Tunnel-based mechanisms for datacenter latency control

Xinpeng Wei

Background: Connections in Datacenter

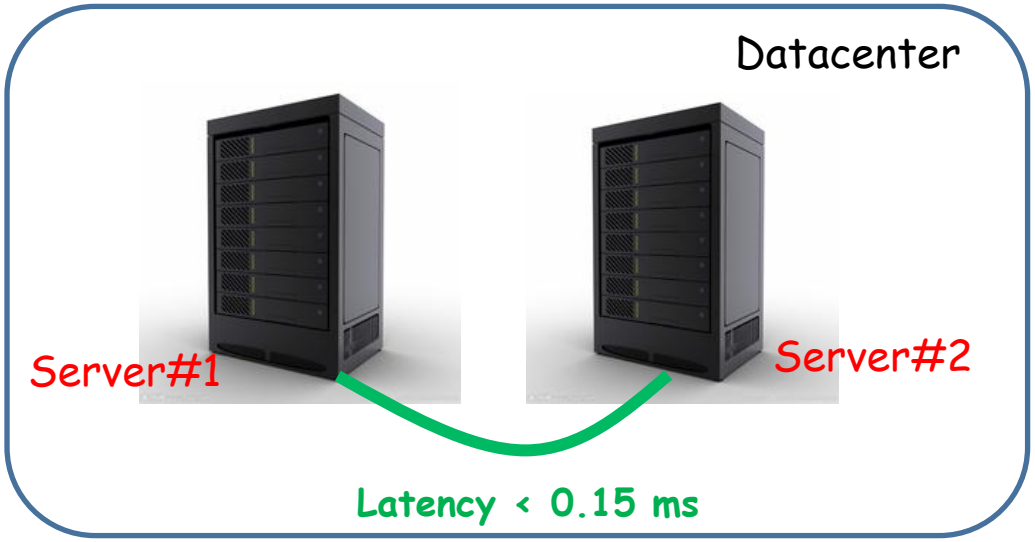

In datacenter, the latency between physical servers is relative small (less than 0.15ms); but in datacenter implementing virtualization, the end-to-end latency is significant, mainly for two reasons:
1) Cost of I/O virtualization (several ms); 2)VM scheduling (tens ms).

The end-to-end latency caused by virtualization is more significant than physical link.

Control the latency induced by virtualization!!!

Background: tunnels in datacenter

Tunnels are widely used in DC scenario!

For example, multi-tenant scenario

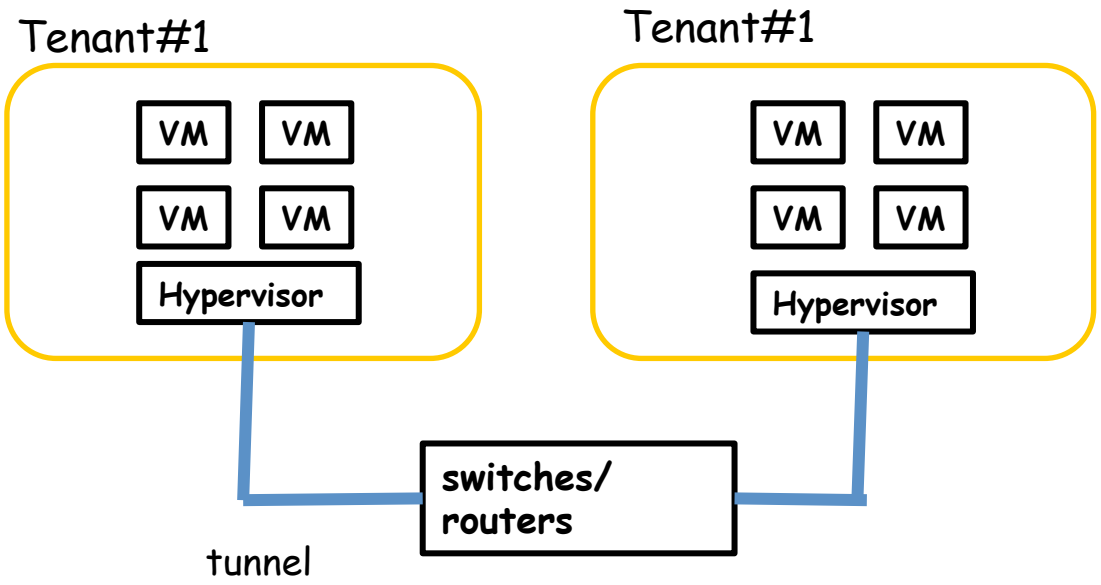

Tunnels end up at hypervisor

Candidate solutions for latency control

- Solution#1: Optimize the performance of I/O virtualization.
- Solution#2: Optimize the performance of VM scheduling mechanism, e.g. by binding VM to physical CPU core.
- Solution#3: Optimize TCP behavior on VM.
- Solution#4: Get hypervisor involved in end-to-end latency control and without or less change to VMs.

Latency control solution

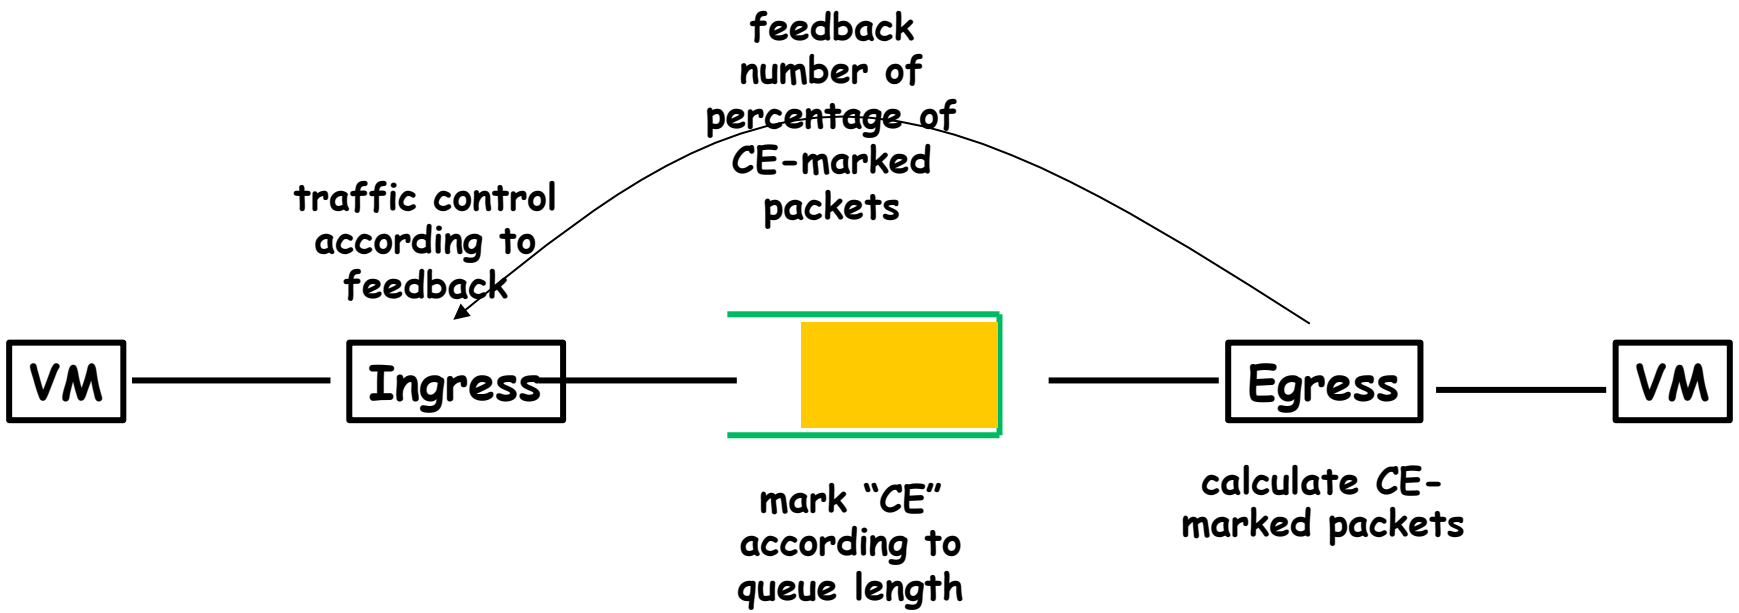

- We assume latency is caused by queue length of router.
- The congestion volume at the egress could be seen as an indication of queue size.

Traffic Control

- Control the traffic entering into the tunnel according to some kinds of policies, for instance,
  - Packet dropping based: drop packets of certain traffic;
  - Packet buffering based: buffer packets of certain traffic to control the end-to-end latency.

Candidate feedback mechanisms

- Extension of tunnel protocol
  - E.g. VXLAN
- Design an dedicated feedback protocol
  - In-band
  - out-band

A simulation (NS3-based)

Ongoing……