SDN in Wide-Area Networks

4th IEEE Intl. Conference on Software-Defined Systems (SDS 2017)
May 8-10, 2017, Valencia, Spain

Oliver Michel

University of Colorado Boulder

Global IP Traffic Growth

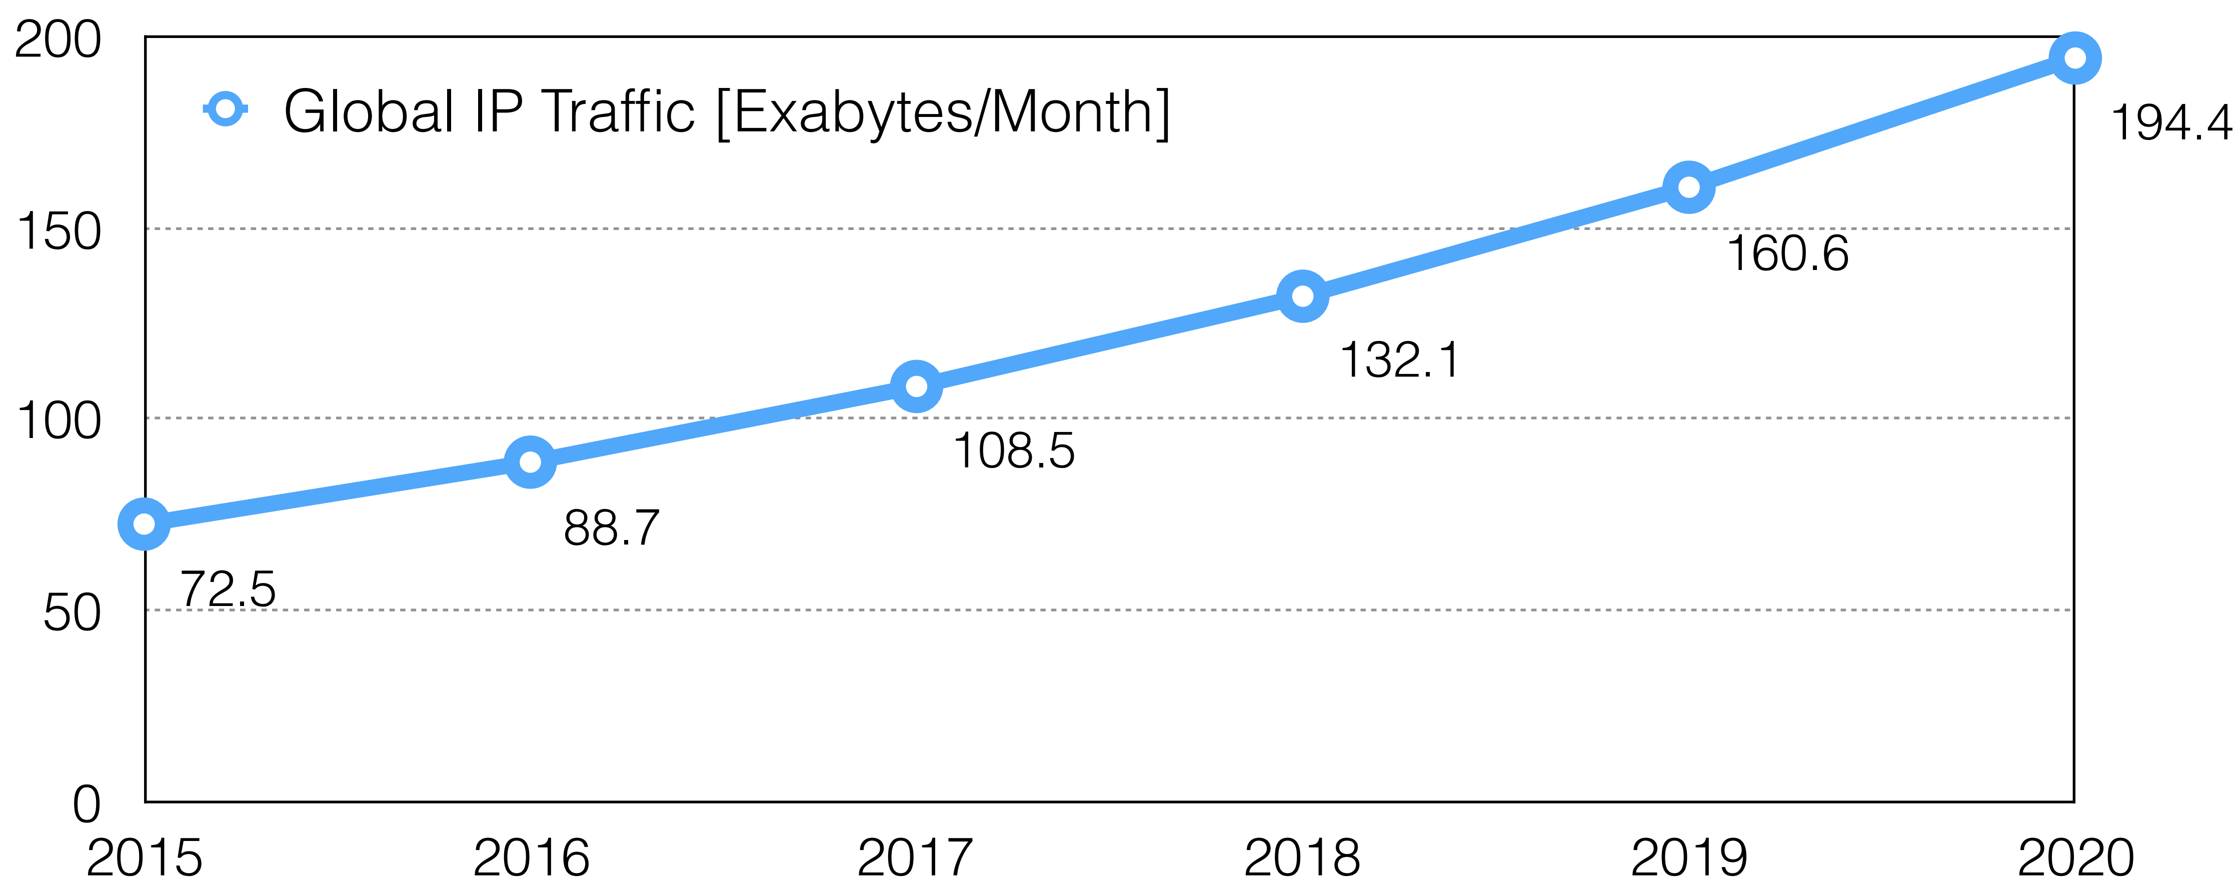

[Cisco® Visual Networking Index (VNI) 2016]

Global IP Traffic Growth

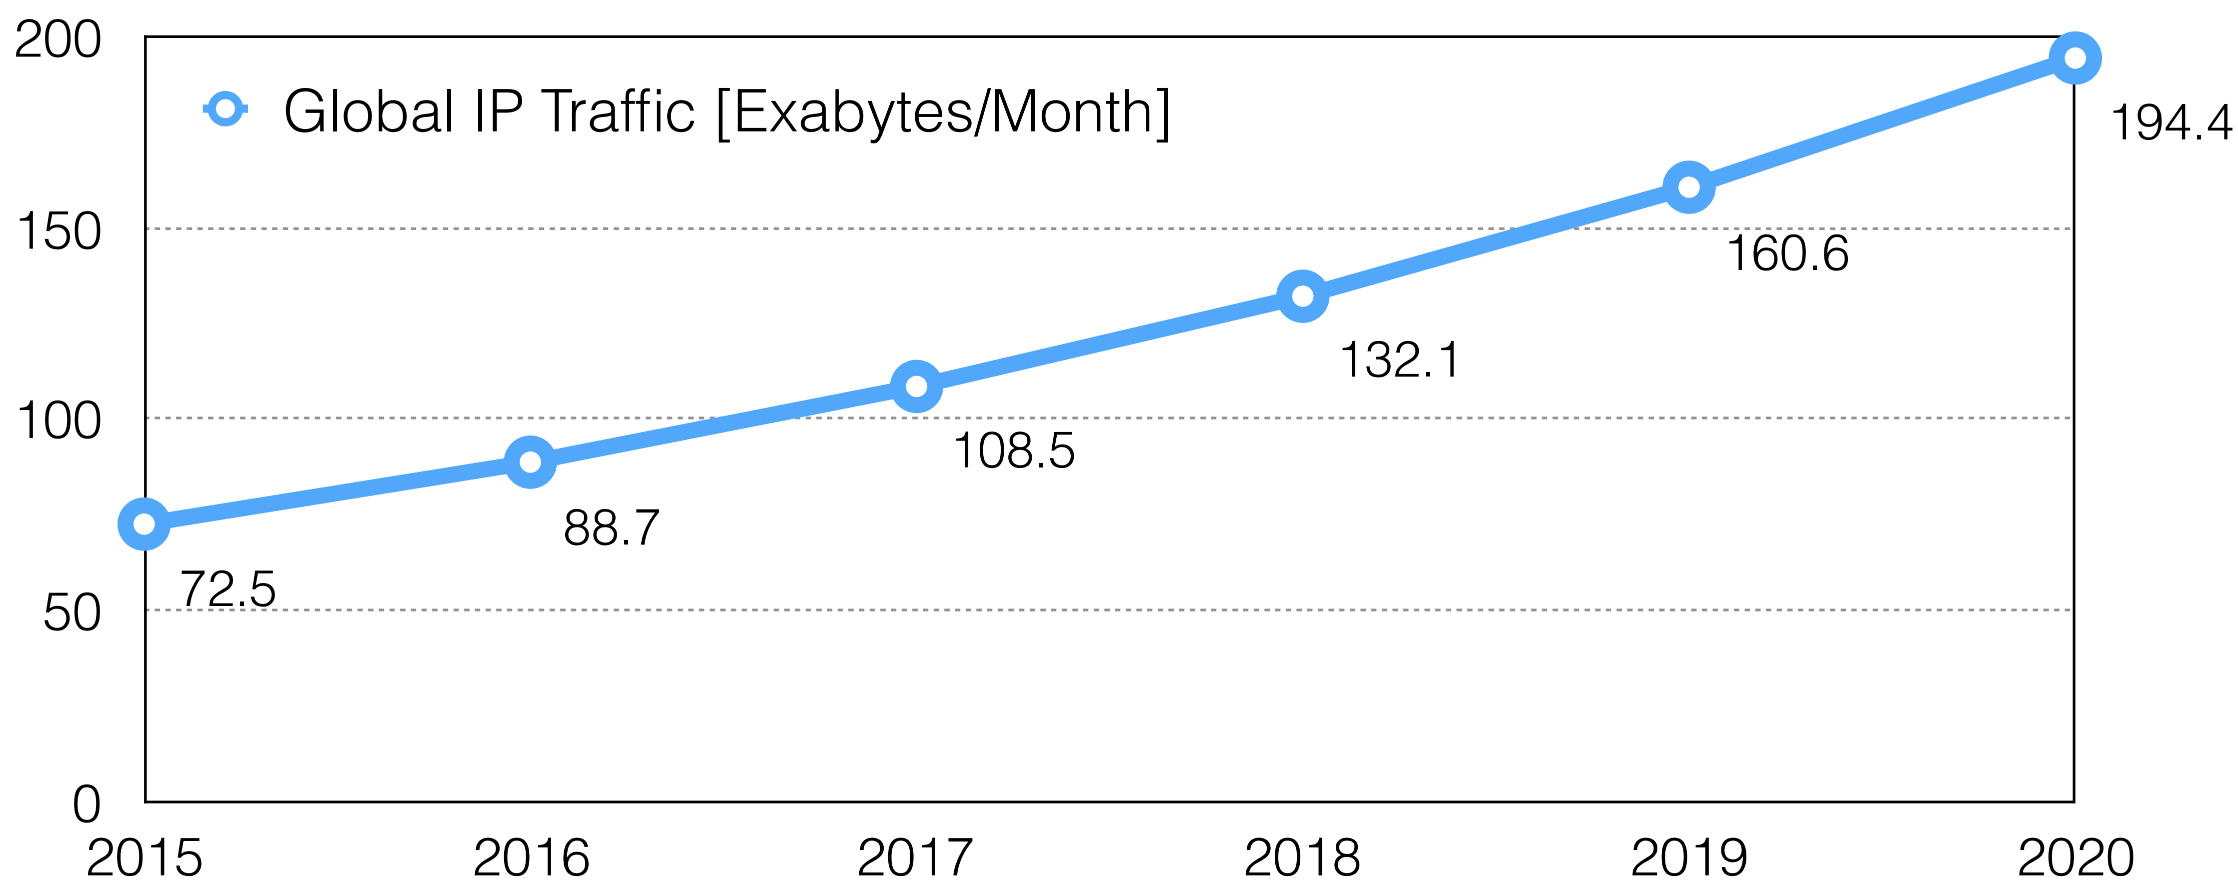

52% → 66%

[Cisco® Visual Networking Index (VNI) 2016]

Global IP Traffic Growth

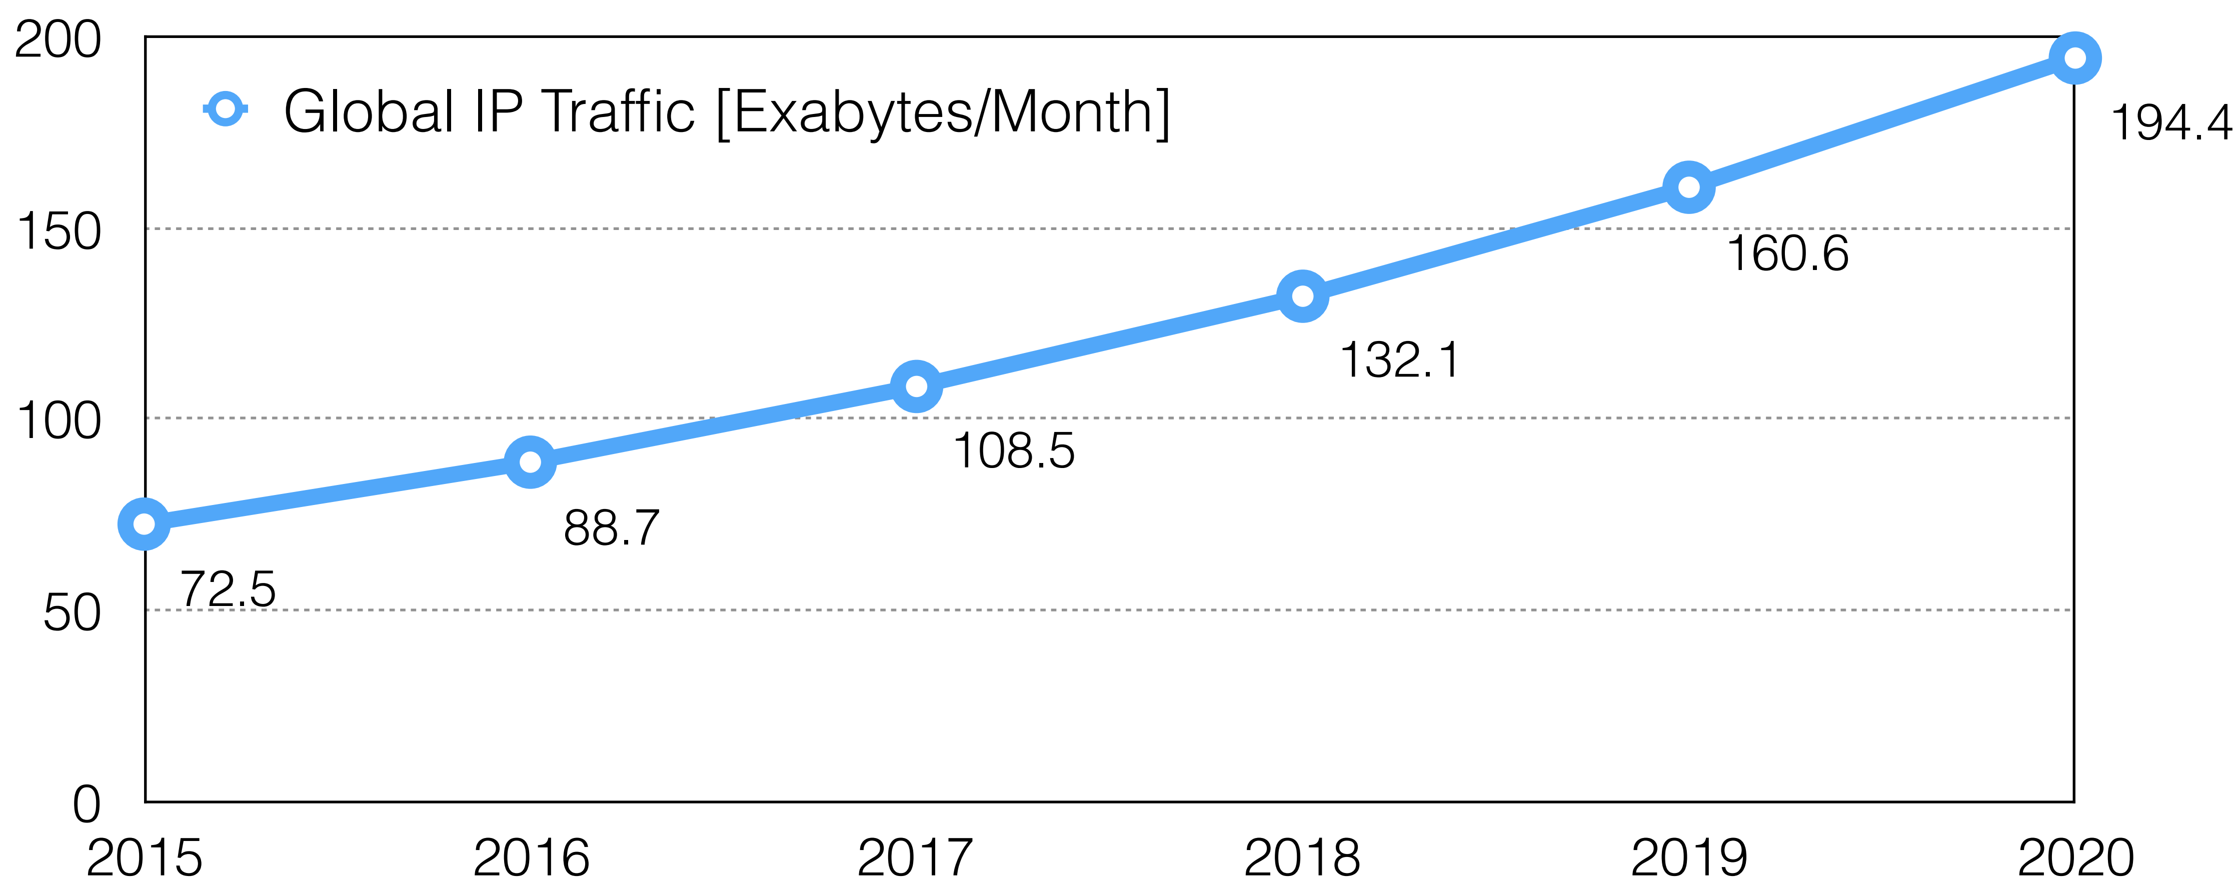

52% → 66%

70% → 82%

[Cisco® Visual Networking Index (VNI) 2016]

SDN in one Slide

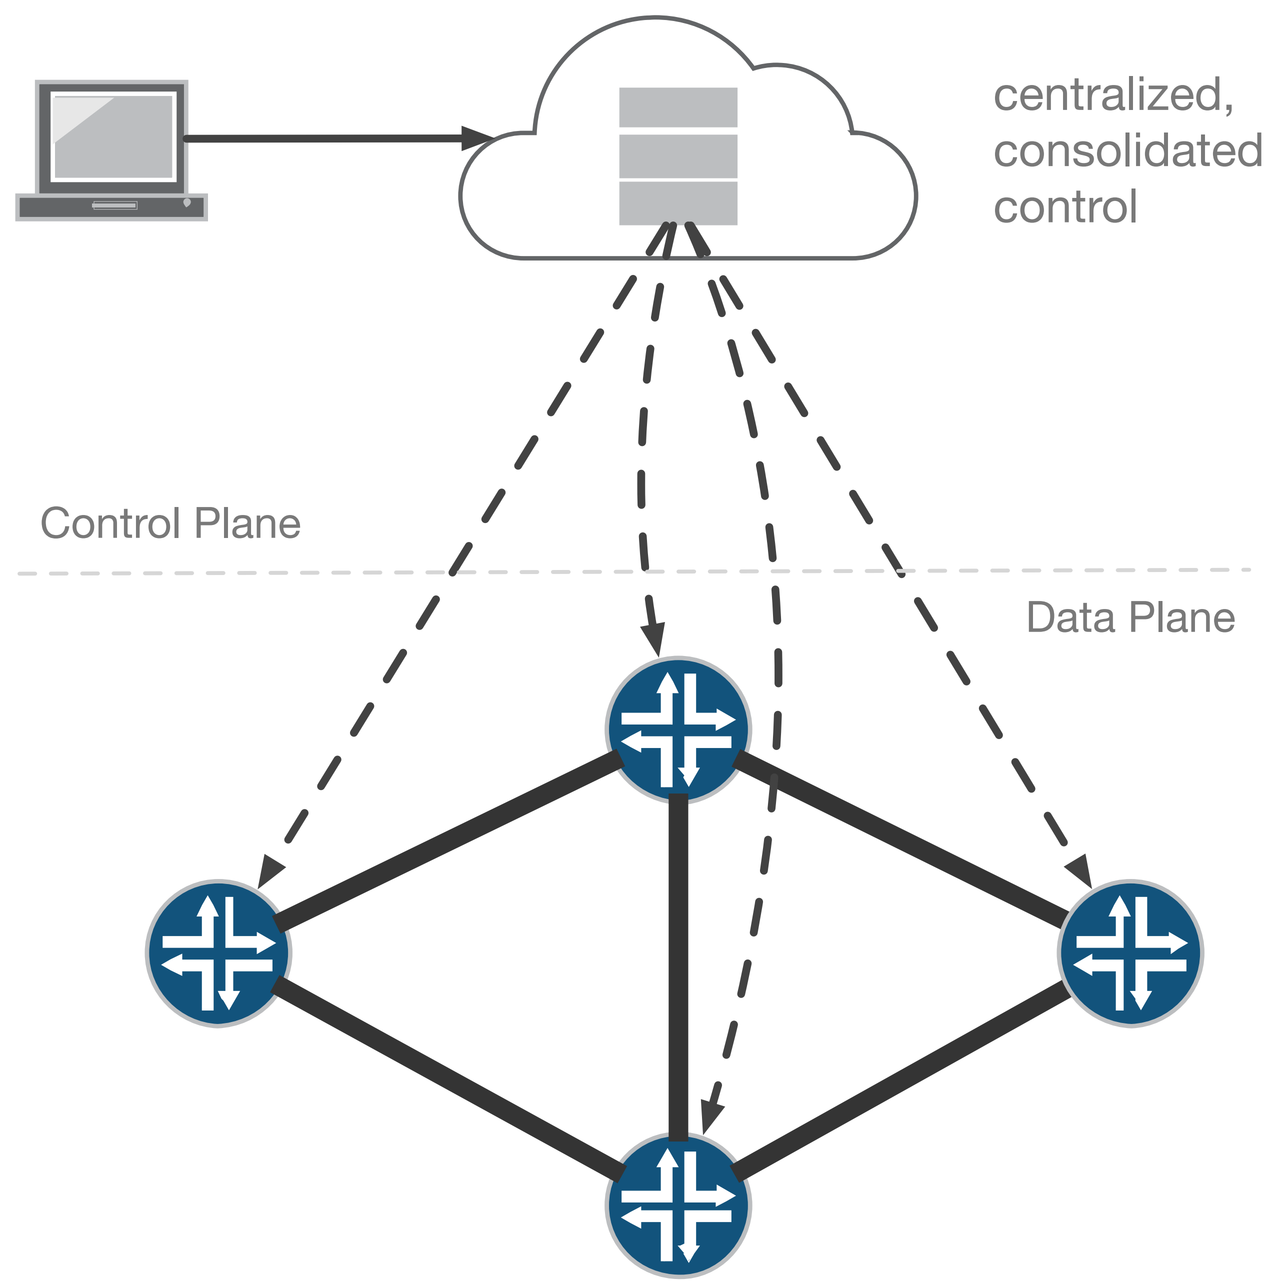

SDN Evolution

SDN Evolution

1. Active Networks
   [Tennenhouse. A survey of active network research. IEEE Comm. '97]

SDN Evolution

1. Active Networks
   [Tennenhouse. A survey of active network research. IEEE Comm. ’97]

2. Data Plane / Control Plane Separation
   [Casado. Ethane. SIGCOMM ’07, Greenberg. 4D. SIGCOMM CCR ’05, Caesar. RCP. ‘05]

SDN Evolution

1. Active Networks
   [Tennenhouse. A survey of active network research. IEEE Comm. ’97]

2. Data Plane / Control Plane Separation
   [Casado. Ethane. SIGCOMM ’07, Greenberg. 4D. SIGCOMM CCR ’05, Caesar. RCP. ‘05]

3. Control Protocols
   [McKeown. OpenFlow. SIGCOMM CCR ’08]

SDN in the Data Center

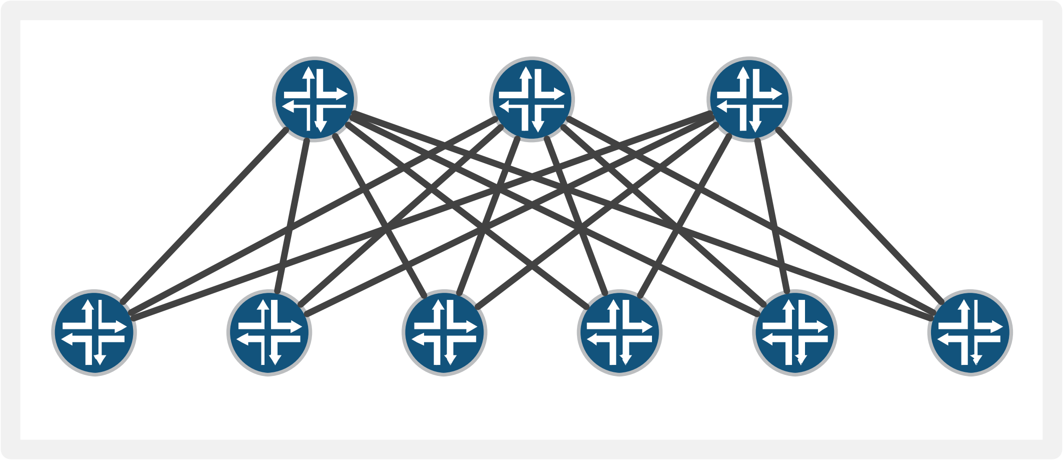

1. Network Virtualization
   [Koponen. Network Virtualization in Multi-Tenant Data Centers. NSDI ’14, Keller. LIME. HotNets ’12, Sherwood. FlowVisor. OSDI ’10]

2. Resource Management
   [Ballani. Oktopus. SIGCOMM ’11, Guo. SecondNet. CoNEXT ‘10]

WAN Challenges

WAN Challenges

- controlled environment
- dedicated control networks
- fewer external factors

WAN Challenges

- controlled environment
- dedicated control networks
- fewer external factors

- fibers in ducts along highways/pipelines
- in-band control

WAN Challenges

- Legacy Equipment, Protocols and Practices
- Different Domains, Stakeholders
- Interoperability Requirements

Facebook
Wedge Platform

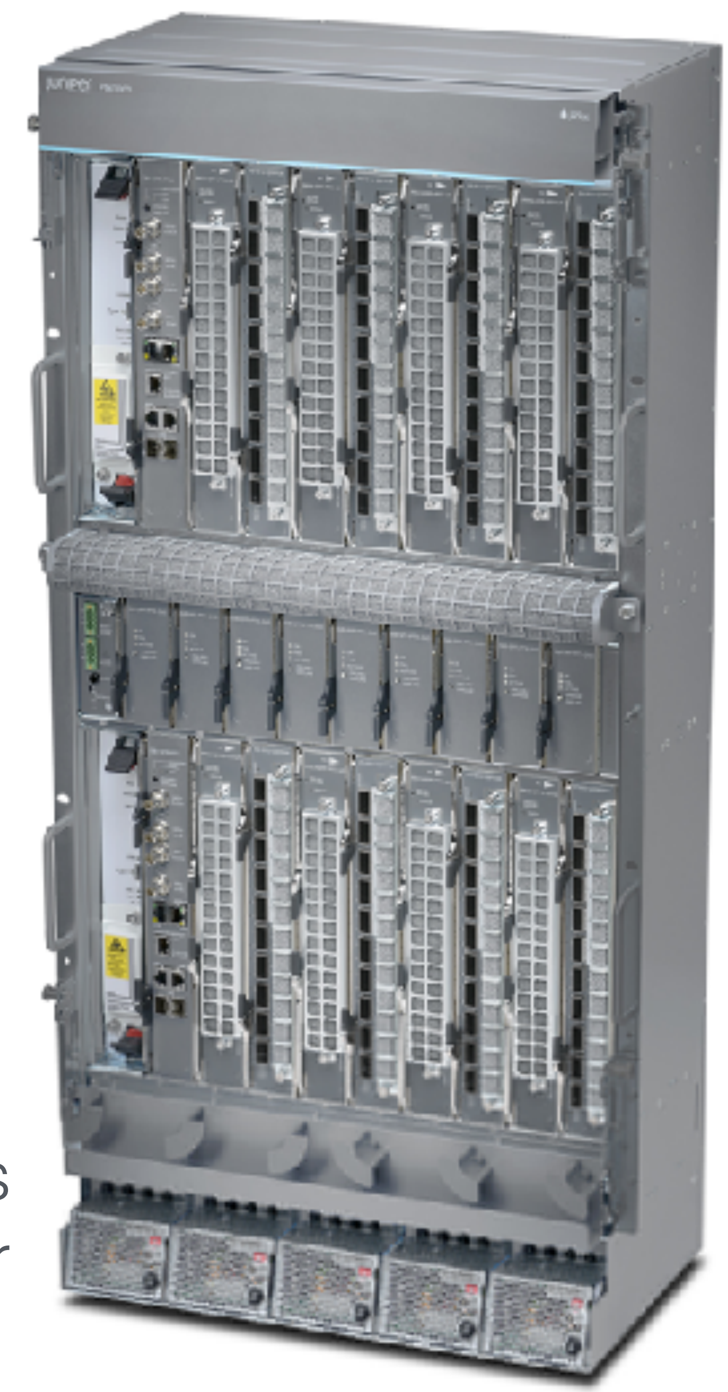

Juniper Networks
PTX3000 Core Router

WAN Opportunities

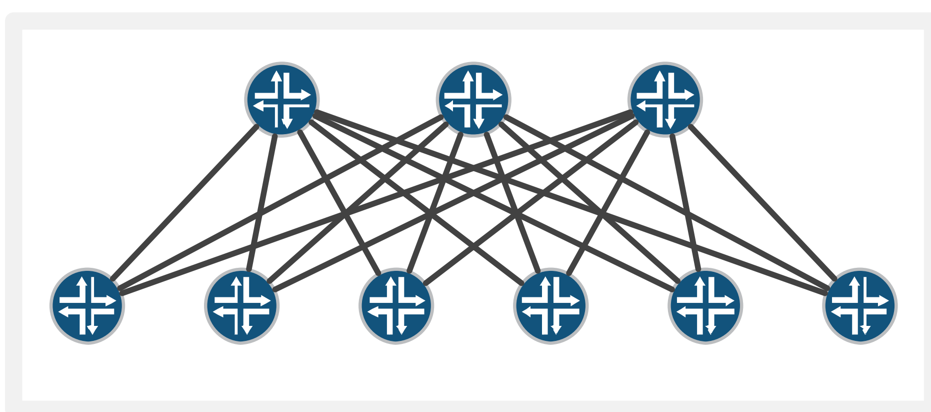

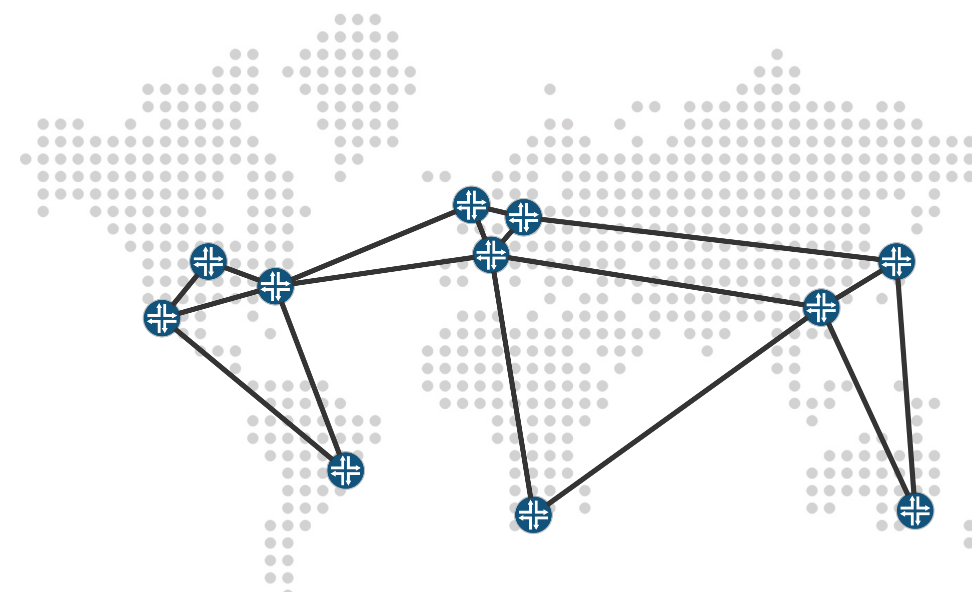

- tree-like networks with high degree of parallel links
- less expensive copper cabling

- mesh network with fewer parallel links
- expensive wide-area fibers and optics

WAN Opportunities

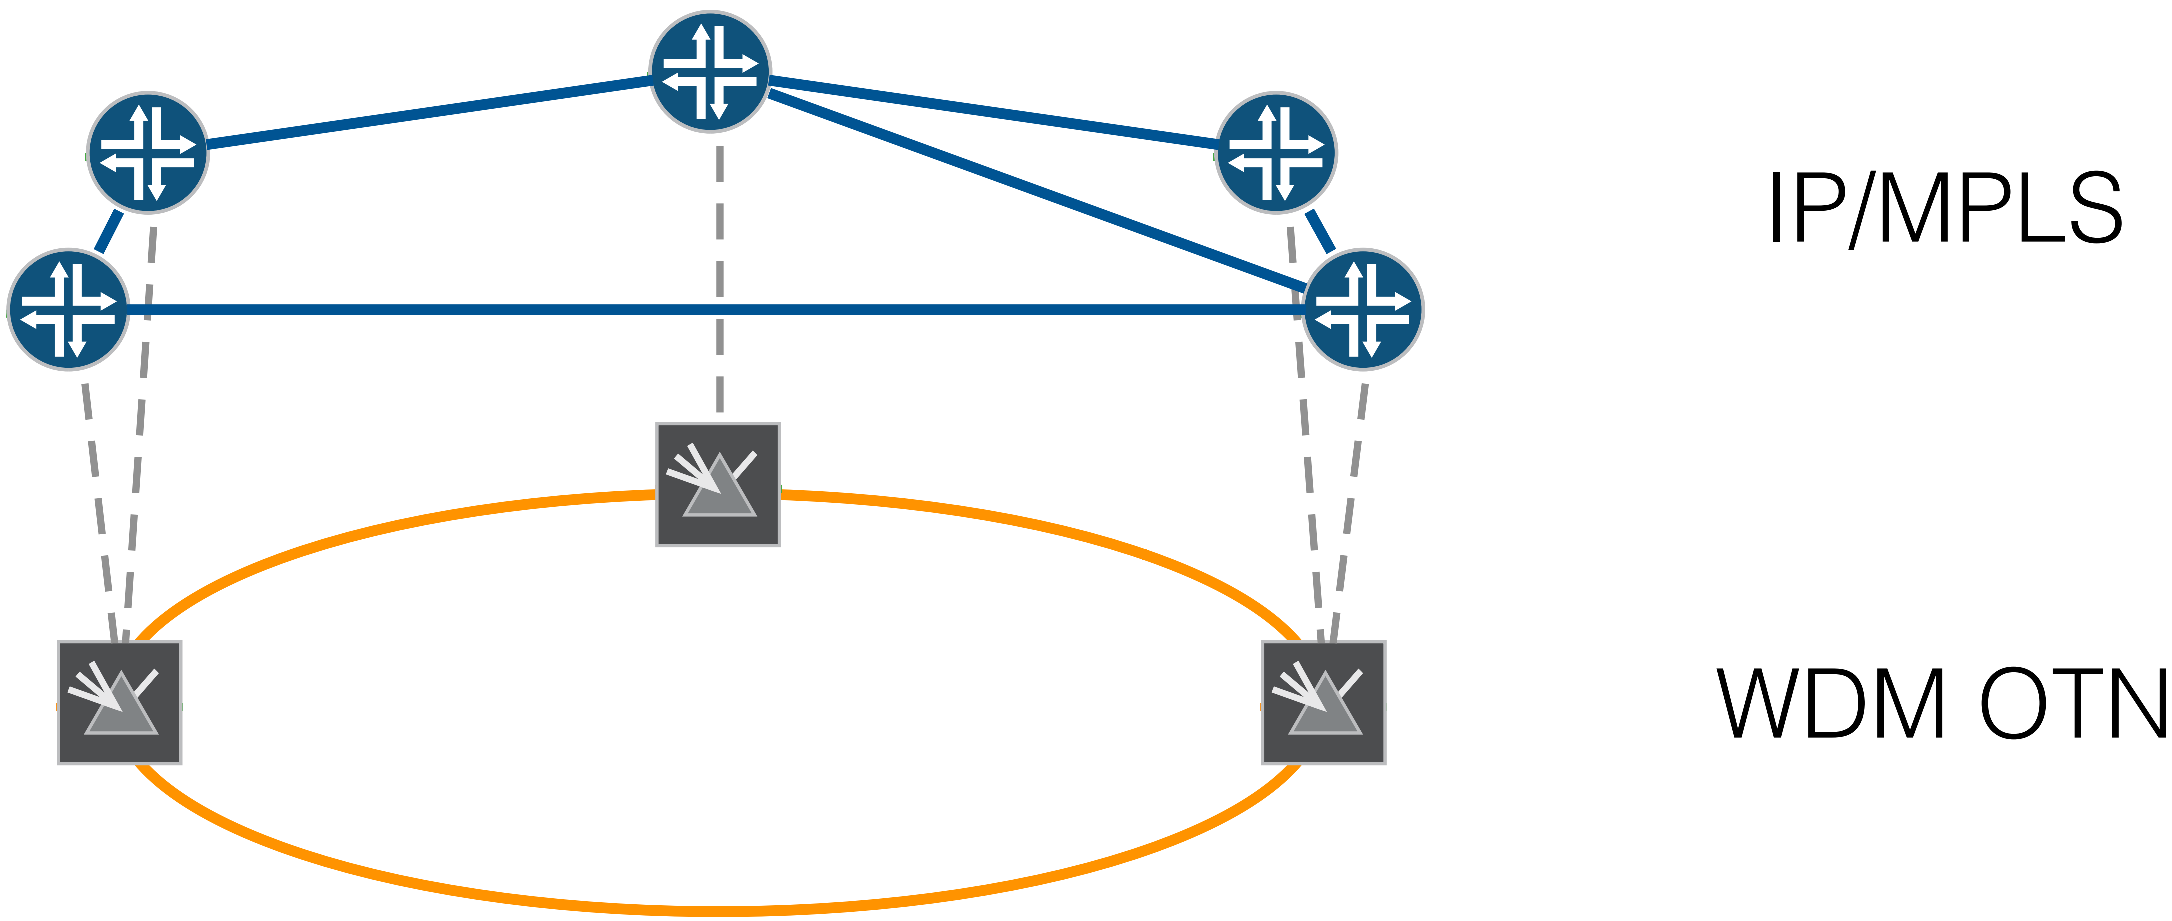

Challenges for SDN in Wide-Area Networks

Challenges for SDN in Wide-Area Networks

1. Distributing SDN Controller State
   [Yeganeh. Kandoo. HotSDN ’12, Dixit. ElastiCon. HotSDN ’12, Berde. ONOS. HotSDN ‘14]

Challenges for SDN in Wide-Area Networks

1. Distributing SDN Controller State
   [Yeganeh. Kandoo. HotSDN ’12, Dixit. ElastiCon. HotSDN ’12, Berde. ONOS. HotSDN ‘14]

2. Placing Controller Instances
   [Heller. Controller Placement. HotSDN ‘12]

Challenges for SDN in Wide-Area Networks

1. Distributing SDN Controller State
   [Yeganeh. Kandoo. HotSDN ’12, Dixit. ElastiCon. HotSDN ’12, Berde. ONOS. HotSDN ‘14]

2. Placing Controller Instances
   [Heller. Controller Placement. HotSDN ‘12]

3. Updating SDN Switches in a consistent Manner
   [Reitblatt. Consistent Updates. SIGCOMM ’12, Jin. Dionysus. SIGCOMM ‘14]

Distributing SDN Controller State

- Kandoo [Yeganeh HotSDN '12]
  - hierarchical model, reduces controller traffic

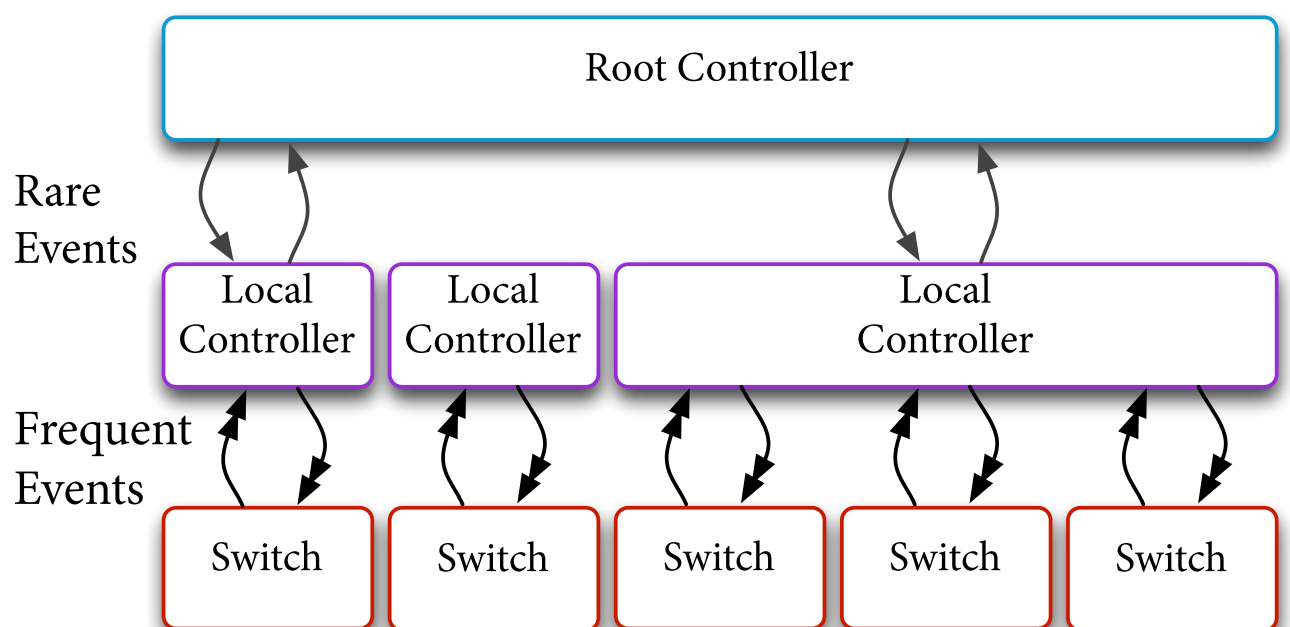

[Yeganeh HotSDN '12]

Distributing SDN Controller State

- Kandoo [Yeganeh HotSDN '12]
    - hierarchical model, reduces controller traffic
- ONOS [Berde. ONOS. HotSDN '14]
    - distributed, eventually consistent network graph through Cassandra backend

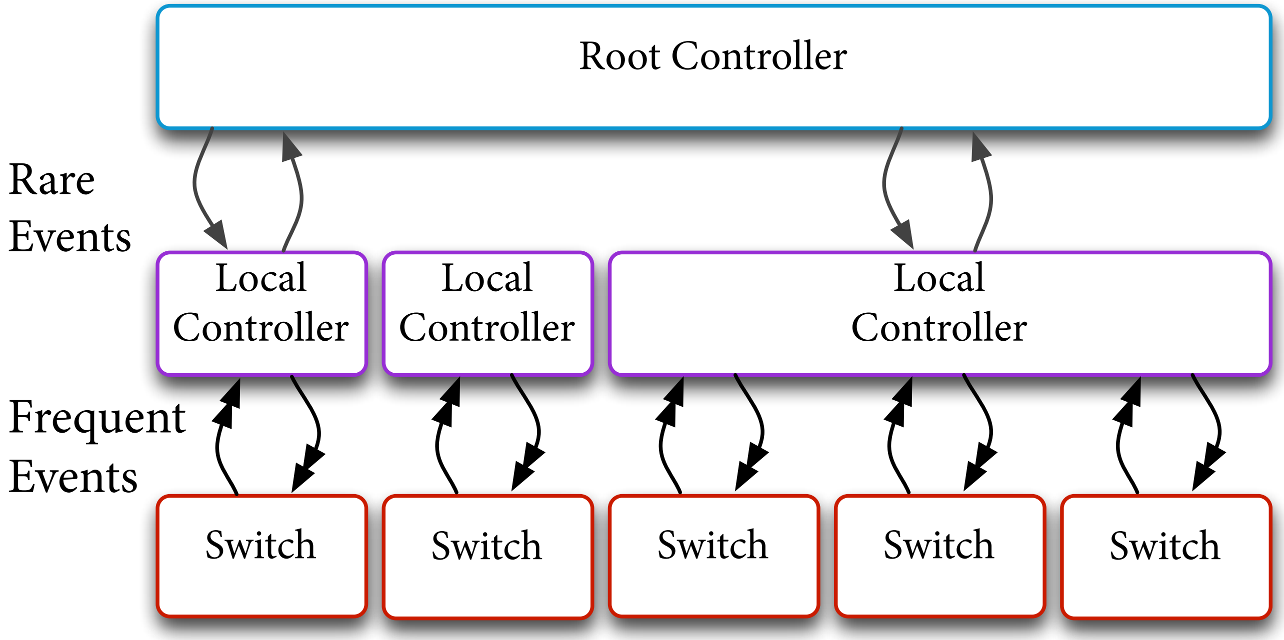

[Yeganeh HotSDN '12]

Placing SDN Controller Instances

The Controller Placement Problem
[Heller ’12]

- 3 fundamental underlying problems
  1. average-case latency
     minimum k-median
  2. worst-case latency
     minimum k-center
  3. nodes within latency bound
     maximum cover
- cost/benefit analysis: single or pair of controllers often enough

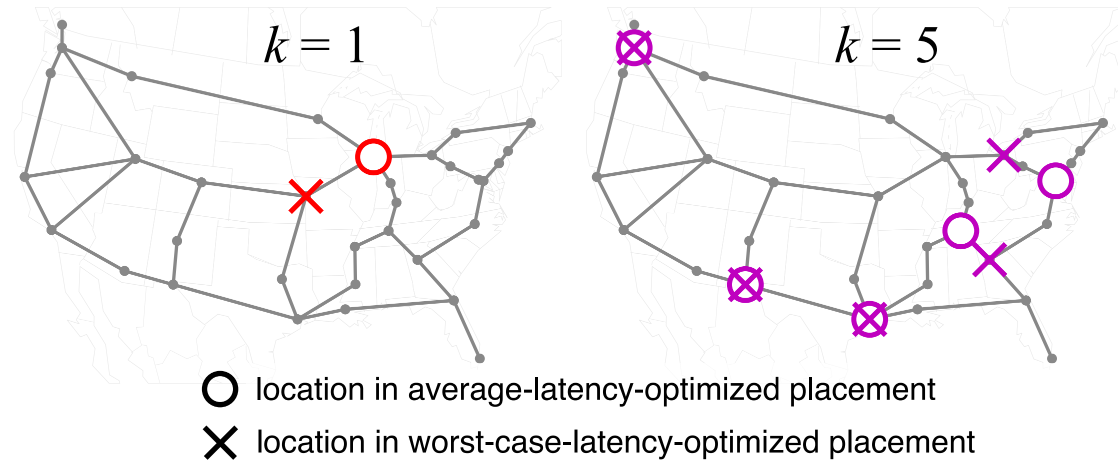

[Heller ’12]

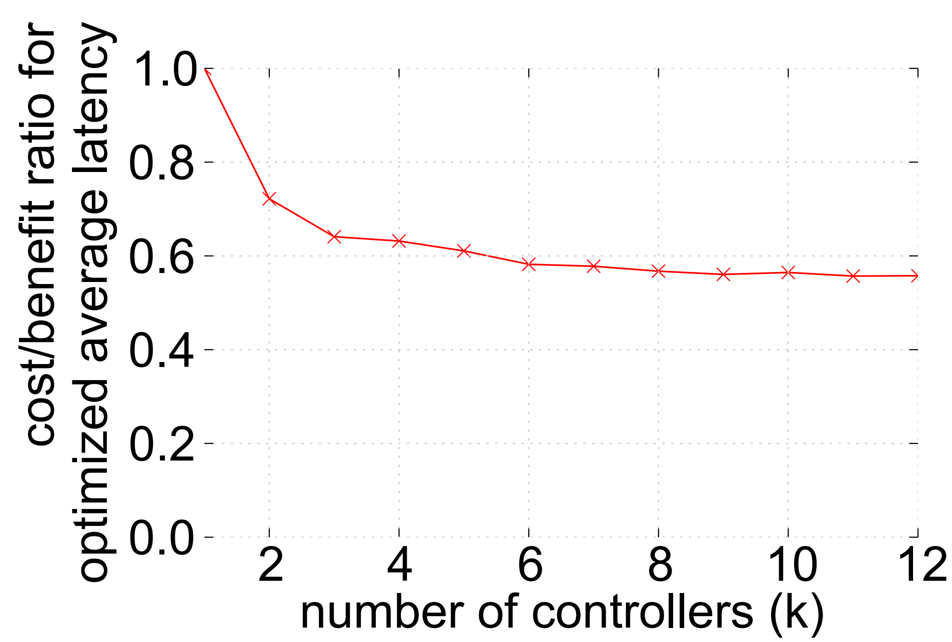

[Heller ’12]

Consistent Data Plane Updates

Problem 1

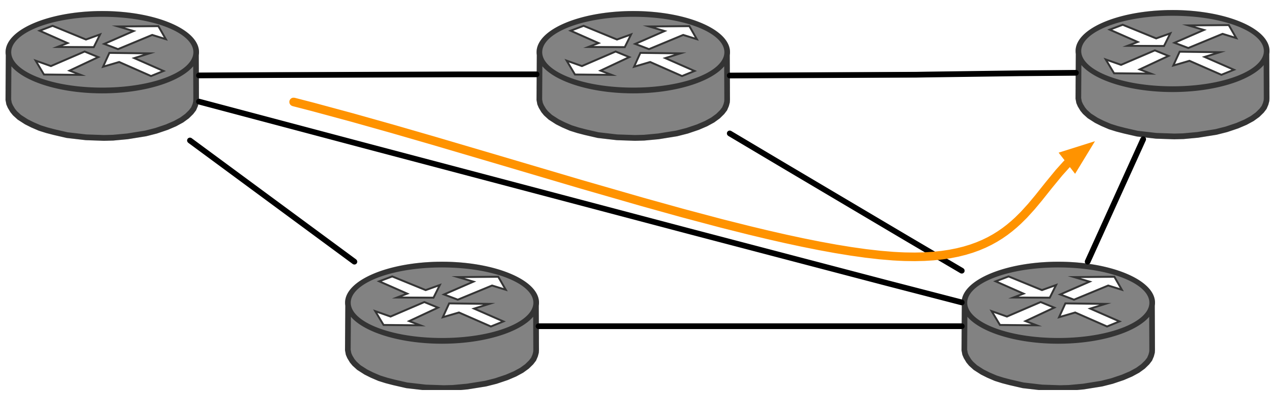

current state

Consistent Data Plane Updates

Problem 1

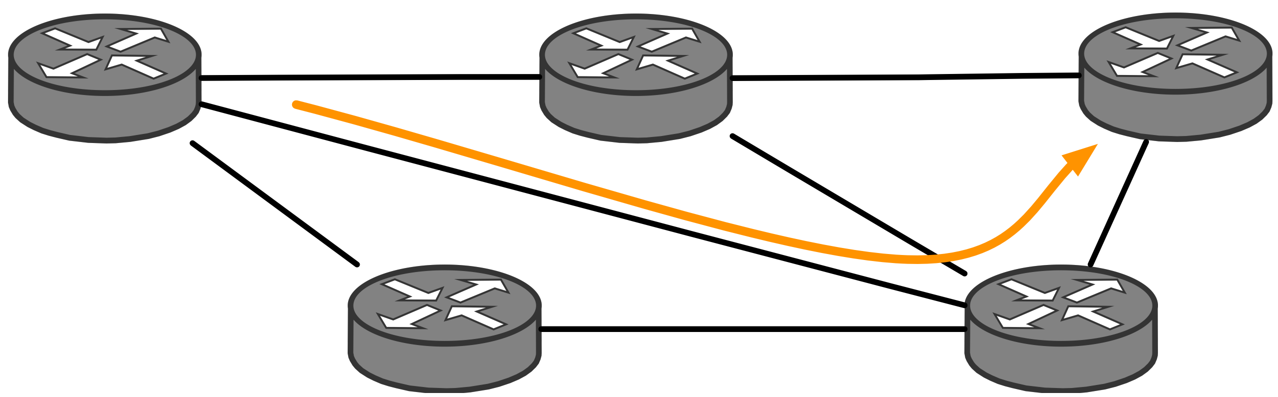

current state

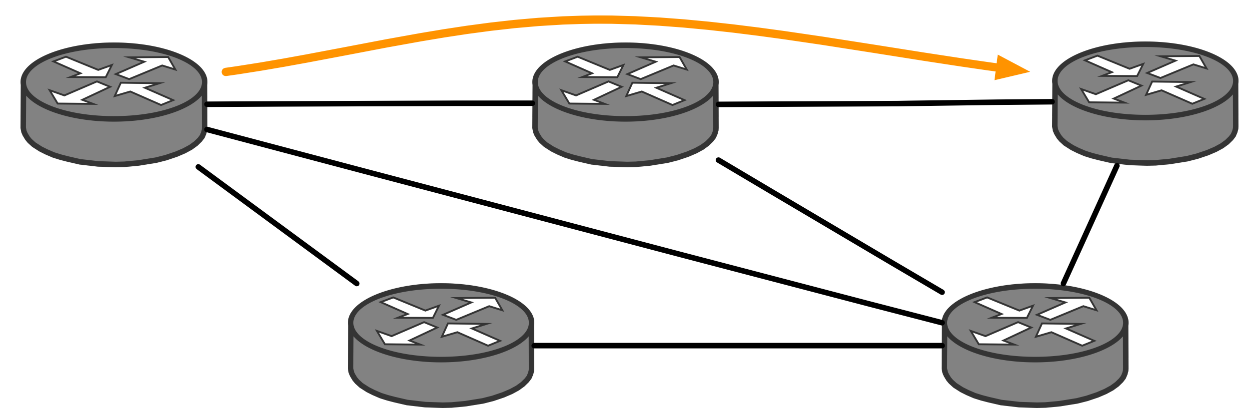

target state

Consistent Data Plane Updates

Problem 1

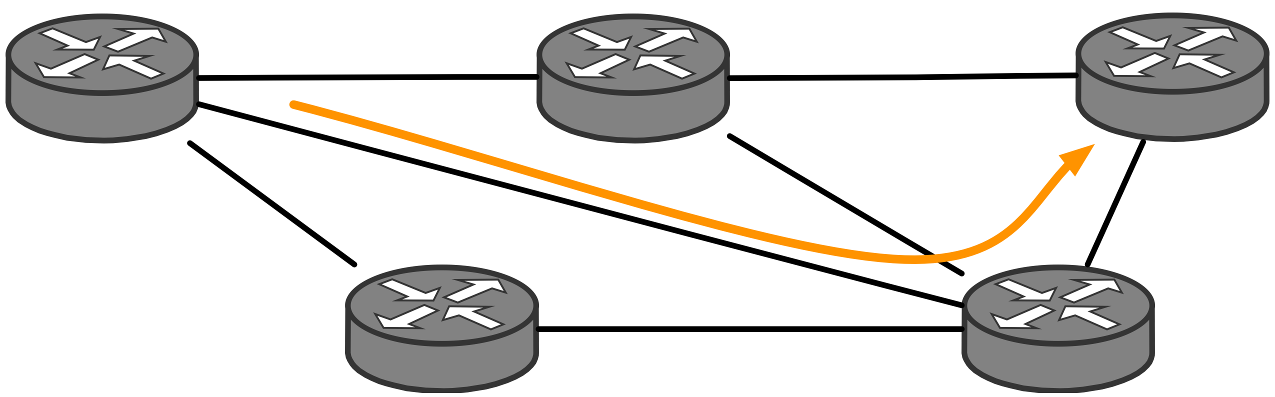

current state

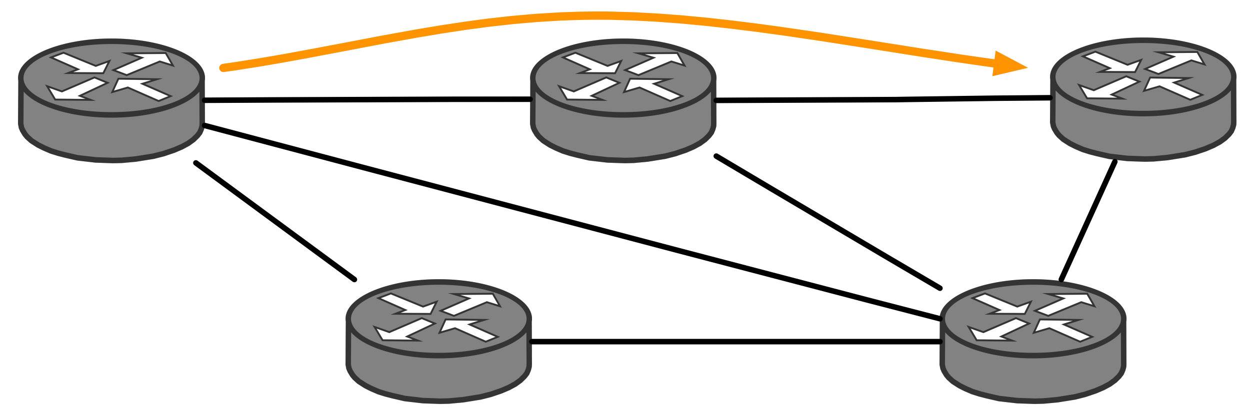

target state

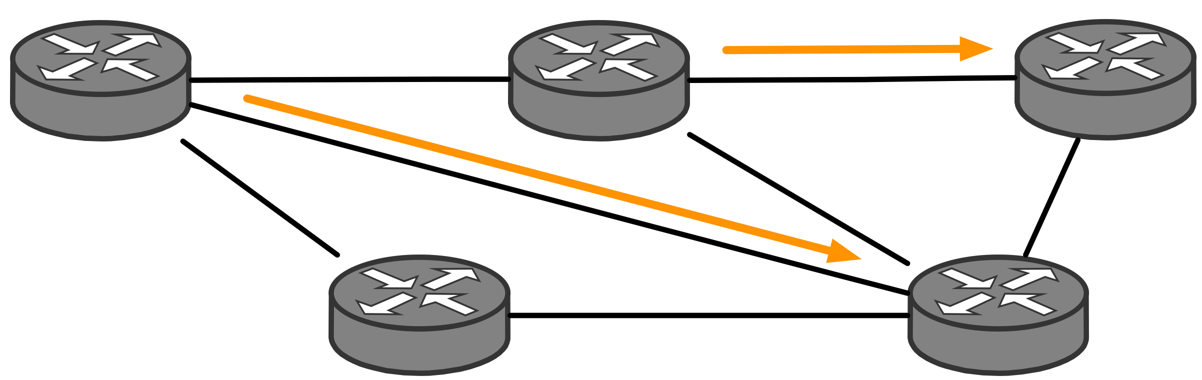

possible intermediate state

Consistent Data Plane Updates

Problem 2

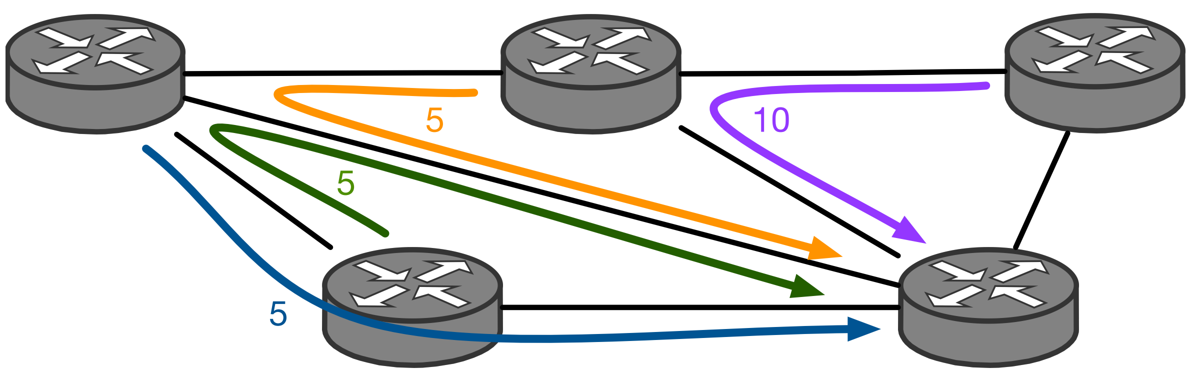

current state

Consistent Data Plane Updates

Problem 2

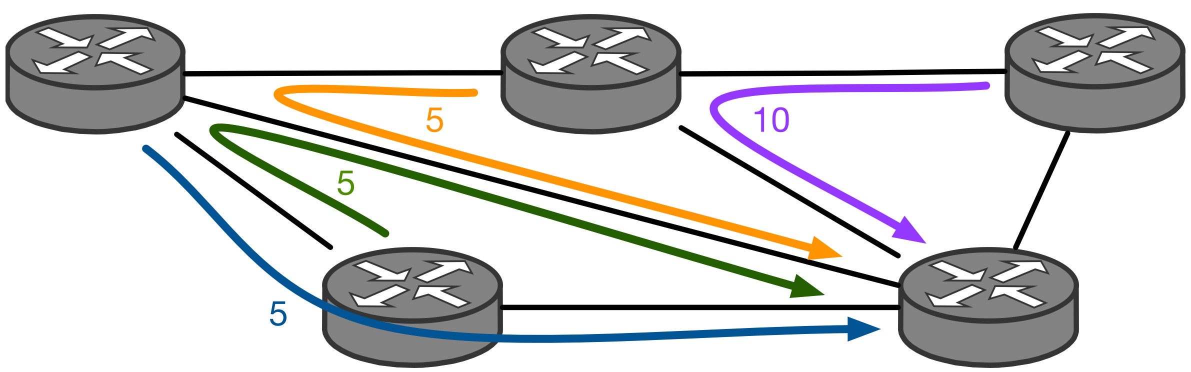

current state

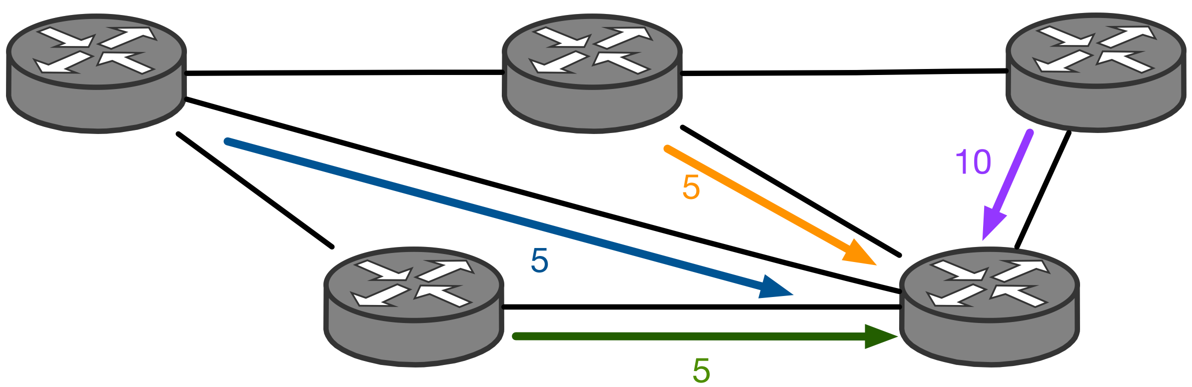

target state

Consistent Data Plane Updates

Problem 2

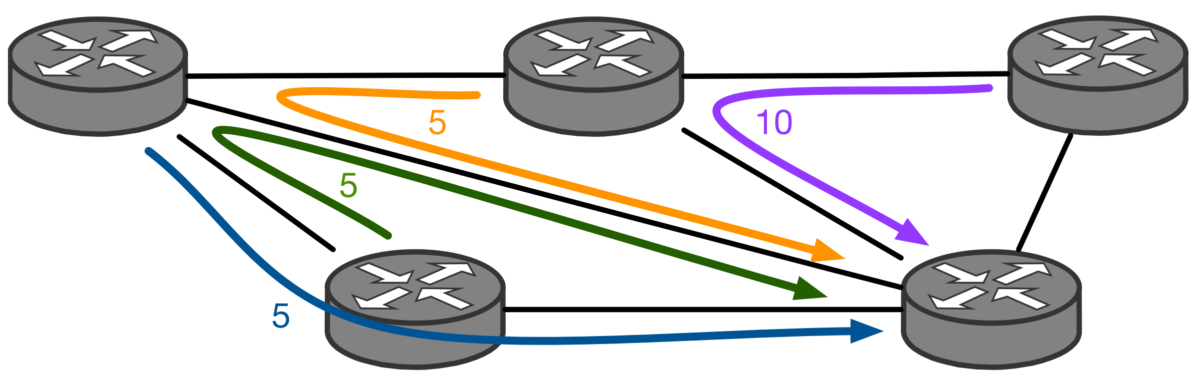

current state

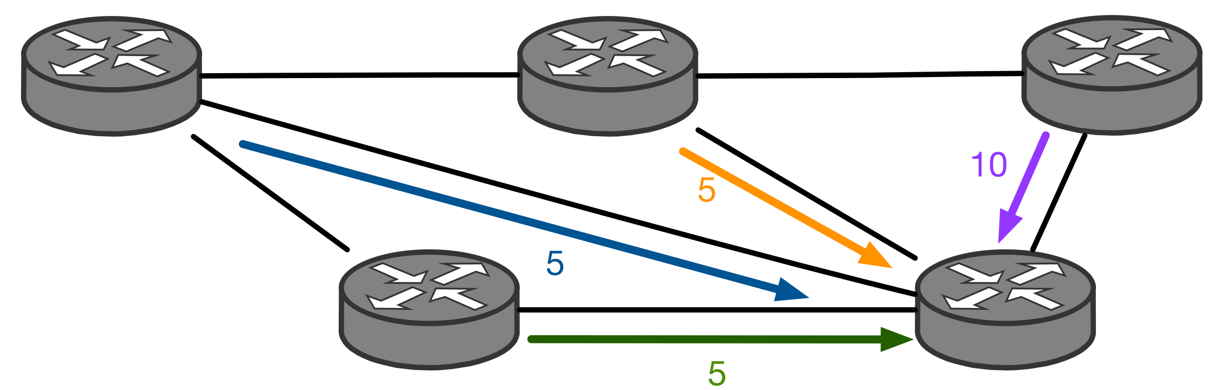

target state

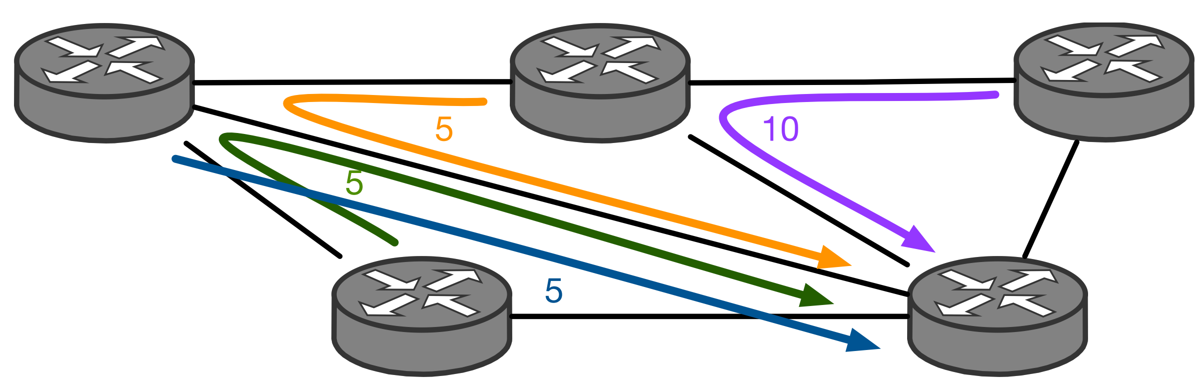

intermediate state

Consistent Data Plane Updates

Consistent Network Updates

[Reitblatt ’12]

- abstract update operation where a set of packets is guaranteed to receive consistent treatment
- per-packet or per-flow consistency
- implementation on top of NOX

Consistent Data Plane Updates

Dynamic Scheduling of Network Updates

[Jin '14]

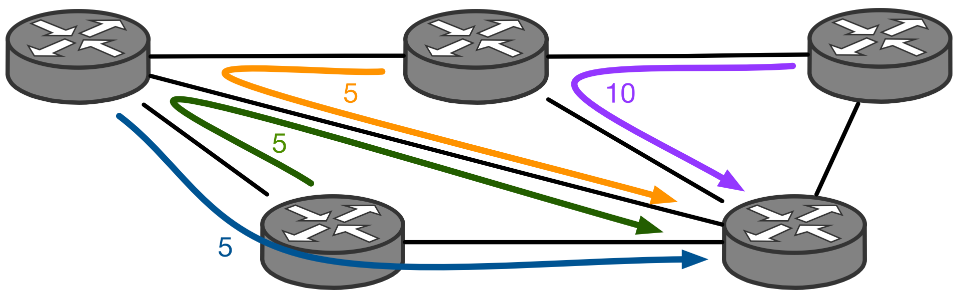

- schedule update order dynamically at runtime accounting for runtime variations
- critical path scheduling through dependency graph

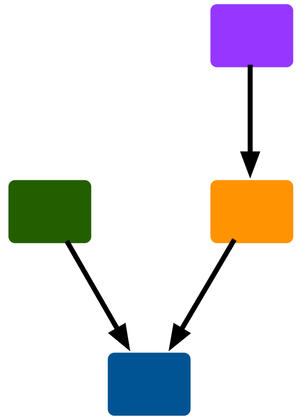

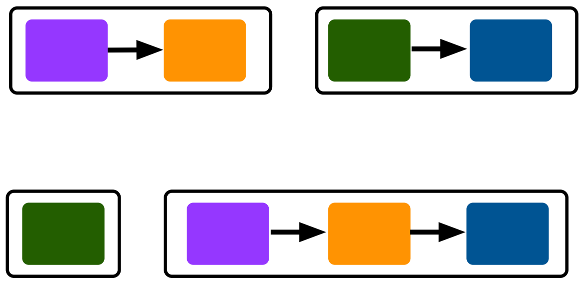

SD-WAN Deployments and Benefits

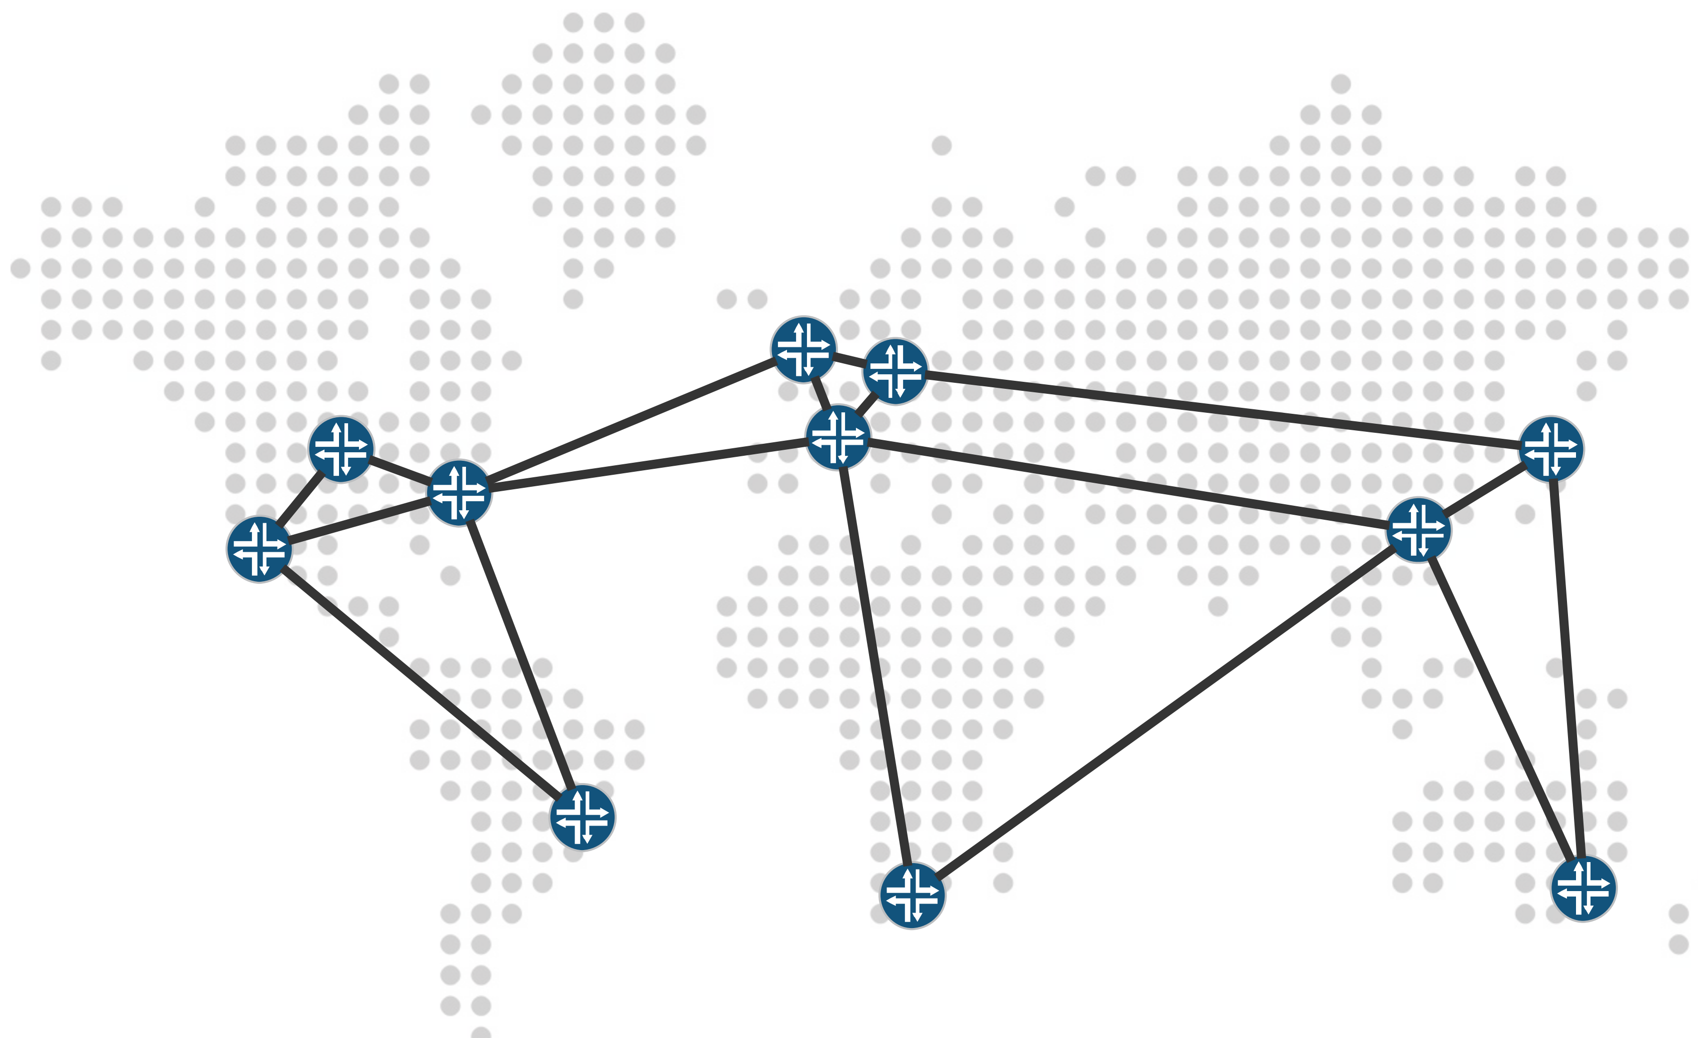

SD-WAN Deployments and Benefits

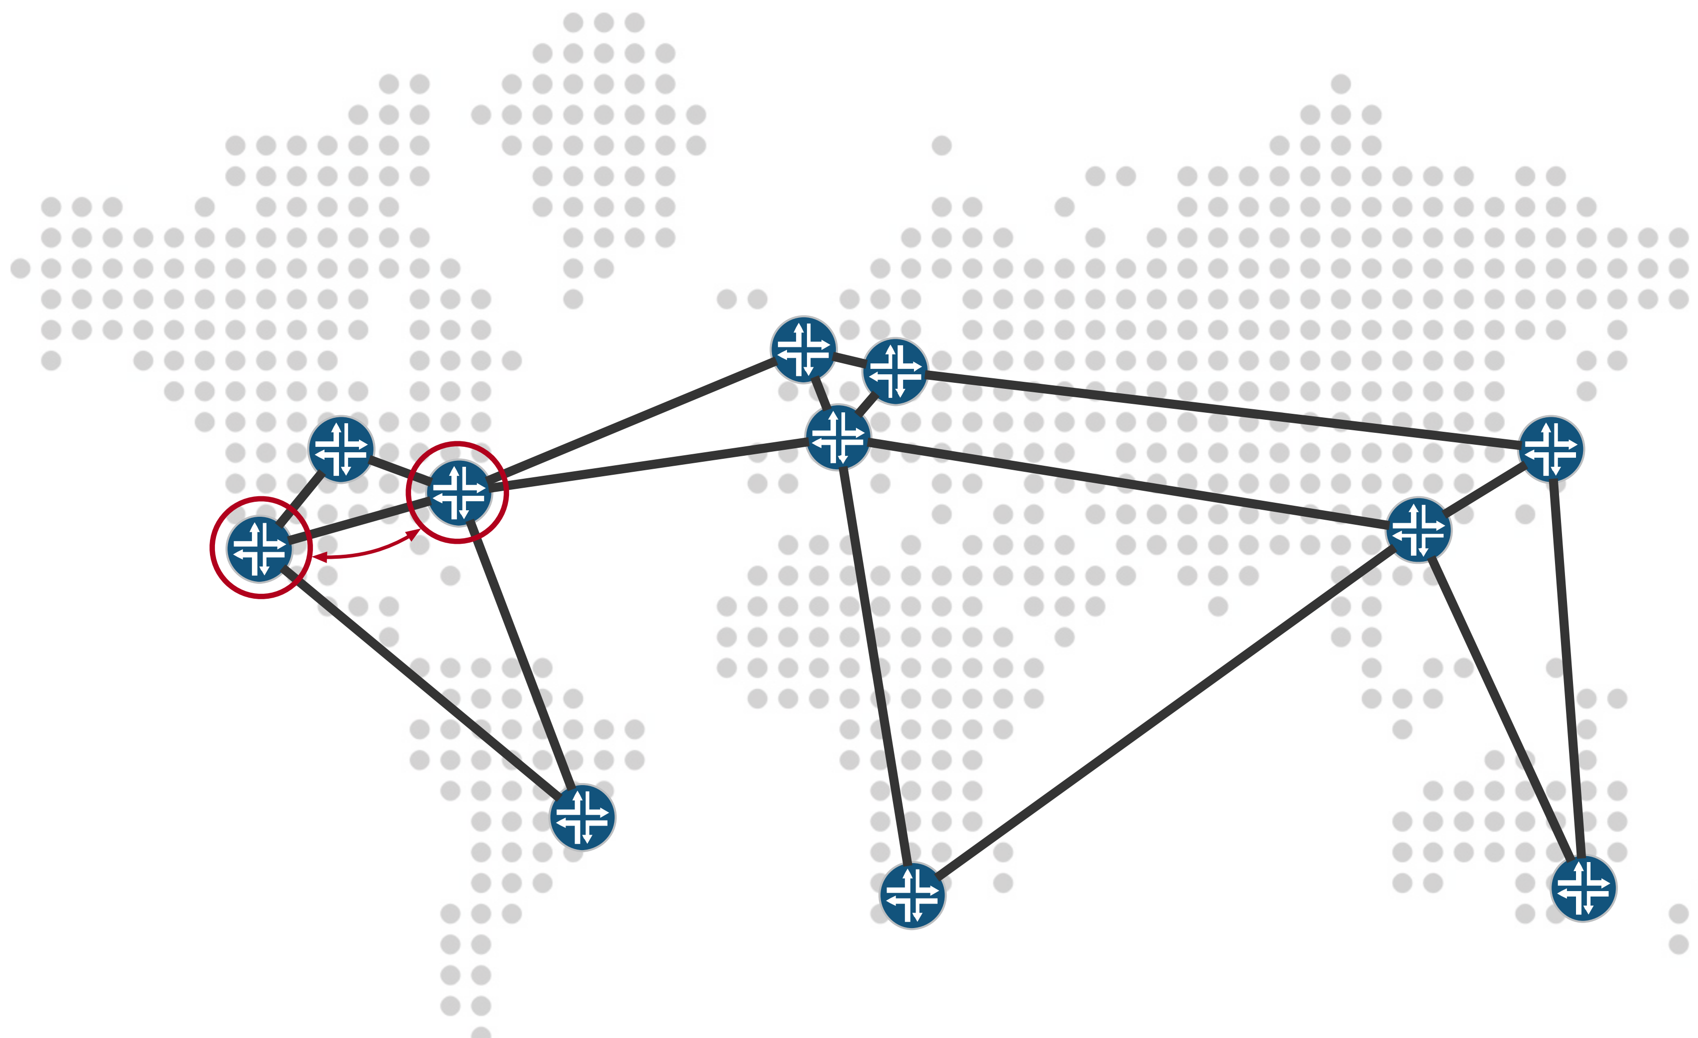

Intra-Domain

- Distributed Applications
- Inter-DC Networks
- Synchronization
- Backup

SD-WAN Deployments and Benefits

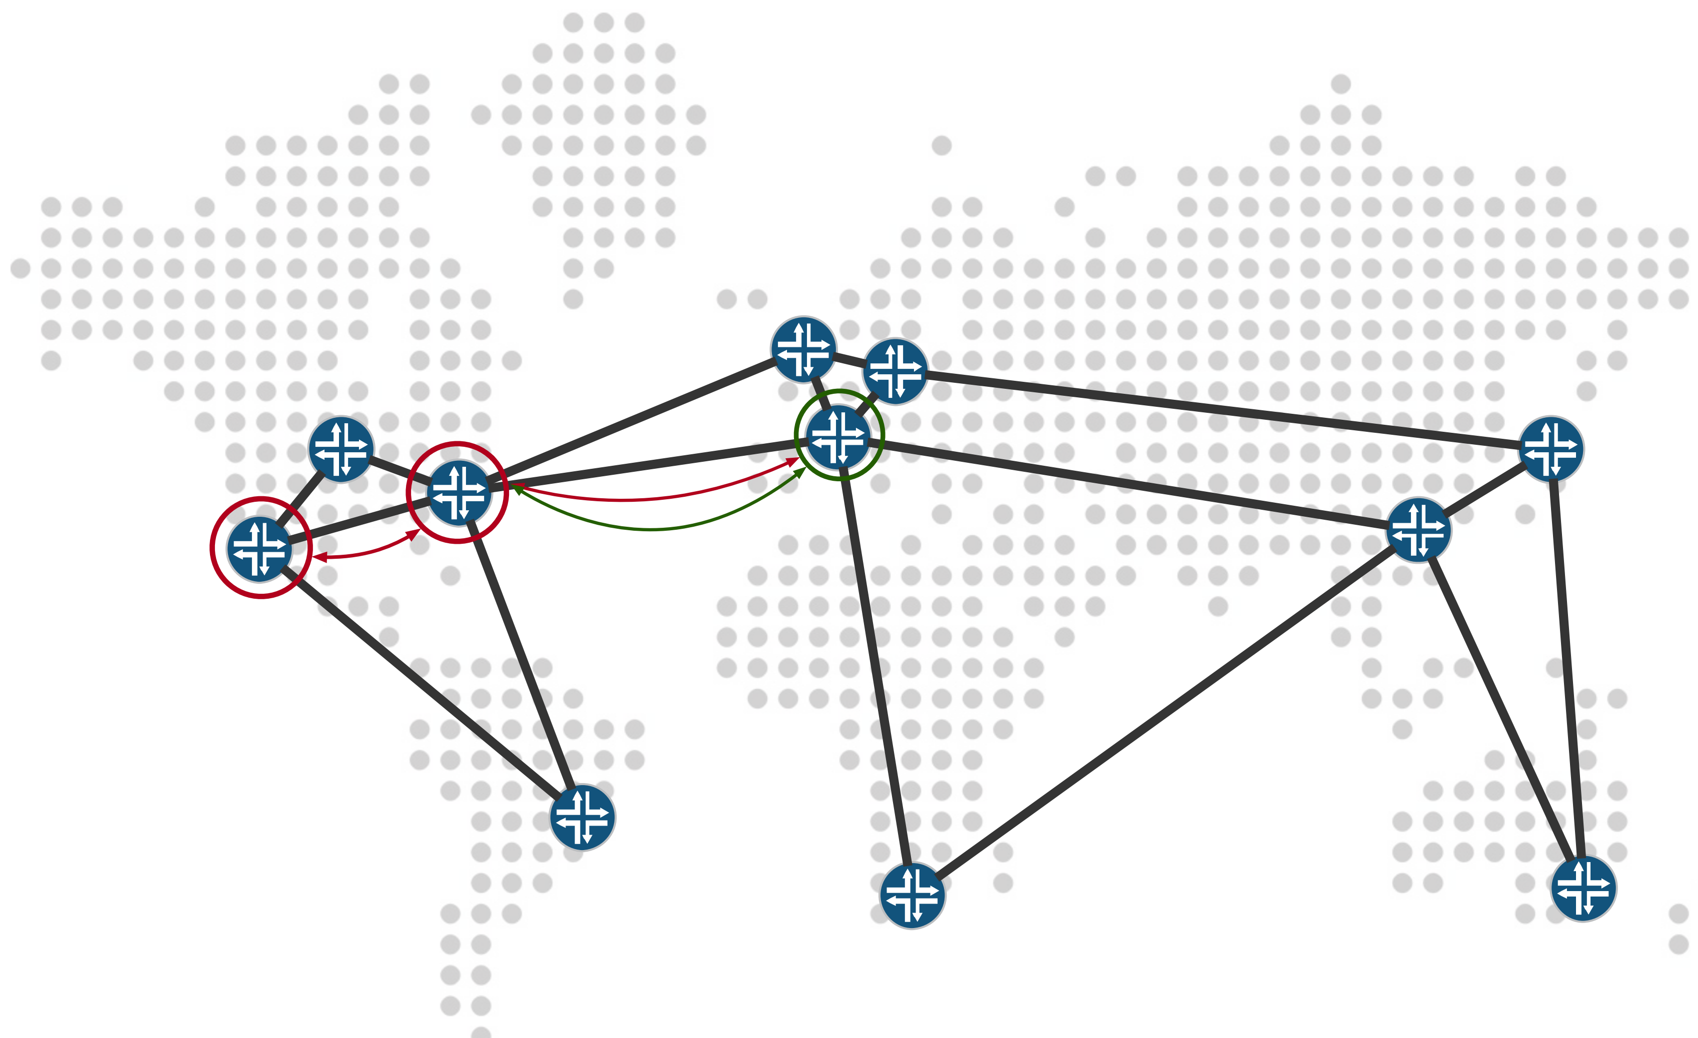

Intra-Domain

- Distributed Applications
- Inter-DC Networks
- Synchronization
- Backup

Inter-Domain

- Content Delivery
- Peering
- BGP inflexibilities

SD-WAN Deployments and Benefits

Achieving High Utilization with Software-Driven WAN

[Hong '13]

- central control of
  - bandwidth allocation for different services
  - centrally computing globally-optimal paths
- frequent data plane updates to maintain high utilization
- congestion-free updates through scratch capacity

SD-WAN Deployments and Benefits

Achieving High Utilization with Software-Driven WAN

[Hong '13]

- central control of
  - bandwidth allocation for different services
  - centrally computing globally-optimal paths
- frequent data plane updates to maintain high utilization
- congestion-free updates through scratch capacity

B4: Experience with a Globally-Deployed Software Defined WAN

[Jain '13]

- integration with legacy routing protocols
- evaluation in production network over three years

Expanding beyond a single Domain

SDX: A Software Defined Internet Exchange

[Gupta '14]

- BGP inflexibilities: indirect control over forwarding
- new use-cases: e.g. application specific peering
- SDN advantages: direct, fine-grained control
- IXPs: natural starting point

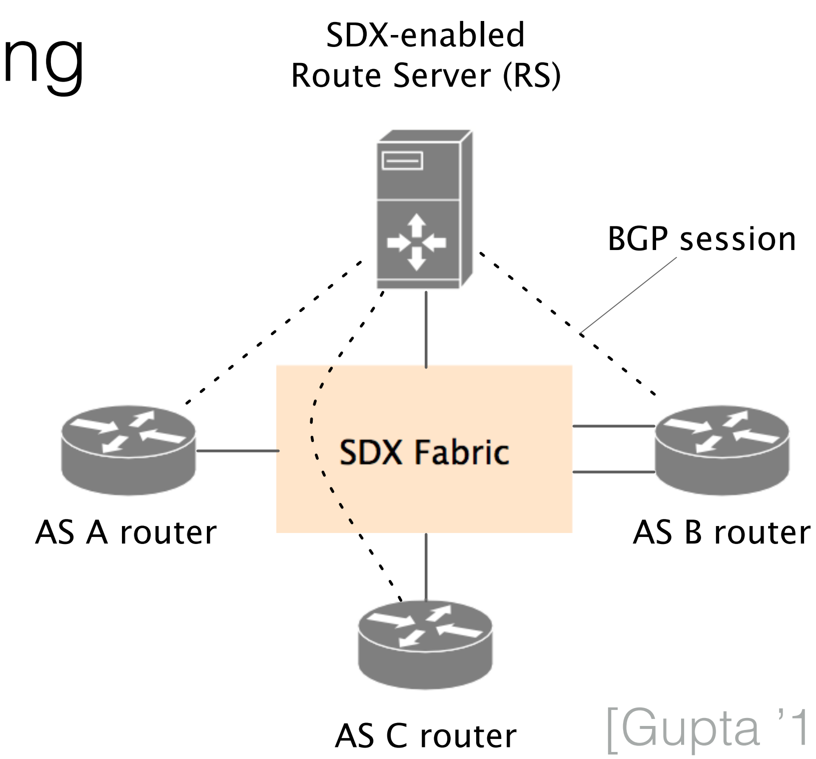

[Gupta '14]

Traffic Engineering, Data Plane Fault Tolerance, and Low-Latency Routing

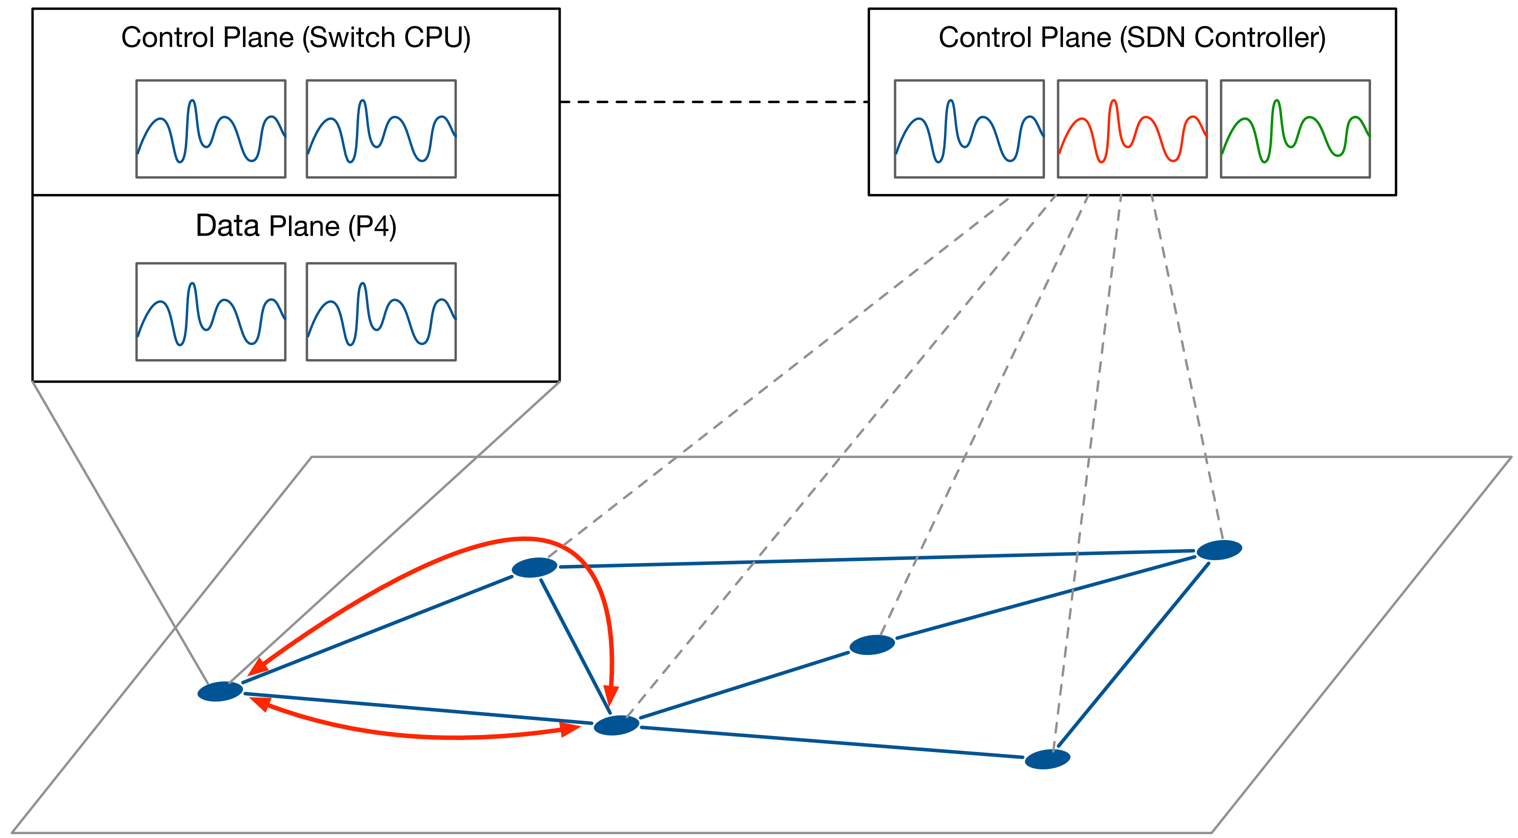

- high uncertainty and randomness in path quality
- active probing and SDN control can help to dynamically change paths
- can in part be done in the data plane (e.g., P4 technologies)

Packet-Optical Convergence

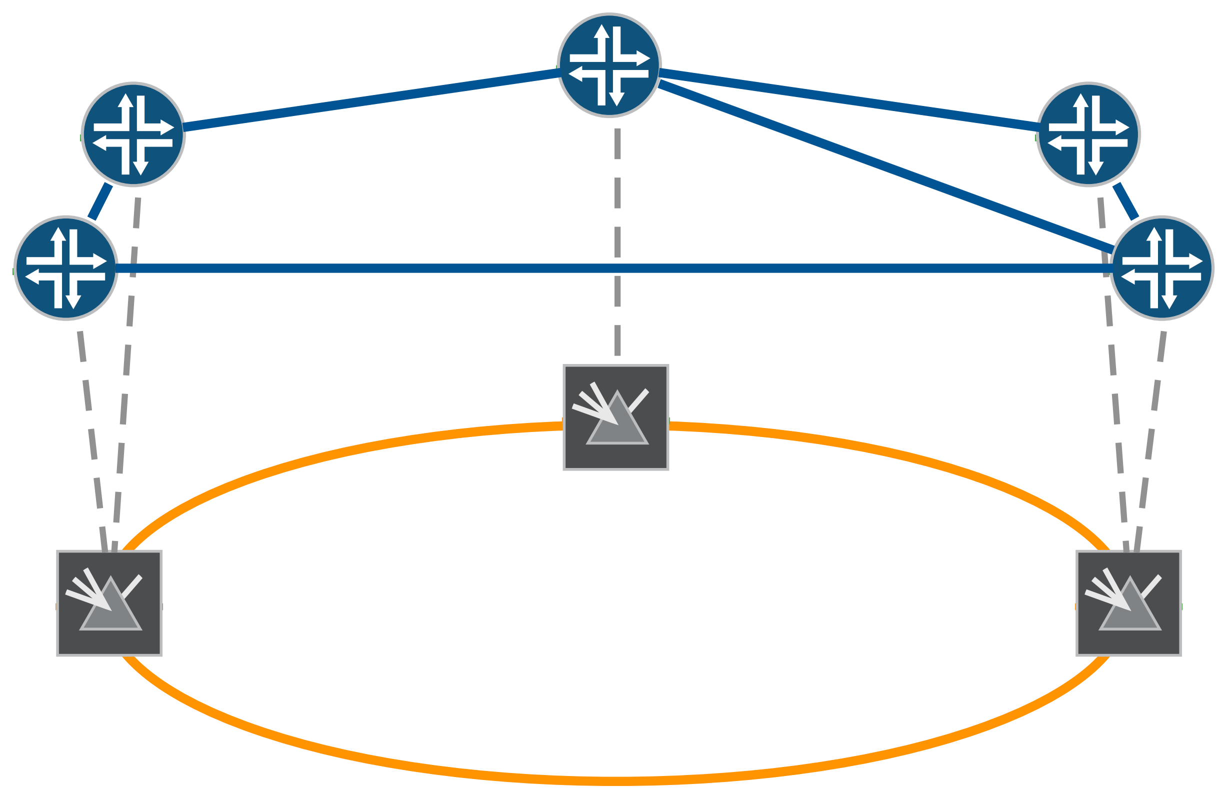

- routing over a more complex topology
- IP layer routing can use transport layer properties for CSPF routing

Internet-Scale Attacks

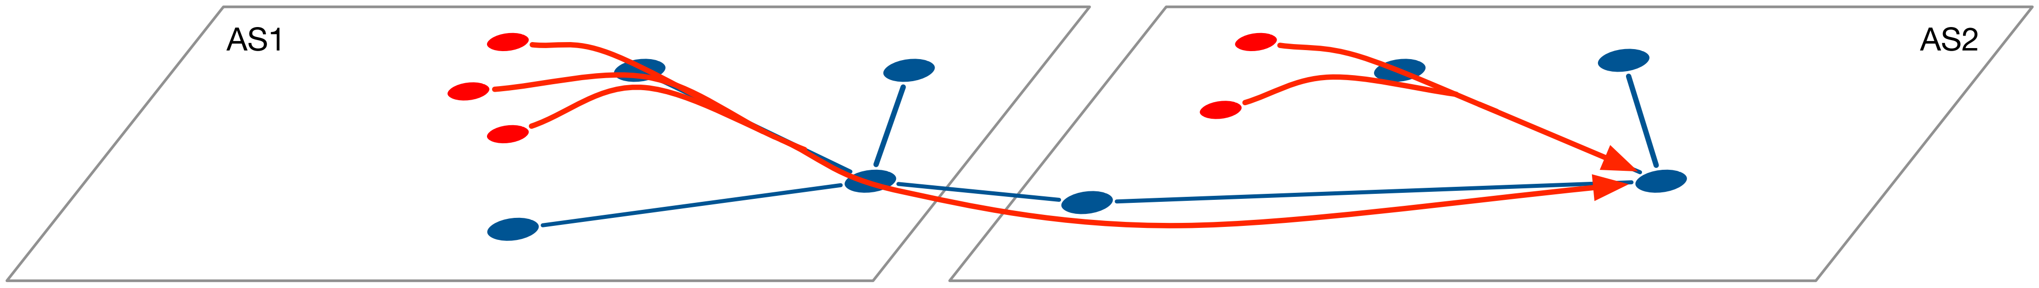

- use centralized logic for analysis and mitigation of Internet-scale attacks across domains
- fine-grained filtering with programmability

Internet-Scale Attacks

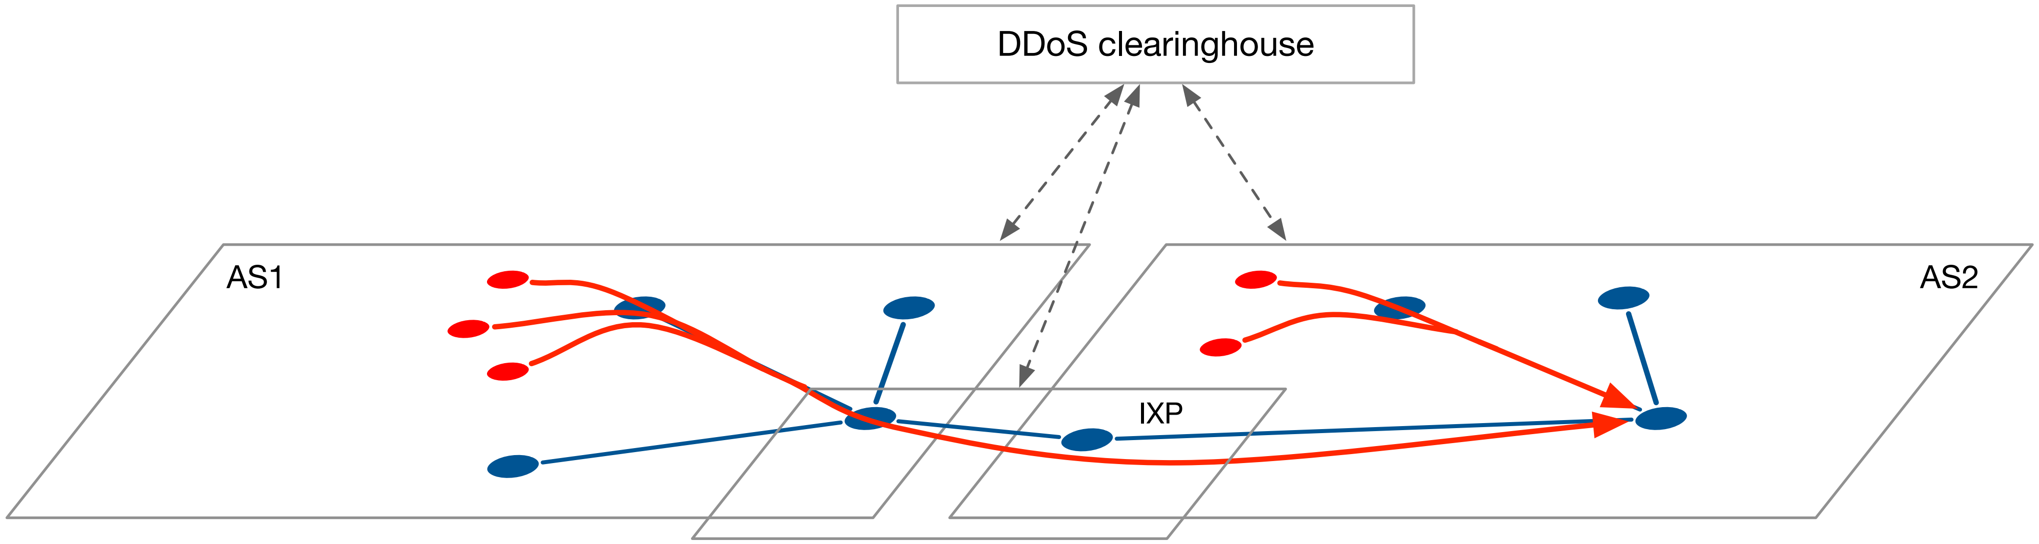

- use centralized logic for analysis and mitigation of Internet-scale attacks across domains
- fine-grained filtering with programmability

Internet-Scale Attacks

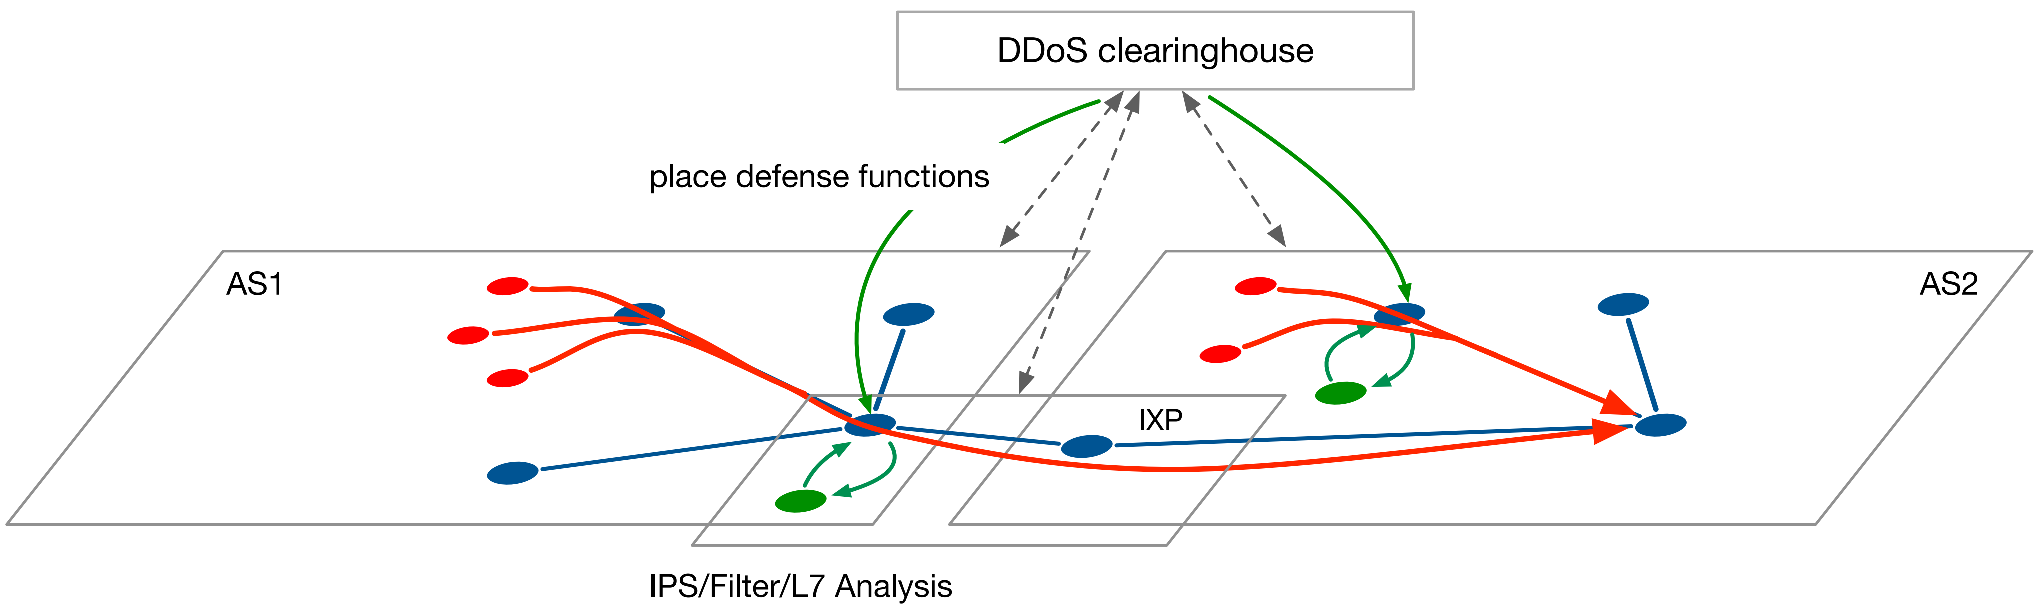

- use centralized logic for analysis and mitigation of Internet-scale attacks across domains
- fine-grained filtering with programmability

Conclusion

- WANs gaining important with mobile traffic rising
- some deployments, typically within domains
- still space for extensive research

Q&A / Discussion

Oliver Michel

oliver.michel@colorado.edu
http://nsr.colorado.edu/oliver

University of Colorado Boulder

BACKUP SLIDES

Active Networks

- rapid traffic growth in mid-‘90s, slow standardization through IETF

Active Networks

- rapid traffic growth in mid-‘90s, slow standardization through IETF
- programmability

Active Networks

- rapid traffic growth in mid-‘90s, slow standardization through IETF
- programmability
- code embedded in packets

Active Networks

- rapid traffic growth in mid-‘90s, slow standardization through IETF
- programmability
- code embedded in packets
- no clear use-cases or applications

Control- and Data Plane Separation

- networks rapidly increasing in size and complexity

Control- and Data Plane Separation

- networks rapidly increasing in size and complexity
- scalability issues

Control- and Data Plane Separation

- networks rapidly increasing in size and complexity
- scalability issues
- manageability issues

Scalability

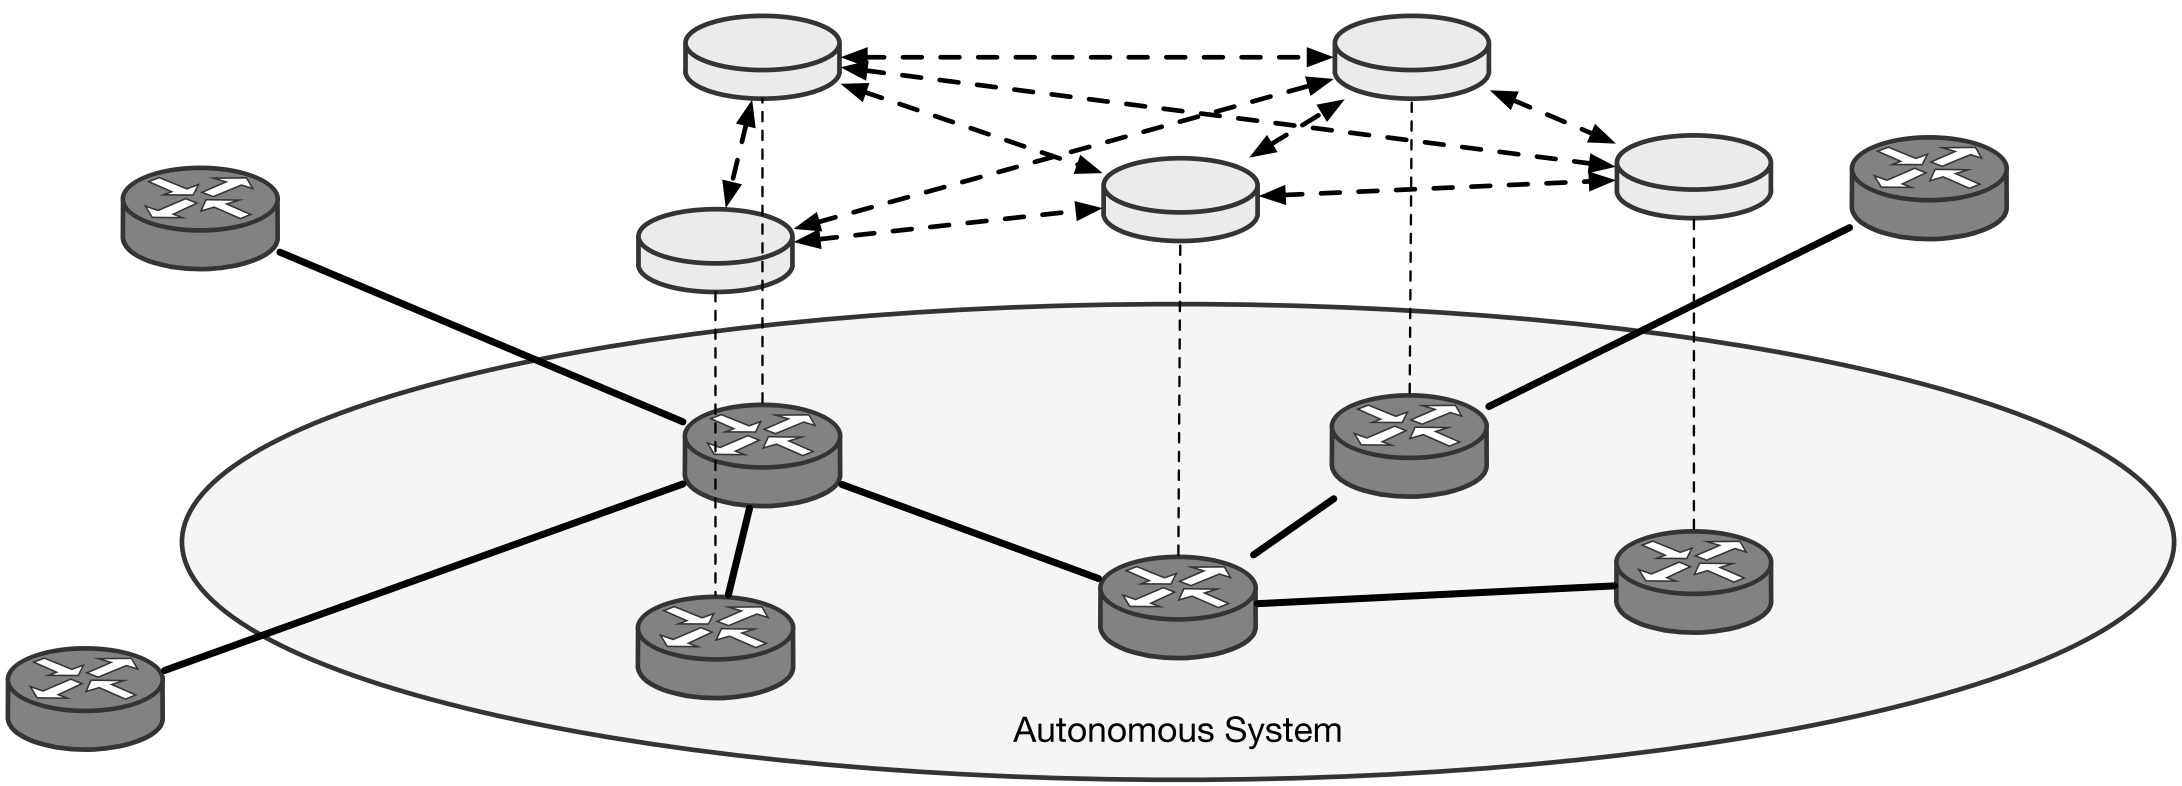

Scalability

Routing Control Platform (RCP)

[Caesar ’05]

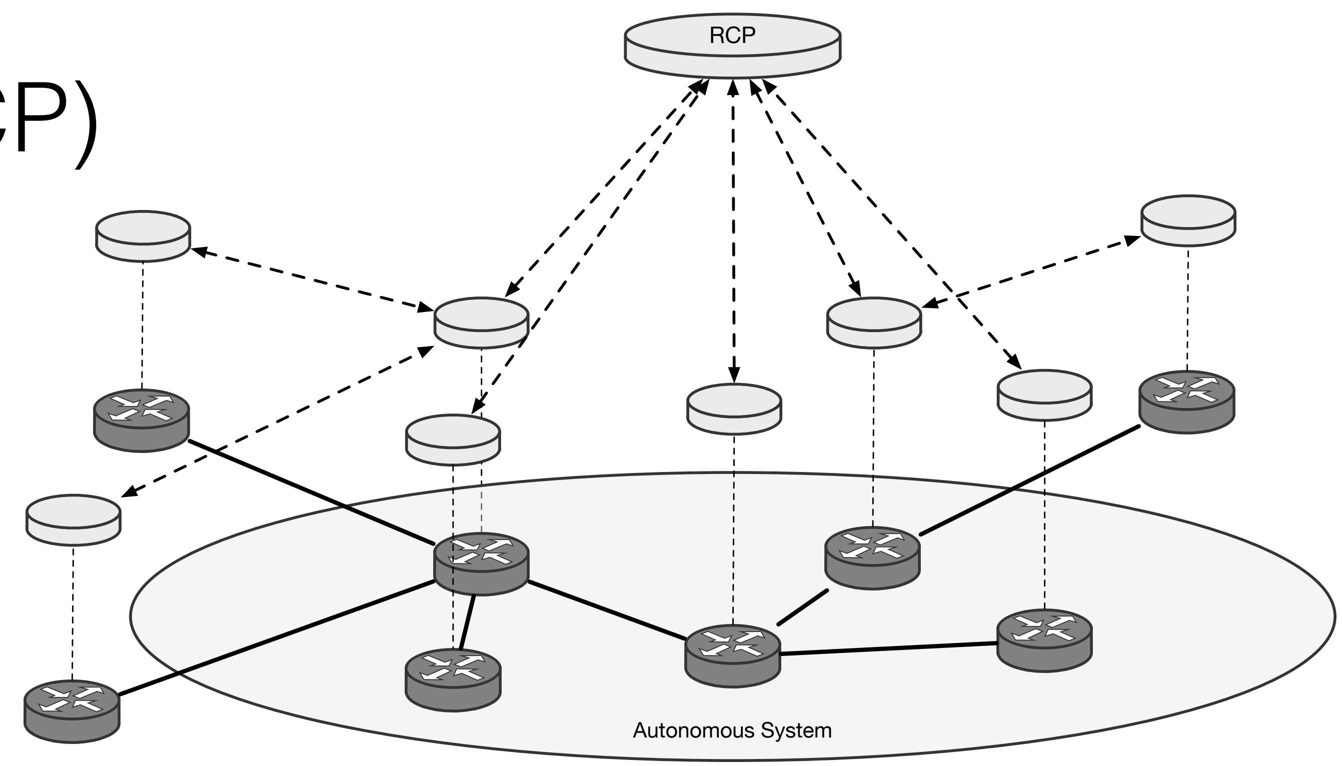

- routers peer with RCP
- mimics full iBGP mesh
- single best route advertised via standard iBGP
- intrinsic correctness of full mesh with scalability of RR
- no route oscillations or forwarding loops

Manageability

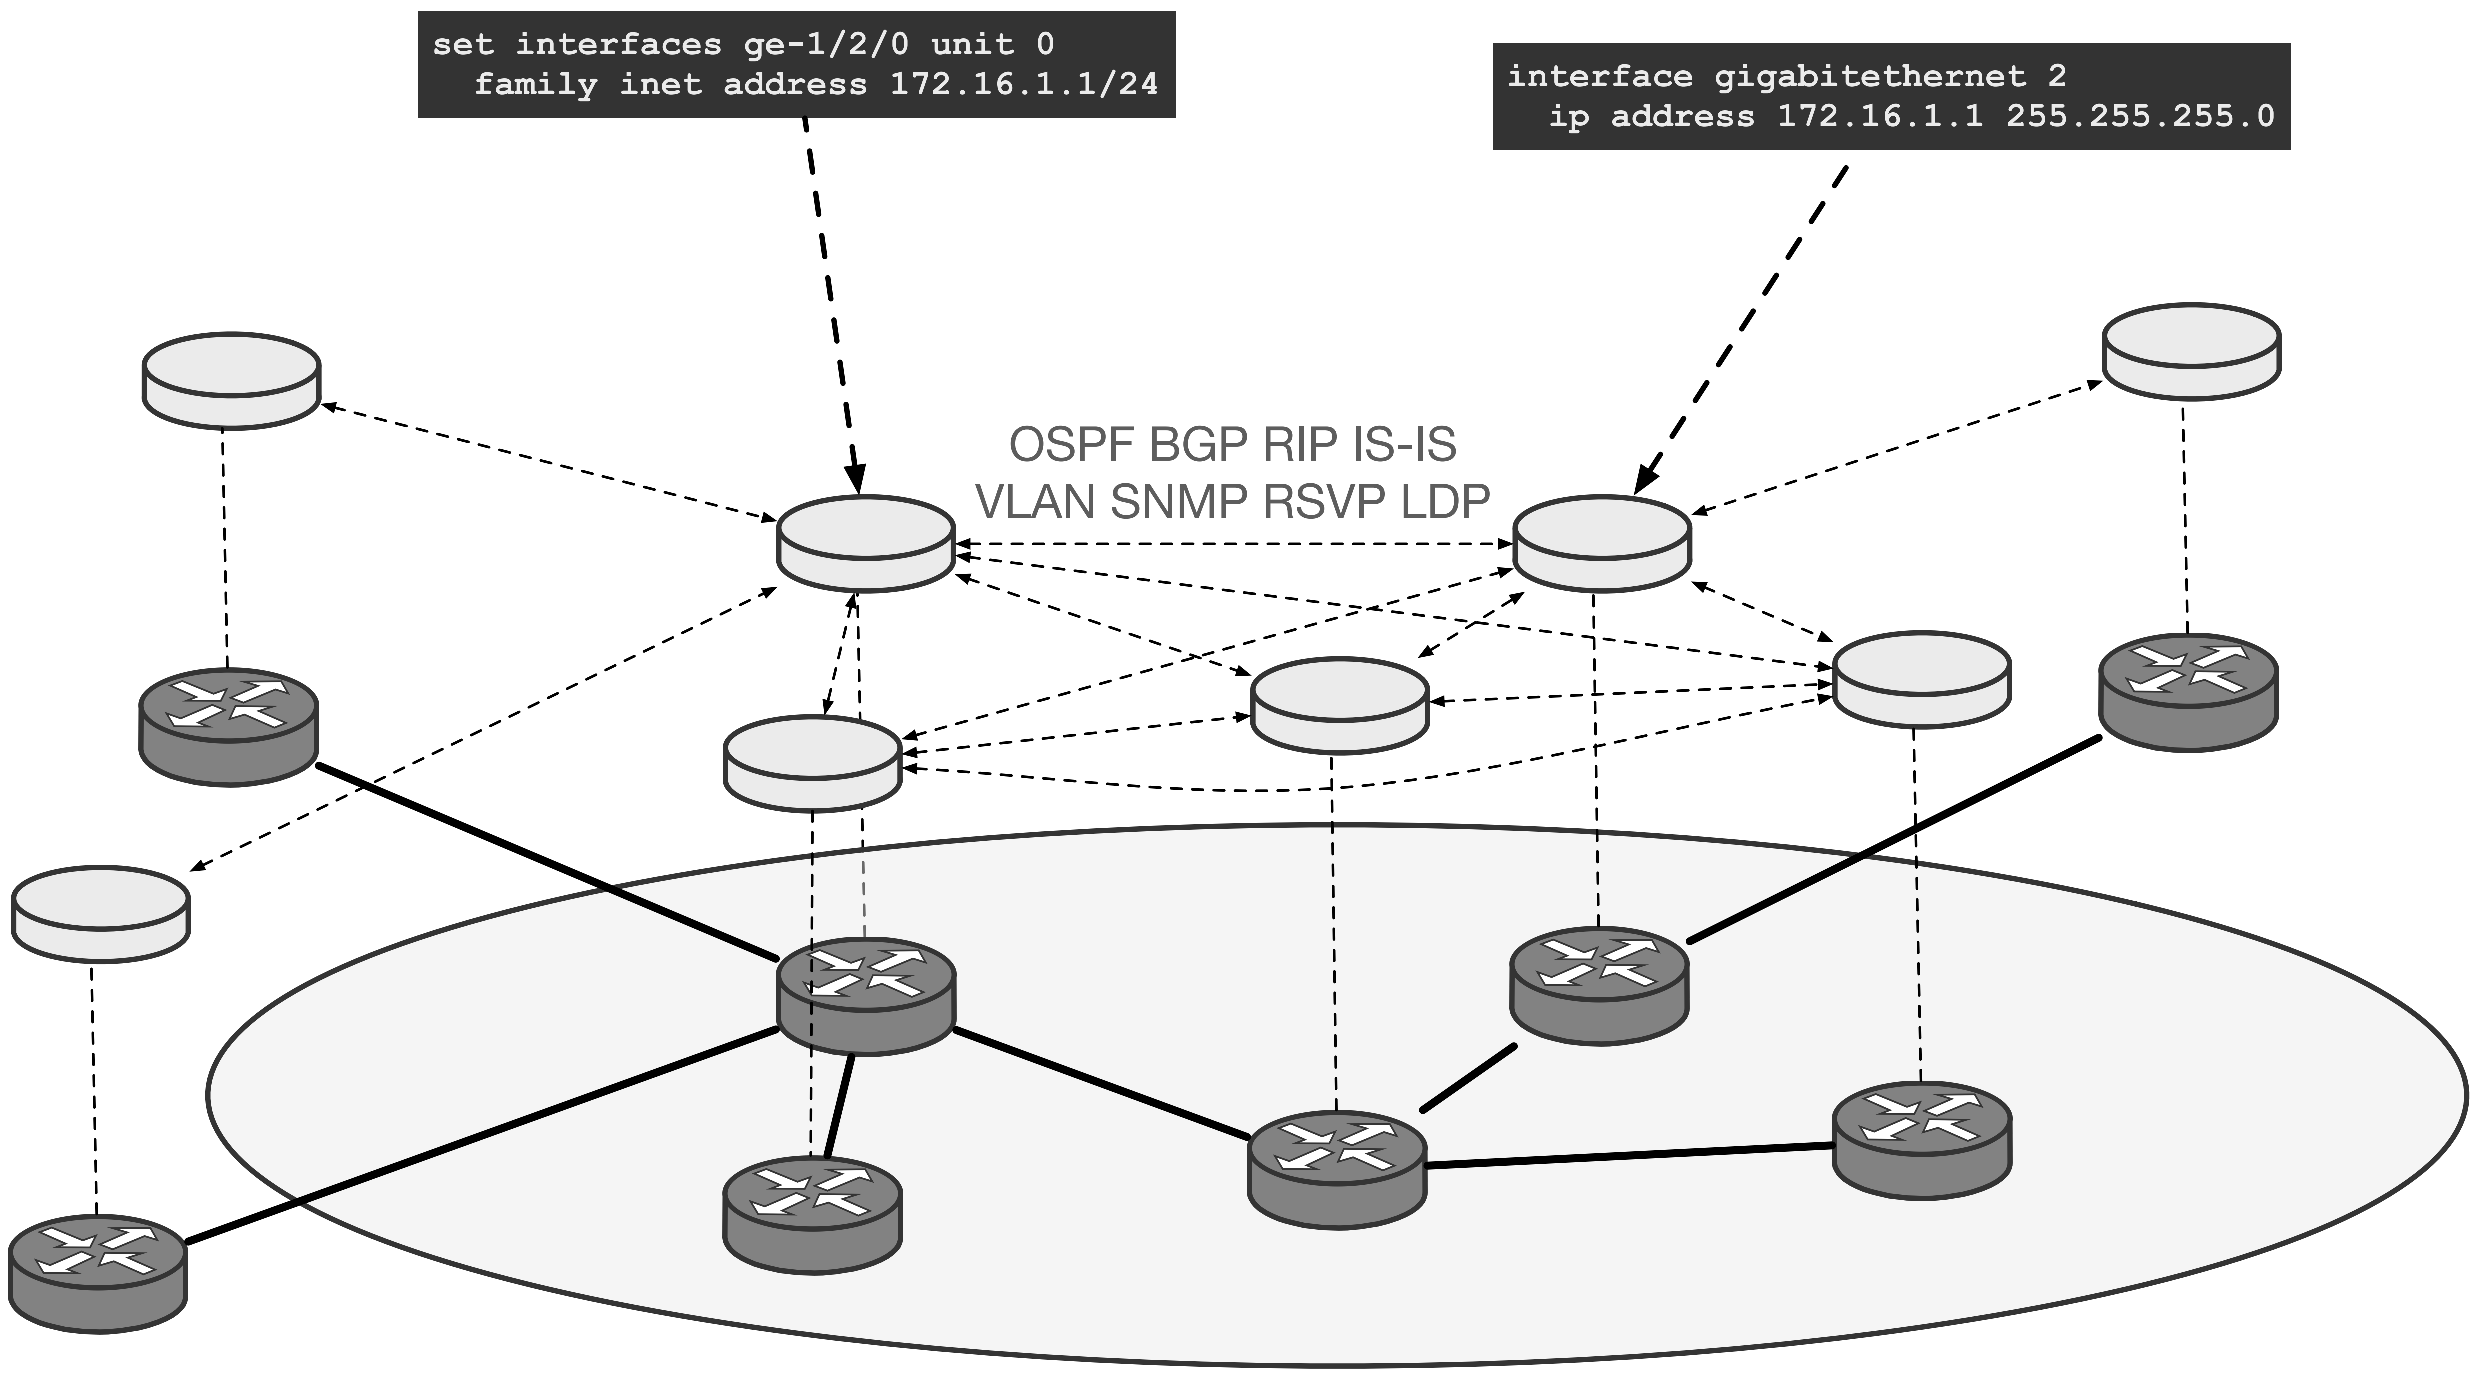

Manageability

- novel architectural pattern for networks based on layers [Greenberg '05]
  - four different layers
  - control/data separation

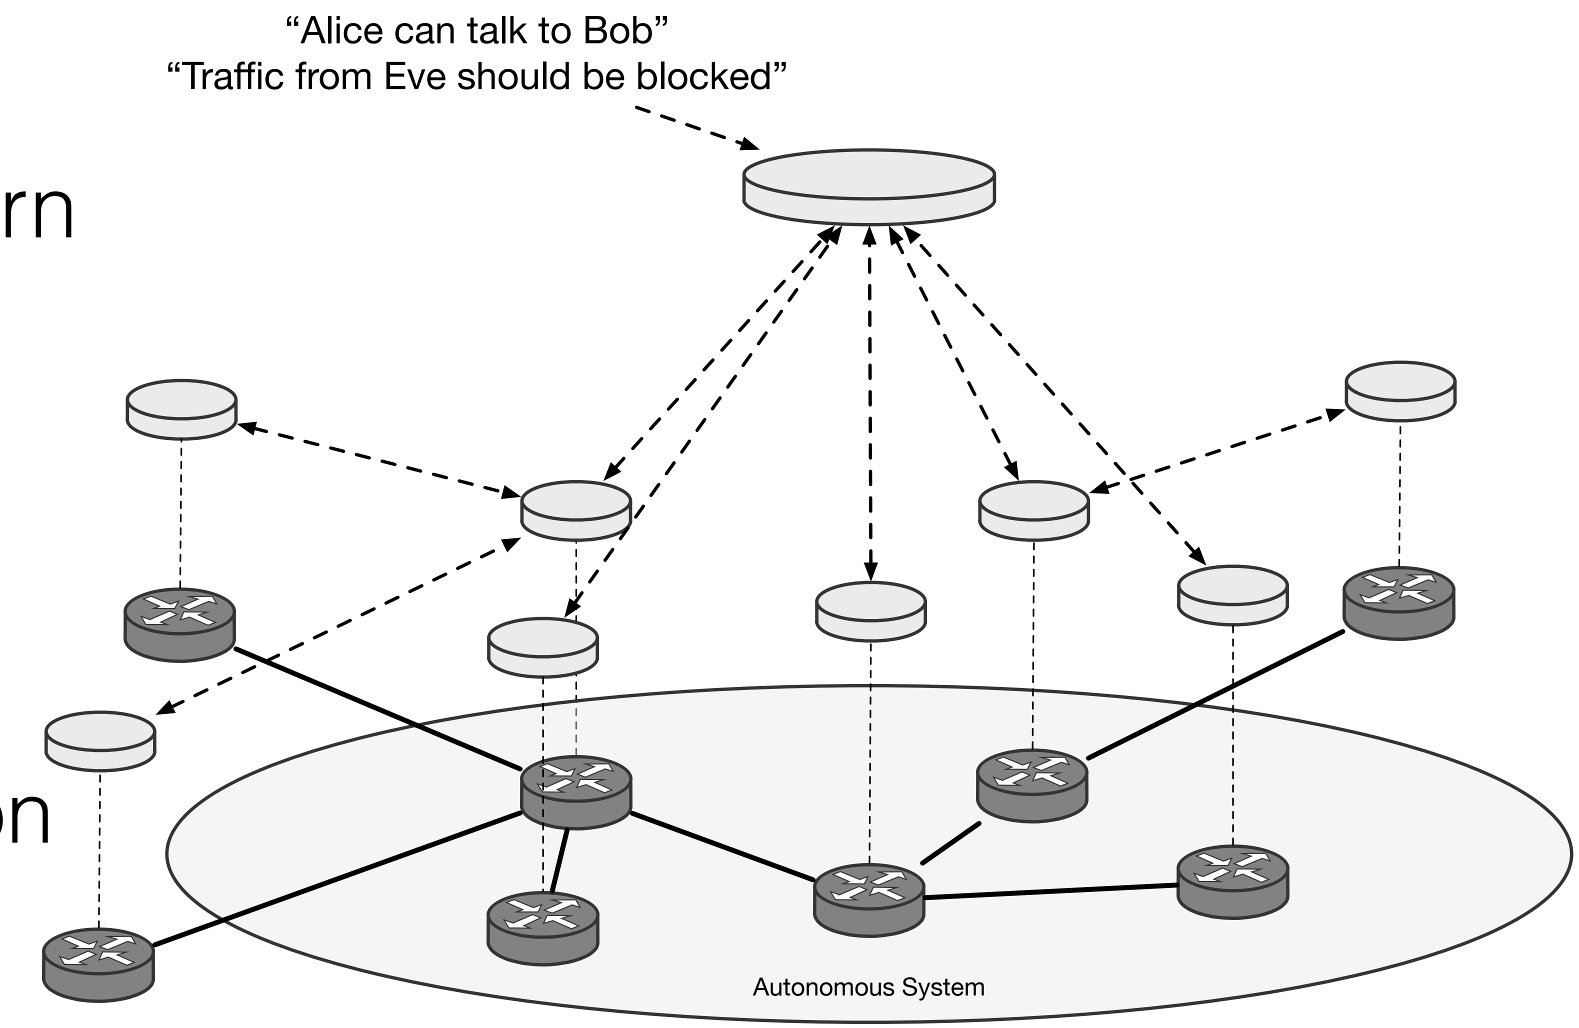

- high-level network policies through centralized controller [Casado '07]
  - simple switch architecture
  - evaluated in real-world deployment

Control Protocols

- need for standardized control between control and data plane
  - generalization of networking equipment

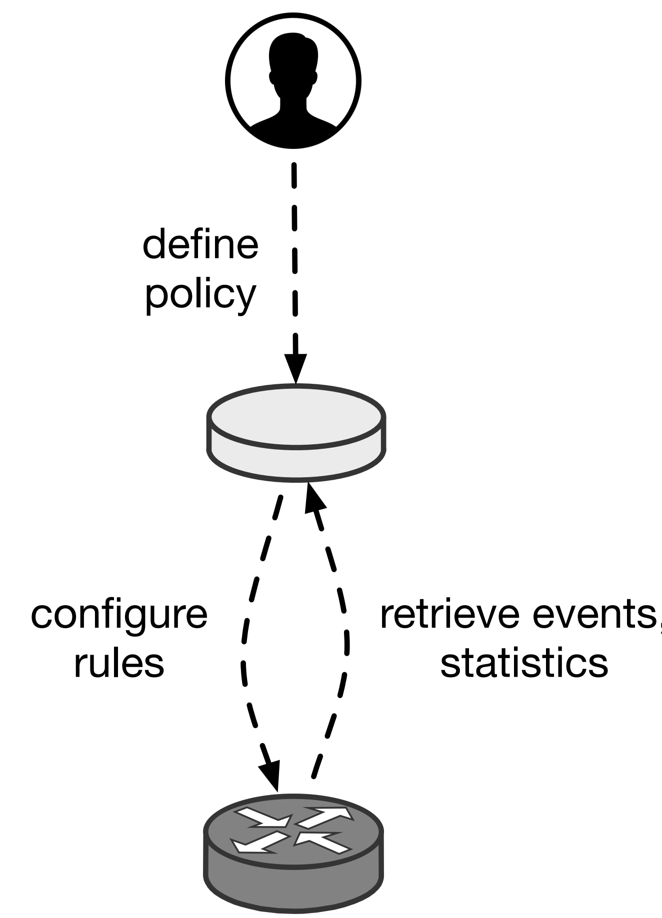

Control Protocols

- OpenFlow [McKeown ’08]
  - open protocol that gives applications control over a switches data plane
  - designed around a set of header match fields and forwarding actions
  - forwarding abstraction balancing…
    1. general match/action (TCAM model)
    2. fixed-function switch ASICs
- not the only protocol

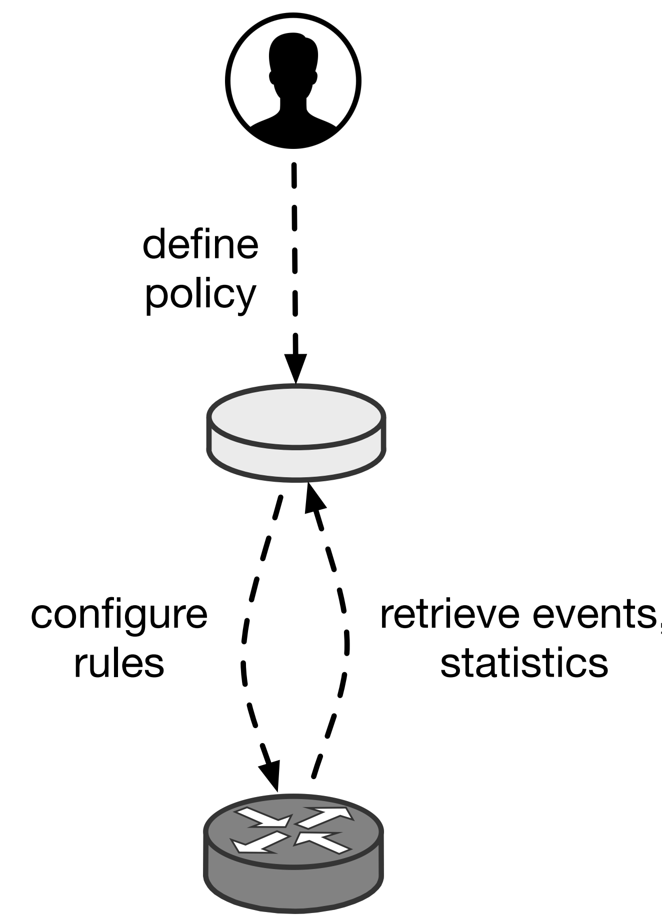

An Industrial-Scale Software Defined Internet Exchange Point

[Gupta ’16]

- fundamental scaling problems in SDX architecture
    - composition of rules requires large state
    - more rules than policies defined needed due to BGP congruence checking

Open Network Operating System [Berde '14]

- global network view shared across all instances
- scale-out and failure resiliency
- each switch connected to primary OF Manager
- new primary selected at failure through consensus protocol by Zookeeper
- distributed, eventually consistent network graph through Cassandra backend

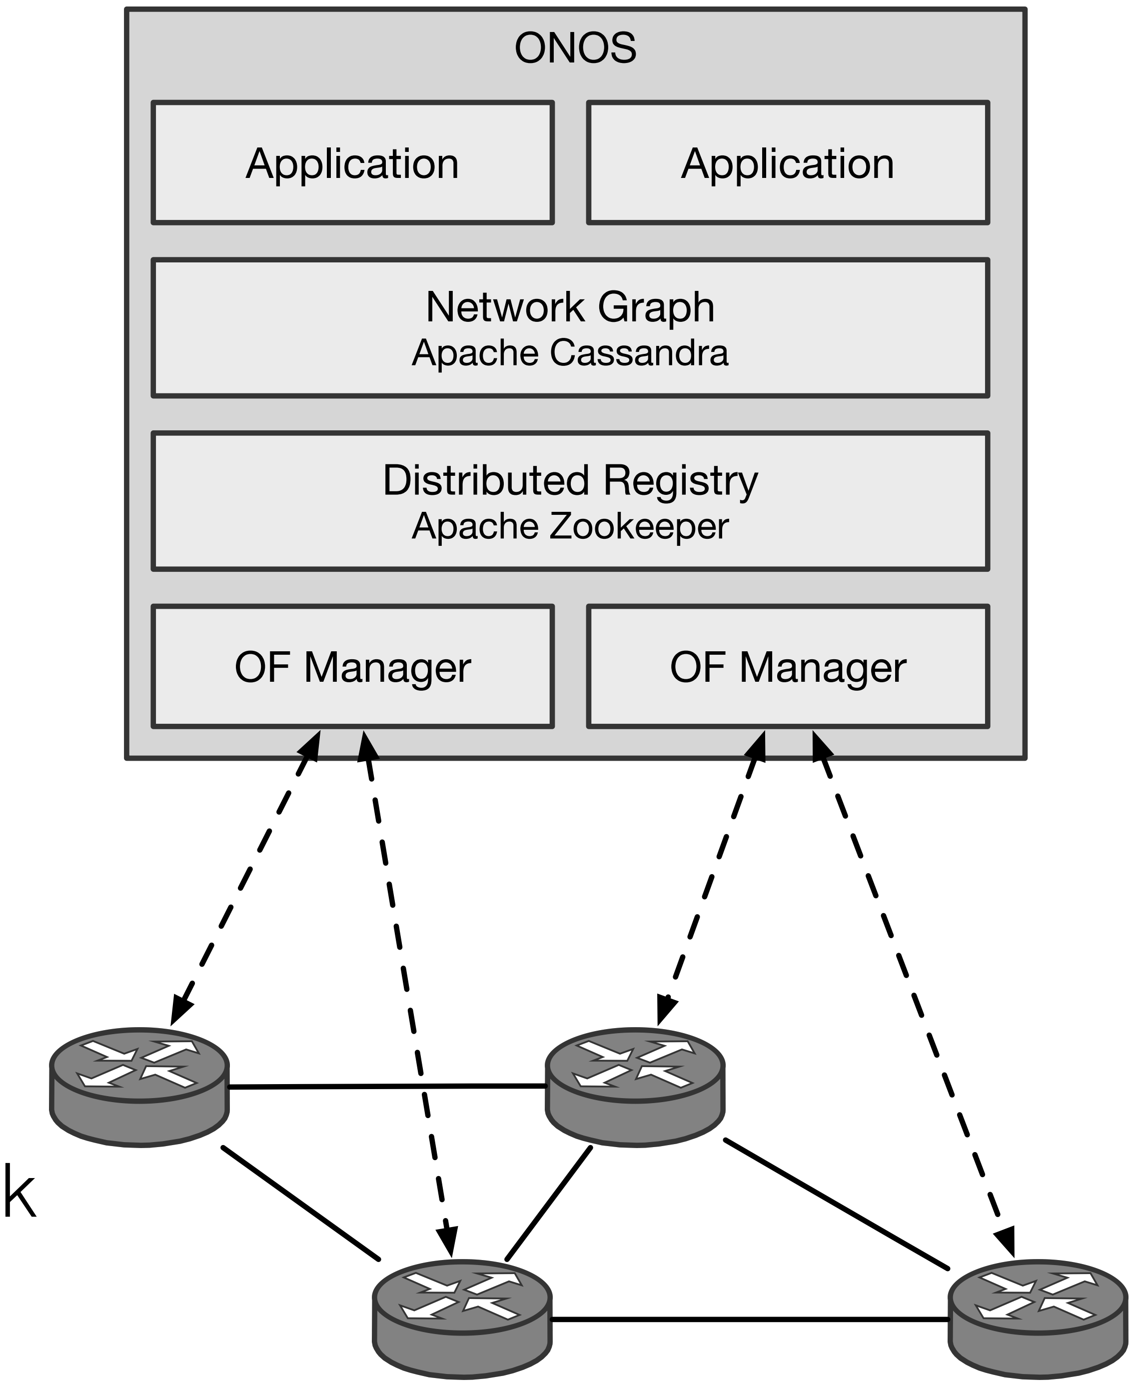

NorthStar Controller

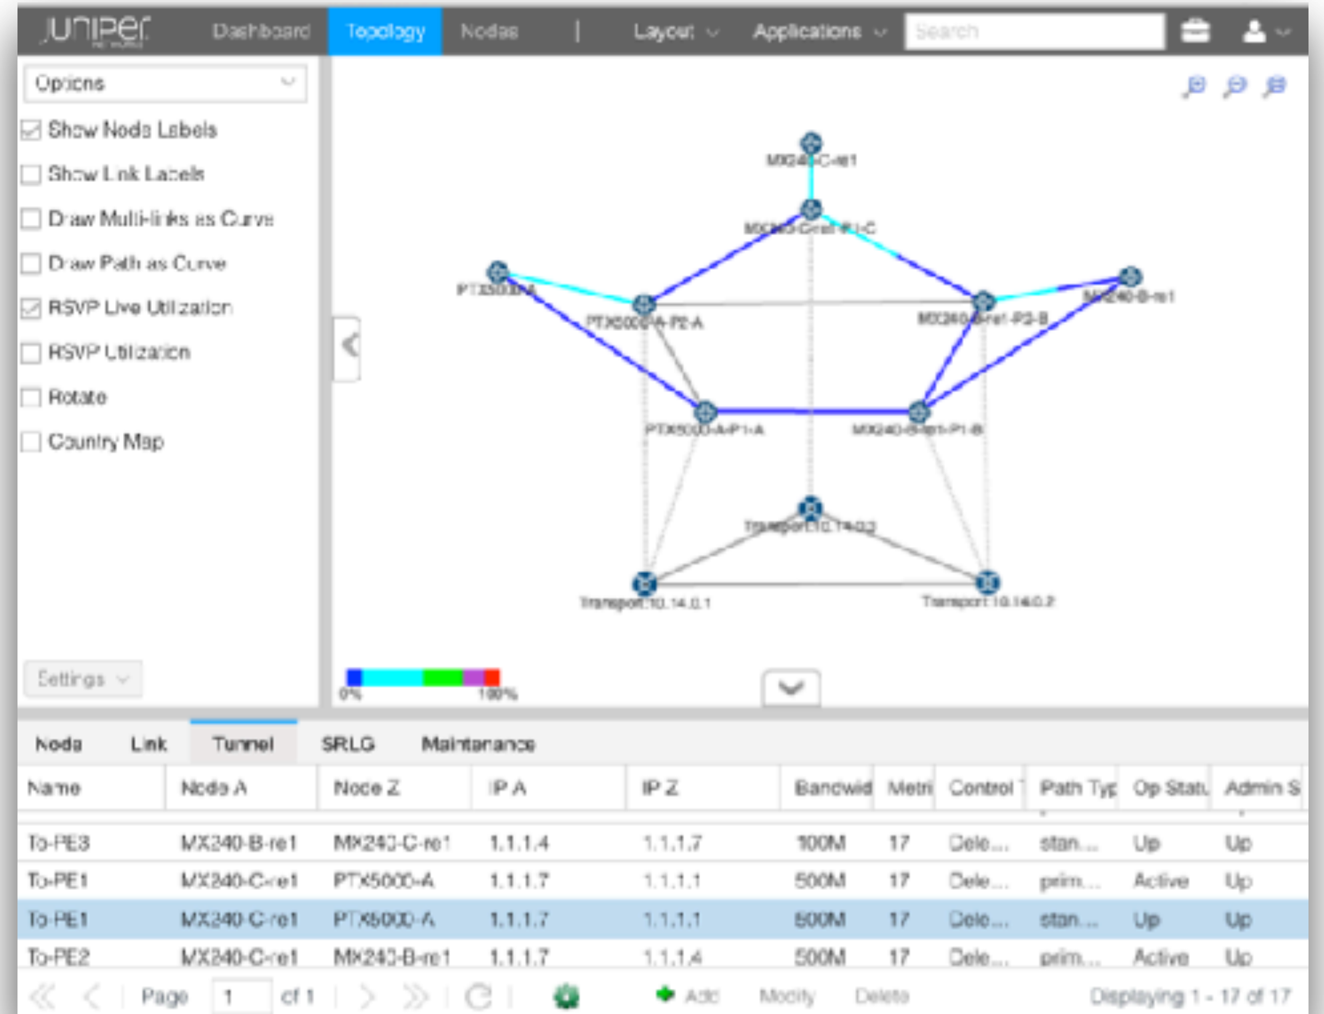

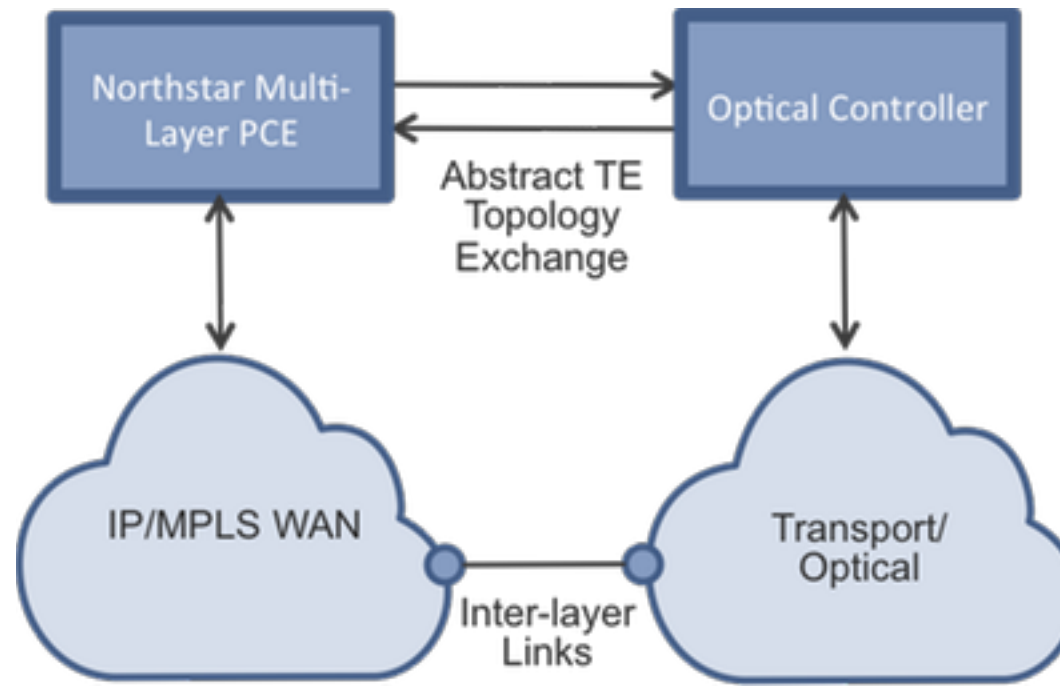

- Multi-Layer WAN Traffic Engineering Solution
- Controller-Controller Interface

Expanding beyond a single Domain

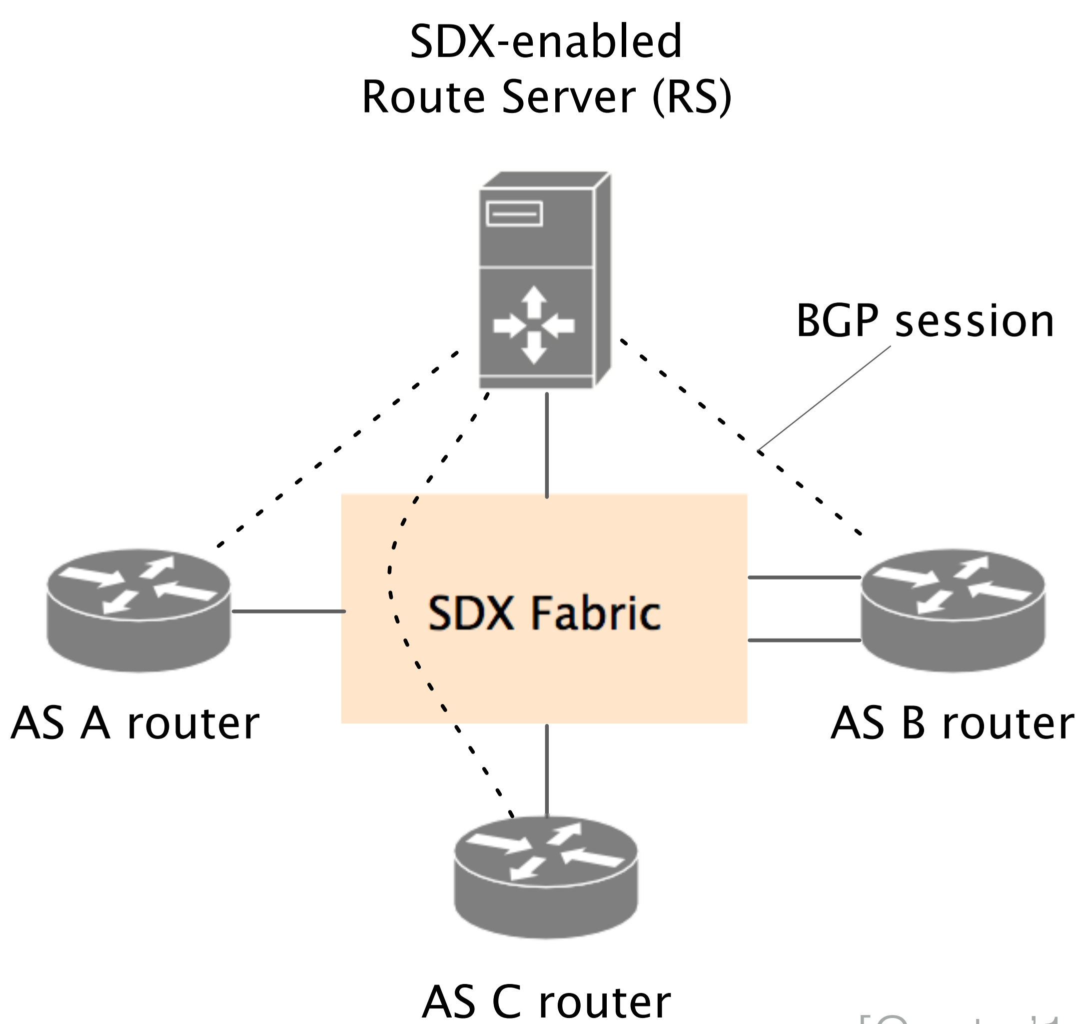

[Gupta '14]

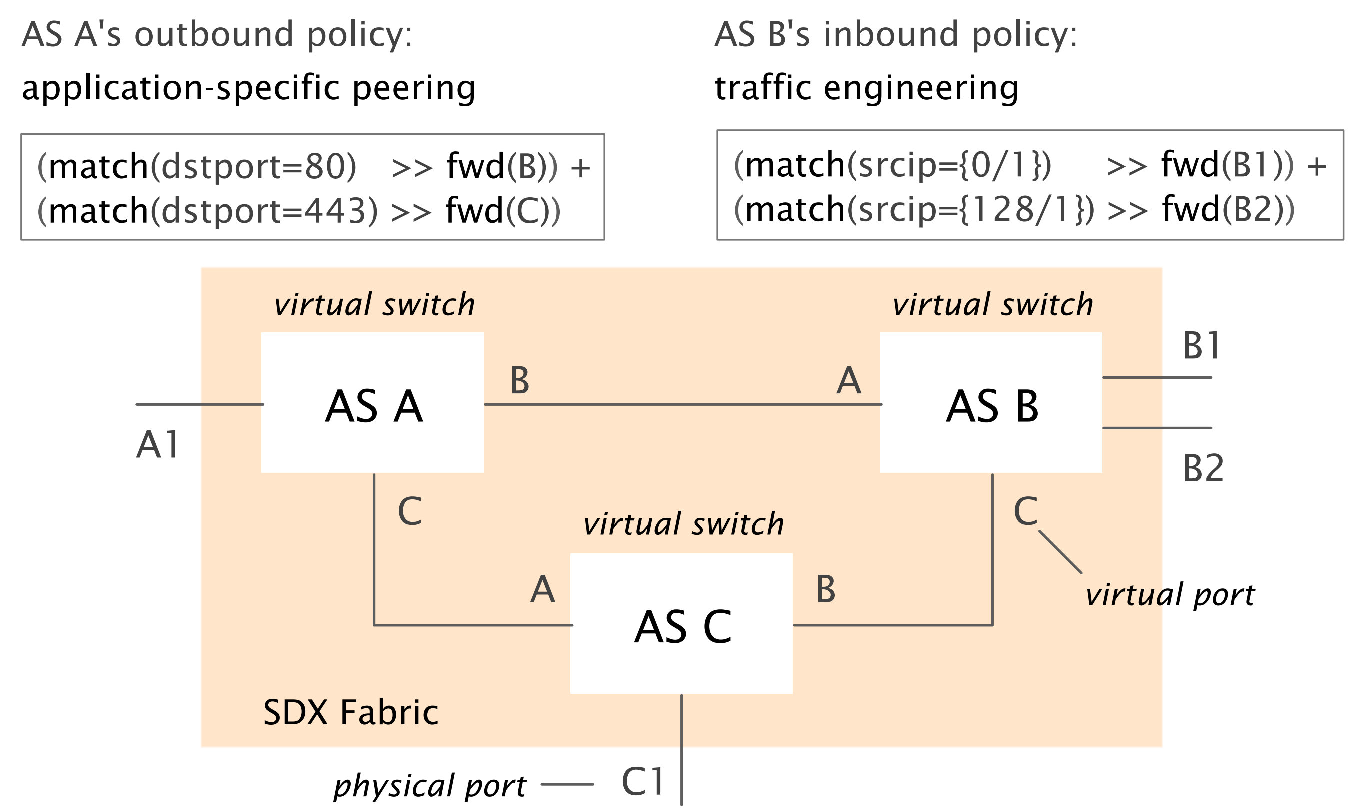

[Gupta '14]