Panopticon: Incremental SDN Deployment in Enterprise Networks

Stefan Schmid
with
Dan Levin, Marco Canini, Fabian Schaffert, Anja Feldmann

https://venture.badpacket.in

I  SDN!
How can I deploy it?

SDN: Where and why?

- Datacenters
- Wide-Area Networks (WANs)
- IXP
- …

SDN in Datacenters

- Characteristics?
- Why SDN?
- How to deploy?

SDN in Datacenters

- Characteristics?
  - Already highly virtualized
  - Quite homogeneous!

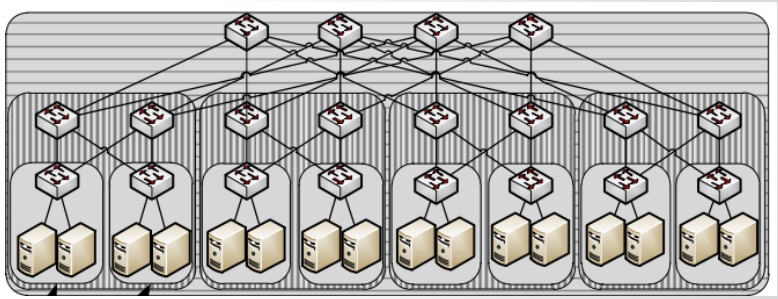

SDN in Datacenters

- Why SDN?
  - Realize fabric abstraction
  - Decouple application from physical infrastructure
  - Own addresses for tenants, VM migration, ...
  - Improve performance

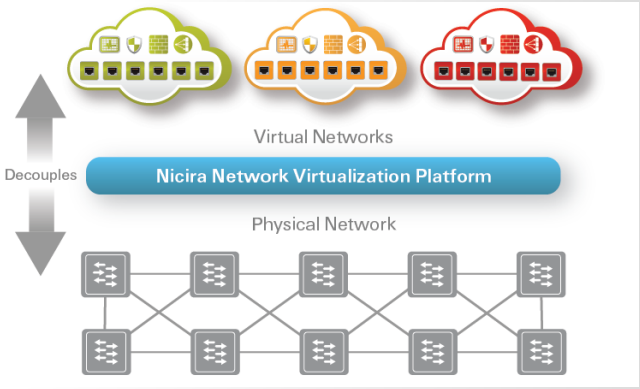

SDN in Datacenters

- How to deploy?
  - Run Open VSwitch on servers!
  - Edge deployment (inside: MPEC)

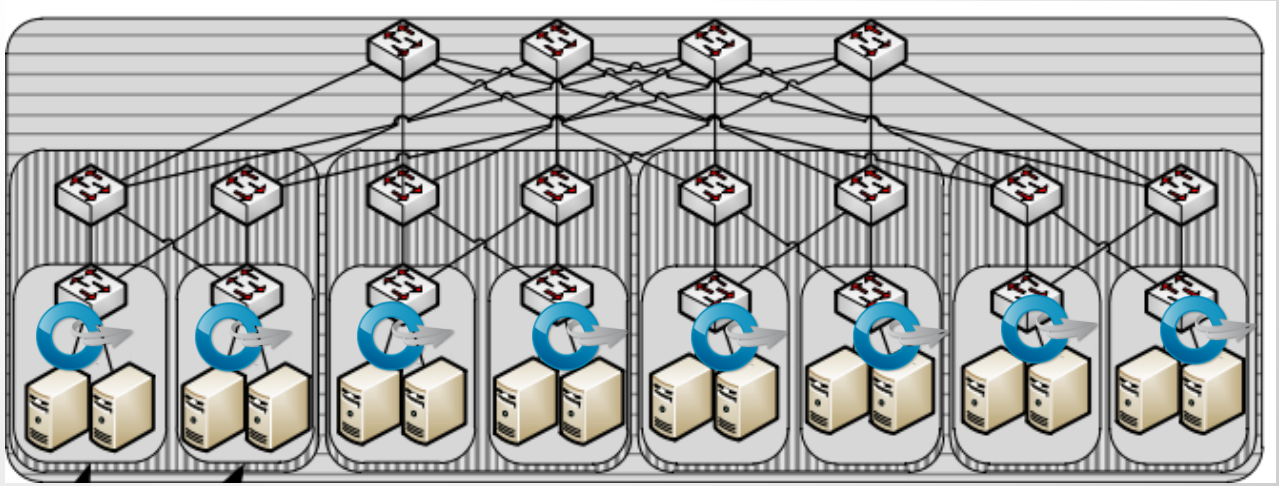

SDN in WANs

- Characteristics?
- Why SDN?
- How to deploy?

SDN in WANs

- Characteristics?
  - Bandwidth precious (WAN traffic grows fastest)
  - Latency matters
  - Probably not so many sites
  - Many different applications and requirements

SDN in WANs

- Why SDN?
  - E.g., Google B4
  - Improve utilization
  - Differentiate applications (latency sensitive Google docs vs datacenter synchronization)

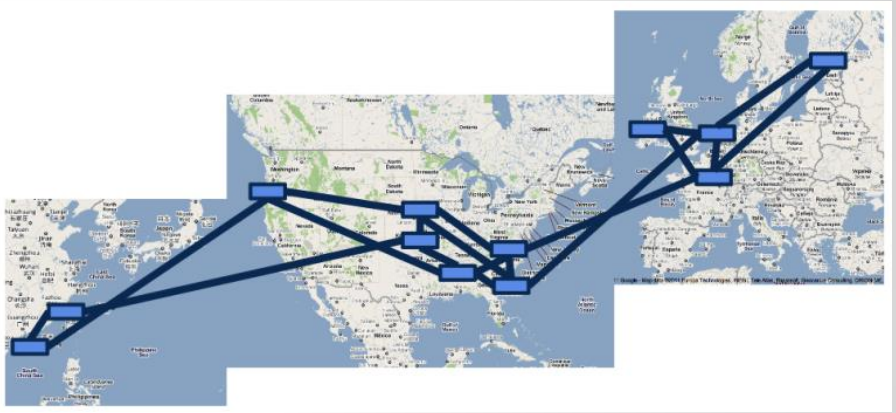

SDN in WANs

- How to deploy?
  - Replace IP “core” routers (running BGP) at border of datacenter (end of long-haul fiber)
  - Gradually replace routers

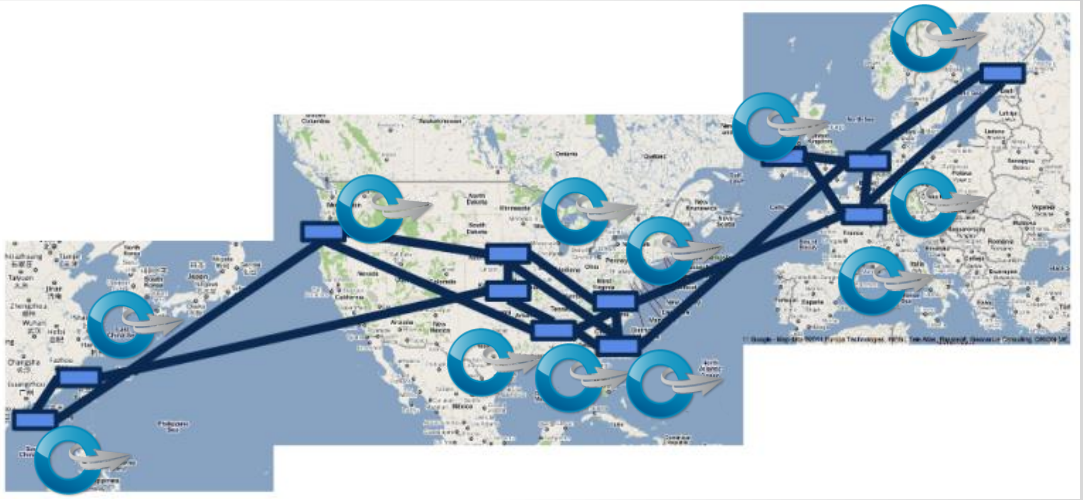

- However, benefits arise only after complete hardware overhaul of network (after years)

SDN in WANs

- Ho
  - F
    r
    a
    (
  - G
  - H
    o

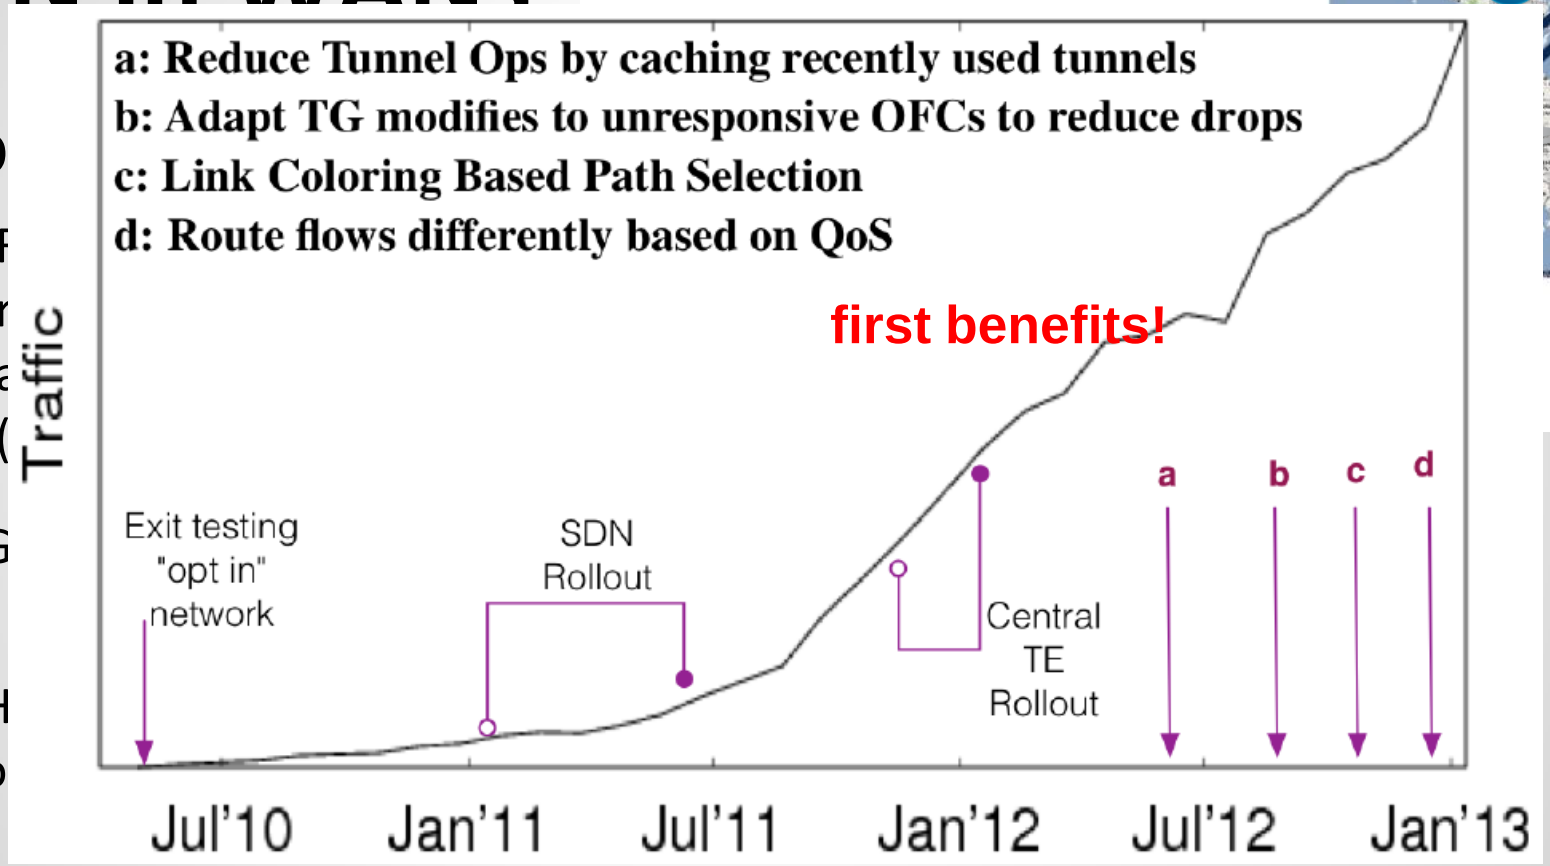

Our Use Case: Enterprise Networks

- What's wrong with enterprise networks?

Problems with traditional networks

Manual, Device-Centric Network Configuration
(CLI, LANmanager)

Un-evolved Best Practices
(tcpdump, traceroute - from the 1990s)

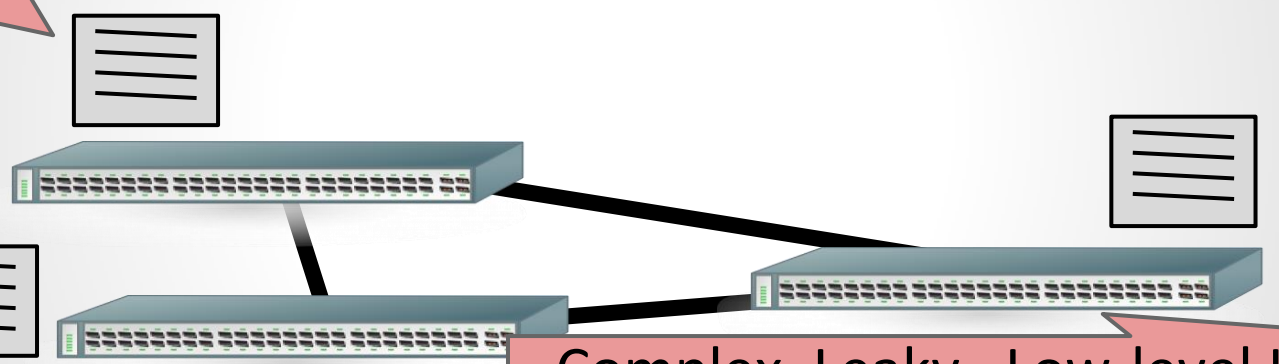

Current Networks

Complex, Leaky, Low-level Interfaces
(VLANs, Spanning Tree, Routing)

Software Defined Networking

Automation und Abstraction for Networks

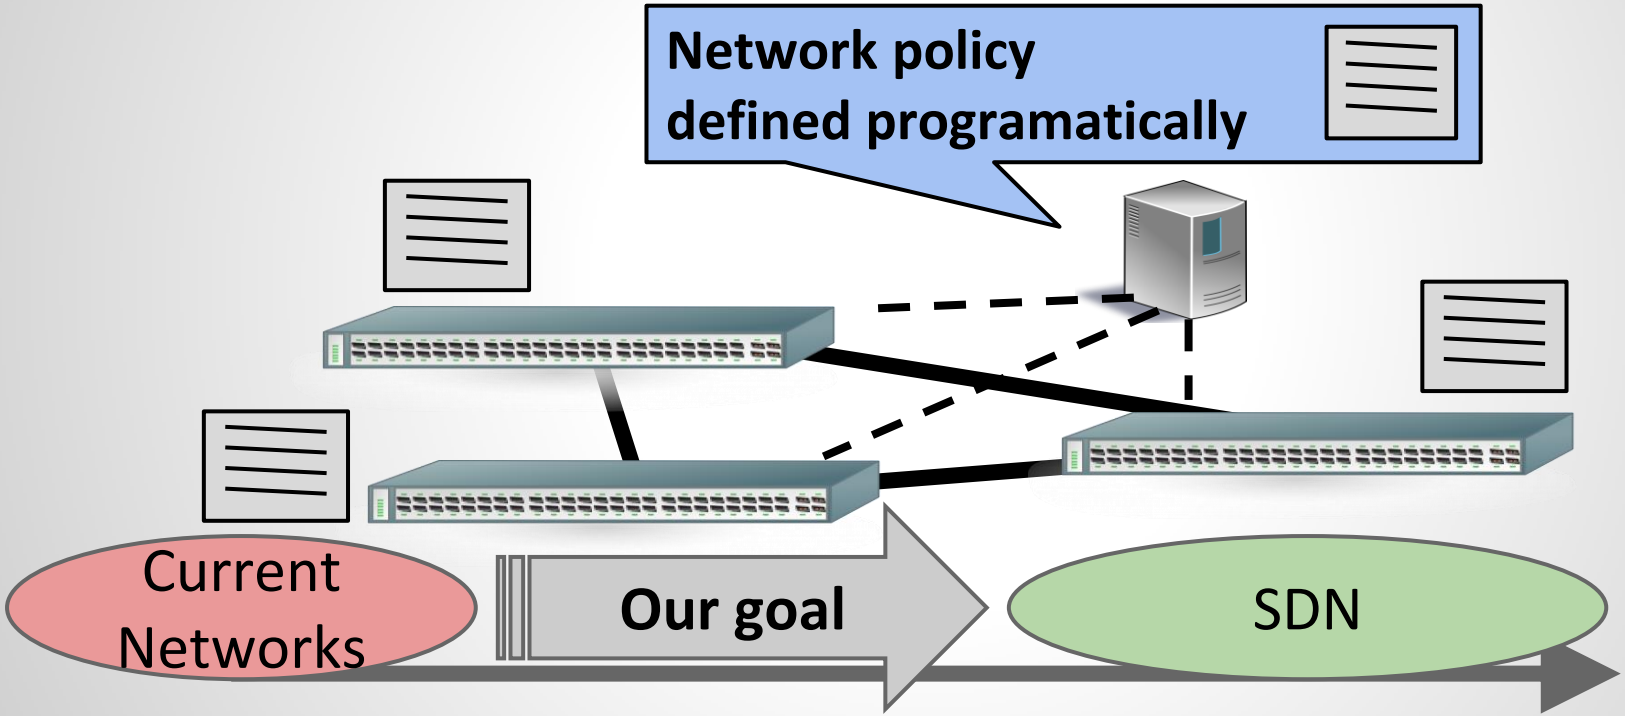

SDN Interface

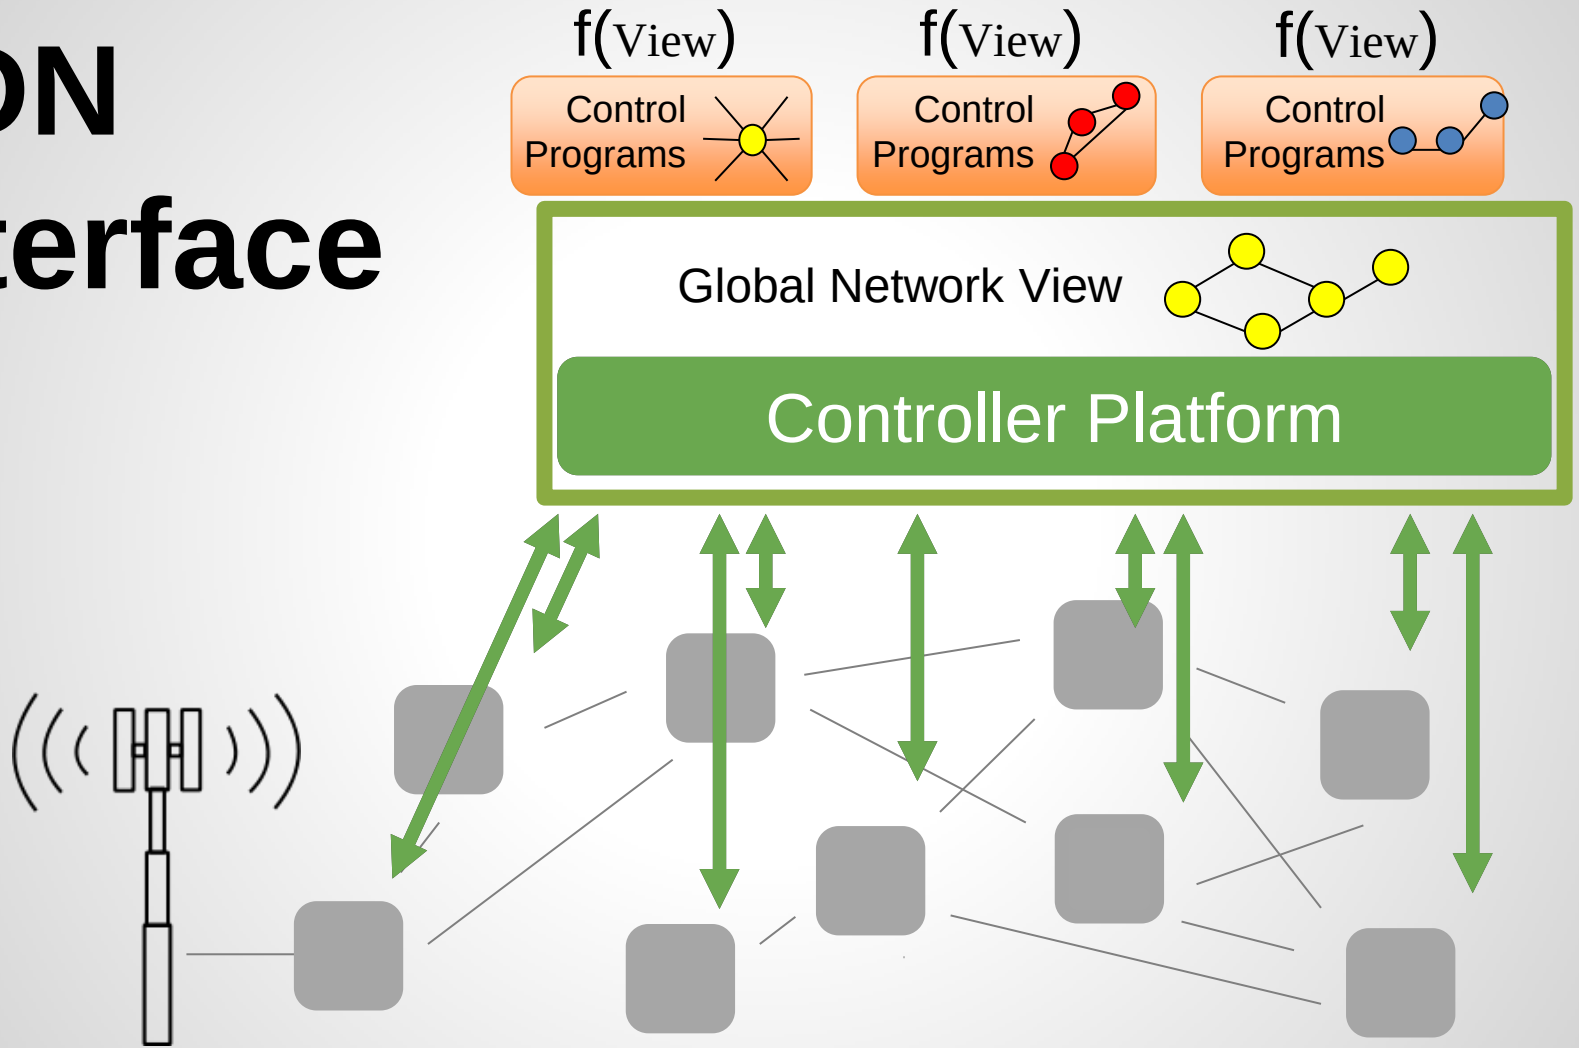

The SDN Deployment Problem

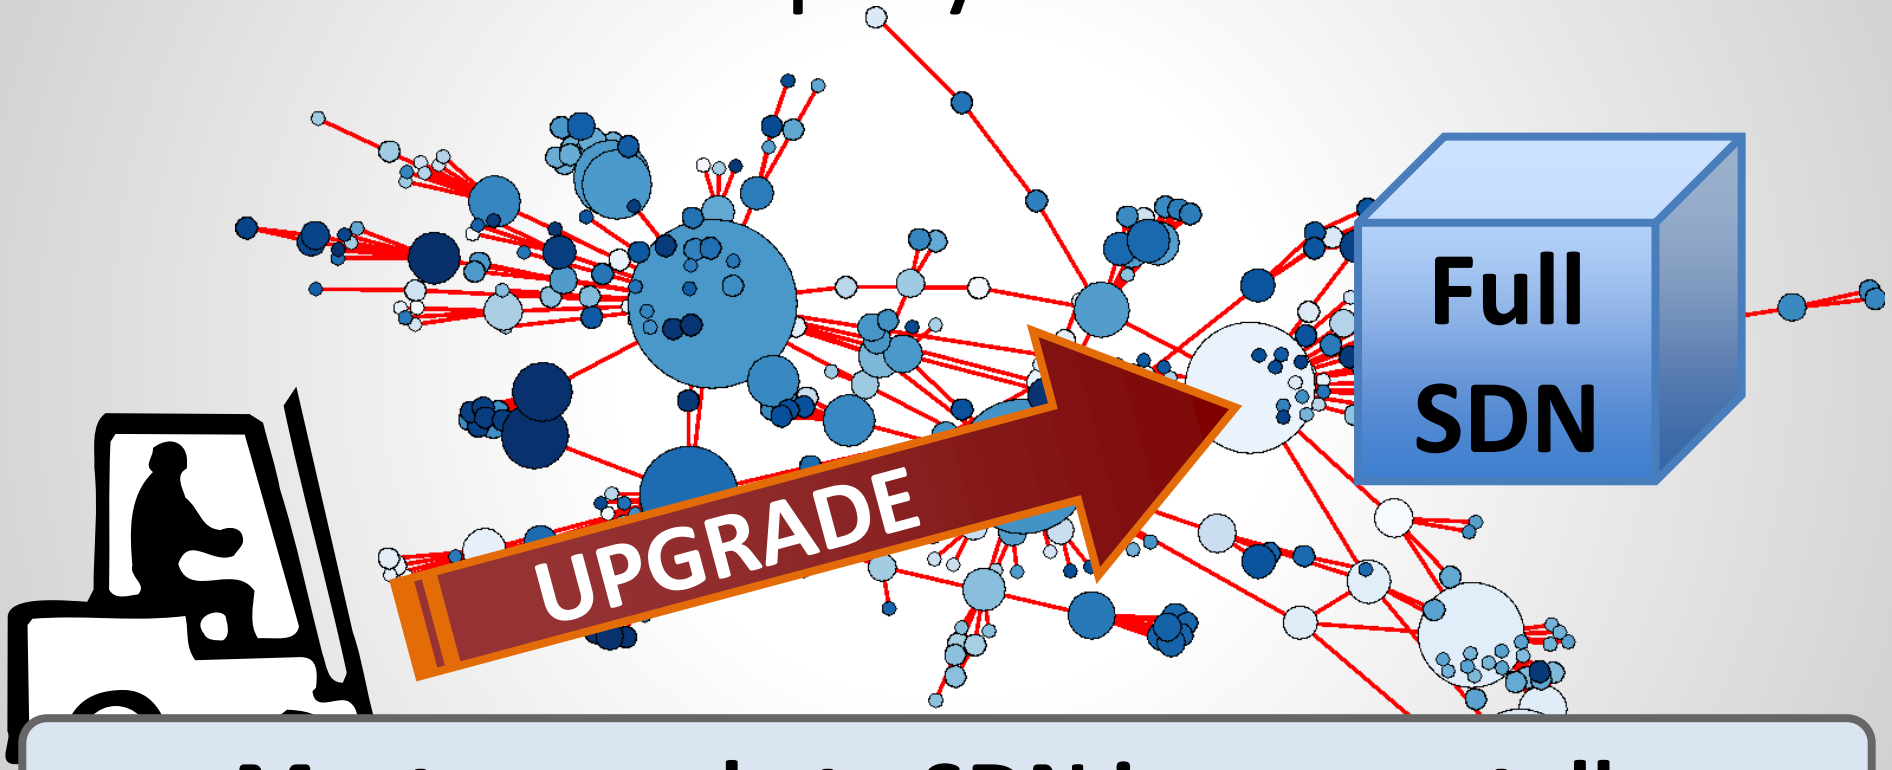

Must upgrade to SDN incrementally

Key Questions

- How can we **incrementally deploy SDN** into enterprise campus networks?

- What **SDN benefits** can be realized in a hybrid deployment?

Current Hybrid Networks

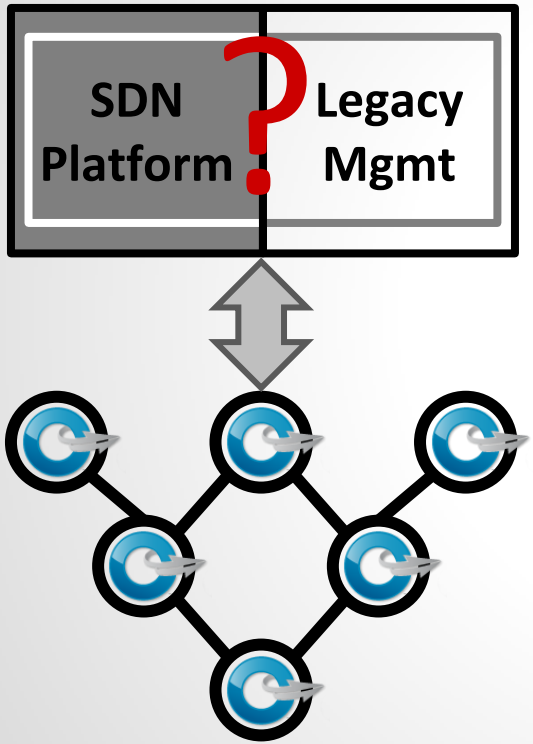

Dual-stack approach

E.g., part of the Stanford network.

Current Hybrid Networks

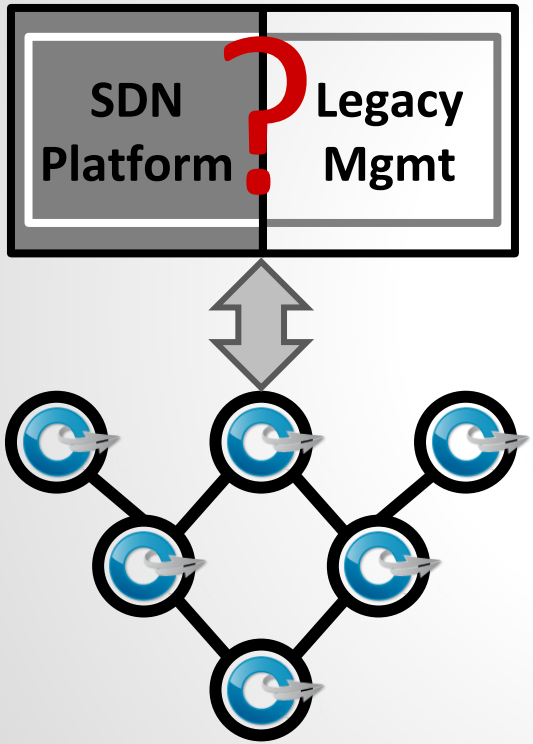

Dual-stack approach

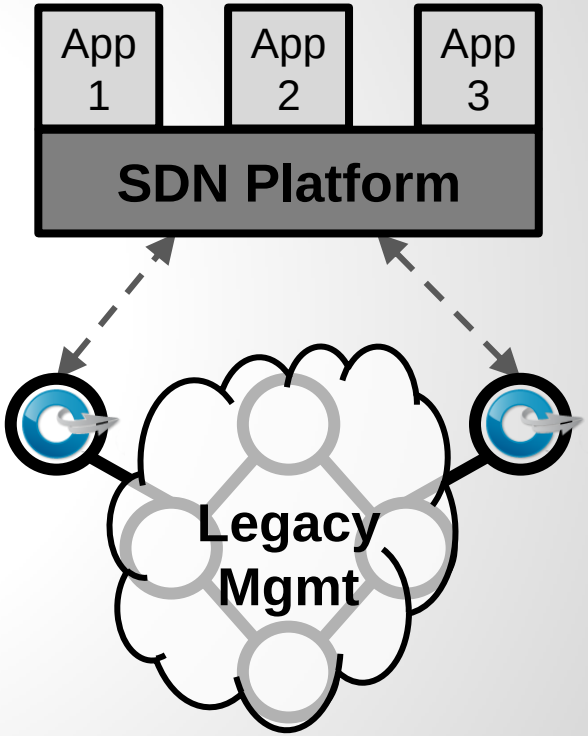

Edge-only approach

The edge is legacy access switches

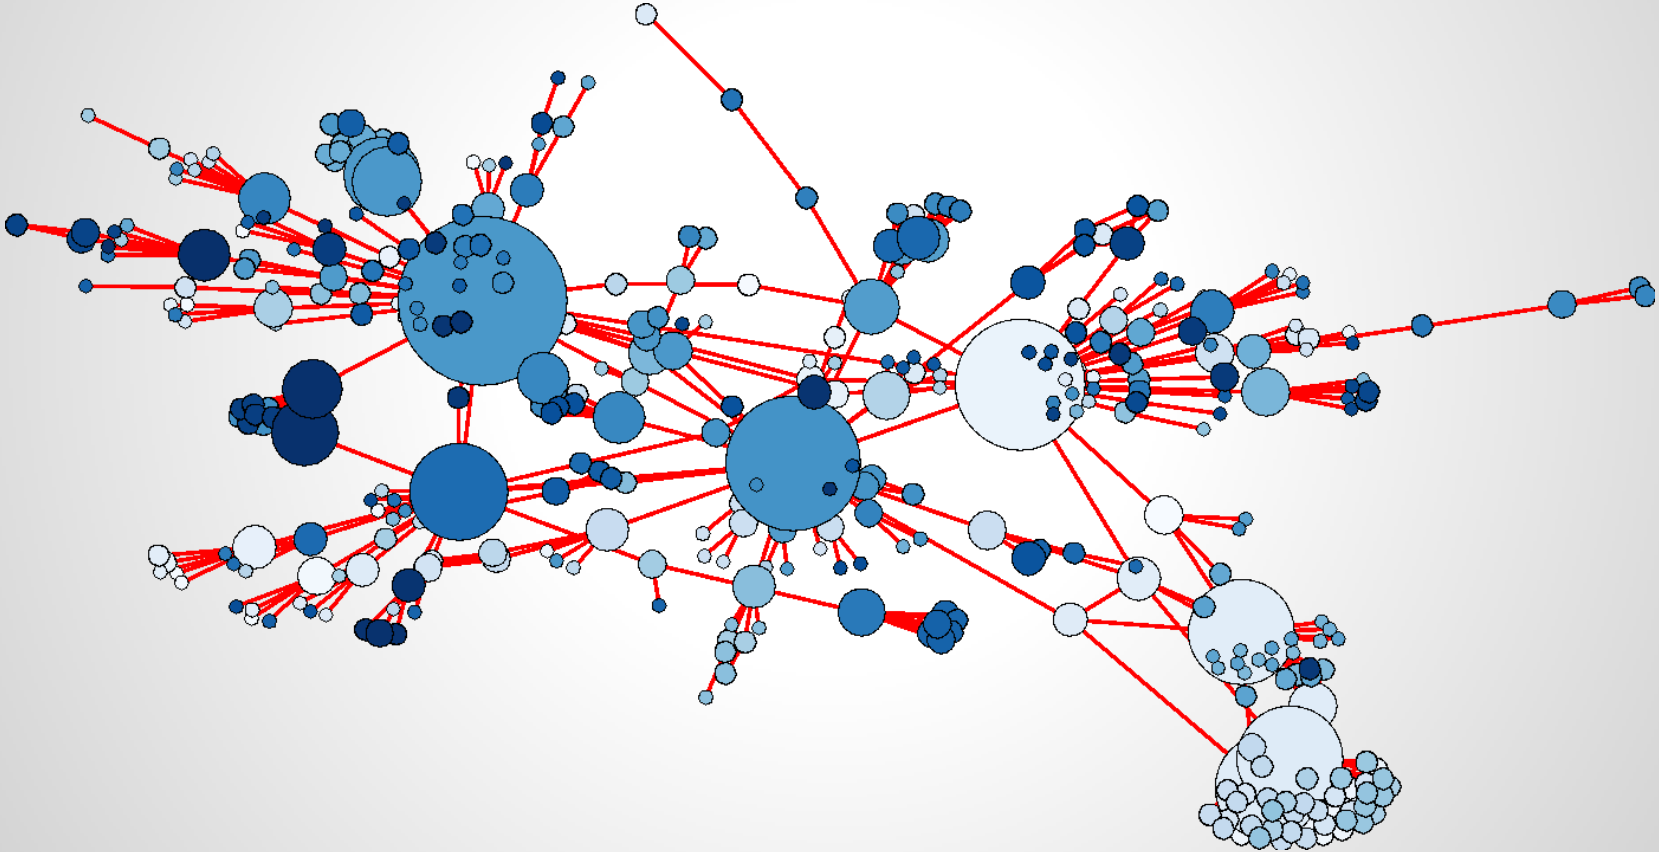

PANOPTICON

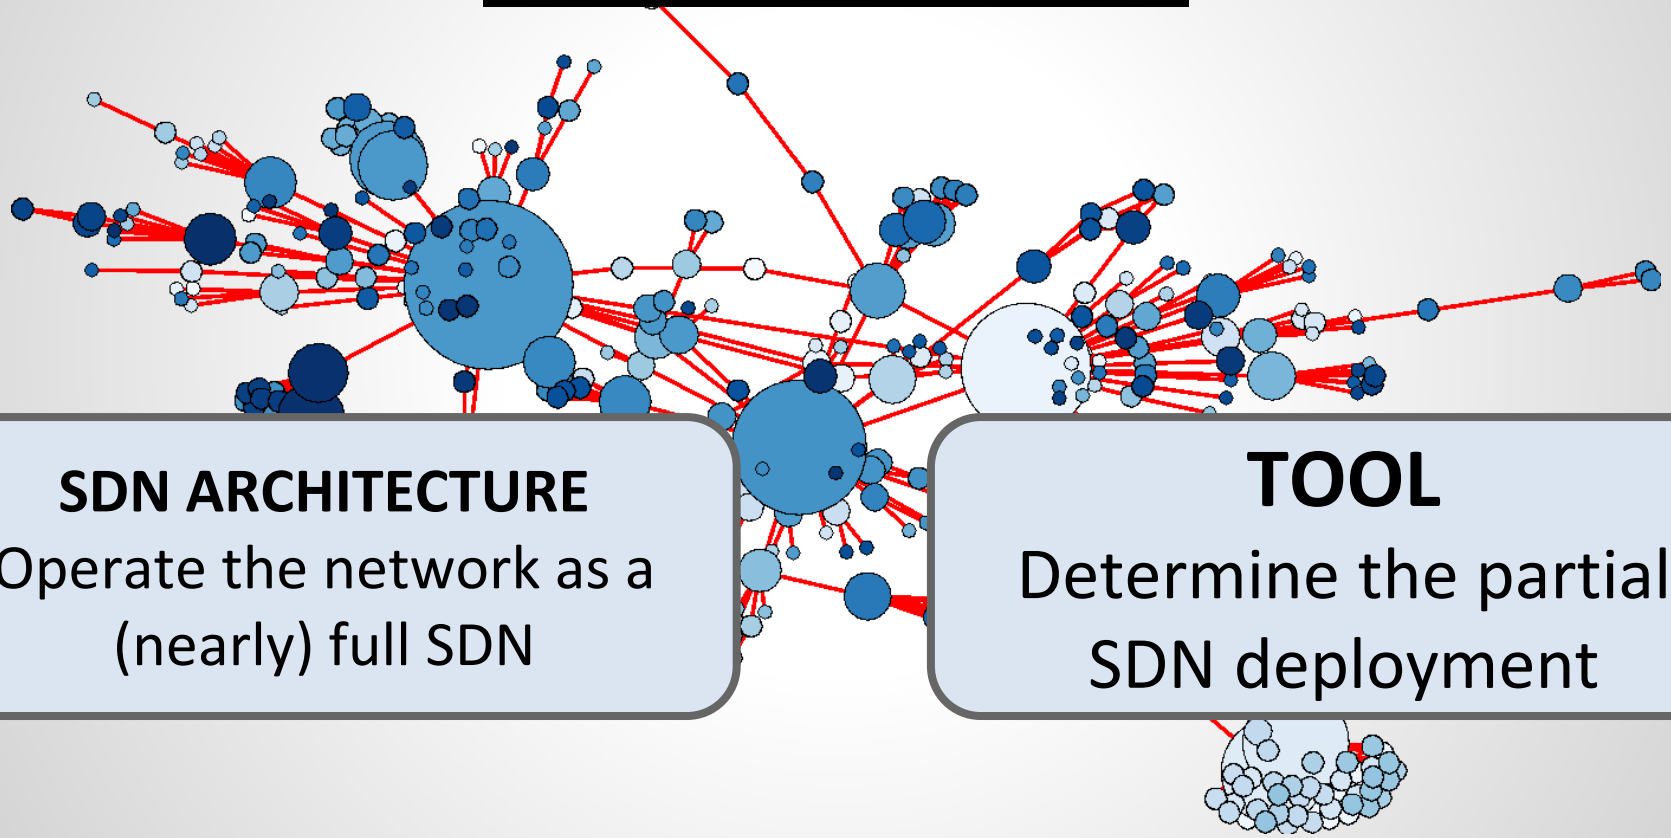

SDN ARCHITECTURE

Operate the network as a (nearly) full SDN

TOOL

Determine the partial SDN deployment

The Existing Network

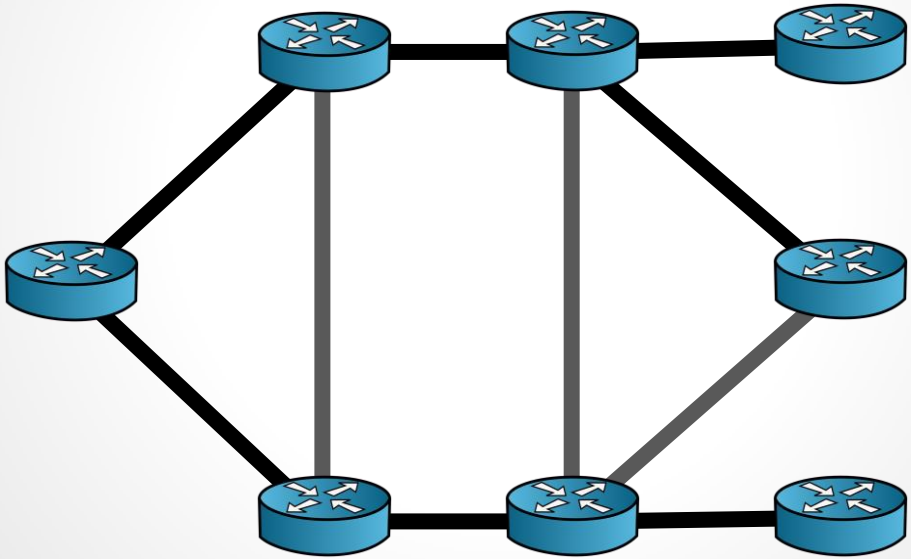

1. Planning the SDN Deployment

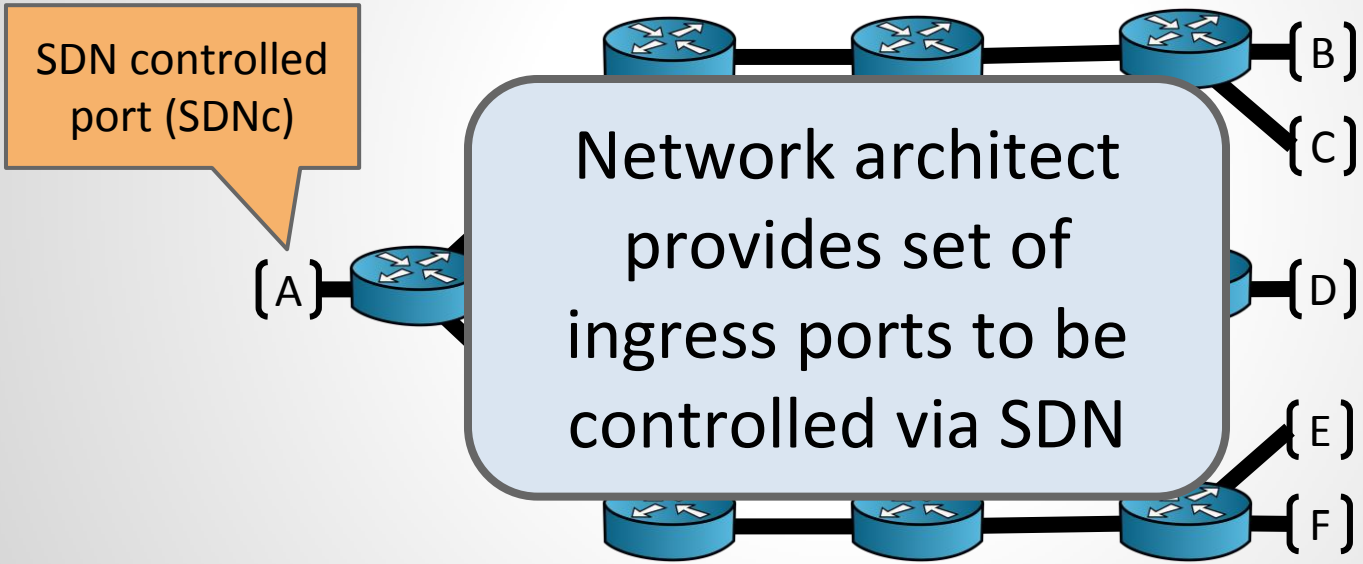

Network topology

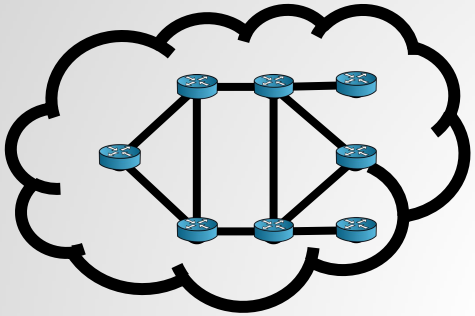

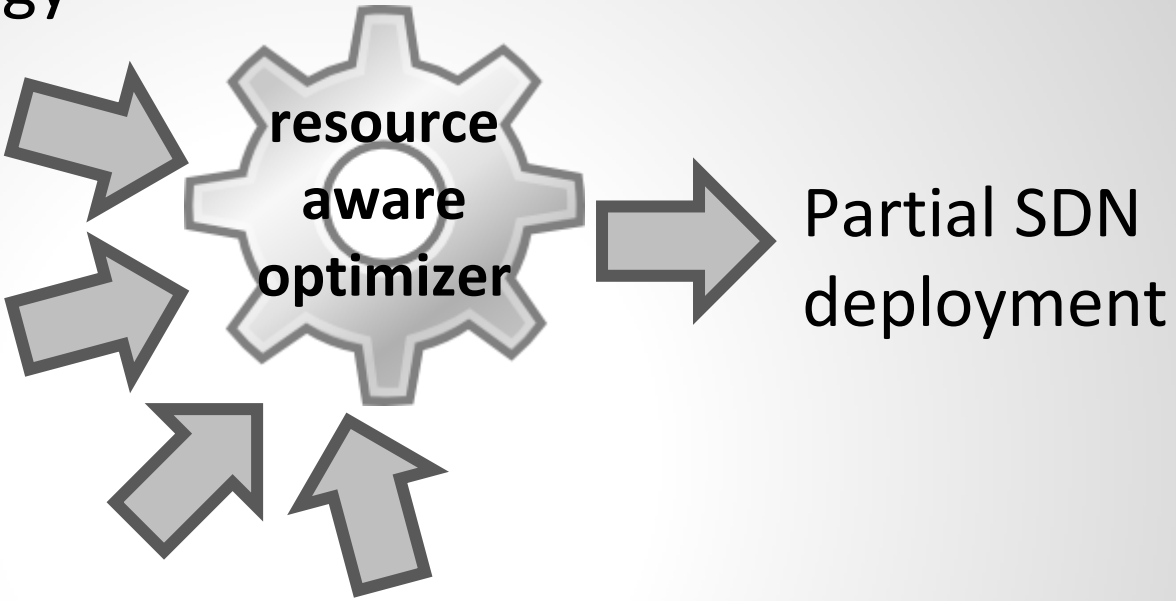

Partial SDN deployment

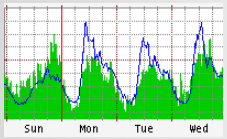

Traffic estimates

Objectives

- Upgrade budget
- Path delay

Tunable parameters

- Price model
- Utilization thresholds
  (link utilization, VLANs, etc.)

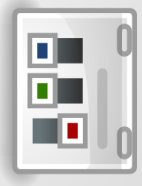

The Partial SDN Deployment ()

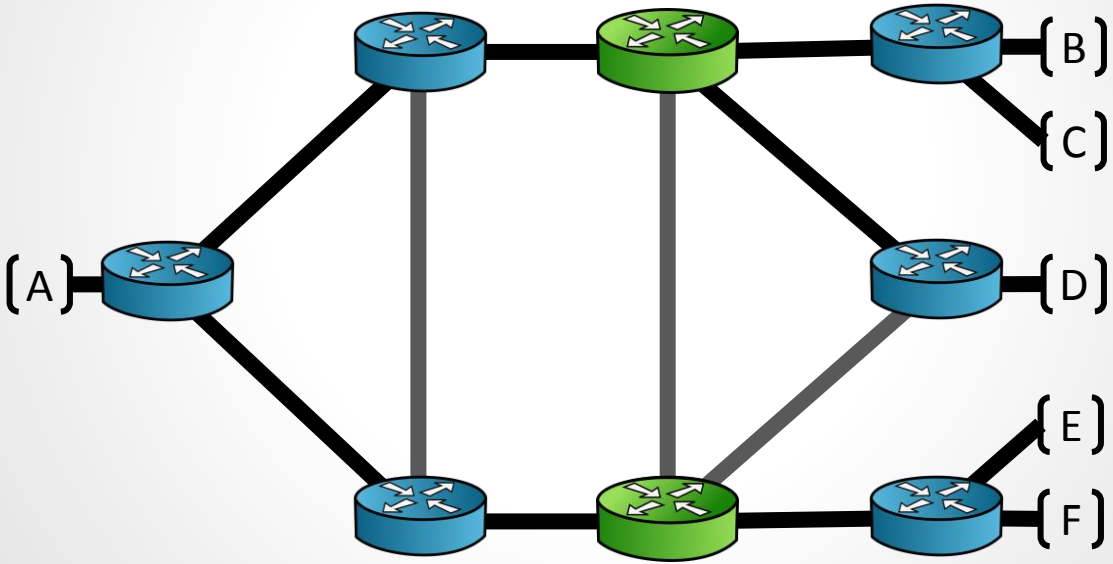

Benefits of Partial SDN Deployment?

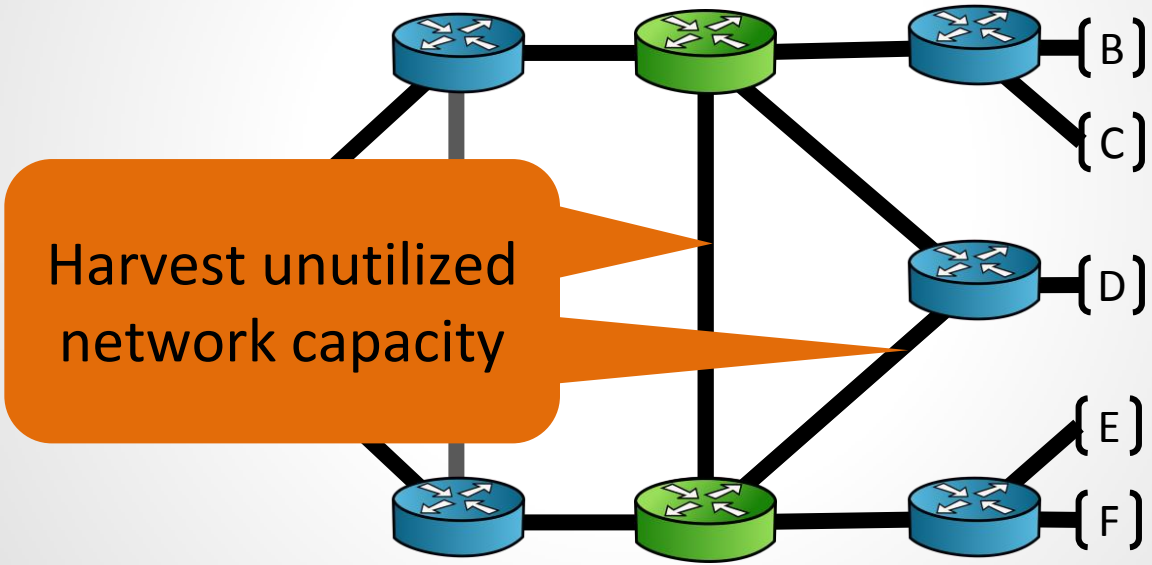

Main benefits of **SDN**= Principled orchestration of the network policy

Can partial SDN deployment still take advantage of *principled network orchestration*

2. Realizing the Benefits of SDN

Insight #1:
≥ 1 SDN switch →
Policy enforcement

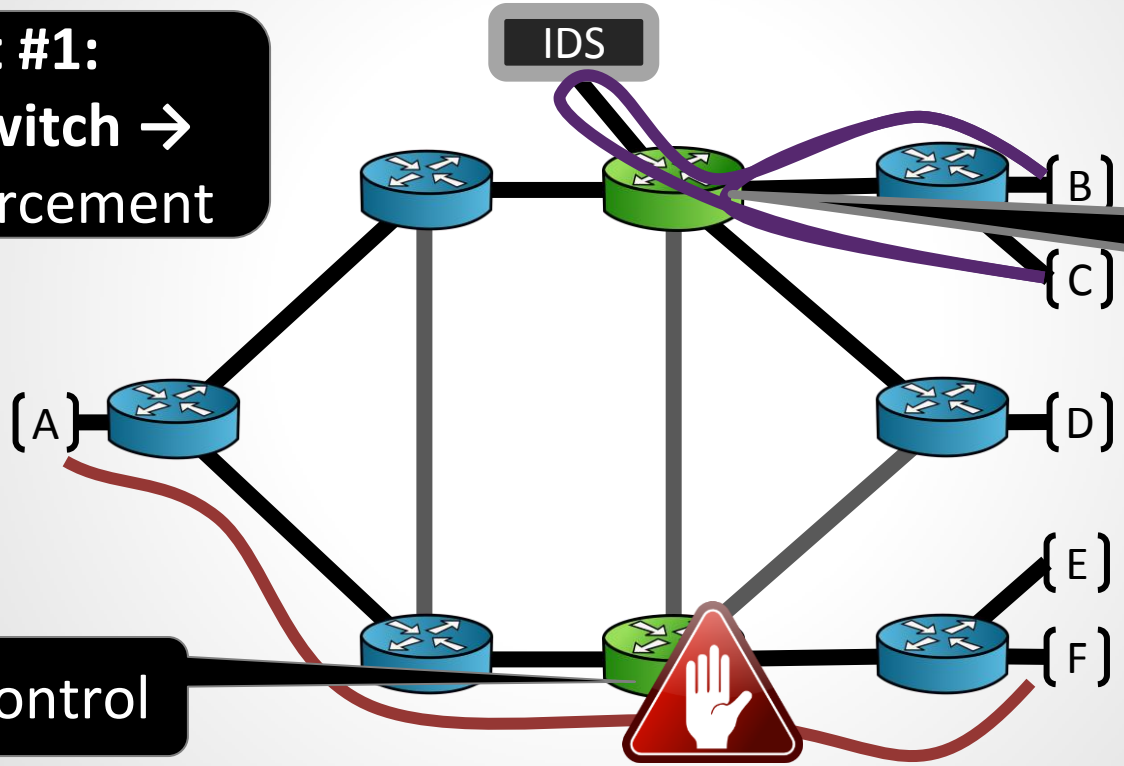

Middlebox traversal

Access control

2. Realizing the Benefits of SDN

Insight #1:
≥ 1 SDN switch →
Policy enforcement

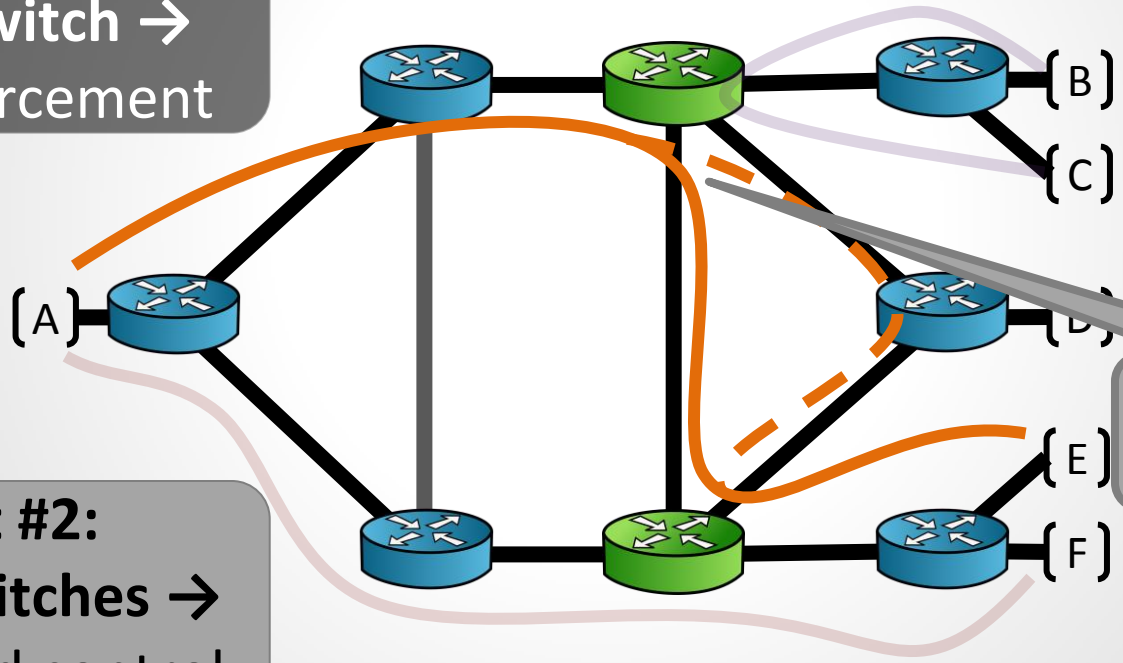

Traffic load-balancing

Insight #2:
≥ 2 SDN switches →
Fine-grained control

Insight #1:
≥ 1 SDN switch →
Policy enforcement

Insight #2:
≥ 2 SDN switches →
Fine-grained control

Ensure that all traffic to/from an SDN-controlled port always **traverses at least one SDN switch**

SDN Waypoint Enforcement

Legacy devices must direct traffic to SDN switches

The **PANOPTICON** SDN Architecture

Conceptually group SDN ports in **Cell Blocks**

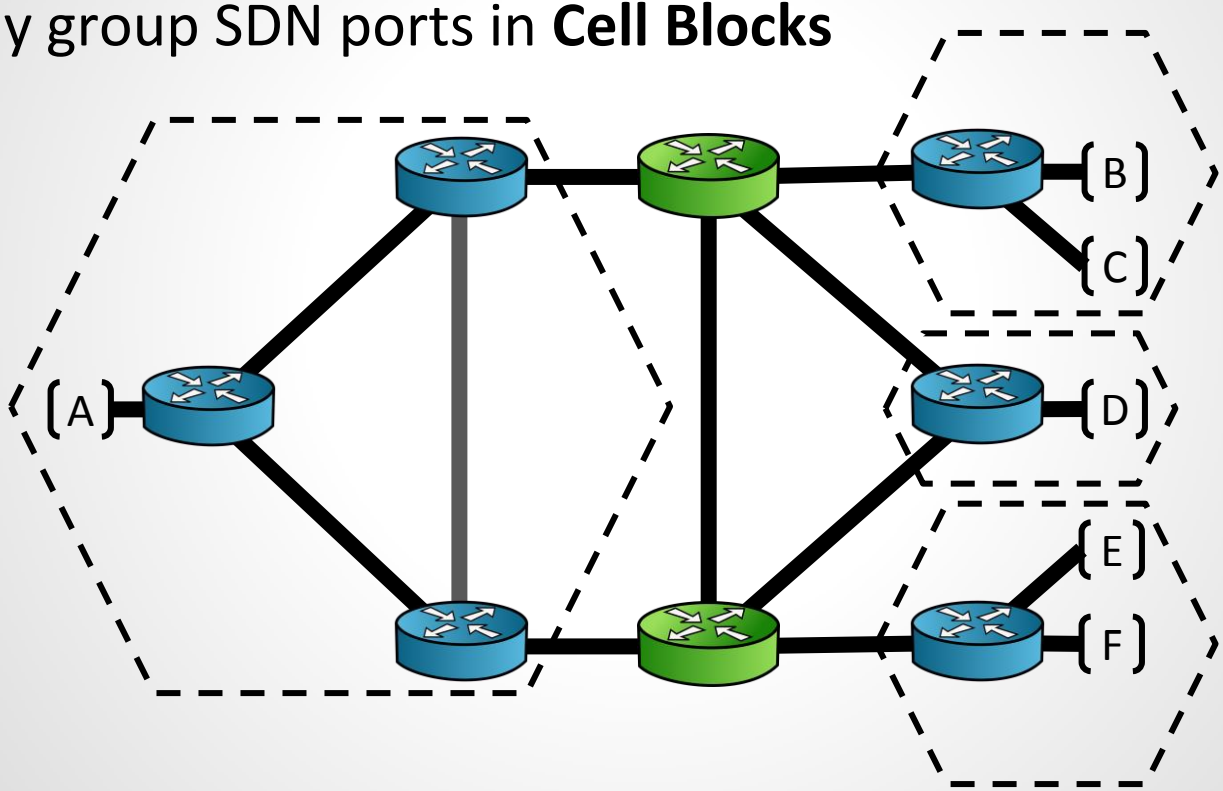

The **PANOPTICON** SDN Architecture

Traffic restricted by using VLANs

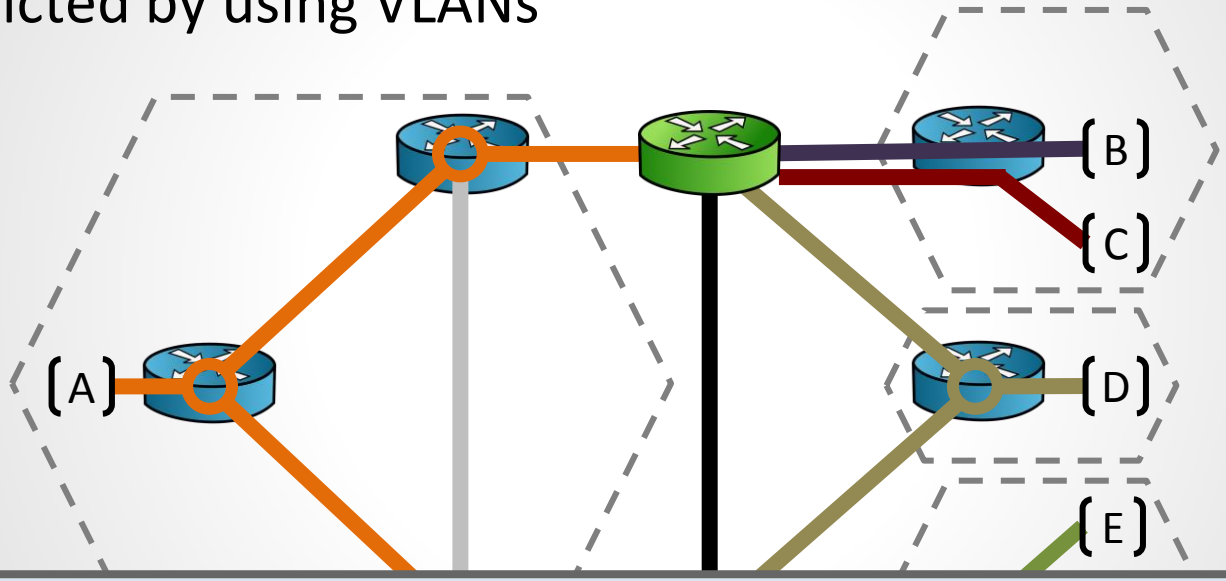

Per-port spanning trees that ensure waypoint enforcement

PANOPTICON

“Logical SDN”

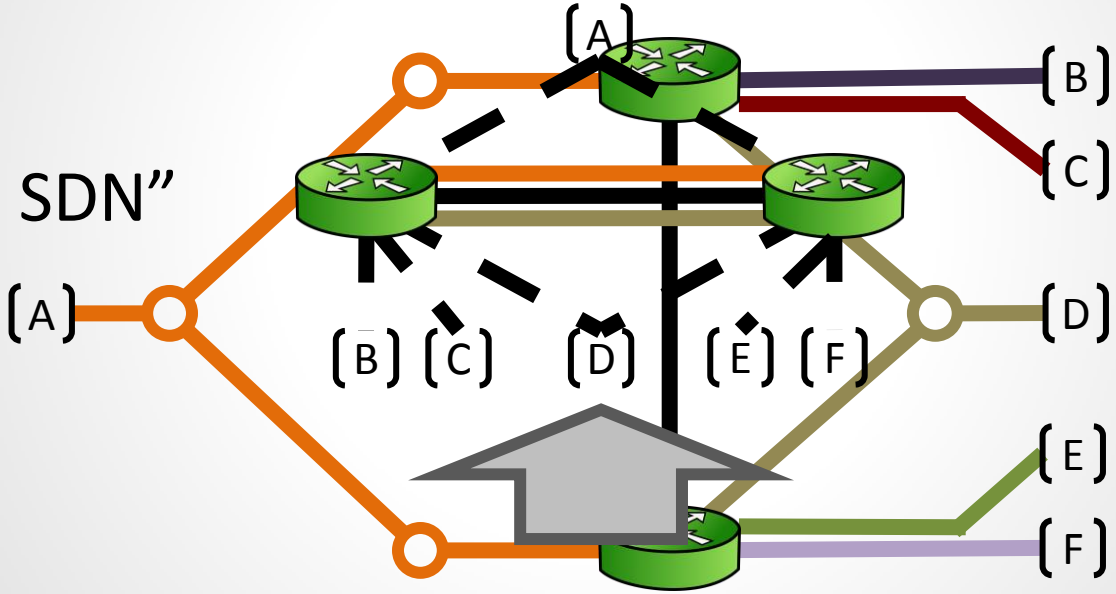

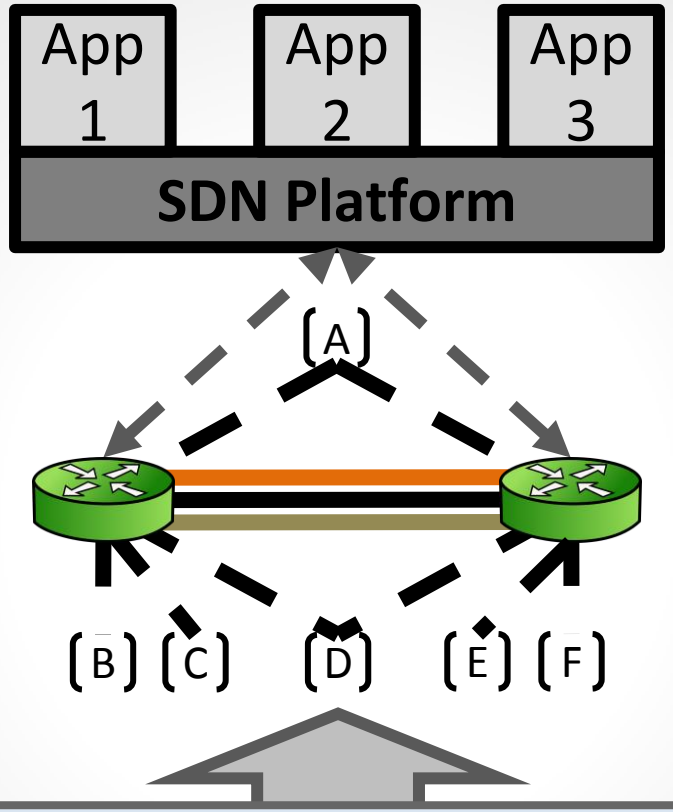

"Logical SDN"

PANOPTICON provides the abstraction of a (nearly) fully-deployed SDN in a partially upgraded network

What is the value of
a logical SDN

Use Case: Planned Maintenance

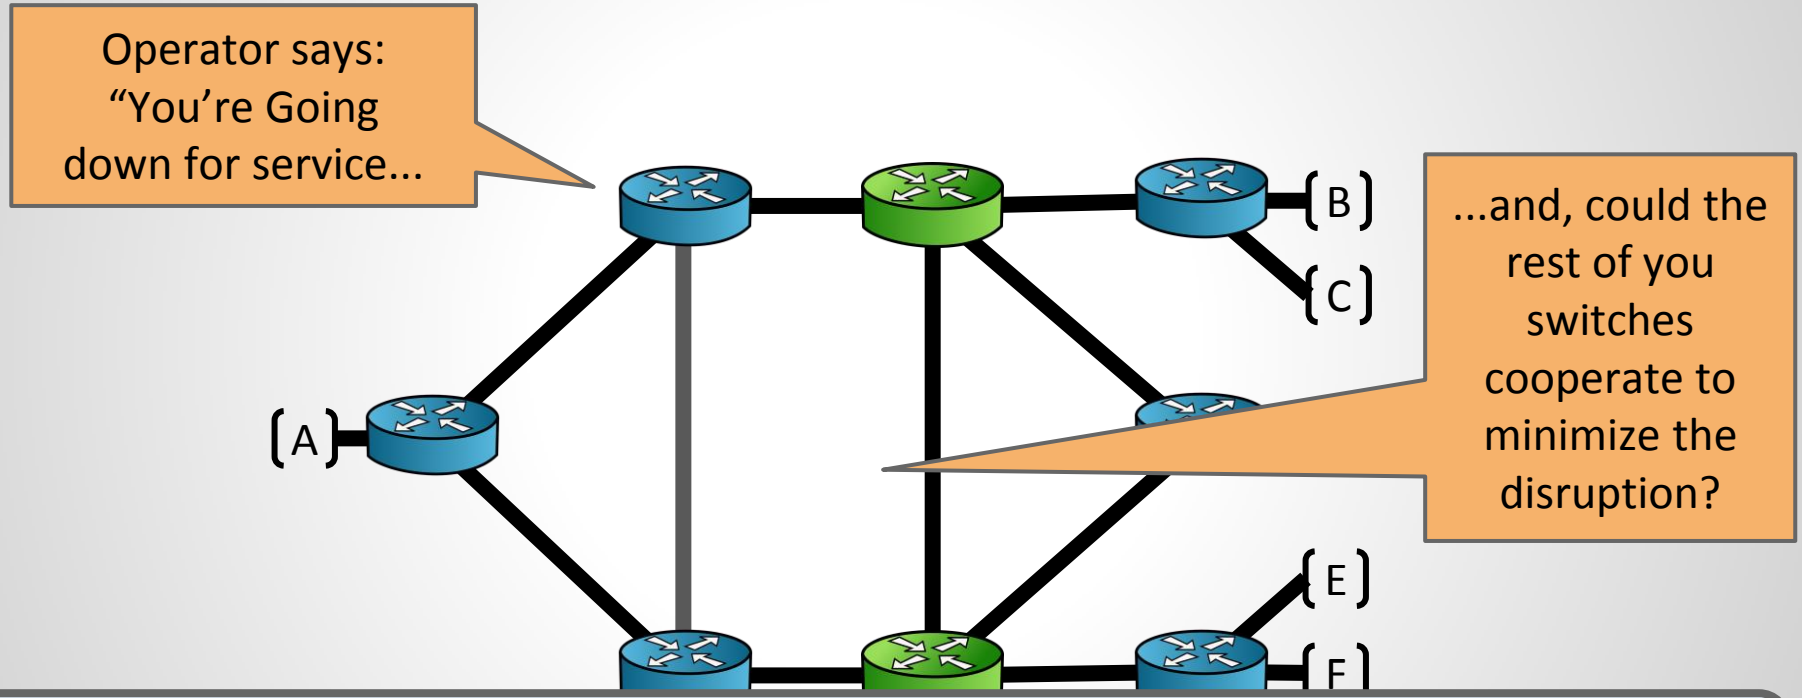

Let software worry about the dependencies, not the human operator!

Use Case Testbed Evaluation

2x NEC IP8800 (OF 1.0)

1x Cisco C3550XL

3x Cisco

2x HP 54

TCP Connection Recovery Time

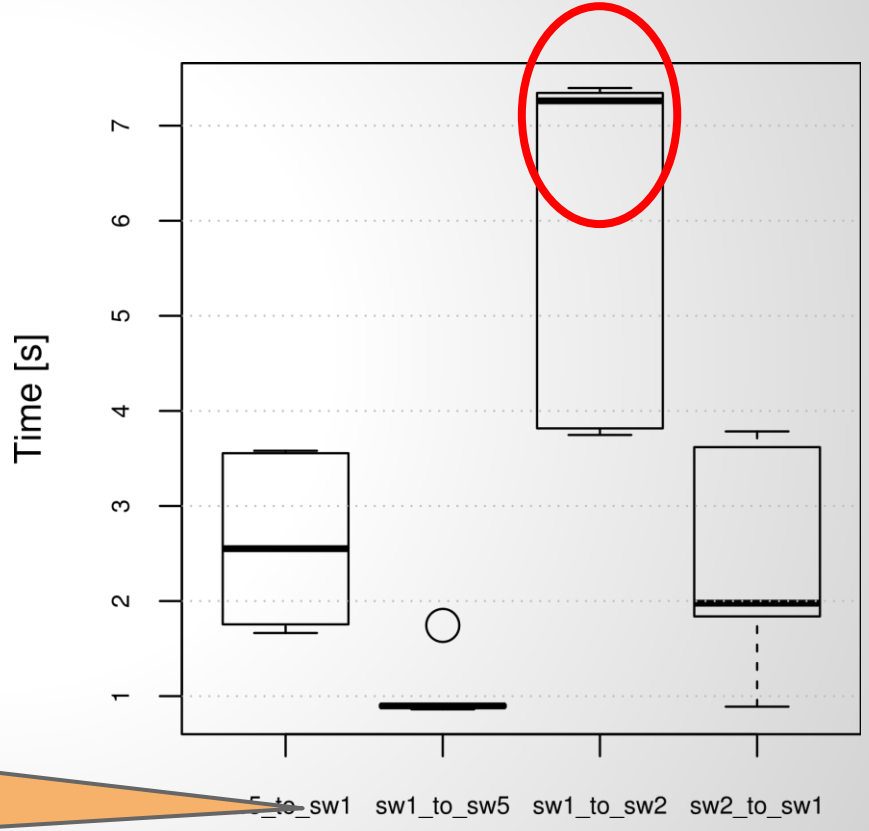

Locations of “port-down” events along one path traversing SDN switch.

Use Case: Planned Maintenance

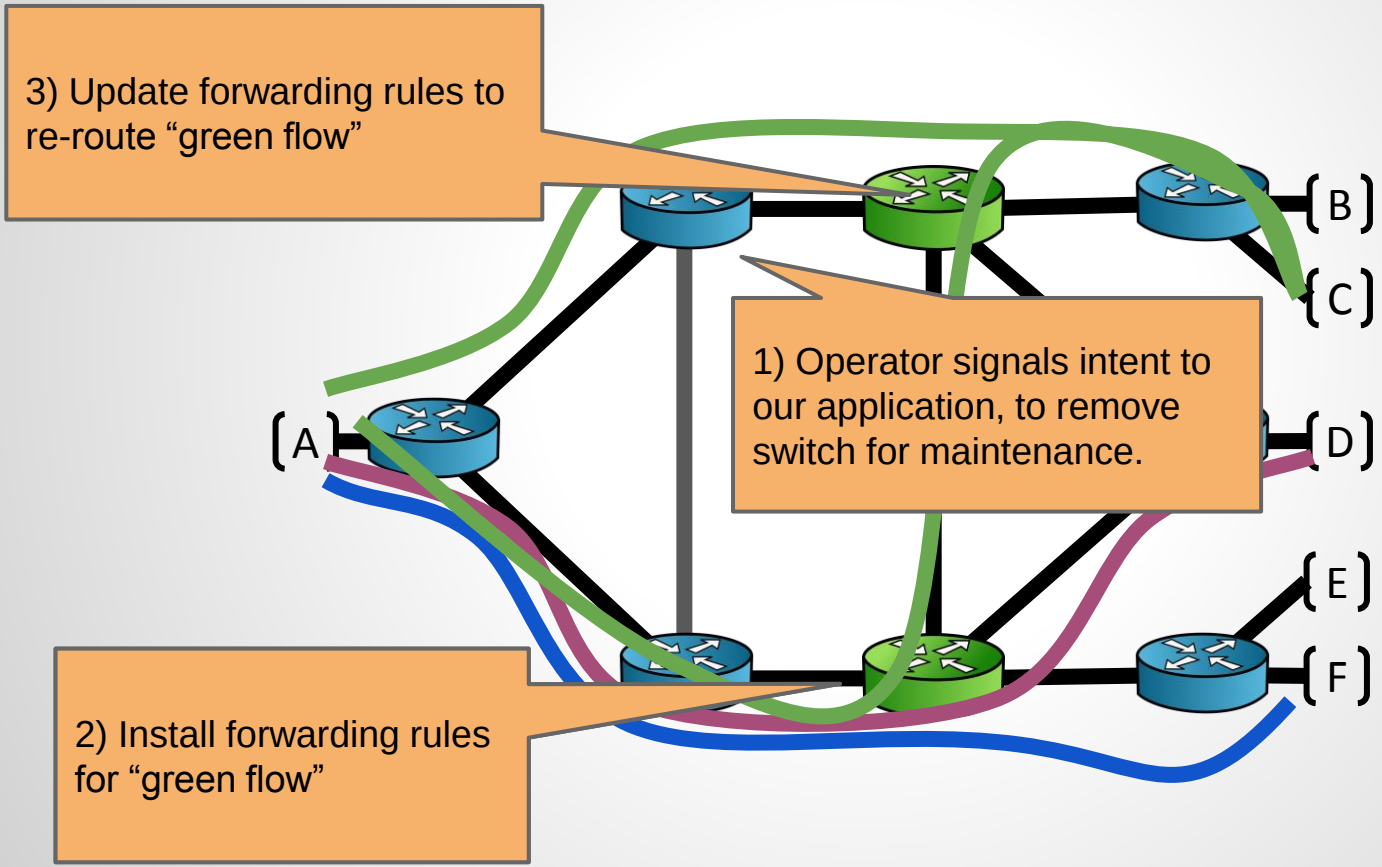

Use Case Testbed Evaluation

2x NEC IP8800 (OF 1.0)

1x Cisco C3550XL

3x Cisco C2960G

2x HP 5406zl

1x Pica8 3290

What is the impact of partial SDN deploymenton the **network traffic**

Potential Impact on Traffic

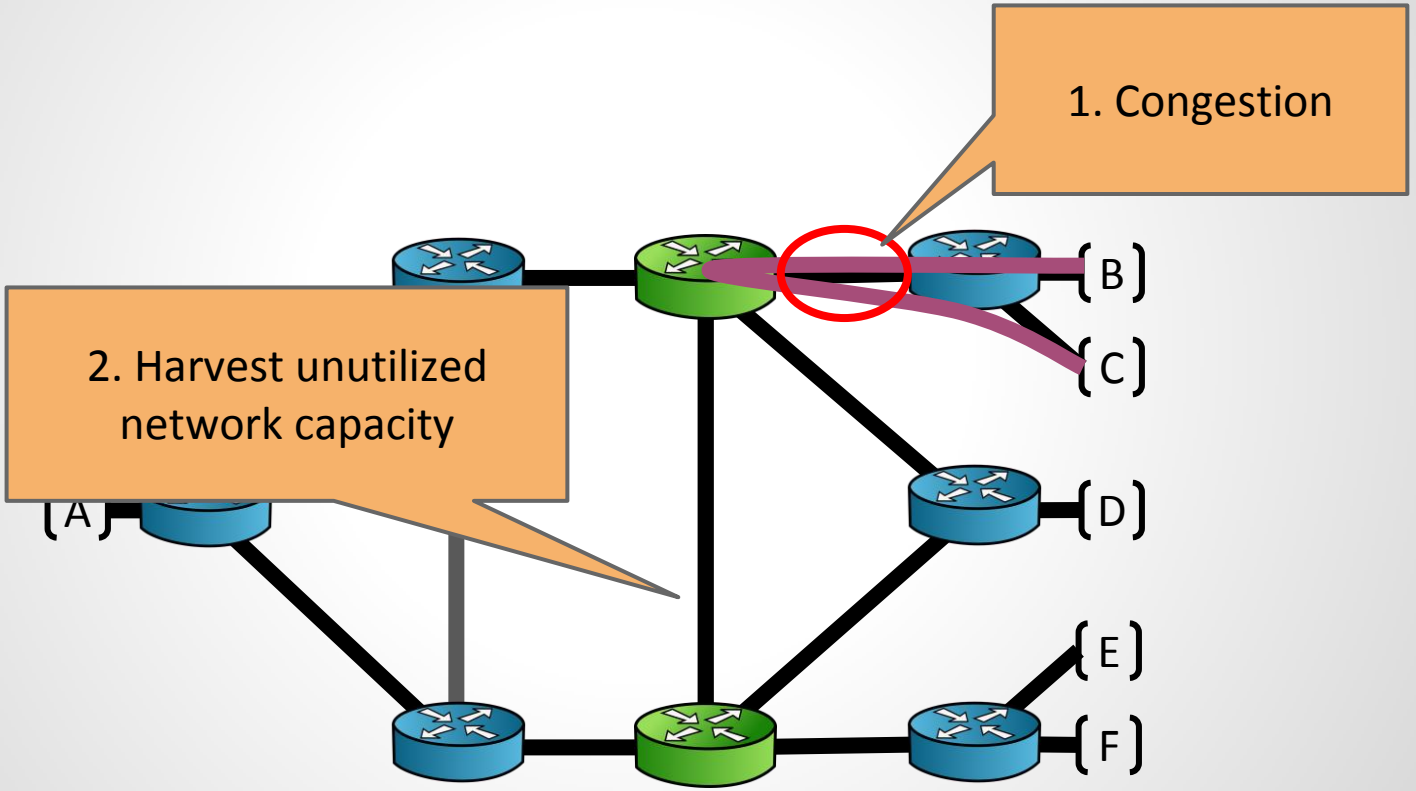

Network Emulation Experiments

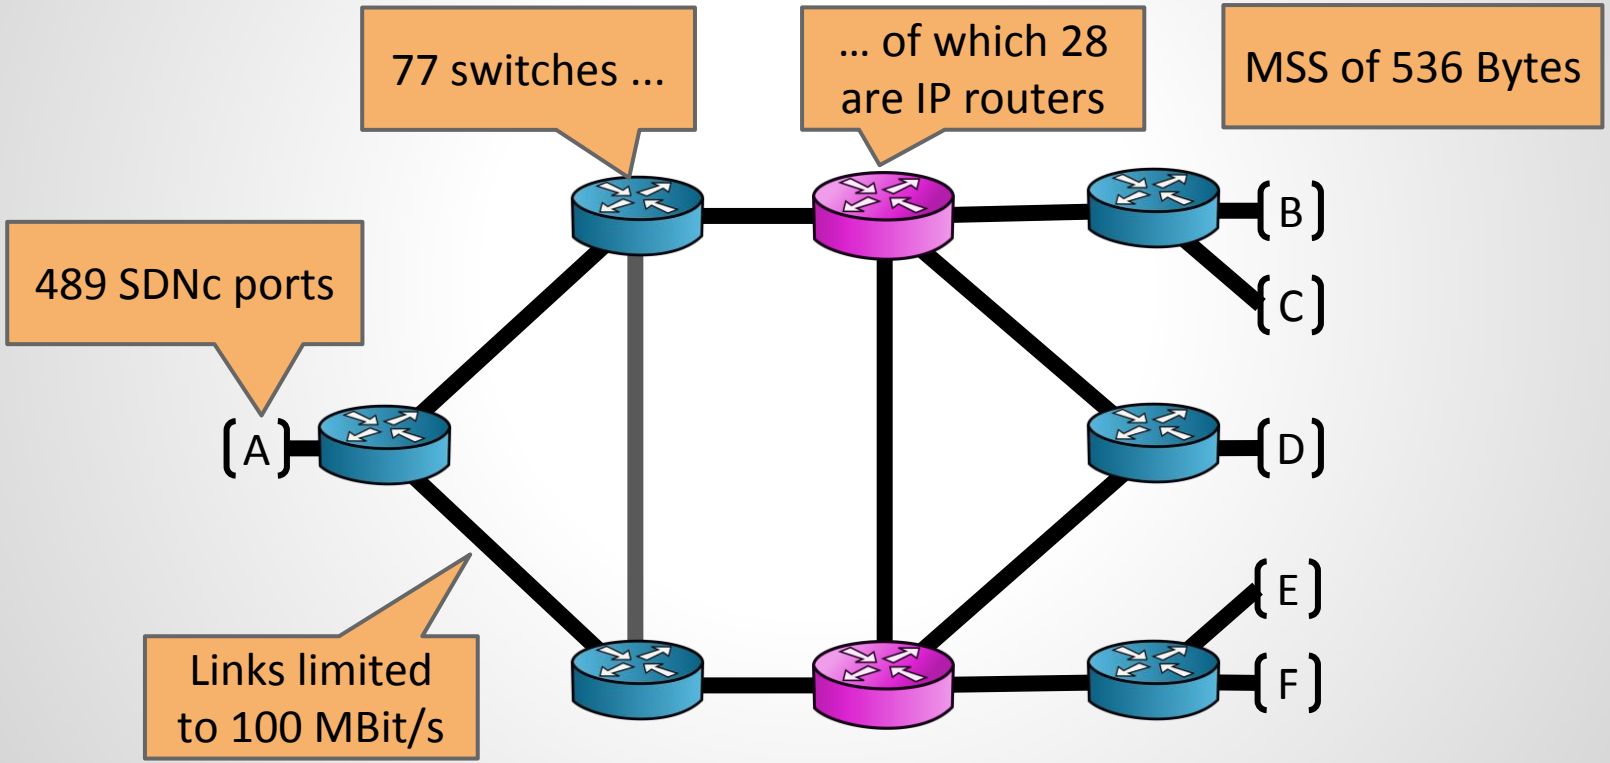

Workload

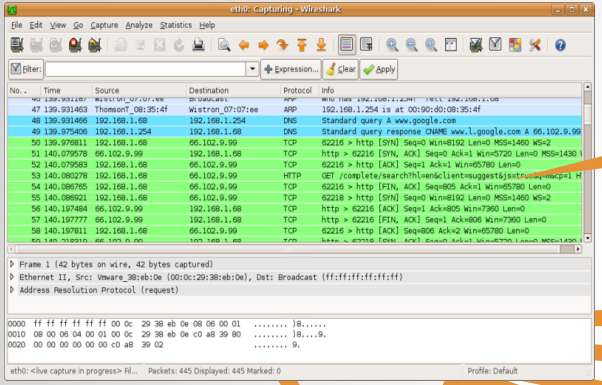

Packet-level enterprise network traffic traces

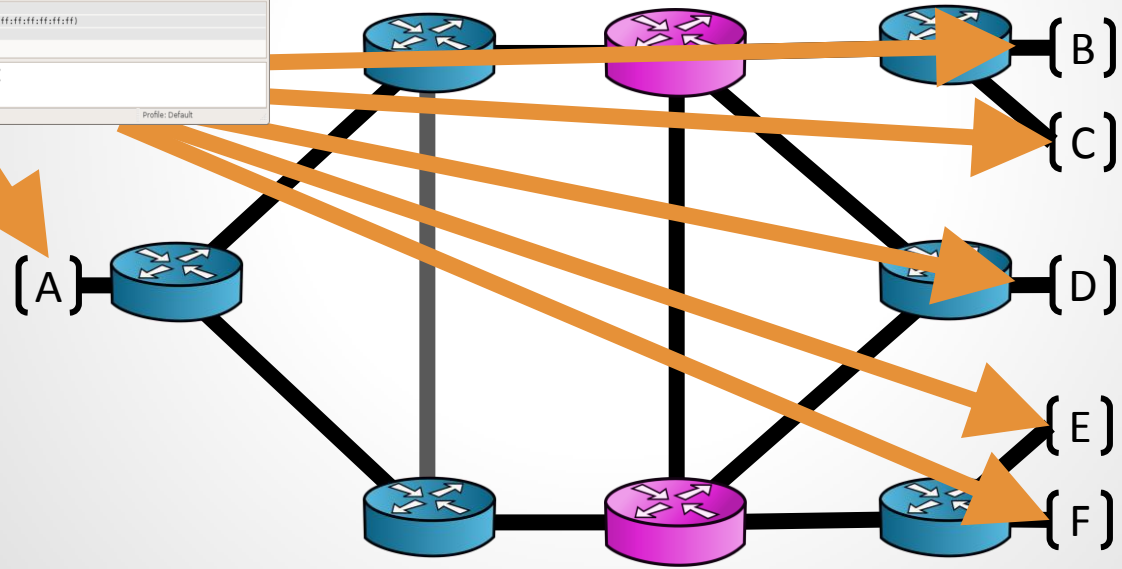

Defines a traffic matrix

Workload

Traffic Matrix: 1024 source-destination communication pairs

Sequentially initiate TCP flows until 100 MB are sent from every Port

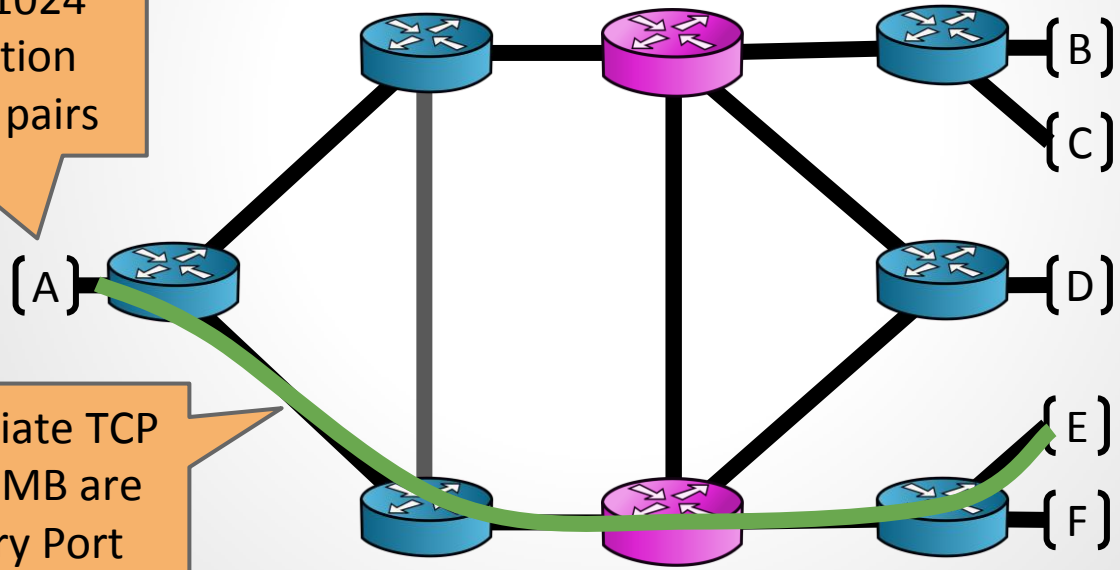

TCP Throughput Comparison

Ethernet switch

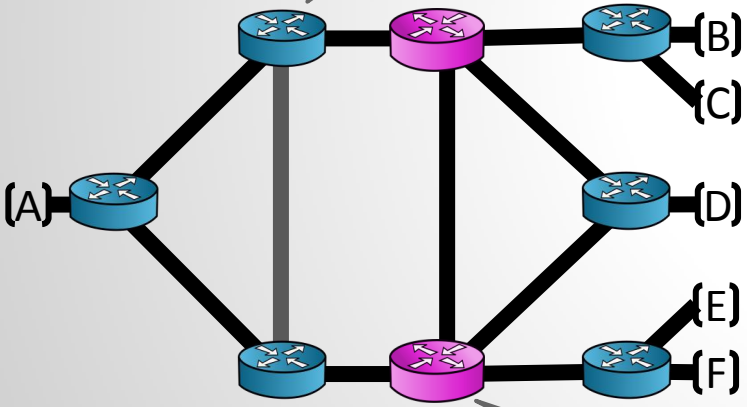

IP router

Legacy

VS.

Vary the number of deployed SDN switches:
5, 10, and 28

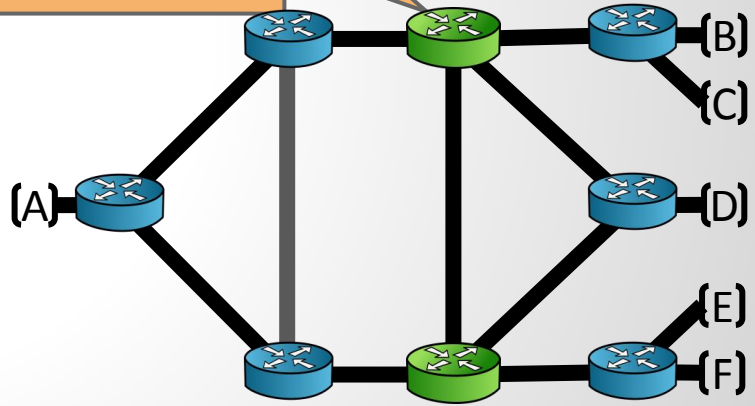

Panopticon

SDN Deployment Size vs. Performance

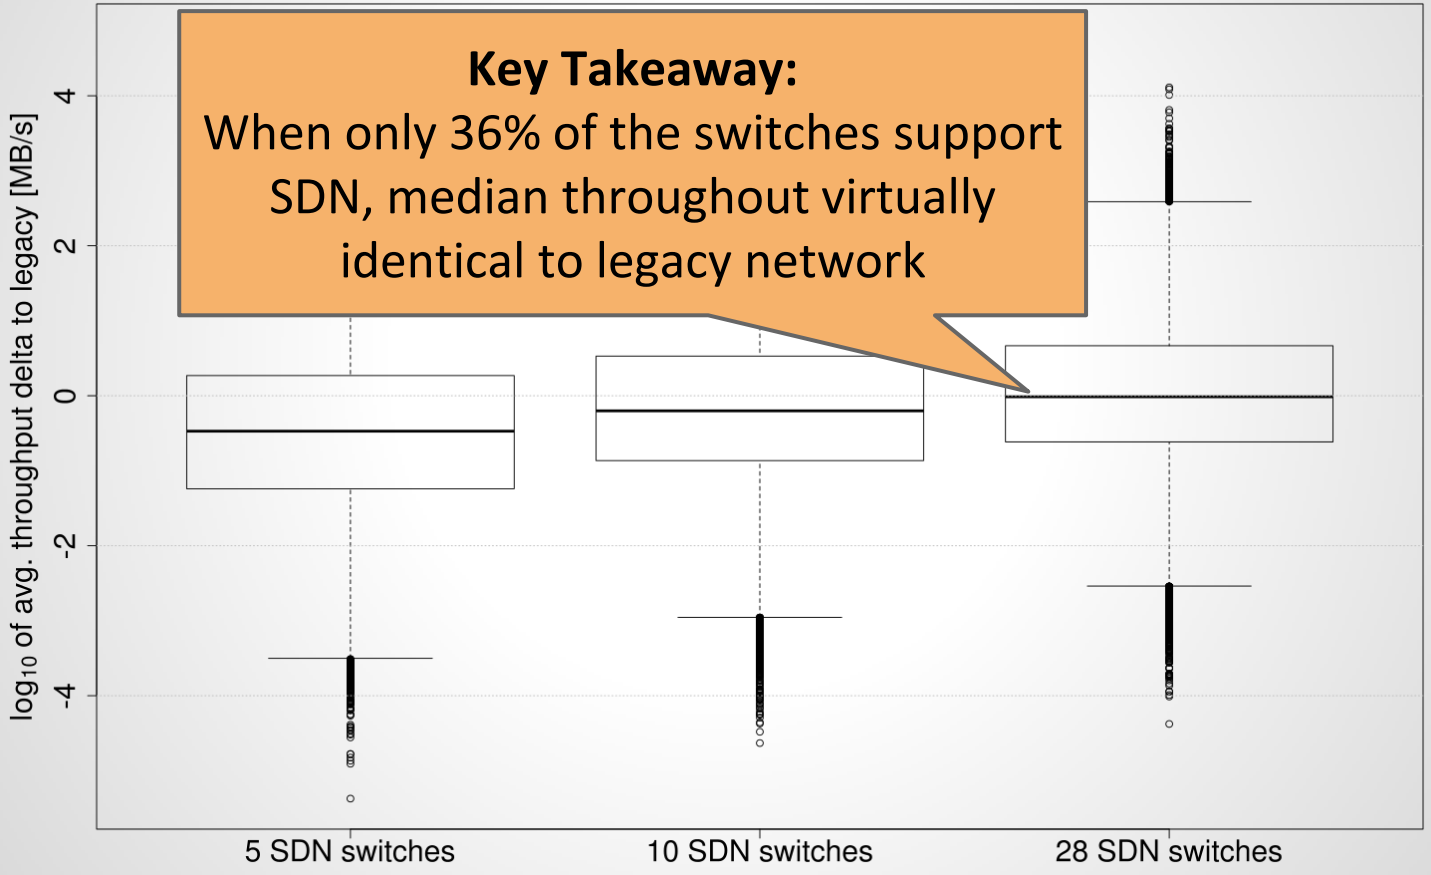

Key Results Highlights

- Evaluated a large campus network (1500+ switches)
- Real topologies and real traffic traces

- Upgrade **2%** of the switches/routers →
  - 100% SDN-controlled ingress ports
  - avg. path stretch < 50%
  - 90th percentile link util. < 25% increase

Summary

SDN ARCHITECTURE
Operate the network as
a (nearly) full SDN

TOOL
Determine the partial
SDN deployment

https://venture.badpacket.in

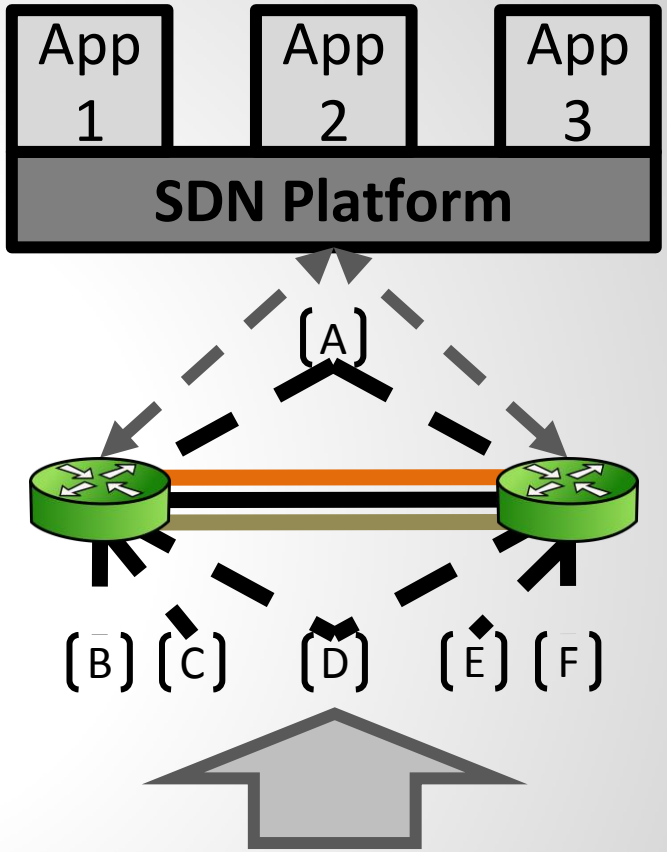

PANOPTICON

Median values of 10 random TM permutations

1296 SDNc Ports

10 policy rules per port (avg)

1024 VLANs

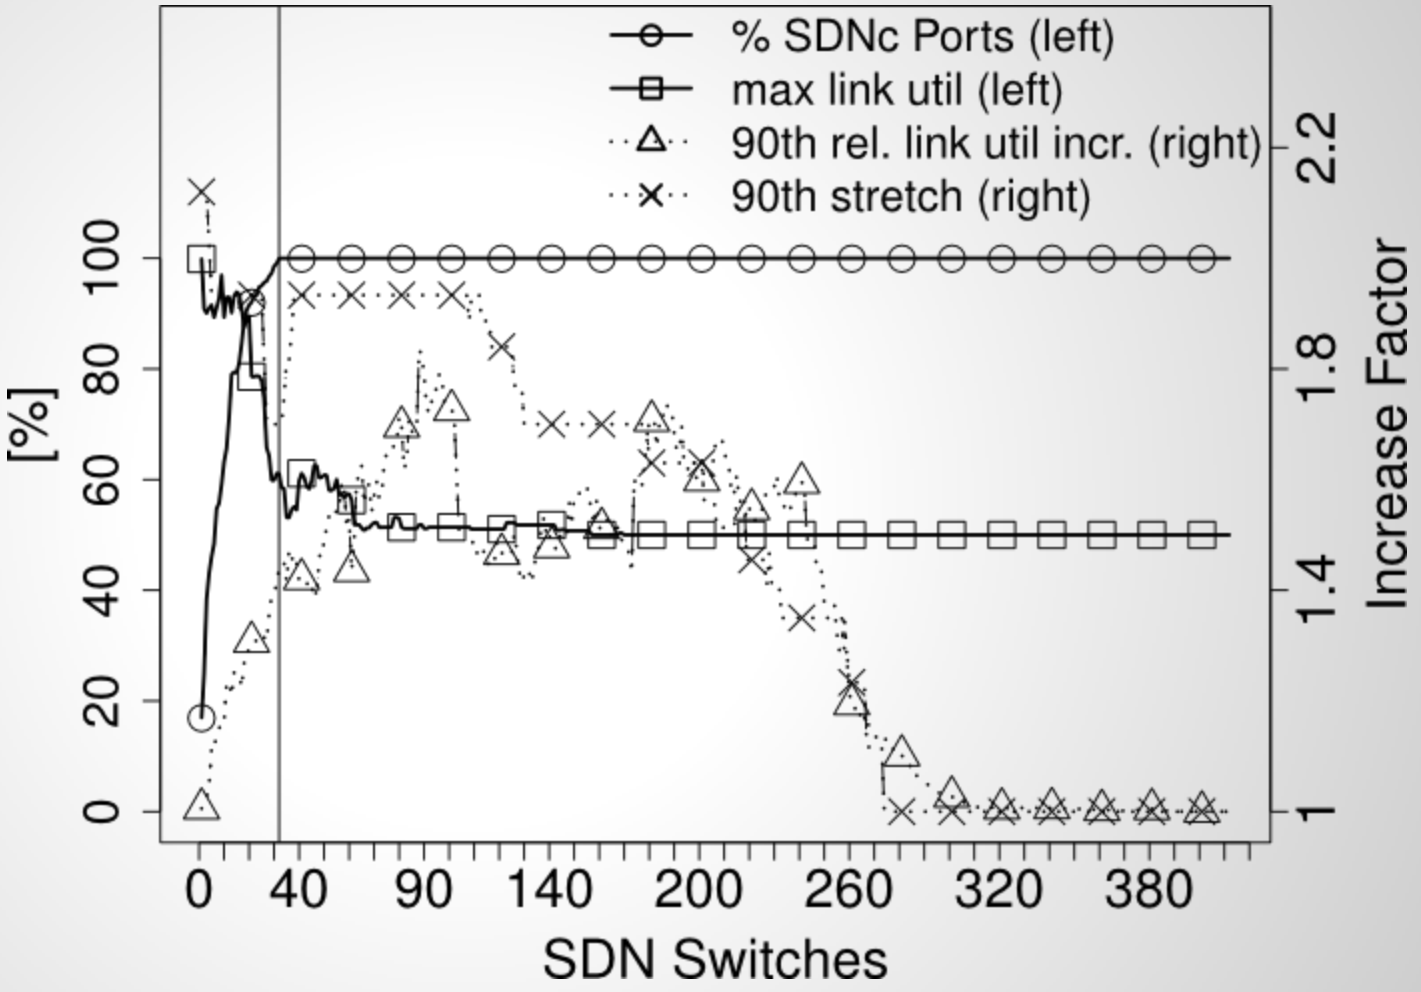

Other Ongoing Projects

- STN: Concurrent Control Plane
- Provably Dataplane Connectivity
- Consistent Network Updates
- AeroFlux: Wifi SDN with near-sighted control