


Why Software Defined Networking (SDN)?

Boyan Sotirov

Agenda


- ▶ Current State of Networking
- ▶ Why
- ▶ What
- ▶ How
- ▶ When

Conventional Networking




- Many complex functions embedded into the infrastructure
 - OSPF, BGP, Multicast, NAT, TE, MPLS, Firewalls...
 - Redundant layers, services
 - Unique “differentiation”
- Mainframe mentality industry
- Functionality → standards → hardware → nodes

Conventional Networking



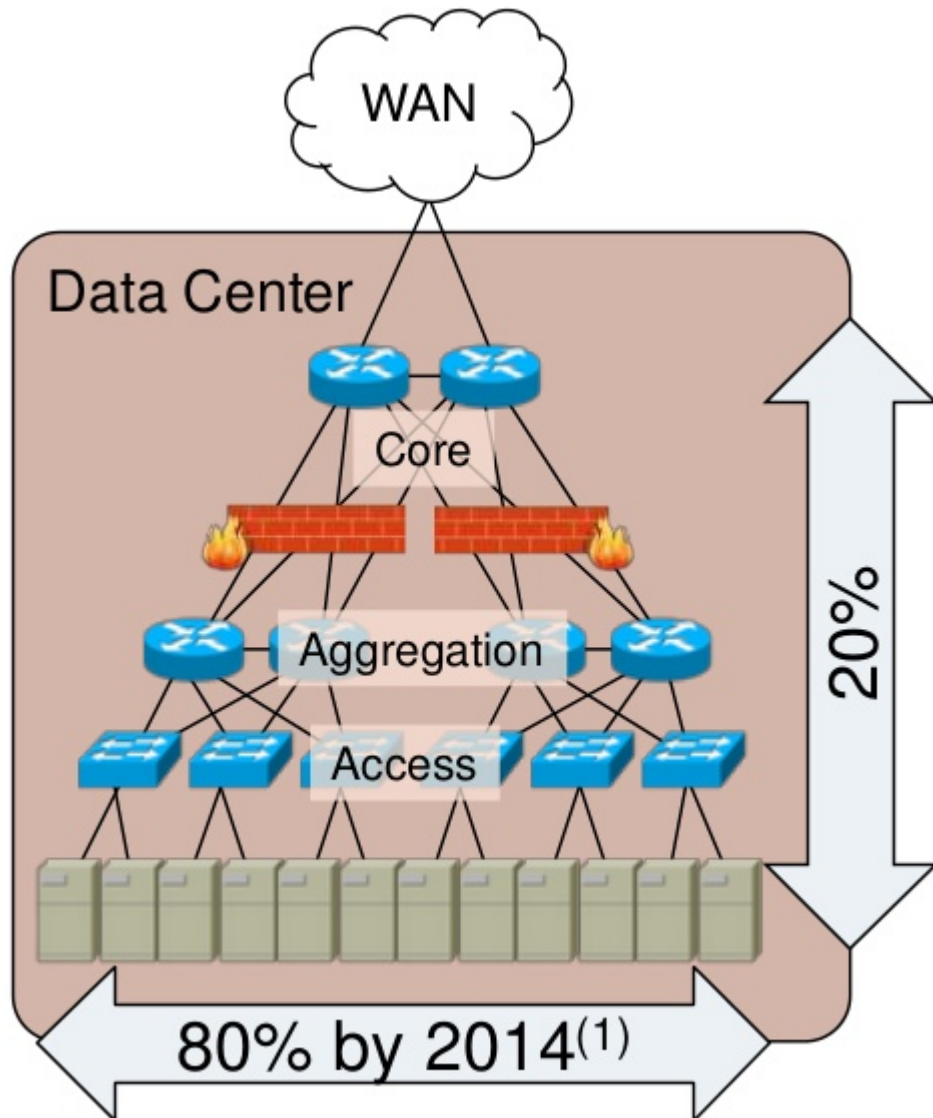
- The networking protocols are distributed among the devices (routers, switches, firewalls and middle boxes)
- The control and data planes are tightly coupled
- No common view of the network
- New networking features are commonly introduced via expensive, specialized and hard-to-configure equipment (aka middle boxes)
- Hard to implement new features and protocols as this means changing the control plane of all devices which are part of the topology
- Each device has to be configured separately which is prone to errors. Many configuration changes are done manually

Evolved Campus Ethernet



- Evolved campus Ethernet into tree structure
 - Core
 - Aggregation
 - Access
- Most traffic is “north-south” (95%)
- Segregated networks at Access to avoid Spanning-Tree problems

The Old Data Center Model



Source: (1) Gartner Synergy Report

- Applied the same model to the Data Center
- Different traffic patterns
 - Majority “east-west”
- Different performance needs
 - Lossless storage traffic
 - Low latency, high bandwidth
- Different service needs to support virtual compute model
 - Static to dynamic
 - Multi-tenancy
 - Workload management
- A new Data Center model exists, called “Spine and Leaf”

Market Drivers

Video and Mobility are transforming business communications

Up to

10X

Increase in network capacity to support new wave of business video applications

INCREASE IN BANDWIDTH REQUIREMENTS

At least

50 Billion

Devices will connect to wireless networks by the year 2020

UNIFIED WIRED AND WIRELESS CAMPUS NETWORKS, IT CONSUMERIZATION

More than

25%

Of all daily business communications will be video or multi-media communications by 2013

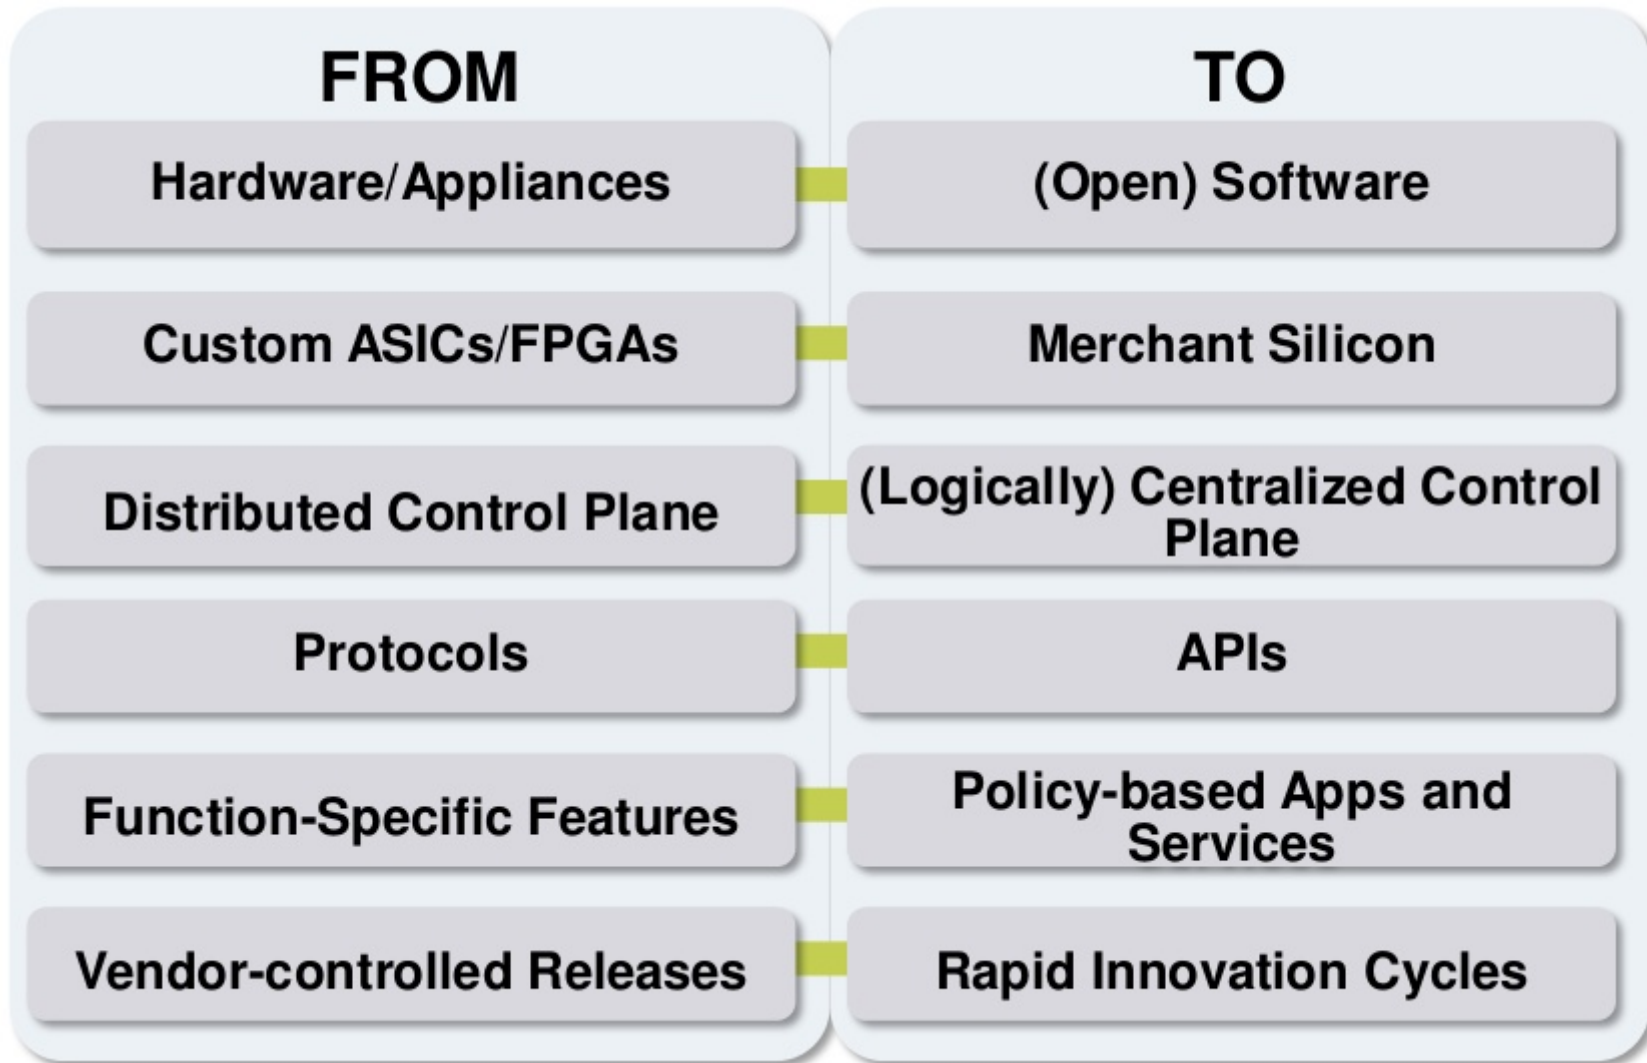
COLLABORATION, TRAINING, PRODUCTIVITY

Source: Gartner – G00207476 Key Technology Analysis
Gartner – G00175764 Key Issues For Communications Strategies, 2010
ONS12

Market Drivers Summarized

- Changing Traffic Patterns
- The “consumerization of IT”
- The rise of cloud services
- “Big Data” means more bandwidth

General Shift in Networking



Limitations of Current Networking Technologies

- Complexity that leads to stasis
- Inconsistent policies
- Scaling the network becomes harder and more complex
- Vendor dependence

Basic SDN Model

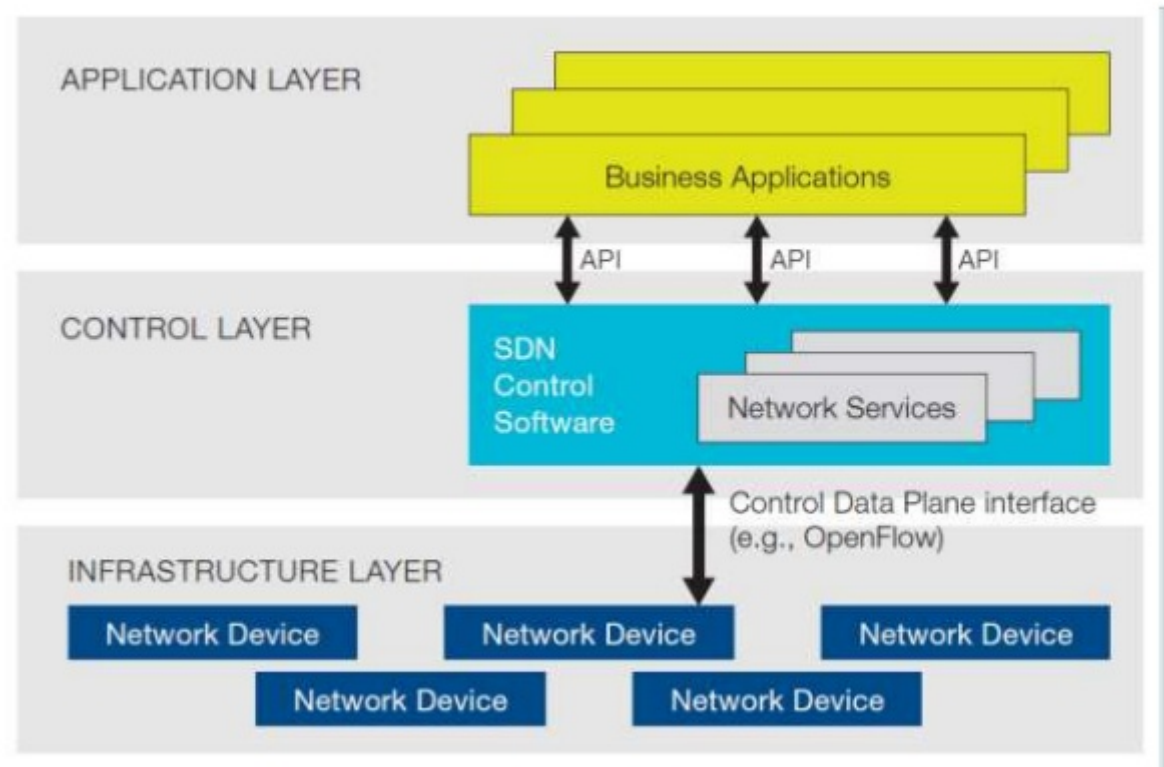
ONF

In the SDN architecture, the control and data planes are decoupled, network intelligence and state are logically centralized and the underlying network infrastructure is abstracted from the applications.

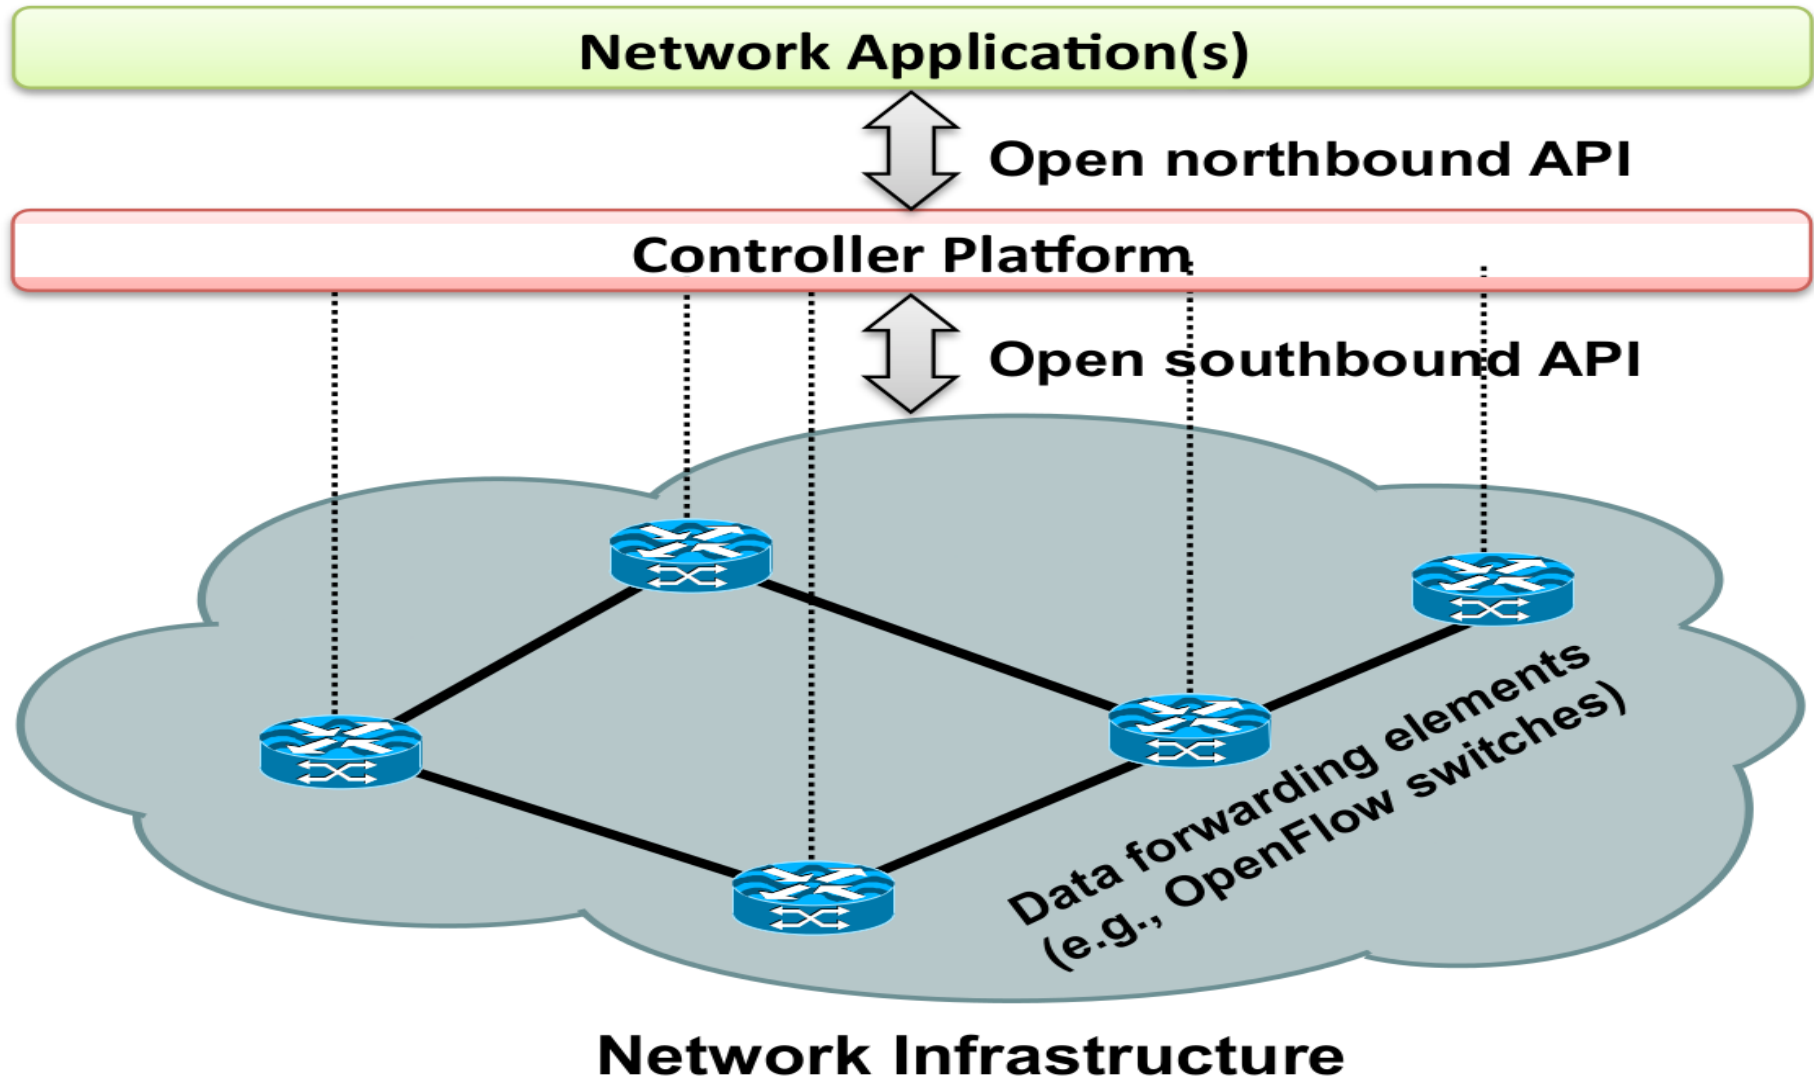
Wikipedia

Software-defined networking (SDN) is an approach to computer networking that allows network administrators to manage network services through abstraction of higher-level functionality. This is done by decoupling the system that makes decisions about where traffic is sent (the control plane) from the underlying systems that forward traffic to the selected destination (the data plane). The inventors and vendors of these systems claim that this simplifies networking

SDN Model




The SDN Model



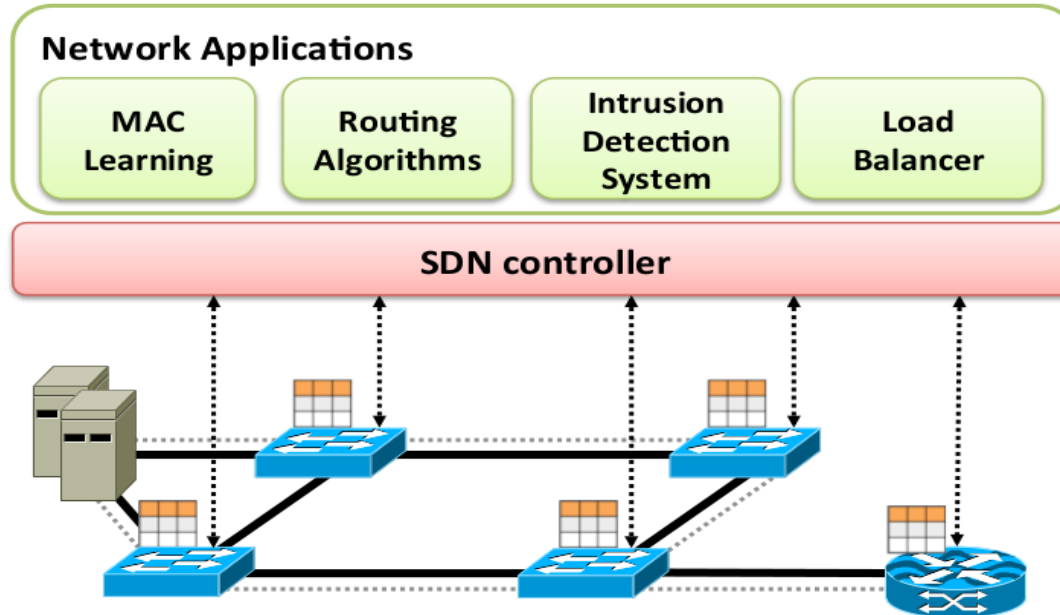
Source: Software-Defined Networking: A Comprehensive Survey, October 8, 2014

SDN vs Conventional Networking

Conventional Networking



Software-Defined Networking



Source: Software-Defined Networking: A Comprehensive Survey, October 8, 2014

SDN Premises and Promises

The premise...

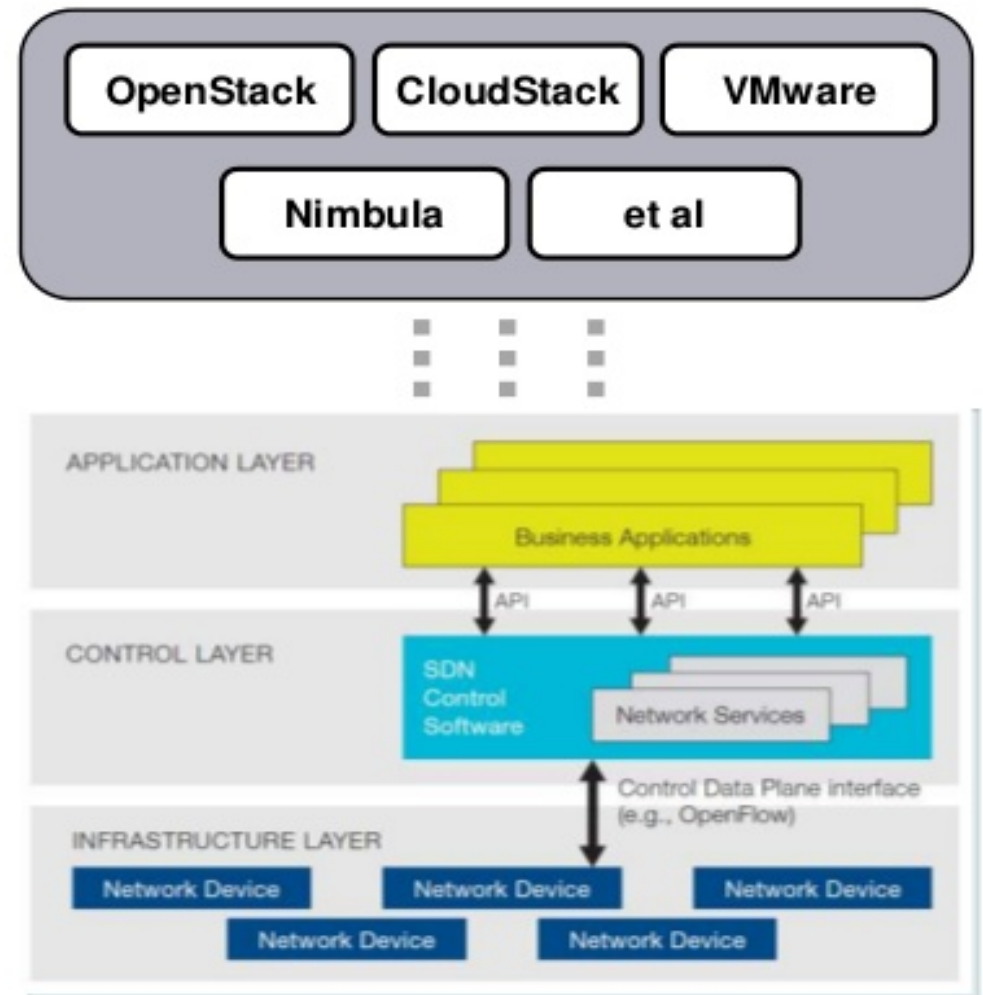
- ▶ Commodity(merchant silicon) solutions can be exploited
- ▶ Control plane can be centralized
- ▶ States can be externalized
- ▶ Acceptable performance can be maintained
- ▶ Standards will evolve
- ▶ Networking manufacturers will adopt SDN-enabled protocols and features

The promise...

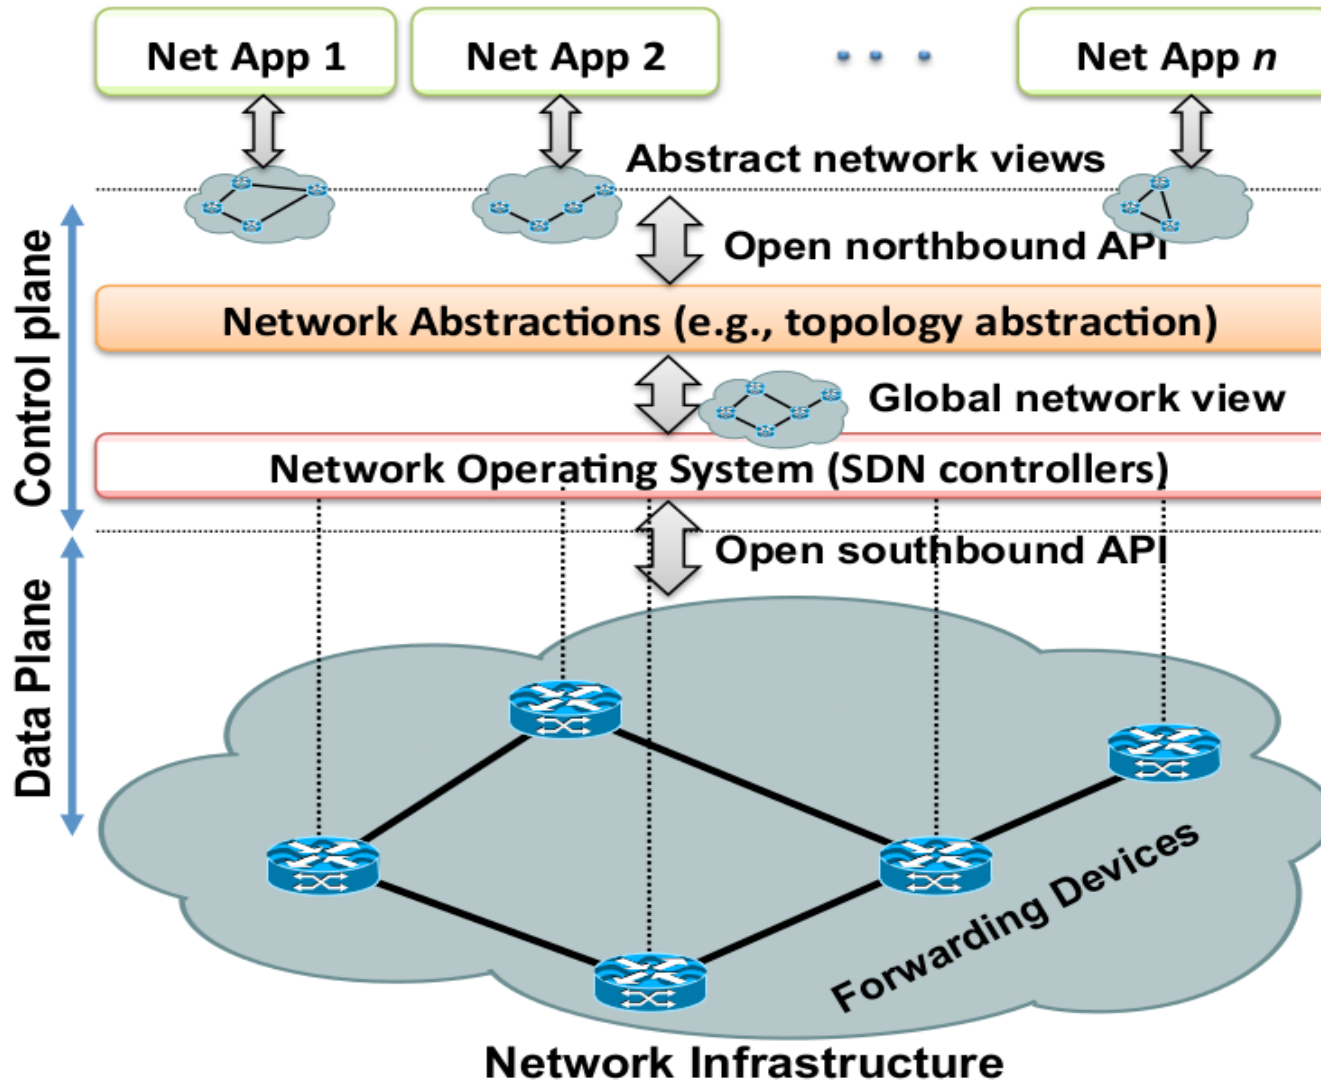
- ▶ Centralized management and control
- ▶ More granular network control
- ▶ Improved automation and management
- ▶ Rapid innovation
- ▶ Programmability
- ▶ Increased network reliability and security
- ▶ Better end-user experience

Essential Elements of SDN

- ▶ Abstraction
- ▶ Polling
- ▶ Orchestration
- ▶ Automation
- ▶ Service insertion
- ▶ Apps
- ▶ Programmability(APIs)



Important SDN Abstractions



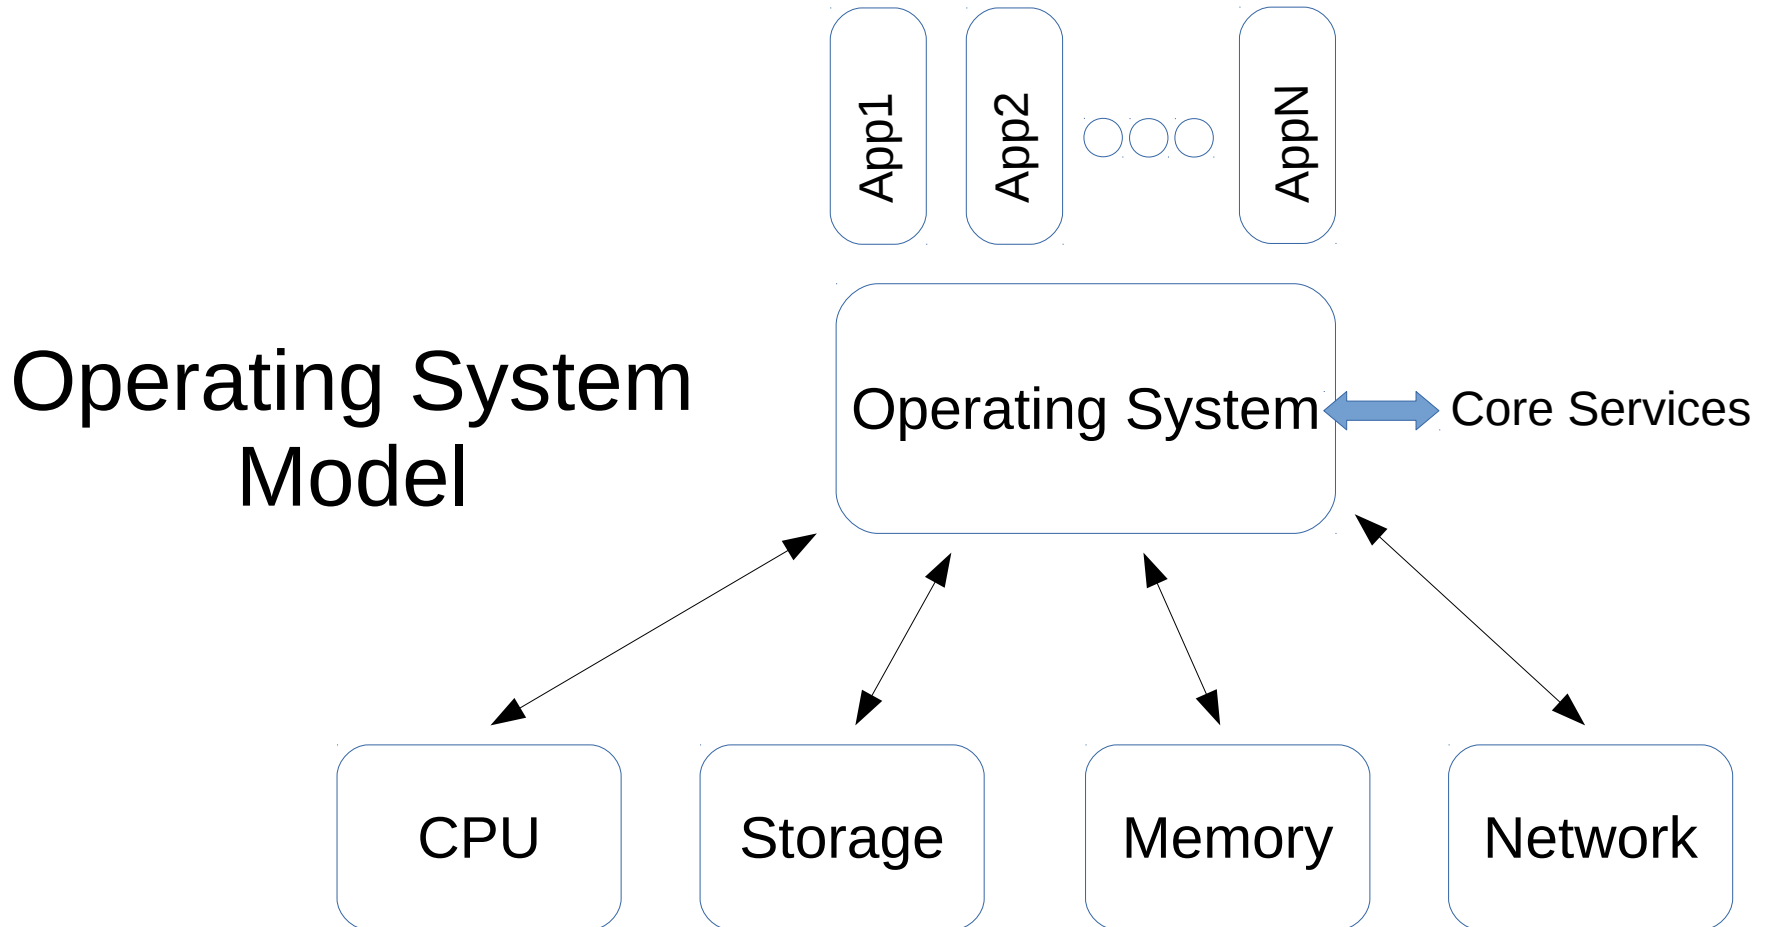
And before we move on...

Let us consider another well known abstraction model as an analogy to better understand SDN abstractions

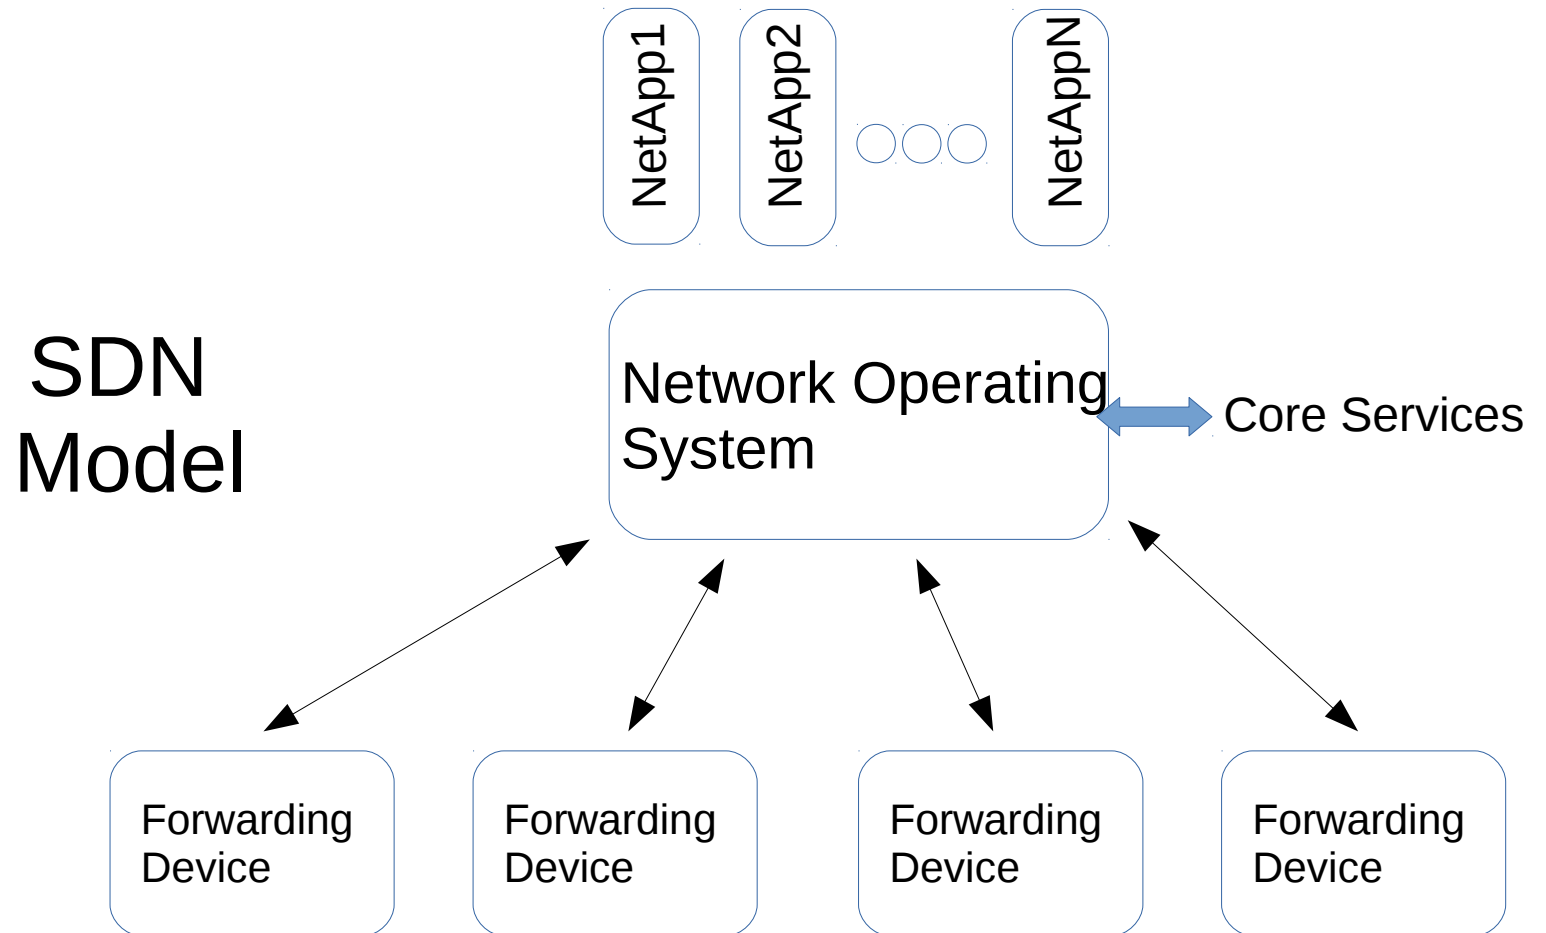
Operating System
Model



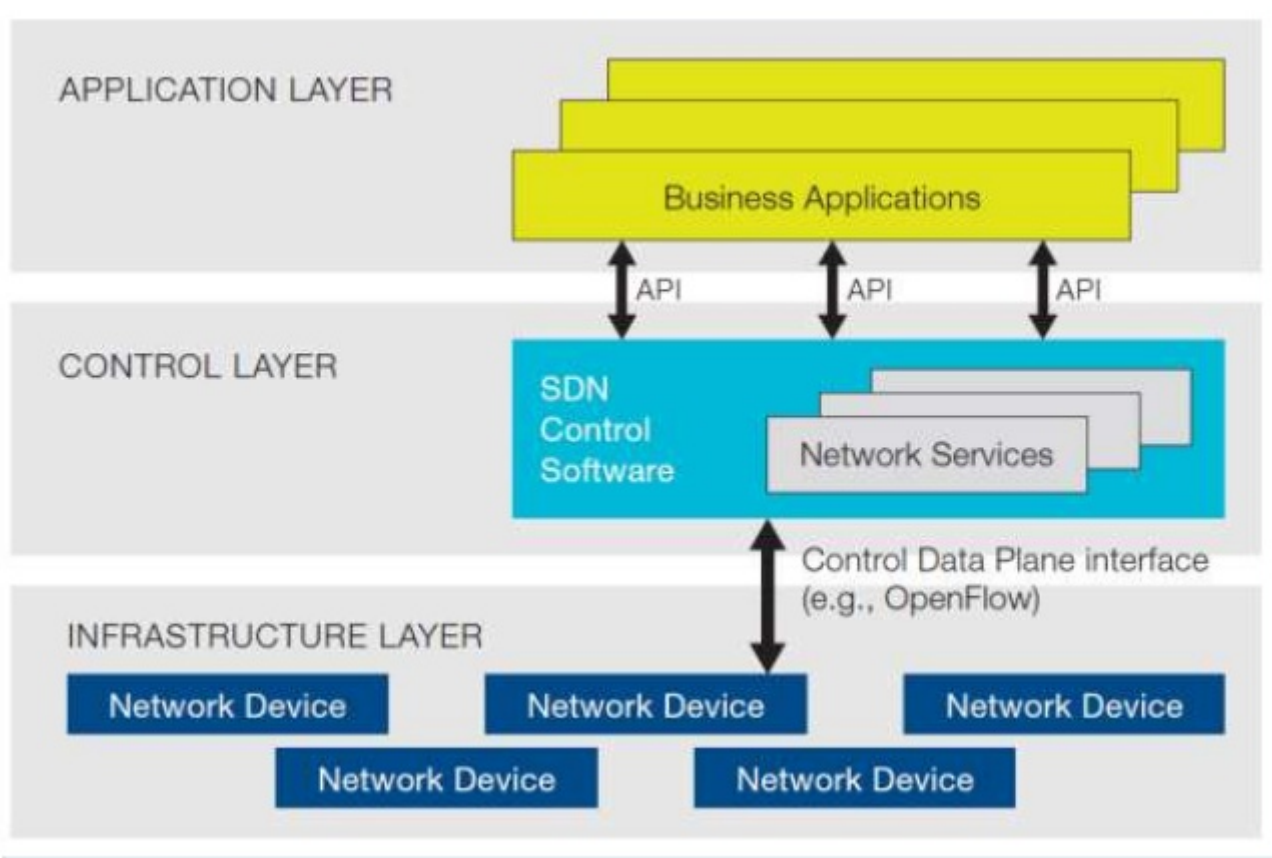
Operating System Model



SDN Model








Inside the Layers

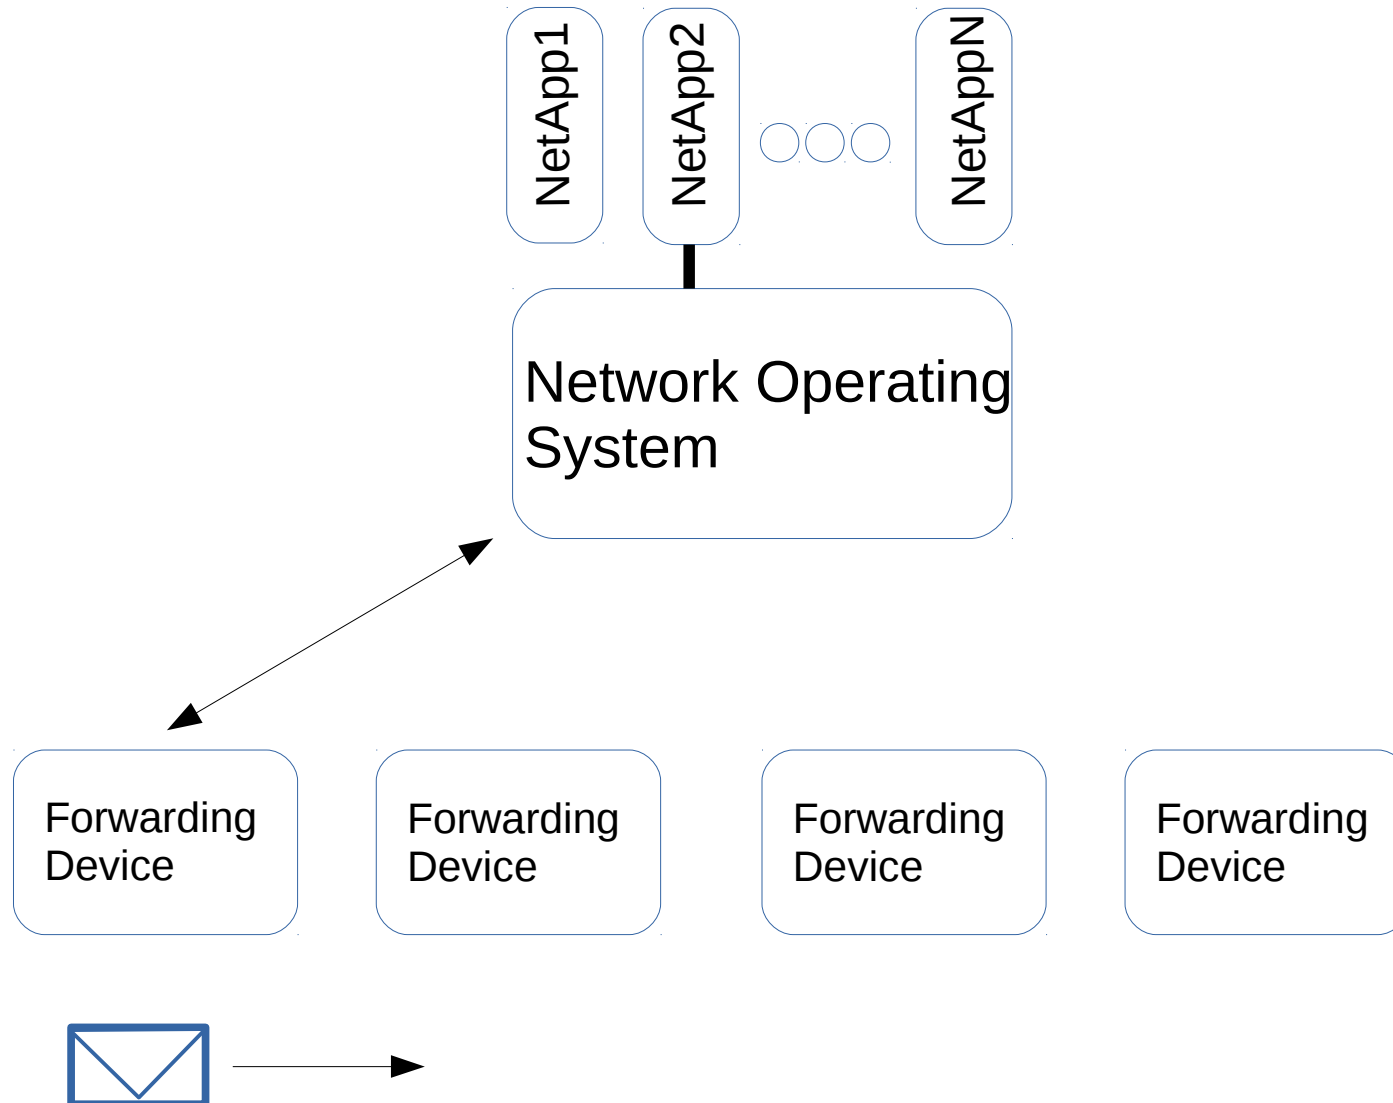


- Virtual network overlays
 - Slicing
 - Tenant-aware broadcast
 - Application-aware packet computation
 - Traffic engineering
 - Network services (FW, LB, Security)
-
- Data plane resource management
 - Common services and libraries
 - Topology
 - State abstraction
-
- Packet forwarding
 - Packet manipulation
 - Statistics gathering

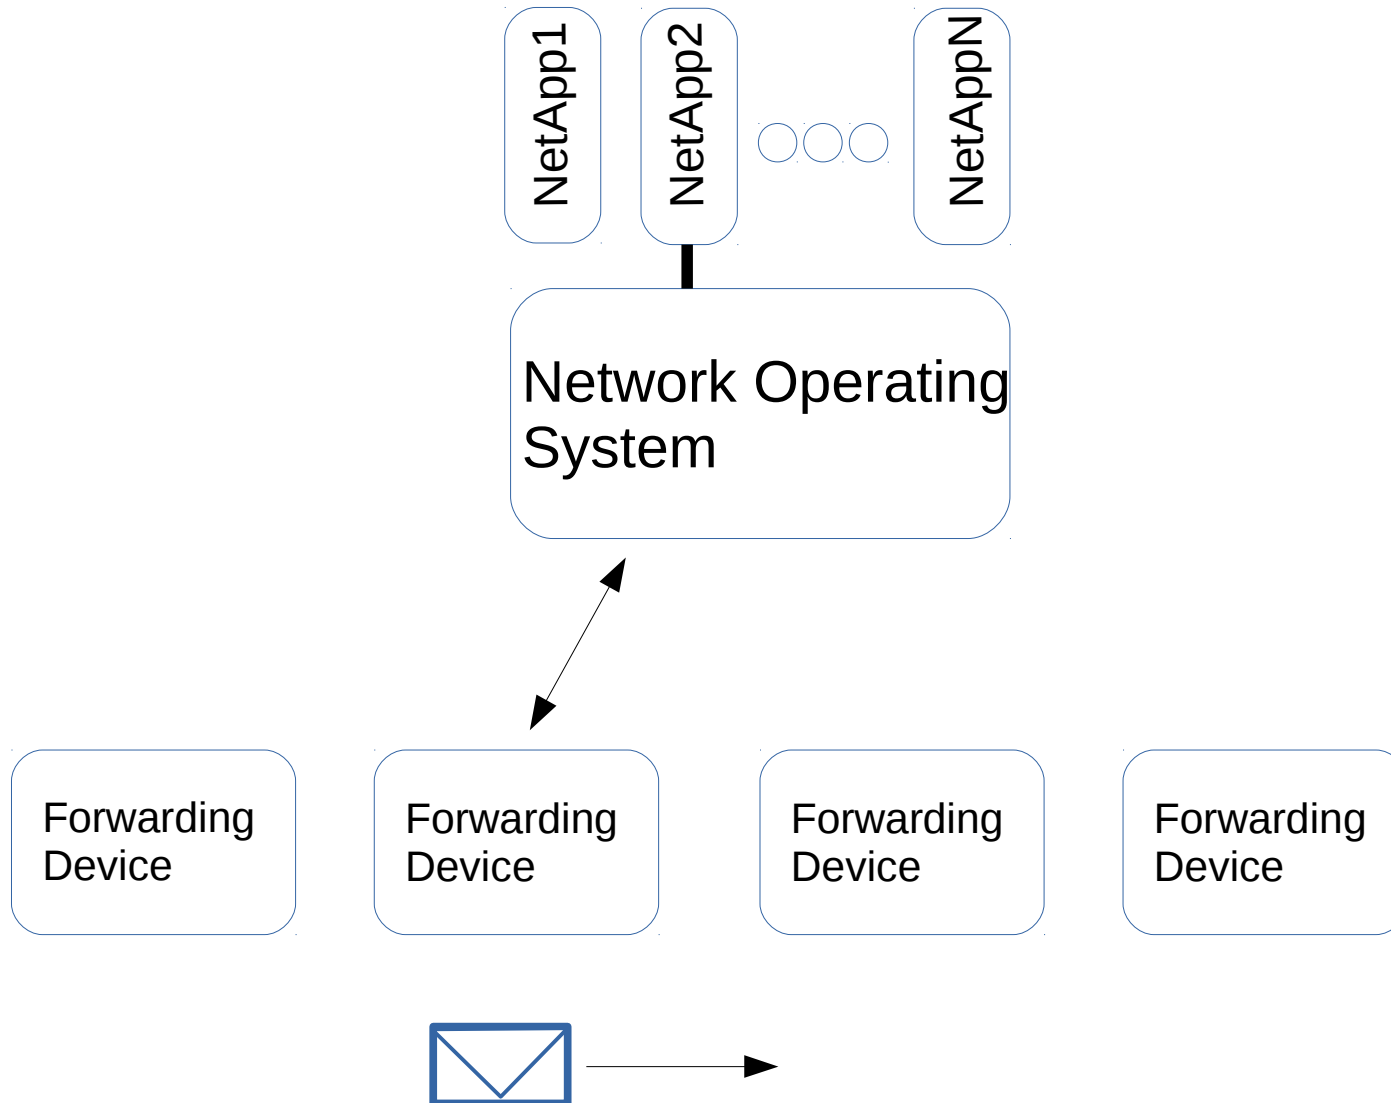
Packet Flow



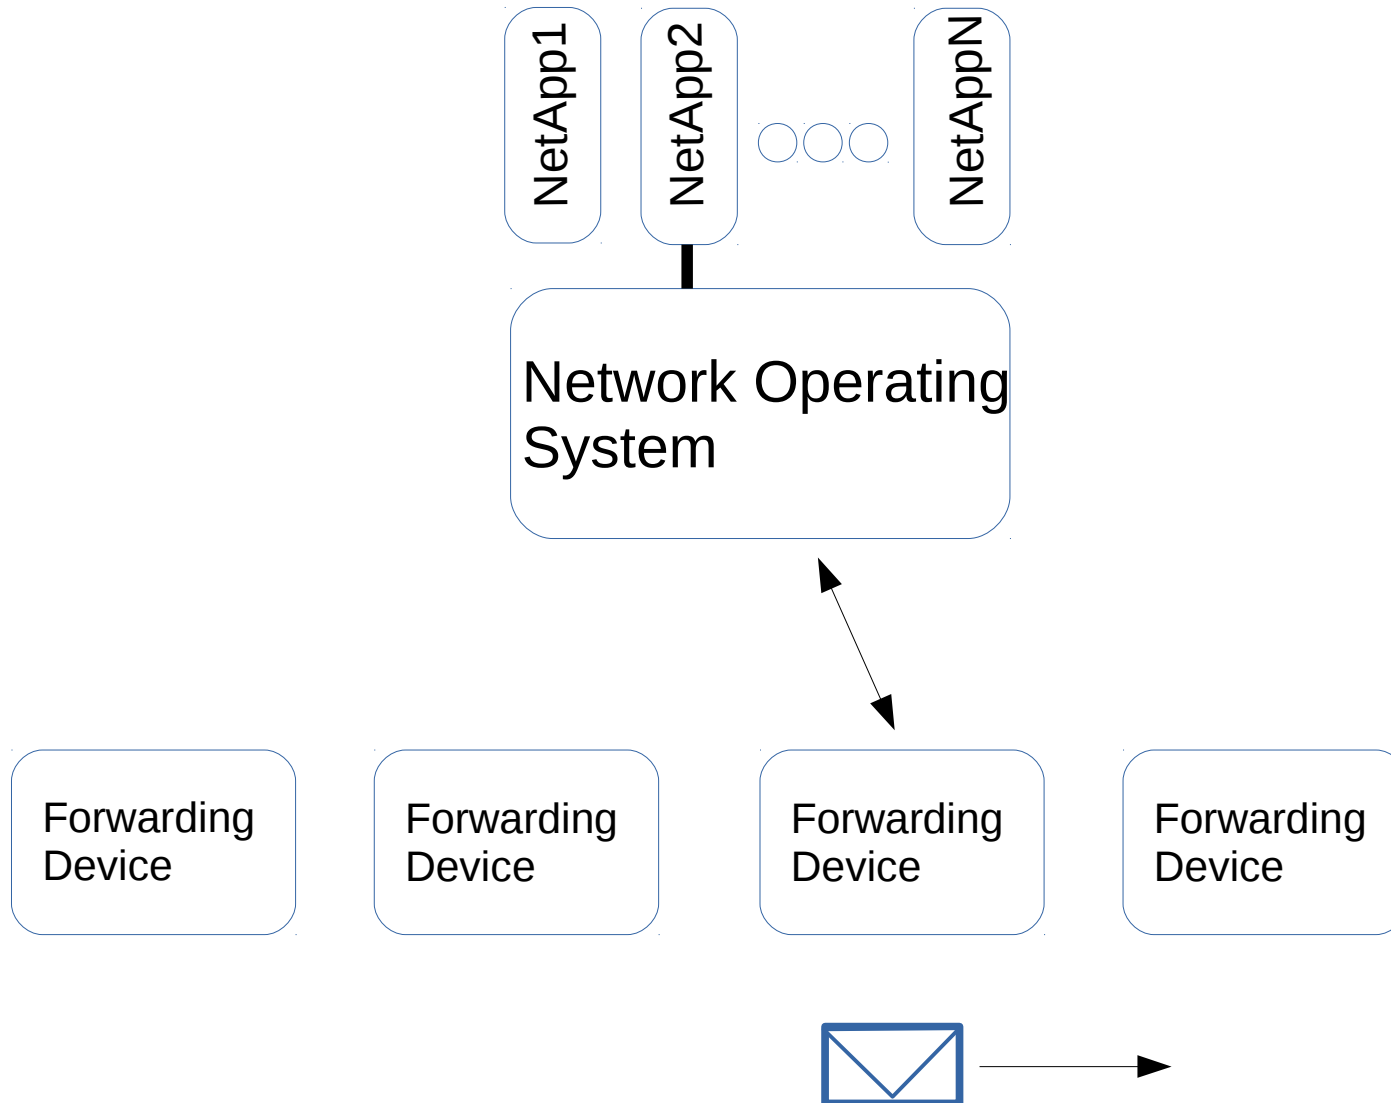
Packet Flow



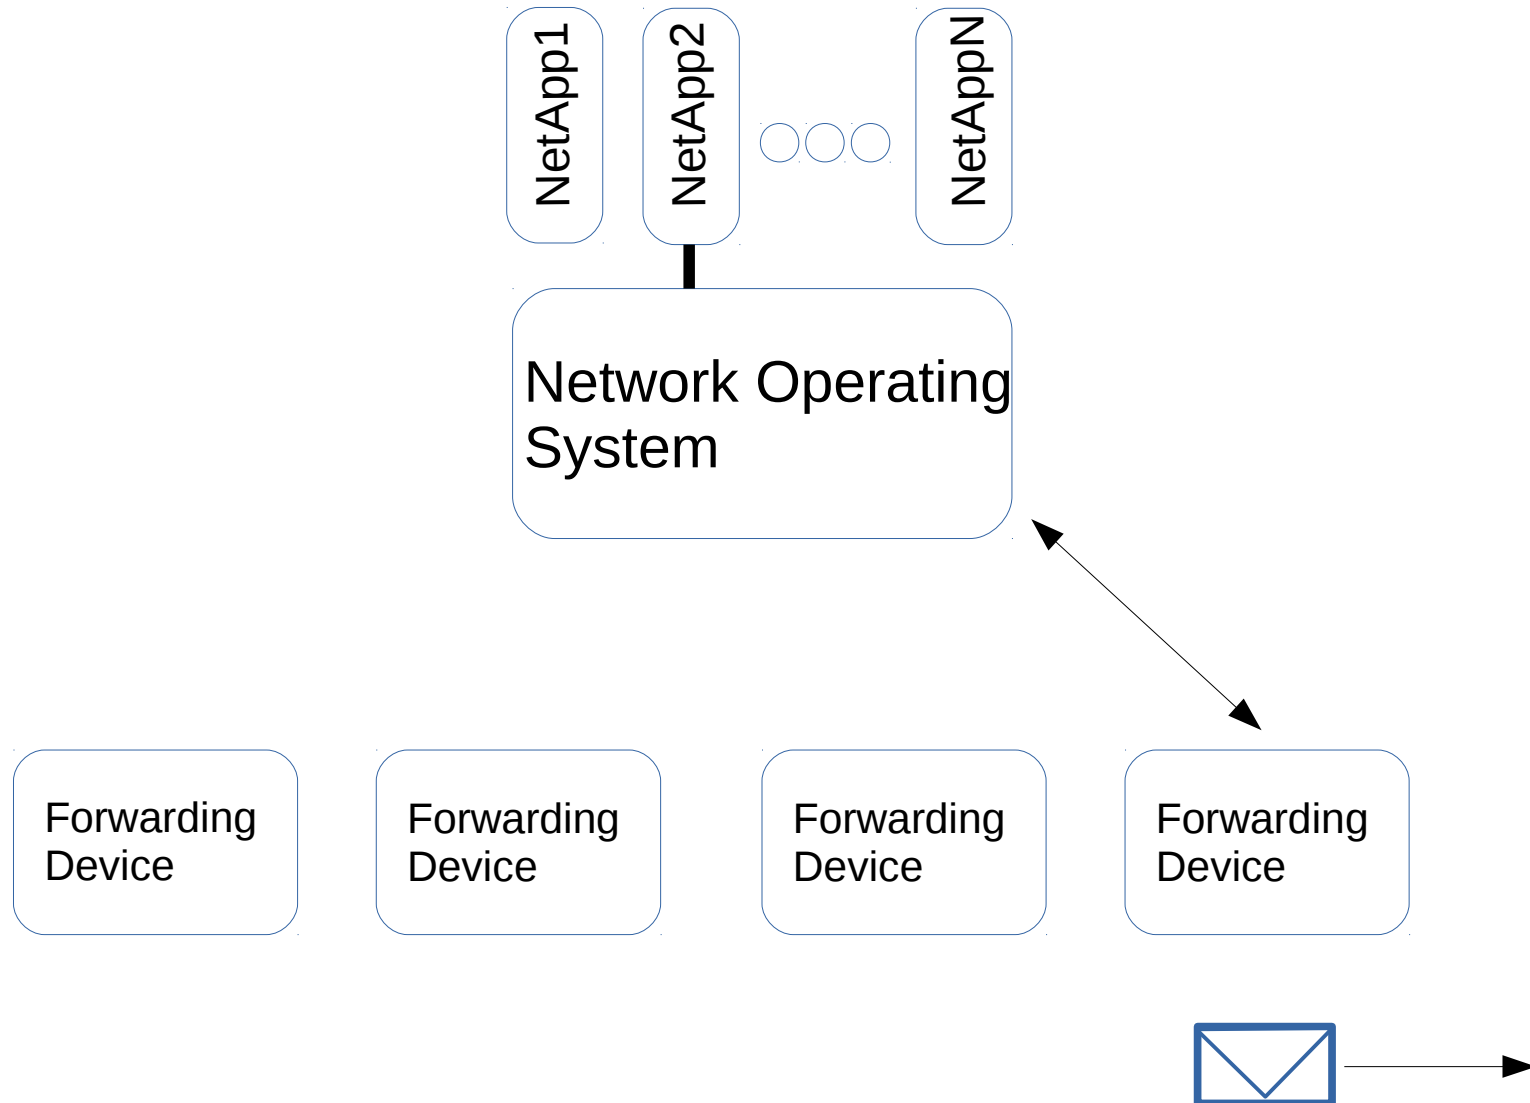
Packet Flow







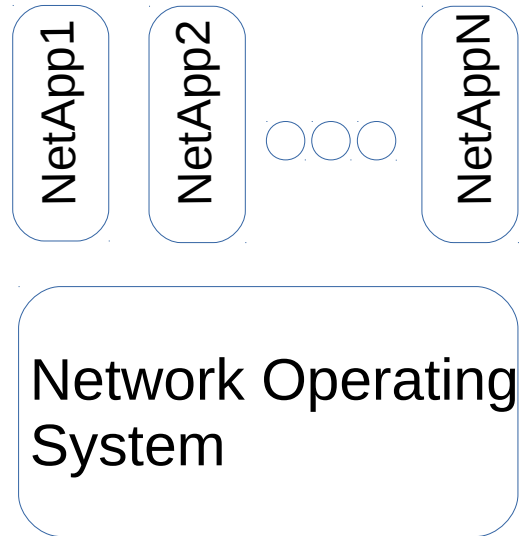
Packet Flow



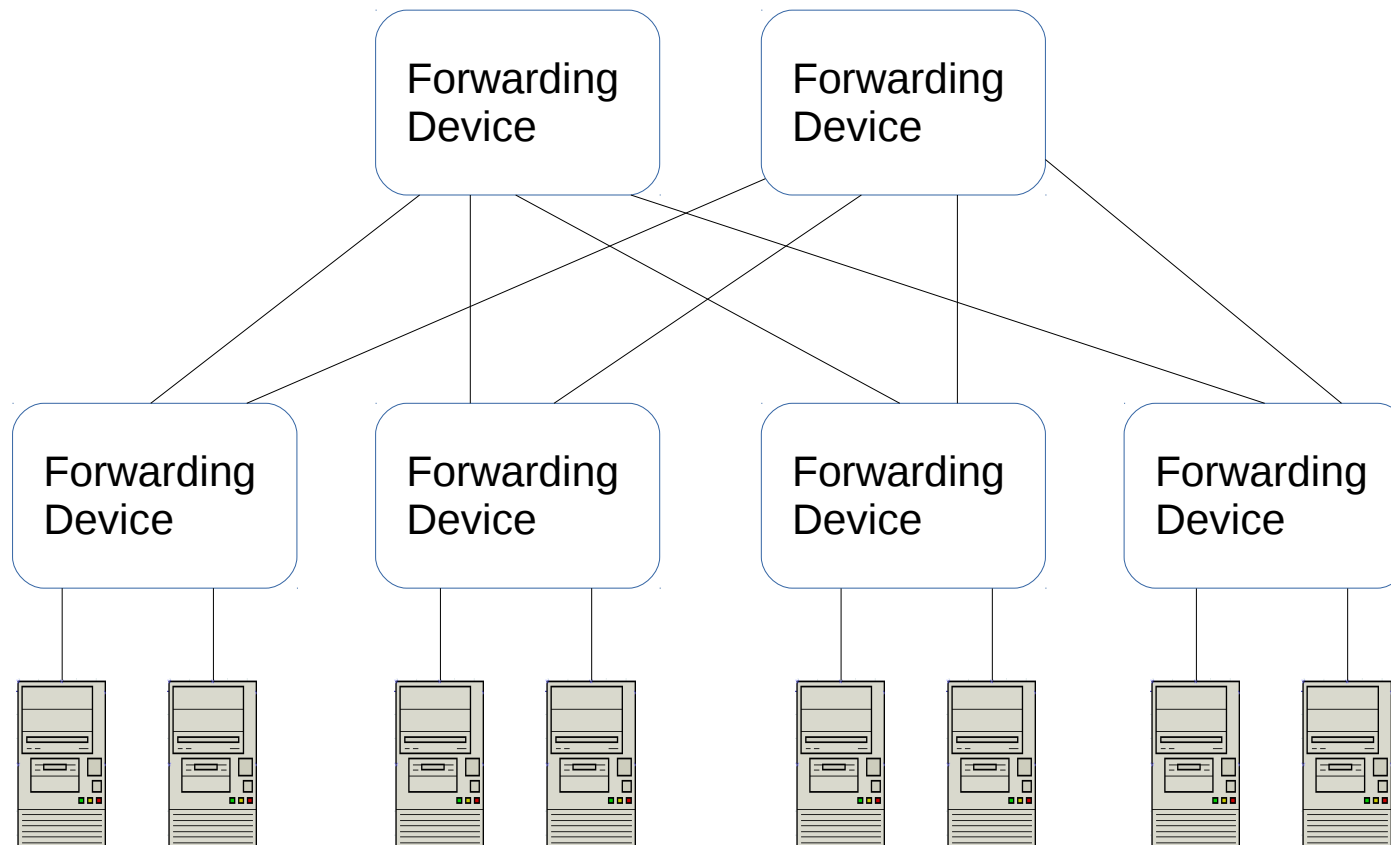
Packet Flow



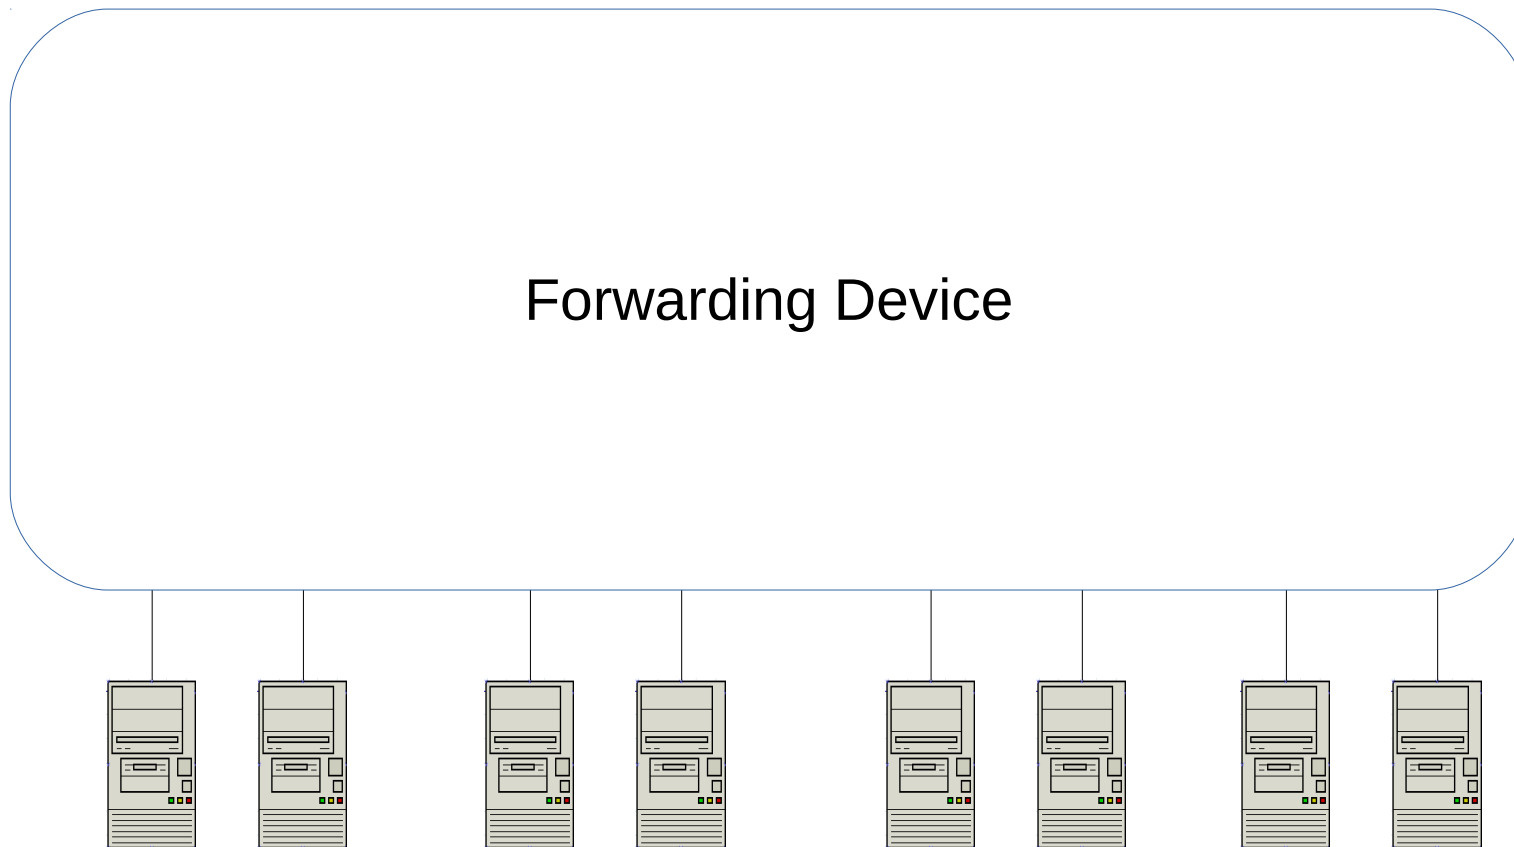
Packet Flow



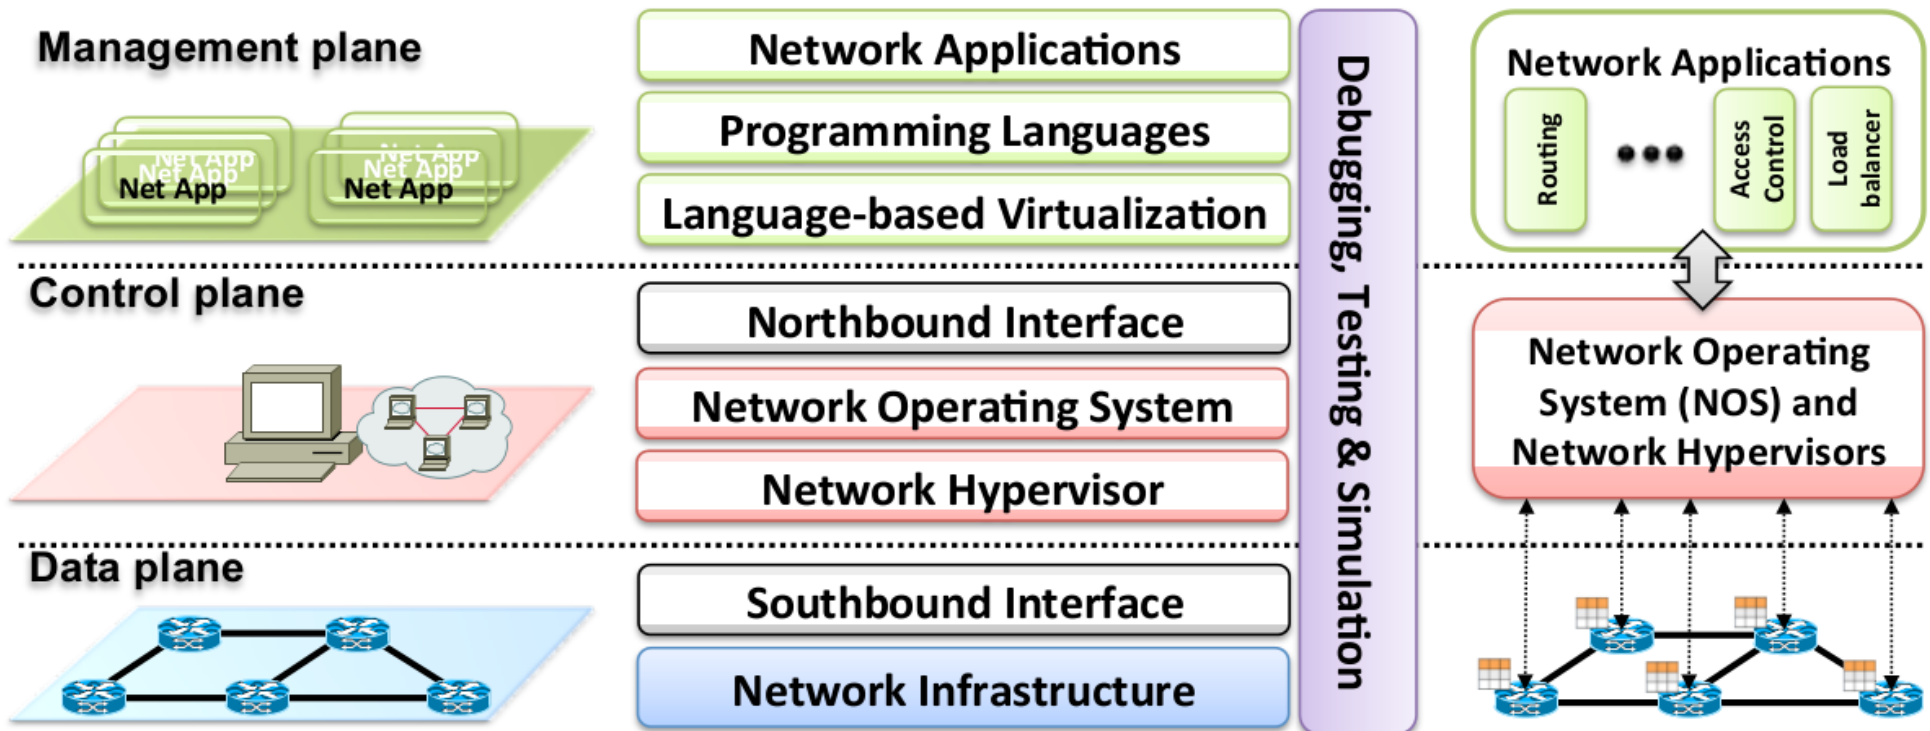
Abstraction for Applications



Abstraction for Applications



SDN Planes Summarized



Typical Use Cases

- ✓ Security
- ✓ Switching
- ✓ Routing
- ✓ Traffic Engineering
- ✓ QoS
- ✓ Network Access Control
- ✓ Load Balancing
- ✓ Monitoring
- ✓ Network Taps
- ✓ Cut-Trough Applications
- ✓ Network Virtualization (Overlays)
- ✓ Multi-Tenancy
- ✓ Campus Slicing
- ✓ New innovations???

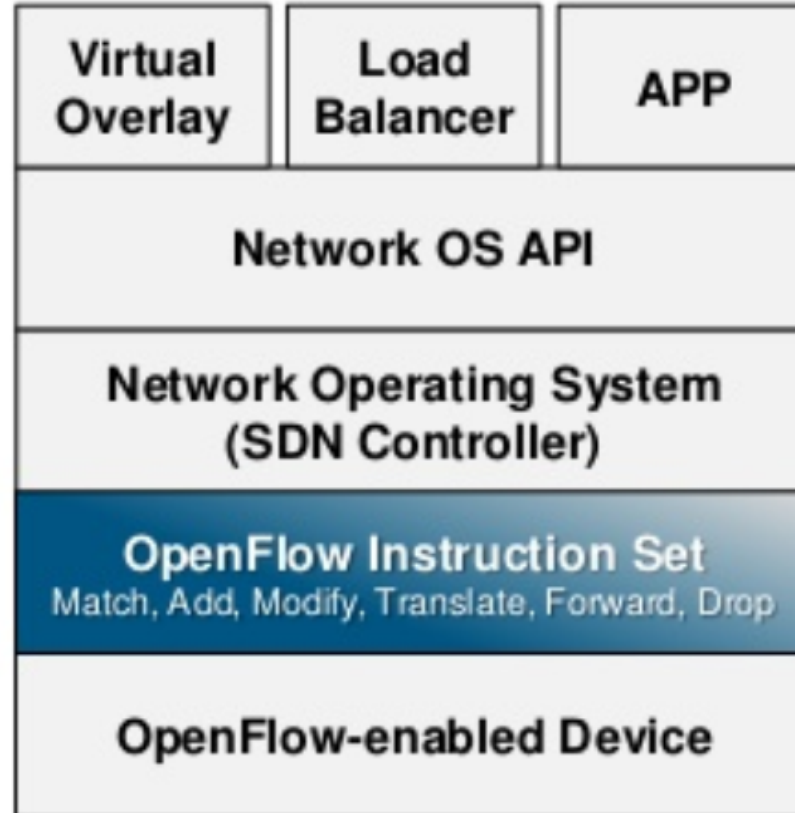
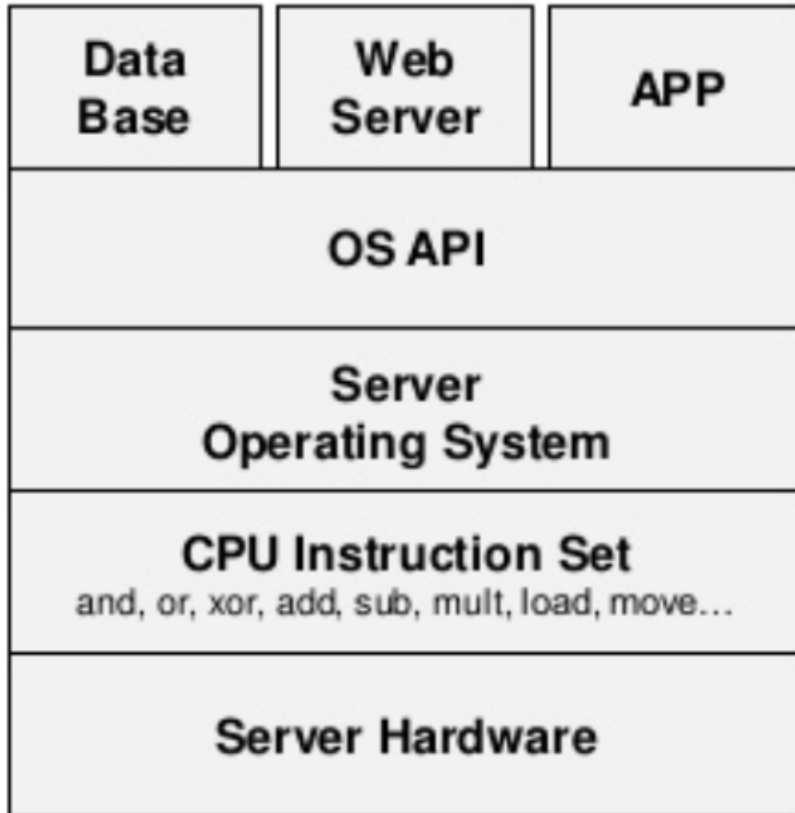
OpenFlow

OpenFlow \neq SDN
(by itself)

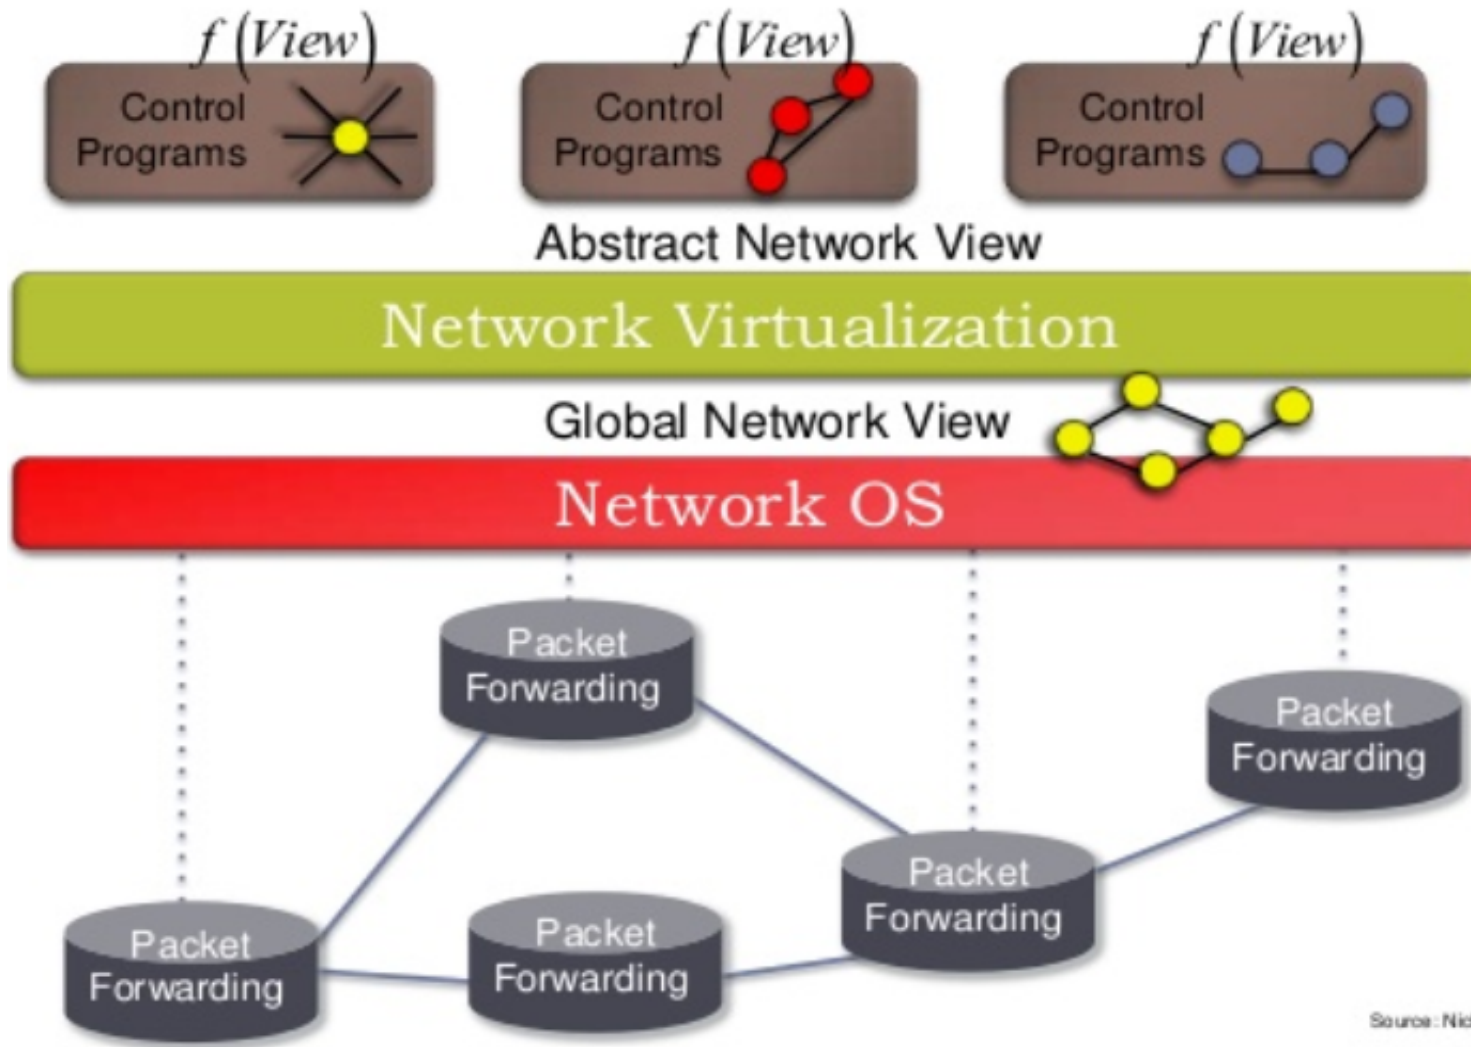
OpenFlow

- A protocol specification
 - Open Networking Foundation
- Requires OpenFlow-enabled devices
 - Switches*
- Defines controller messages
 - PACKET_IN, PACKET_OUT, REMOVE_FLOW etc.
- Enables construction of Flow Tables
 - Match/Action

OpenFlow Analogy

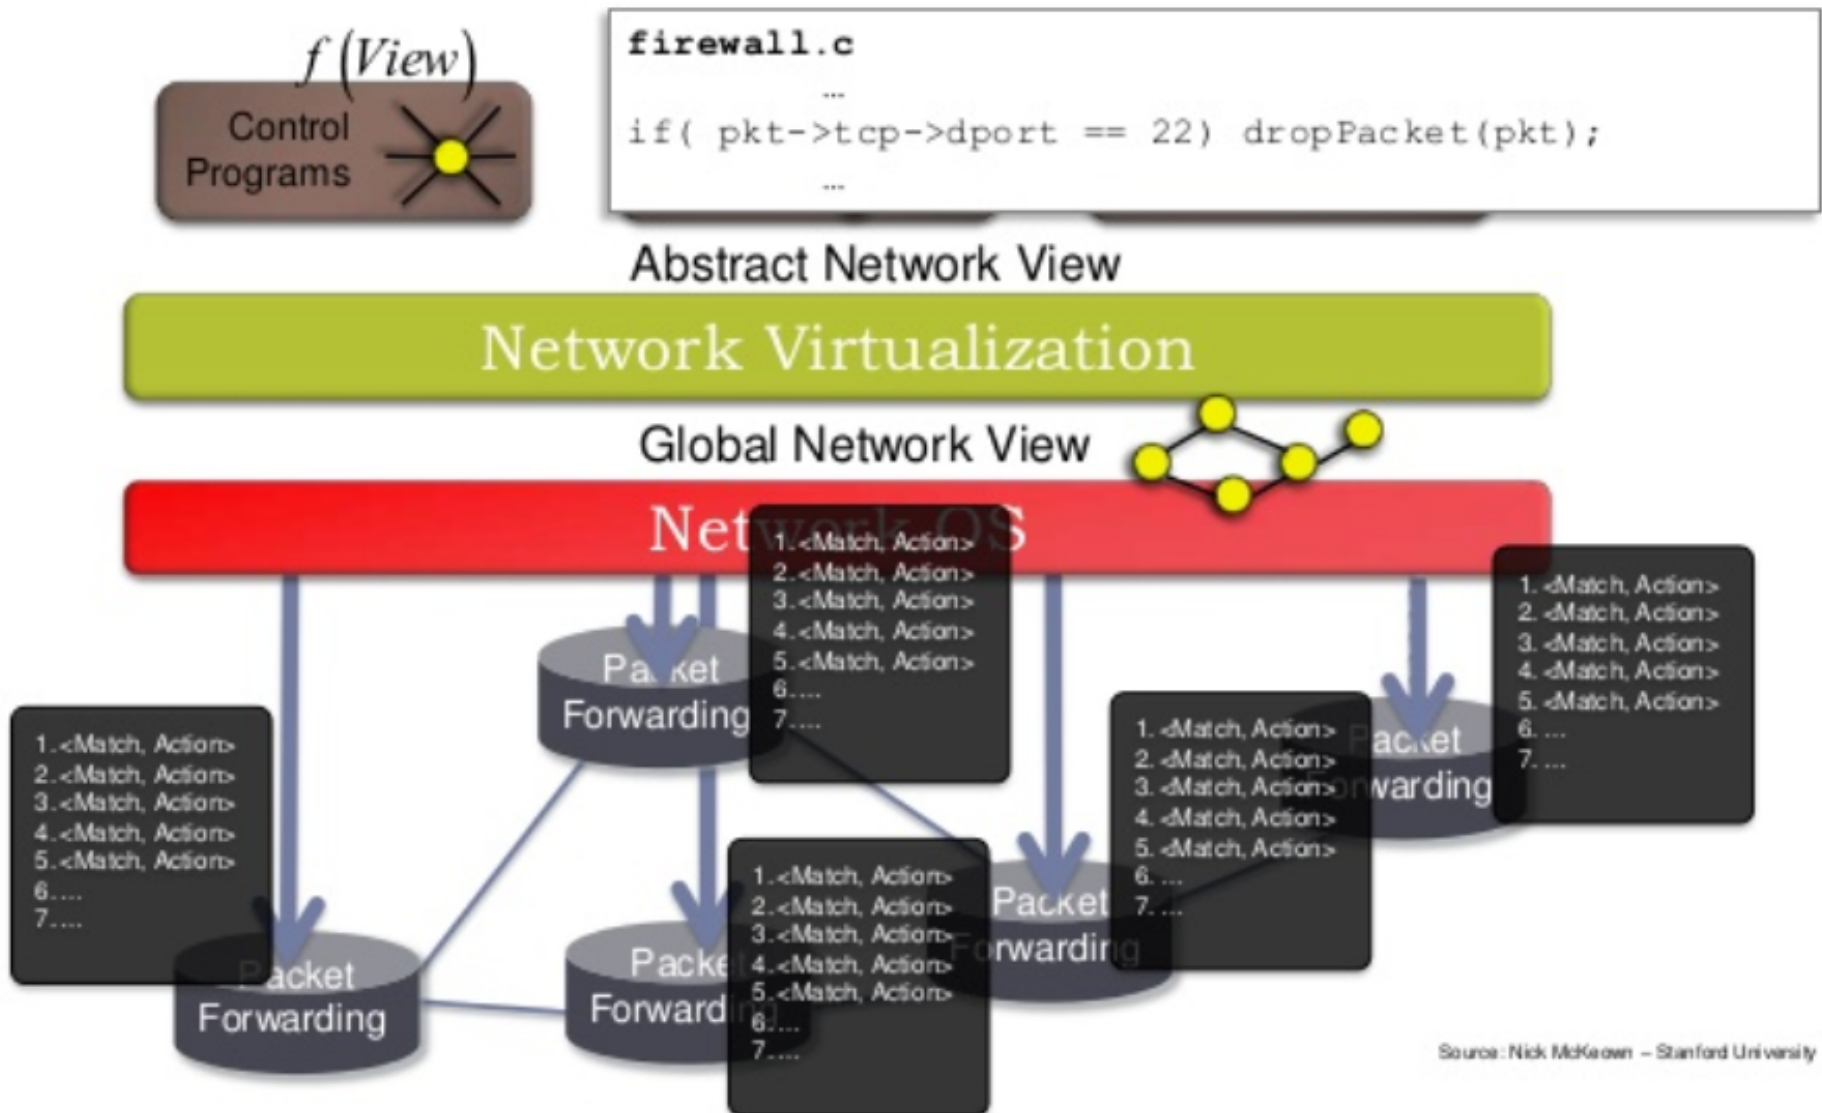


Simple OpenFlow-enabled Example



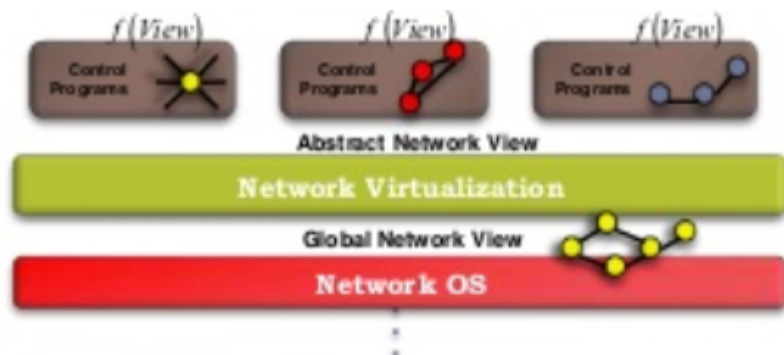
Source: Nick McKeown - Stanford University

Simple OpenFlow-enabled Example



Source: Nick McKeown - Stanford University

Flow Table Example



OpenFlow-enabled switch

Flow Table

MAC src	MAC dst	IP Src	IP Dst	TCP dport	...	Action	Count
*	10:20:..	*	*	*	*	port 1	250
*	*	*	5.6.7.8	*	*	port 2	300
*	*	*	*	25	*	drop	892
*	*	*	192.*	*	*	local	120
*	*	*	*	*	*	controller	11

Flow Table

Generic primitive that sits on top of (virtual) switch TCAM, designed to match well with common ASICs.

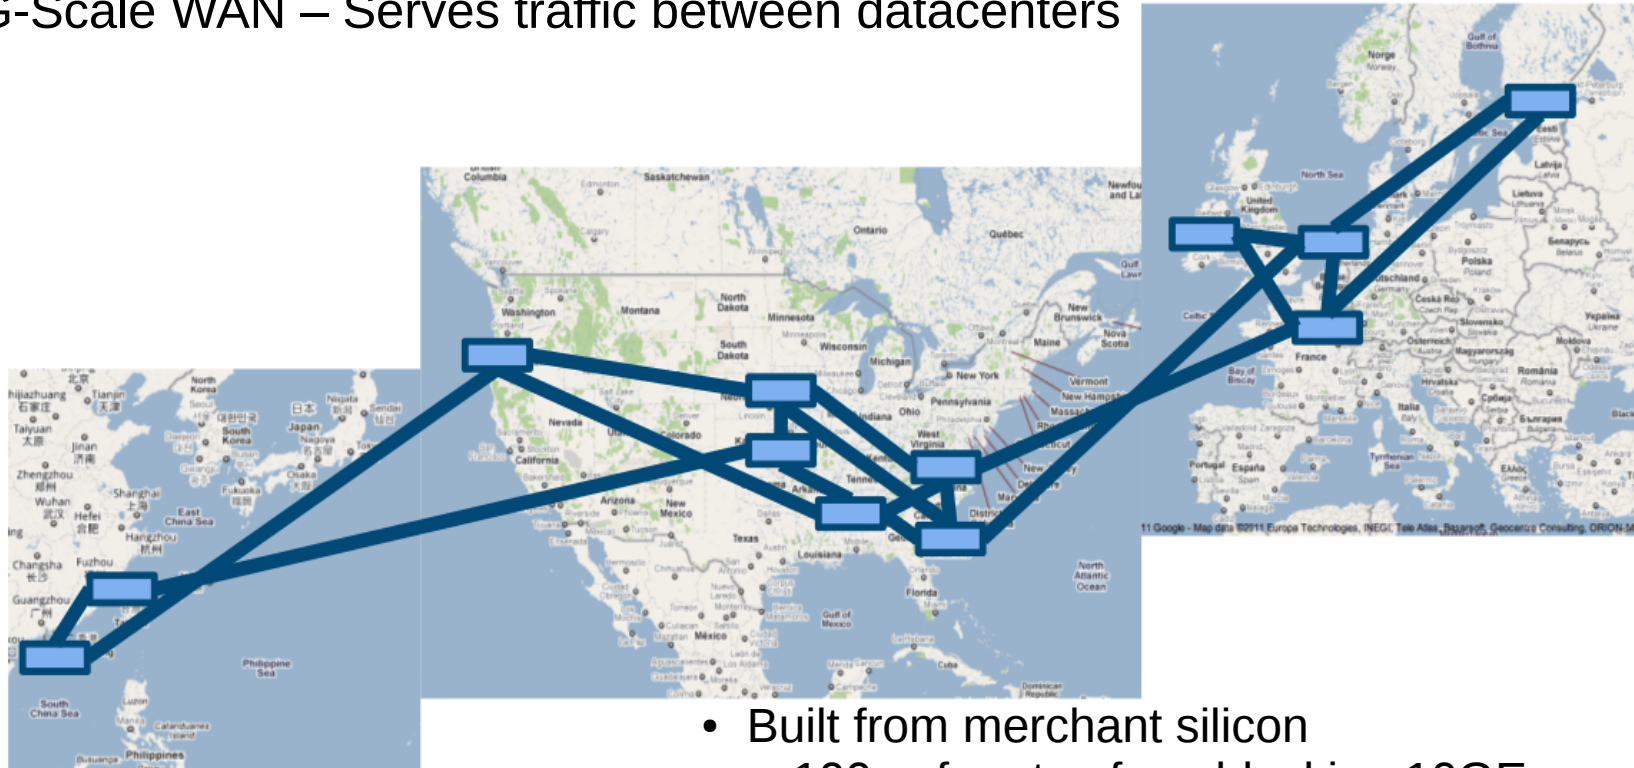
Example actions:

1. Switching and routing (port)
2. Firewall (drop)
3. Use with switch's non OpenFlow logic (local)
4. Send to controller for processing (controller)

Foundation network functions are split between switch and high-level decisions at the controller

Real World G-Scale Example

G-Scale WAN – Serves traffic between datacenters




- Built from merchant silicon
 - 100s of ports of nonblocking 10GE
- OpenFlow support
- Open source routing stacks
- Multiple chassis per site
 - Fault tolerance
 - Scale to multiple Tbps



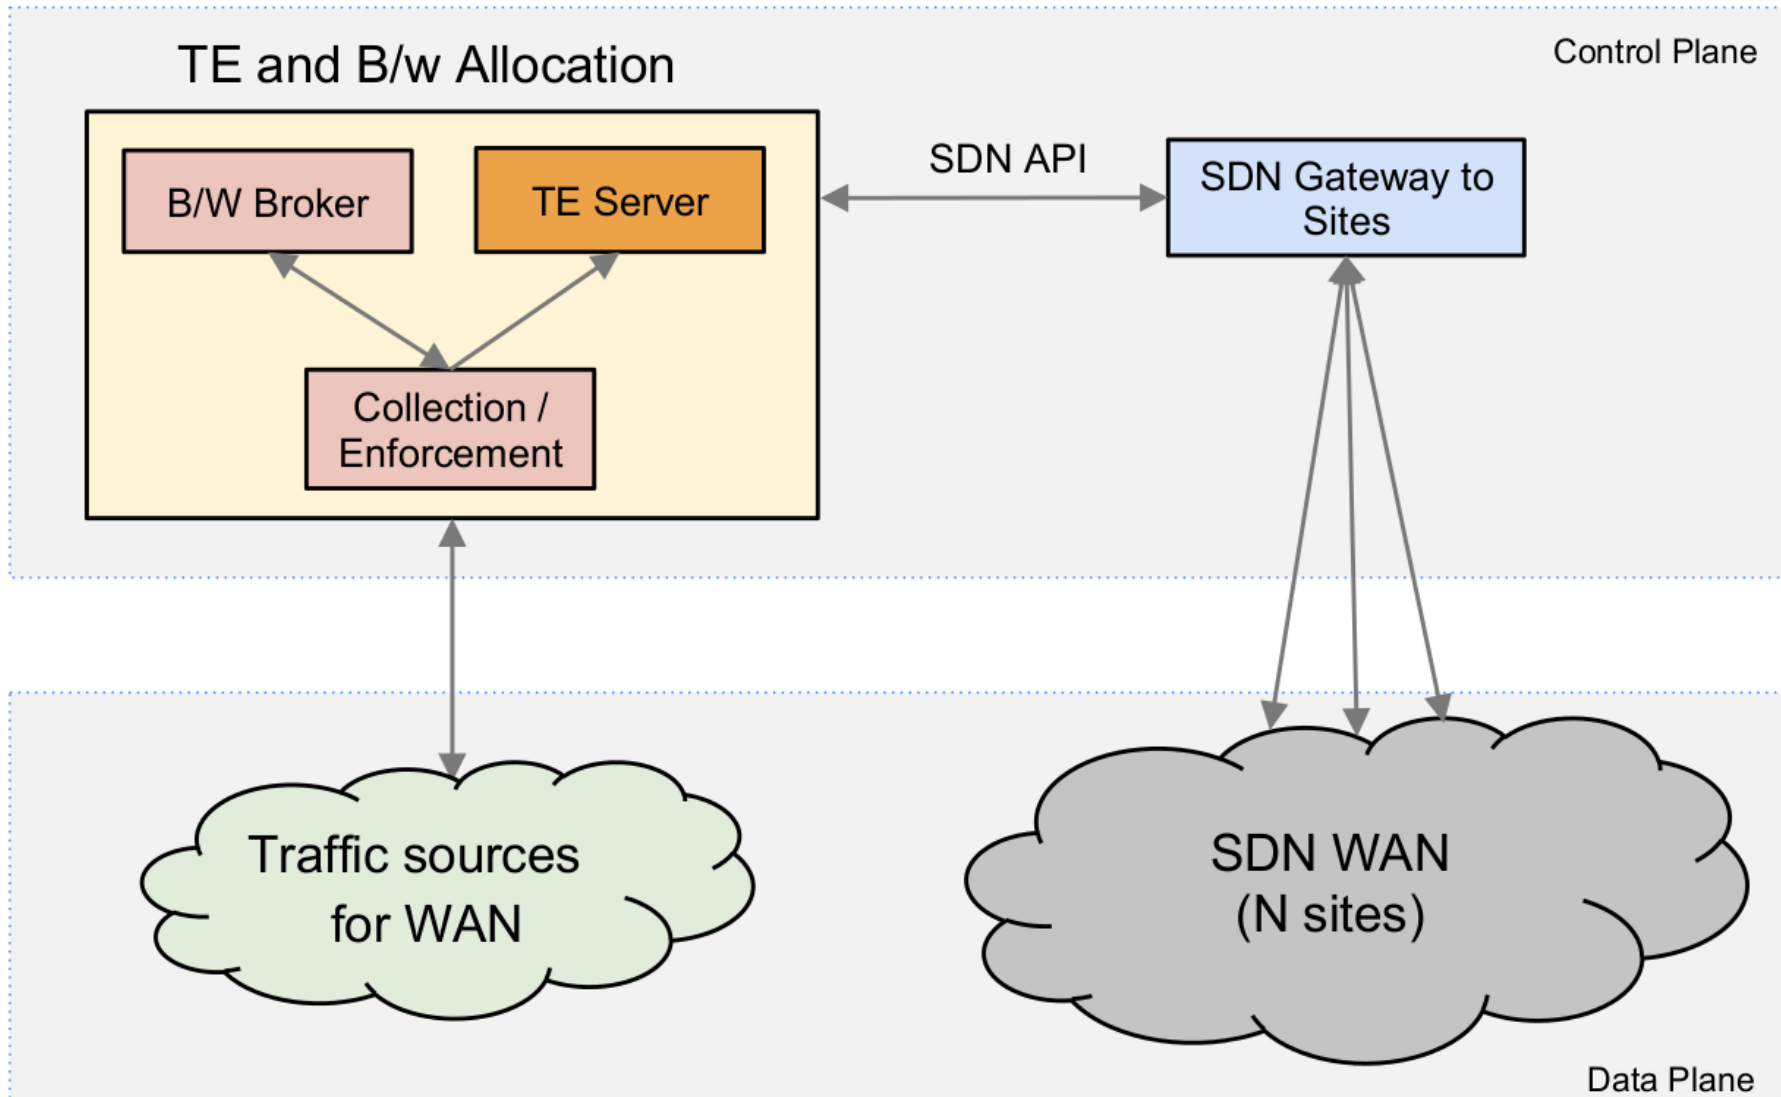
Real World G-Scale WAN Deployment



- Multiple switch chassis in each domain
 - Custom hardware running Linux
- Open source routing stacks
- OpenFlow support



Centralized TE in G-Scale WAN



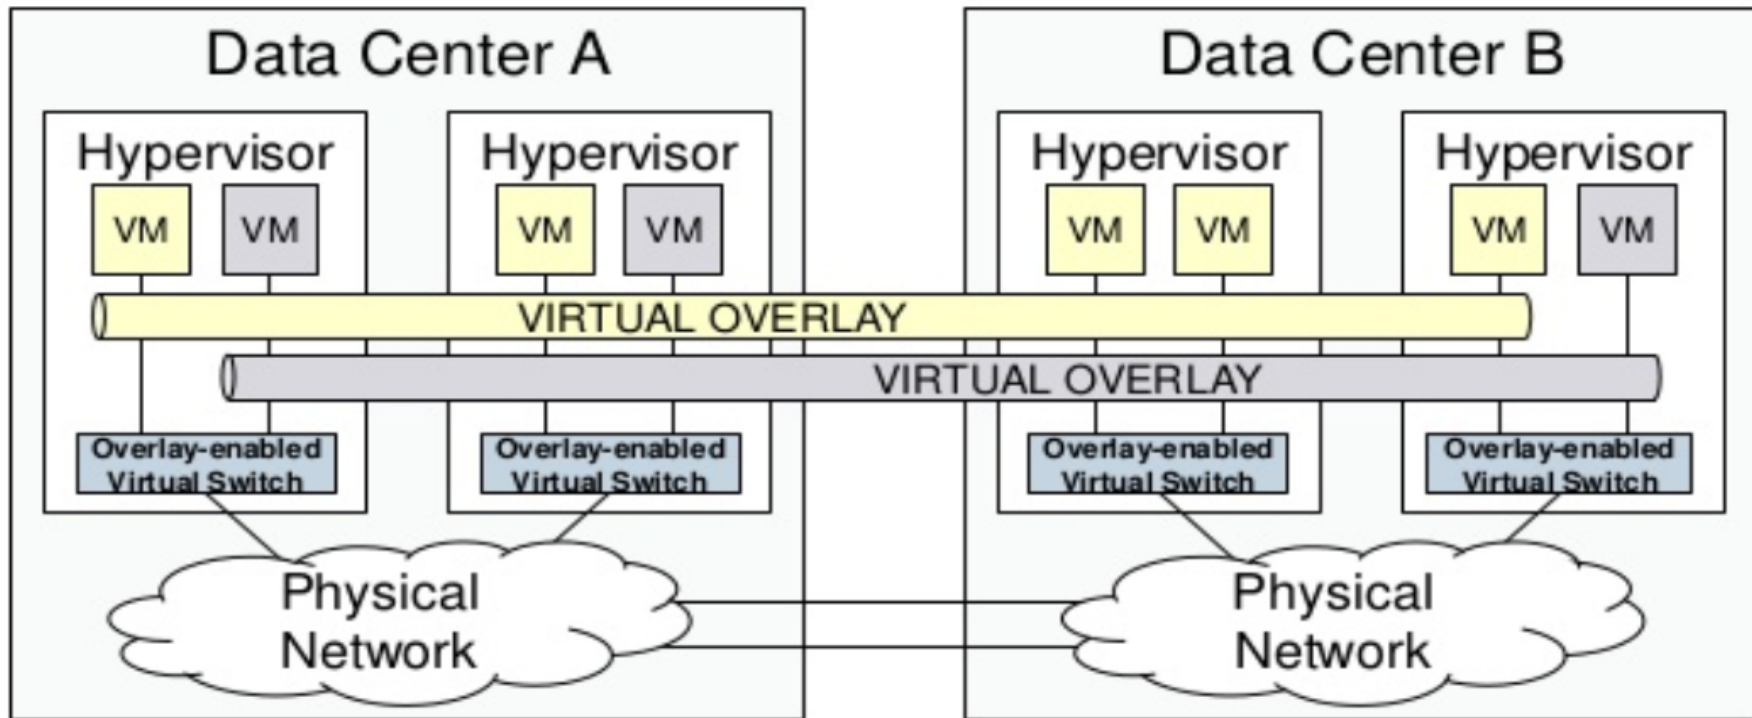
Benefits of SDN for G-Scale WAN

- Unified view of the network fabric
 - Simplifies configuration, management and provisioning
- High utilization – up to 95% utilization of the network
- Faster failure handling
 - Systems converge faster to target optimum and behavior is predictable
- Faster time to market/deployment
 - Only features needed are developed and rigorous testing helps accelerate deployment
- Hitless upgrade

Data Center/Cloud Networking Issues

- VLAN limits (4094)
- Spanning Tree Protocol disabled links
- Reconfiguration to extend VLANs
- MAC address contention
- MAC address table size in ToR switches
- Layer 3 address contention
- Security “choke points”
- ...

Virtual Overlays Using IP-encapsulation



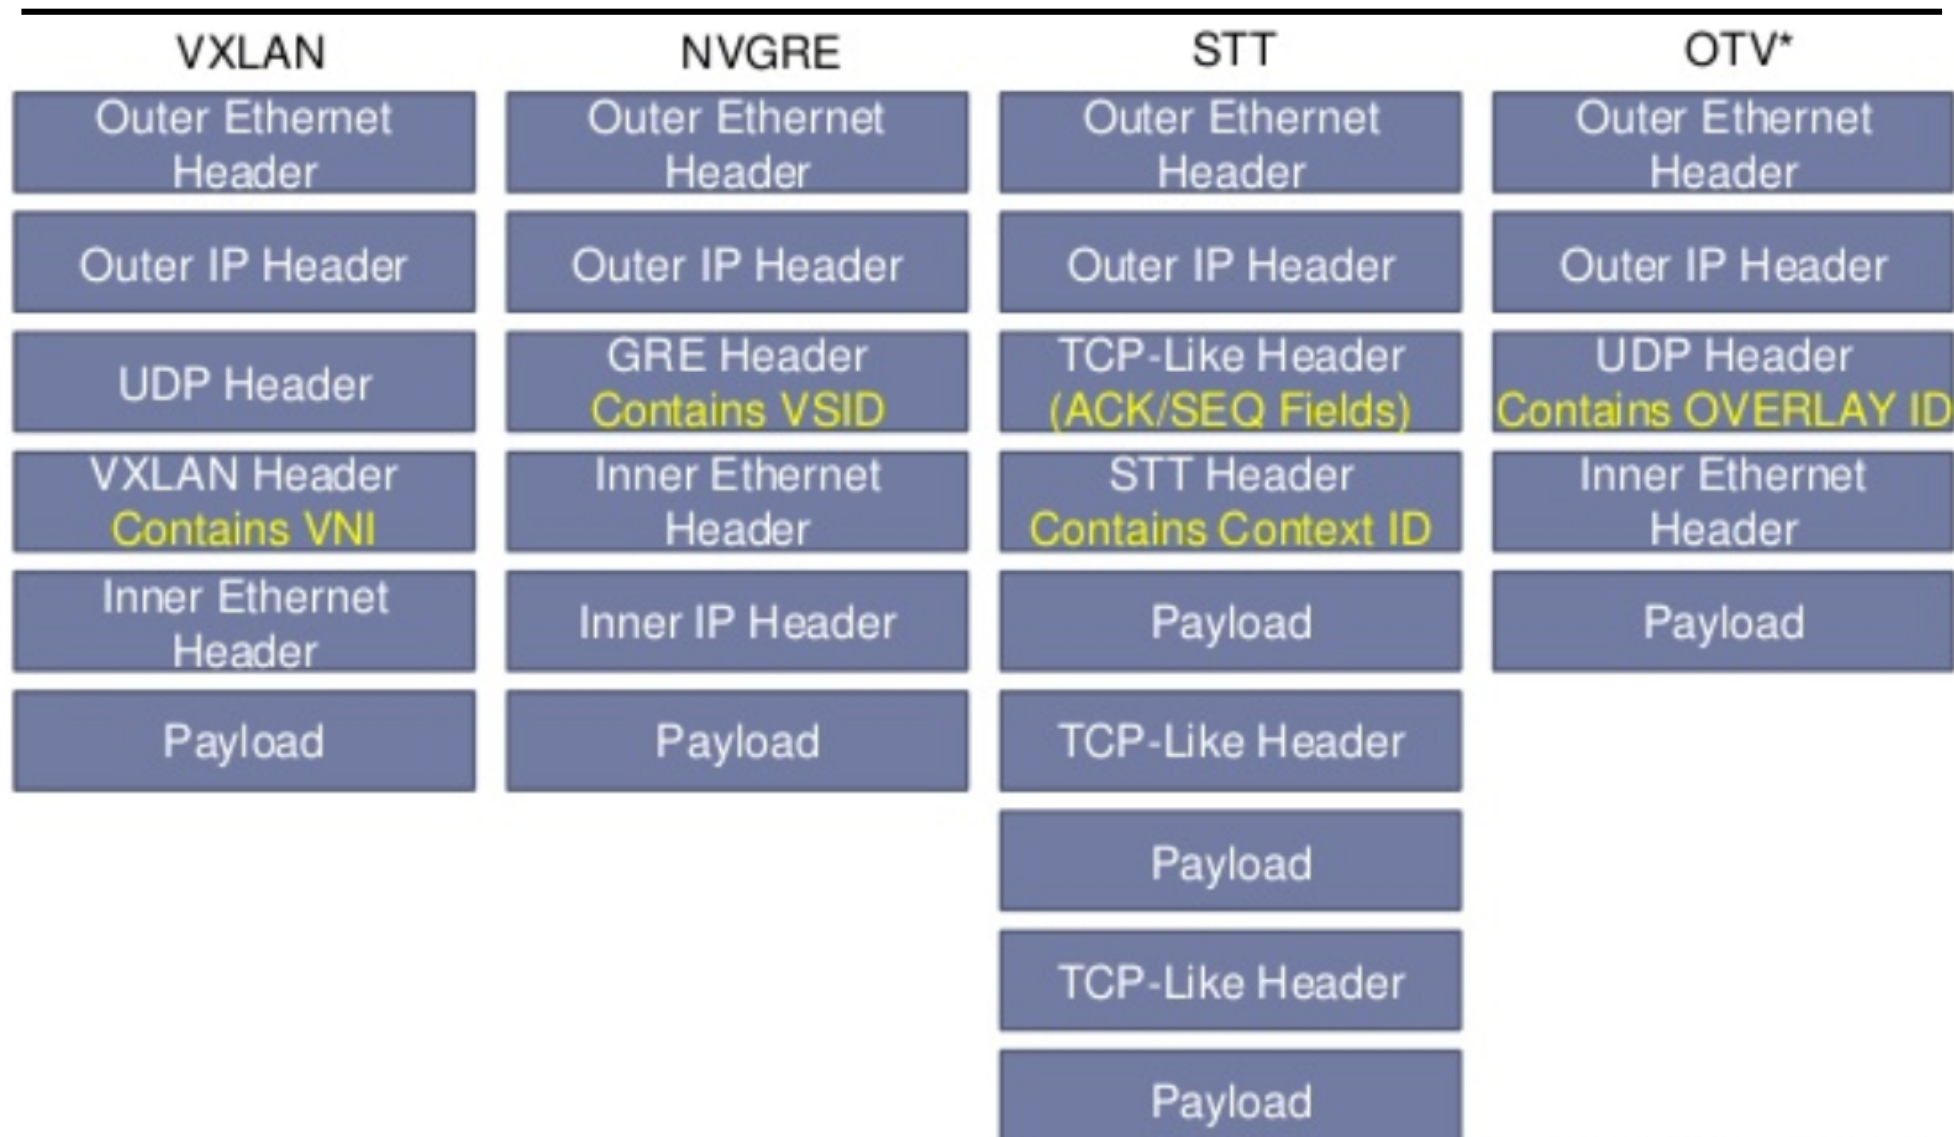
- Similar to other tunneling methods (L2TPv3, AtOM, VPLS)
- Encapsulation via tunnel end-points
- Not dependent on specific transports
- Layer 2 over Layer 3
- Easier to set up “customer edge”

Common IP-based Encapsulation Methods

Method	Full Name	Sponsors	Approach
DOVE	Distributed Overlay Virtual Ethernet	IBM	Leverages OTV and VxLAN
NVGRE	Network Virtualization using Generic Routing Encapsulation	Arista Networks, Broadcom, Dell, HP, Intel, Microsoft	24-bit Virtual Subnet Identifier (VSI) in GRE Header
OTV	Overlay Transport Virtualization	Cisco	VLAN extension via GRE/MPLS (Nexus 7000)
STT	Stateless Transport Tunneling	Nicira(VMWare)	64-bit context ID in STT header, "TCP like" header, leverages NIC hardware resources, not ratified
VxLAN	Virtual Extensible Local Area Network	Arista Networks, Broadcom, HP, Cisco, Citrix, Red Hat, VMWare	24-bit VxLAN Network Identifier (VNI) in VxLAN header inside UDP packet

- Different approaches to destination endpoint identification
- Different approaches to load balancing for efficiency
- Can be negatively impacted by "middle boxes"
- Some increased exposure to Mac-over-IP security threats


Encapsulation Headers



*According to Cisco specification it's a UDP header. In reality it turns out to be GRE+MPLS header that is being used.

Where would be more likely to find SDN in the recent future?

- Data Center
- Campus and Branch
- Access and Aggregation
- WAN
- Core
- Edge



What Lays Ahead?

Data Plane

▶ State of specifications

- ▶ Maturity
- ▶ Changes across releases

▶ Silicon Concerns

- ▶ Specifications outpace silicon development
- ▶ Merchant silicon not optimized for OpenFlow

▶ Performance

- ▶ Scalability of Flow-Matches (limited by TCAM size)
- ▶ Cost driver excludes rich multi-core xPU ecosystem

Control Plane

▶ Scalability

- ▶ Centralized vs Distributed
- ▶ State coherence between control and data plane

▶ Interoperability

- ▶ SDN to non-SDN
- ▶ Inter-Controller
- ▶ Multi-orchestrator conflicts
- ▶ Virtual overlays

Credits

- Software-Defined Networking: A Comprehensive Survey, October 8, 2014
- ONF White Paper Software-Defined Networking: The New Norm for Networks, April 13, 2012
- Software-Defined Networking (SDN): Unleashing the Power of the Network, presentation by Robert Keahey
- Introduction to SDN (Software-defined Networking), David Mahler, November 3, 2014
- SDN@Google presentation, Amit Agrawal