

Software-Defined Internet Exchange

5590: software defined networking

anduo wang, Temple University

T 17:30-20:00

SDX: A Software Defined Internet Exchange

today's Internet routing

border gateway protocol (BGP), many problems

- (IP) destination based routing
- influence only direct neighbors
- indirect expression of policy

Internet exchange point (IXP)

layer 2 location where multiple networks meet to exchange traffic

- already taken root in Europe
- expected to emerge in North America

peering tension with emerging applications (e.g., video traffic)

- has the need

right place (front line) to innovate

- has the incentive

wider range of routing decision

IXP+SDN = SDX

BGP

- (IP) destination based routing
- influence only direct neighbors
- indirect expression of policy

SDX

- far more flexible routing decision?

wider range of routing decision

application specific peering

- two neighboring ASes exchange traffic only for certain applications (HTTP, HTTPS)

BGP

- (IP) destination based routing
- influence only direct neighbors
- indirect expression of policy

solution?

wider range of routing decision

application specific peering

- two neighboring ASes exchange traffic only for certain applications (HTTP, HTTPS)

BGP

- (IP) destination based routing
- influence only direct neighbors
- indirect expression of policy

solution?

SDX

- install custom rules for groups of flows corresponding to a specific application

wider range of routing decision

inbound traffic engineering

- control how traffic enter a network

BGP

- (IP) destination based routing
- influence only direct neighbors
- indirect expression of policy

solution?

wider range of routing decision

inbound traffic engineering

- control how traffic enter a network

BGP

- (IP) destination based routing
- influence only direct neighbors
- indirect expression of policy

solution?

SDX

- install rules at an exchange point that directly control inbound traffic according to source IP or port number

wider range of routing decision

wider-area load balancing

- content providers balance client requests across multiple servers

BGP

- (IP) destination based routing
- influence only direct neighbors
- indirect expression of policy

solution?

wider range of routing decision

wider-area load balancing

- content providers balance client requests across multiple servers

BGP

- (IP) destination based routing
- influence only direct neighbors
- indirect expression of policy

solution?

SDX

- announces an anycast IP to clients, rewrite the client requests in the middle of the network

IXP+SDN = SDX?

but SDN

- limited to intradomain: by definition, an SDN controller has purview of the entire network
- SDN is only a platform, not the solution

SDX goals/challenges

- compelling applications
- programming abstraction
- scalable operation
- realistic deployment

IXP+SDN = SDX?

but SDN

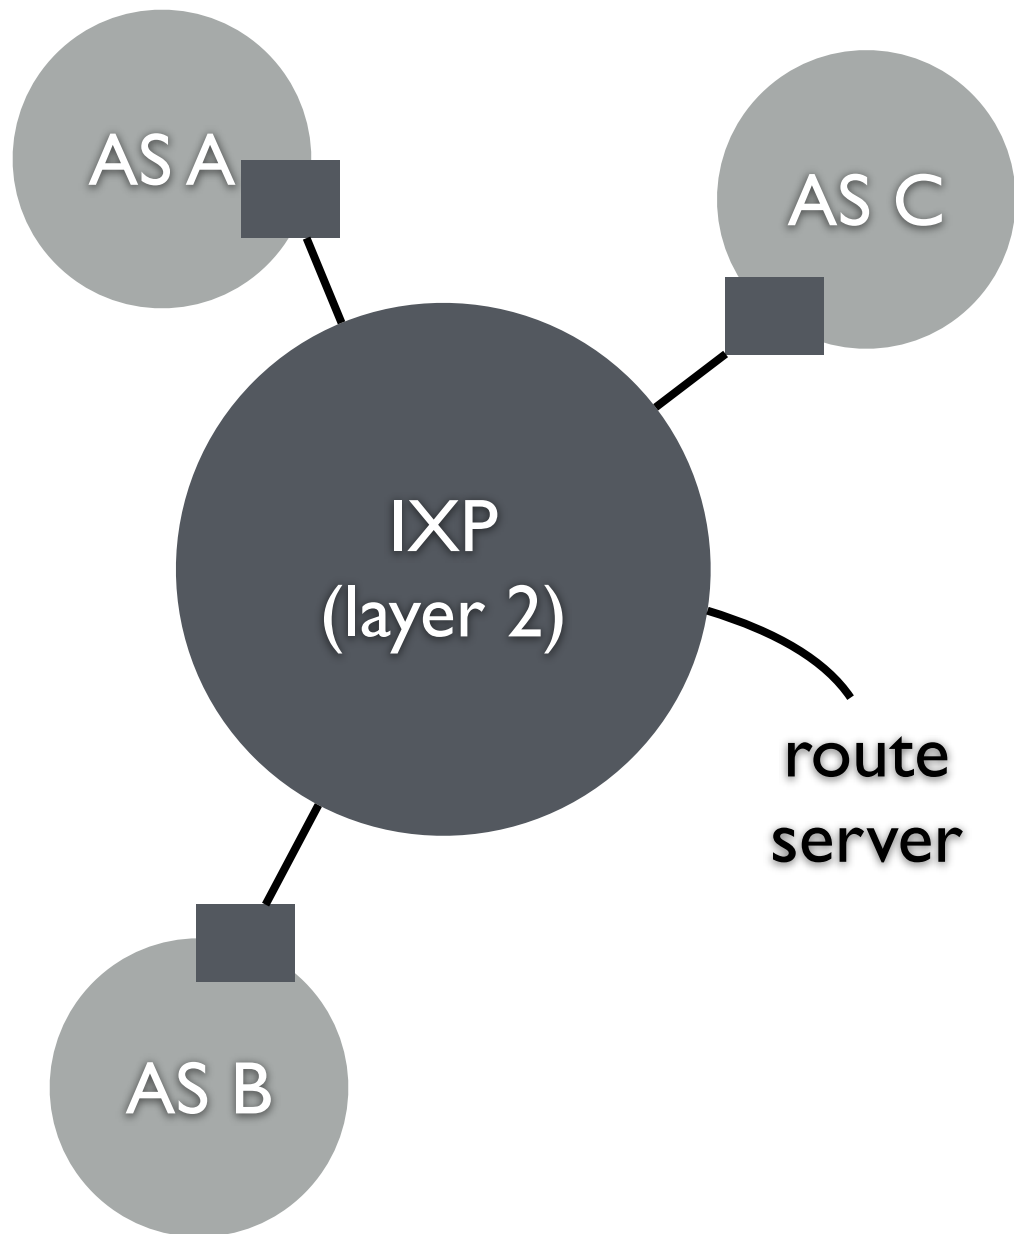
- limited to intradomain: by definition, an SDN controller has purview of the entire network
- SDN is only a platform, not the solution

SDX goals/challenges


- compelling applications
- programming abstraction
- scalable operation
- realistic deployment

programming abstraction

traditional IXP




SDX

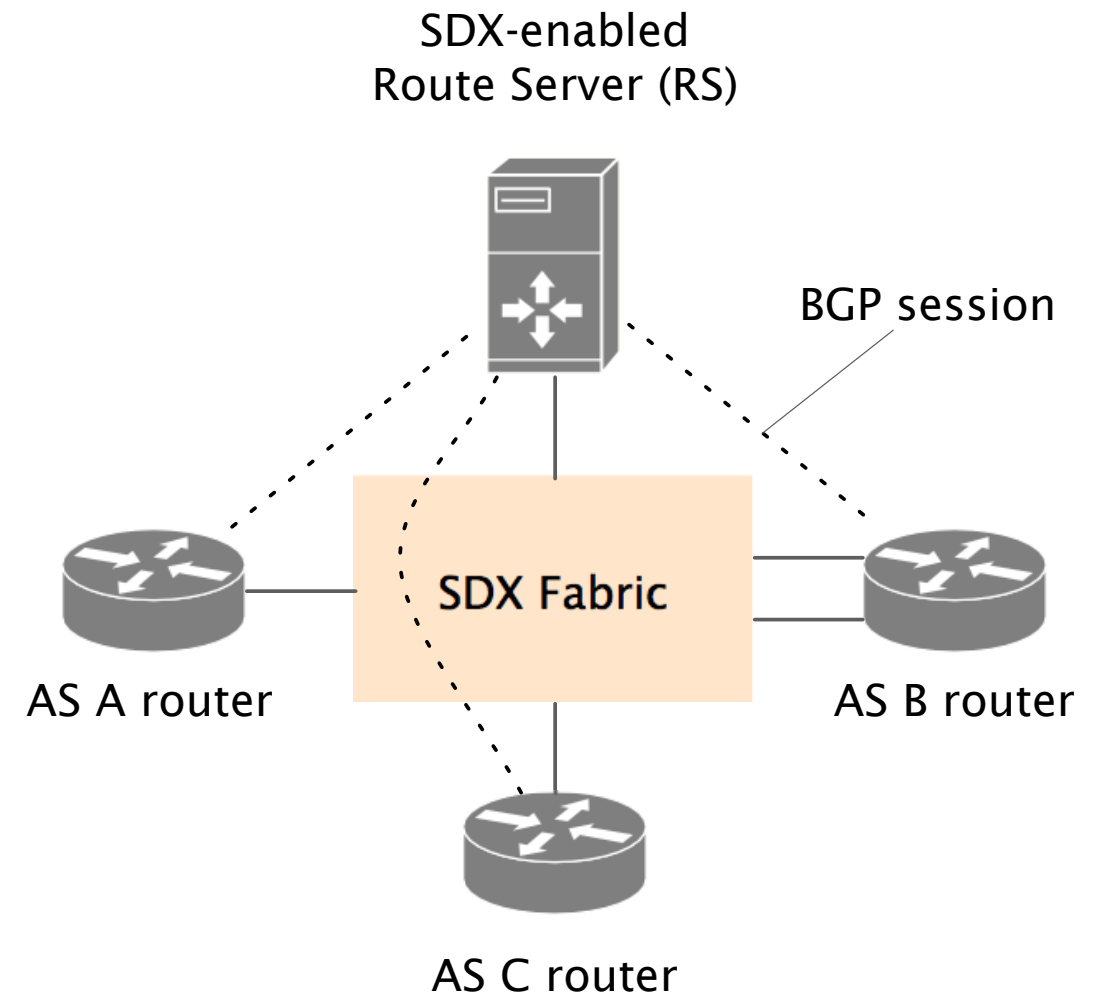


programming abstraction

traditional IXP



SDX

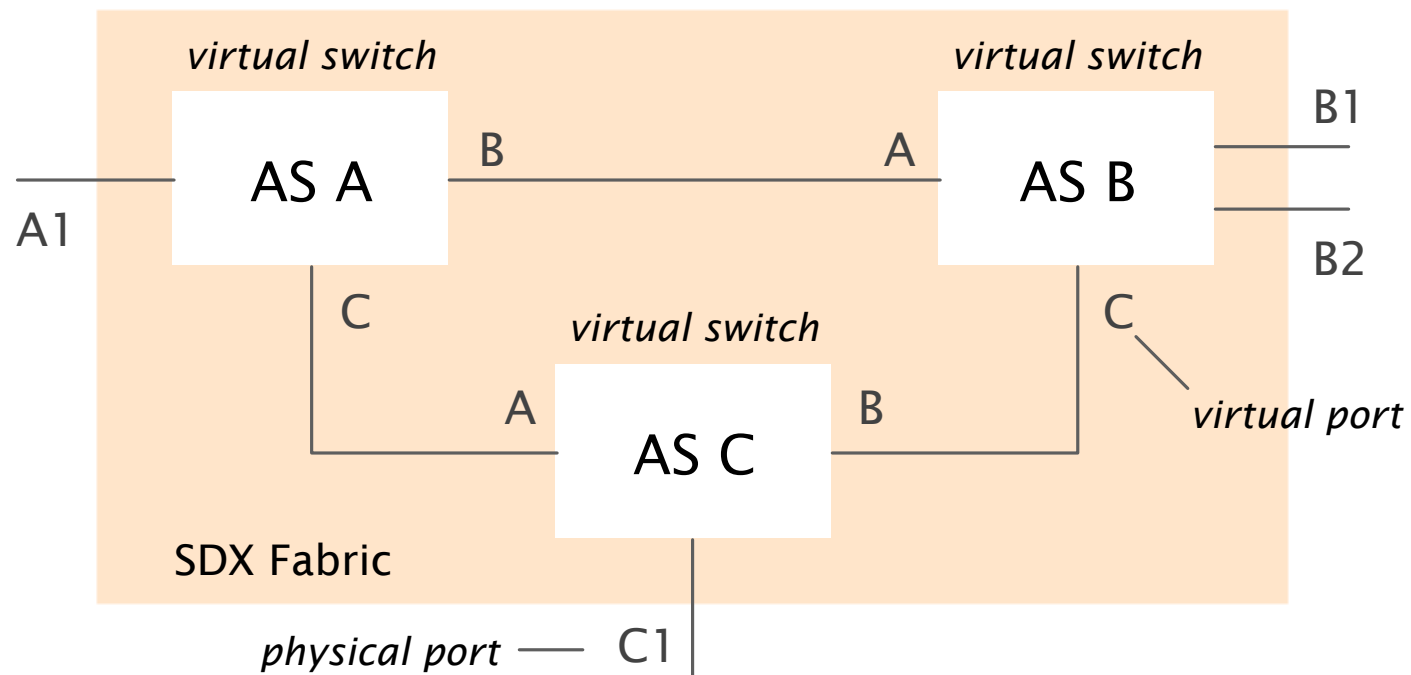


- virtual SDX switch
- gives each AS the illusion of its own virtual SDN switch

virtual SDX switch abstraction

SDX

- gives each AS the illusion of its own virtual SDN switch




virtual SDX switch abstraction

AS A's outbound policy:
application-specific peering

```
(match(dstport=80) >> fwd(B)) +  
(match(dstport=443) >> fwd(C))
```

application specific
peering




virtual SDX switch abstraction

AS B's inbound policy:
traffic engineering

```
(match(srcip={0/1}) >> fwd(B1)) +  
(match(srcip={128/1}) >> fwd(B2))
```

inbound traffic engineering

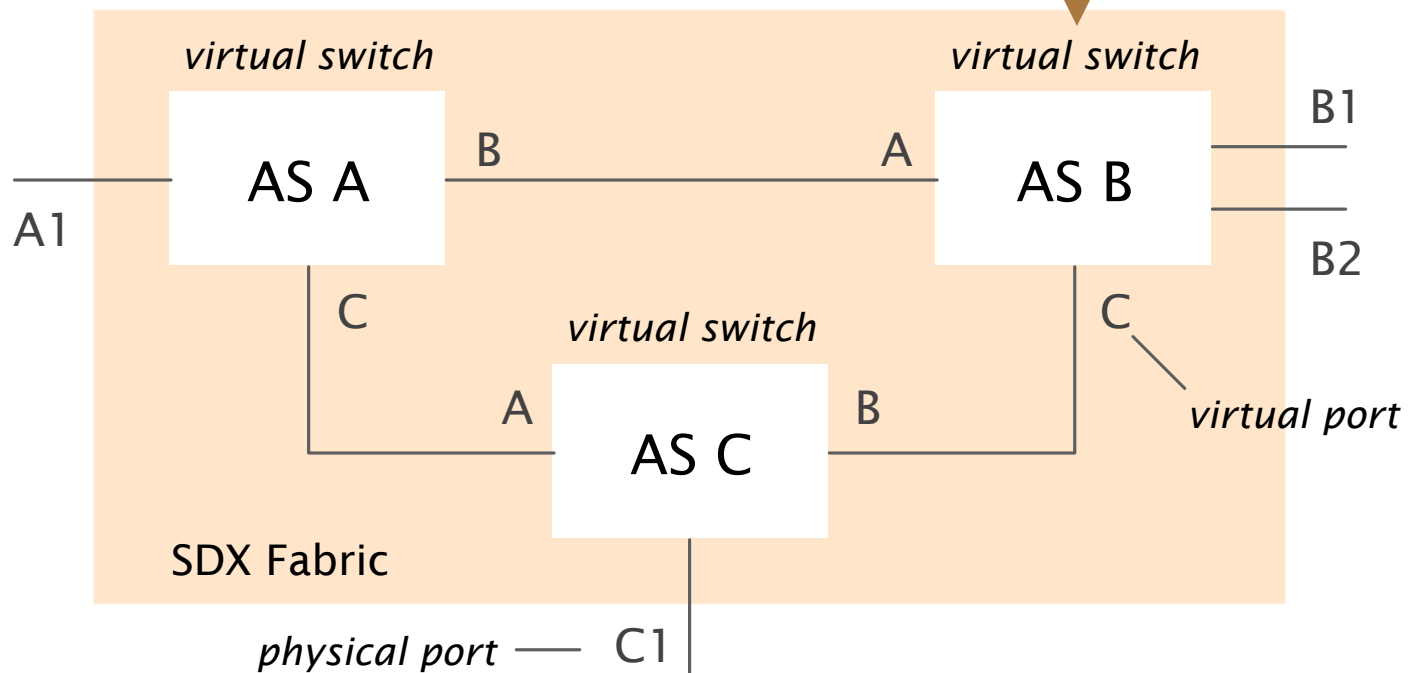


virtual SDX switch abstraction

SDX compiles As's
outbound policy with
B's inbound policy

```
(match(srcip = {0.0.0.0/1}) >> fwd(B1)) +  
(match(srcip = {128.0.0.0/1}) >> fwd(B2))
```

```
(match(dstport = 80) >> fwd(B)) +  
(match(dstport = 443) >> fwd(C))
```



```
(match(port=A1, dstport=80,  
srcip={0.0.0.0/1}) >> fwd(B1)) +  
(match(port=A1, dstport=80,  
srcip={128.0.0.0/1}) >> fwd(B2))
```

virtual SDX switch abstraction

```
match (dstip=74.125.1.1) >>  
  (match (srcip=96.25.160.0/24) >>  
    mod (dstip=74.125.224.161)) +  
  (match (srcip=128.125.163.0/24) >>  
    mod (dstip=74.125.137.139))
```


wider area load
balancing

integration with interdomain routing

BGP routes for RS

prefix	received
p1	C , B
p2	C , B
p3	B , C
p4	C
p5	A

elected routes



SDX route server

- maintains routes on behalf of all participants
- overriding default BGP routes
- forwarding only along BGP-advertised routes
- grouping traffic based on BGP attributes