

(Demystifying)

Datacenter Network Monitoring and Management

Vojislav Dukic
vojislav.dukic (a t) [inf.ethz.ch](mailto:vojislav.dukic@inf.ethz.ch)

Desislava Dimitrova
desislava.dimitrova (a t) inf.ethz.ch

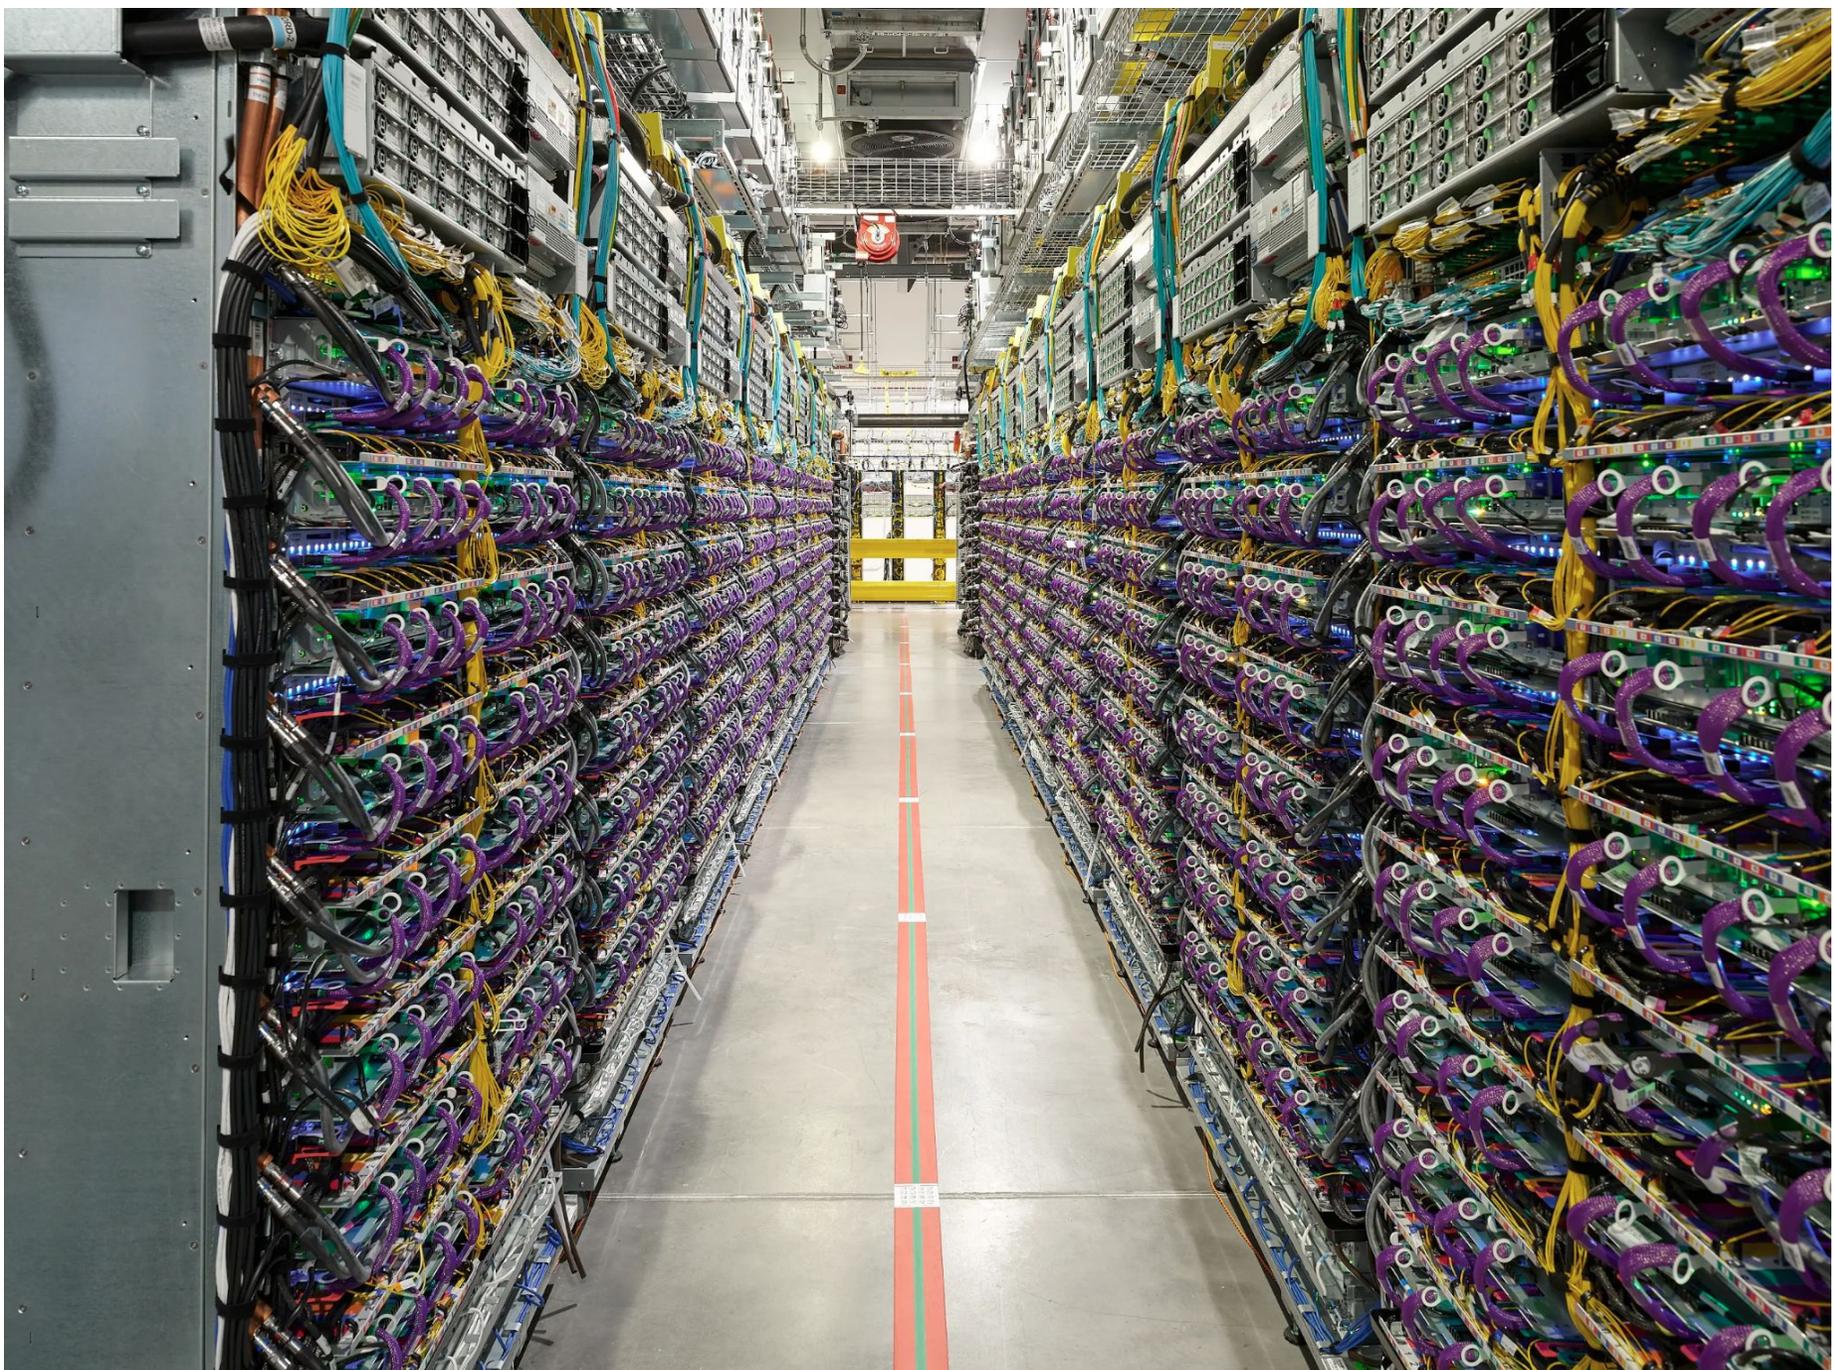
Course schedule and materials

<https://vojislav.dukic.ch/teaching/2026-dcnet/>

Grading

1 (2) report(s) every week +
one time 15 minute presentation

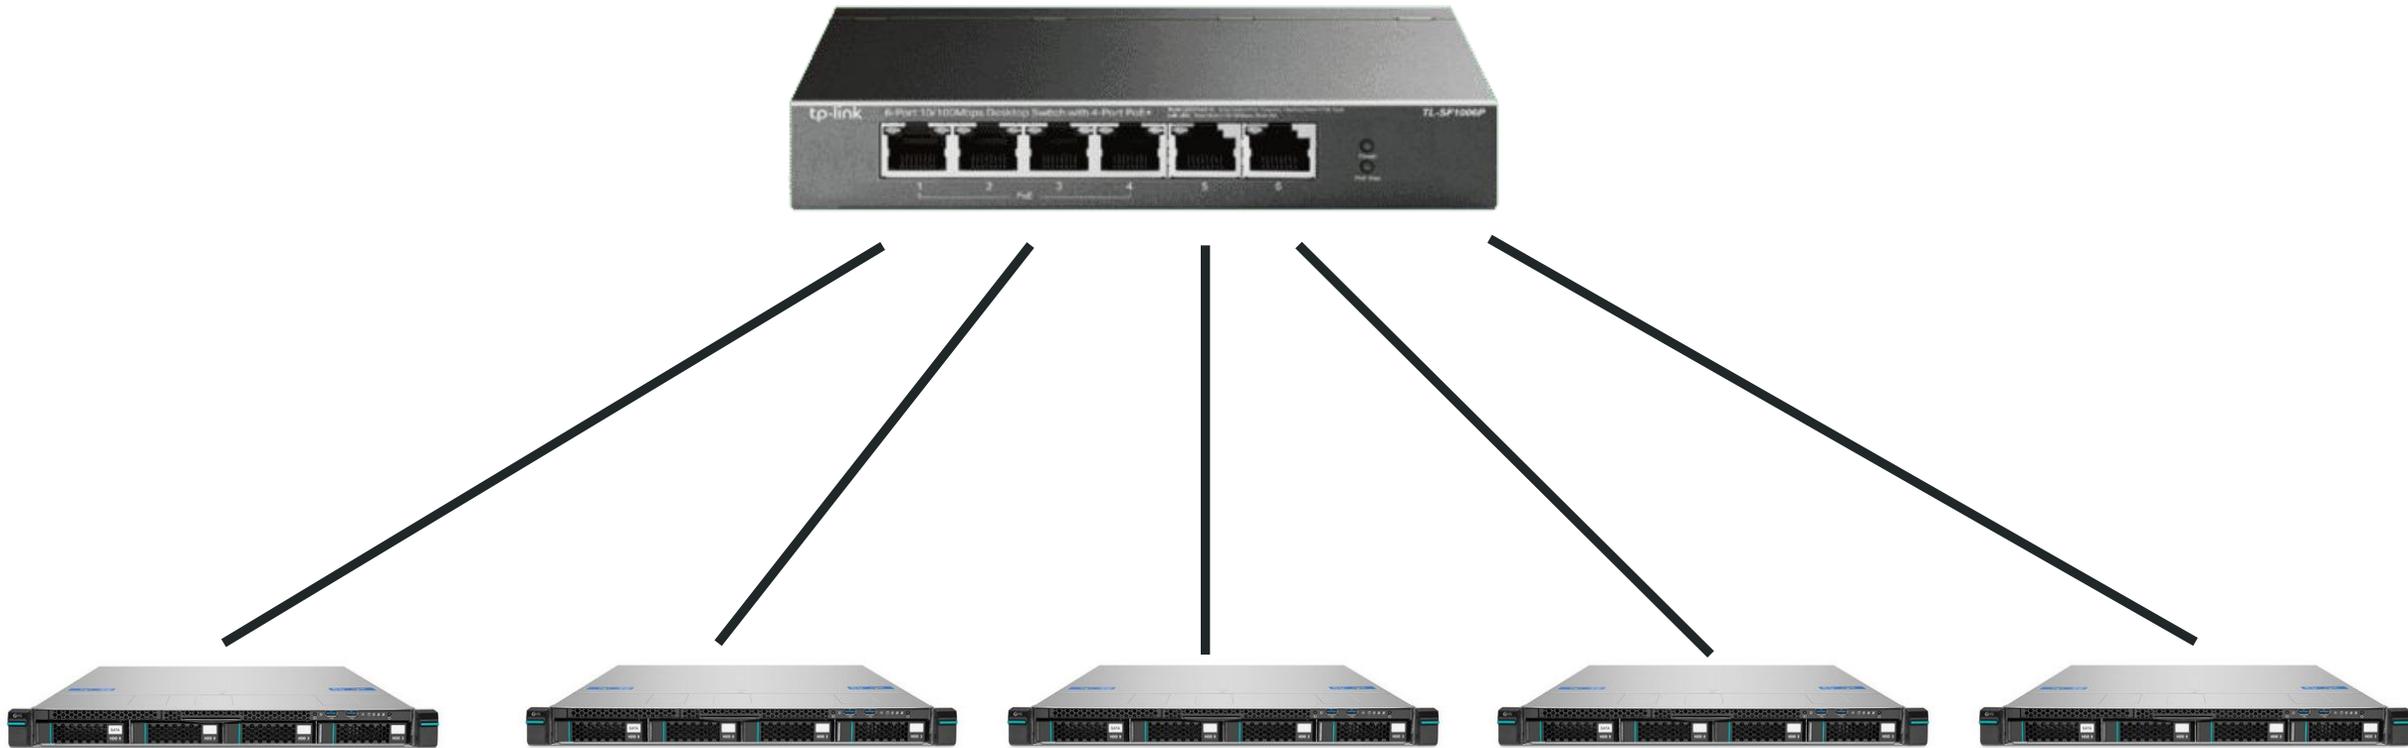
presentation in groups of 2 (3)



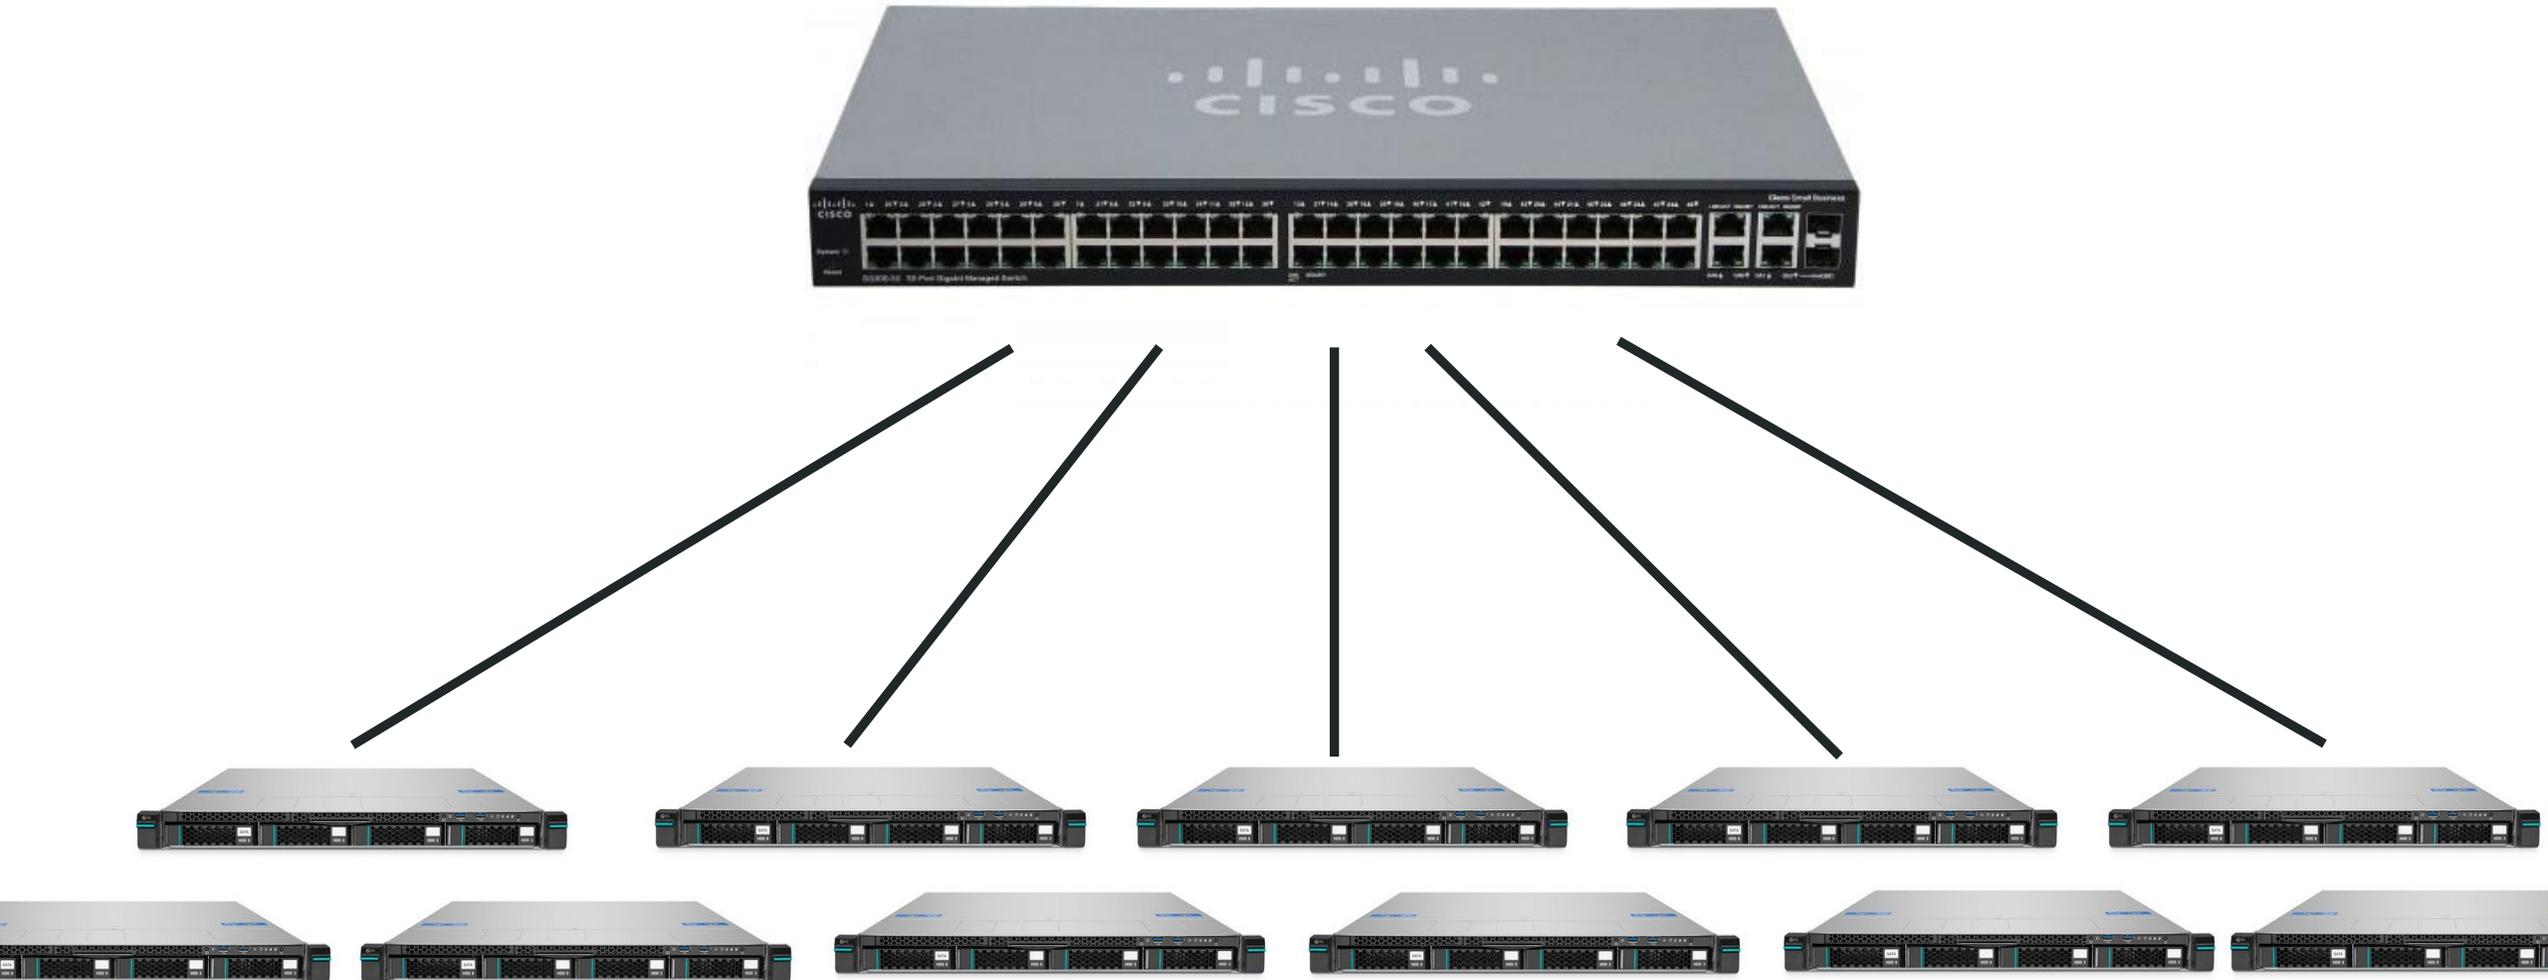
Commodity hardware



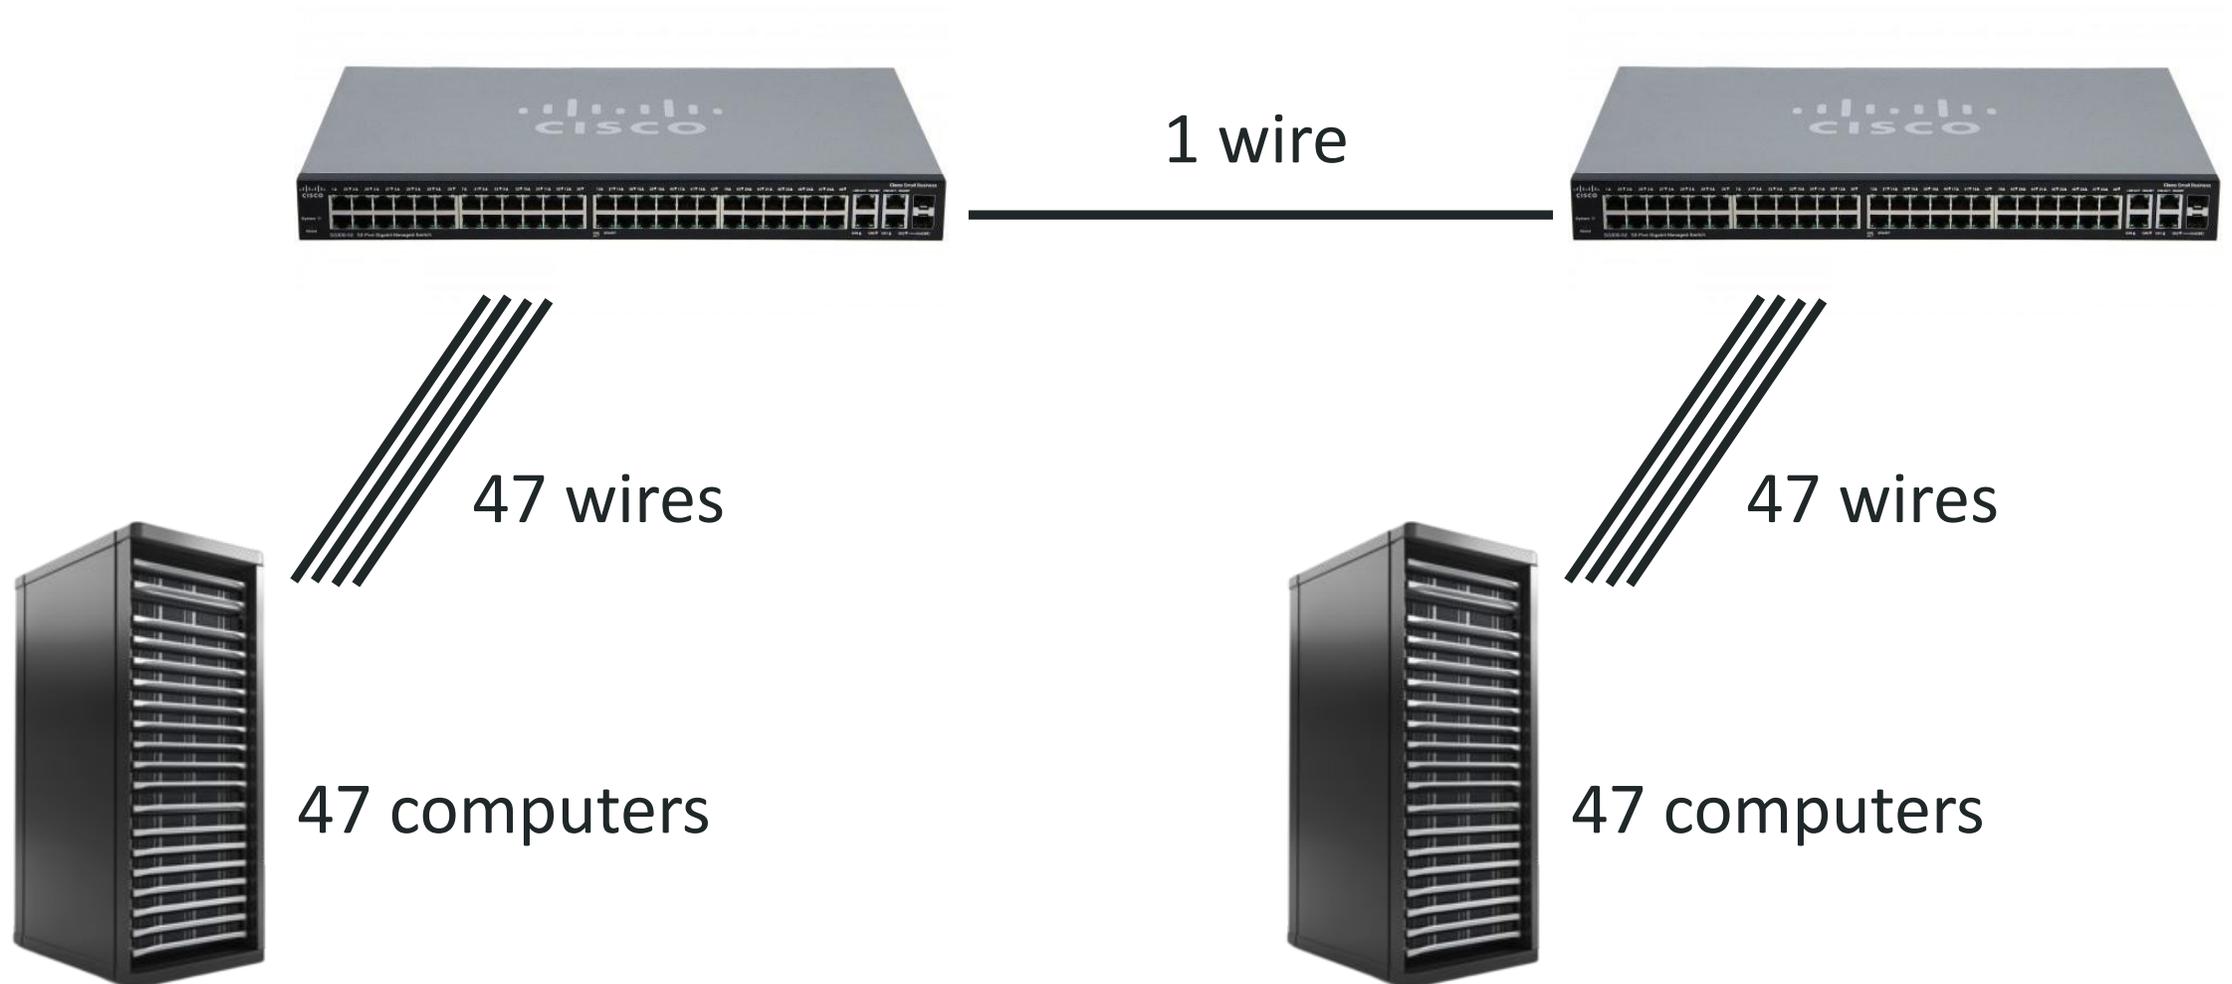
One machine – not enough. How about 5 machines (Rack)?



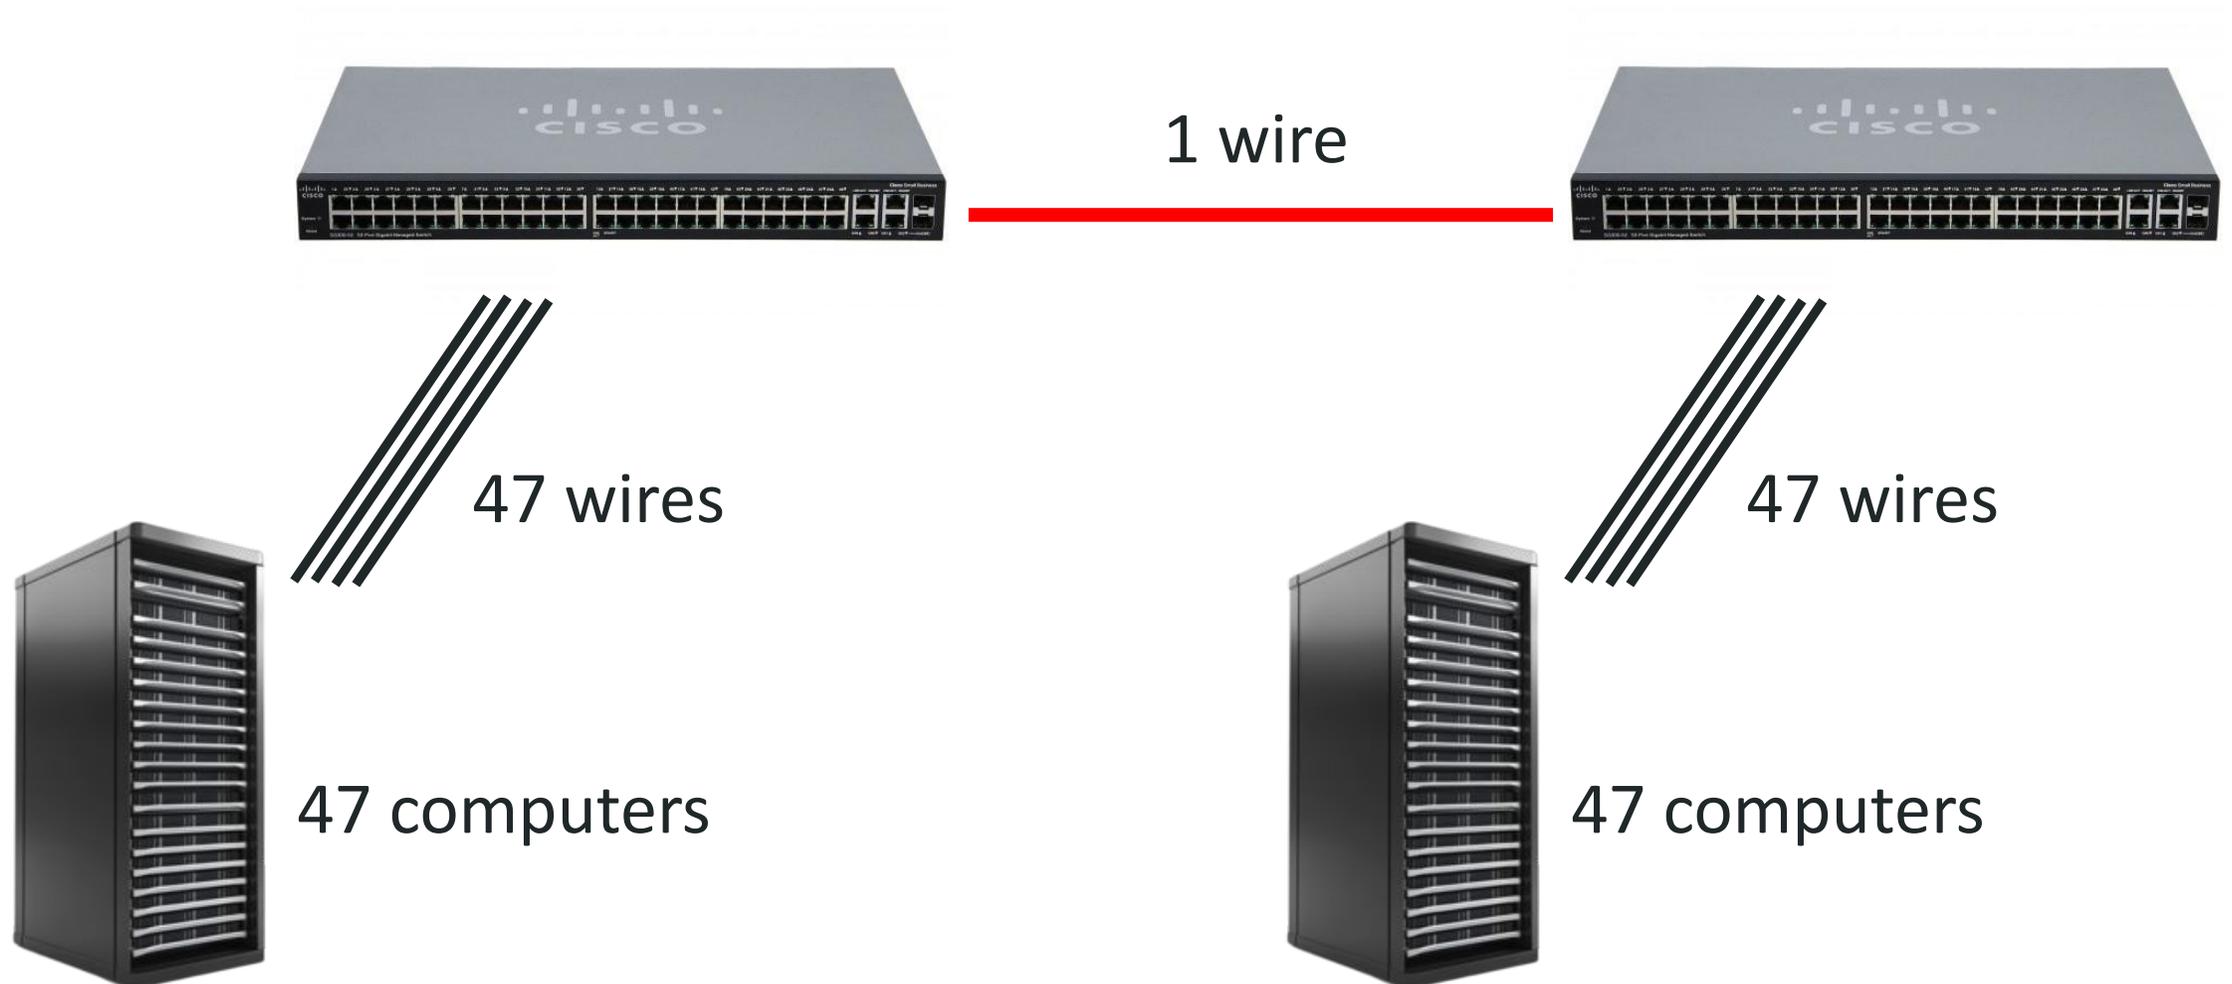
5 not enough? How about 48?
Buy a bigger switch approach



Beyond 48? They don't sell bigger switches.
How about connecting two switches?

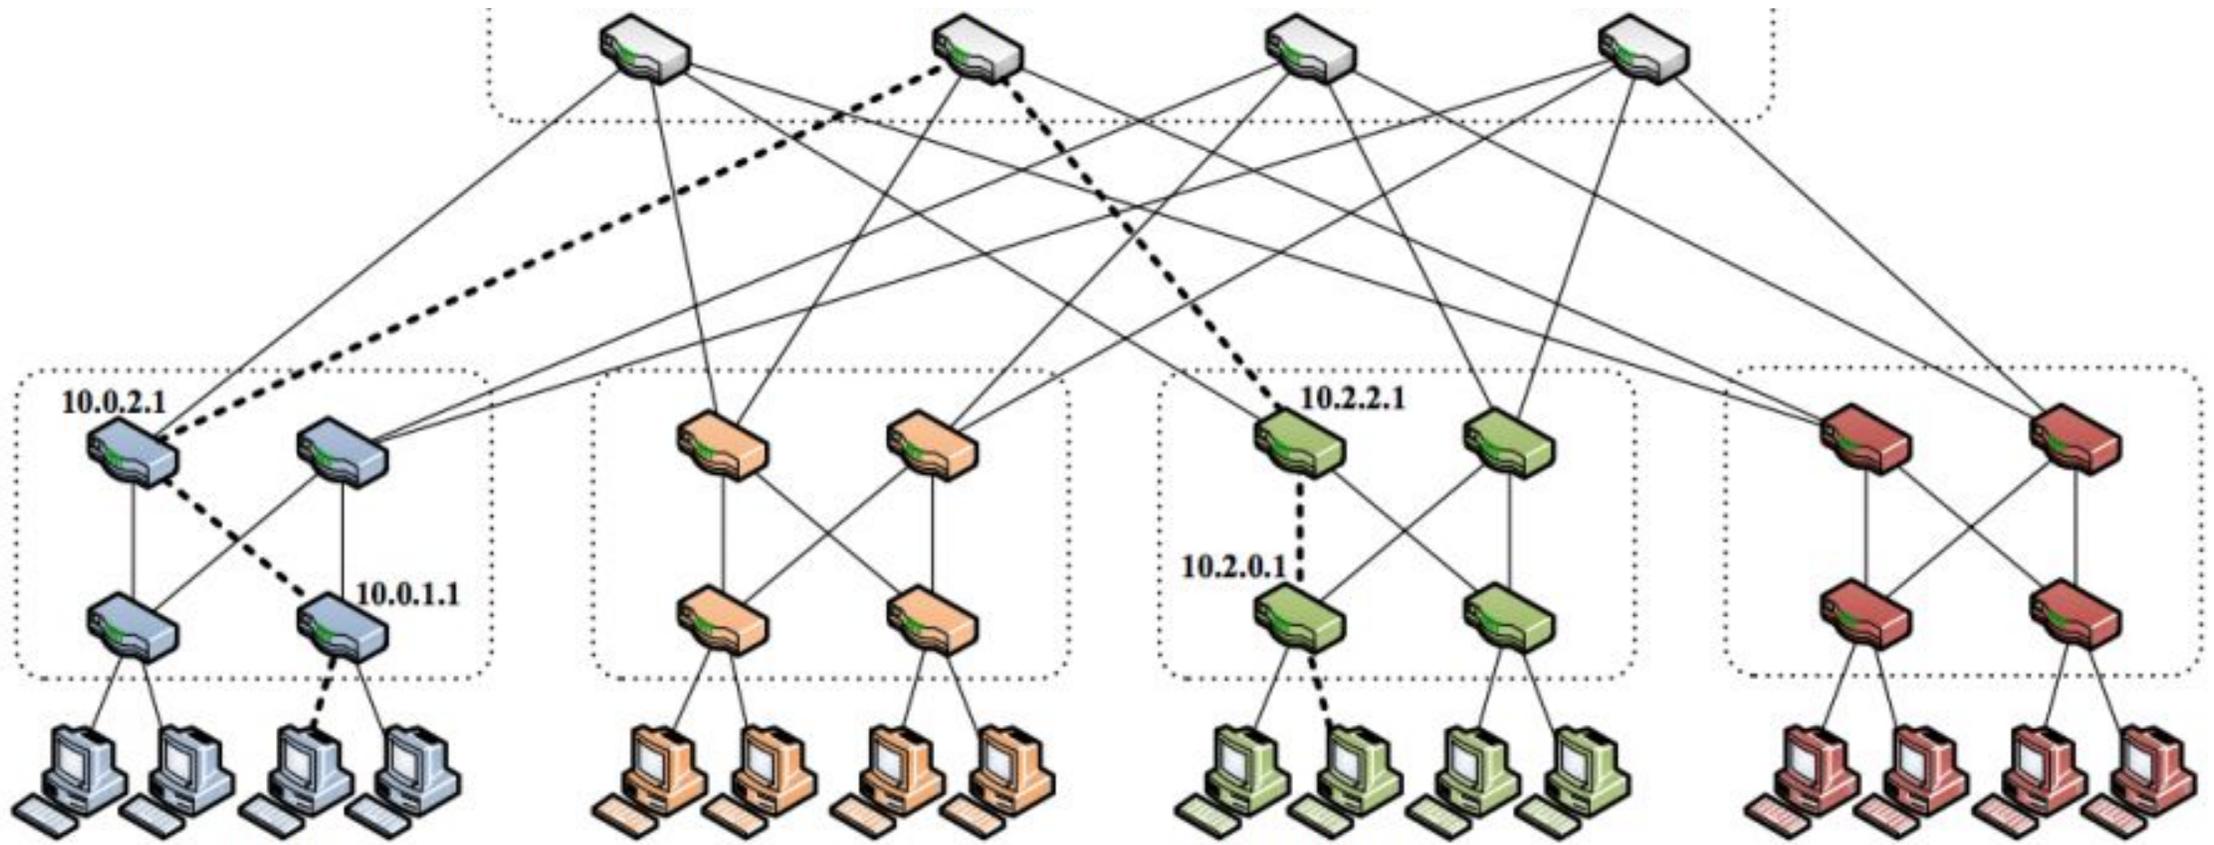


Bandwidth problem: The bottleneck link prevents congestion-free communication between machines



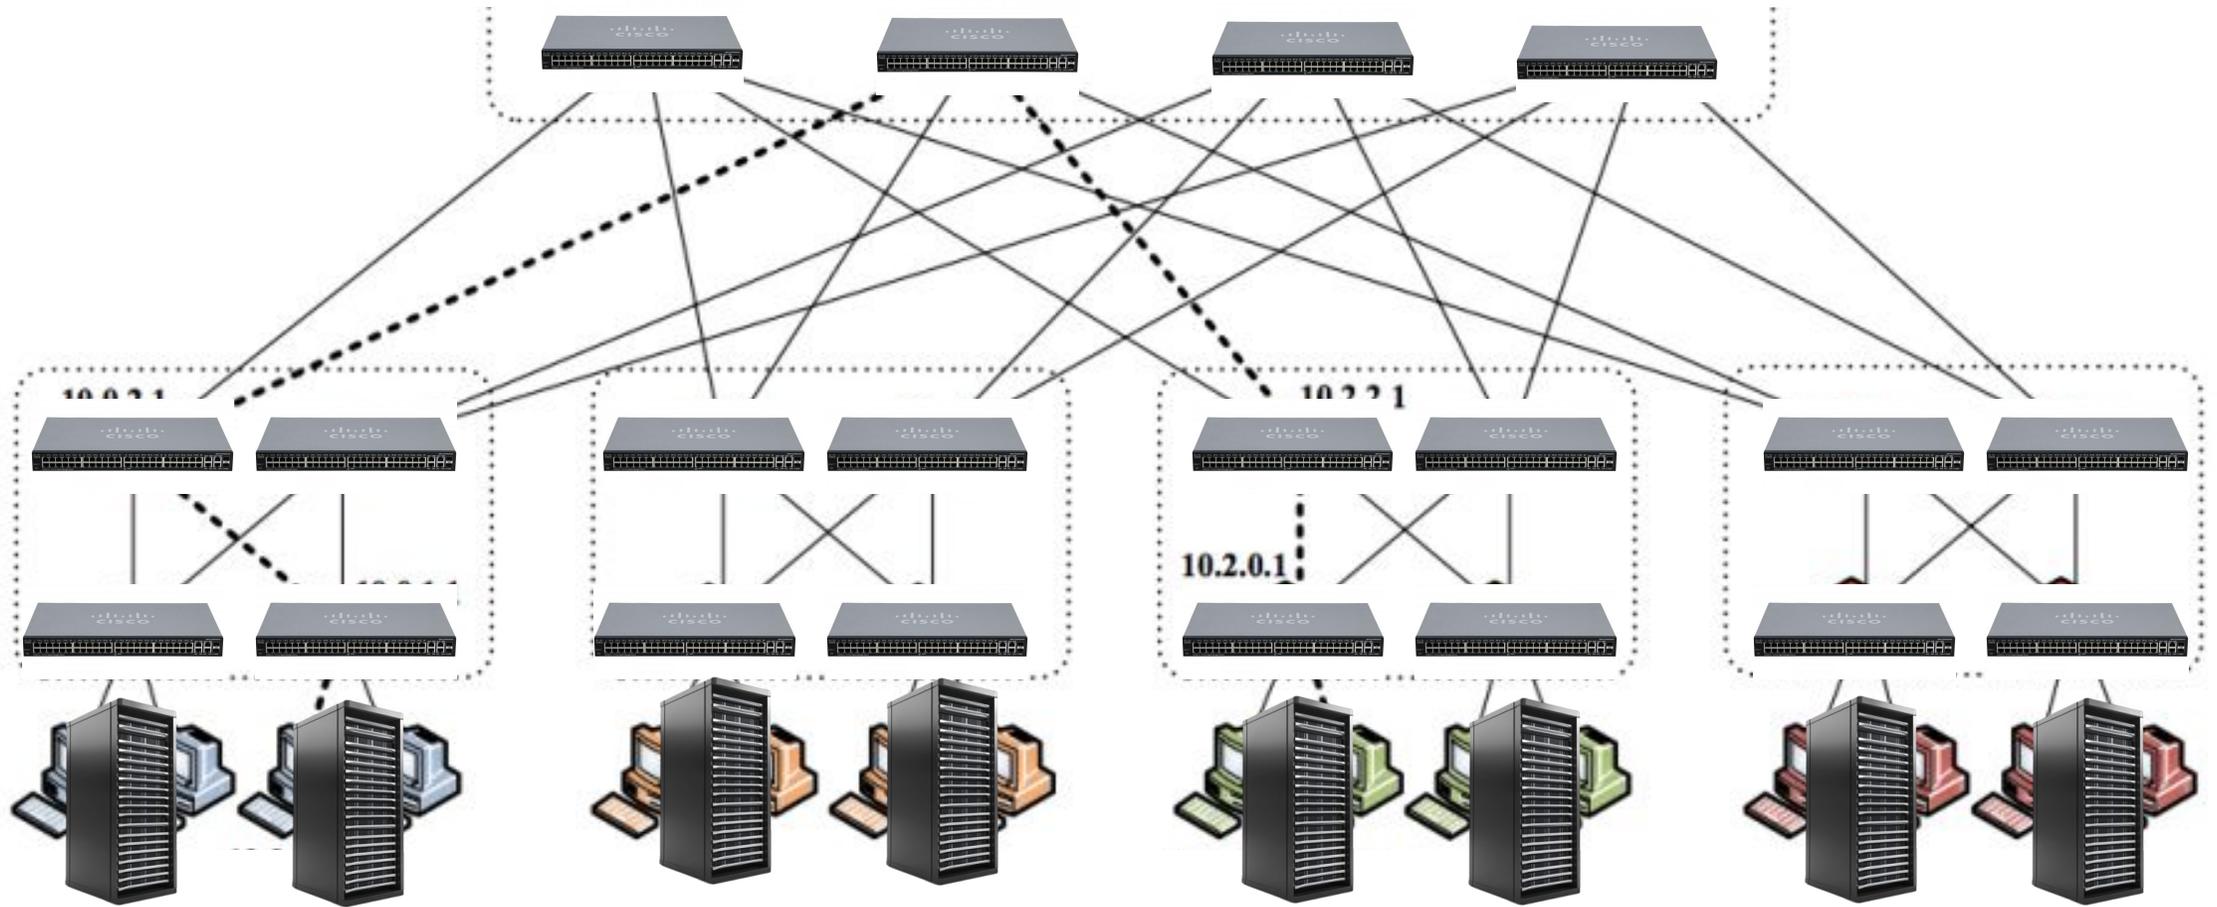
Solution – Fat Tree

non-blocking Clos topology

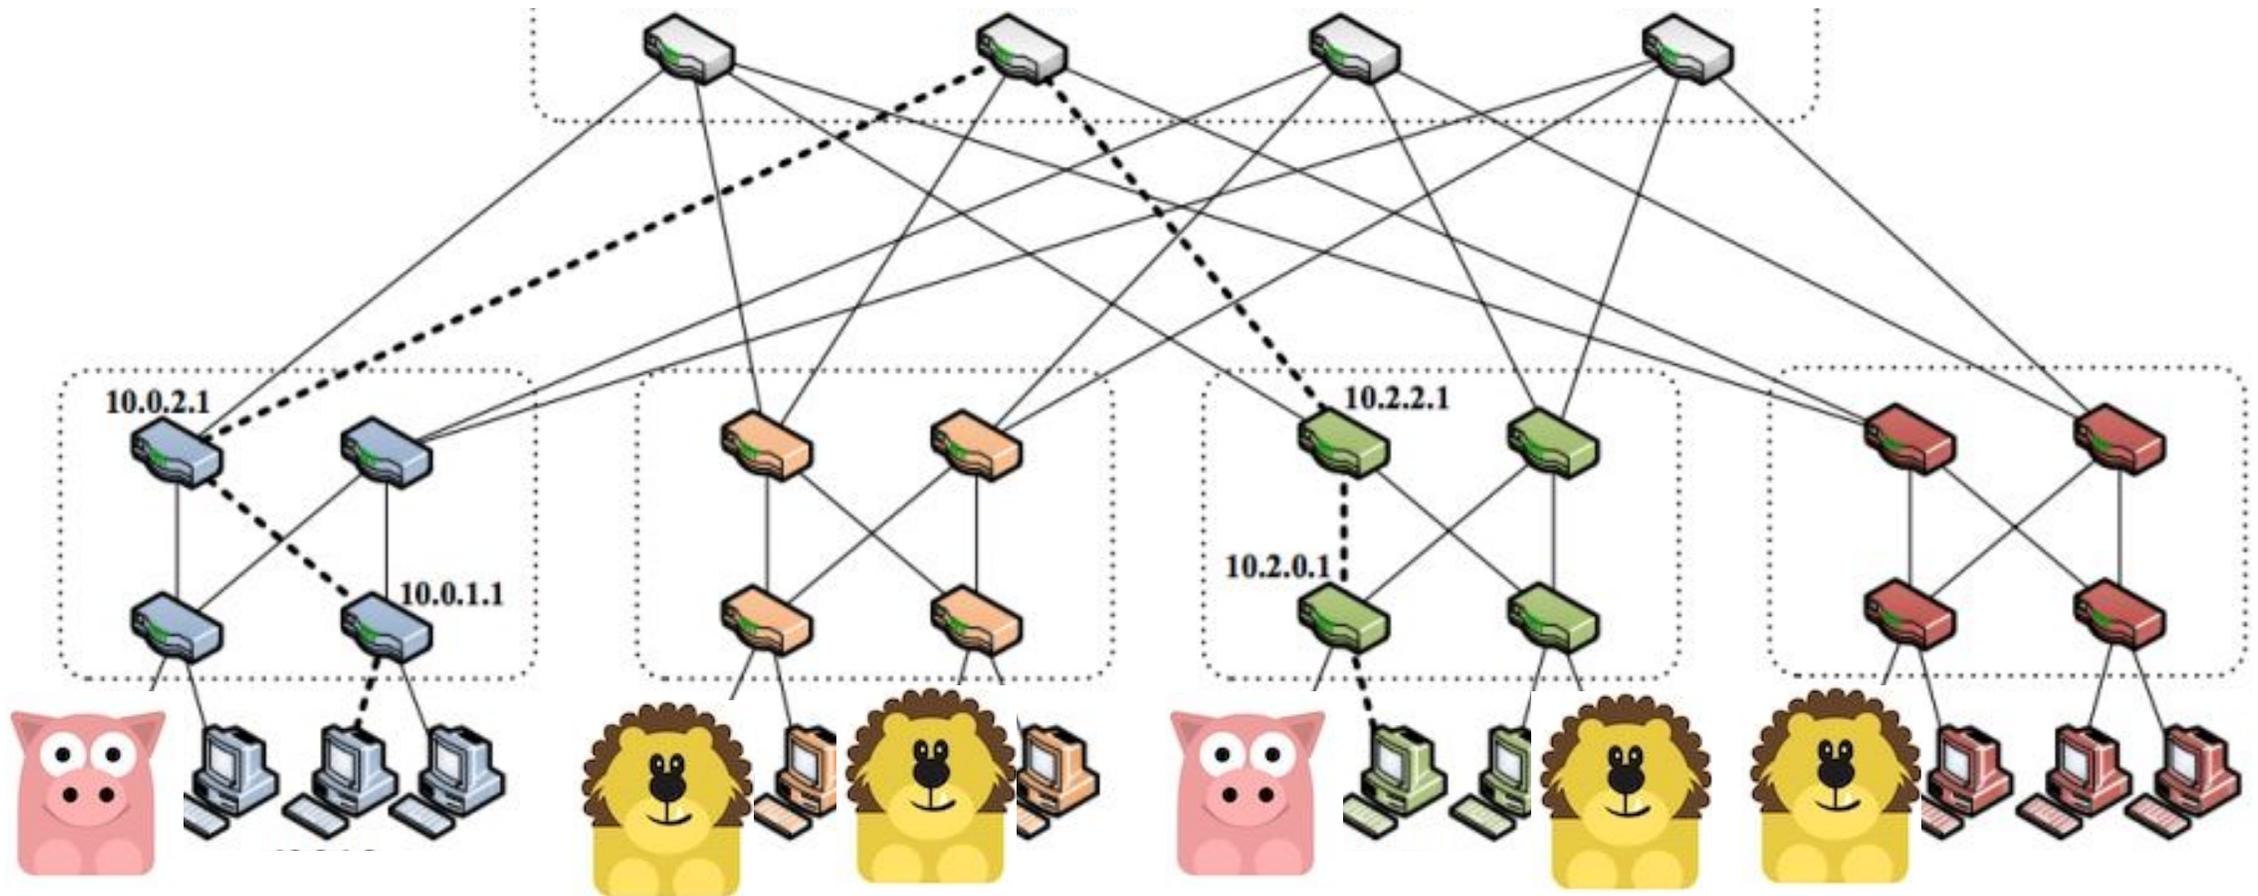


* A scalable, commodity data center network architecture, SIGCOMM '08

Solution – Fat Tree **non-blocking** Clos topology

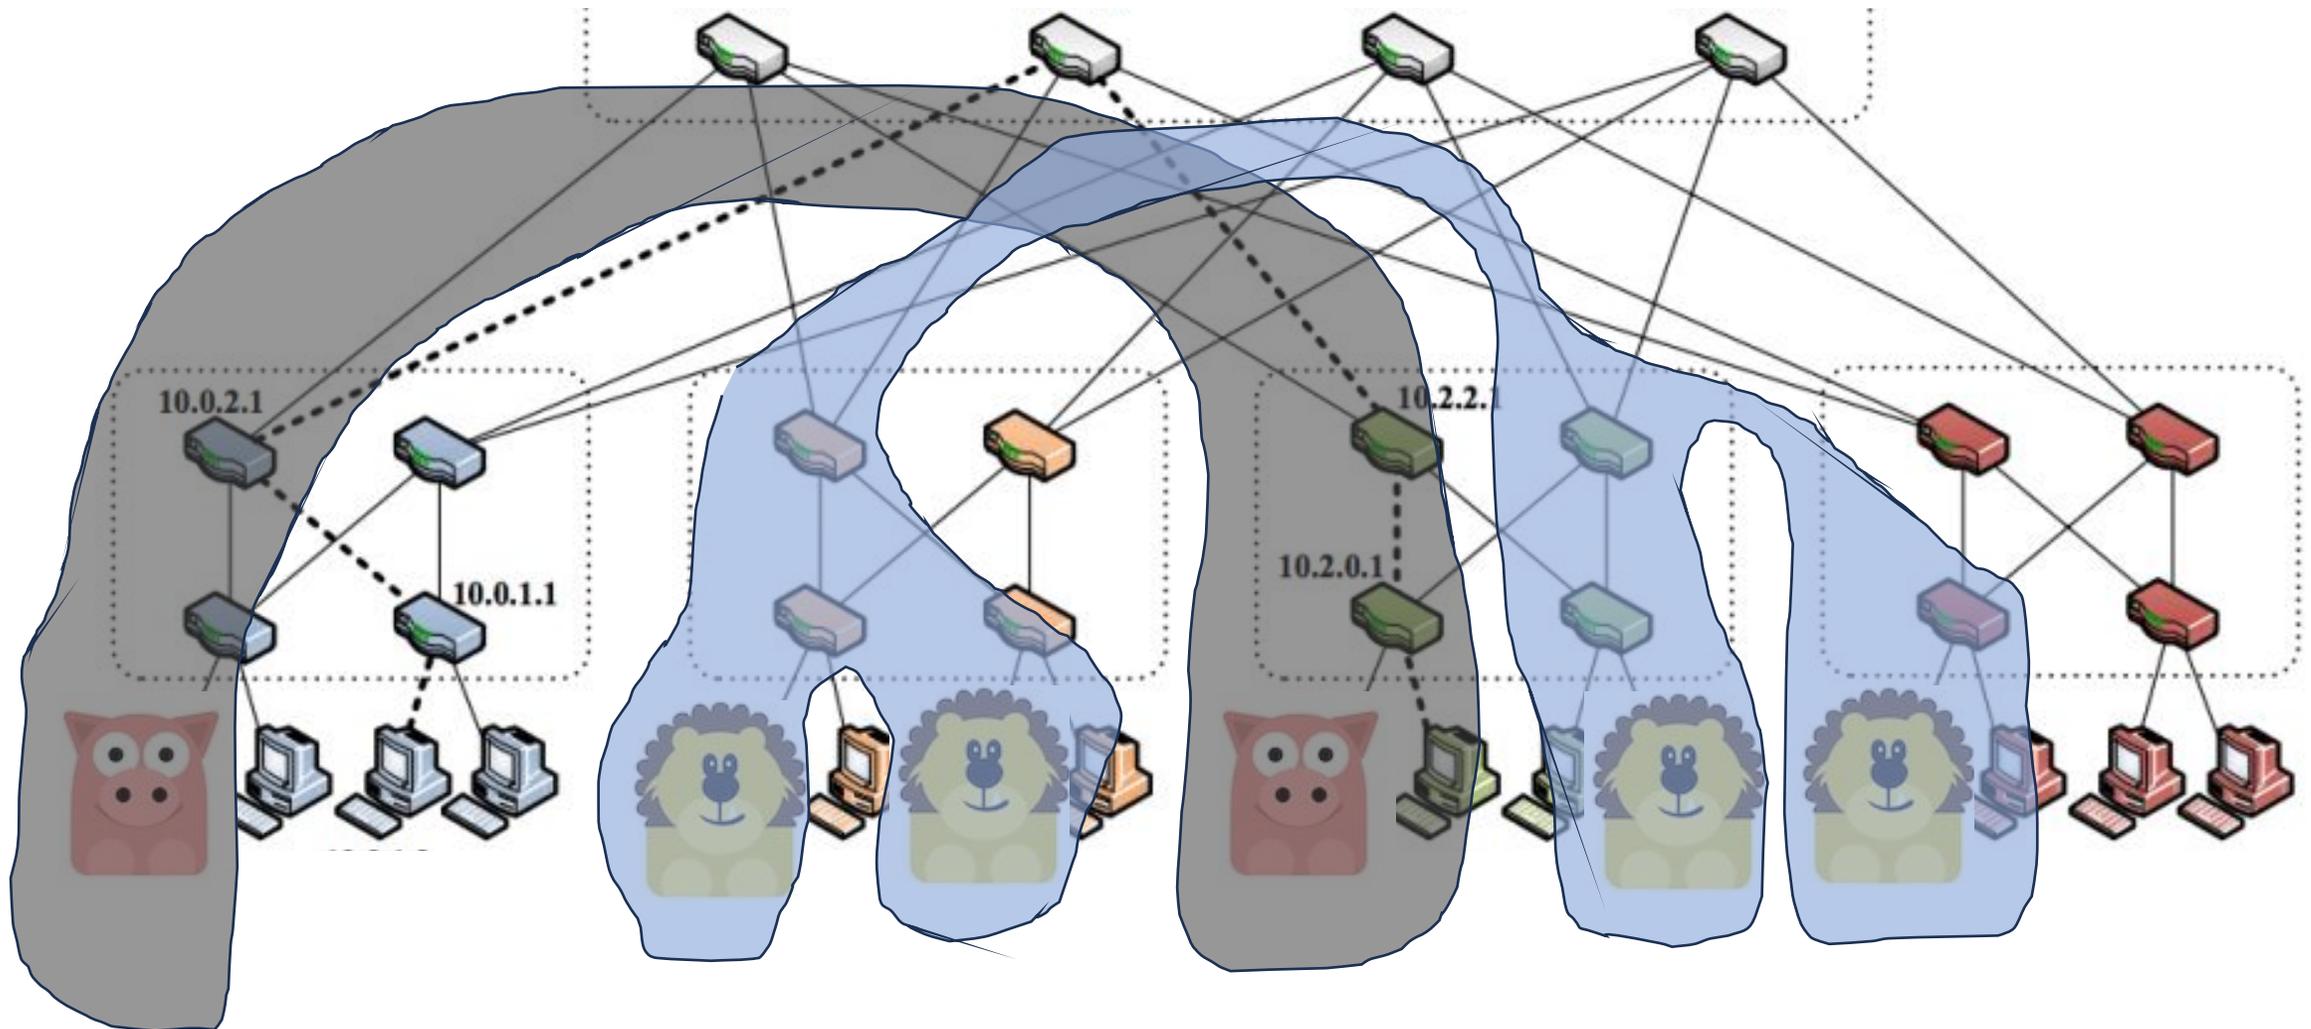


Problem: The infrastructure is too large for a single customer – sell it to multiple tenants

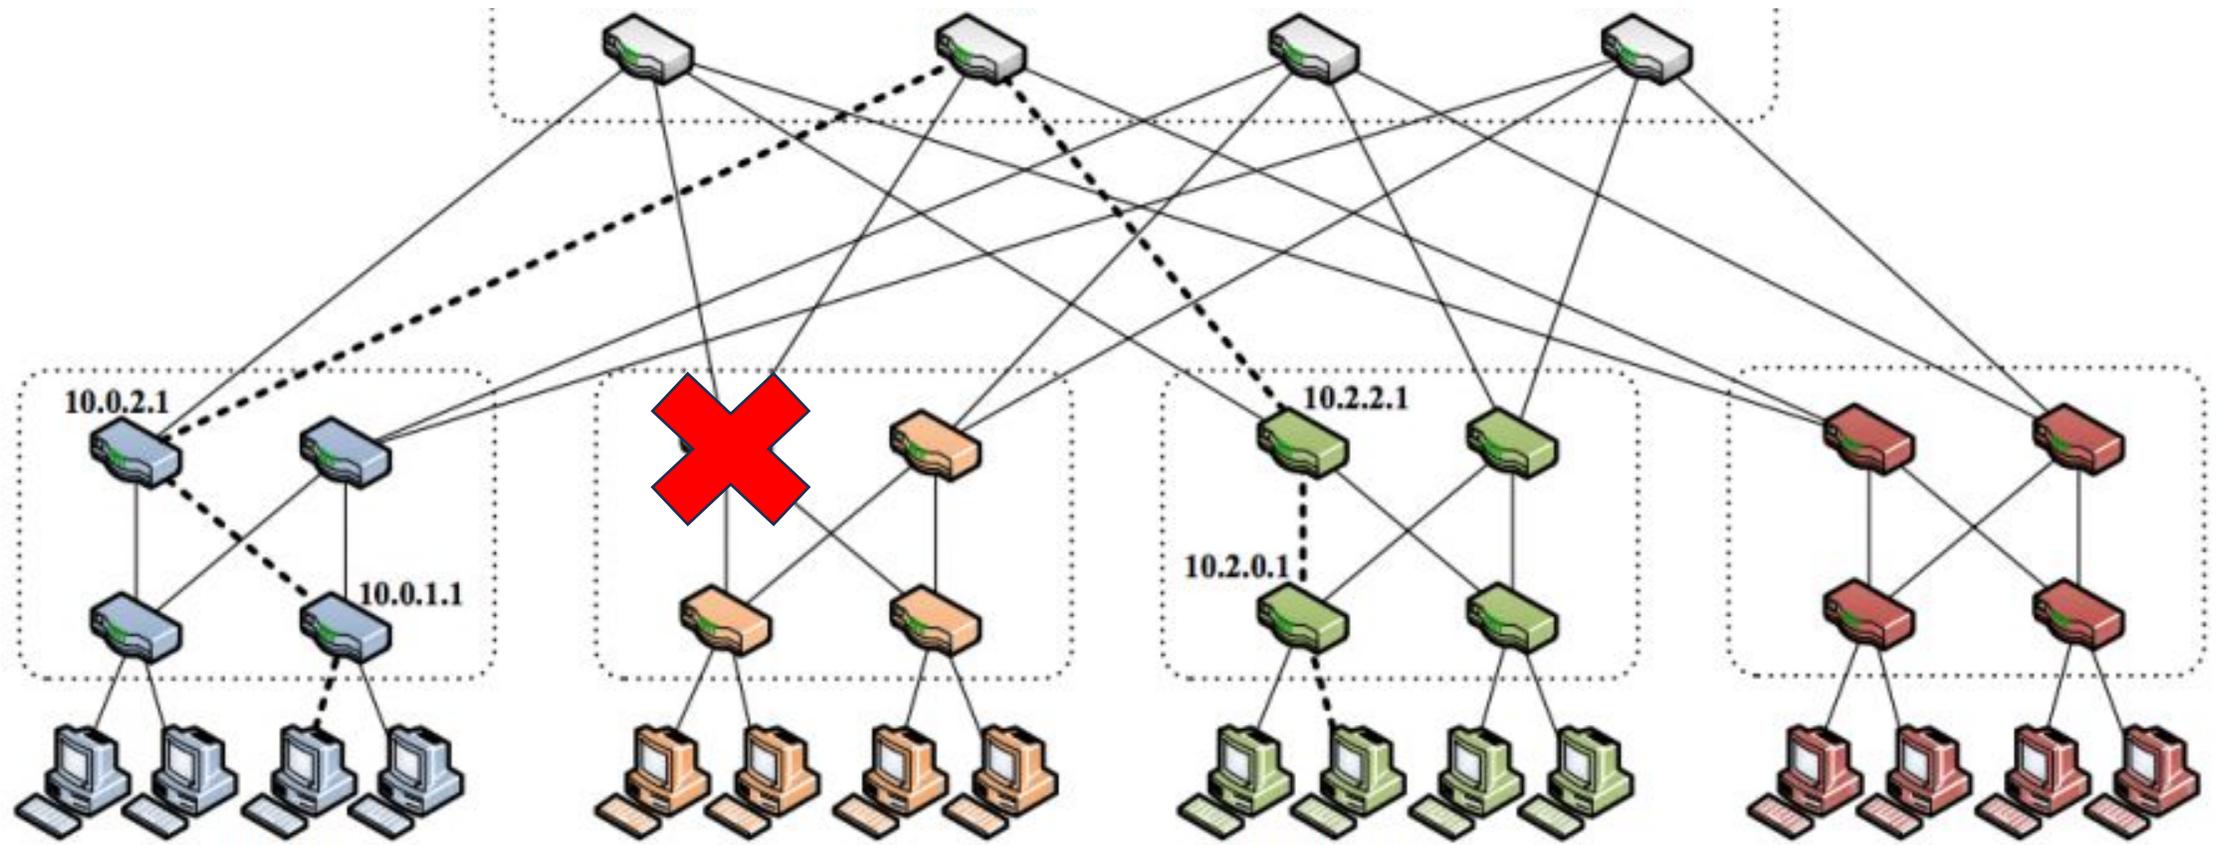


Solution: Virtual Private Cloud (VPC)

A logical layer on top of the physical network

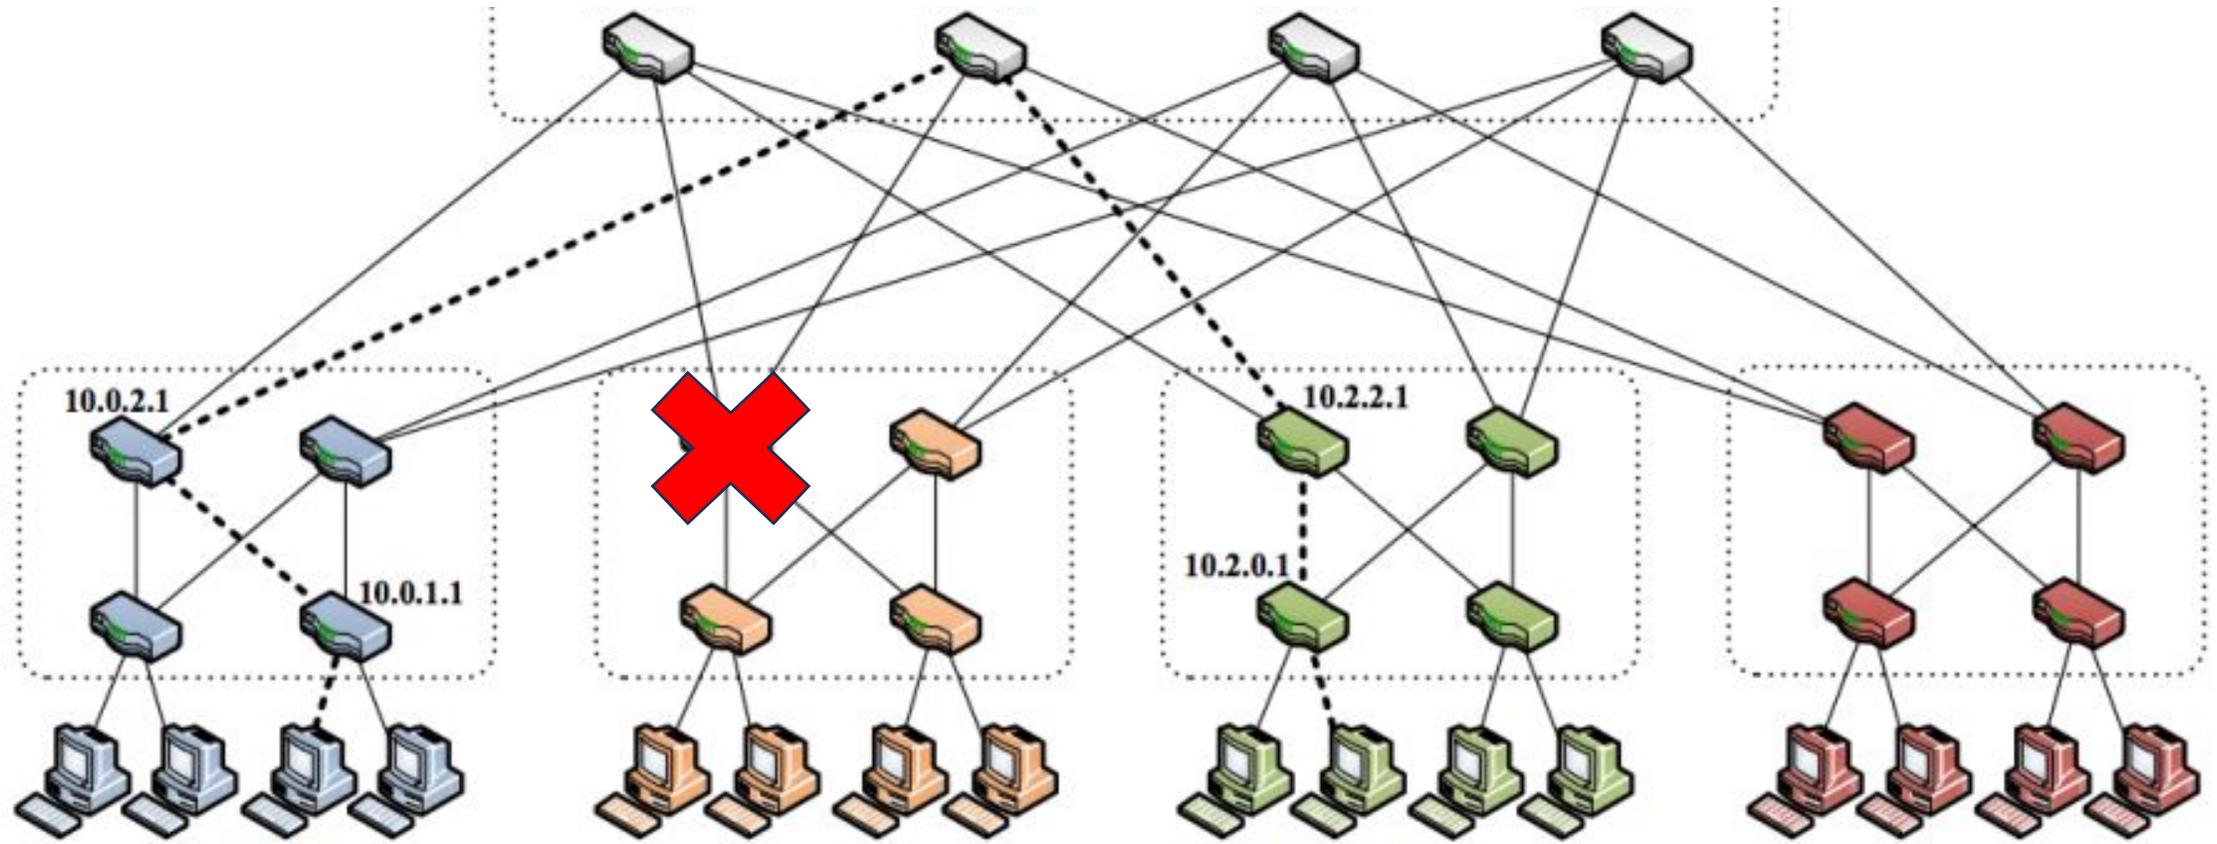


Problem: What if something fails?

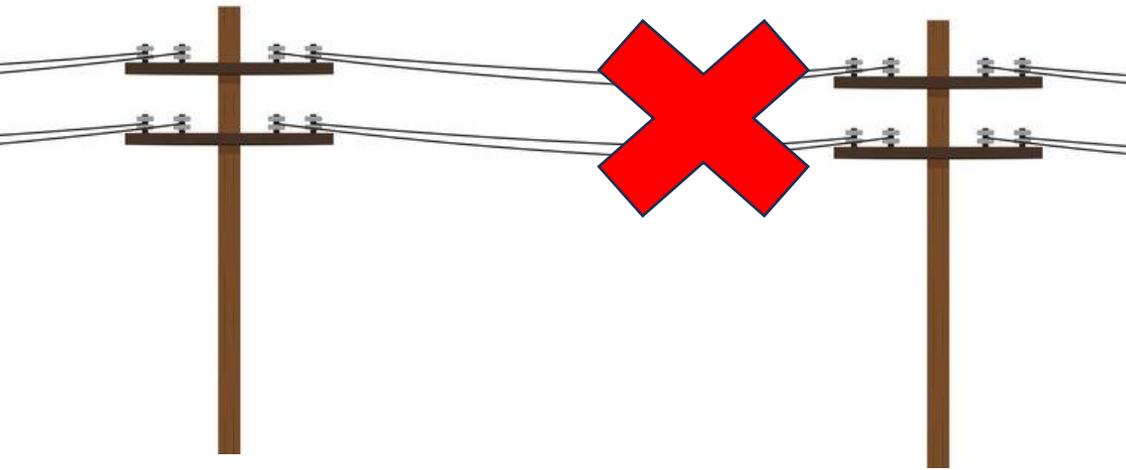


Problem: What if something fails?

Solution: Fat-tree has enough redundancy



Problem: What if a data center fails?

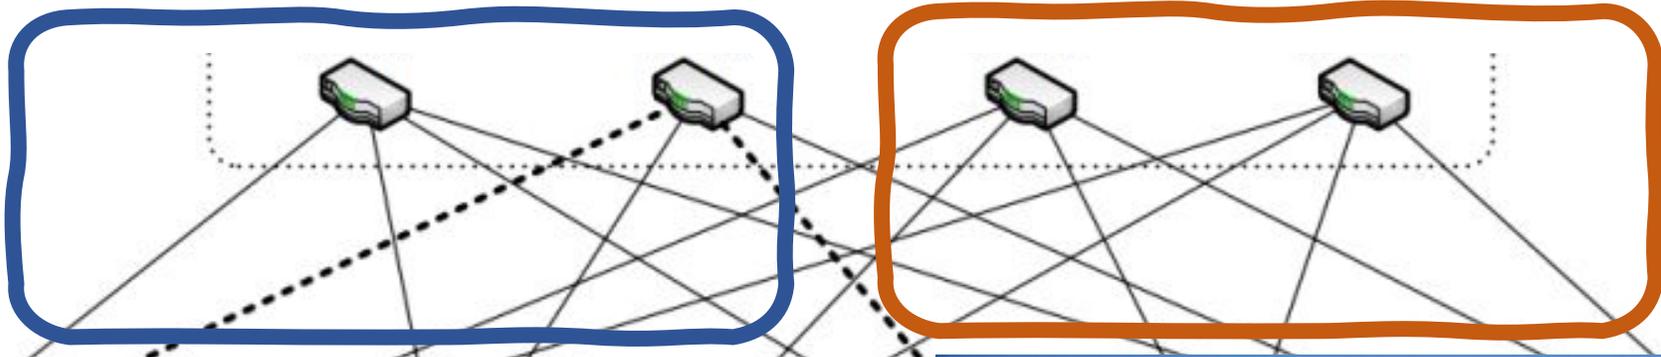


Diesel generators



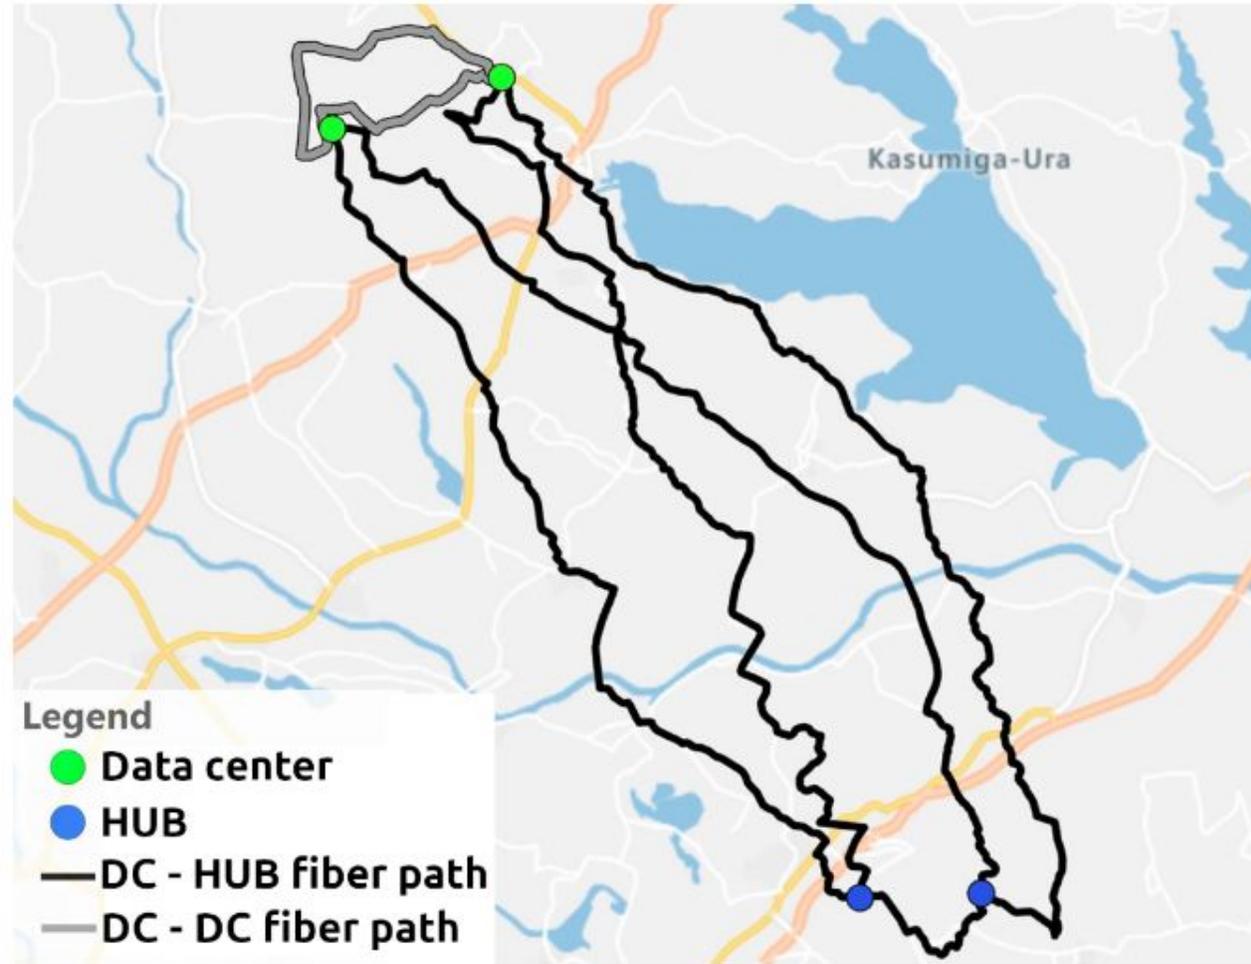
What if the generators fail?

Solution: Availability Zones – run your app in multiple AZs



10

Data Center Interconnect (DCI)

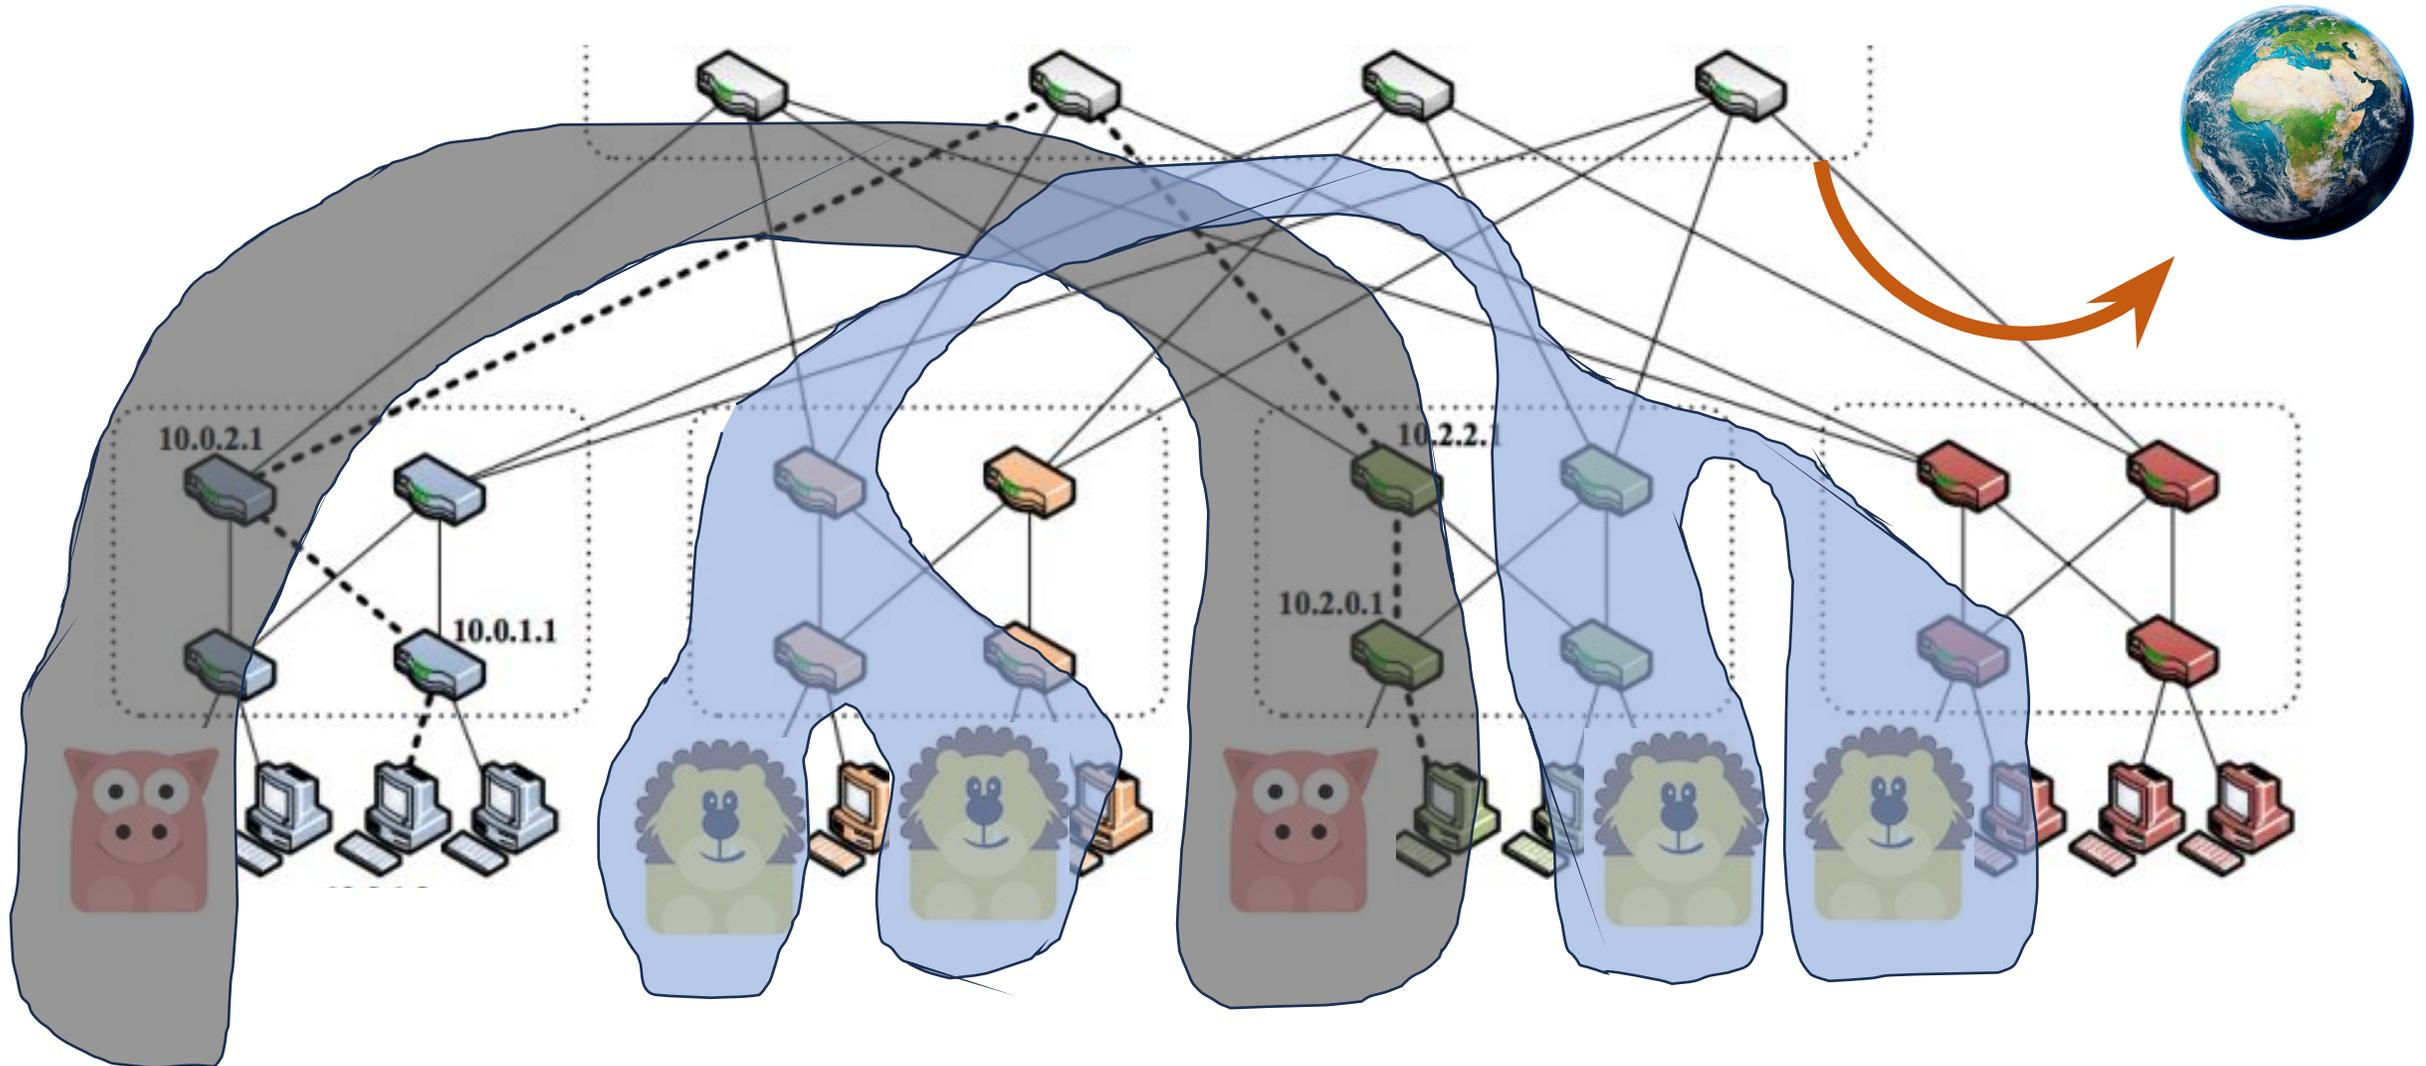


Data Center Interconnect (DCI)

Fiber huts



Problem: Users want to connect to the Internet



Solution: Gateway – you need to enable it for your applications

- Access right
- Virtual Private Cloud
- Firewall

Core network switch



Cisco Firewall
(NFV more generally)



Internet



Problem: Users are far away

Limit – c speed

Speed of light through a fiber cable – $2/3 c$ speed

Real path = 3 x Geodesic



Solution 1: Move your workload closer to the user

Cloud providers have infrastructure in multiple regions across the world



Solution 2: Use Content Delivery Networks

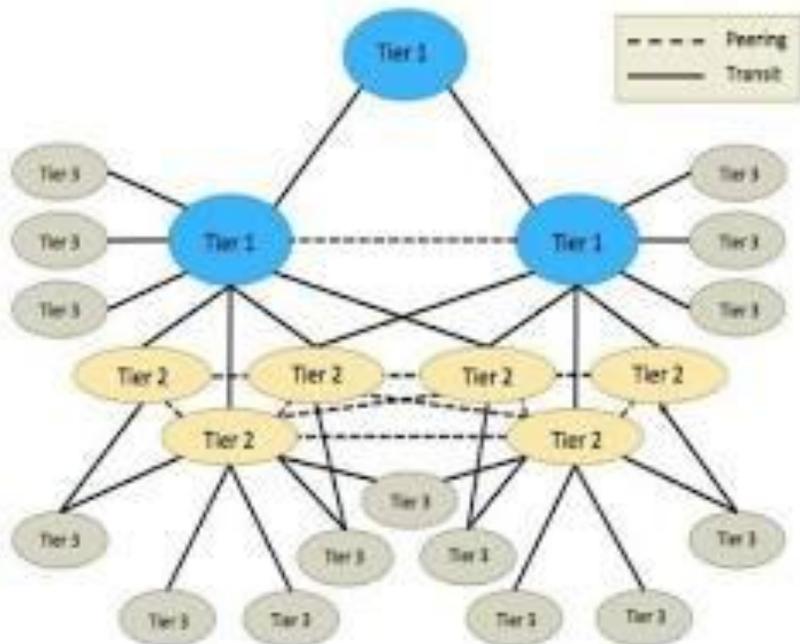
100s of regions across the world
Mostly for static content



Problem: Internet bandwidth is expensive

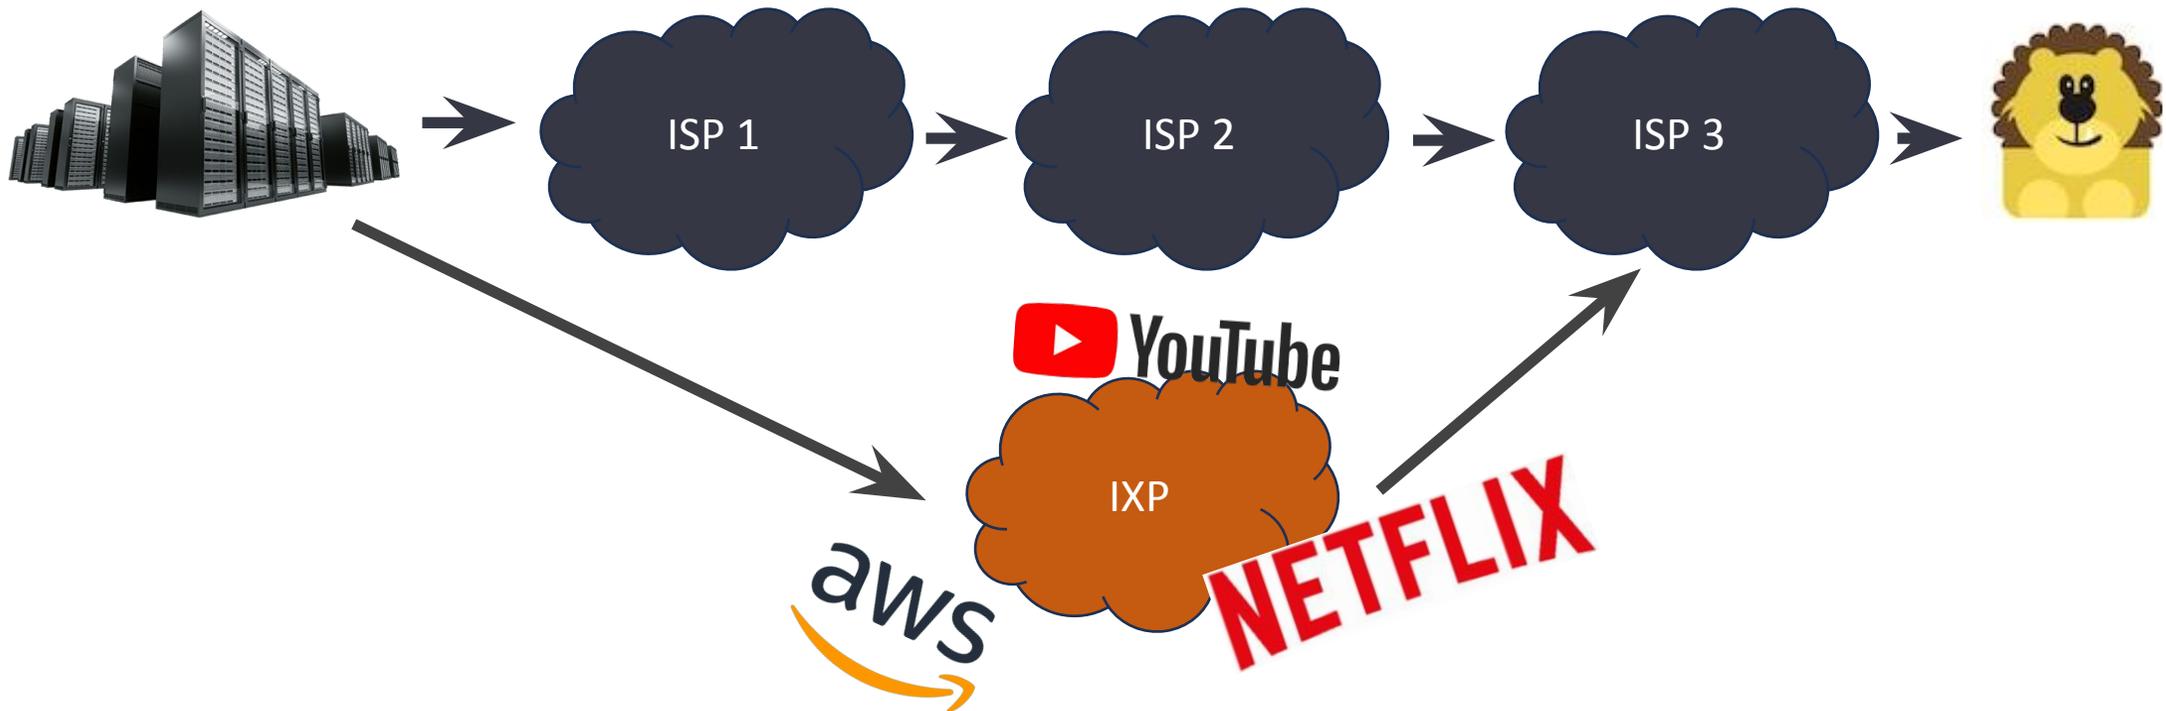


Problem: Internet bandwidth is expensive



ISP Tiering – Tier 1, Tier 2, and Tier 3

(One possible) Solution: Internet Exchange Points (IXP)



The final path: From the cloud to the user

