CS 168 Introduction to the Internet: Architecture and Protocols
Spring 2024 Sylvia Ratnasamy & Rob Shakir Discussion 9:
Software-Defined Networking

1 SDN: Overview

You are a chief network engineer at CS168DB, a company that specializes in databases. Currently, your networking infrastructure still relies on hardware switches that converge on routing with various L2/L3 protocols. You have tens of thousands of servers that need to communicate with each other. As the chief network engineer, you are thinking about making the switch to software defined networking.

1. What are your current control plane abstractions?
   Solution: None!

2. What are the required control plane's abstractions if you make the switch to SDN?
   Solution: There are three. The first is an abstraction of the forwarding model of switches, OpenFlow is an example. The second is an abstraction of network state, in the form of a global network view via a Network Operating System. The third abstraction is an abstraction of the specification for network behavior, which is used to specify the *goals* of an operator on an abstract view of the network.

3. What happens when a link goes down in your new SDN setup?
   Solution: The software switches connected to the link will communicate the topology change to the controller. The controller will then recompute routing state and disseminate that information to the routers.

4. Is OpenFlow equivalent to SDN?
   Solution: No.

2 Traffic Engineering

Consider the following network topology with servers, virtual machines, routers and their associated addresses. Note that all links have a cost of 1 and 100Gbps of bandwidth unless otherwise specified.

Assume the following for traffic engineering (TE):

- All encapsulation for traffic engineering occurs at the servers, though decapsulation occurs as needed at the routers.
- Each router's label for the TE header is Rx where x is its router number (ie, Router 2's label for TE is "R2")
- When TE is enabled, all packets must have a TE header with some label

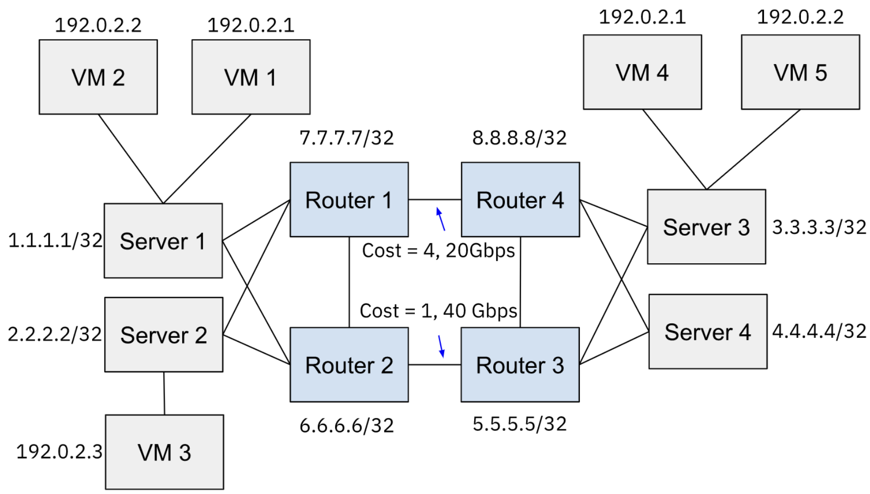

(1) Consider a case with no traffic engineering. Assume that some workloads are running on the network with the following constant traffic between servers (the traffic originally comes from the VMs, but for now we are just considering the aggregate of the traffic from each server).

- Workload I: S1 to S4 = 30Gbps
- Workload II: S2 to S4 = 20Gbps

What route(s) is used for both workloads?

Solution: R2-R3-S4

(2) Is there congestion in the network for the workloads in part (1)? What paths should each workload take to use the resources more efficiently? Assume for now that workloads cannot have their traffic split across different paths.

Solution: Yes there is congestion.

Workload I: S1-R2-R3-S4

Workload II: S2-R1-R4-S4

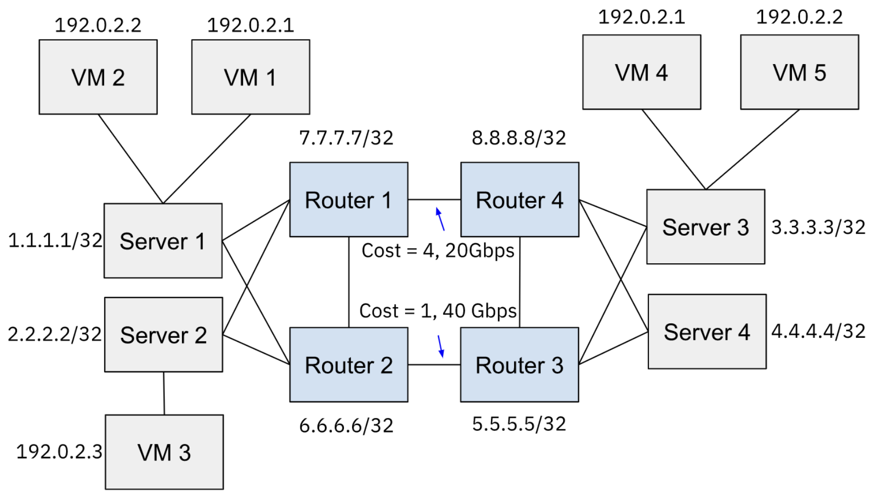

(3) Now, we want to add traffic engineering to implement a policy to enable the route(s) in part (2). Fill in the table with the header fields necessary to enforce the chosen paths. Fill in only the headers that would be seen at each location on the path specified.

Workload	Location	TE Header Label	Destination IP
I	Leaving Server 1		
I	Arriving at Server 4		
II	Leaving Server 2		
II	Arriving at Server 4		

Solution:

Workload	Location	TE Header Label	Destination IP
I	Leaving Server 1	R2 or R3	4.4.4.4
I	Arriving at Server 4		4.4.4.4
II	Leaving Server 2	R4	4.4.4.4
II	Arriving at Server 4		4.4.4.4

(4) Assume that the traffic on the network changes. Consider the traffic between VMs rather than servers.

- Workload I: VM1 to VM5 = 40Gbps
- Workload II: VM3 to VM5 = 10Gbps

Perform traffic engineering to put each workload on the best path for overall utilization. Fill in the table with the header fields necessary to perform such TE and to implement the necessary network virtualization. Fill in only the headers that would be seen at each location on the path specified.

Workload	Location	TE Label	Underlay Dst Address	Overlay Dst Address
I	Leaving VM 1			
I	Leaving Server 1			
I	Arriving at Server 3			
I	Arriving at VM 2			
II	Leaving VM 3			
II	Leaving Server 2			
II	Arriving at Server 3			
II	Arriving at VM 5			

Solution:

Workload	Location	TE Label	Underlay Dst Address	Overlay Dst Address
I	Leaving VM 1			192.0.2.2
I	Leaving Server 1	R2 or R3	3.3.3.3	192.0.2.2
I	Arriving at Server 3		3.3.3.3	192.0.2.2
I	Arriving at VM 2			192.0.2.2
II	Leaving VM 3			192.0.2.2
II	Leaving Server 2	R4	3.3.3.3	192.0.2.2
II	Arriving at Server 3		3.3.3.3	192.0.2.2
II	Arriving at VM 5			192.0.2.2