

6.1800 Spring 2025

Lecture #14: Datacenter Networks

networking in “the cloud”

DPPR due Friday. Your grade is coming from your WRAP instructor. You will get feedback from you recitation instructor but *not a grade*.

DPPR due Friday. Your grade is coming from your WRAP instructor. You will get feedback from you recitation instructor but *not a grade*.

Exam 1 the week after spring break

DPPR due Friday. Your grade is coming from your WRAP instructor. You will get feedback from you recitation instructor but *not a grade*.


Exam 1 the week after spring break

if you need to take a makeup exam, you need to let us know by **this Friday**

**the website will tell you how to do this: <http://web.mit.edu/6.1800/www/assignments/exam-1.shtml>
(spoilers, it's a google form)**

1970s: ARPANet 1978: flexibility and layering early 80s: growth → change late 80s: growth → problems 1993: commercialization

hosts.txt distance-vector routing TCP, UDP OSPF, EGP, DNS congestion collapse (which led to congestion control) policy routing CIDR



CAIDA's IPv4 AS Core,
January 2020

(<https://www.caida.org/projects/cartography/as-core/2020/>)

today: turning our attention away from the Internet to datacenter networks. what's different in this environment, and why does it matter?

application

the things that actually generate traffic

transport

sharing the network, reliability (or not)
examples: TCP, UDP

network


naming, addressing, routing
examples: IP

link


communication between two directly-connected nodes
examples: ethernet, bluetooth, 802.11 (wifi)



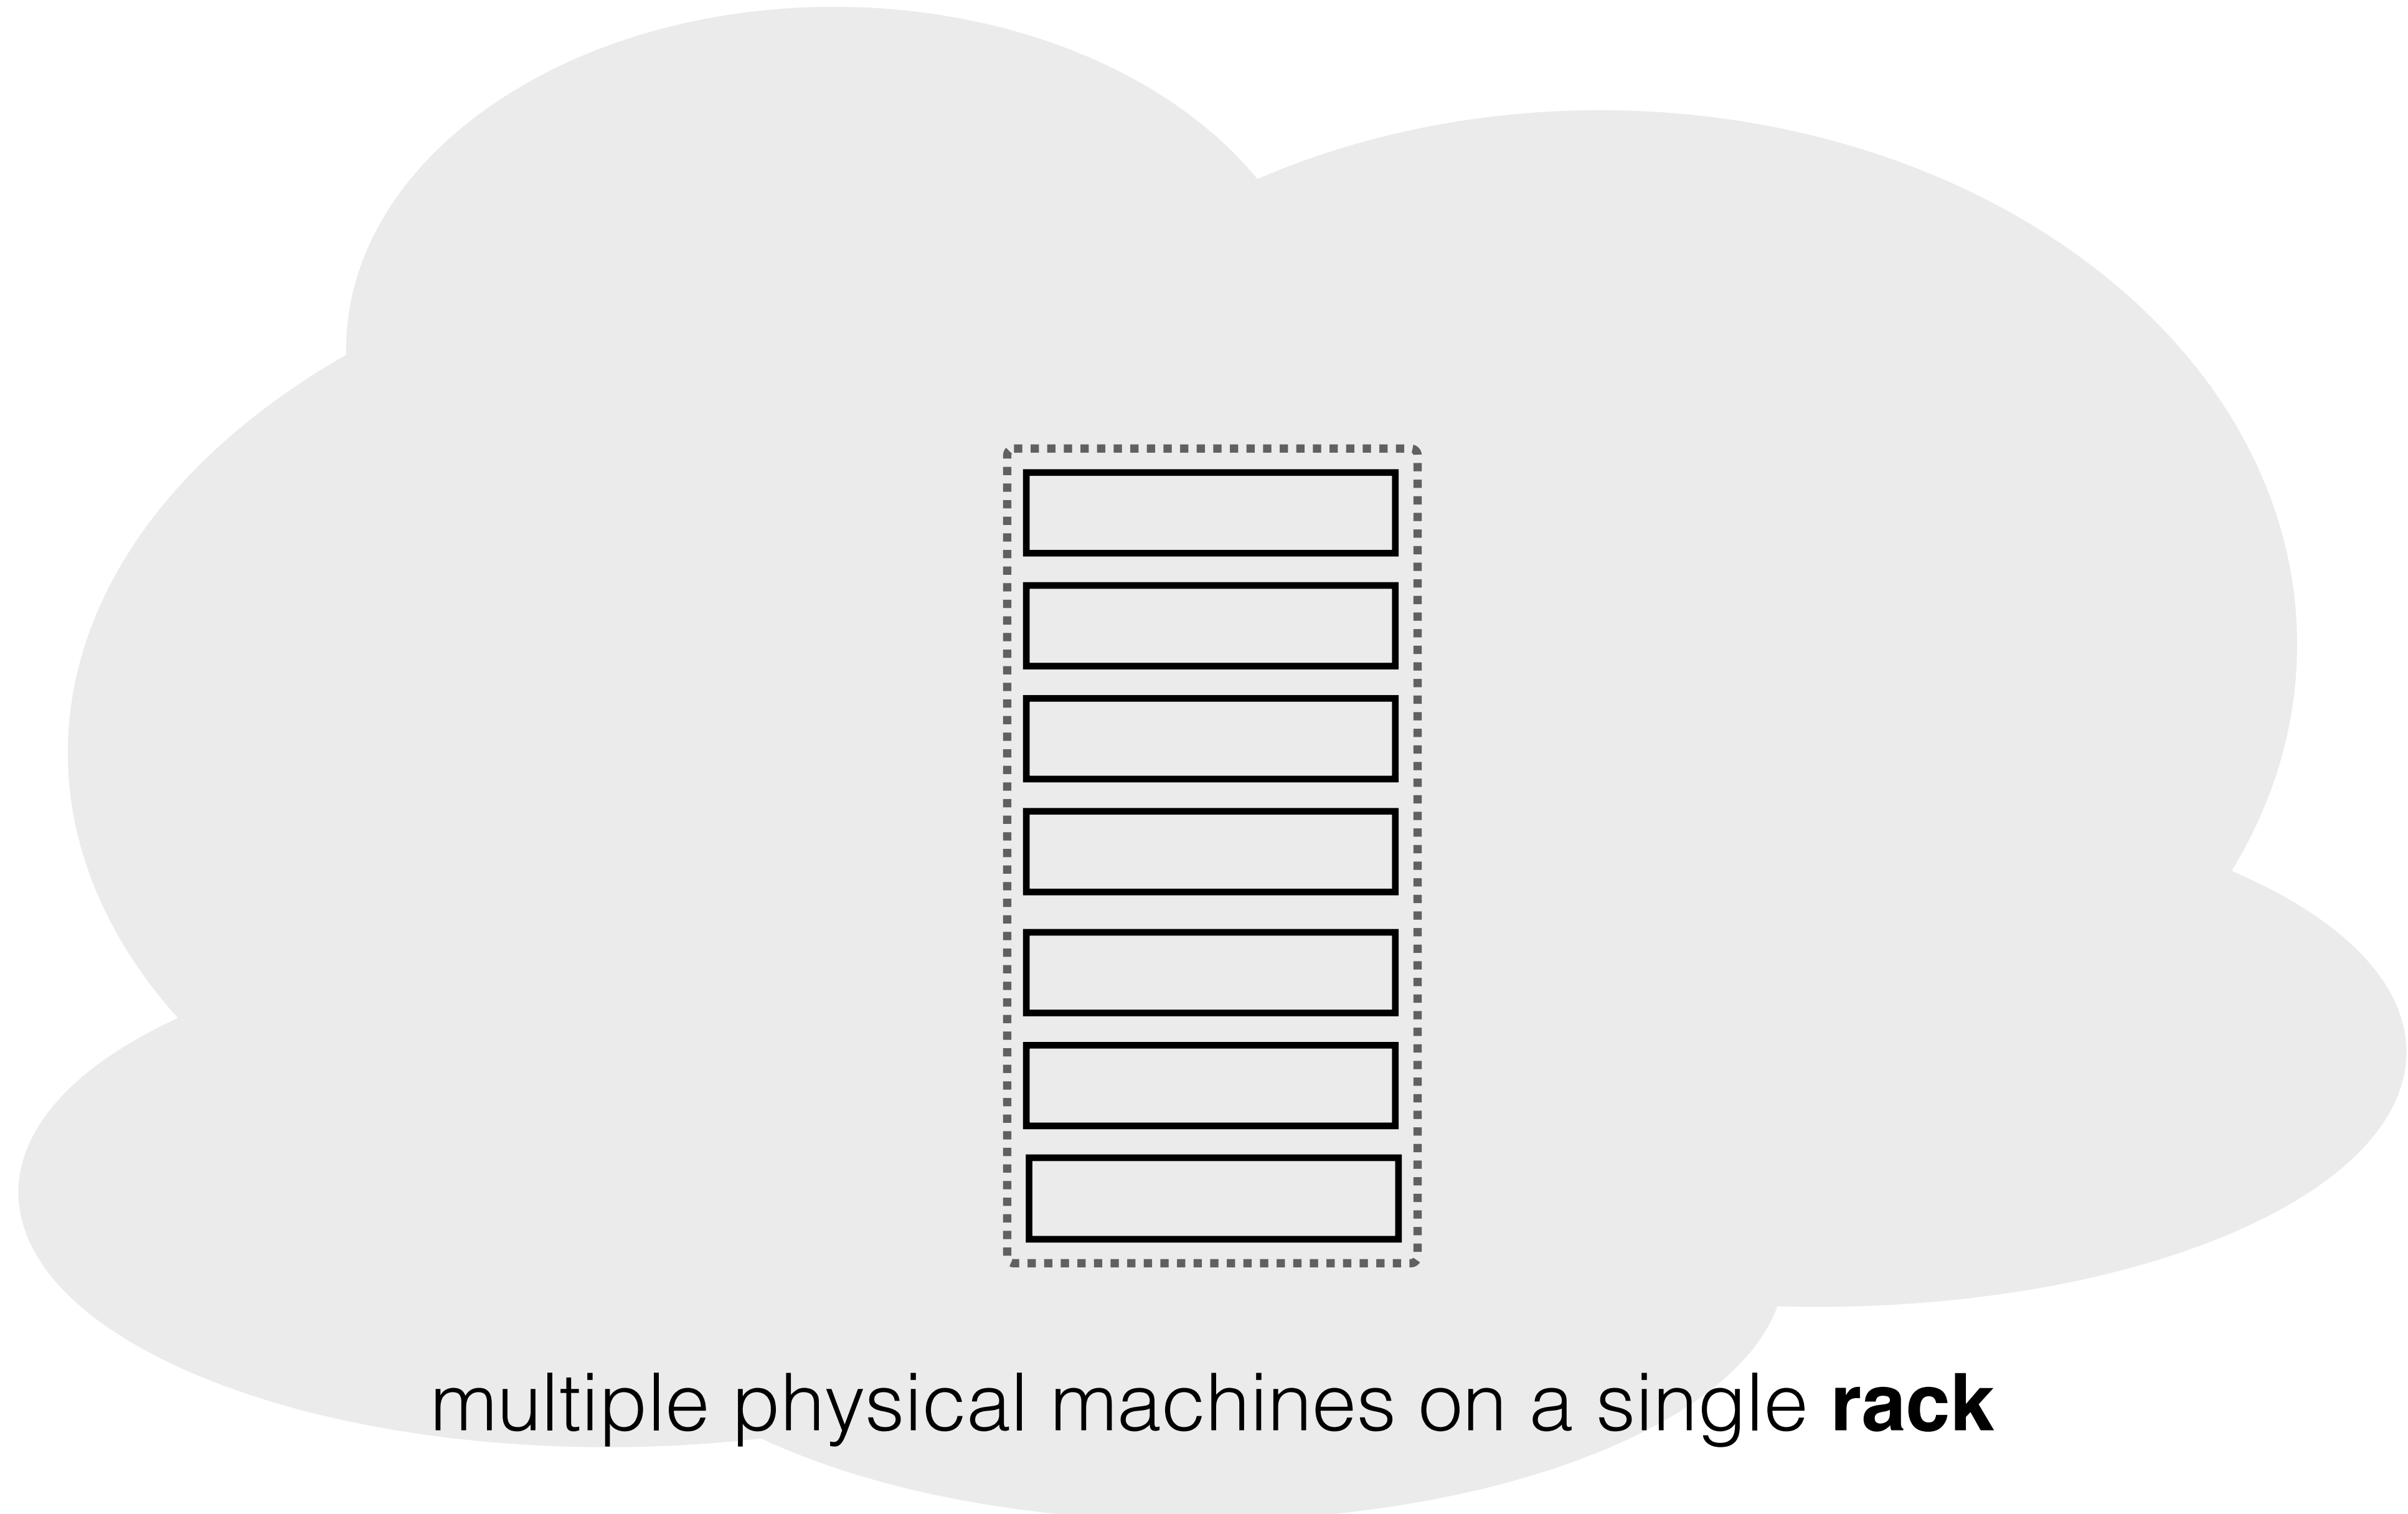
datacenter networks back many of the services you use every day



datacenter networks back many of the services you use every day




datacenter networks back many of the services you use every day




datacenter networks back many of the services you use every day



datacenter networks back many of the services you use every day




datacenter networks back many of the services you use every day




we need a way to communicate across racks. we control this network, so we can design its topology

datacenter networks back many of the services you use every day




we need a way to communicate across racks. we control this network, so we can design its topology

datacenter networks back many of the services you use every day




question: are there any downsides to this particular topology?

datacenter networks back many of the services you use every day



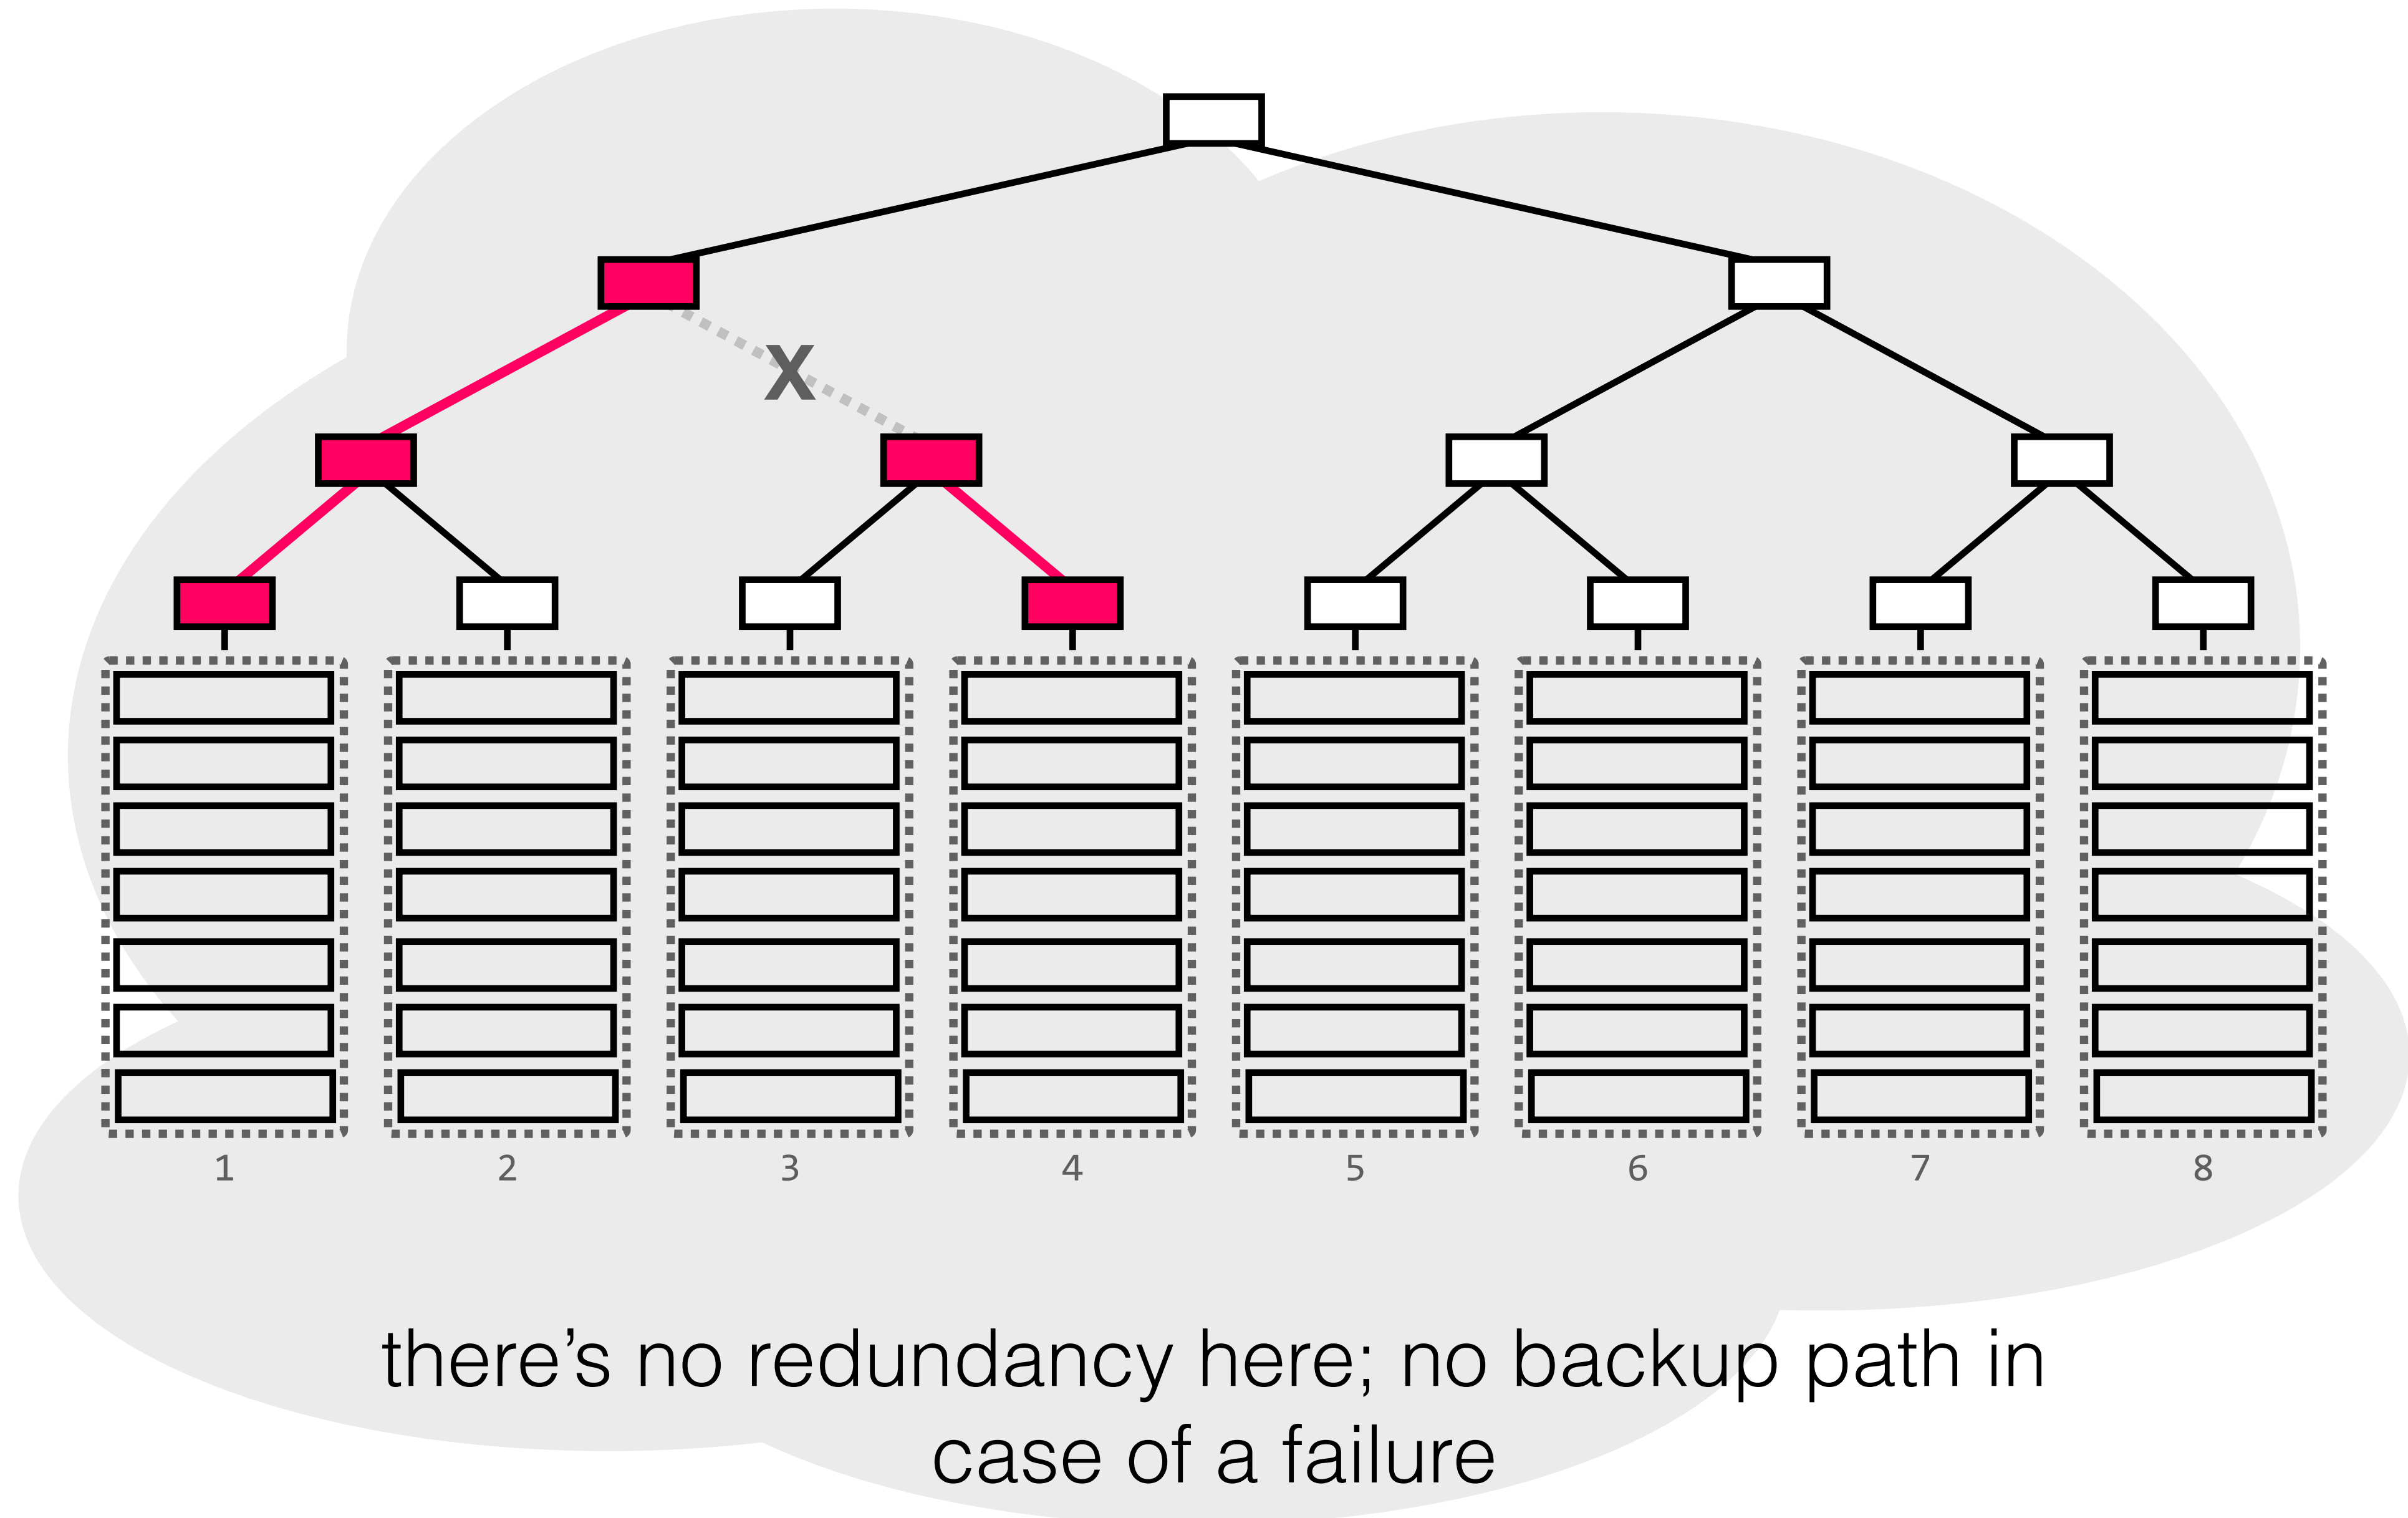
there's no redundancy here; no backup path in case of a failure

datacenter networks back many of the services you use every day




there's no redundancy here; no backup path in case of a failure

datacenter networks back many of the services you use every day




datacenter networks back many of the services you use every day




certain topologies can add a lot of redundancy
this is an example of a *clos* topology

datacenter networks back many of the services you use every day




certain topologies can add a lot of redundancy
this is an example of a *clos* topology

datacenter networks back many of the services you use every day




certain topologies can add a lot of redundancy
this is an example of a *clos* topology

datacenter networks back many of the services you use every day




certain topologies can add a lot of redundancy
this is an example of a *clos* topology

datacenter networks back many of the services you use every day




certain topologies can add a lot of redundancy
this is an example of a *clos* topology

datacenter networks back many of the services you use every day




certain topologies can add a lot of redundancy
this is an example of a *clos* topology

datacenter networks back many of the services you use every day




certain topologies can add a lot of redundancy
this is an example of a *clos* topology

datacenter networks back many of the services you use every day




certain topologies can add a lot of redundancy
this is an example of a *clos* topology

datacenter networks back many of the services you use every day




certain topologies can add a lot of redundancy
this is an example of a *clos* topology

datacenter networks back many of the services you use every day




certain topologies can add a lot of redundancy
this is an example of a *clos* topology

datacenter networks back many of the services you use every day




standard routing protocols will pick a single path
and stick to it until something changes; **multi-path
routing** can load-balance across paths

datacenter networks back many of the services you use every day




standard routing protocols will pick a single path
and stick to it until something changes; **multi-path
routing** can load-balance across paths

datacenter networks back many of the services you use every day




standard routing protocols will pick a single path
and stick to it until something changes; **multi-path
routing** can load-balance across paths

datacenter networks back many of the services you use every day




question: suppose we used round-robin scheduling to send packets from a single TCP flow across these two paths. what might happen?

datacenter networks back many of the services you use every day



question: suppose we used round-robin scheduling to send packets from a single TCP flow across these two paths. what might happen?


datacenter networks back many of the services you use every day



multipath routing can help us load-balance, but we need to be careful about how we divide traffic across the paths

e.g., dividing a single TCP flow across multiple paths will make congestion control more difficult

datacenter networks back many of the services you use every day




multipath routing can help us load-balance, but we need to be careful about how we divide traffic across the paths

e.g., dividing a single TCP flow across multiple paths will make congestion control more difficult

datacenter networks back many of the services you use every day

one approach is to **keep each TCP flow on a single path**, but utilize different paths for different flows
“ECMP” is one protocol that does this




multipath routing can help us load-balance, but we need to be careful about how we divide traffic across the paths

e.g., dividing a single TCP flow across multiple paths will make congestion control more difficult

datacenter networks back many of the services you use every day

one approach is to **keep each TCP flow on a single path**, but utilize different paths for different flows
“ECMP” is one protocol that does this

however, these protocols are often not “congestion aware”; we could still end up with way more traffic on one path than the other




multipath routing can help us load-balance, but we need to be careful about how we divide traffic across the paths

e.g., dividing a single TCP flow across multiple paths will make congestion control more difficult

datacenter networks back many of the services you use every day

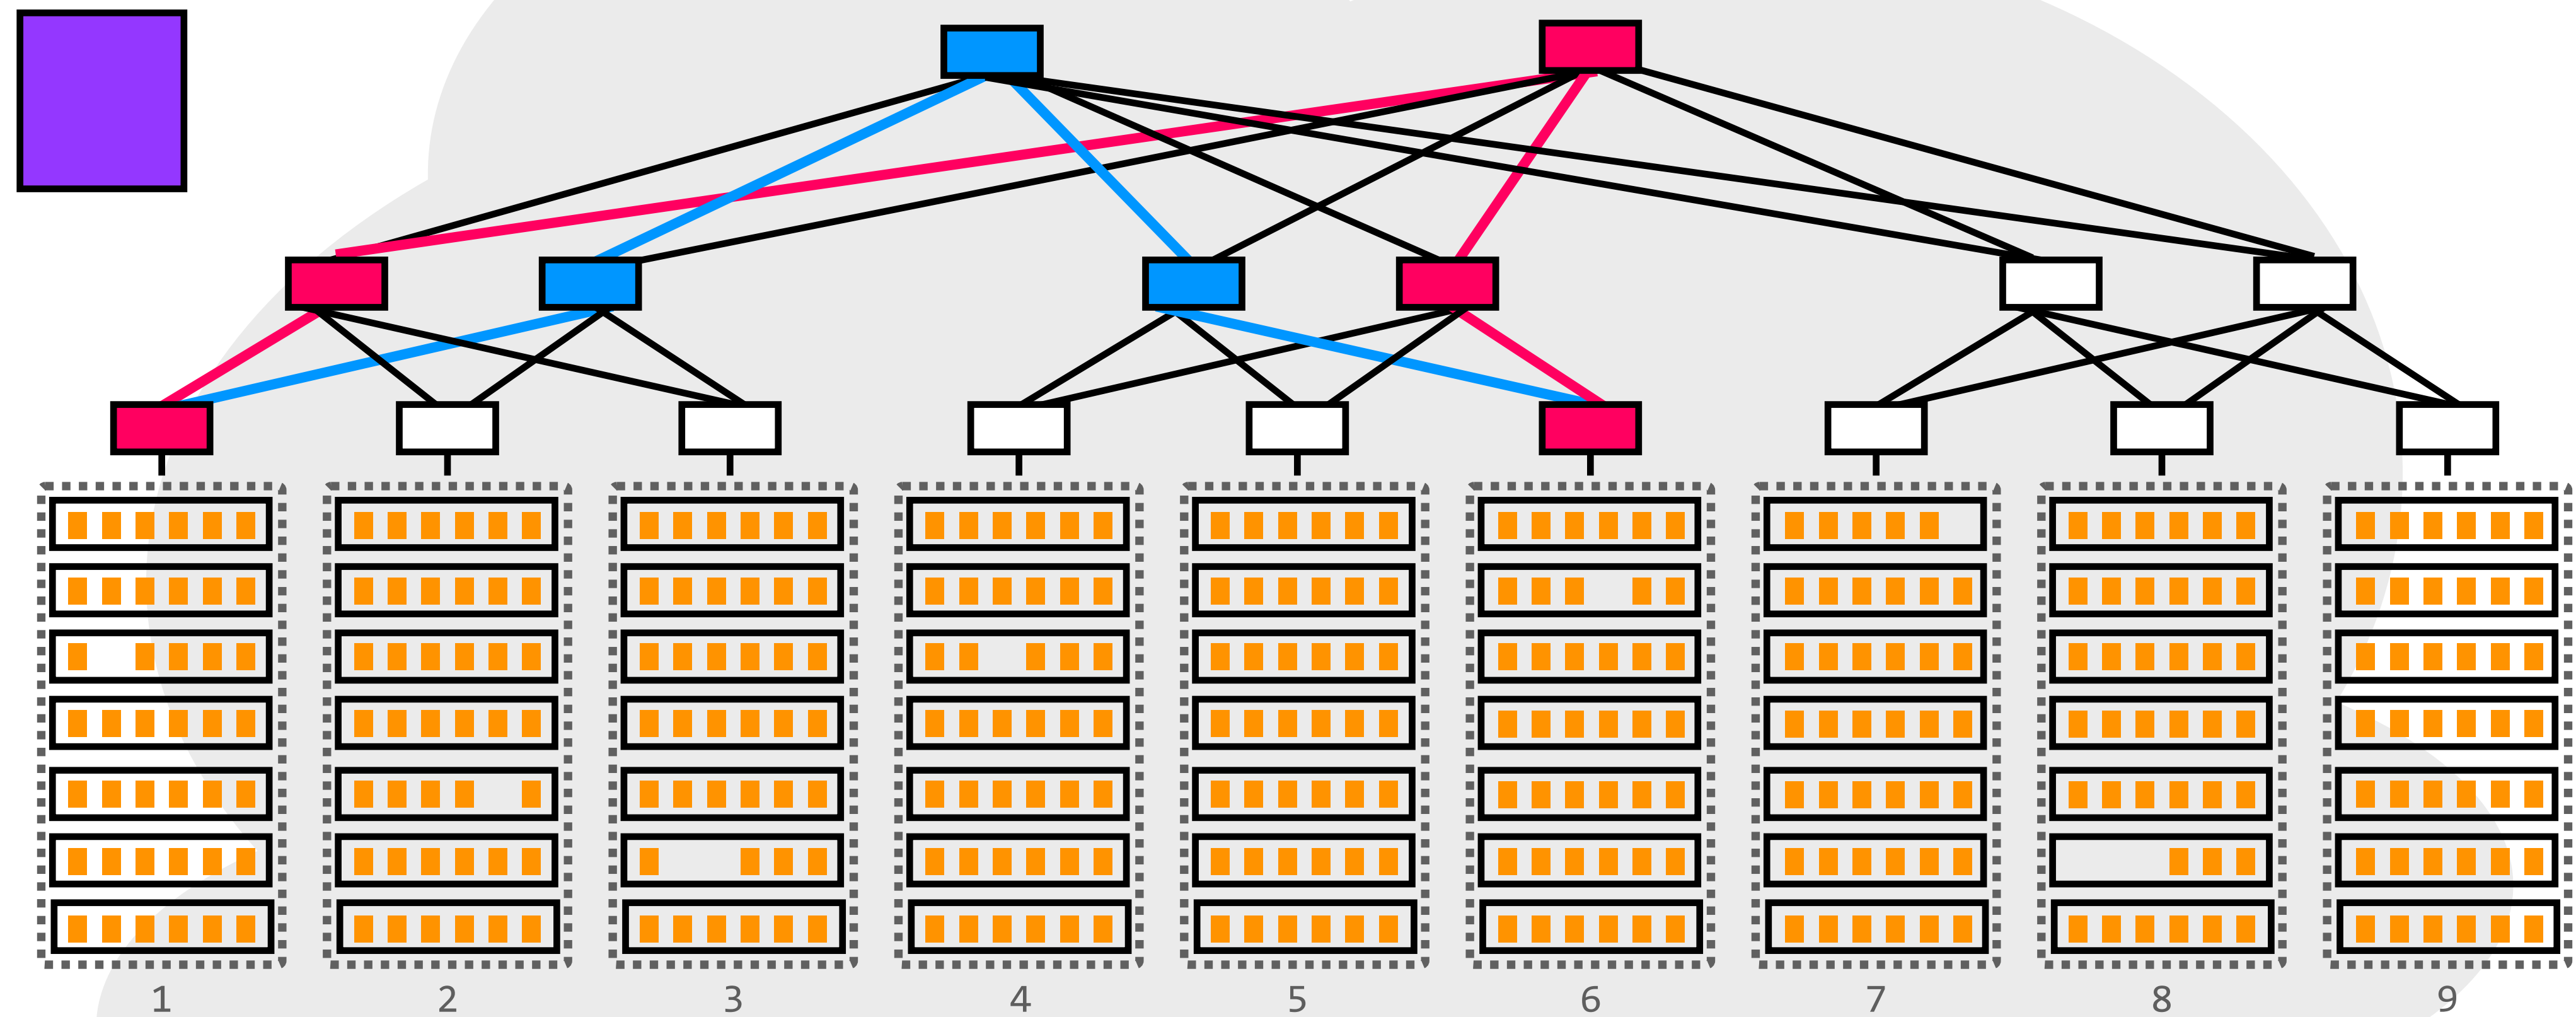
many datacenters use a **centralized controller** to manage routing and other things



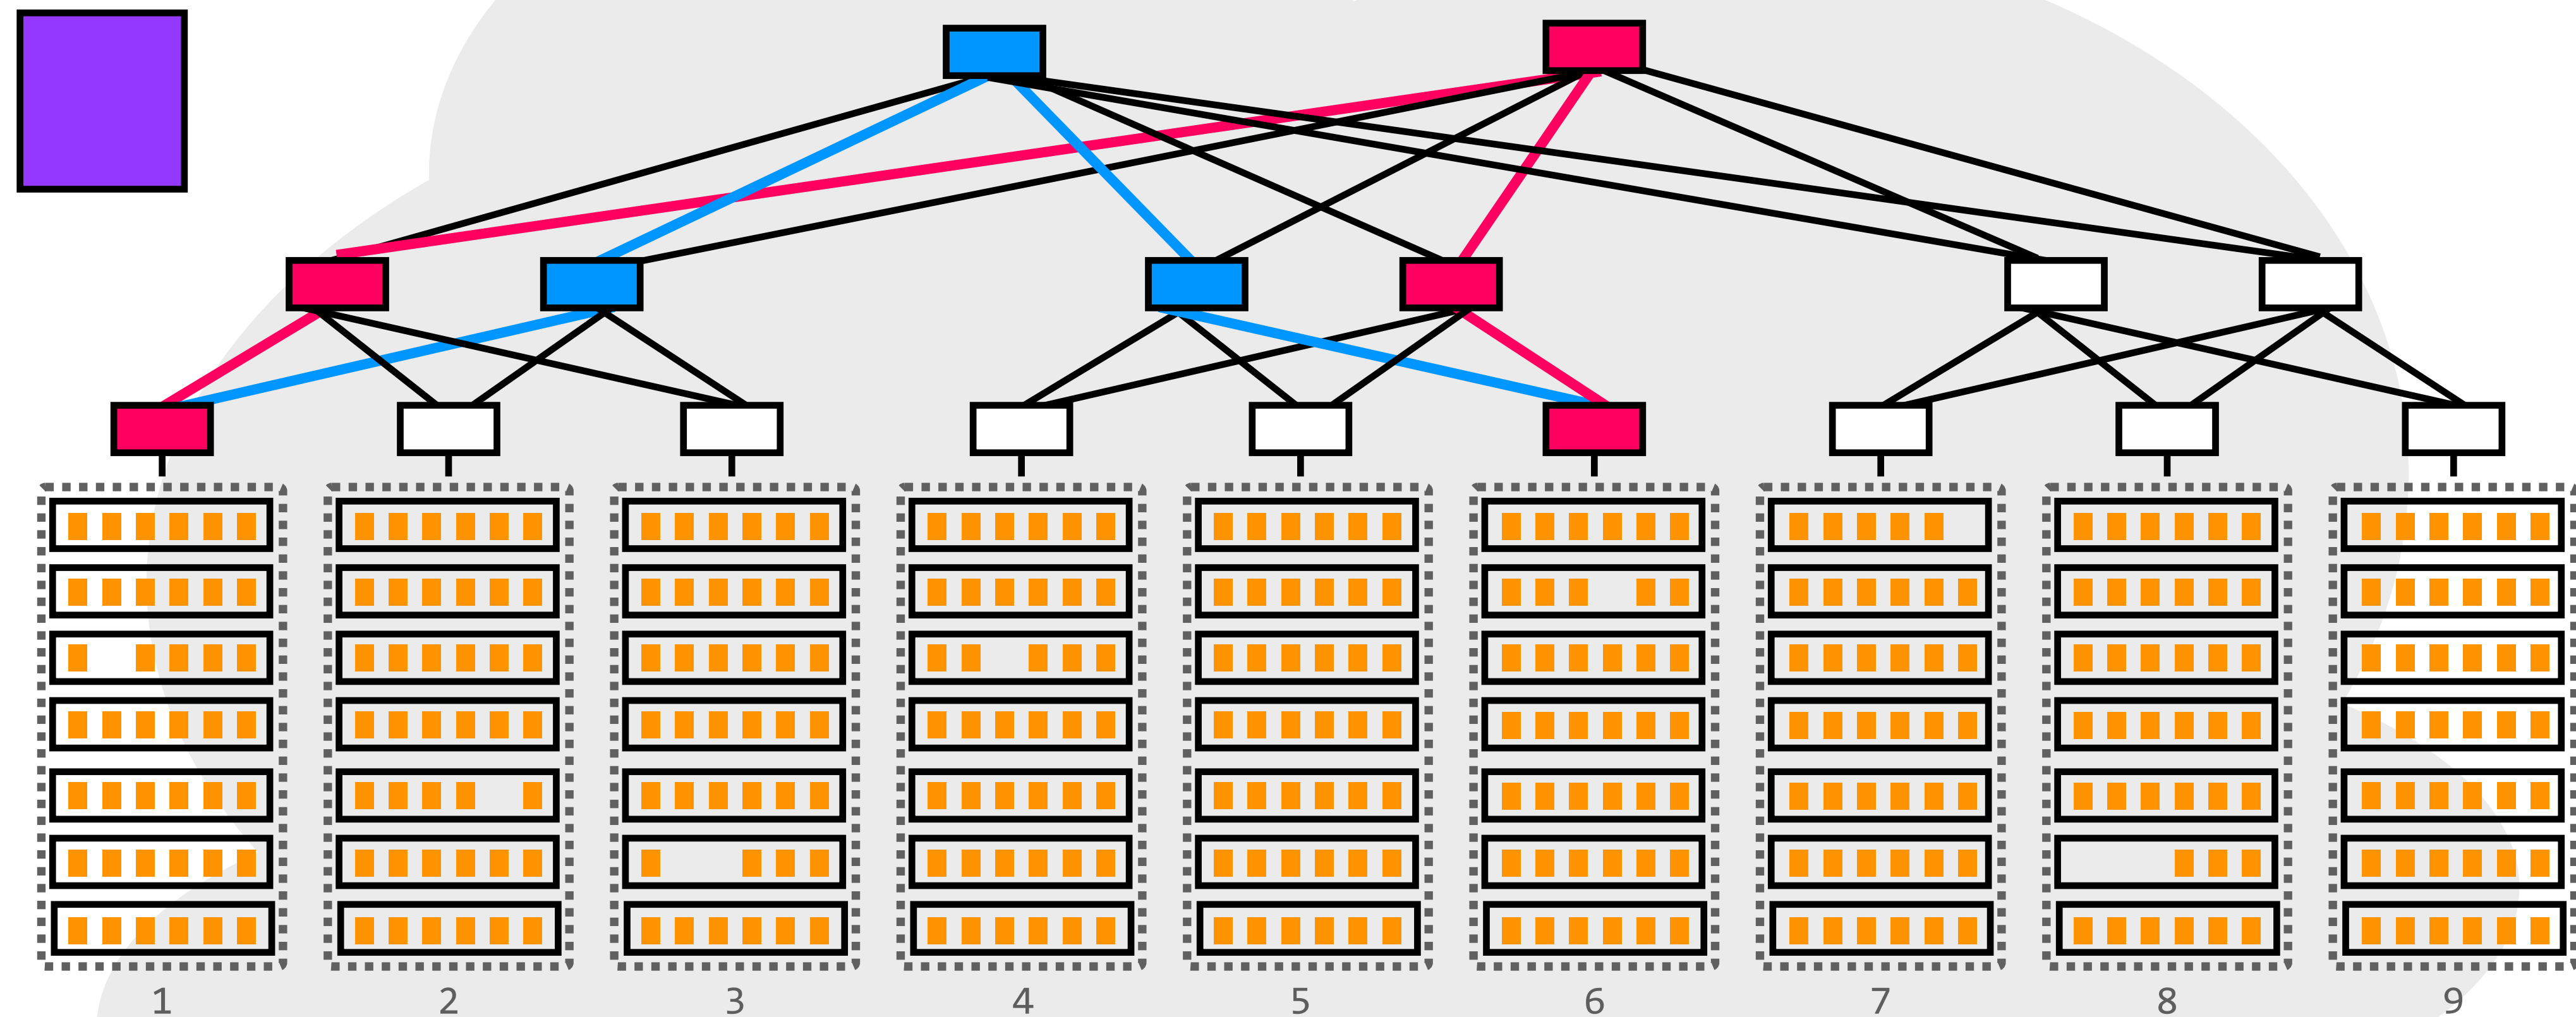
multipath routing can help us load-balance, but we need to be careful about how we divide traffic across the paths

e.g., dividing a single TCP flow across multiple paths will make congestion control more difficult

datacenter networks back many of the services you use every day

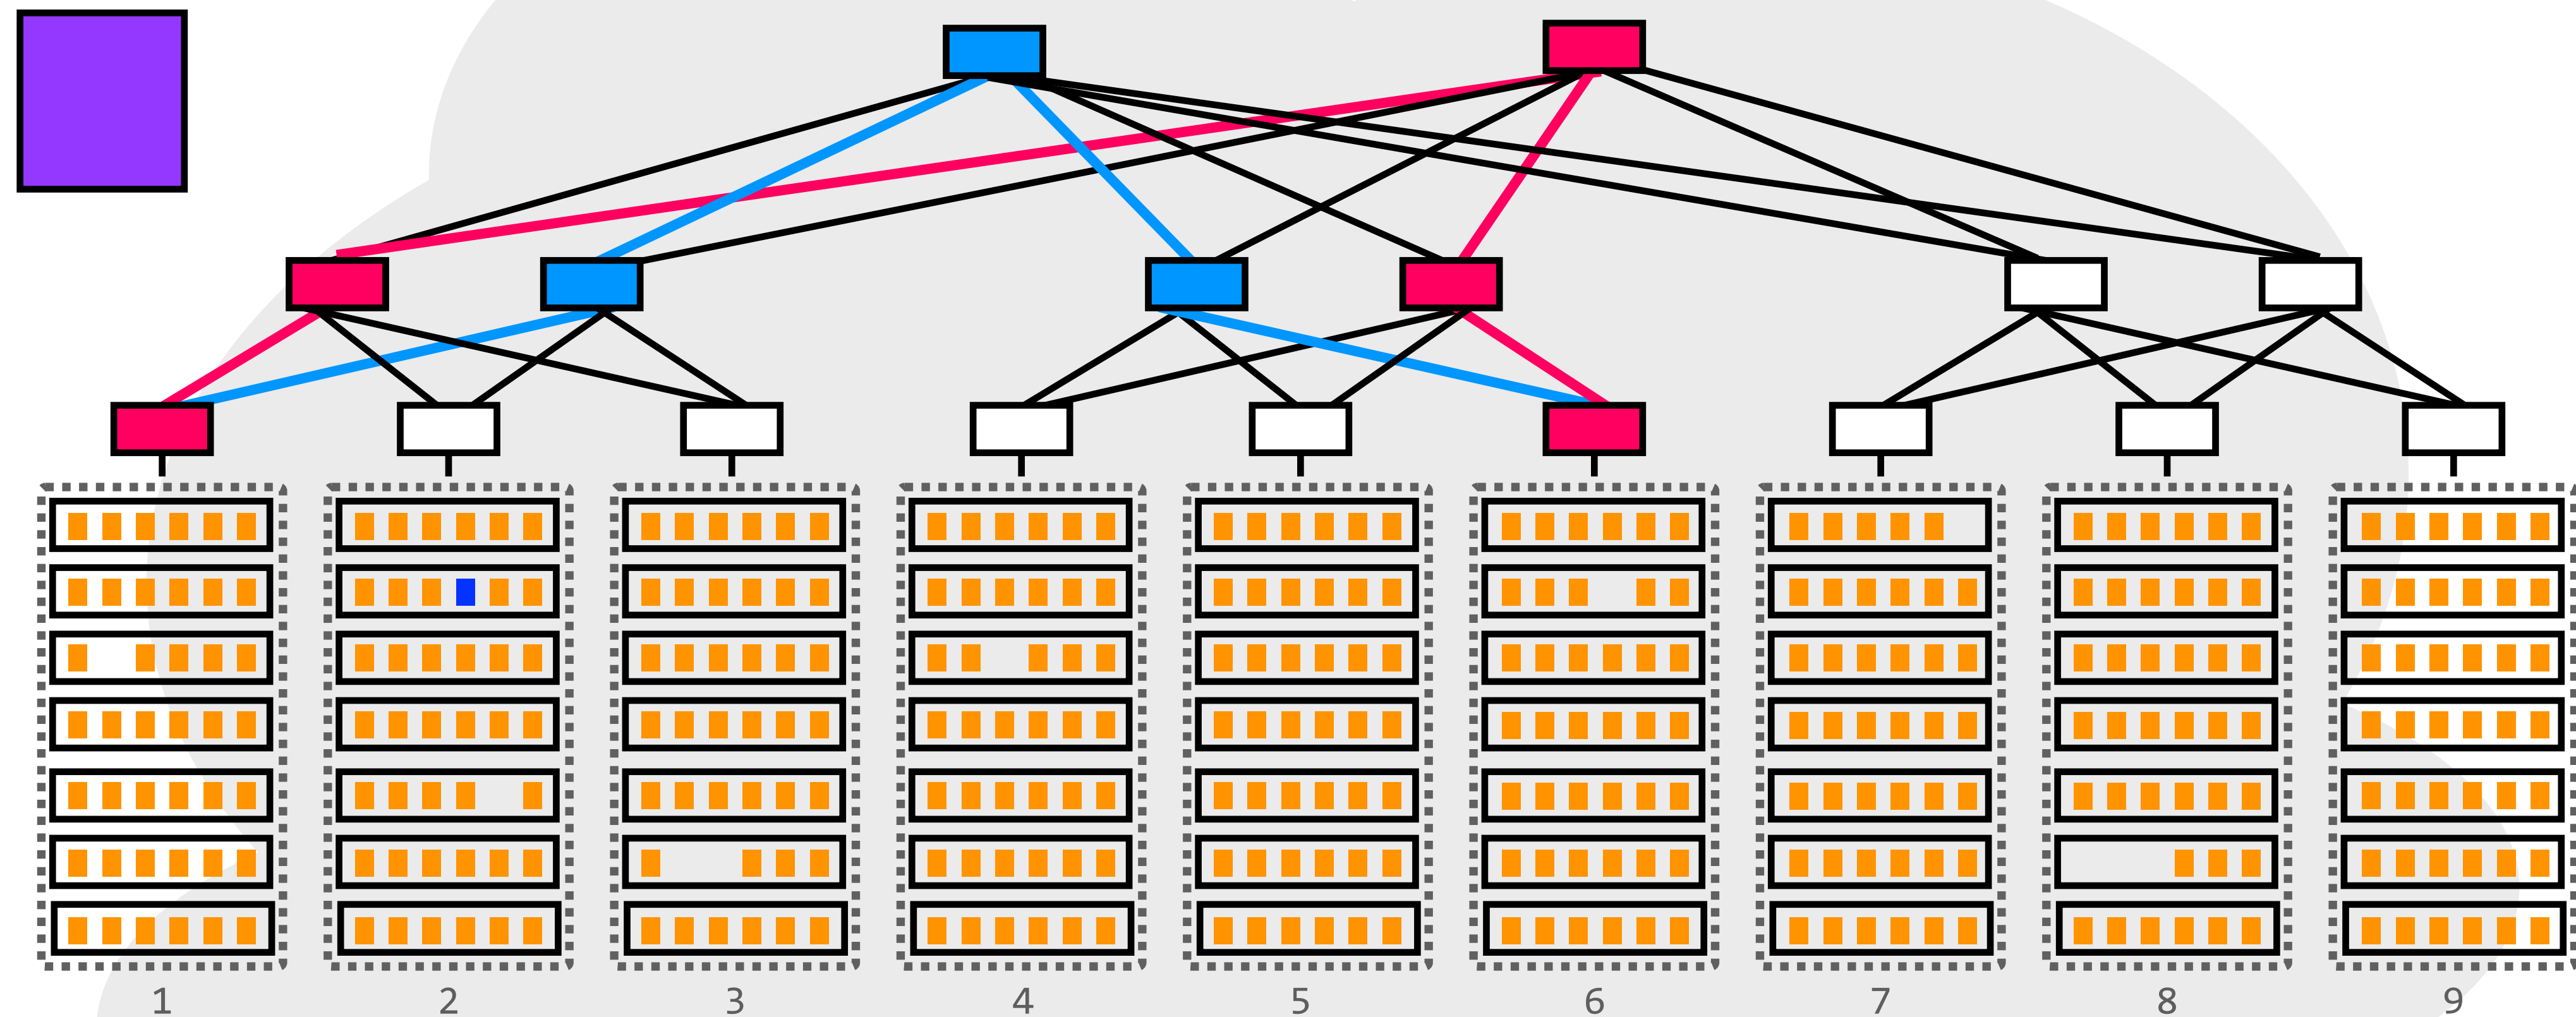


datacenter networks back many of the services you use every day



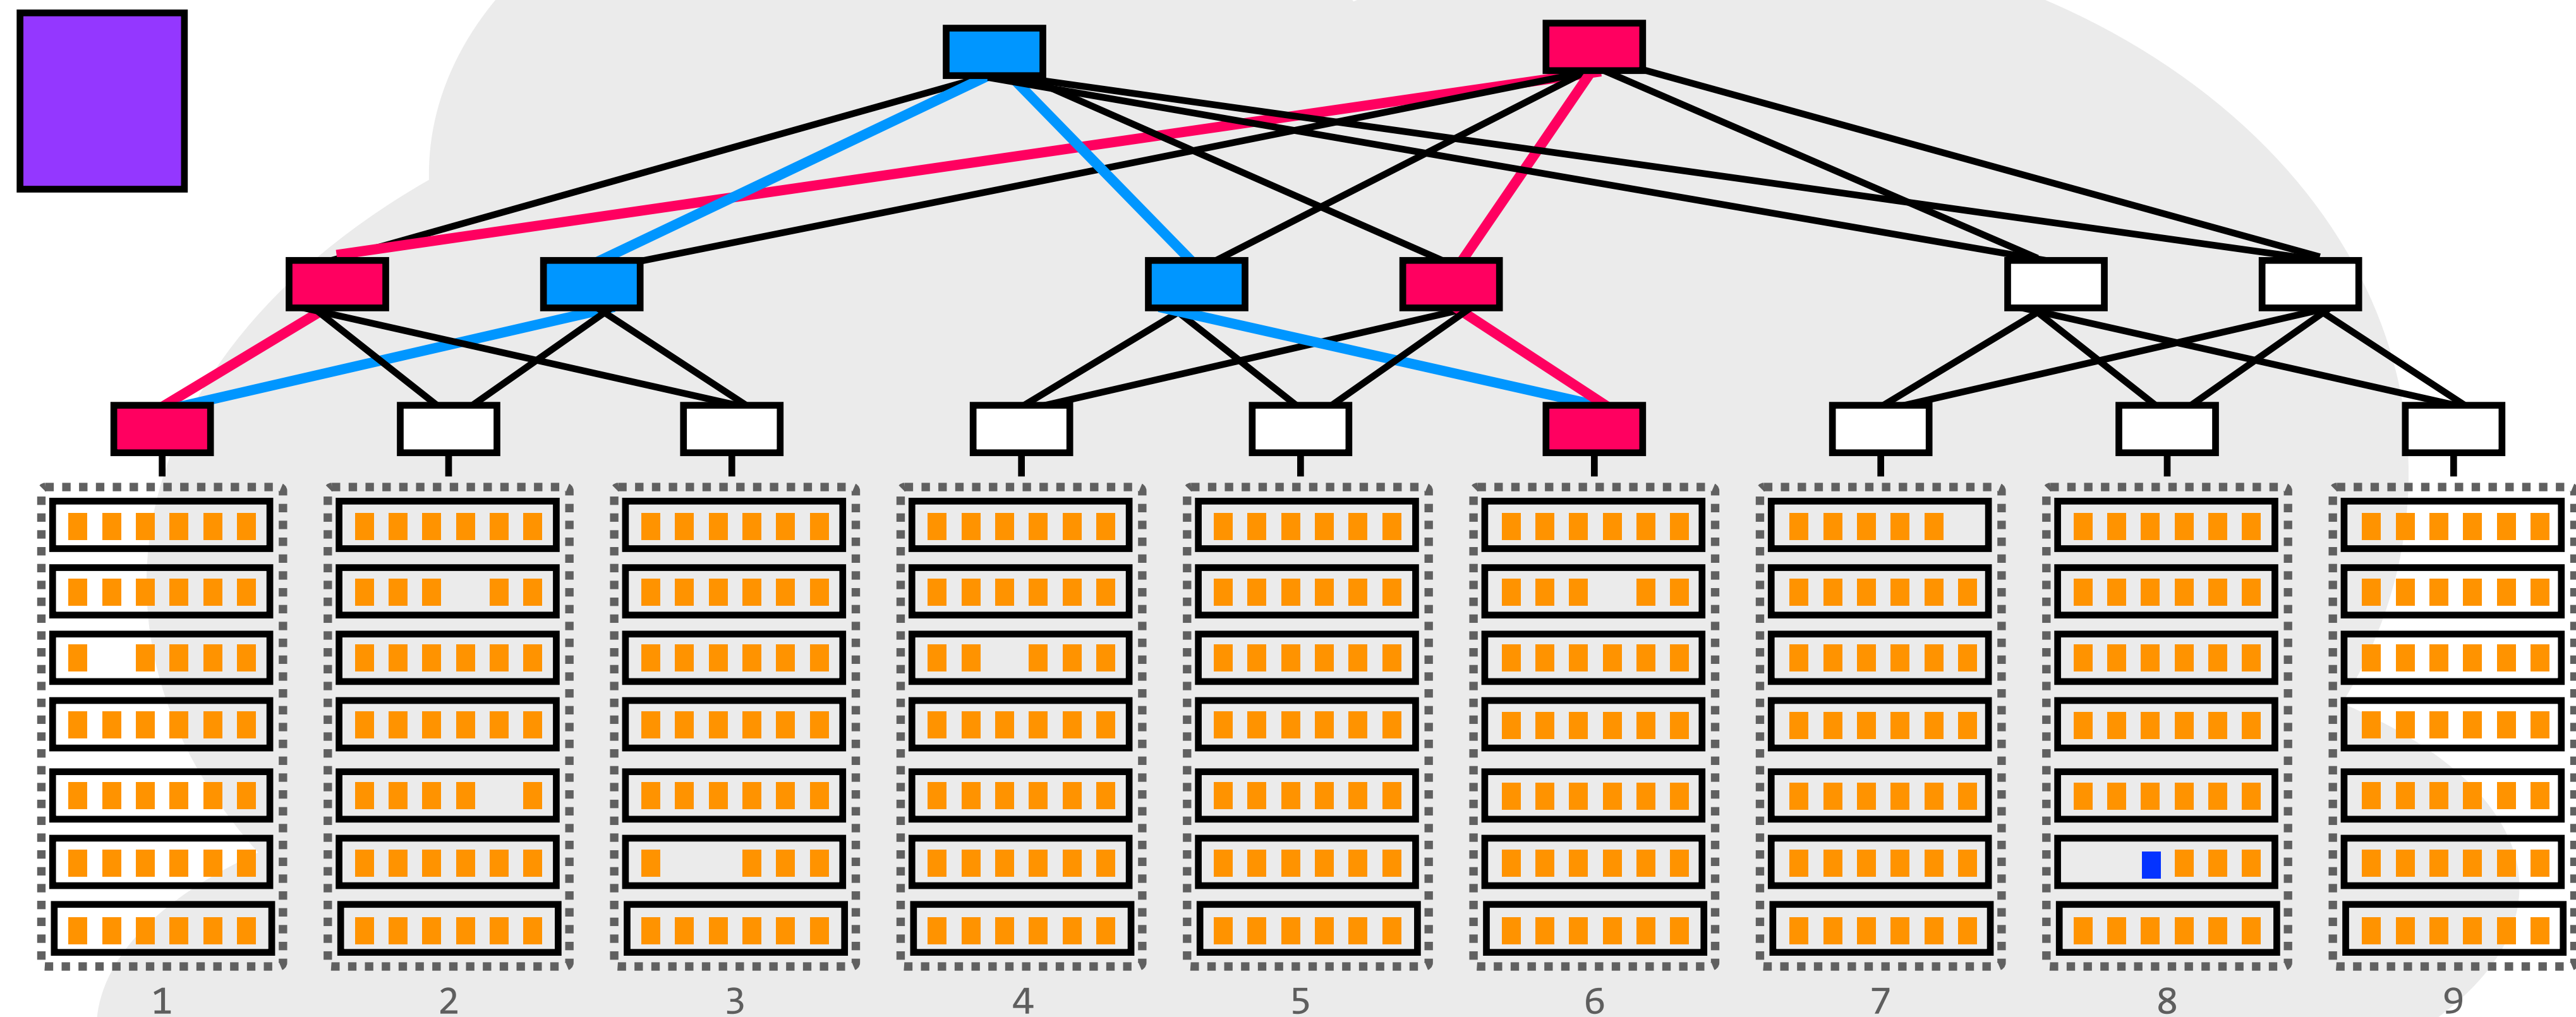
each physical machine can host multiple **virtual machines**, which sometimes need to be moved around in the network

datacenter networks back many of the services you use every day



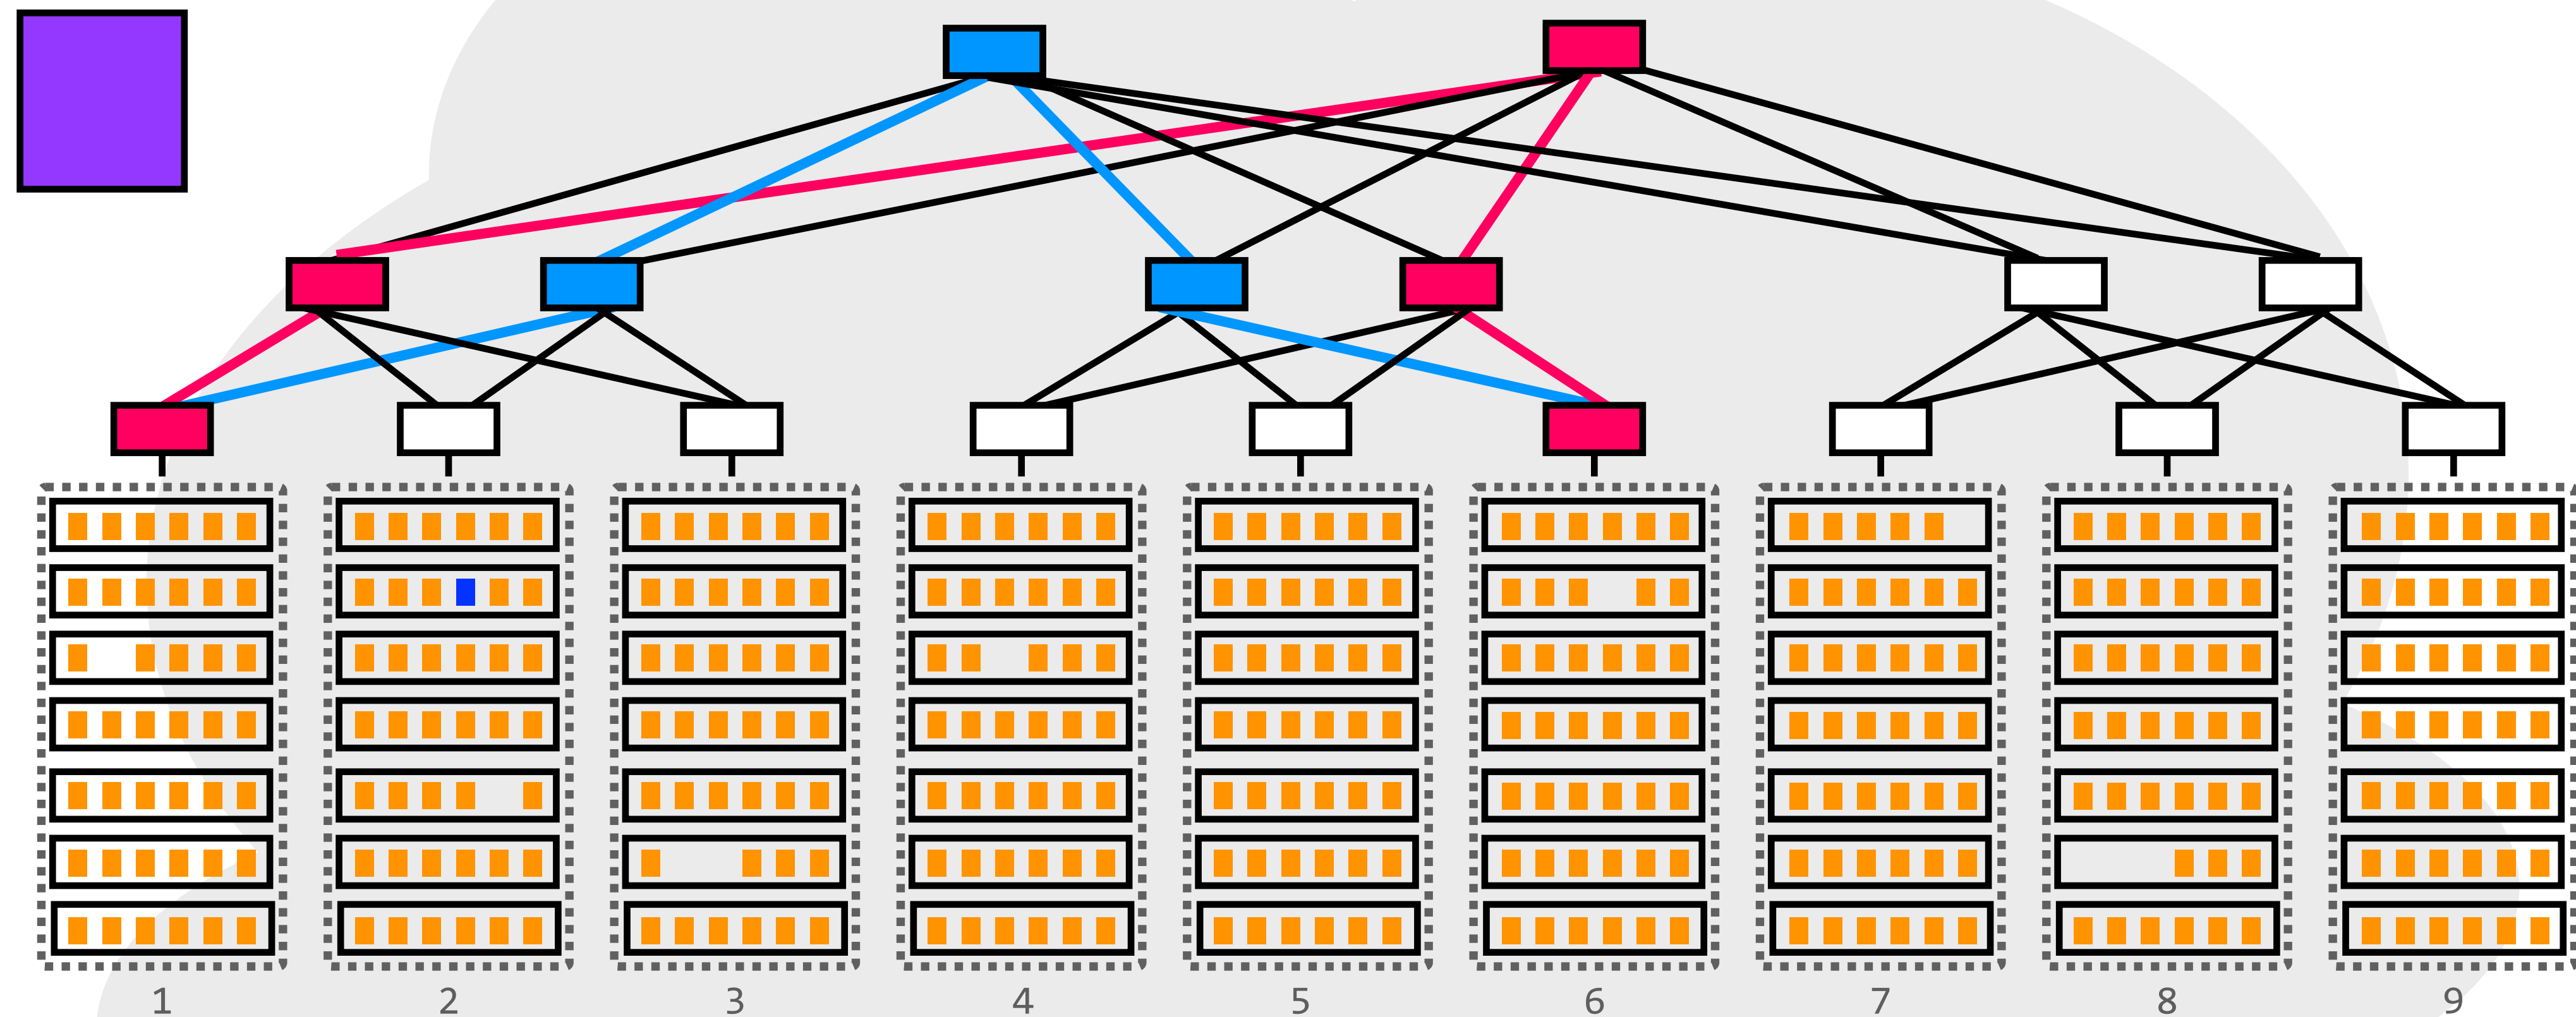
each physical machine can host multiple **virtual machines**, which sometimes need to be moved around in the network

datacenter networks back many of the services you use every day



each physical machine can host multiple **virtual machines**, which sometimes need to be moved around in the network

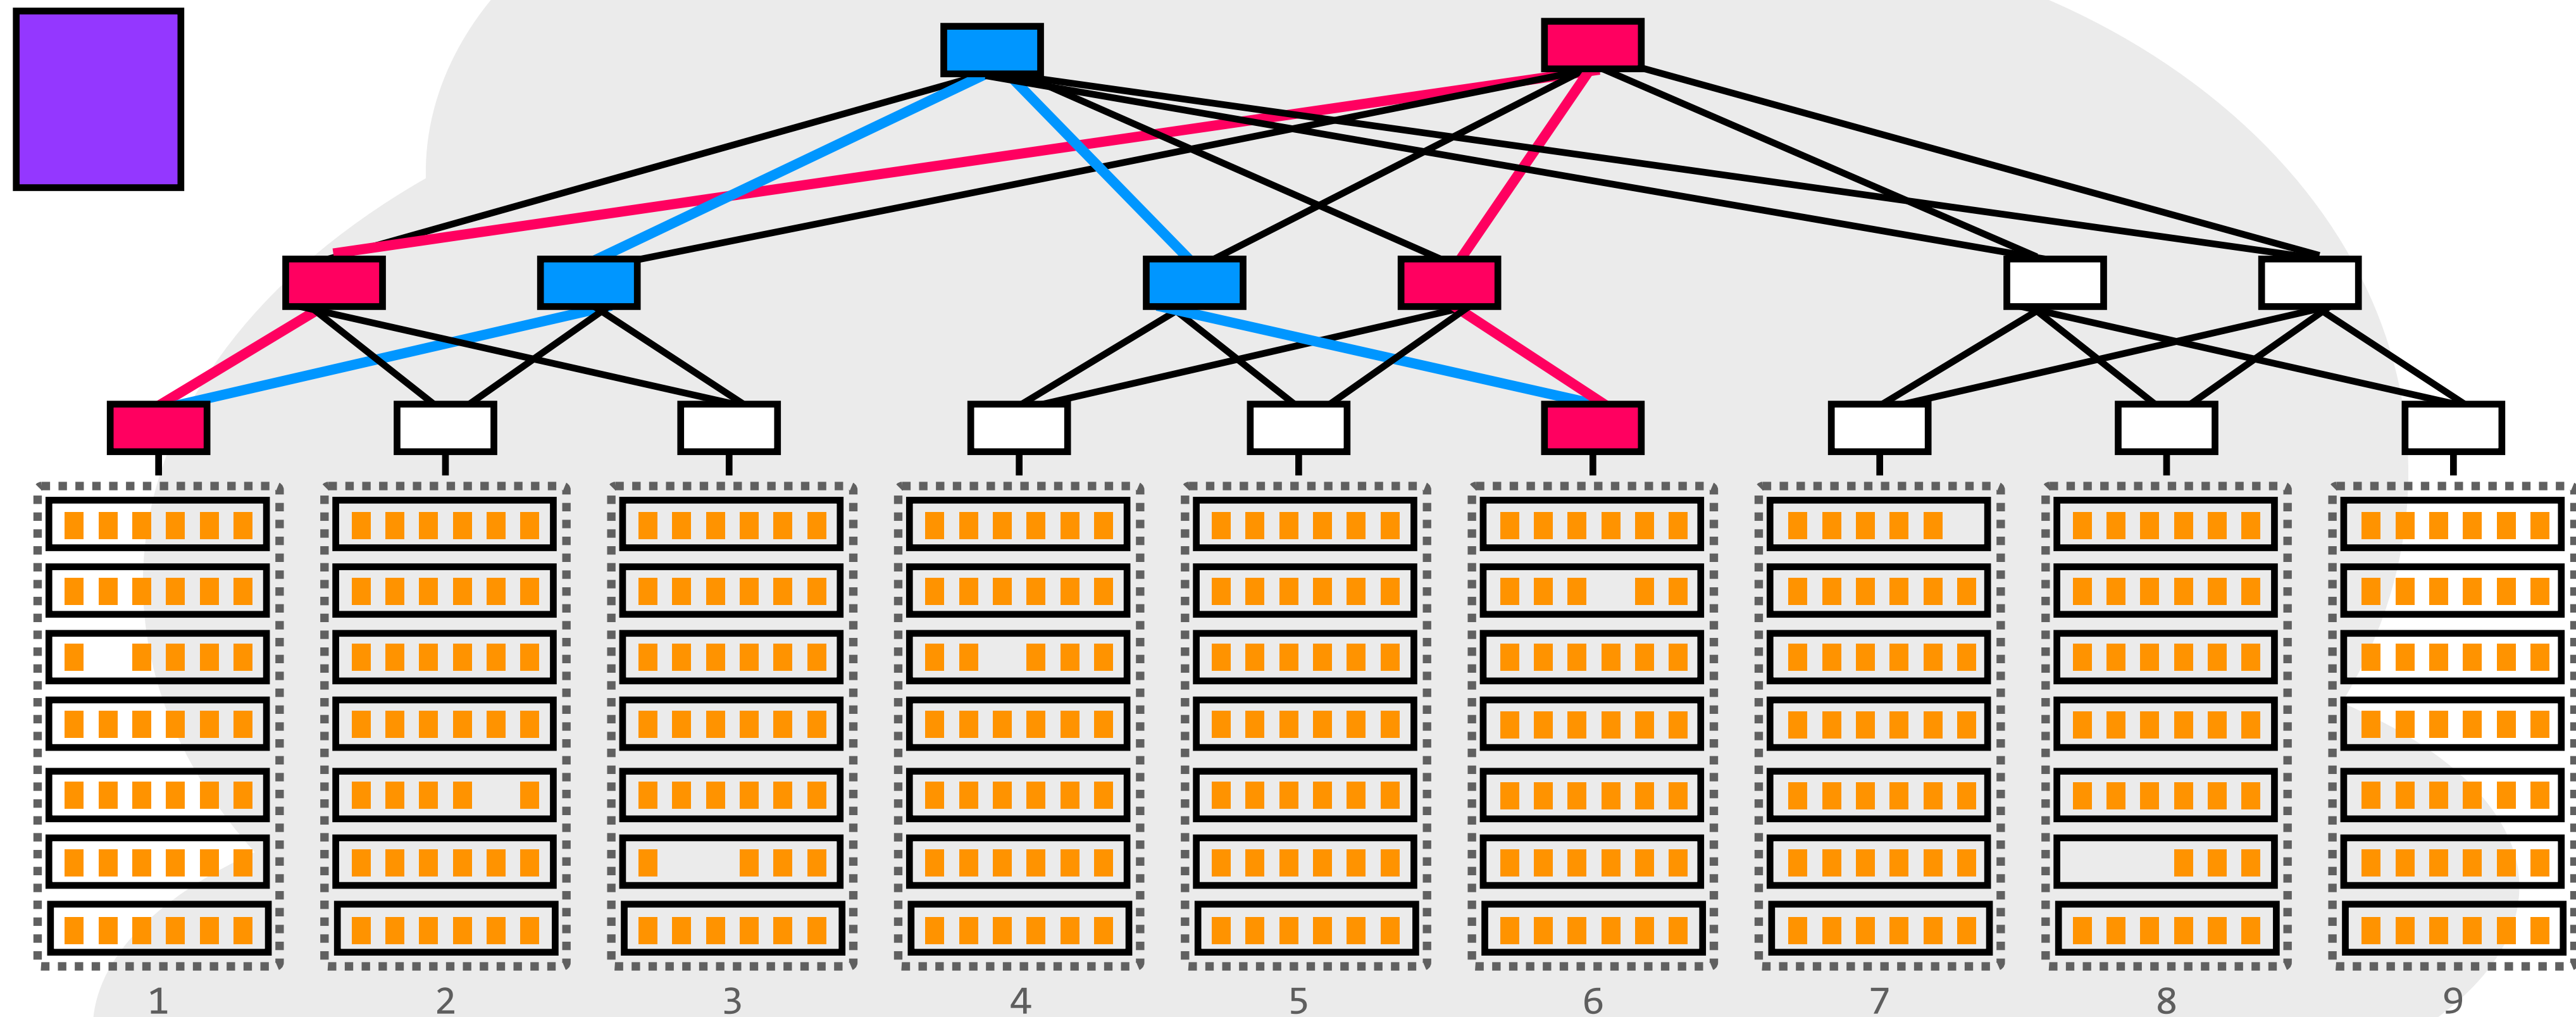
datacenter networks back many of the services you use every day



each physical machine can host multiple **virtual machines**, which sometimes need to be moved around in the network


datacenters need to decouple a VM's name from its physical location in order to make this work

datacenter networks back many of the services you use every day



because datacenter networks are under the control of a single administrative entity, we have a level of control over the network that we simply don't have on the Internet

6.1800 in the news



A New Surge in Power Use Is Threatening U.S. Climate Goals

A boom in data centers and factories is straining electric grids and propping up fossil fuels.

By Brad Plumer and Nadja Popovich March 14, 2024

6.1800 in the news

For much of the 20th century, America's electricity use increased steadily and utilities built plenty of coal, gas and nuclear plants in response. But starting in the mid-2000s, demand flattened. The economy and population kept expanding, but factories, lightbulbs and even refrigerators became much more energy efficient.


Now demand is rising again, for several reasons.

The growth of remote work, video streaming and online shopping has led to a frenzied expansion of data centers across the nation. The rise of artificial intelligence is poised to accelerate that trend: By 2030, [electricity demand at U.S. data centers could triple](#), using as much power as 40 million homes, according to Boston Consulting Group.

In Northern Virginia, one of the nation's largest data center hubs, at least 75 facilities have opened since 2019 and Dominion Energy, the local utility, says data center capacity could double in just five years.

1970s: ARPAnet 1978: flexibility and layering early 80s: growth → change late 80s: growth → problems 1993: commercialization

hosts.txt distance-vector routing TCP, UDP OSPF, EGP, DNS congestion collapse policy routing CIDR



CAIDA's IPv4 AS Core,
January 2020

(<https://www.caida.org/projects/cartography/as-core/2020/>)

different networking environments give us different opportunities and impact applications in different ways

application

the things that actually generate traffic

transport

sharing the network, reliability (or not)

examples: TCP, UDP

network

naming, addressing, routing

examples: IP

link

communication between two directly-connected nodes

examples: ethernet, bluetooth, 802.11 (wifi)