


Data-Center Traffic Management

COS 597E: Software Defined Networking

Jennifer Rexford
Princeton University
MW 11:00am-12:20pm

Cloud Computing

Cloud Computing



- Elastic resources
 - Expand and contract resources
 - Pay-per-use
 - Infrastructure on demand
- Multi-tenancy
 - Multiple independent users
 - Security and resource isolation
 - Amortize the cost of the (shared) infrastructure
- Flexible service management

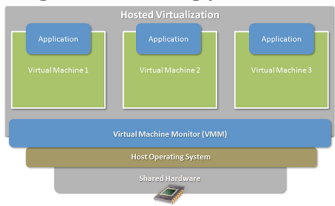
Cloud Service Models

- Software as a Service
 - Provider licenses applications to users as a service
 - E.g., customer relationship management, e-mail, ...
 - Avoid costs of installation, maintenance, patches...
- Platform as a Service
 - Provider offers platform for building applications
 - E.g., Google’s App-Engine
 - Avoid worrying about scalability of platform

Cloud Service Models

- Infrastructure as a Service
 - Provider offers raw computing, storage, and network
 - E.g., Amazon’s Elastic Computing Cloud (EC2)
 - Avoid buying servers and estimating resource needs

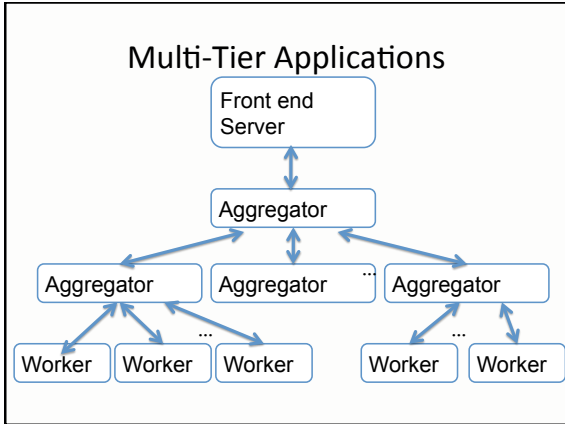
Enabling Technology: Virtualization



- Multiple virtual machines on one physical machine
- Applications run unmodified as on real machine
- VM can migrate from one computer to another

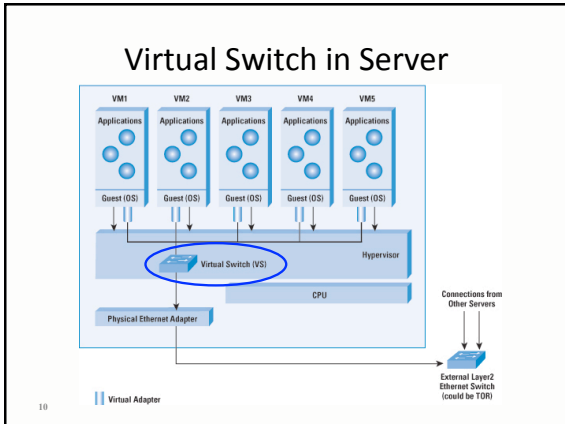
Multi-Tier Applications

- Applications consist of tasks
 - Many separate components
 - Running on different machines
- Commodity computers
 - Many general-purpose computers
 - Not one big mainframe
 - Easier scaling



Data Center Network

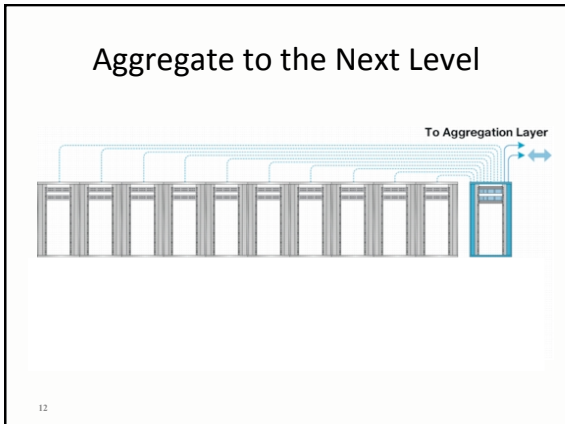
9



Top-of-Rack Architecture


- Rack of servers
 - Commodity servers
 - And top-of-rack switch
- Modular design
 - Preconfigured racks
 - Power, network, and storage cabling

11




Modularity, Modularity, Modularity

- Containers

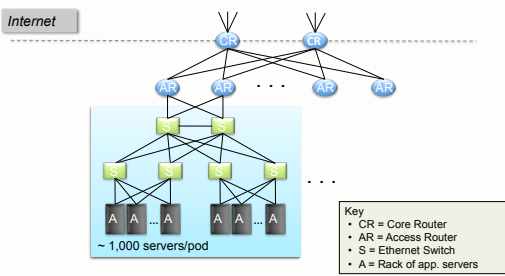


- Many containers



13

Data Center Network Topology



Internet

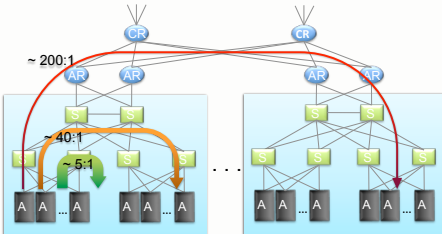
~ 1,000 servers/pod

Key

- CR = Core Router
- AR = Access Router
- S = Ethernet Switch
- A = Rack of app. servers

14

Capacity Mismatch



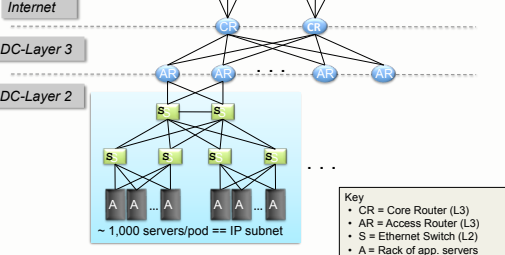
~ 200:1

~ 40:1

~ 5:1

15

Data-Center Routing



Internet

DC-Layer 3

DC-Layer 2

~ 1,000 servers/pod == IP subnet

Key

- CR = Core Router (L3)
- AR = Access Router (L3)
- S = Ethernet Switch (L2)
- A = Rack of app. servers

16

Traffic Management

Hedera and HONE

17

Traffic Management Challenges

- High volumes of "east-west traffic"
- Low bisection bandwidth
- Volatile traffic patterns
- Elephant flows
- TCP incast
- Naïve application programmers
- Performance problems due to stragglers
- Difficulty of collecting measurement data

18

Traffic Management Opportunities

- Low latencies within the data center
 - Small TCP round-trip times
 - Easier to use central controller
- End-to-end control
 - Applications, servers, and switches
- Greater visibility
 - Monitoring on the end hosts and soft switches
- Green-field deployments
- VM placement and migration
- Simple, symmetric topologies

19

Discussion

- Granularity of monitoring and control
 - Individual flows?
 - Larger traffic aggregates?
- End host vs. network
 - Where to measure?
 - Where to exercise control?
- Integrating end hosts with the controller

20