Data Center Networking

Amin Vahdat CSE 222a July 1, 2009

vahdat@cs.ucsd.edu

Current Team

- Amin Vahdat
- Mohammad Al-Fares
- Nathan Farrington
- Nelson Huang
- Brian Kantor
- Harsha Madhyastha
- Pardis Miri

- Radhika Niranjan
- Andreas Pamporis
- Sivasankar Radhakrishnan
- Erik Rubow
- Aram Shahinfard
- Vikram Subramanya



Walmart's Data Center



Blueprints for 200k sq. ft. Data Center in OR



Google's 36 Worldwide Data Centers



- Commoditization in the data center
 - Inexpensive, commodity PCs and storage devices
 - But network still highly specialized
- Data center is not a "small Internet"
 - One admin domain, not adversarial, limited policy routing, etc.
- Bandwidth is often the bottleneck
 - "Cloud" Computing
 - Service-oriented Architectures
 - Data Analysis (MapReduce)

Network Design Goals

- Scalable interconnection bandwidth
 - Full bisection bandwidth between all pairs of hosts
 Aggregate bandwidth = # hosts × host NIC capacity
- Economies of scale
 - Price/port linear with number of hosts
- Single network fabric
 - Support Ethernet and IP without end host modifications
- Management
 - Modular design
 - Avoid actively managing 100's-1000's network elements

Current Data Center Topologies

- Edge hosts connect to 1G Top of Rack (ToR) switch
- ToR switches connect to 10G End of Row (EoR) switches
- Large clusters: EoR switches to 10G core switches
 - Oversubscription of 2.5:1 to 8:1 typical in guidelines
- No story for what happens as we move to 10G to the edge



Key challenges: performance, cost, routing, energy, cabling