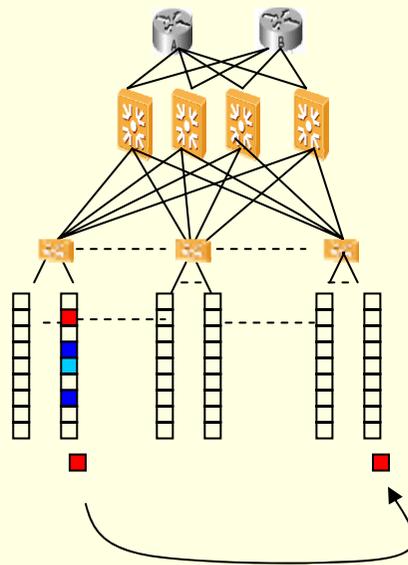

Problem statement on address resolution in VM migration

Yizhou Li (liyizhou@huawei.com)
Nie ShiZhong(nieshzh@ctbri.com.cn)

Assumption for VM migration

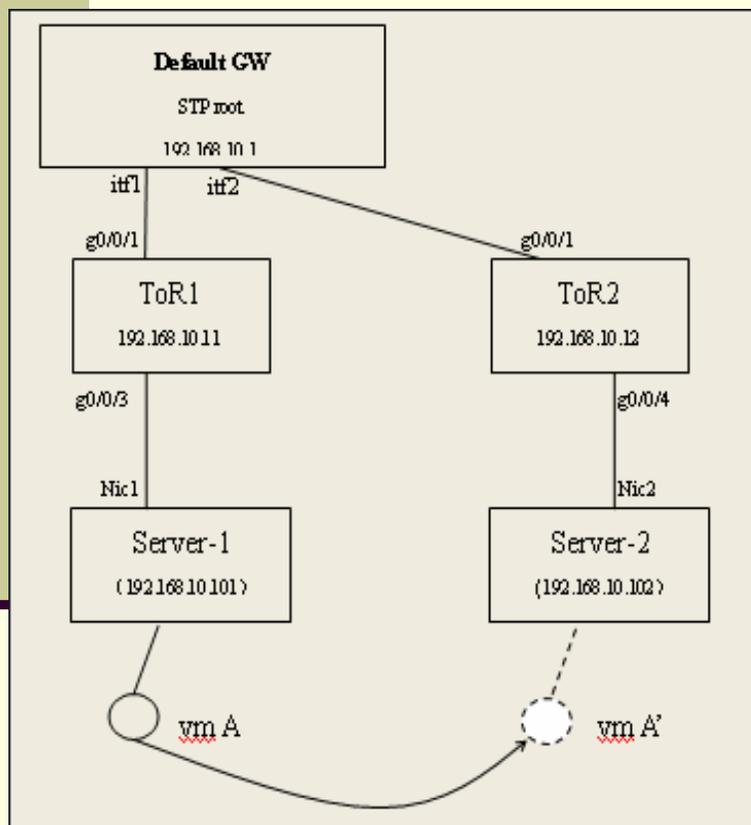
- Same MAC and IP address
- Service provided by VM should not be interrupted.



Problems on address resolution in VM migration

- No signaling message to indicate VM having left server
 - No “ARP leave” message to make other nodes forget the learned address/port
 - Blackhole before entry ageing.
 - Keepalive?
 - De-registration?
 - Management plane (like vCenter) knows vm is migrating. Can it inform network management system?
- Uncertainty of signaling message after VM starts to operate at new location
 - gratuitous ARP request, gratuitous ARP reply, reverse ARP
 - Different implementation to handle different ARP type, result may not be always as we expected (see next slide)
- Difficulty of traffic redirection after migration
 - all switches in the network should be able to correctly send the frame to VM’s new location
 - Redirect packet in cache at old location to new location may be needed?
- Different optimization methods
 - How to tell it is a new VM or migrated one? Blind flooding or optimized multicast?

Examples for VM migration



#	packet sent aft VM migration	Is VM's interface updated to itf2 on GW?
1	std gratuitous ARP	Y
2	broadcast ARP reply	N
3	RARP	N
4	ARP request with GW as target IP	Y
5	ARP request with other host as target IP	N
6	unicast ARP reply with GW as destination	Y
7	unicast ARP reply with other host as destination	N

Security concerns

- Some existing mechanism may not be applicable
 - MAC locking: locking a MAC address to a specific physical port of the switch.
 - DHCP snooping: binding IP/MAC by snooping DHCP ACK to port of switch.