# Policy-preserving Middlebox Placement in SDN-Enabled Data Centers

Bin Tang
Computer Science Department

California State University Dominguez Hills

Some slides are from
www.cs.berkeley.edu/~randy/Courses/CS268.F08/lectures/22-policy_switching.ppt, and
*www.cs.yale.edu/homes/yu-minlan/talk/sigcomm13.**ppt**x*

# Overview

- What is middlebox?
- What is SDN (Software Defined Network) and NFV (Network Function Virtulization)?
- Policy-preserving middlebox placement problem in data centers
  - Problems and preliminary solutions
- Conclusions

# Middleboxes

- A **middlebox,** or network appliance, is a computer networking device that transforms, inspects, filters, or otherwise manipulates traffic for purposes other than packet forwarding.
  - Intermediaries in-between the communicating hosts
  - Often without knowledge of one or both parties
- Examples
  - Network address translators
  - Firewalls
  - Load balancers
  - Intrusion detection systems
  - Transparent Web proxy caches

# Problem: Middleboxes are hard to deploy

- Place on network path


- On path placement fails to achieve

| | |
|---|---|
| Flexibility | (Re)configurable network topology |
| Efficiency | No middlebox resource wastage |
| Correctness | Guaranteed middlebox traversal |

# Common data center topology

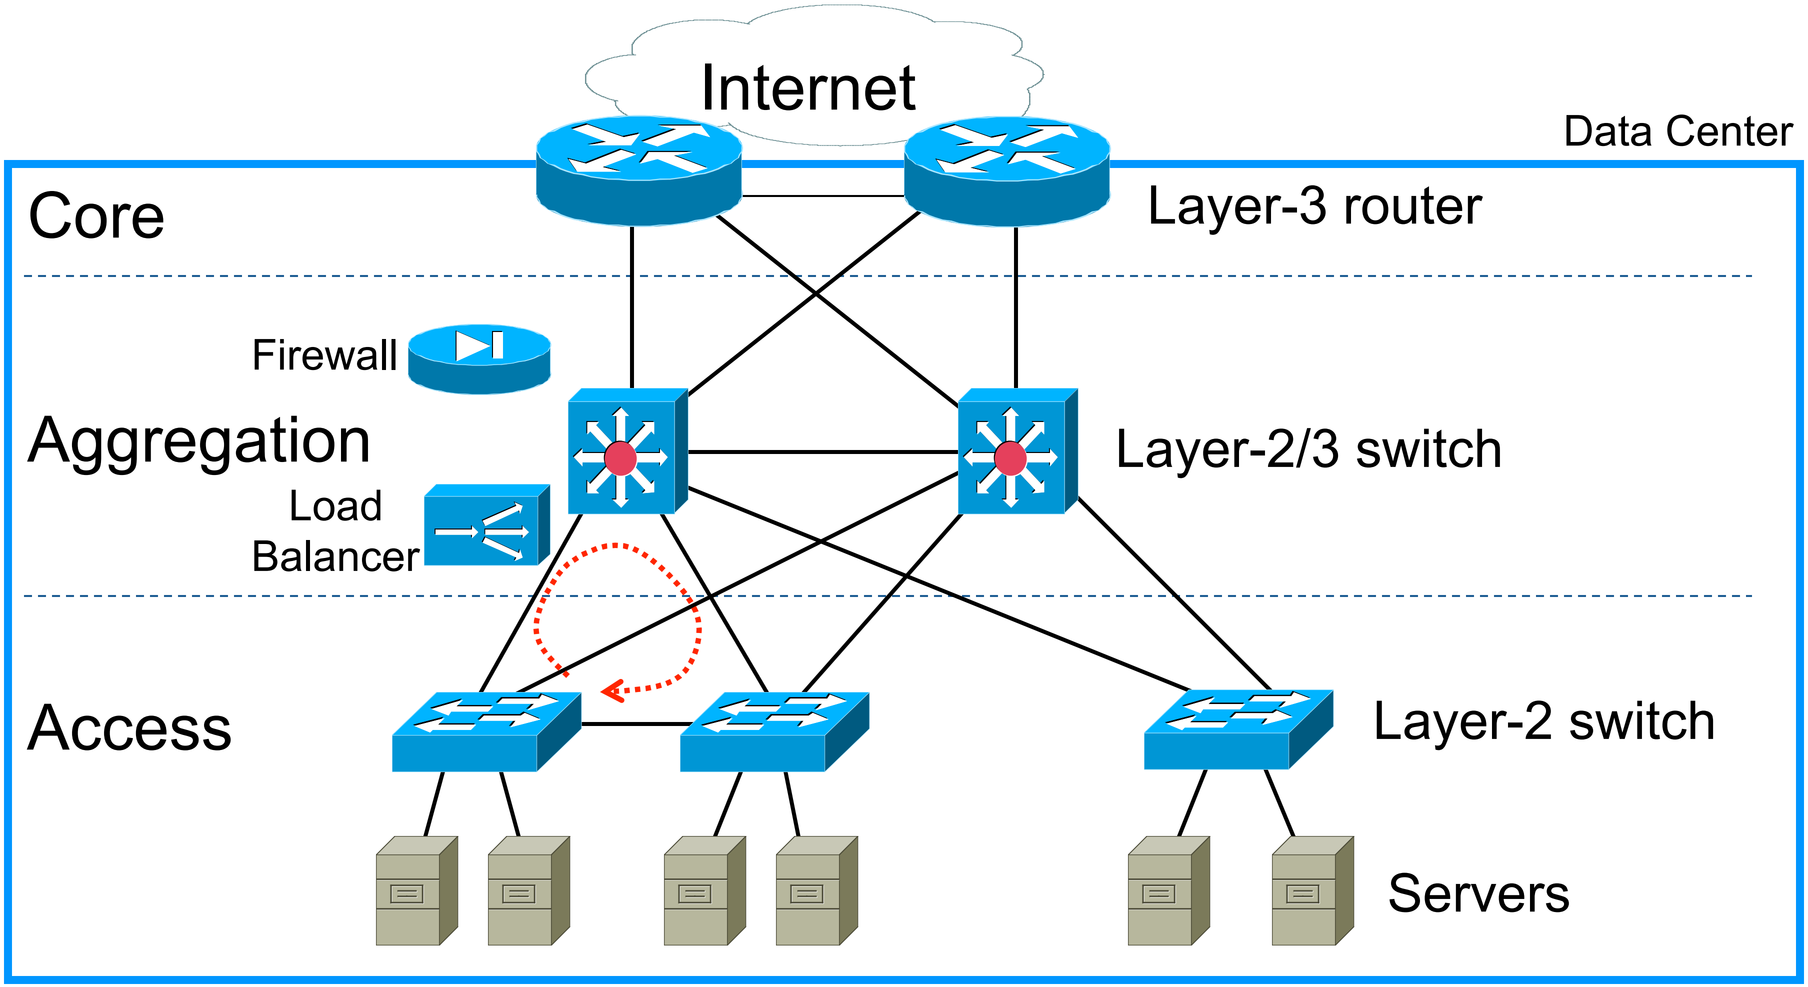

# Inflexible topology

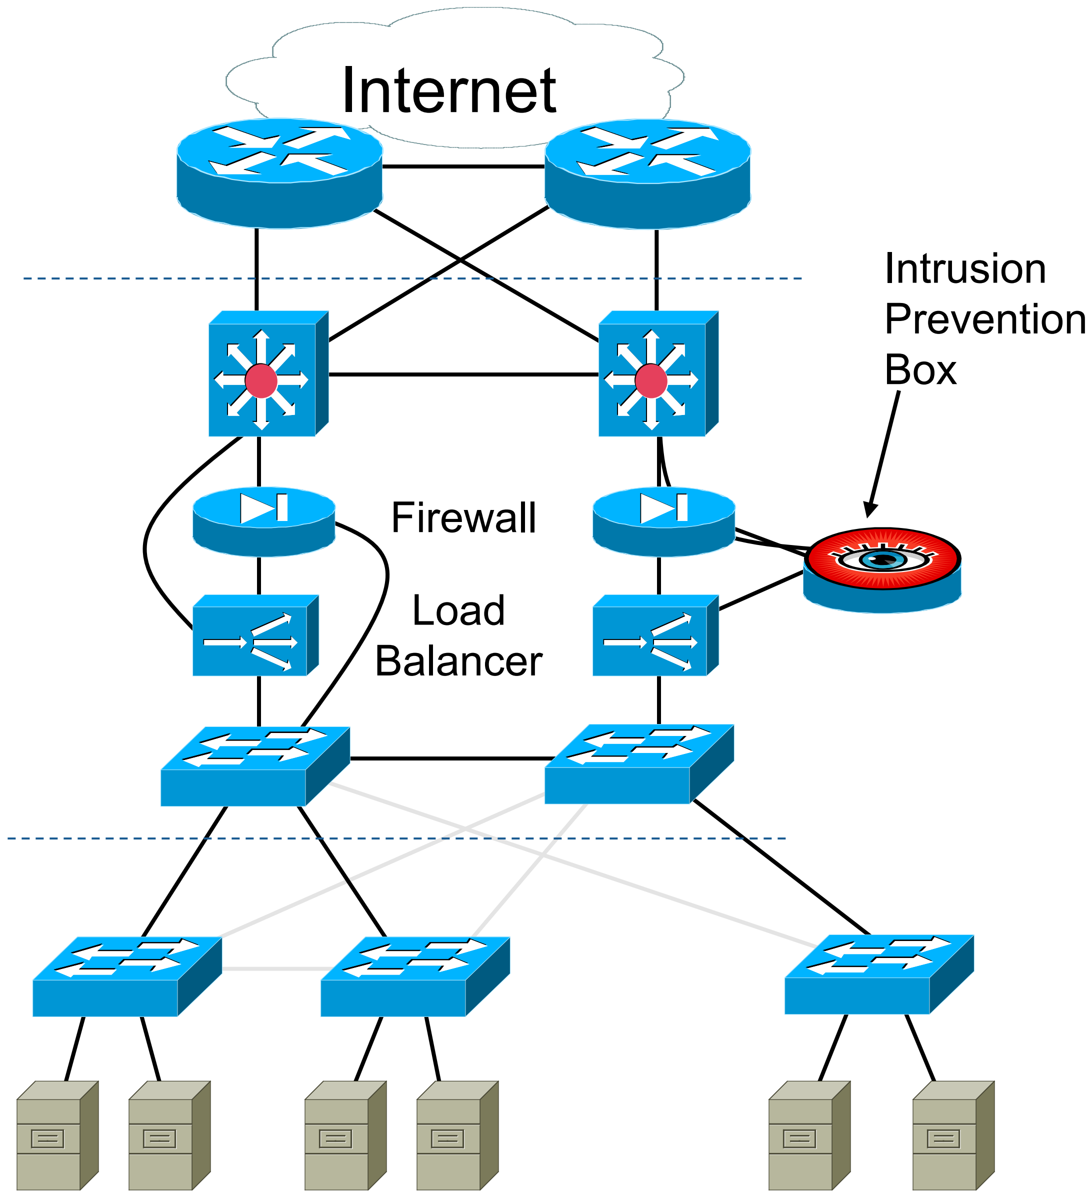

# Inefficient - middlebox resource wastage

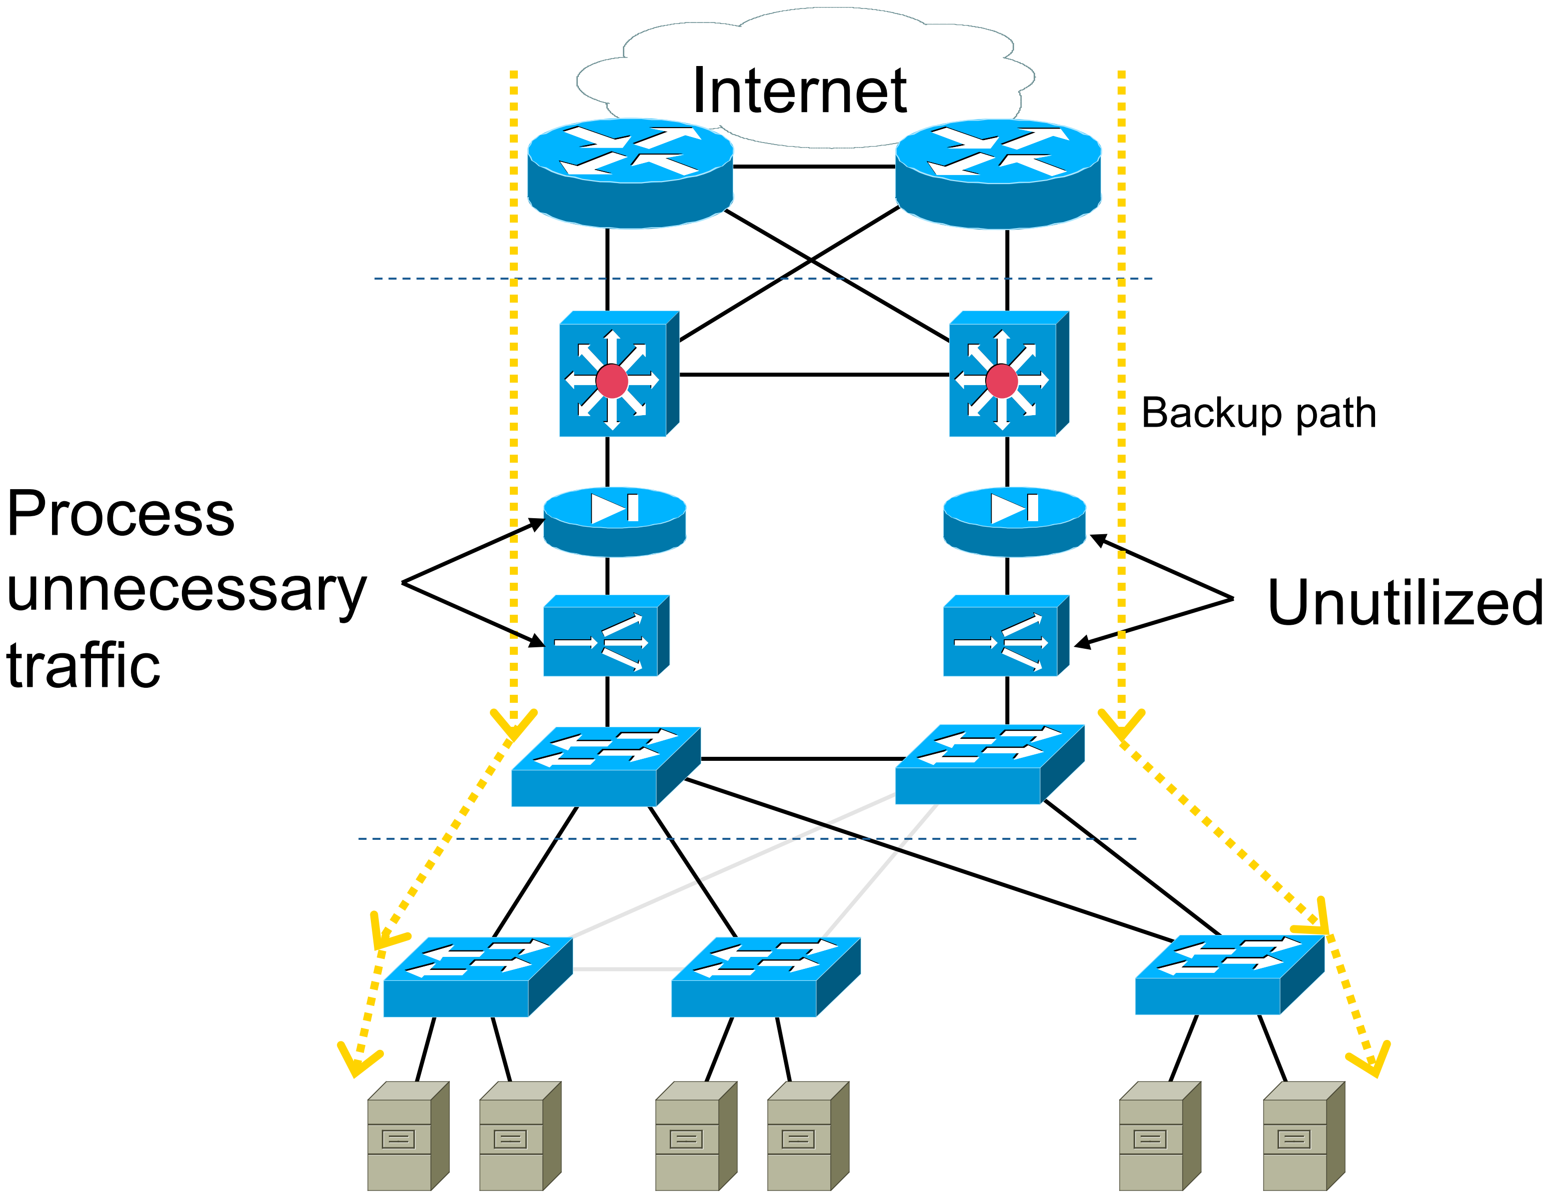

# Policy-Preserving of MBs

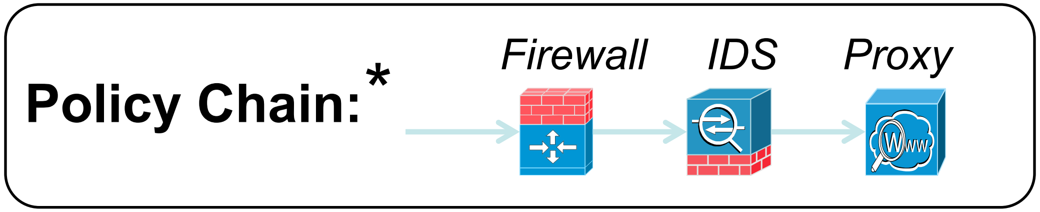


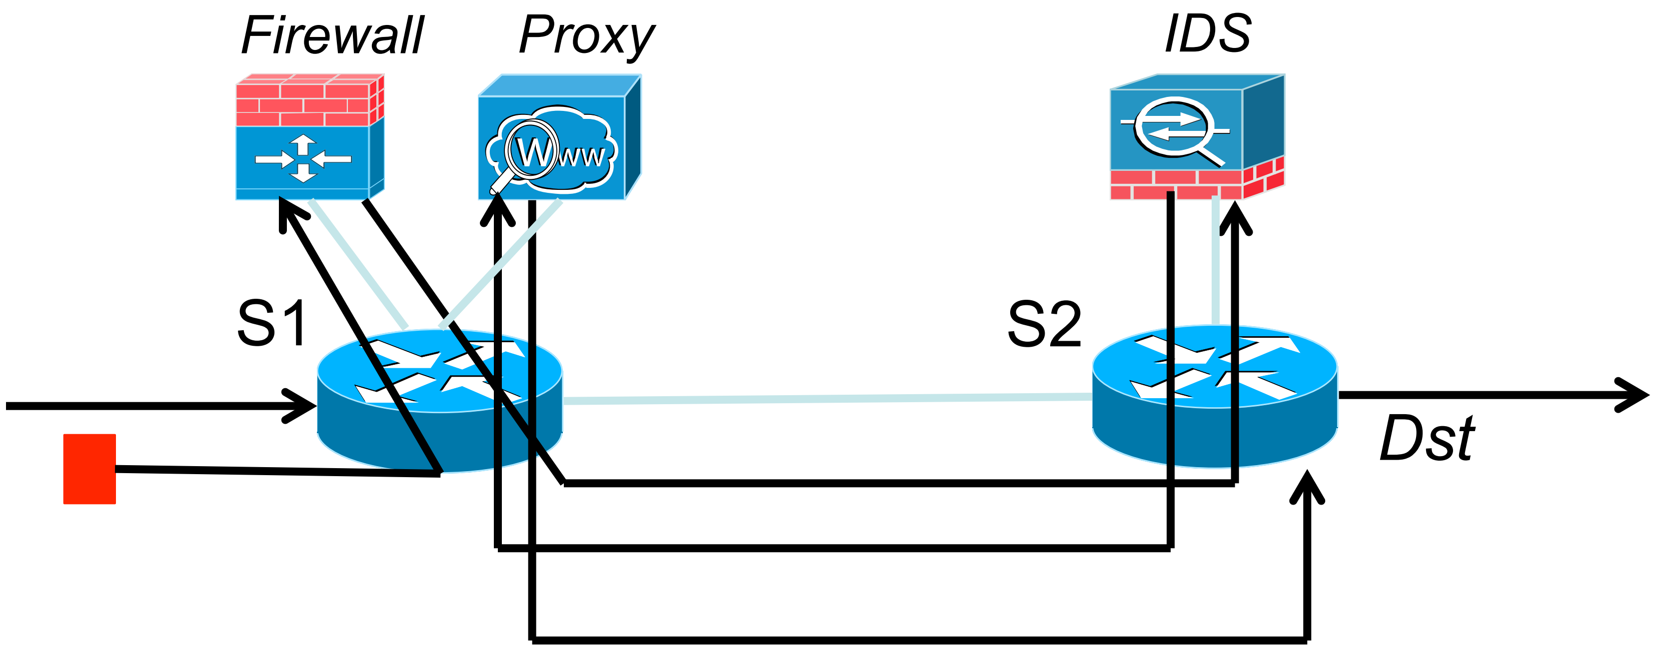

# The Internet: A Remarkable Story

- Tremendous success
  - From research experiment to global infrastructure
- Brilliance of under-specifying
  - Network: best-effort packet delivery
  - Hosts: arbitrary applications
- Enables innovation in applications
  - Web, P2P, VoIP, social networks, virtual worlds
- But, change is easy only at the edge… ☹

# Inside the 'Net: A Different Story…

- Closed equipment
  - Software bundled with hardware
  - Vendor-specific interfaces
- Over specified
  - Slow protocol standardization
- Few people can innovate
  - Equipment vendors write the code
  - Long delays to introduce new features

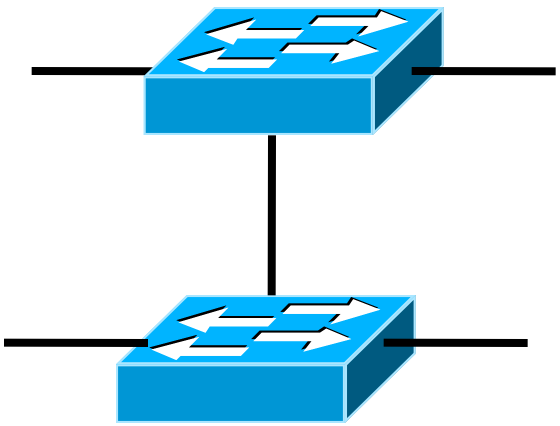

**Impacts performance, security, reliability, cost…**

# Networks are Hard to Manage

- Operating a network is expensive
  - More than half the cost of a network
  - Yet, operator error causes most outages
- Buggy software in the equipment
  - Routers with 20+ million lines of code
  - Cascading failures, vulnerabilities, etc.
- The network is "in the way"
  - Especially a problem in data centers
  - … and home networks

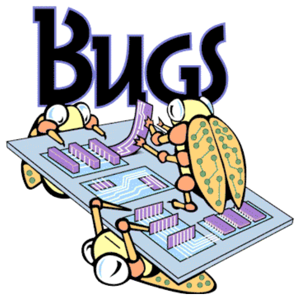


# Traditional Computer Networks

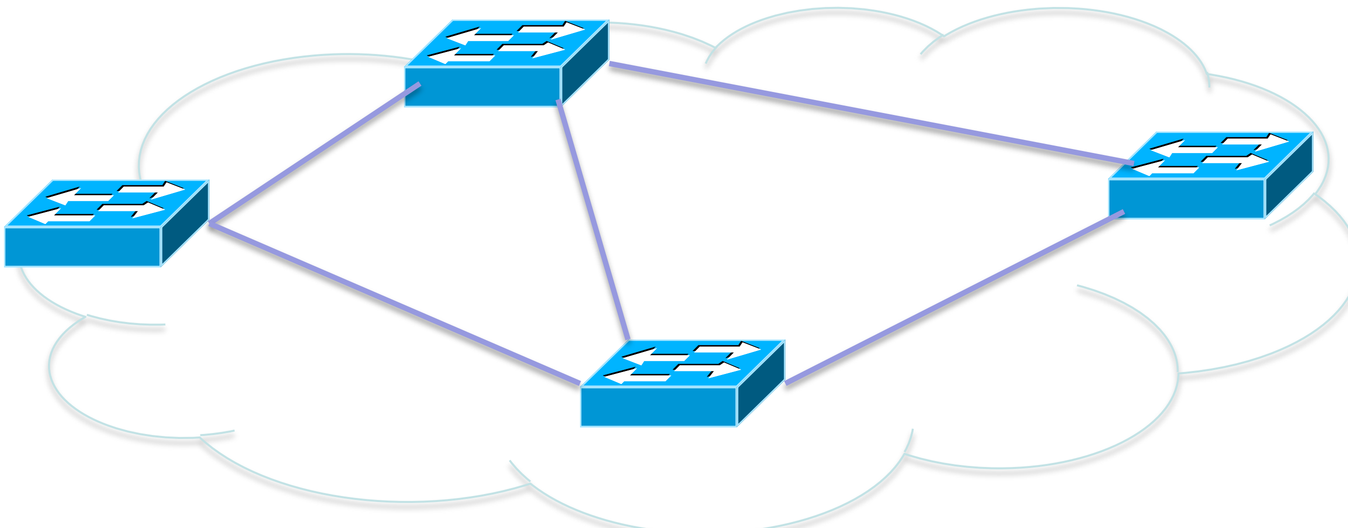

**Data plane:**
**Packet streaming**

Forward, filter, buffer, mark, rate-limit, and measure packets

# Traditional Computer Networks

**Control plane:**
**Distributed algorithms**

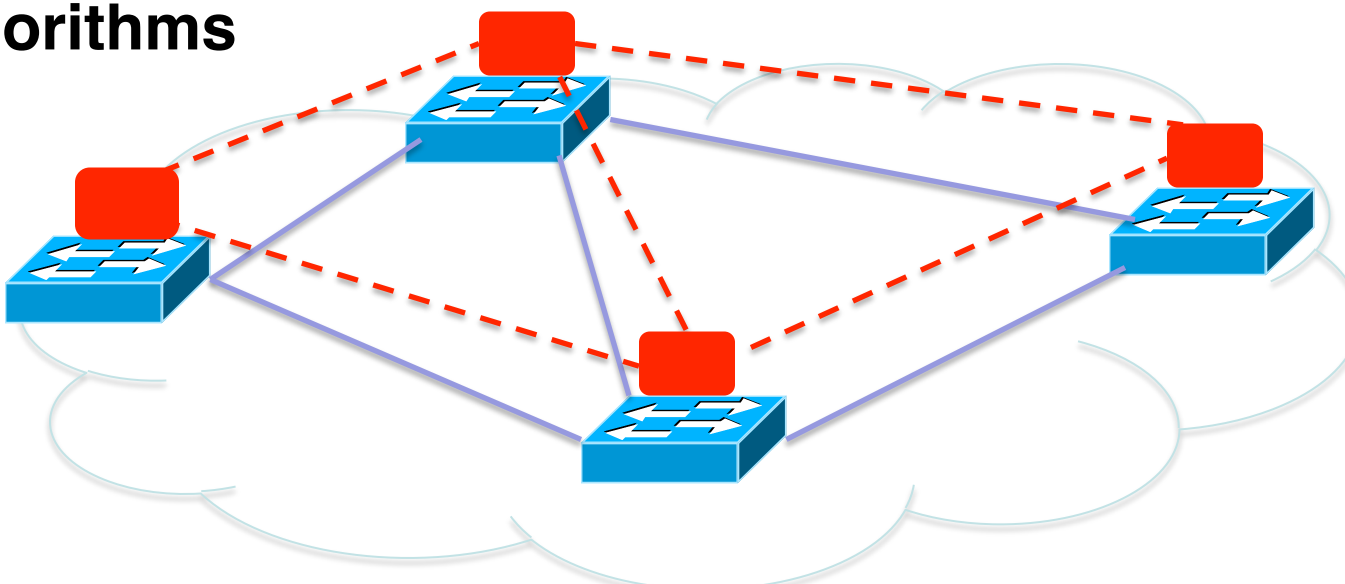

Track topology changes, compute routes, install forwarding rules

# Software Defined Networking (SDN)

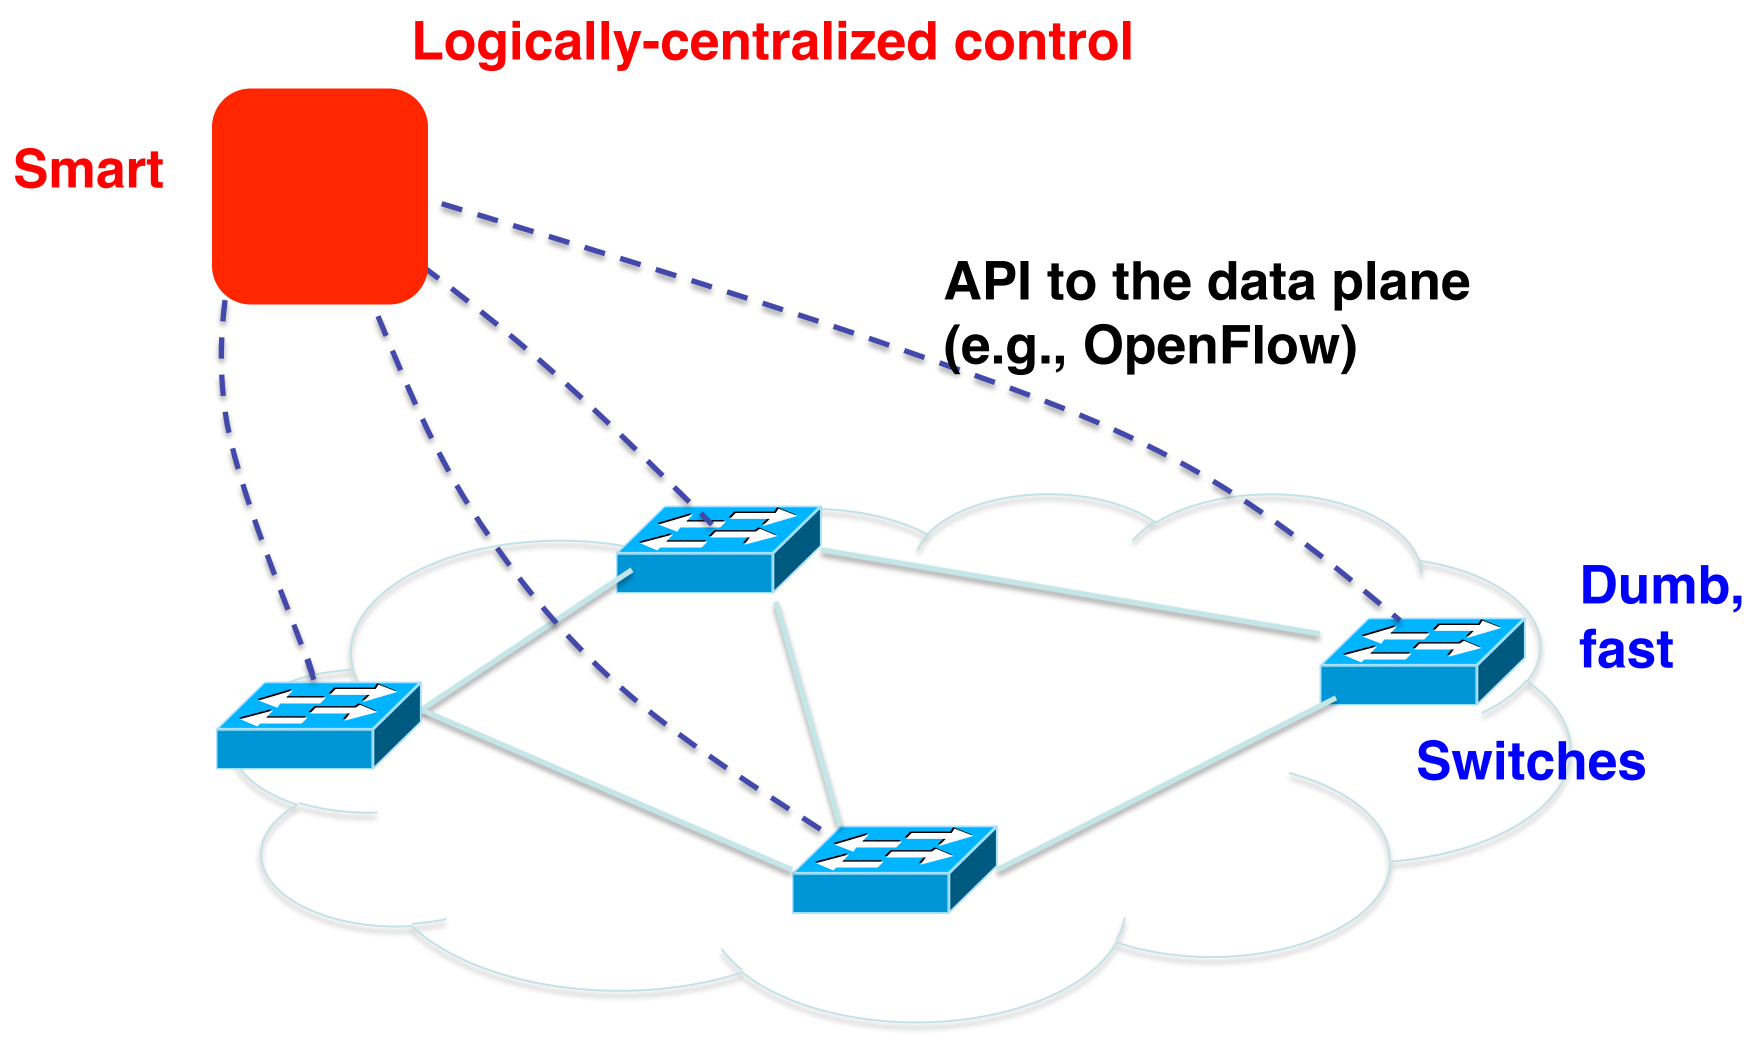

# 3 Complementary but Independent Networking Developments

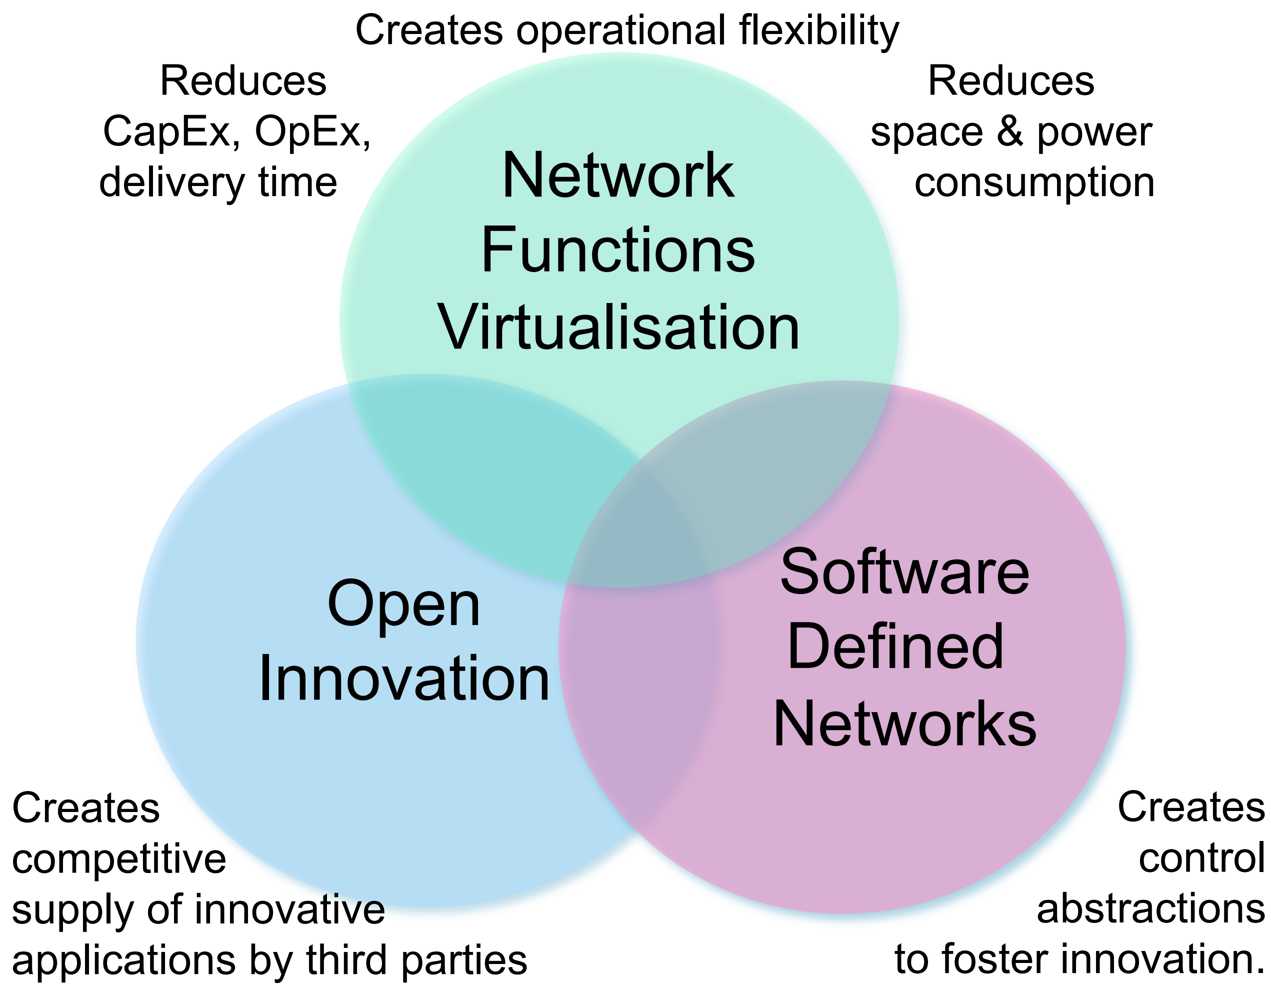

# Network Functions Virtualisation: Vision

## Classical Network Appliance Approach

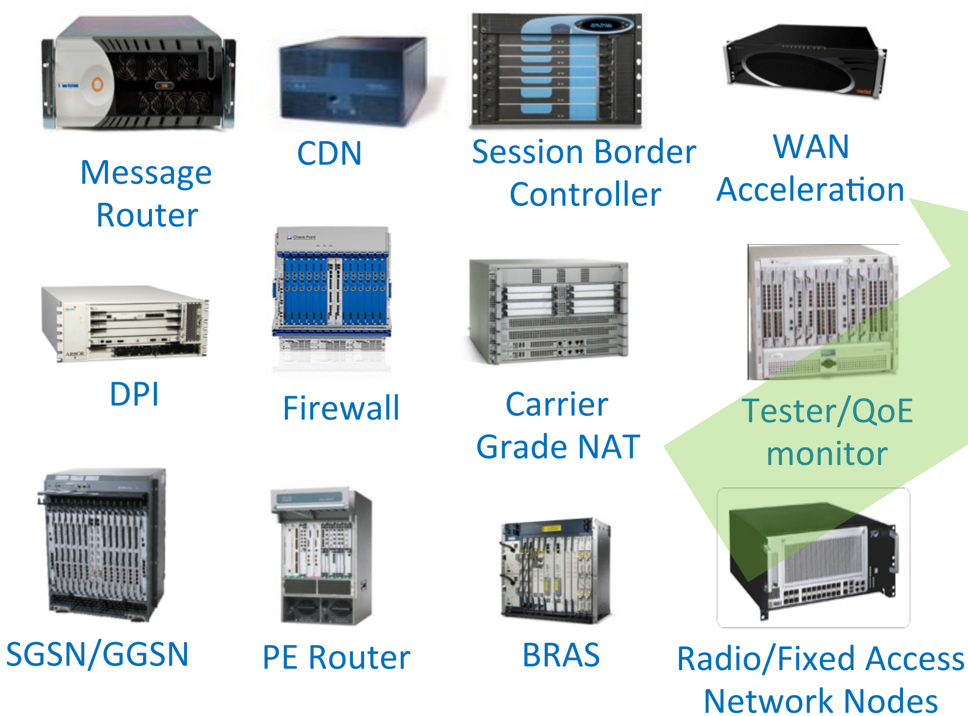


- Fragmented, purpose-built hardware.
- Physical install per appliance per site.
- Hardware development large barrier to entry for new vendors, constraining innovation & competition.

## Network Functions Virtualisation Approach

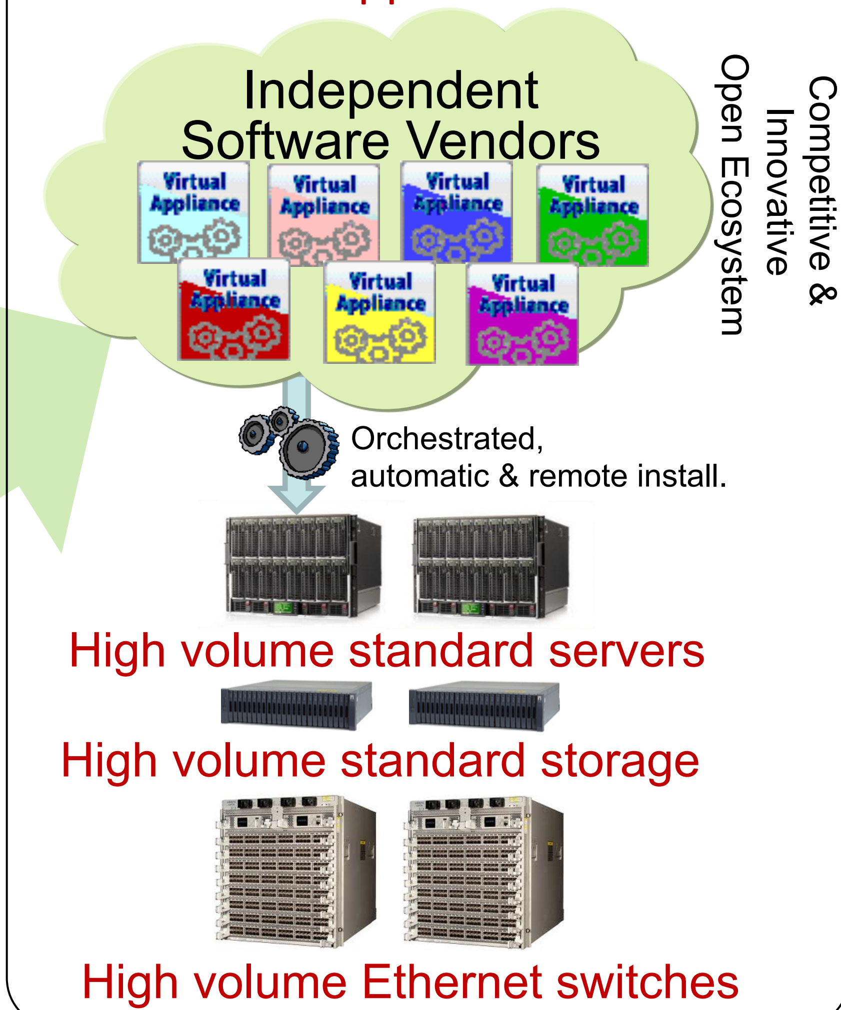

# Policy-Preserving MB Placement Problem in Data Centers

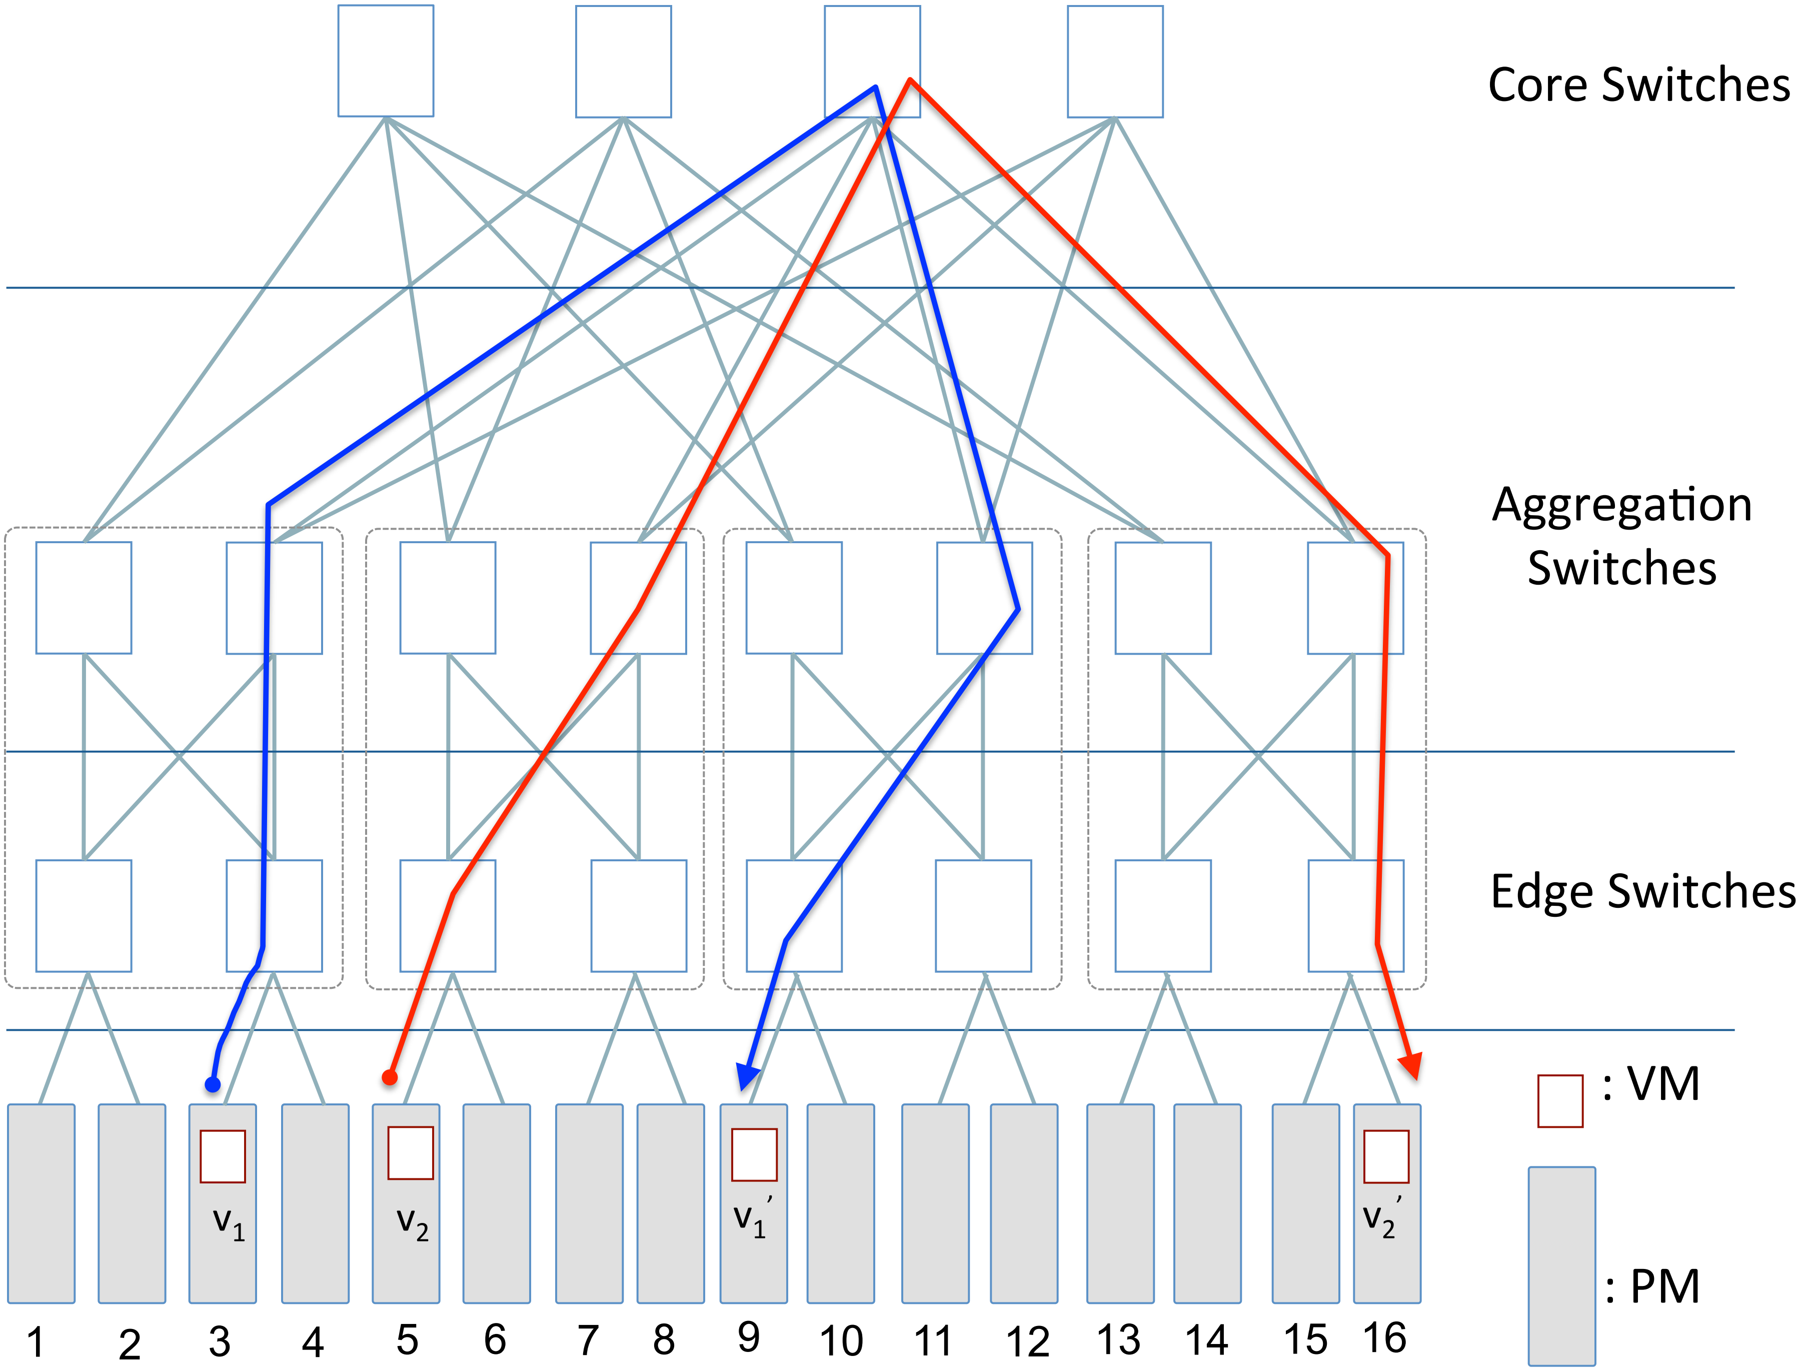
Core Switches
Aggregation Switches
Edge Switches
: VM
: PM
$v_1$
$v_2$
$v_1'$
$v_2'$
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16

# MB Placement Problems

- Many communication pairs in the network
- Single MB Type
  - One MB type, say firewall, but multiple instances
- Multiple MBs Type
  - each has one instance
  - Ordered Service Chaining
  - Unordered Server Chaining
- Goal: Minimize total communication cost
- Constraint: Capacity of MB (each can only process limited number of pairs)

# Single MB Case

- Given a data center graph $G(V,E)$
- There are m instances of a MB, placed at different node in V
- A set of p communicating node pairs P, each pair (s,t) in P needs to traverse to an instance of a MB
- Each middlebox can only be traversed by at most k pairs
- When p = (s,t) traverses an MB instance m, its cost c(p,m) = d(s,sw(m) ) + d(sw(m),t)
- Goal: assign all the pairs in P, each traverses one MB instance, s.t. the total cost is minimized, subject to that each MB instance takes at most k pairs.

# Solution – minimum cost flow

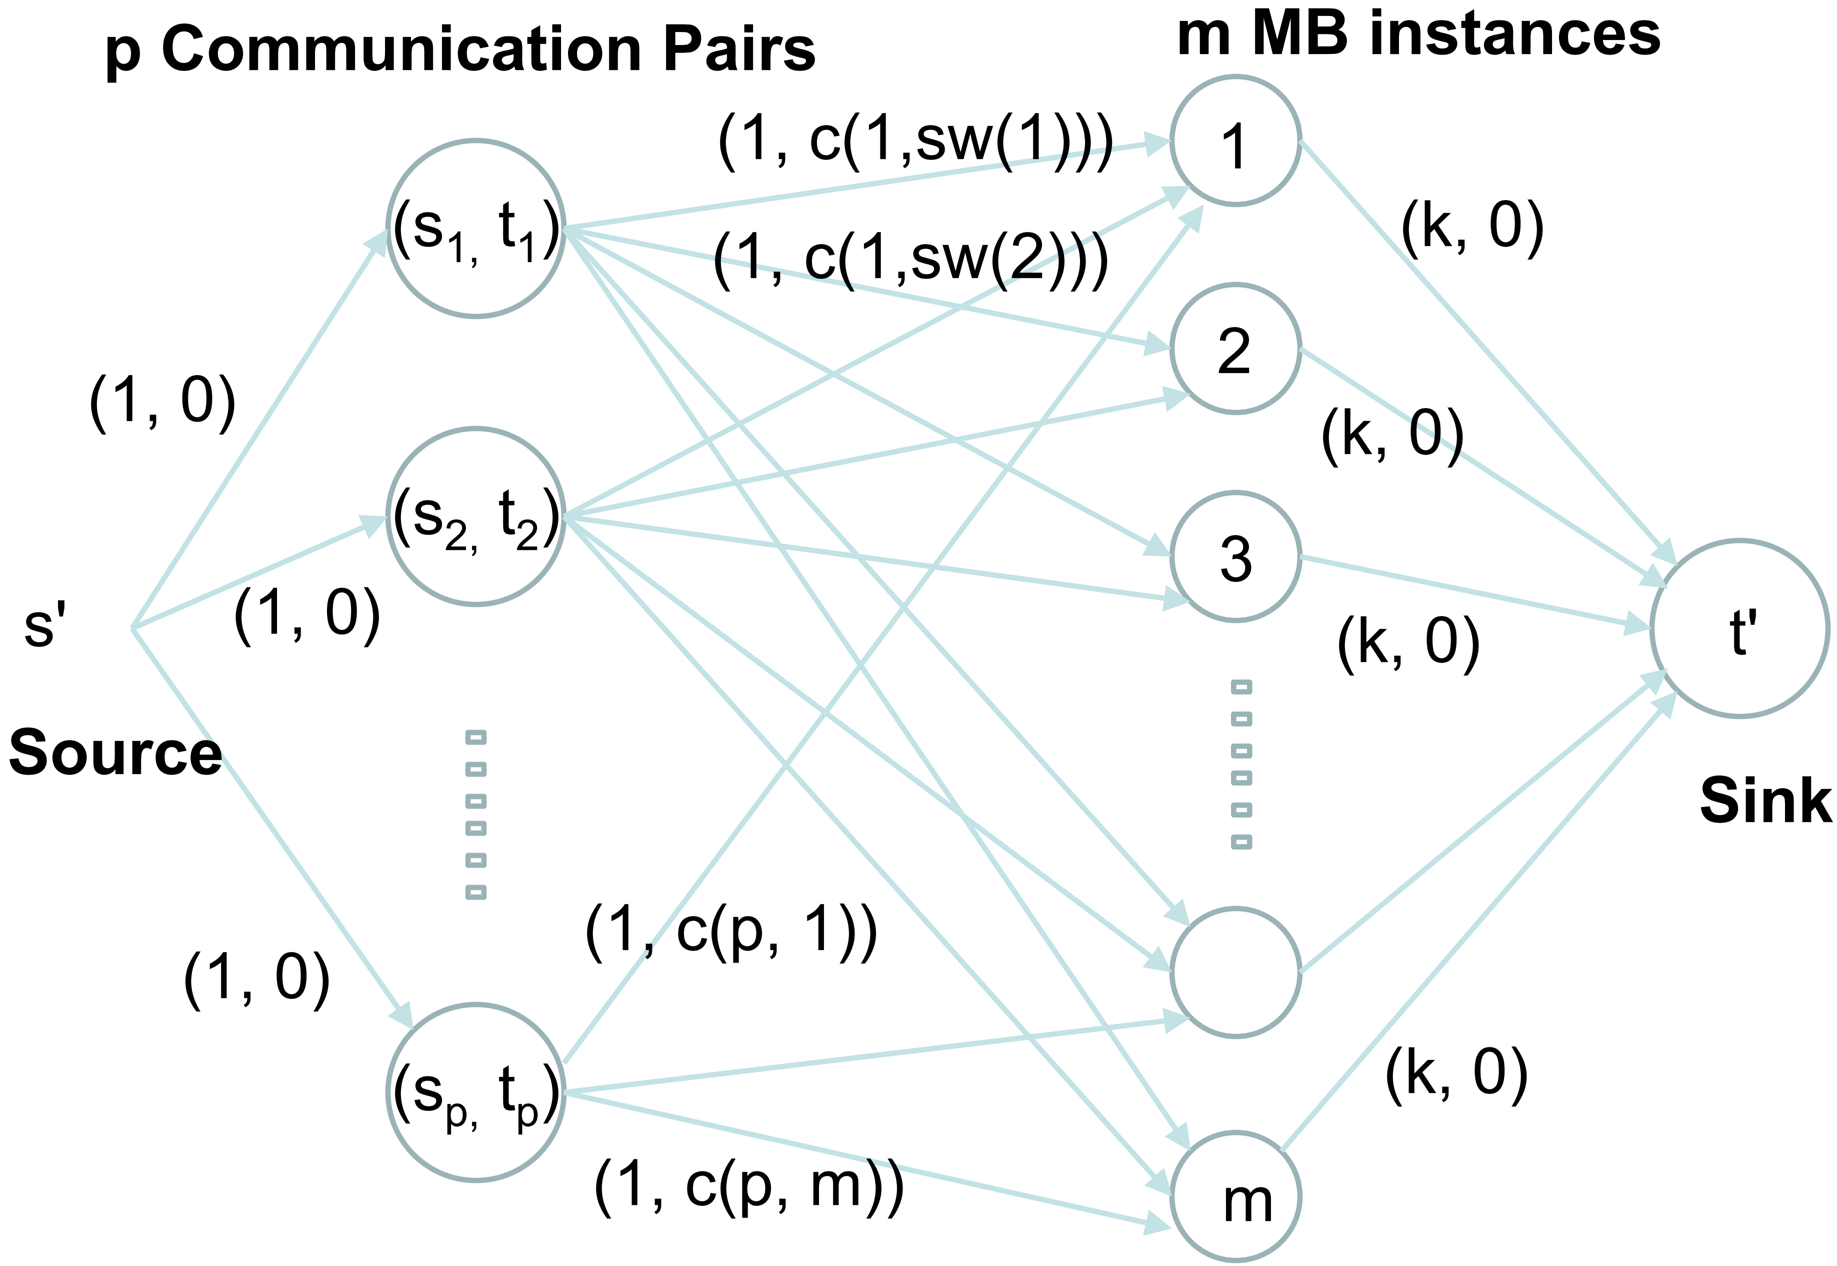
p Communication Pairs
m MB instances
(1, c(1,sw(1)))
(1, c(1,sw(2)))
1
2
3
m
(s1, t1)
(s2, t2)
(sp, tp)
(1, 0)
(1, 0)
(1, 0)
s'
Source
(k, 0)
(k, 0)
(k, 0)
(k, 0)
t'
Sink
(1, c(p, 1))
(1, c(p, m))

# Ordered Multiple MBs Case

- Given a data center graph *G(V,E)*
- There are m MBs $M=\{mb_1, mb_2, \ldots, mb_m\}$ to be placed inside the data center
- A set of p communicating node pairs P, each pair (s,t) in P needs to traverse $mb_1, mb_2, \ldots, mb_m$ in that order
- The cost for p = (s,t) is $c(p) = d(s, mb_1) + d(mb_1, mb_2) + \ldots + d(mb_{m-1}, mb_m) + d(mb_m, t)$
- Goal: where to place the m MBs, s.t. the total cost of all p pairs is minimized

# Ordered Multiple MBs Case: Solution

- NP-hard
- Random: randomly place the m MBs inside the data center
- Greedy: takes place in m rounds
  - In round i, it places $mb_i$ at a node that minimizes the total communication cost so far
- Load Balancing: each switch can only accommodate limited number of communication pairs

# Un-Ordered Multiple MBs Case

- Given a data center graph $G(V,E)$
- There are m MBs $M=\{mb_1, mb_2, \ldots, mb_m\}$ $_{to}$ to be placed inside the data center
- A set of p communicating node pairs P, each pair (s,t) in P needs to traverse $mb_1, mb_2, \ldots, mb_{m\,,}$ but not necessarily in that order
- The cost for p = (s,t) is c(p) = $d(s, mb_{i,1}) + d(mb_{i,1}, mb_{i,2}) + \ldots + d(mb_{i,m-1}, mb_{i,m}) + d(mb_{i,m}, t)$
- Goal: where to place the m MBs, s.t. the total cost of all p pairs is minimized

# Un-Ordered Multiple MBs Case: Solution

- Even more complicated that Ordered Multiple MB case

# MB Migration Problems

- Many communication pairs in the network
- Move MBs from their initial location to other locations
- Goal: Minimize total communication cost
- Constraint: Capacity of MB (each can only process limited number of pairs)

# MB Replication Problems

- Many communication pairs in the network
- Multiple MB types, each has one instance
- Goal: How to replicate the MBs, in order to minimize total communication cost
- Constraint: Capacity of switch (each can only store limited number of MB instances)

# Conclusions

- Deploying middleboxes  is hard, but SDN and NFV makes it easier
- Middleboxes management in SDN-enabled data center is a new and exciting research fields
- Many new algorithmic problems that have not been solved
- Need your participation!

# Questions?