

# VL2: A Scalable and Flexible Data Center Network (SIGCOMM 2009)

Presented by Ankur Dave  
CS294-110, Fall 2015

# Motivation

- Conventional datacenter networks are hierarchical
- Oversubscription forces services to reserve hosts for locality
- Redundant links don't help performance




Figure 1: Common data center interconnect topology. Host to switch links are GigE and links between switches are 10 GigE.

# Goals


- Want *agility*: assign any server to any service
  - For VM migration and demand spikes
- Requirements:
  1. *High capacity* regardless of location
  2. *Flat addressing* for easy migration
- Ideal: Whole datacenter is a big Ethernet

# Solution

1. High capacity: Clos topology + Valiant Load Balancing
2. Flat addressing: Directory service


# Clos Topology

- Provides many paths between any pair of hosts



# Valiant Load Balancing

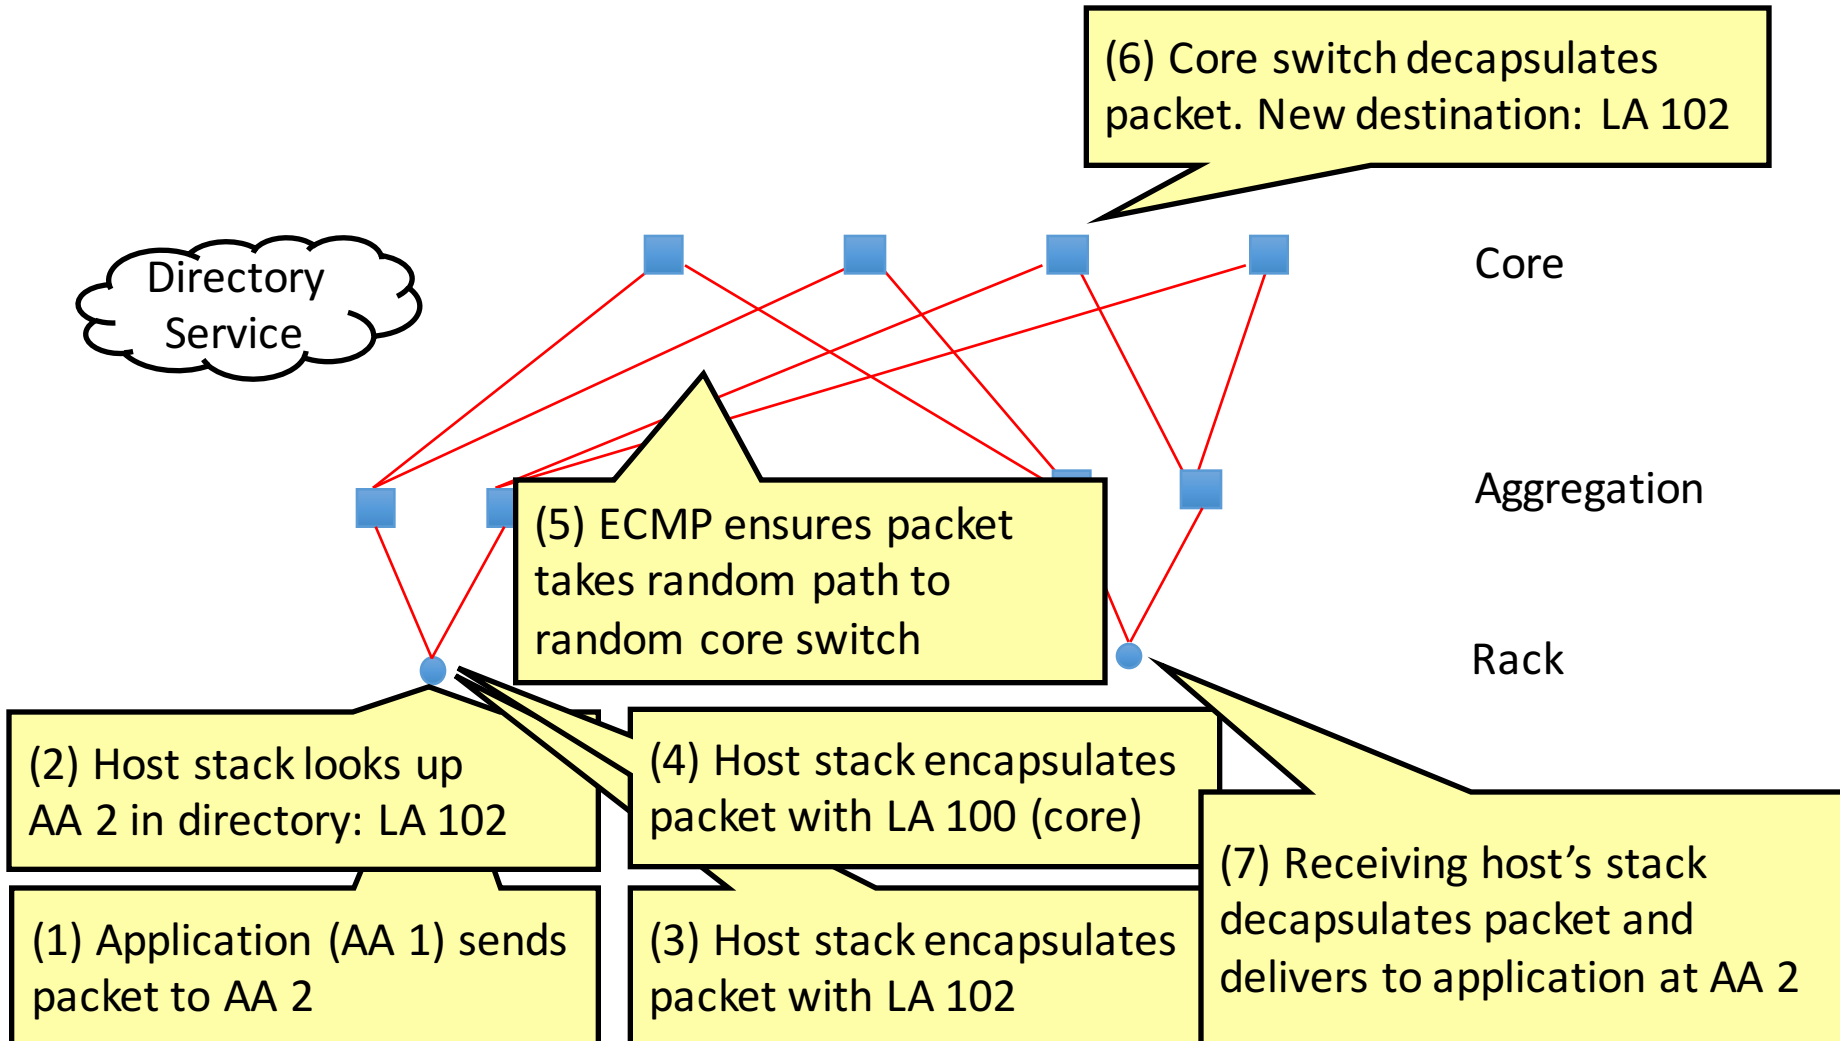
- Pick a random path for each flow using ECMP
- Custom host networking stack rewrites packets
  - Using IP encapsulation
- Provides high capacity



# Directory Service

- Apps see application-specific IP addresses (AAs)
- Switches route location-specific IP addresses (LAs)
- Directory service stores AA-LA mapping
  - And can enforce access control policies
- Custom host networking stack rewrites packets

# Path of a Packet





# Evaluation

- Valiant Load Balancing works
  - 94% efficiency, 0.98 fairness, isolation, failures
  - Assumptions: flows capped at ~1 GB, network links 10x faster than host links
- Directory service is fast enough
  - 60 servers can handle 100K hosts

# Discussion

- Which is easier to change, hosts or switches?
- Flat addresses and routing: ROFL, Disco
  - Map to hierarchical vs. route directly on flat