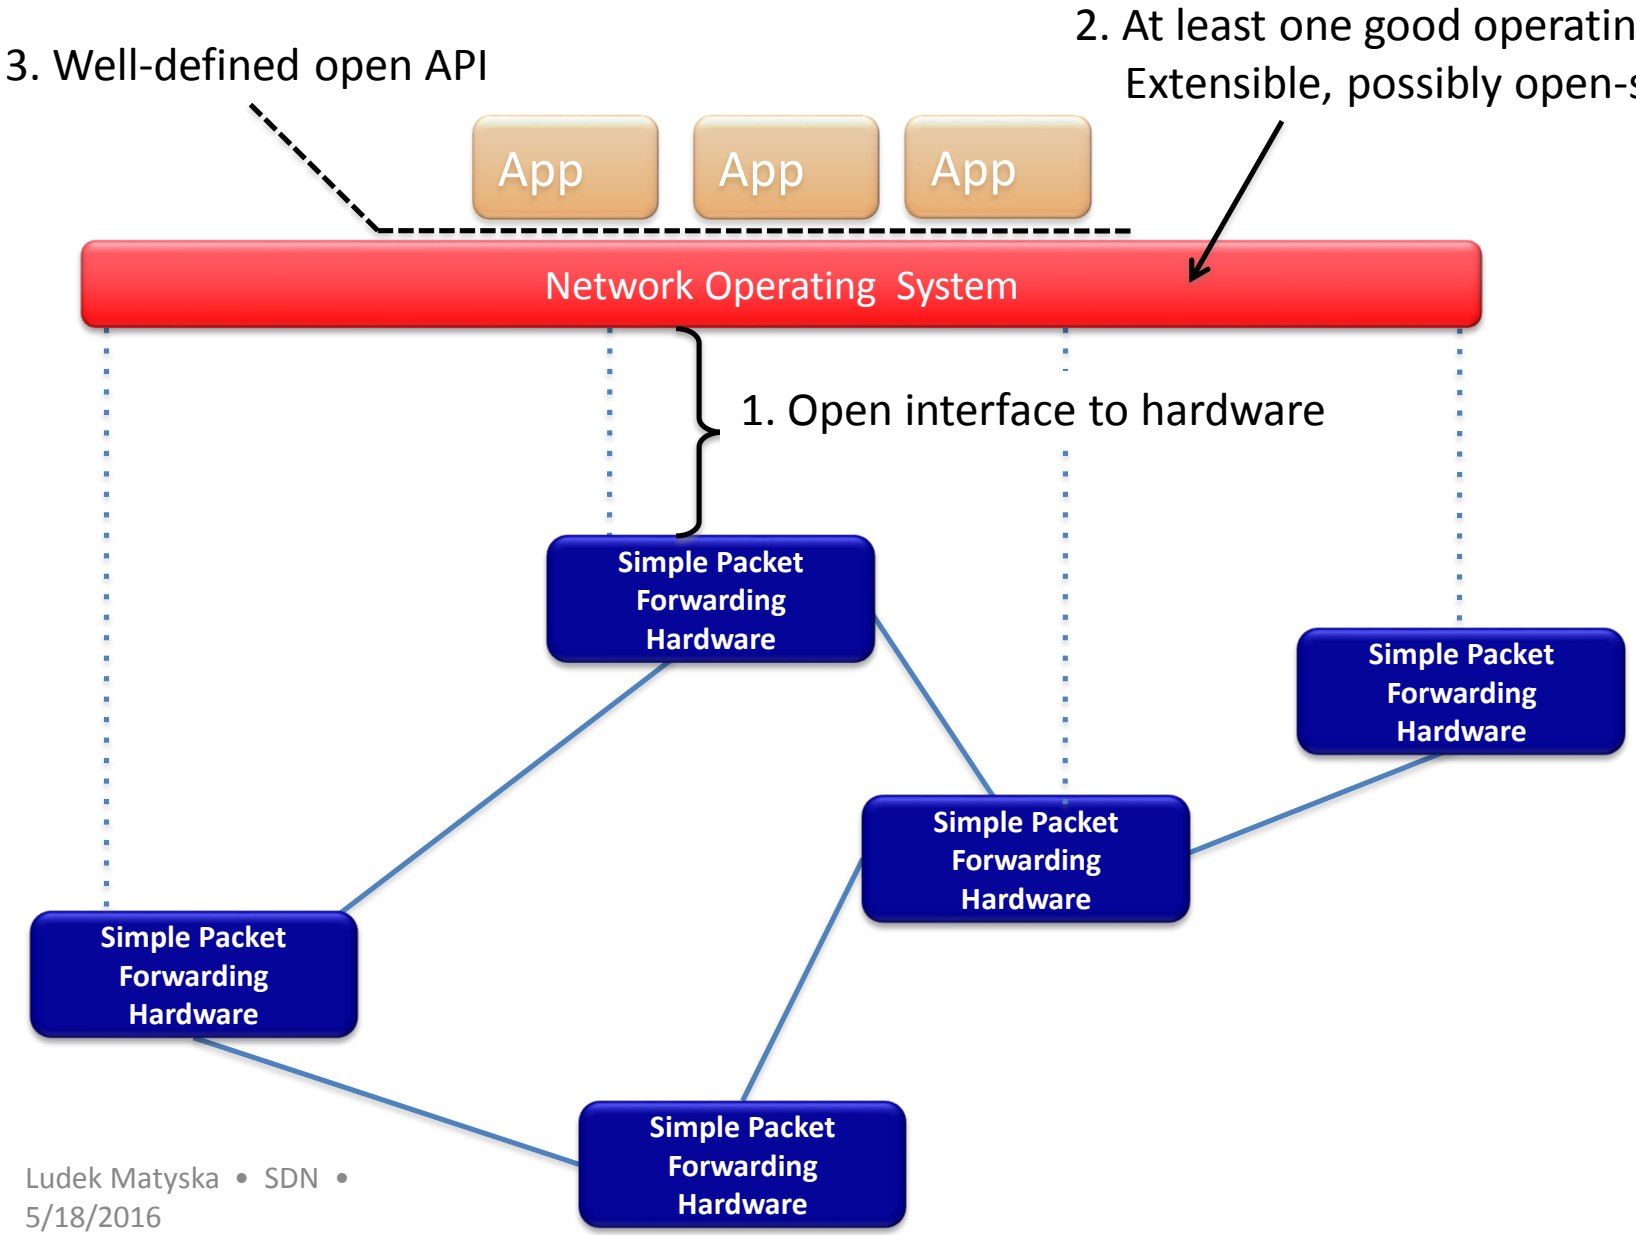

Software-defined network (SDN)

(An Attempt At) SDN Defined

SDN Essentials

- SDN Classic Definition (Open SDN)
  - *A technology to networking which allows centralized, programmable control planes so that network operators can control and manage directly their own virtualized networks.*
- Basic Concepts
  - Separation of control and data planes
  - Centralized, programmable control planes of network equipment
  - Support of multiple, isolated virtual networks
  - Networks must adjust and respond dynamically
  - Newly added features must not disrupt the network
  - Alleviate the need for manual configuration of individual devices
- Reality Today
  - Vendors and customers have morphed the original definition
  - Now there are flavors of SDN that fit into the general framework to varying degrees

Origins of SDN

- SDN originated from OpenFlow
- Centralized Controller
  ⇒ Easy to program
  ⇒ Change routing policies on the fly
  ⇒ Software Defined Network (SDN)
- Initially, SDN=
  - Separation of Control and Data Plane
  - Centralization of Control
  - OpenFlow to talk to the data plane
- Now the definition has changed significantly.

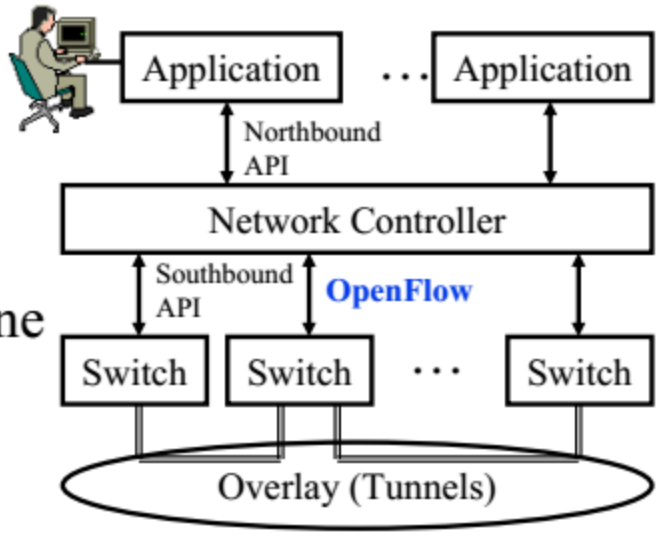

What is SDN?

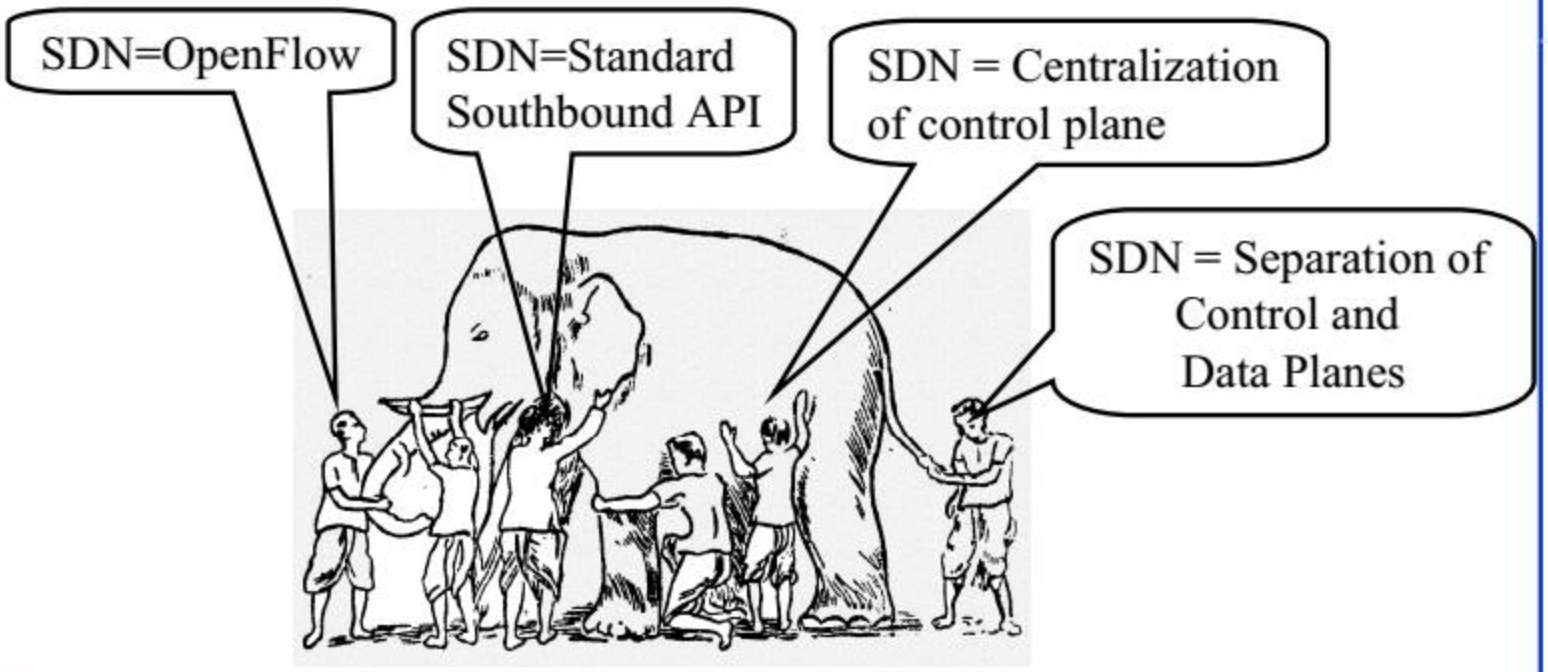

- All of these are mechanisms.
- SDN is *not* a mechanism.
- It is a framework to solve a set of problems ⇒ Many solutions

Original Definition of SDN

"*What is SDN?*

The physical separation of the network control plane from the forwarding plane, and where a control plane controls several devices."

1. Directly programmable
2. Agile: *Abstracting control from forwarding*
3. Centrally managed
4. Programmatically configured
5. Open standards-based vendor neutral

The above definition includes *How*.

Now many different opinions about *How*.

⇒SDN has become more general. Need to define by *What*?

Ref: https://www.opennetworking.org/index.php?option=com_content&view=article&id=686&Itemid=272&lang=en

What = Why We need SDN?

1. **Virtualization**: Use network resource without worrying about where it is physically located, how much it is, how it is organized, etc.
2. **Orchestration**: Should be able to control and manage thousands of devices with one command.
3. **Programmable**: Should be able to change behavior on the fly.
4. **Dynamic Scaling**: Should be able to change size, quantity
5. **Automation**: To lower OpEx minimize manual involvement
   - Troubleshooting
   - Reduce downtime
   - Policy enforcement
   - Provisioning/Re-provisioning/Segmentation of resources
   - Add new workloads, sites, devices, and resources

Why We need SDN? (Cont)

6. **Visibility**: Monitor resources, connectivity
7. **Performance**: Optimize network device utilization
   - Traffic engineering/Bandwidth management
   - Capacity optimization
   - Load balancing
   - High utilization
   - Fast failure handling
8. **Multi-tenancy**: Tenants need complete control over their addresses, topology, and routing, security
9. **Service Integration**: Load balancers, firewalls, Intrusion Detection Systems (IDS), provisioned on demand and placed appropriately on the traffic path

Why We need SDN? (Cont)

10. Openness: Full choice of "How" mechanisms
⇒ Modular plug-ins
⇒ Abstraction:

- Abstract = Summary = Essence = General Idea
  ⇒ Hide the details.
- Also, abstract is opposite of concrete
  ⇒ Define tasks by APIs and not by how it should be done.
  E.g., send from A to B. Not OSPF.

Ref: http://www.networkworld.com/news/2013/110813-onug-sdn-275784.html

Ref: Open Data Center Alliance Usage Model: Software Defined Networking Rev 1.0,"
http://www.opendatacenteralliance.org/docs/Software_Defined_Networking_Master_Usage_Model_Rev1.0.pdf

SDN Definition

- SDN is a *framework* to allow network administrators to *automatically* and dynamically manage and control a *large number* of network devices, *services*, topology, traffic paths, and packet handling (quality of service) policies using high-level languages and APIs. Management includes provisioning, operating, *monitoring*, optimizing, and managing FCAPS (faults, configuration, accounting, *performance*, and security) in a *multi-tenant* environment.
- Key: Dynamic ⇒ Quick
  Legacy approaches such as CLI were not quick particularly for large networks

OpenFlow protocol

What is OpenFlow?

- OpenFlow is a protocol that enables programmability of the forwarding plane across the network as a whole
- OpenFlow is leveraged at the Southbound Interface between SDN Controller and OpenFlow switch
- OpenFlow attempts to abstract the implementation details of networks and forwarding elements using simple messaging

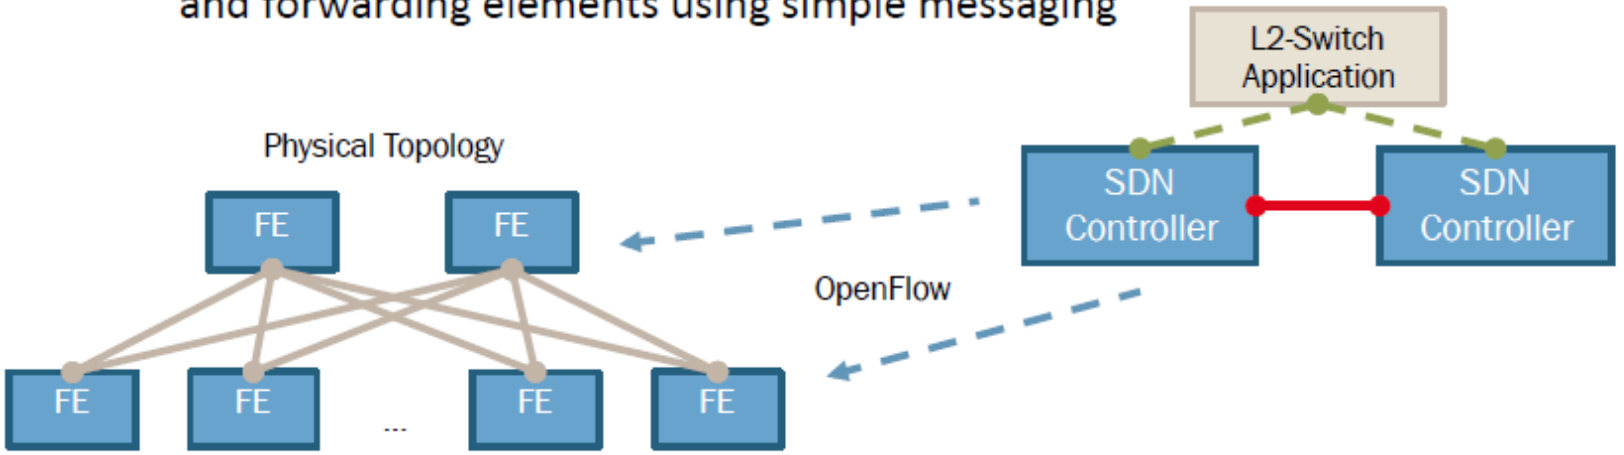

What is OpenFlow?

SDN Essentials

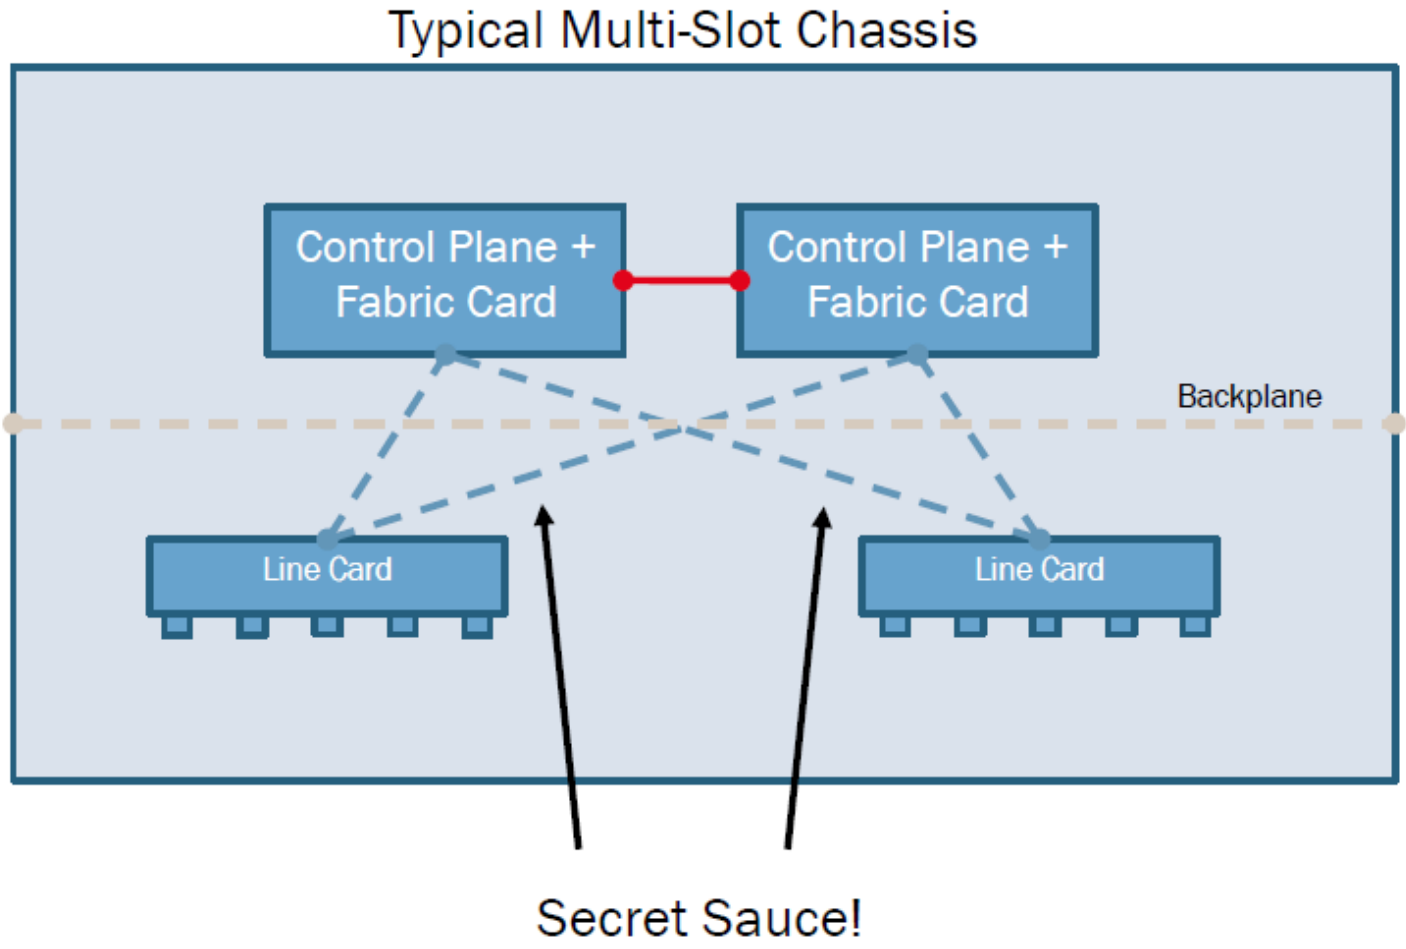

What is OpenFlow?

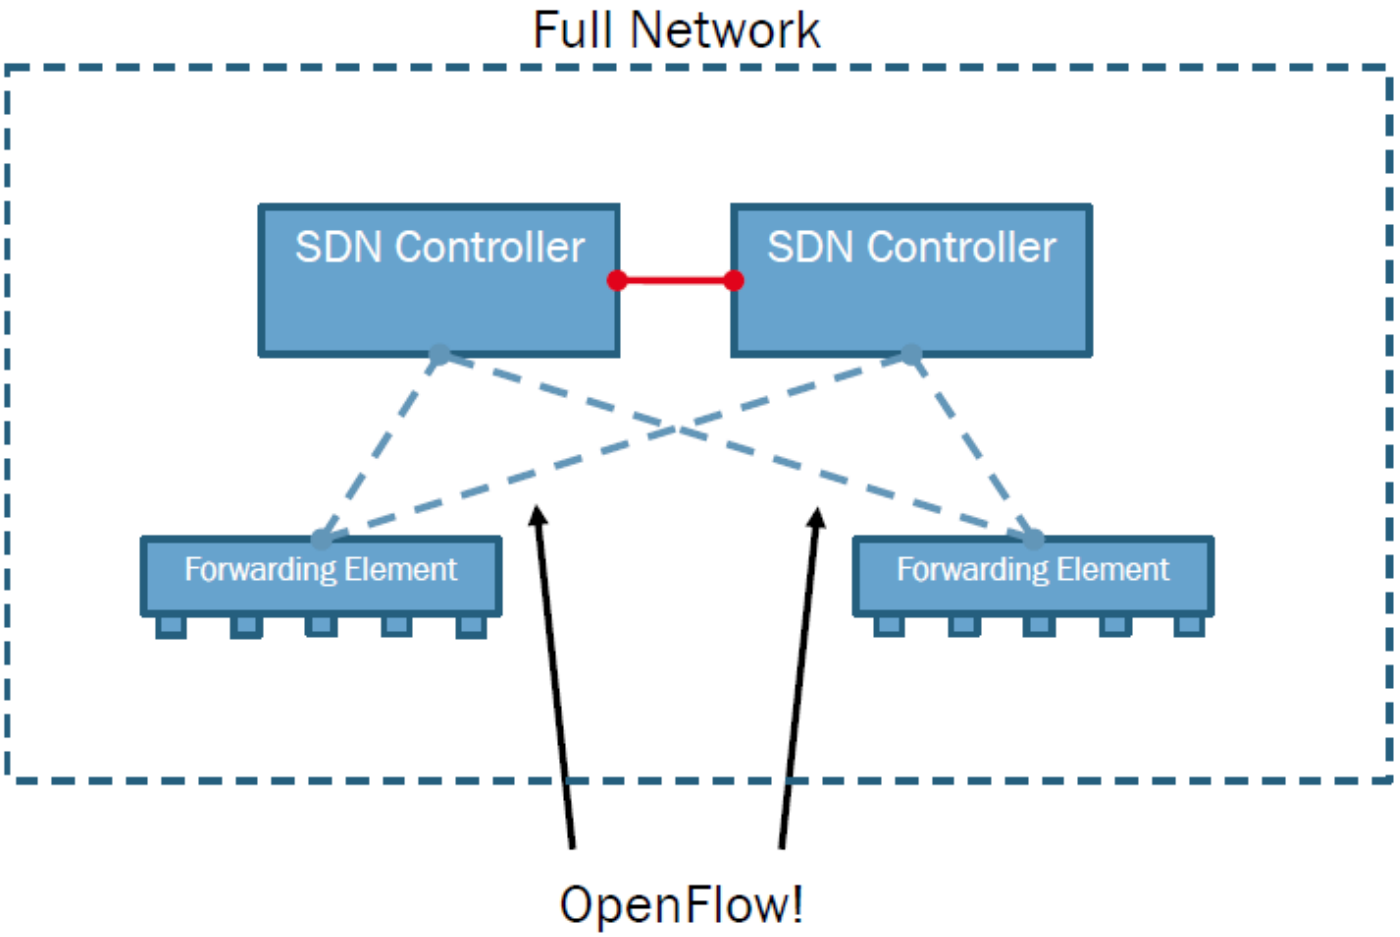

Components of OpenFlow Network

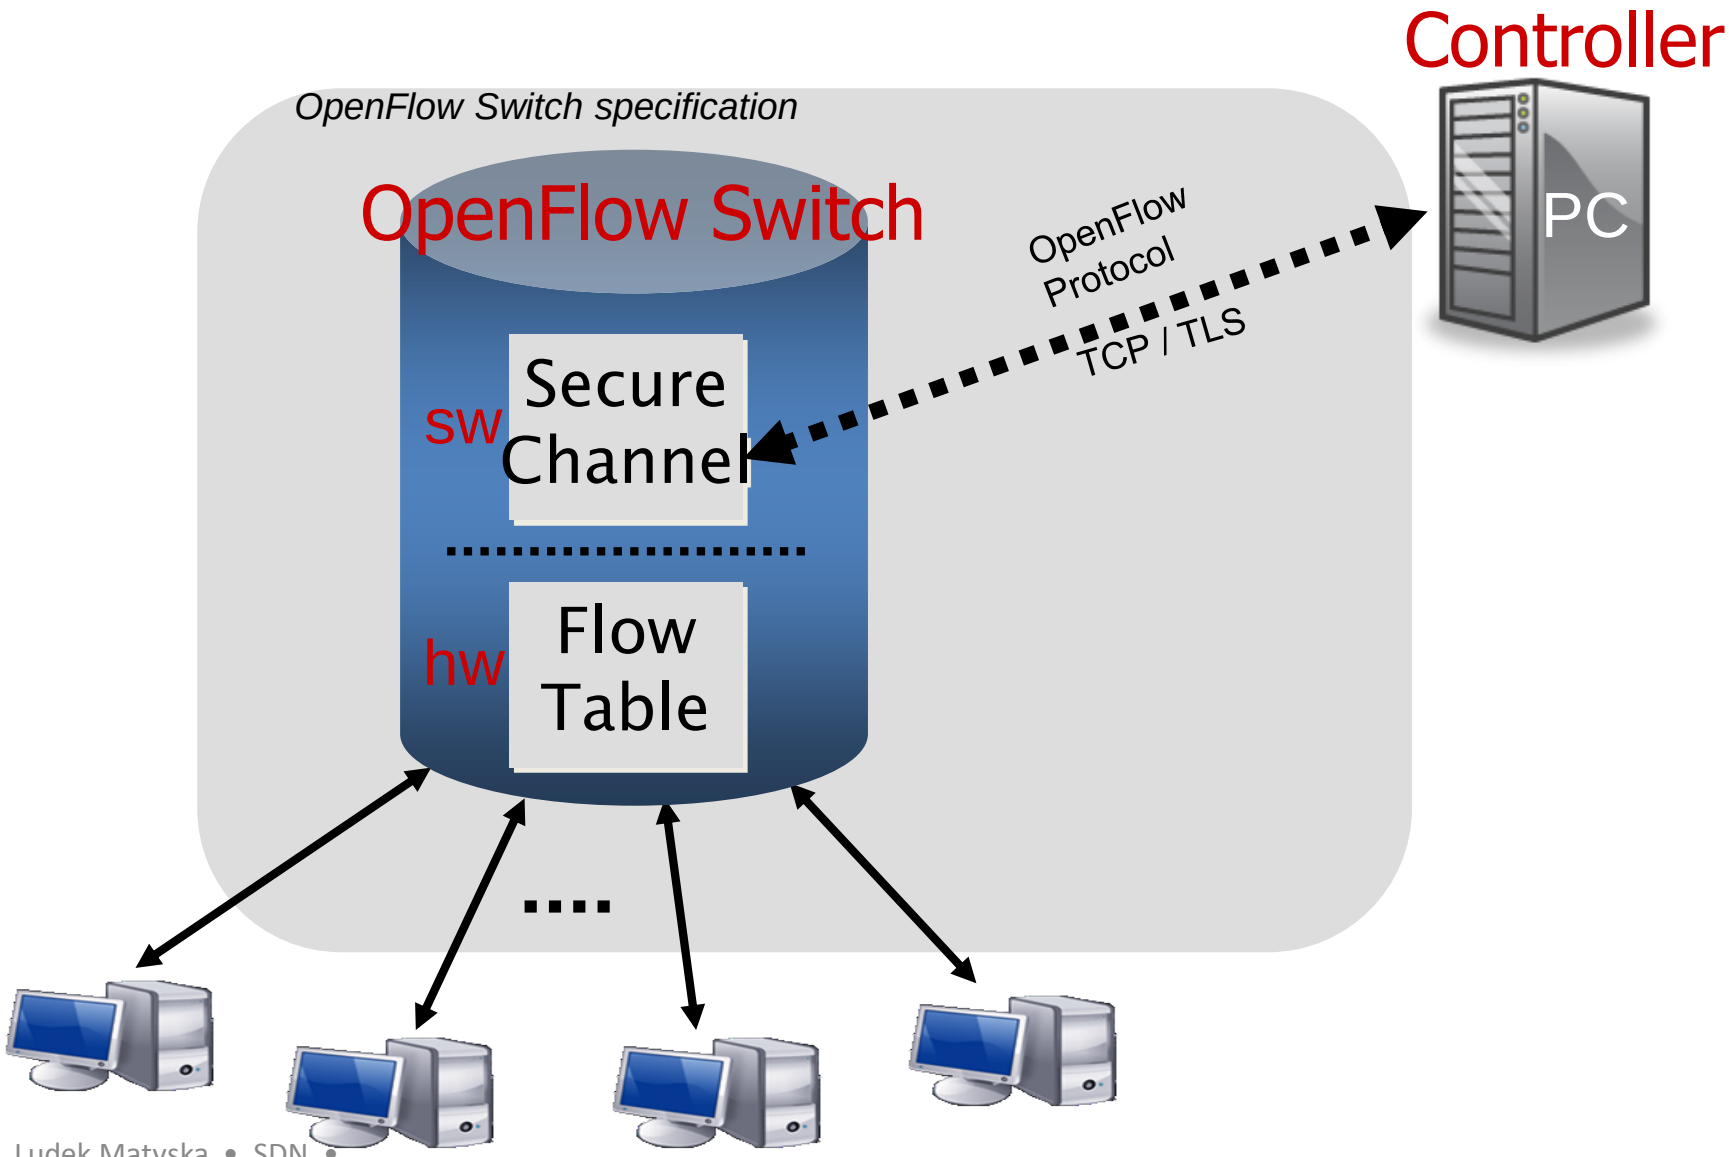

* Figure From OpenFlow Switch Specification

How does OpenFlow work?

OpenFlow Example

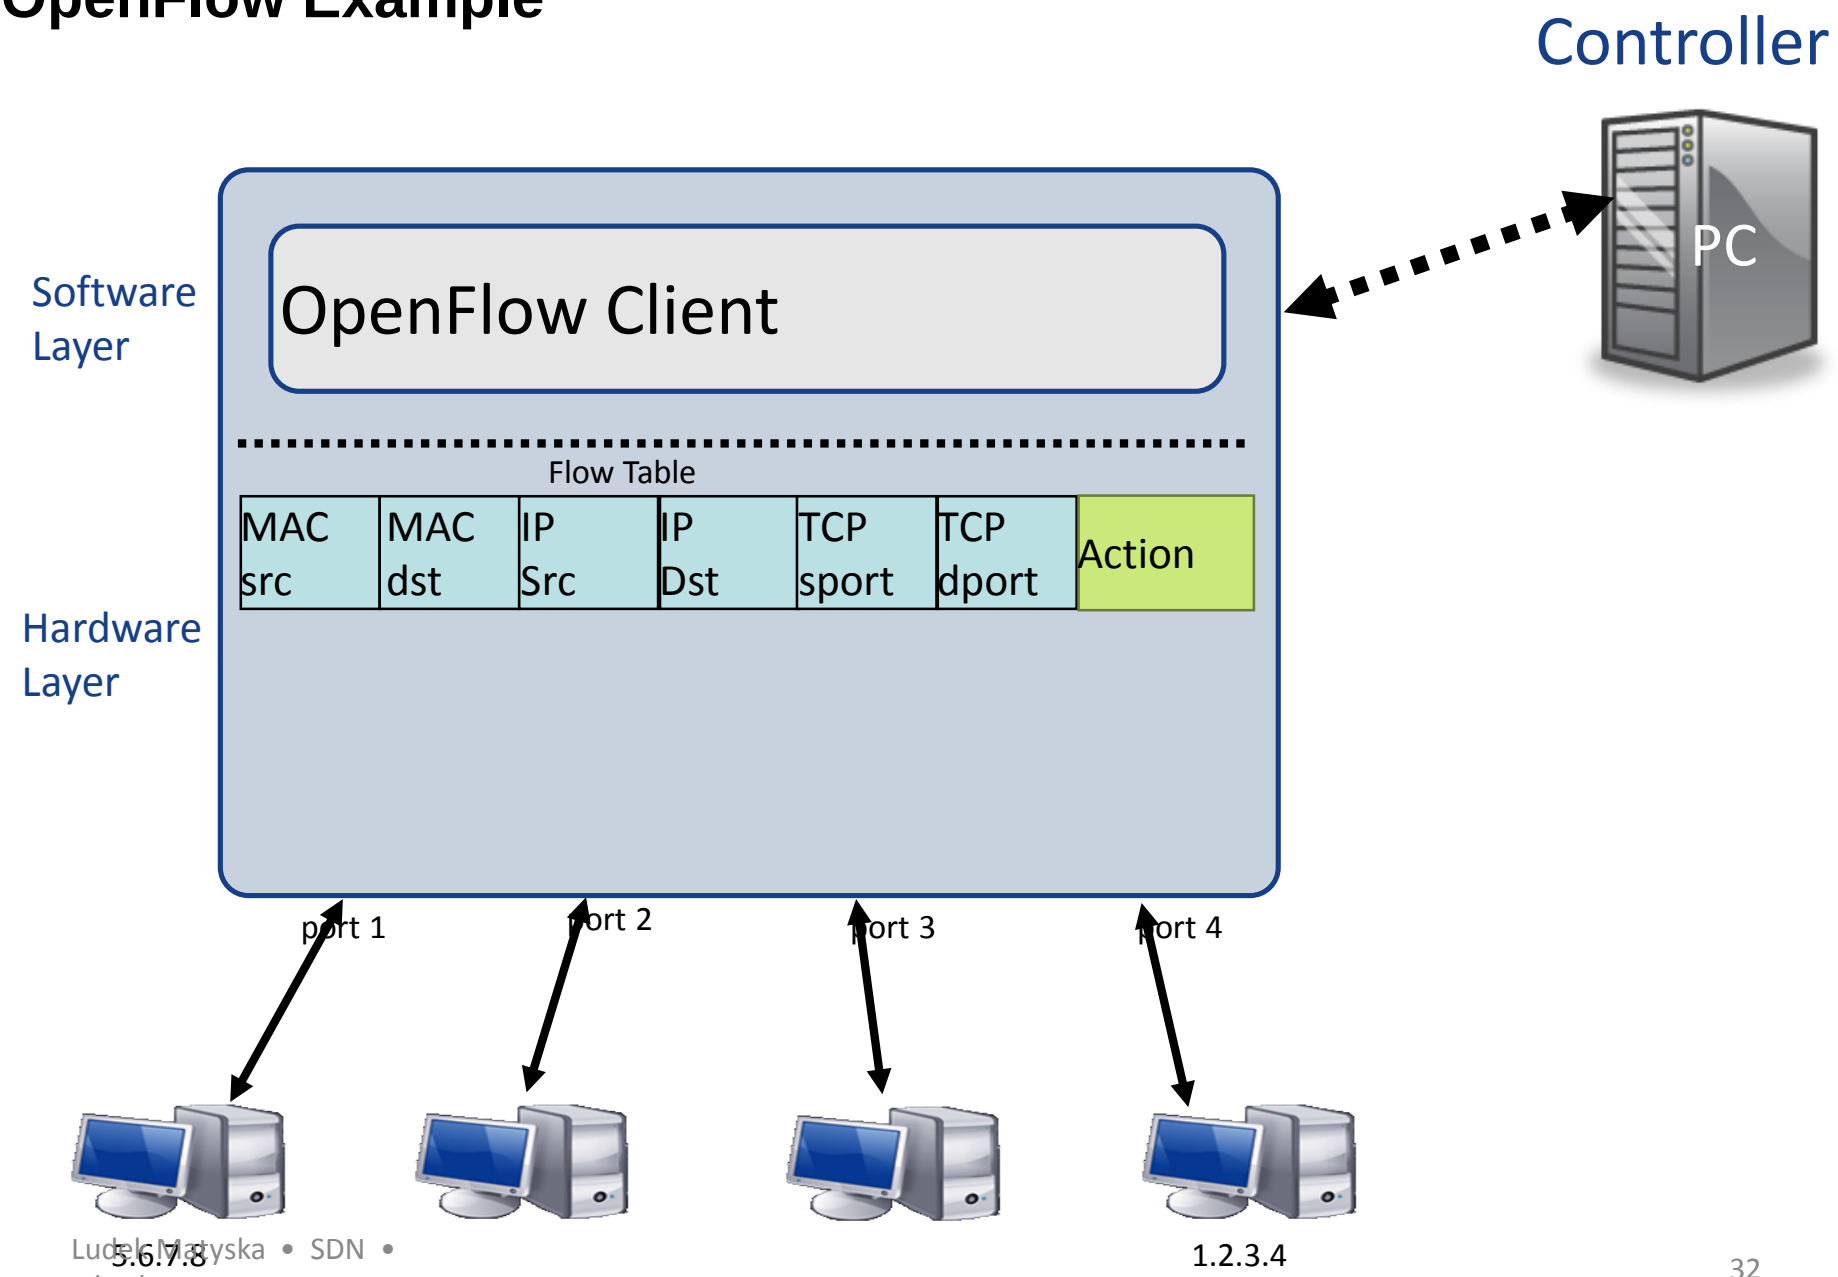

OpenFlow Example

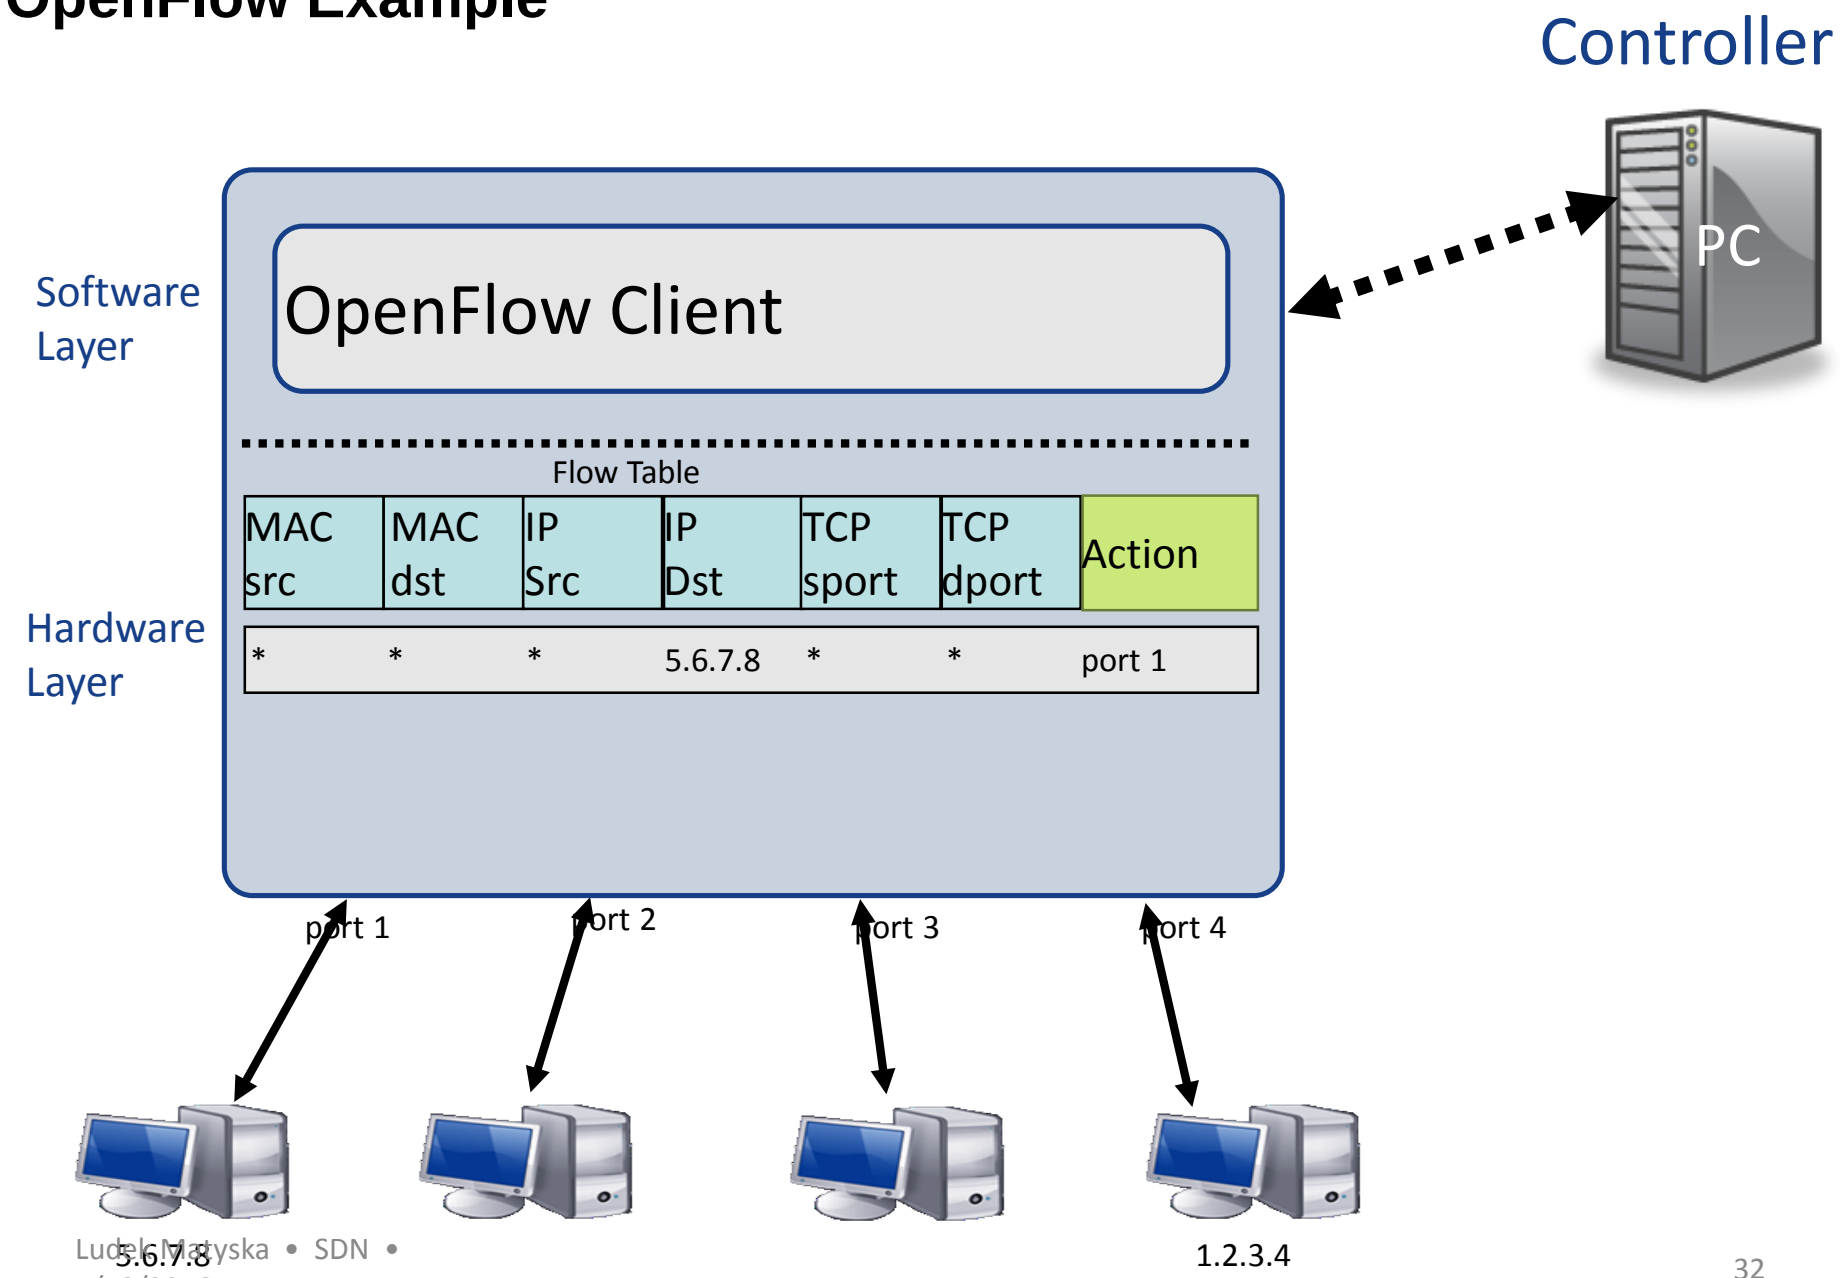

OpenFlow usage

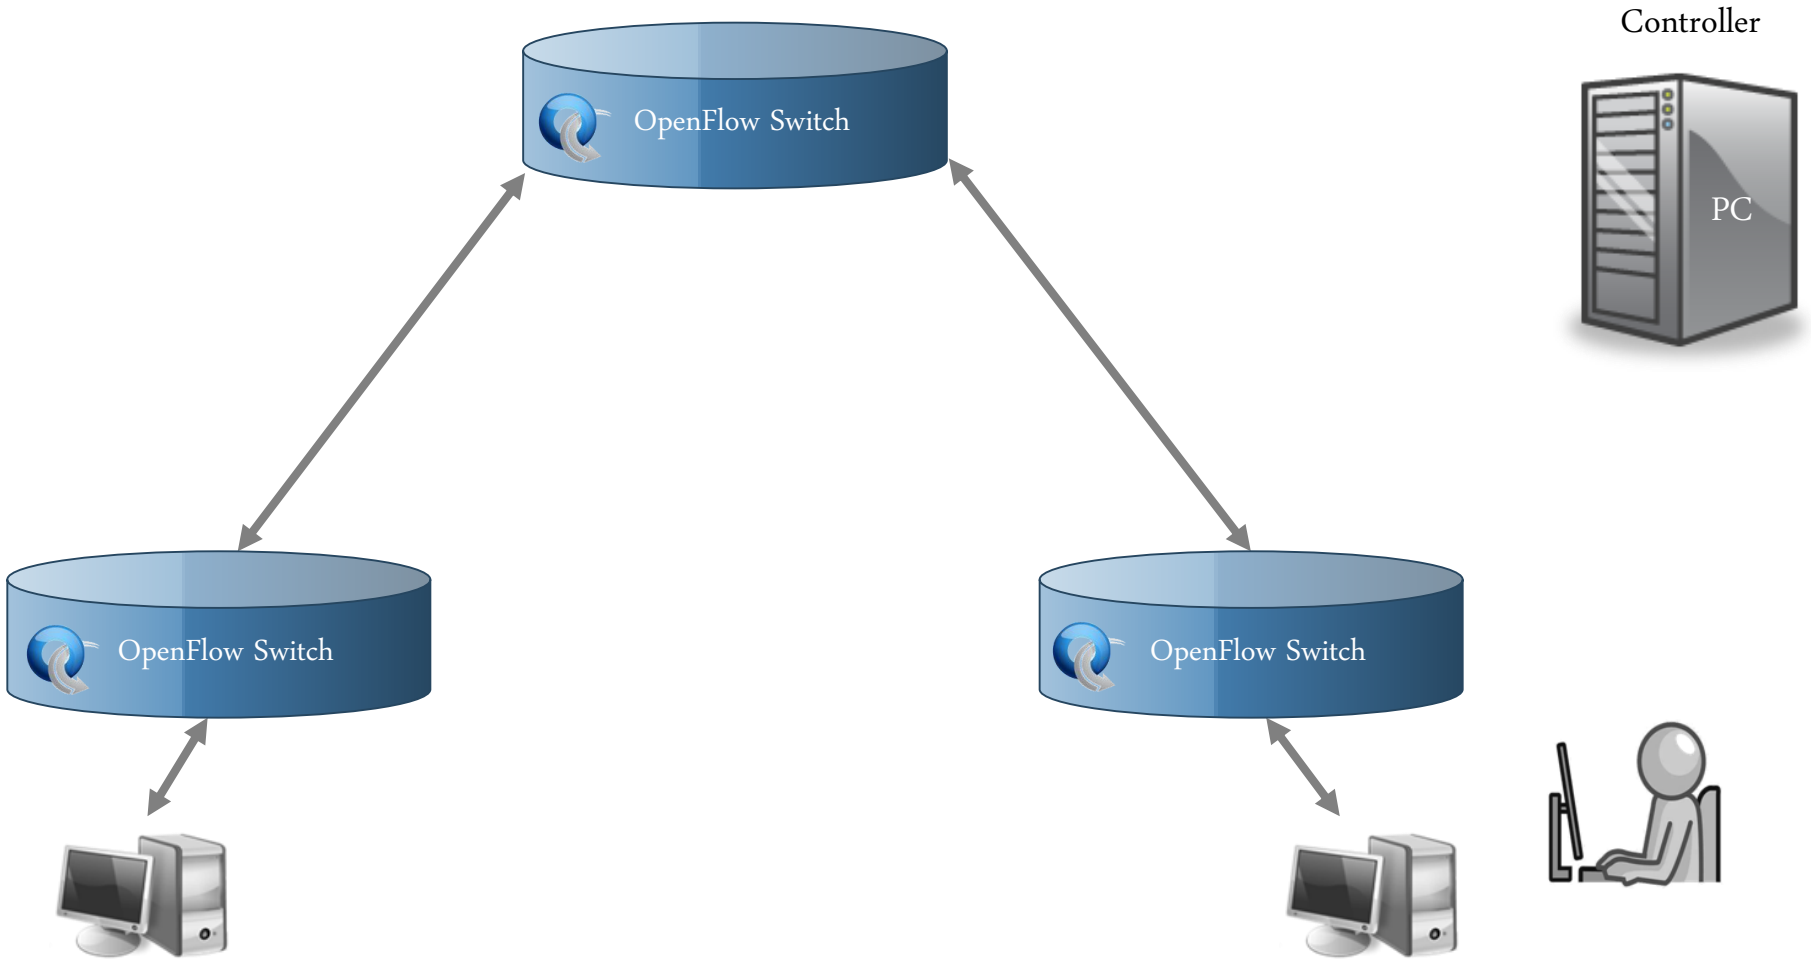

OpenFlow offloads control intelligence to a remote software

OpenFlow usage

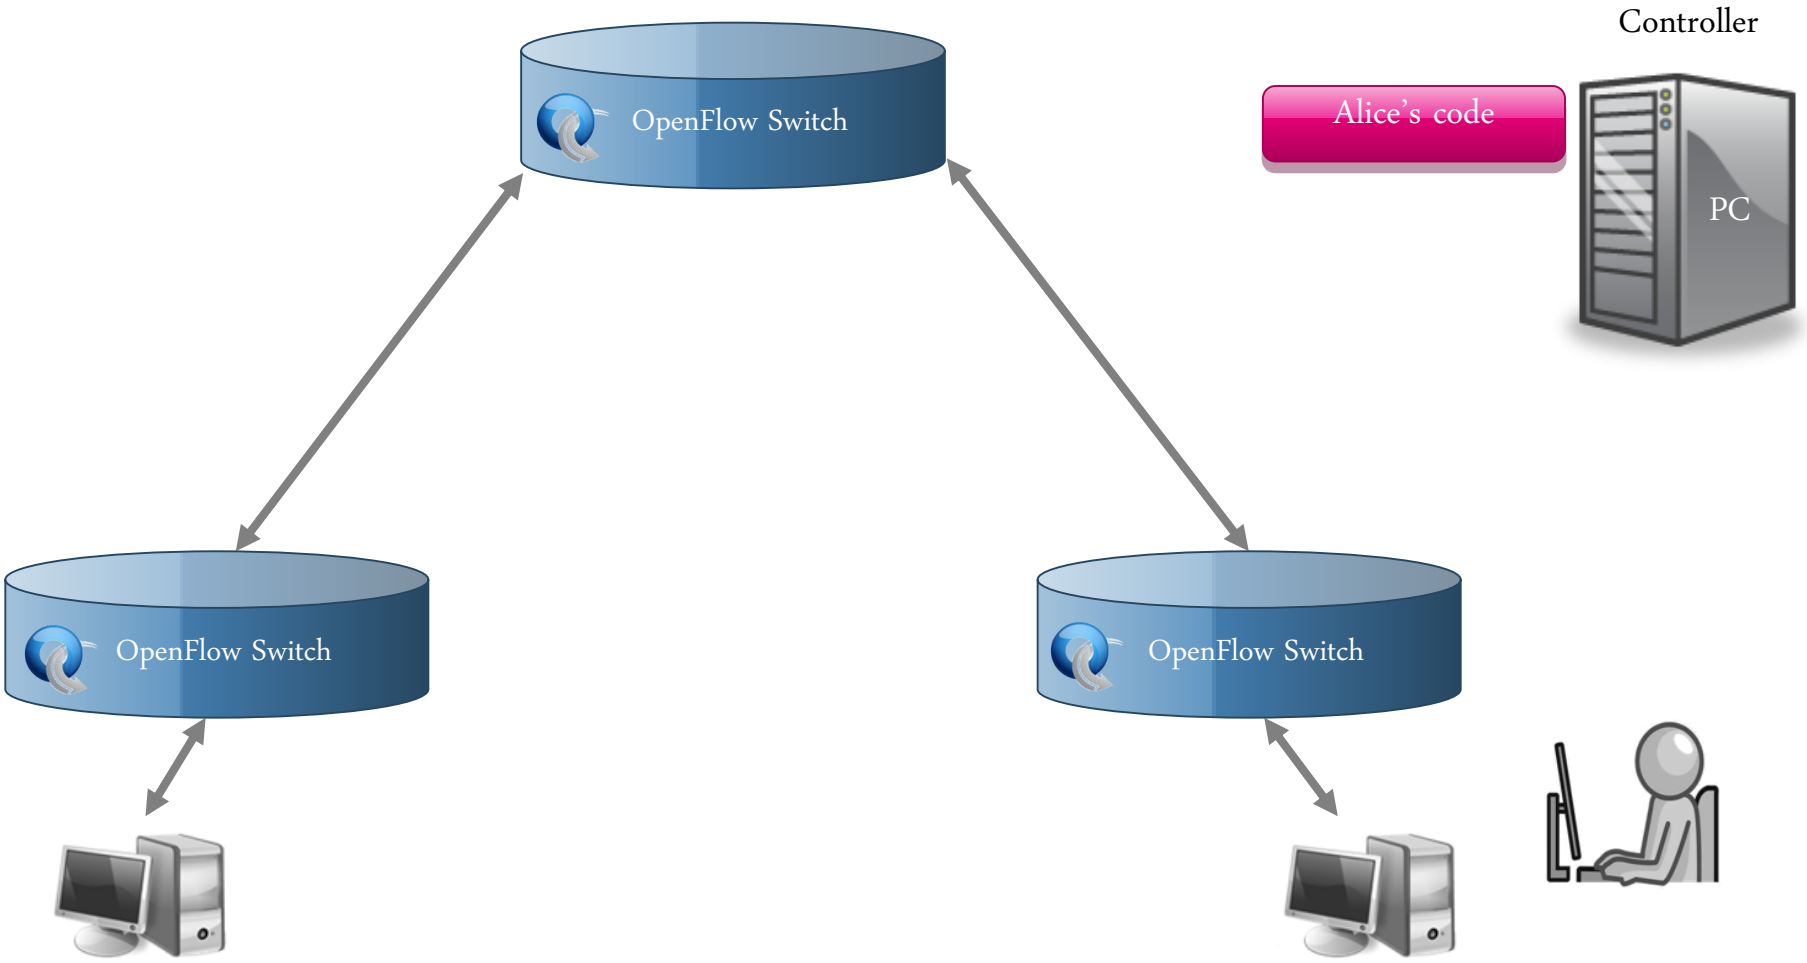

OpenFlow offloads control intelligence to a remote software

OpenFlow usage

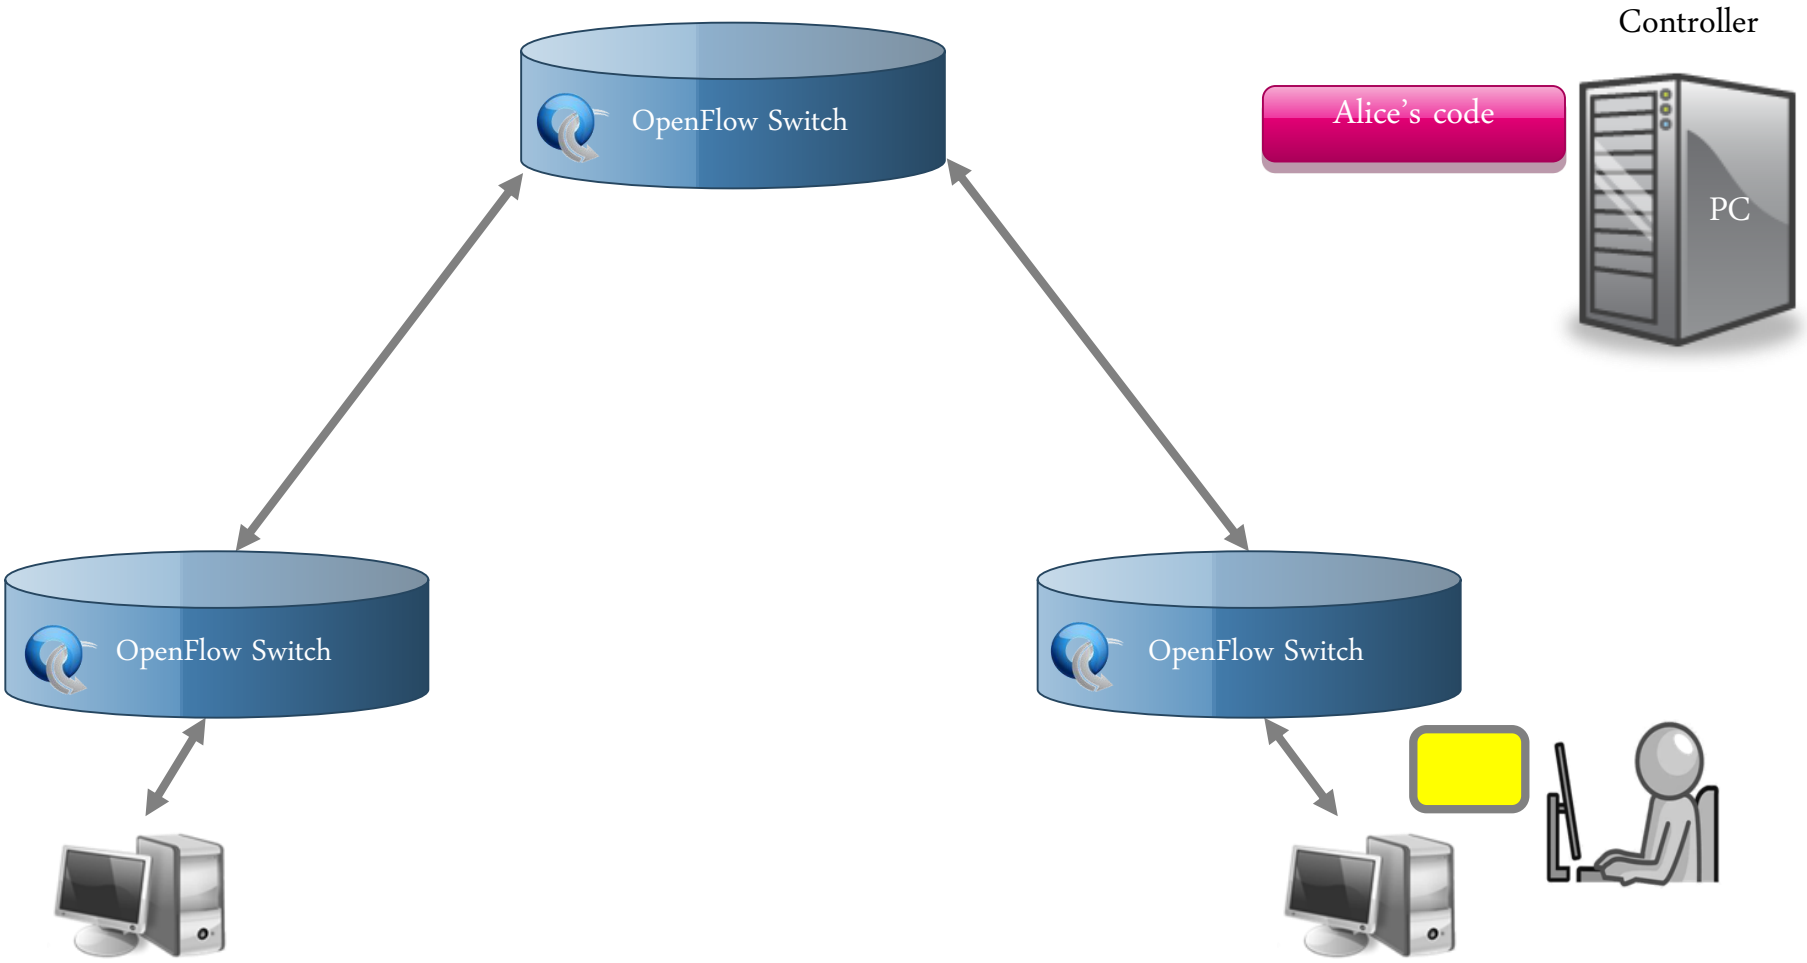

OpenFlow offloads control intelligence to a remote software

OpenFlow usage

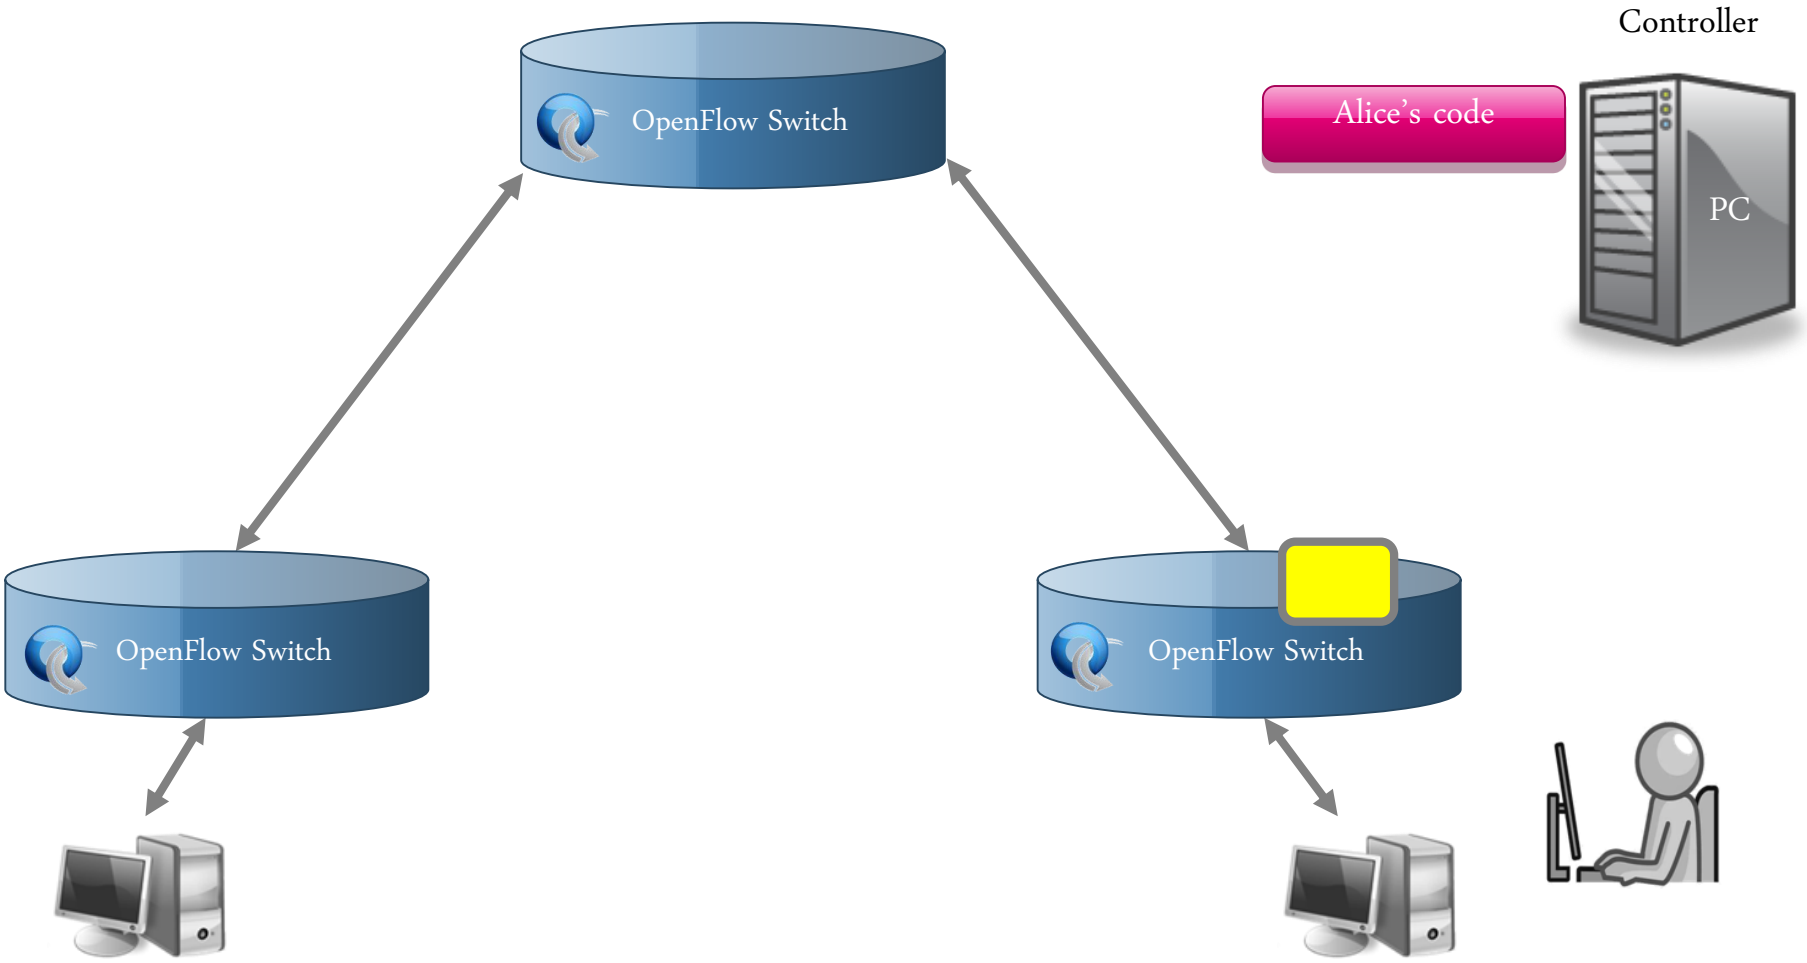

OpenFlow offloads control intelligence to a remote software

OpenFlow usage

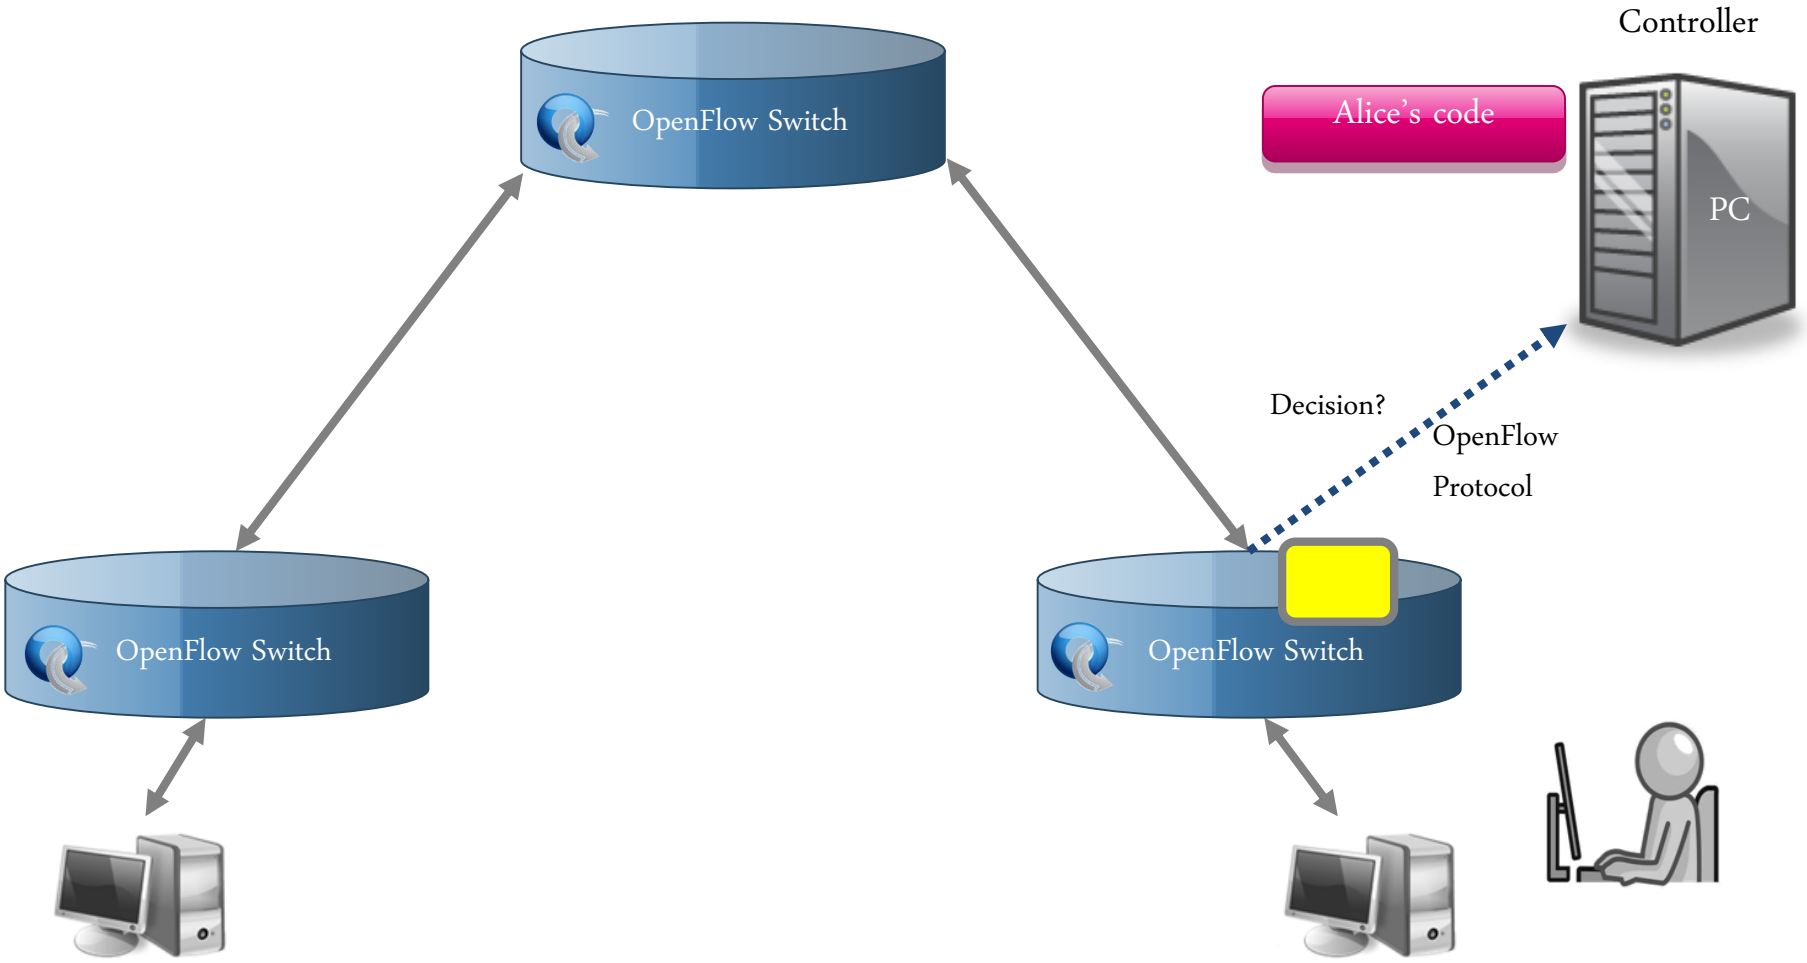

OpenFlow offloads control intelligence to a remote software

OpenFlow usage

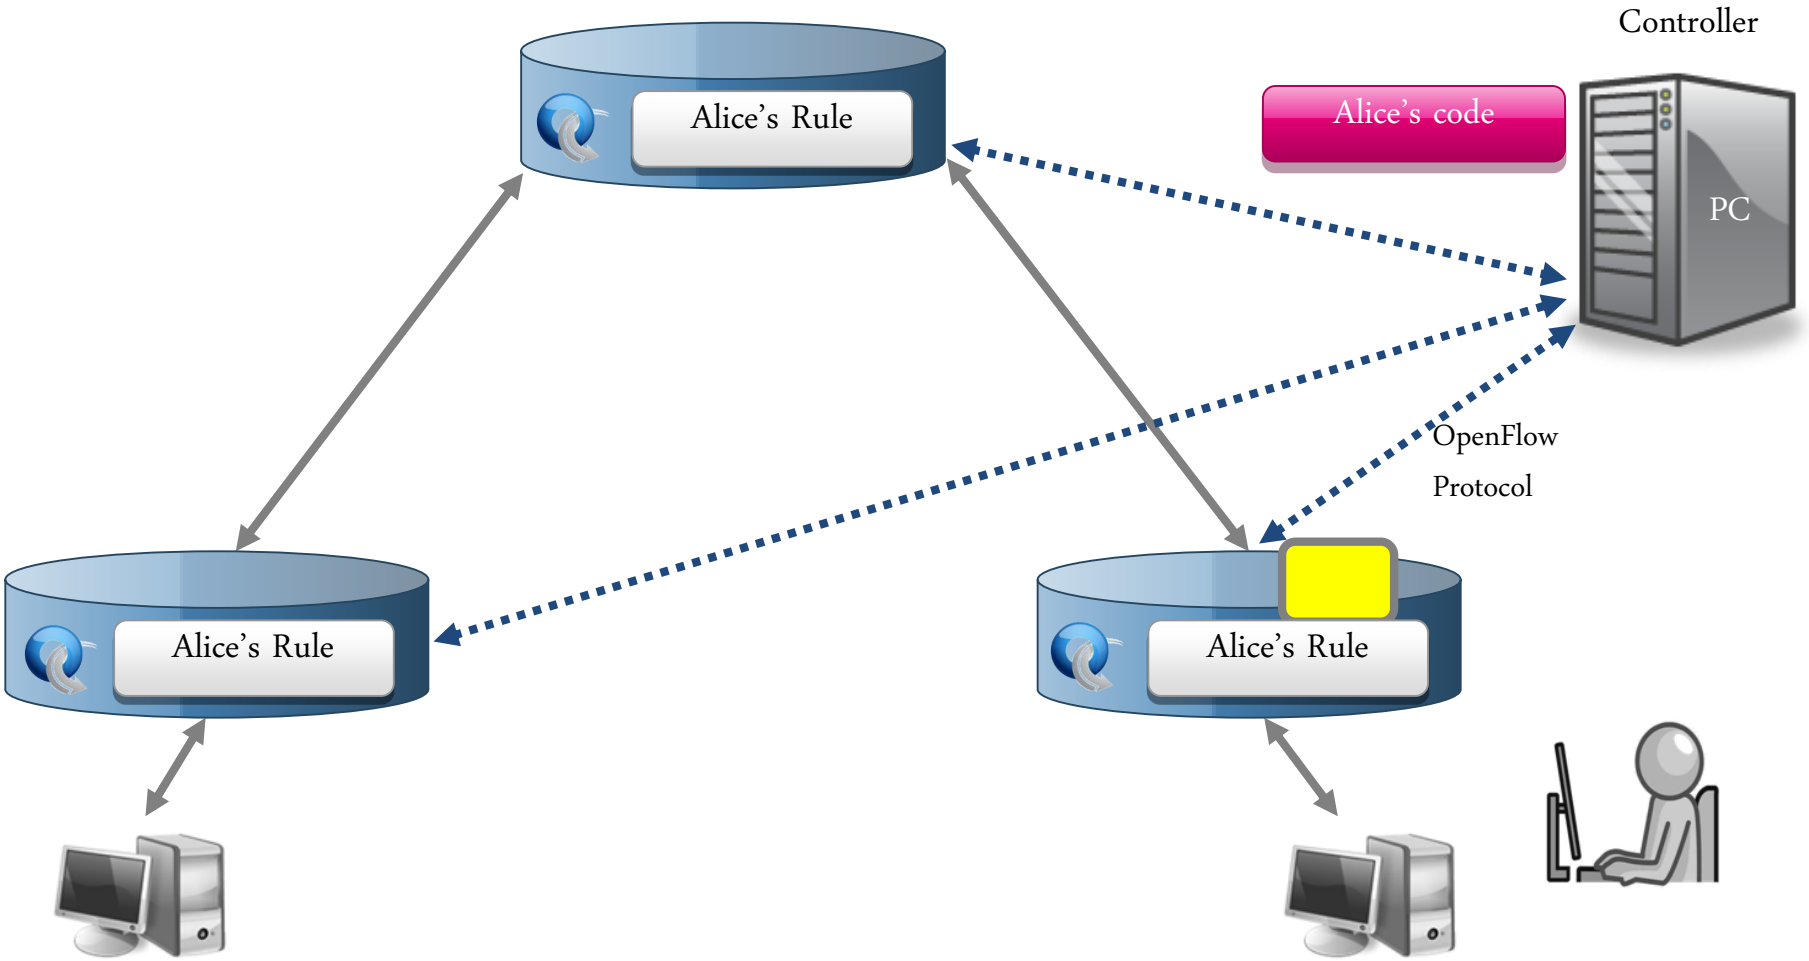

OpenFlow offloads control intelligence to a remote software

OpenFlow usage

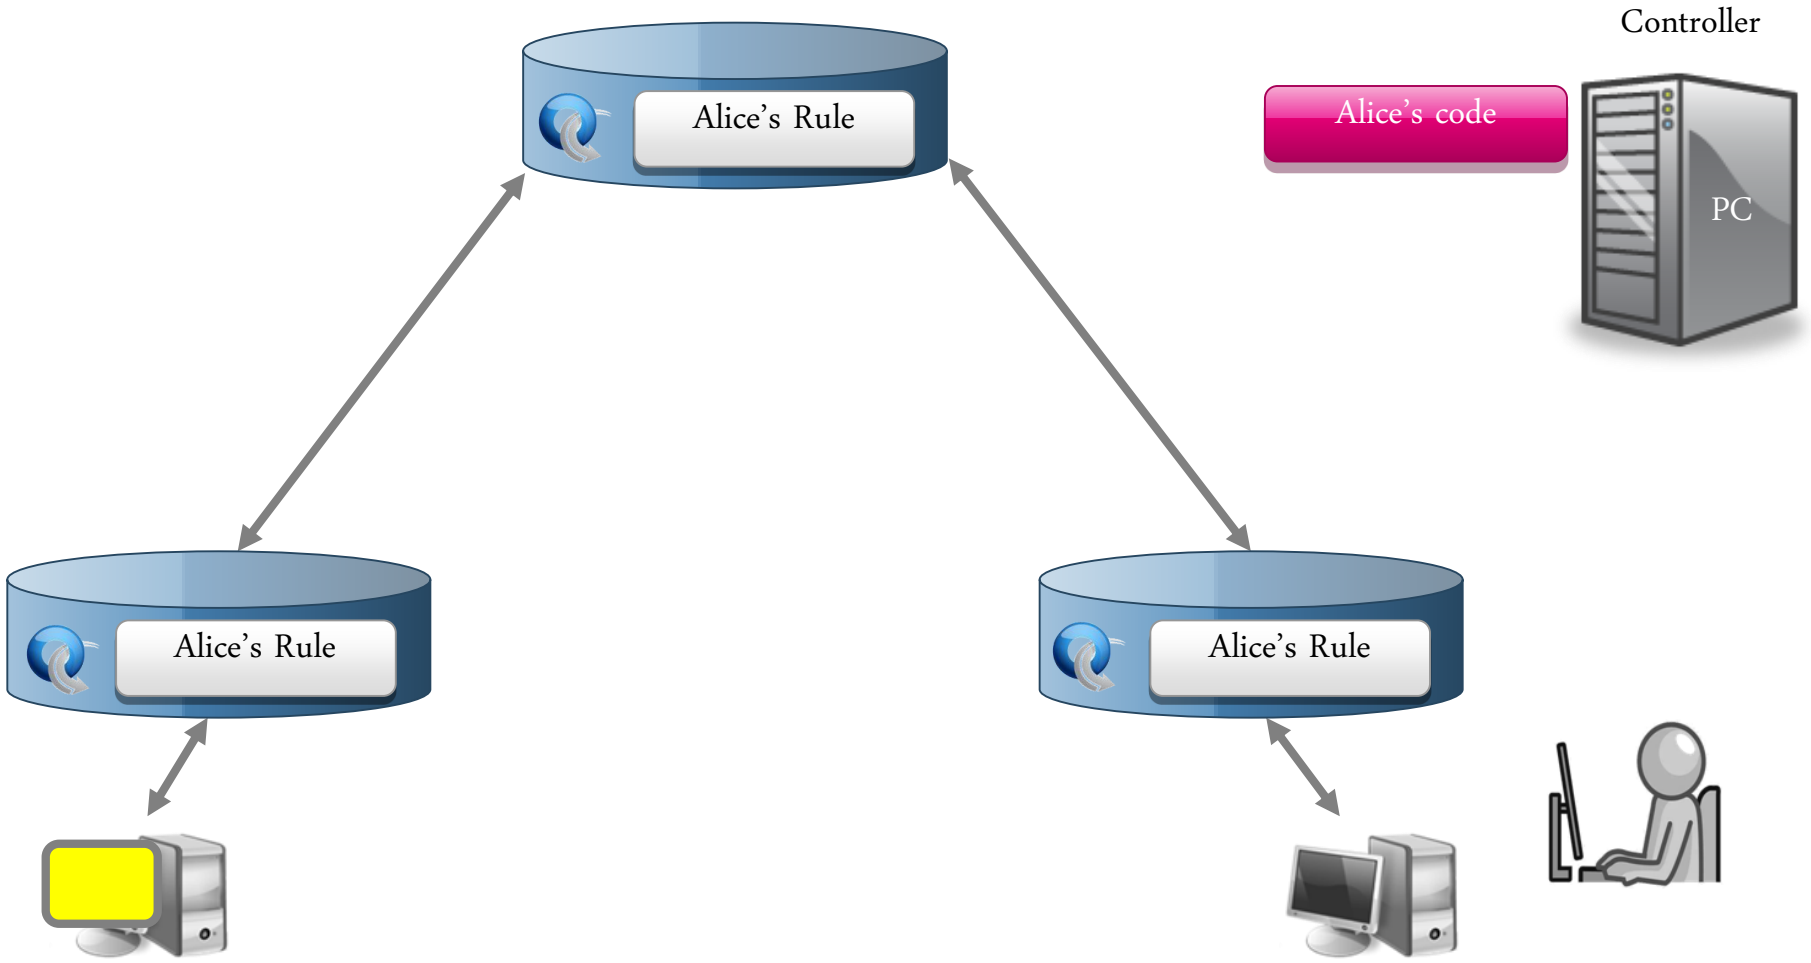

OpenFlow offloads control intelligence to a remote software

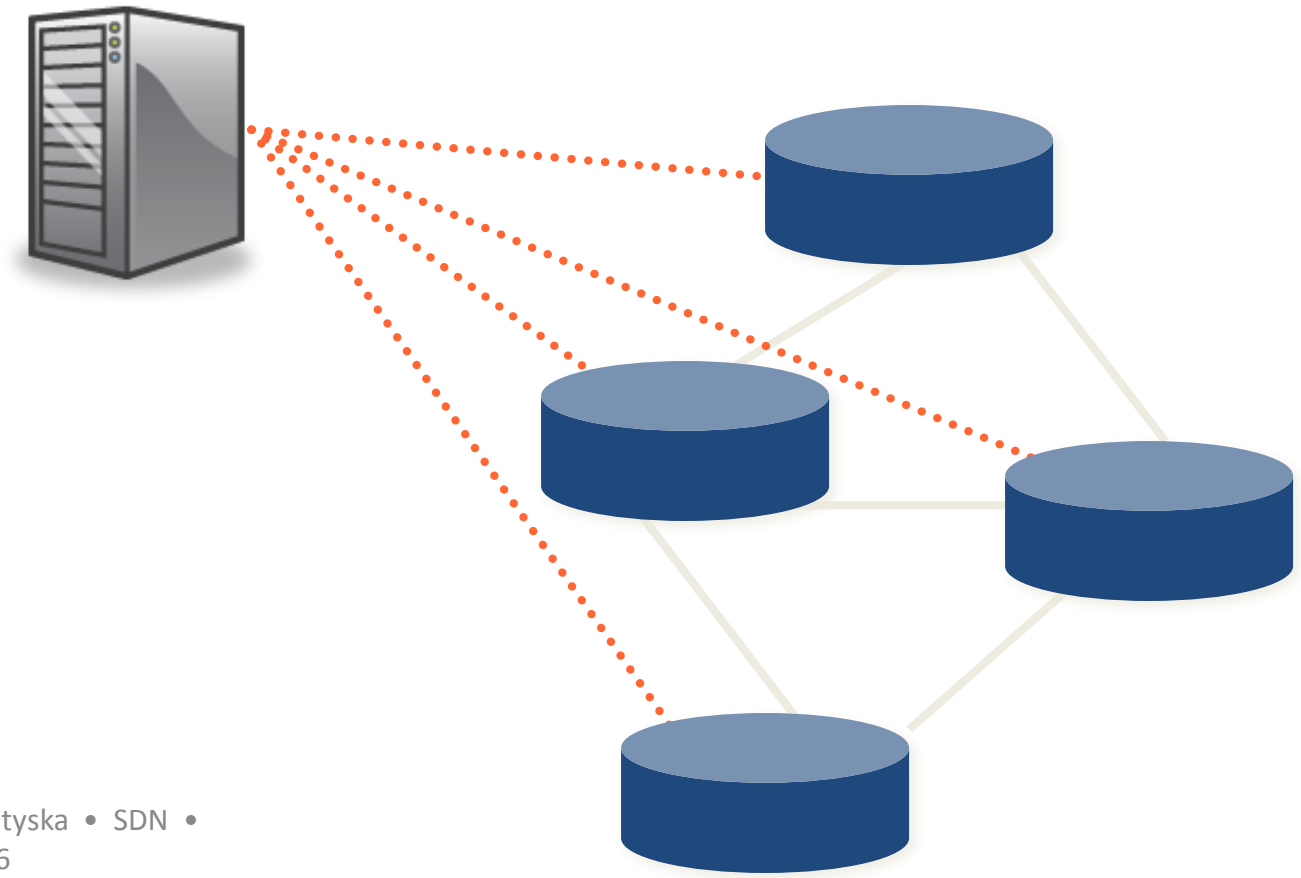

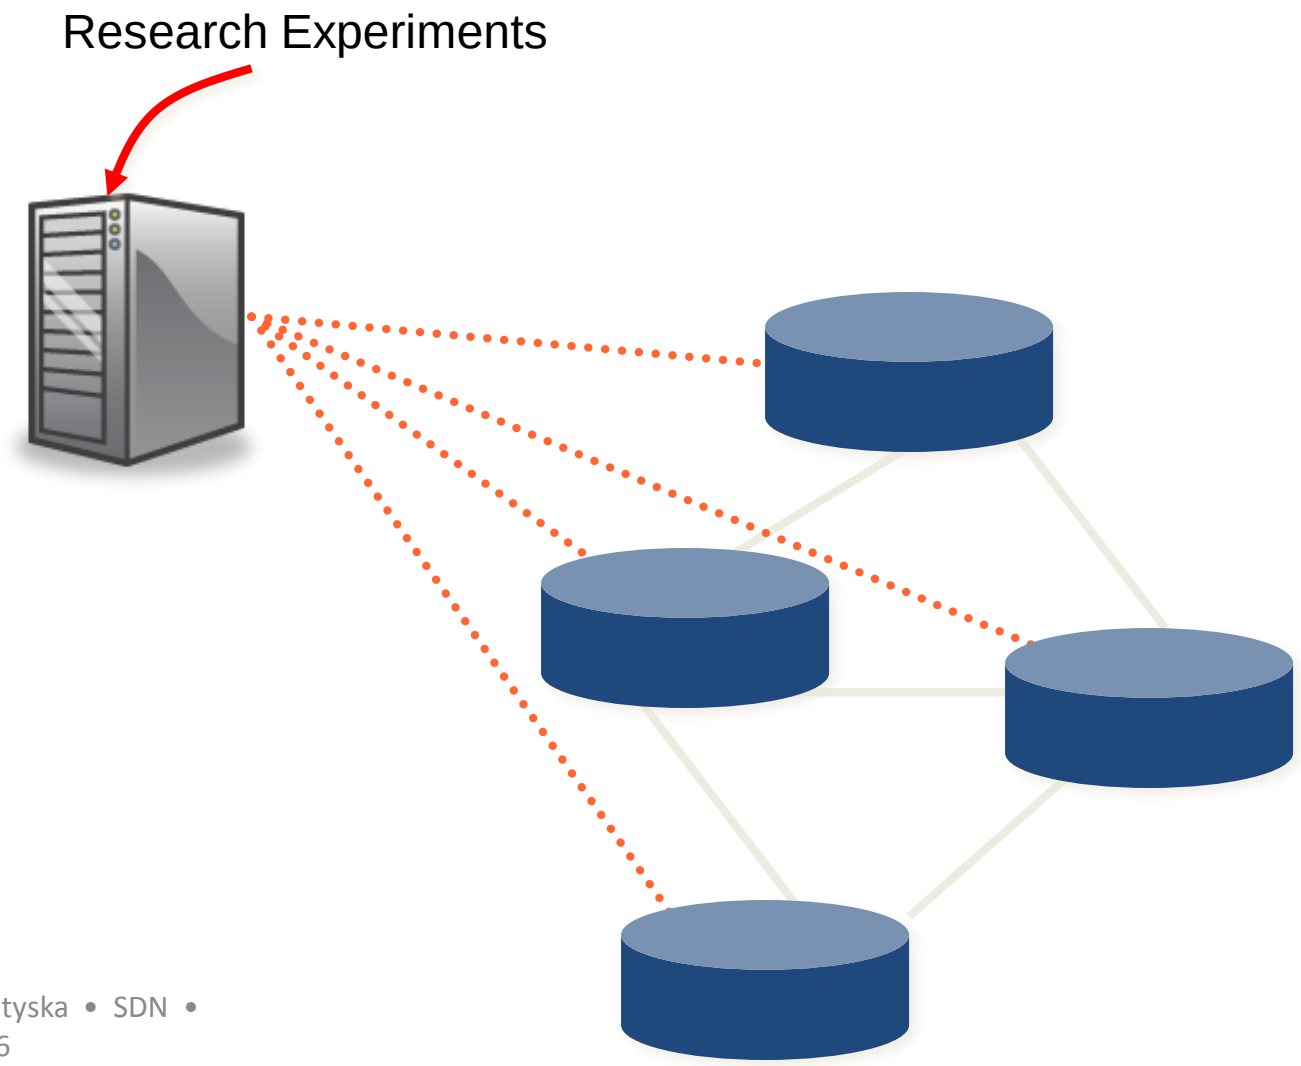
Research Experiments

“If header = ***x***, send to port 4”

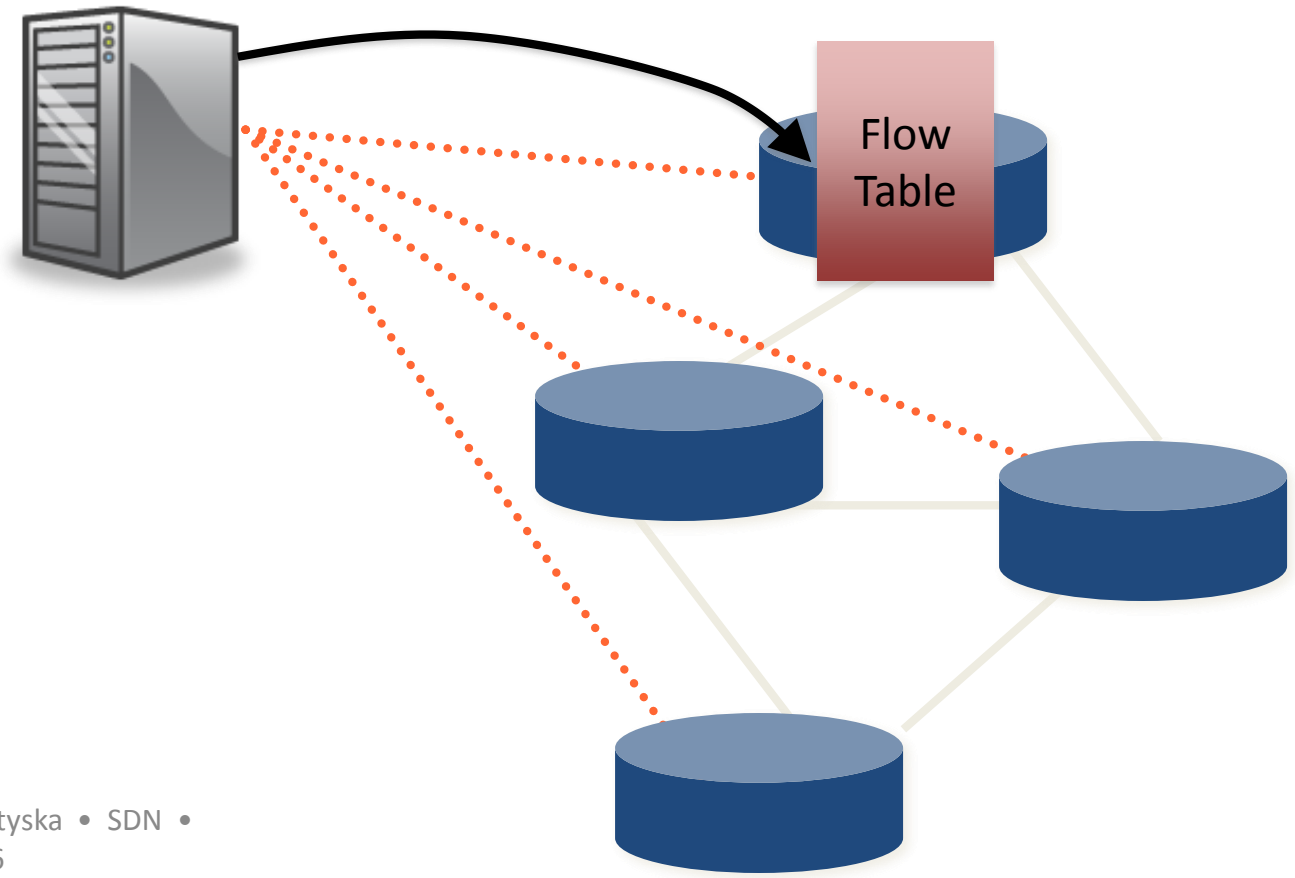

“If header = ***x***, send to port 4”
“If header = ***y***, overwrite header with ***z***, send to ports 5,6”

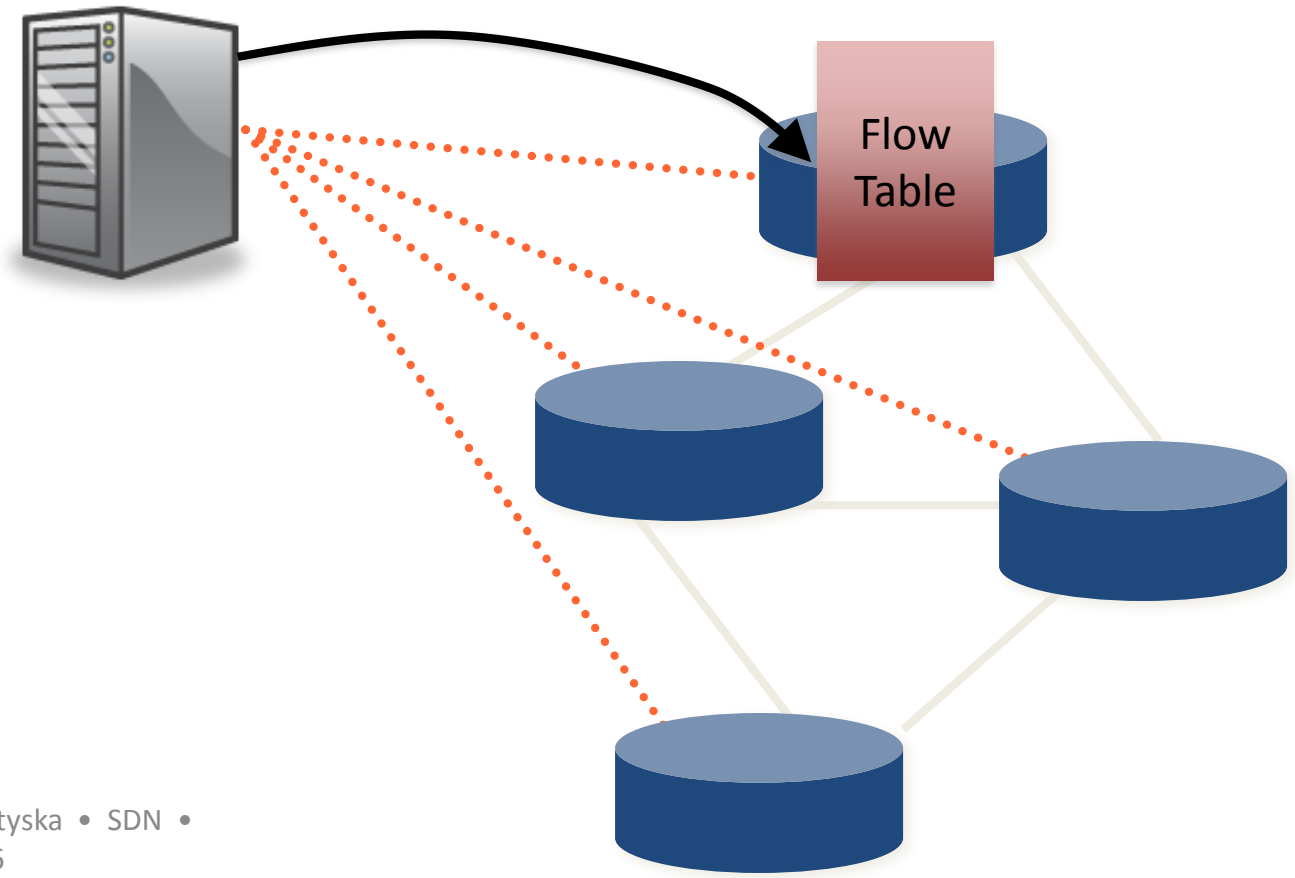

“If header = ***x***, send to port 4”
“If header = ***y***, overwrite header with ***z***, send to ports 5,6”
“If header = ***?***, send to me”

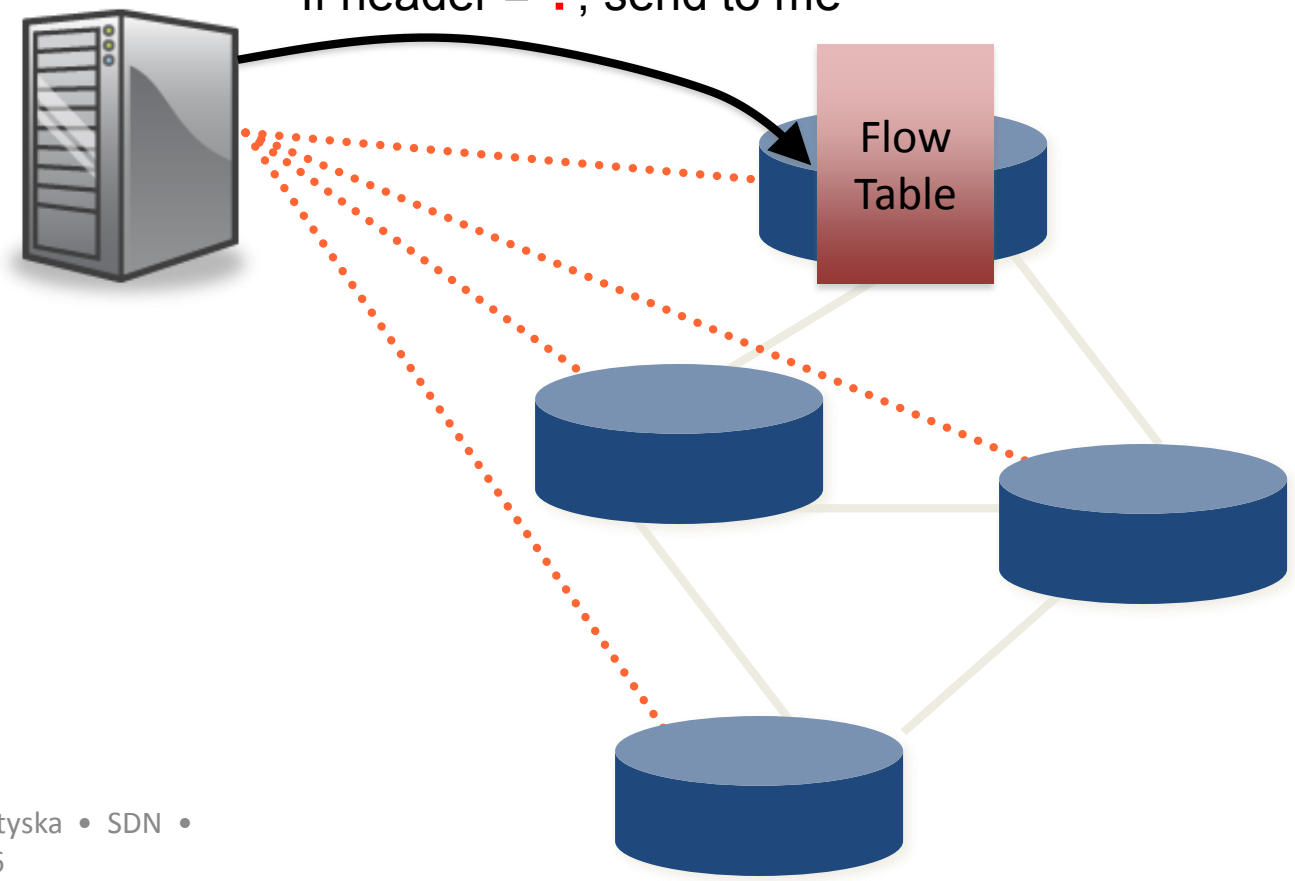

How does OpenFlow work?

SDN Essentials

- The steps below illustrate a simplified example interaction between an SDN controller and OpenFlow switch:
  - Step 1: Connection setup between OpenFlow switch and SDN Controller
  - Step 2: Proactive flow programming
  - Step 3: Topology discovery via LLDP
  - Step 4: Control plane maintenance and reactive flow programming
- The goal is not to exhaustively teach every OpenFlow message type
- Instead, this provides an illustration of how OpenFlow may operate to simplify a network use case (L2-Switch)

OpenFlow Basics

Flow Table Entries

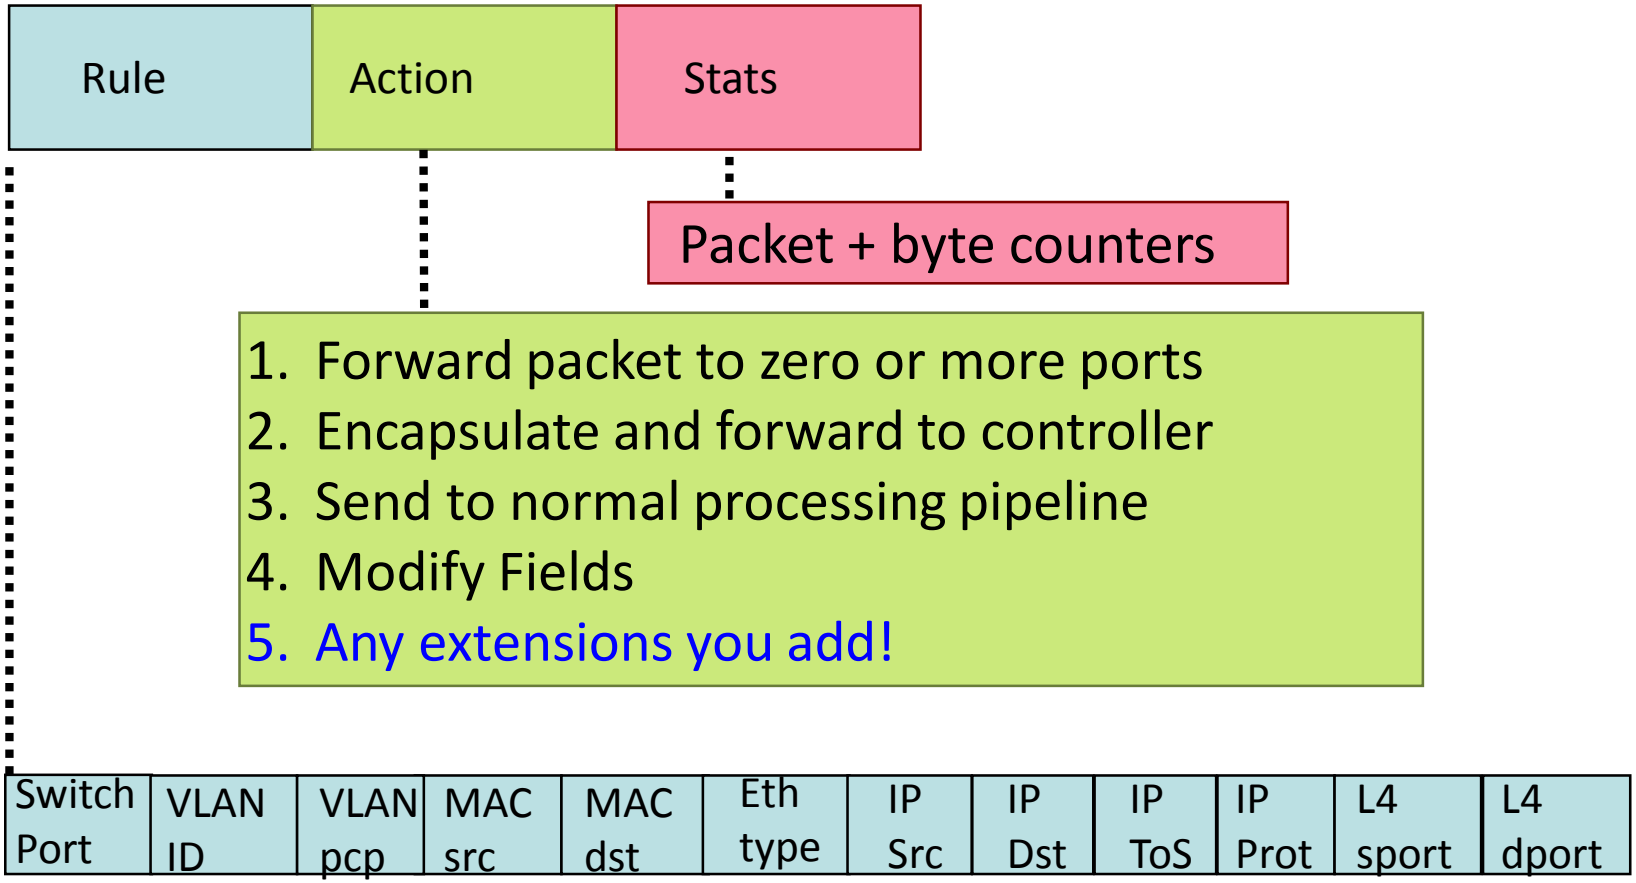

+ mask what fields to match

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	00:20..	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	Action
*	*	*	*	*	*	*	*	*	22	drop

Scalability & Robustness

Centralized vs Distributed Control

Both models are possible with OpenFlow

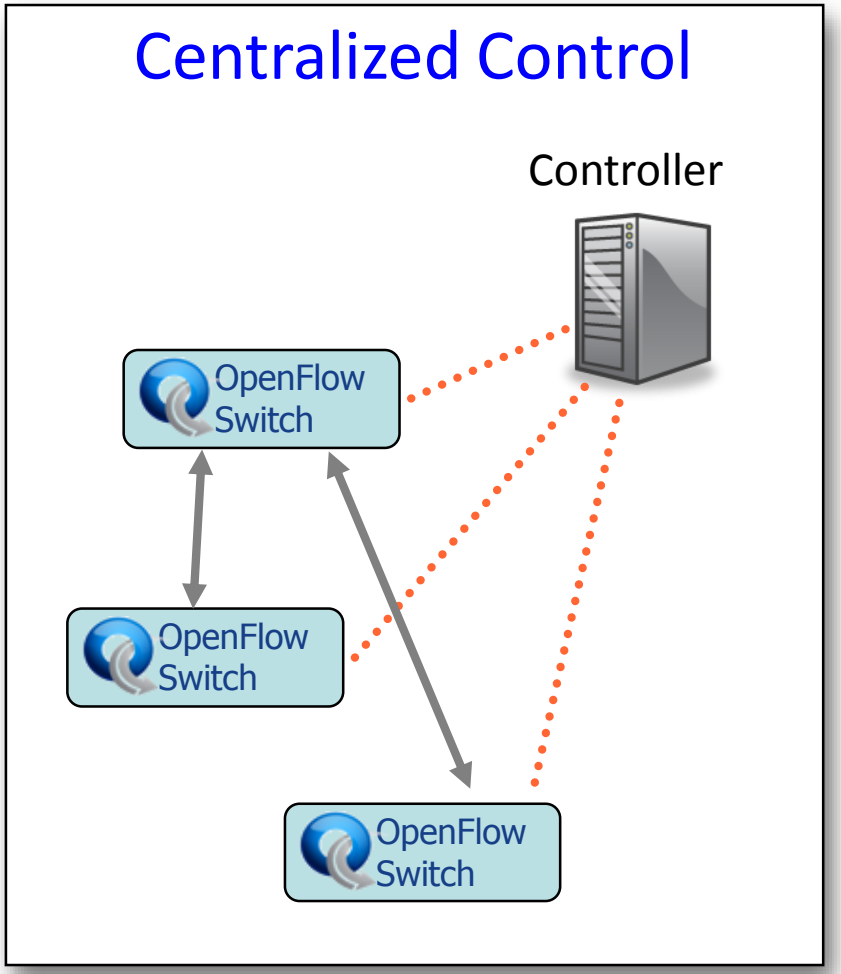

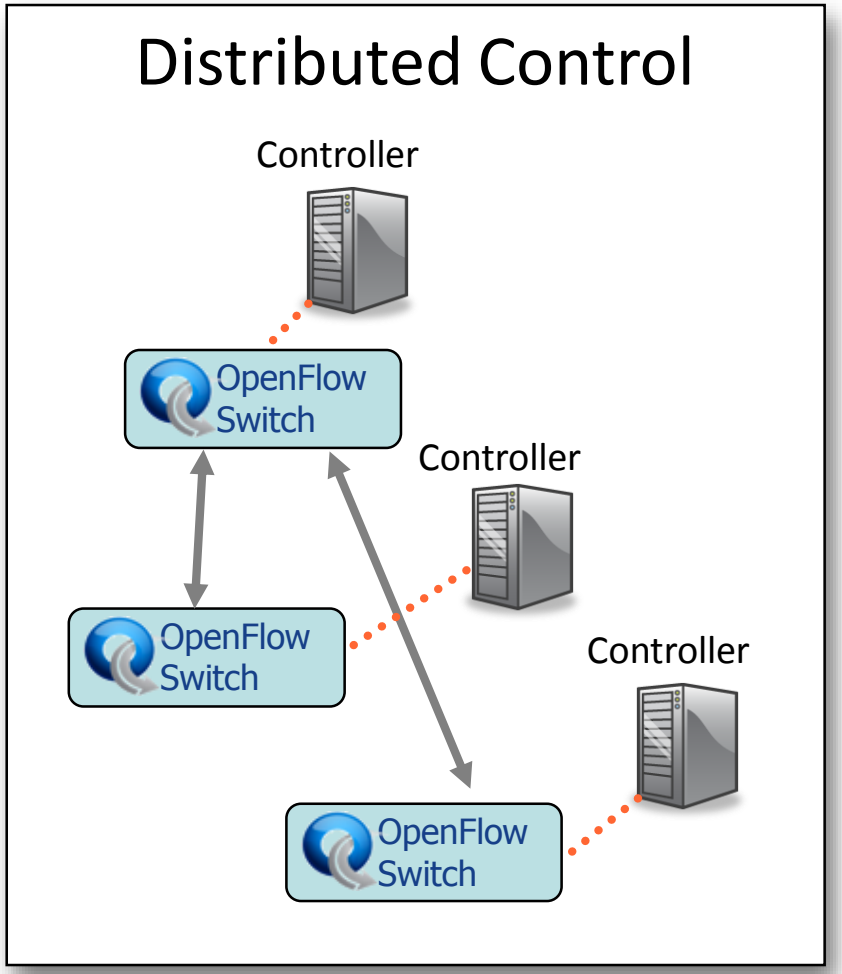

Flow Routing vs. Aggregation

Both models are possible with OpenFlow

Flow-Based

- Every flow is individually set up by controller
- Exact-match flow entries
- Flow table contains one entry per flow
- Good for fine grain control, e.g. campus networks

Aggregated

- One flow entry covers large groups of flows
- Wildcard flow entries
- Flow table contains one entry per category of flows
- Good for large number of flows, e.g. backbone

Reactive vs. Proactive (pre-populated)

Both models are possible with OpenFlow

Reactive

- First packet of flow triggers controller to insert flow entries
- Efficient use of flow table
- Every flow incurs small additional flow setup time
- If control connection lost, switch has limited utility

Proactive

- Controller pre-populates flow table in switch
- Zero additional flow setup time
- Loss of control connection does not disrupt traffic
- Essentially requires aggregated (wildcard) rules

Network Function Virtualization (NFV)

Today's network infrastructure

- Diverse network functions (NF).
- Providing desired overall functionality or service.
- Adding new services
  - New service instances take weeks to activate
  - New service types may take months up to years
  - New service types require either new equipment or upgrading of existing equipment
- We have to simplify network design, increase agility, speed up deployment of new services.

What is NFV ?

- NFV – Network Functions Virtualization
- Likewise VM
- Virtualization – NF and part of the infrastructure is implemented as a software.
- Result – from dedicated proprietary appliance to COTS hardware

NFV in nutshell

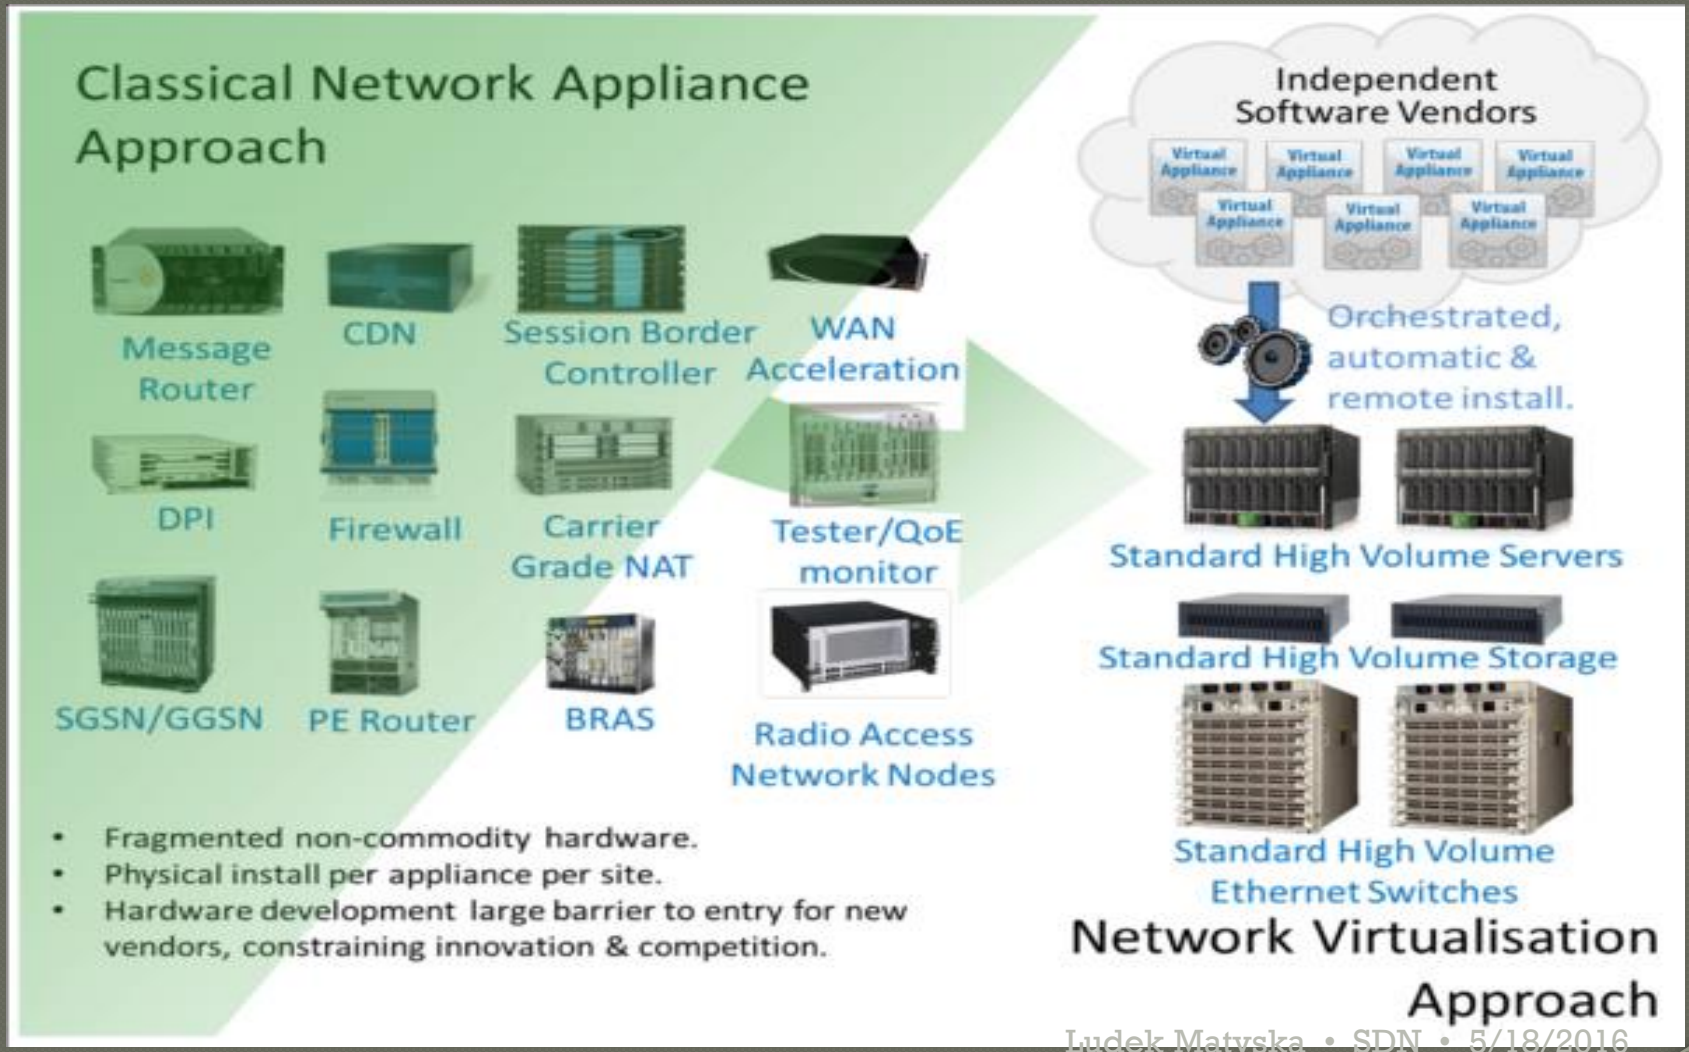

NFV in nutshell

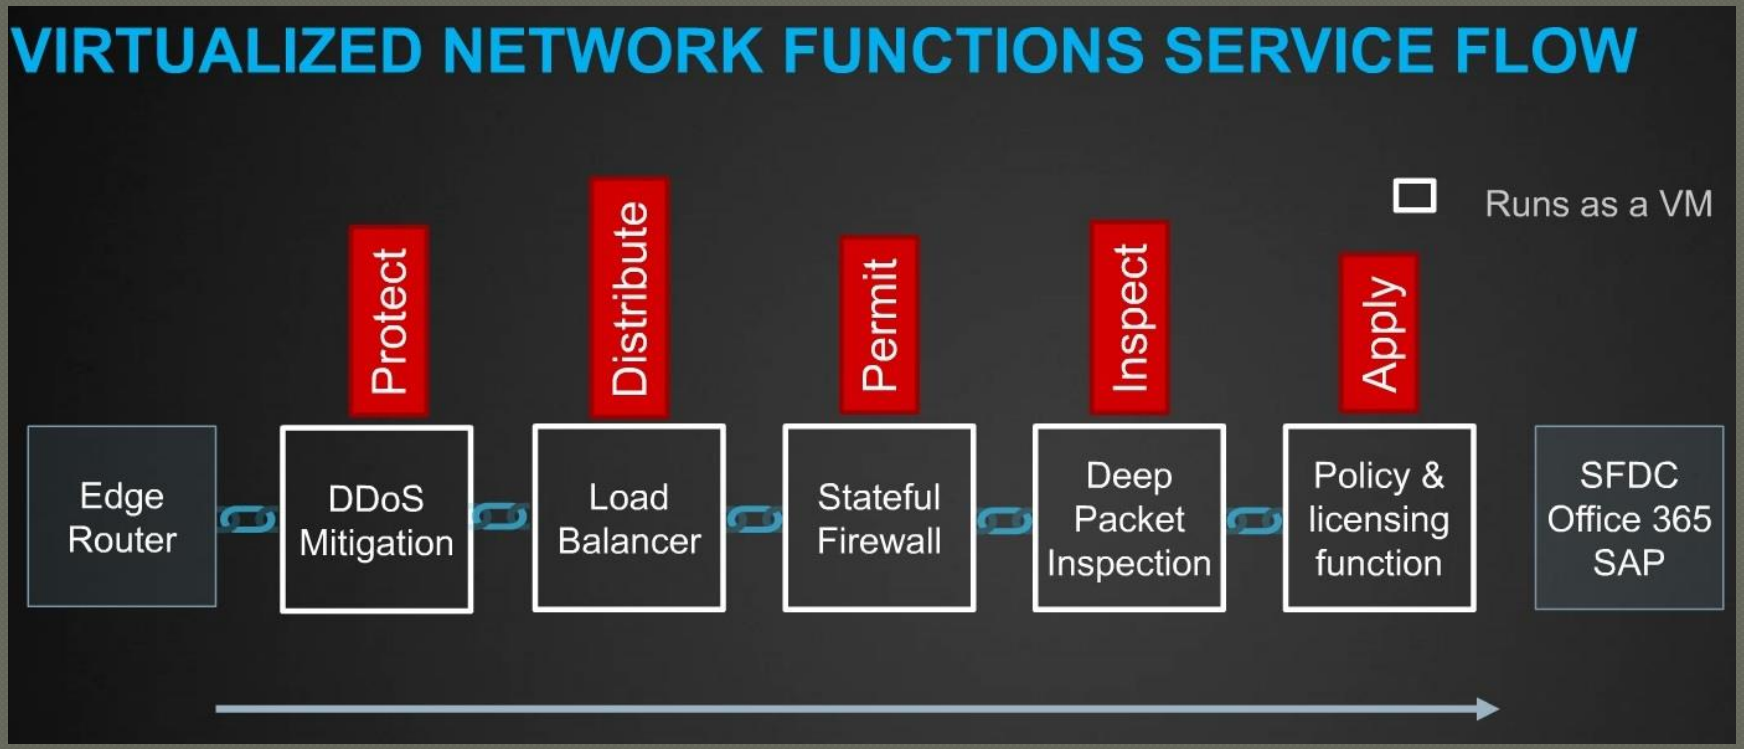

NFV in nutshell

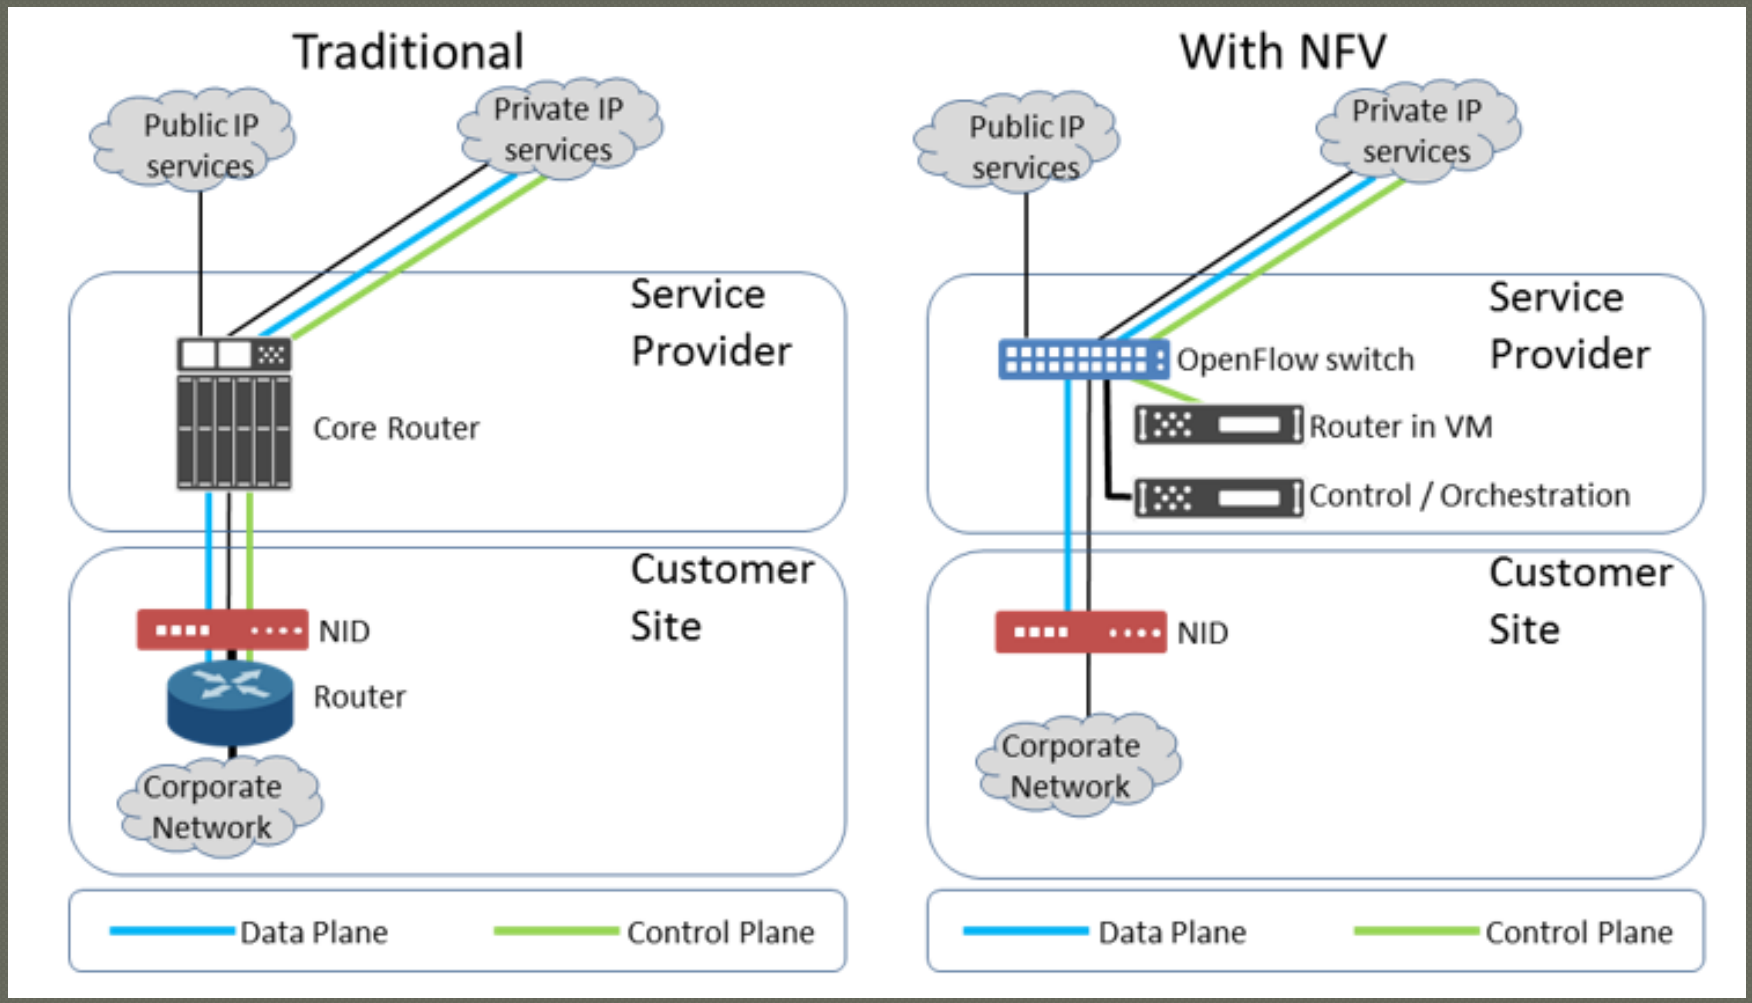

SDN controllers

SDN controllers

Common objectives:

- **multiple Southbound interface protocol support**
- **well-defined Northbound API support**
- **programmability**
- **high availability & performance**
- **security**

Open-source vs. commercial

2016 Controller Landscape – OPEN-SOURCE

Active	Not Active (Apparently)
Floodlight	Beacon
LOOM	FlowER
OpenContrail*	NOX/POX
OpenDaylight*	NodeFlow
OpenMUL	
ONOS*	
Ryu*	
Trema	

* - more prominent

2016 Controller Landscape - COMMERCIAL

ODL-based	ODL-friendly	Non-ODL Based
ADVA	NEC	Big Switch
Avaya	Nokia/Nuage Networks	Juniper (Contrail/Northstar)
Brocade	Oracle	Midokura
Ciena (also proprietary)	Pluribus	Plexxi
Cisco (also proprietary)		PLUMgrid
Coriant		Sonus (Vello Systems)
Dell		VMware NSX
Ericsson		
Extreme		
Fujitsu		
HPE (also proprietary)		
Huawei (also proprietary)		
Inocybe		

Updated in 2016 Feb from original source: https://www.sdxcentral.com/reports/sdn-controllers-2015/

OpenDaylight architecture illustration

SDN Essentials

OPEN DAYLIGHT - Lithium

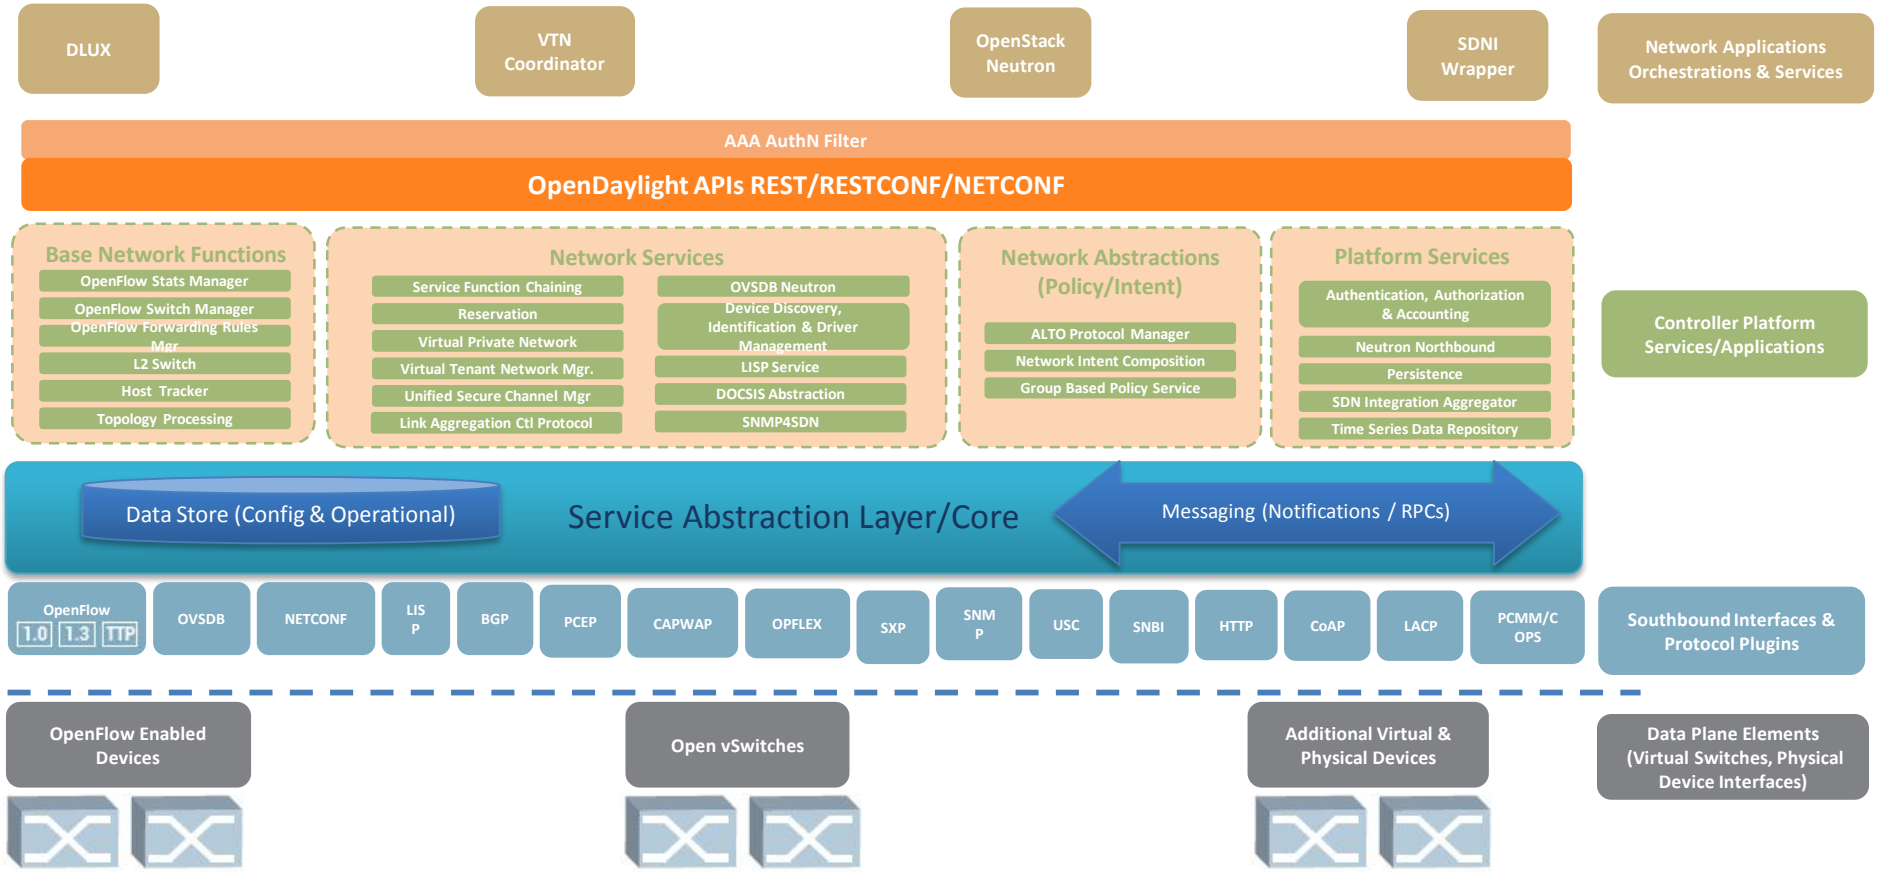

NETCONF & YANG

different SDN view...

NETCONF and YANG in Context

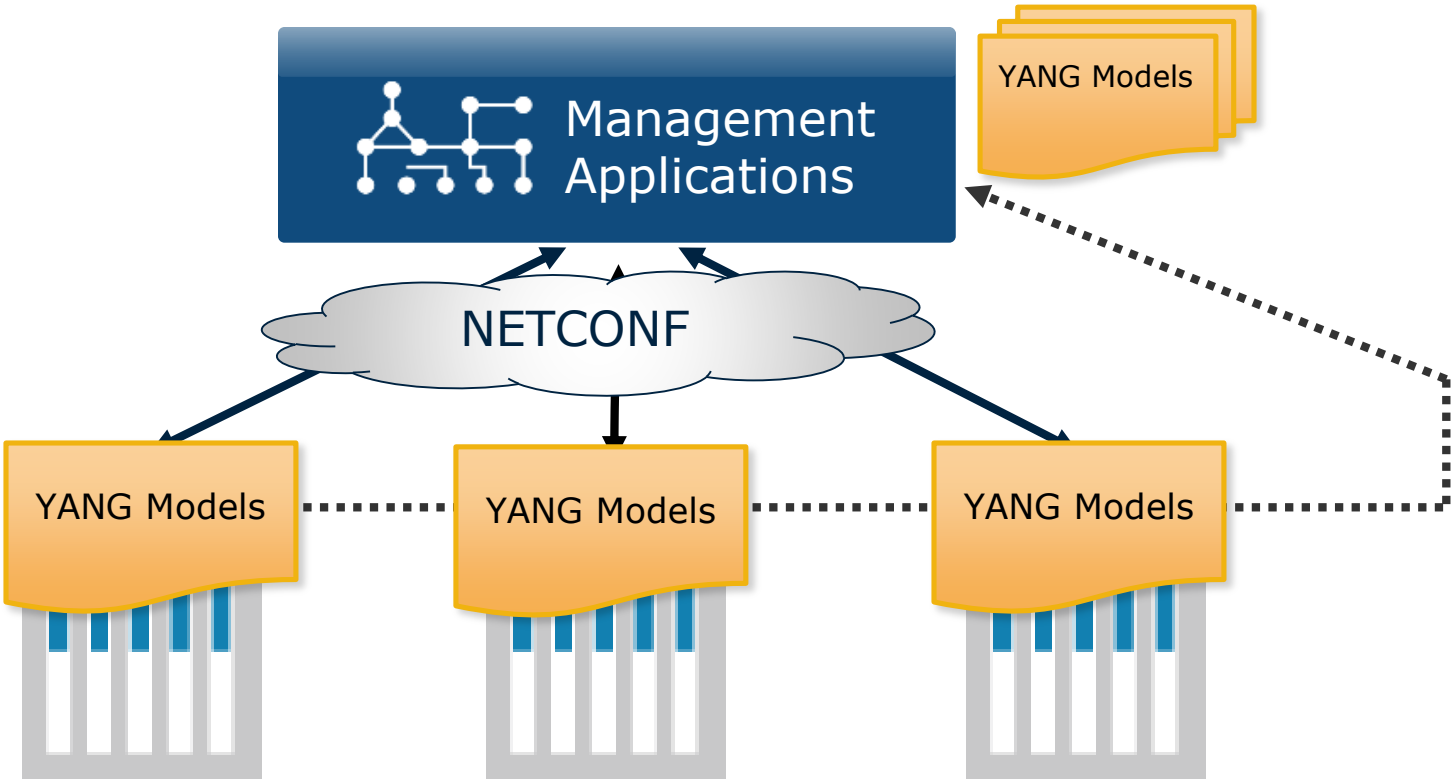

What is a Data-Model? What is a Network Management Protocol?

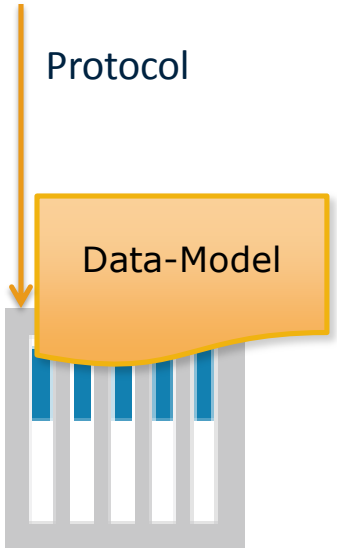

- Data-Model
  - A data-model explicitly and precisely determines the structure, syntax and semantics of the data…
  - …that is *externally* visible
  - Consistent and complete
- Protocol
  - Remote primitives to view and manipulate the data
  - Encoding of the data as defined by the data-model

Confusion and Comparison

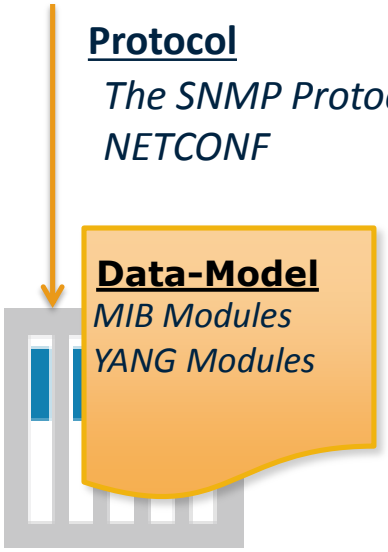

- Data-Models and information Models
  - Information models are for humans
    - Not everything
    - Not always detailed and precise
    - Starting-point for data-model
- Protocol
  - Confusion between domain-specific network management protocols and general RPC mechanisms
    - NETCONF vs. CORBA, SOAP, REST, …

Standards background, motivation and history

RFC 3535: Operators' problems and requirements on network management

Informational RFC 3535

Abstract

This document provides an overview of a workshop held by the Internet Architecture Board (IAB) on Network Management. The workshop was hosted by CNRI in Reston, VA, USA on June 4 thru June 6, 2002. The goal of the workshop was to continue the important **dialog** started between **network operators** and protocol developers, and to guide the IETFs focus on future work regarding network management.

- SNMP had failed
  - For configuration, that is
  - Extensive use in fault handling and monitoring
- CLI scripting
  - "Market share" 70%+

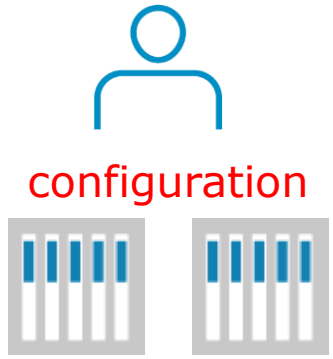

Operator Requirement #1/14

1. **Ease of use** is a key requirement for any network management technology from the operators point of view.

Maybe not assume integrators and software developers for any addition or change

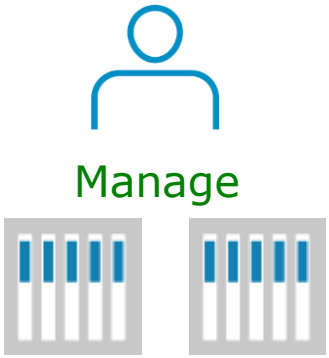

Operator Requirement #2-3/14

2. It is necessary to make a **clear distinction** between **configuration data**, data that describes **operational state and statistics**.

3. It is required to be able to **fetch separately configuration data**, operational state data, and statistics from devices, and to be able to compare these between devices.

- Clearly separating configuration
- Ability to compare across devices

```
$show running-config
```

Operator Requirement #4-5/14

4. It is necessary to enable operators to concentrate on the **configuration of the network** as a whole rather than individual devices.

5. Support for **configuration transactions** across a number of devices would significantly simplify network configuration management.

- Service and Network management, not only device management
- Network wide transactions

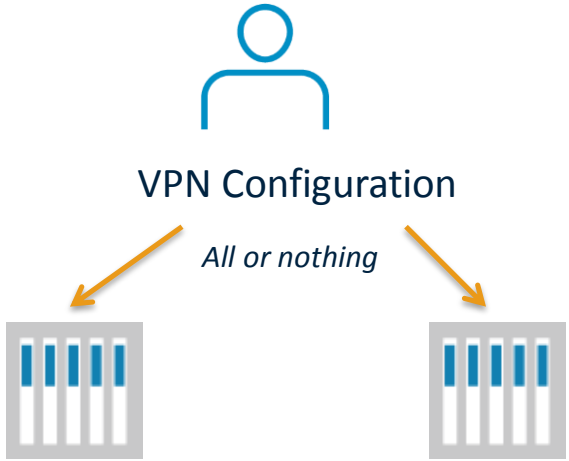

Operator Requirement #6-7/14

6. Given configuration A and configuration B, it should be possible to generate the **operations necessary to get from A to B** with minimal state changes and effects on network and systems. It is important to minimize the impact caused by configuration changes.

7. A mechanism to dump and restore configurations is a primitive operation needed by operators. Standards for **pulling and pushing configurations** from/to devices are desirable.

- Devices figure out ordering
- No unnecessary changes
- Finally: backup/restore of configuration

The litmus-test of a management interface

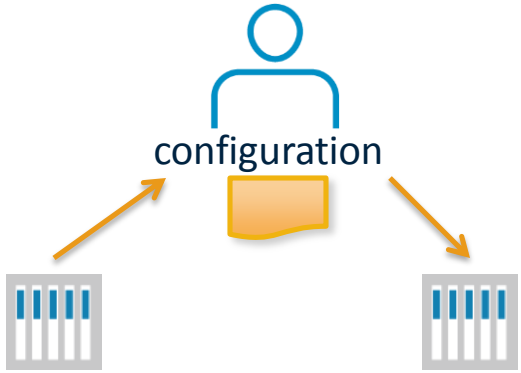

Operator Requirement #8, 10/14

8. It must be easy to do **consistency** checks of configurations over time and between the ends of a link in order to determine the changes between two configurations and whether those configurations are consistent.

10. It is highly desirable that **text** processing tools such as diff, and version management tools such as RCS or CVS, can be used to process configurations, which implies that devices should not arbitrarily reorder data such as access control lists.

- Validation of configuration
- Validation at network level
- Text based configuration

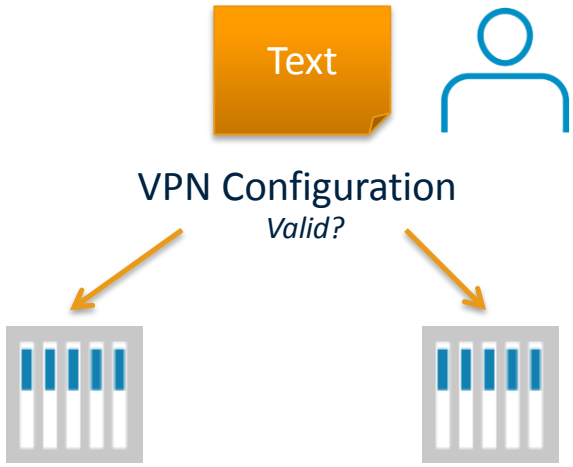

Operator Requirement #9/14

9. Network wide configurations are typically stored in central master databases and transformed into formats that can be pushed to devices, either by generating sequences of CLI commands or complete configuration files that are pushed to devices. There is no **common database schema** …, although the models used by various operators are probably very similar.

It is desirable to extract, document, and standardize the common parts of these network wide configuration database schemas.

- Standardized data models

Interfaces Data-Model

Operator Requirement #13/14

13. It is important to distinguish between the distribution of configurations and the activation of a certain configuration.

Devices should be able to hold multiple configurations.

- Support for multiple configuration sets
- Delayed, orchestrated activation

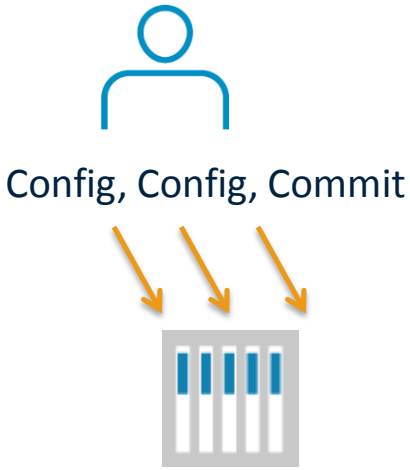

NETCONF was designed to conform to RFC 3535.

Today many operators require NETCONF and YANG in devices.

NETCONF makes a difference on the bottom line.

What makes NETCONF/YANG different?

	SNMP	NETCONF	SOAP	REST
Standard	IETF	IETF	W3C	-
Resources	OIDs	Paths		URLs
Data models	Defined in MIBs	YANG Core Models		
Data Modeling Language	SMI	YANG	(WSDL, not data)	Undefined, (WSDL), WADL, text…
Management Operations	SNMP	NETCONF	In the XML Schema, not standardized	HTTP operations
Encoding	BER	XML	XML	XML, JSON,…
Transport Stack	UDP	SSH TCP	SSL HTTP TCP	SSL HTTP TCP

"RESTConf"

What makes NETCONF/YANG different?

SNMP

- GET
- GET-NEXT
- SET
- TRAP
- …

… so what?

NETCONF

- <get-config>
- <edit-config>
- <copy-config>
- <delete-config>
- <get>
- <lock>
- …

… same same?

What makes NETCONF/YANG different?

This is where the difference is:
In the supported use cases!

Use Case	SNMP	NETCONF
Get collection of status fields	Yes	Yes. Bulk xfer up to 10x faster. Really.
Set collection of configuration fields	Yes, up to 64kB	Yes
Set configuration fields in transaction	No	Yes
Transactions across multiple network elements	No	Yes
Invoke administrative actions	Well…	Yes
Send event notifications	Yes	Yes, connected
Backup and restore configuration	Usually not	Yes
Secure protocol	v3 is fair	Yes
Test configuration before final commit	No	Yes

YANG ?

- Data modeling language
  - Configuration data
  - State data
- Tree structure
- Data and Types

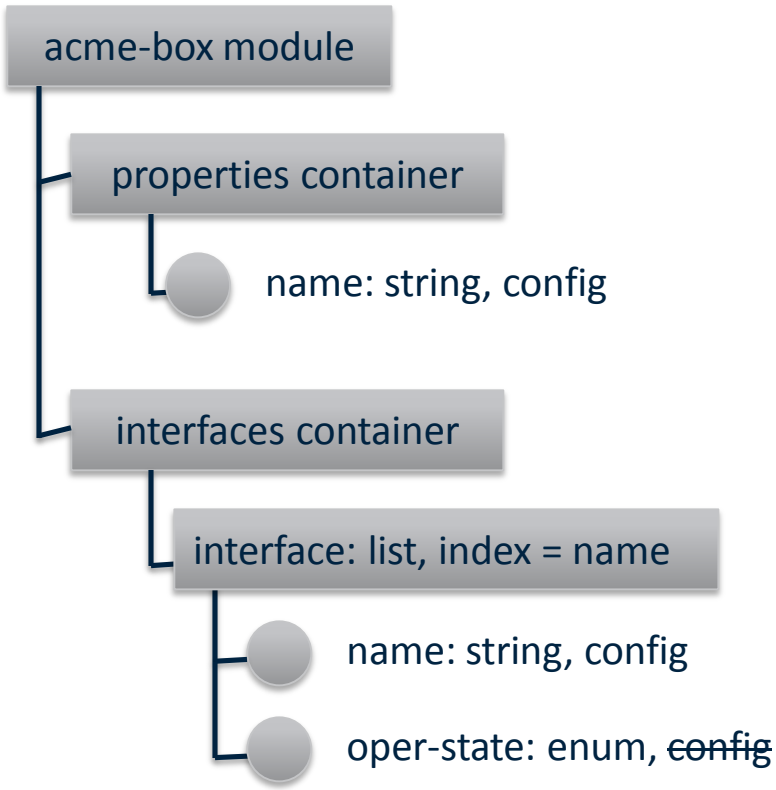

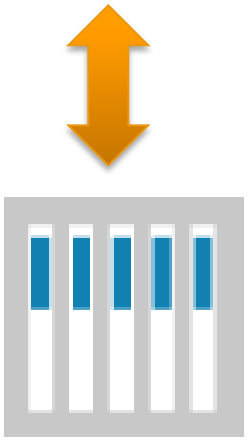

YANG Module Contents

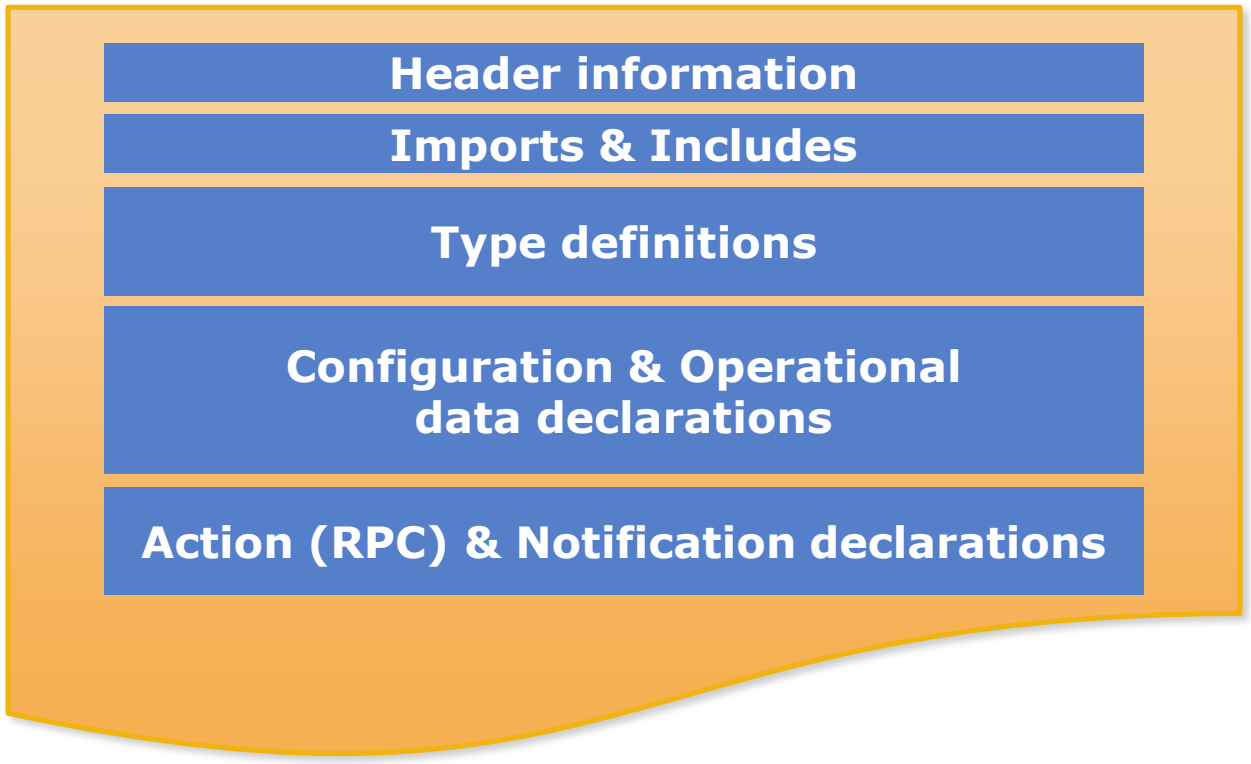

YANG Header

```
module acme-module {
  namespace "http://acme.example.com/module";
  prefix acme;

  import "ietf-yang-types" {
    prefix yang;
  }
  include "acme-system";

  organization "ACME Inc.";
  contact joe@acme.example.com;
  description "Module describing the ACME products";
  revision 2007-06-09 {
    description "Initial revision.";
  }
```

Thank you for your attention!