CSC4200/5200 – COMPUTER NETWORKING

Instructor: Susmit Shannigrahi

GLOBAL INTERNET

sshannigrahi@tntech.edu

GTA: dereddick42@students.tntech.edu

Tennessee TECH

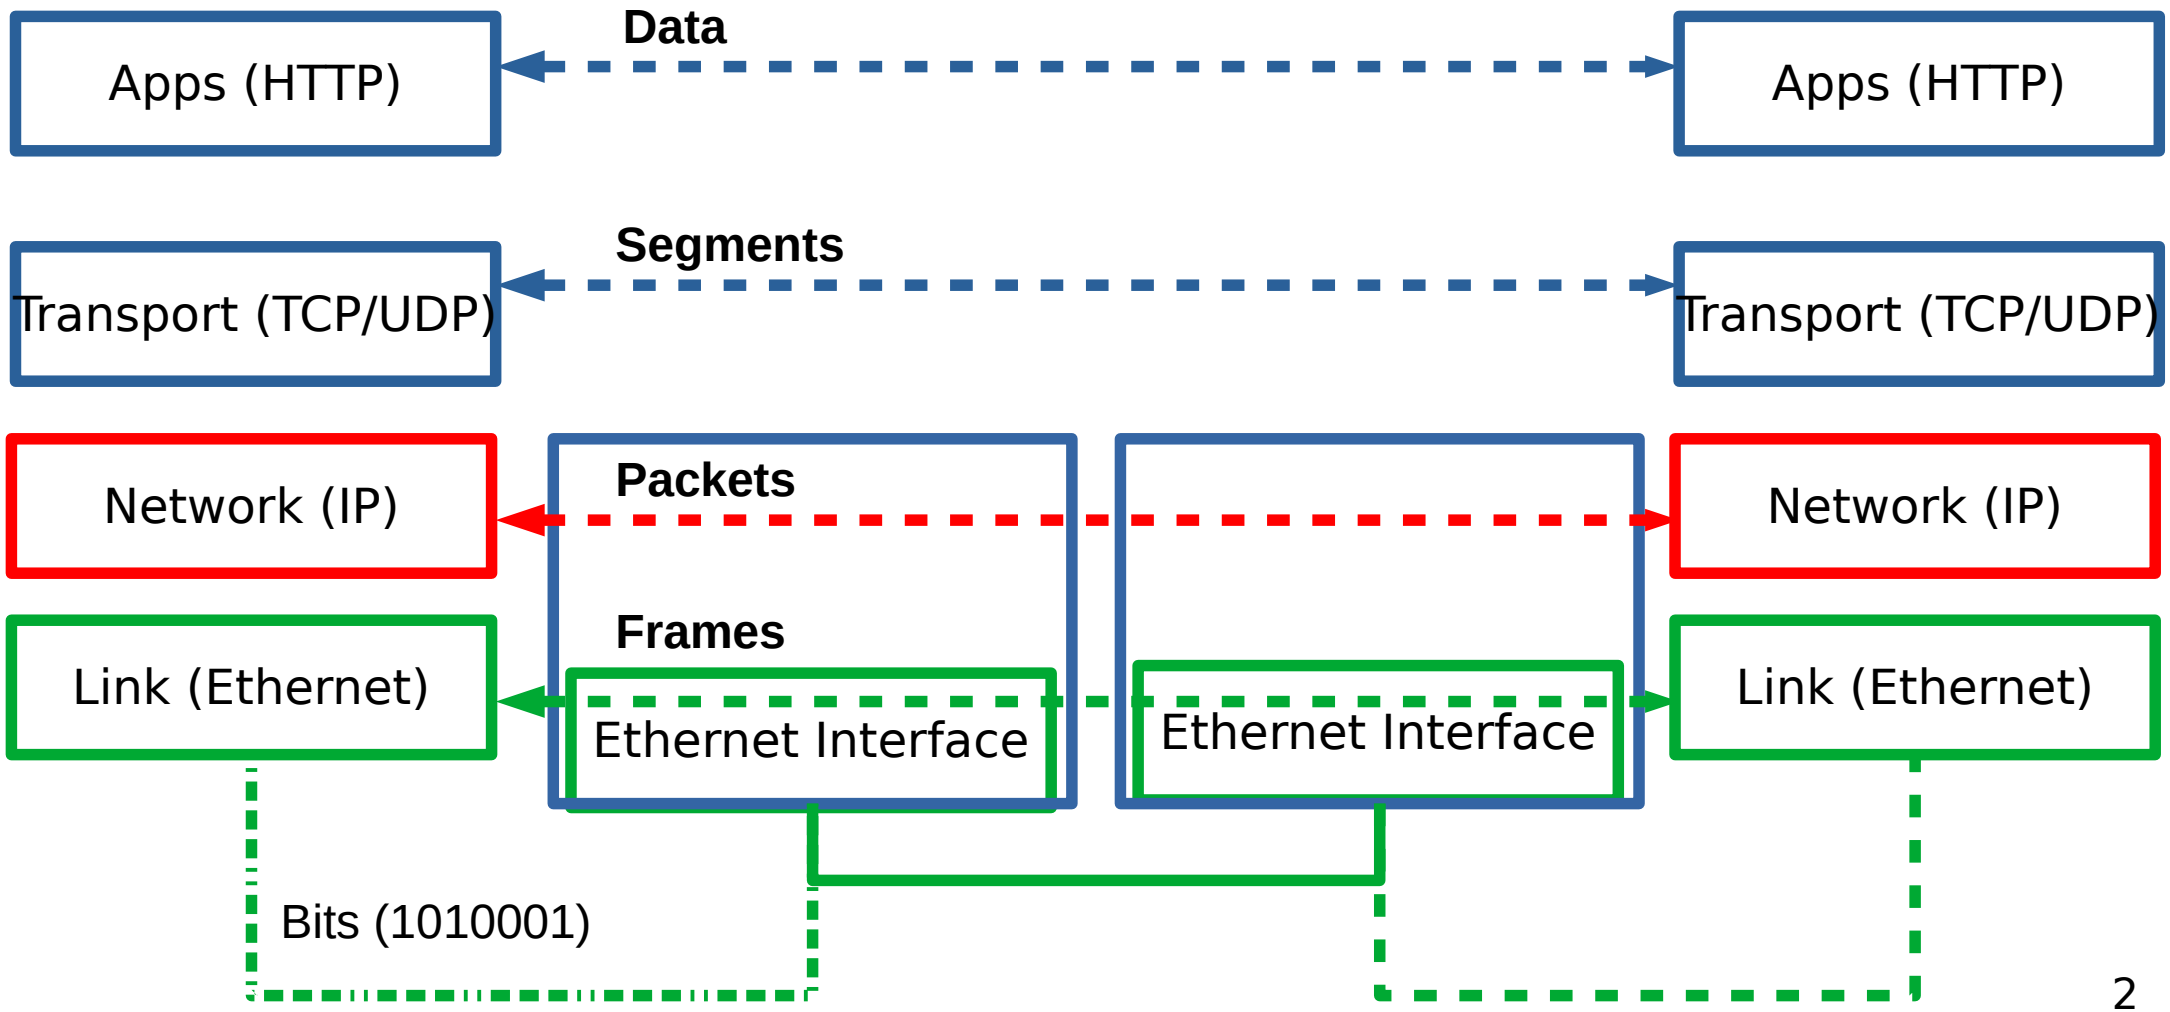
Apps (HTTP)
Data
Apps (HTTP)
Transport (TCP/UDP)
Segments
Transport (TCP/UDP)
Network (IP)
Packets
Network (IP)
Link (Ethernet)
Frames
Ethernet Interface
Ethernet Interface
Link (Ethernet)
Bits (1010001)

Logistics

- PA2 will post tomorrow – due March 31st.
- Those who asked for PA1 extension – email me the code.
- Homework2 will post next week.
- Second exam – end of March.

Back to Addressing

- A 32 bit number in quad-dot notation
- Identifies an ***Interface***
  - ***A host might have several interfaces!!!***
- *129.82.138.254*

  10000001.01010010.10001010.11

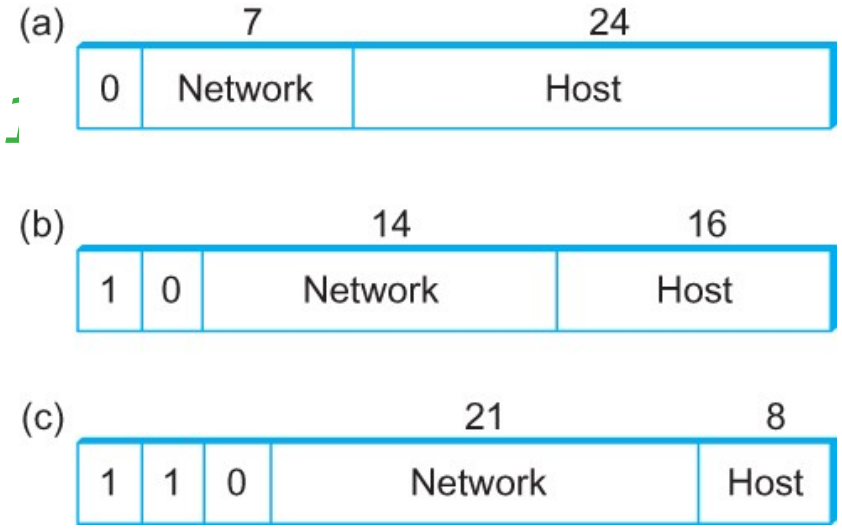

Subnets Revisited

Recipe:

- Create isolated networks – *subnets*

- No longer need to know individual Ips – knowing the subnet is enough
  - *223.1.1.0/14 → Interface 2*

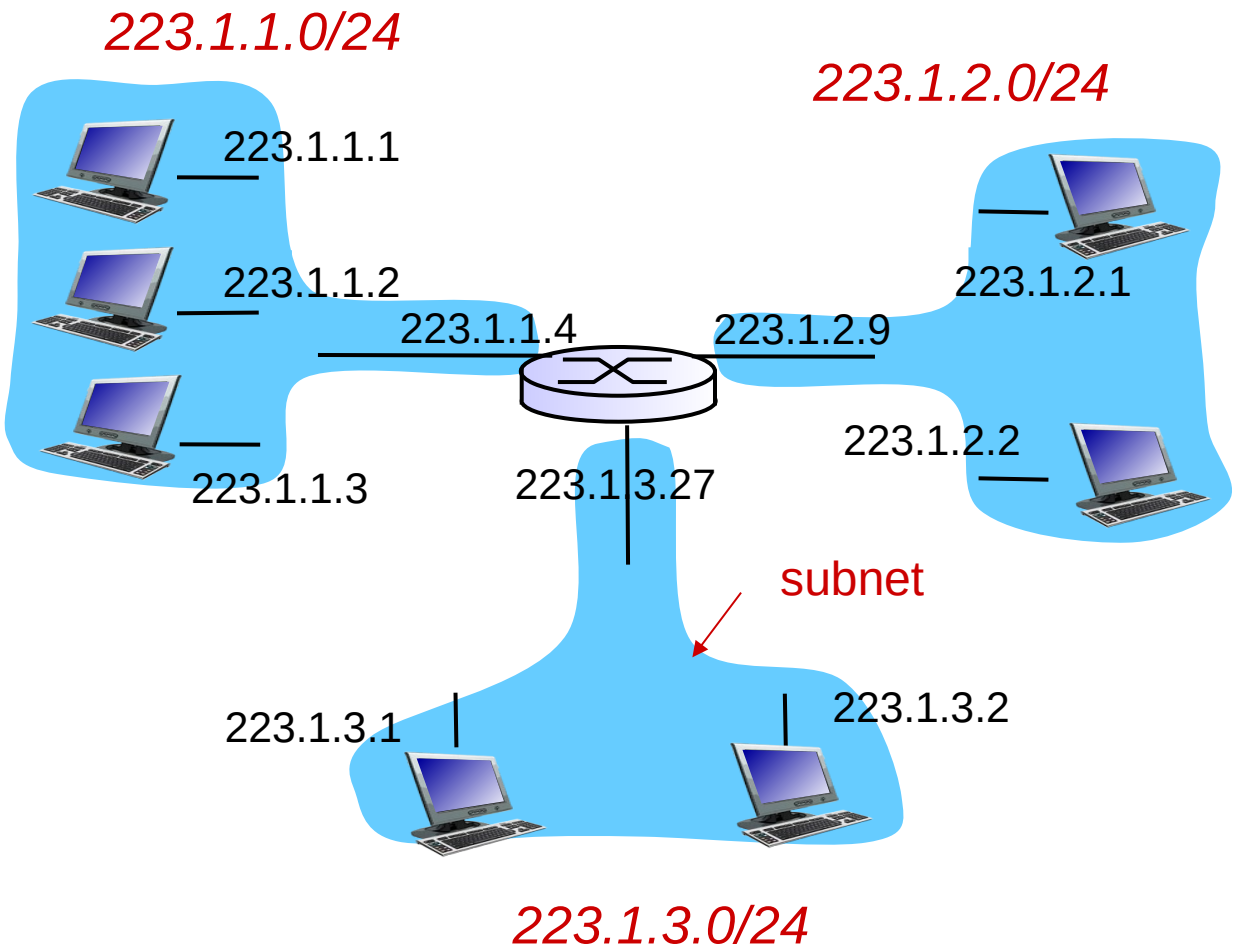

subnet mask: /24

Subnets – Class based

- Originally – only rigid boundaries
  - Class A – 5.0.0.0/8 - 0*
  - Class B – 149.149.0.0/16 - 10*
  - Class C – 129.82.138.0/14 - 110*
  - Class D – 224.0.0.0 – 1110*
  - Class E - reserved

Subnets – Classless CIDR

- No rigid boundaries
  - *129.82.138.0/25*

    10000001.01010010.10001010.10000000

Subnets (Prefixes) scales the Internet

- Addresses are allocated in contiguous prefixes (tntech 149.149.0.0/16)
- Routing protocols operate based on prefixes (how do I reach 149.149.0.0/16)?

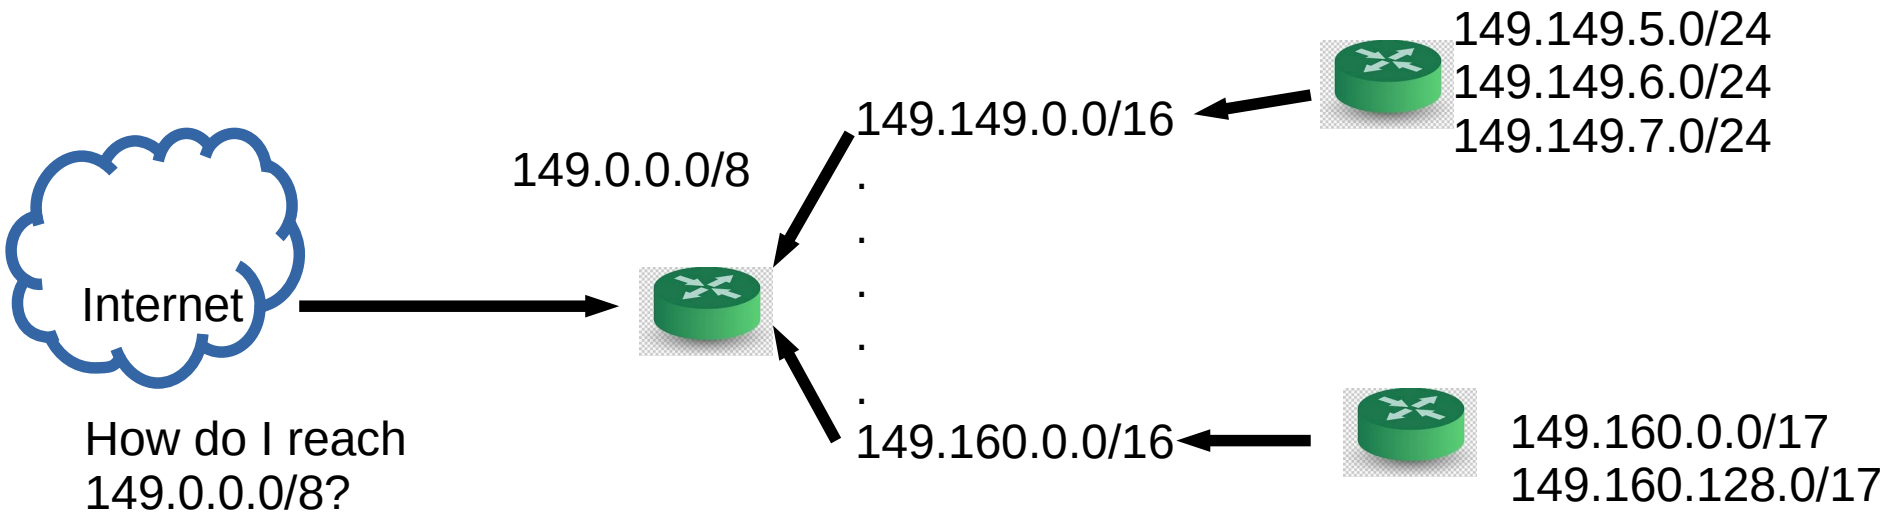

Not

How do I reach 149.149.5.0/24

How do I reach 149.149.6.0/24

Who gets what prefix?

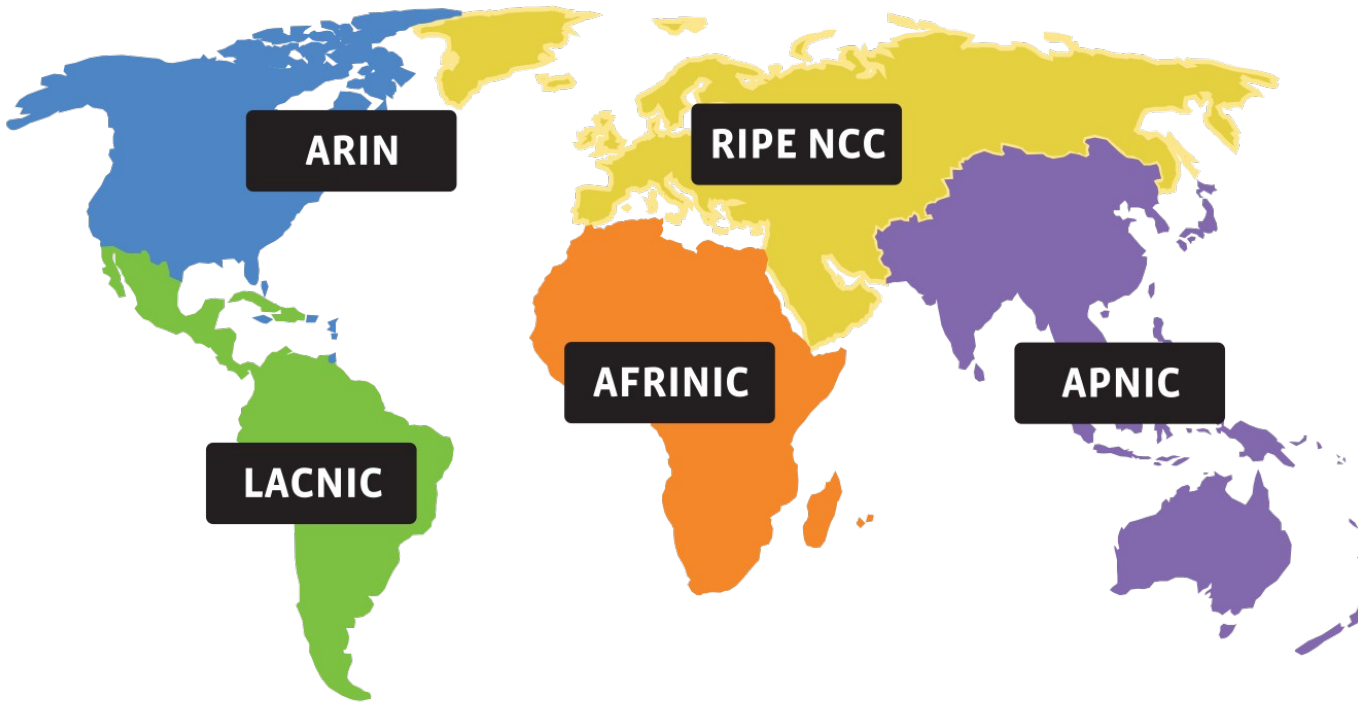

0. Internet Corporation for Assigned Names and Numbers (ICANN) – Decides which RIRs get what address

1. Regional Internet Registries (RIRs) – Which orgs get what address

2. ISPs – Which customers get which address

How do you know who has a prefix? “whois”

$ whois tntech.edu

Domain Name: TNTECH.EDU

Registrant:
Tennessee Technological University
Information Technology Service
1010 N. Peachtree Street
Cookeville, TN 38505
USA

Domain record activated: 09-Sep-1992
Domain record last updated: 26-Sep-2019
Domain expires: 31-Jul-2020

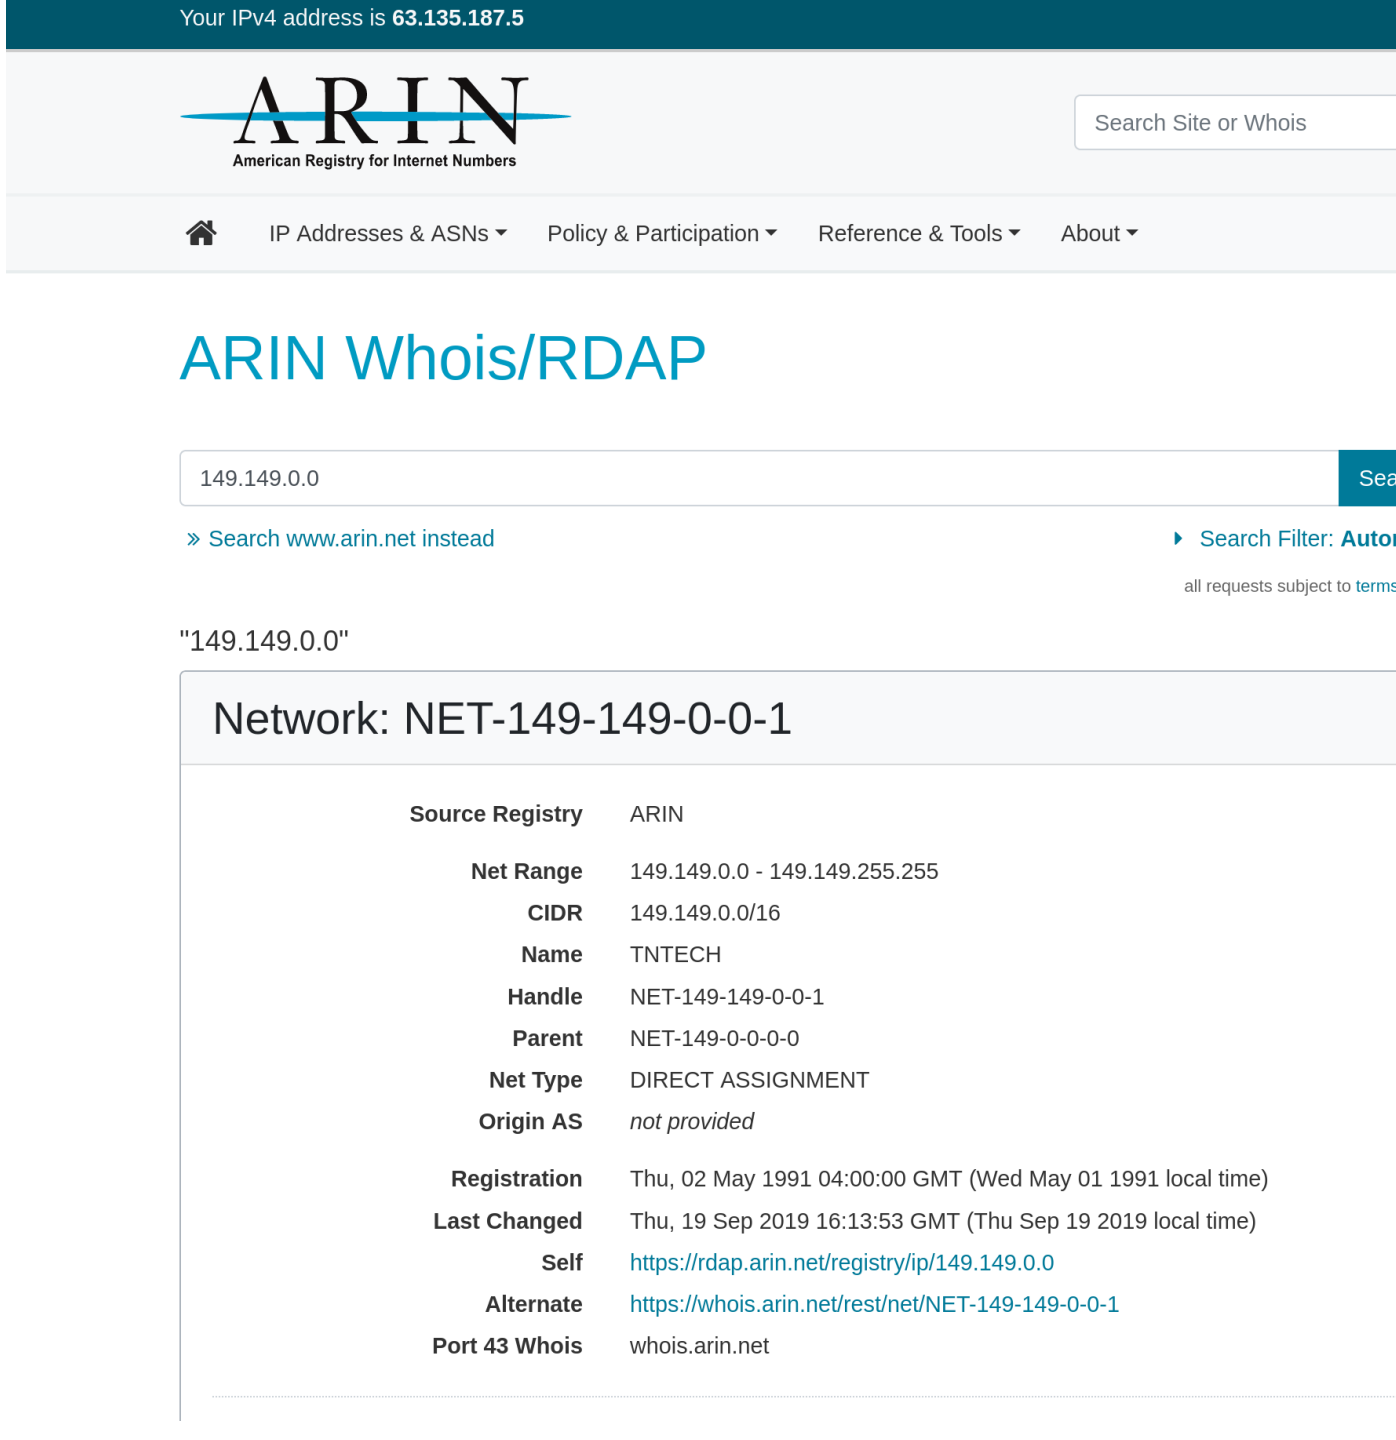

How many prefixes are there?

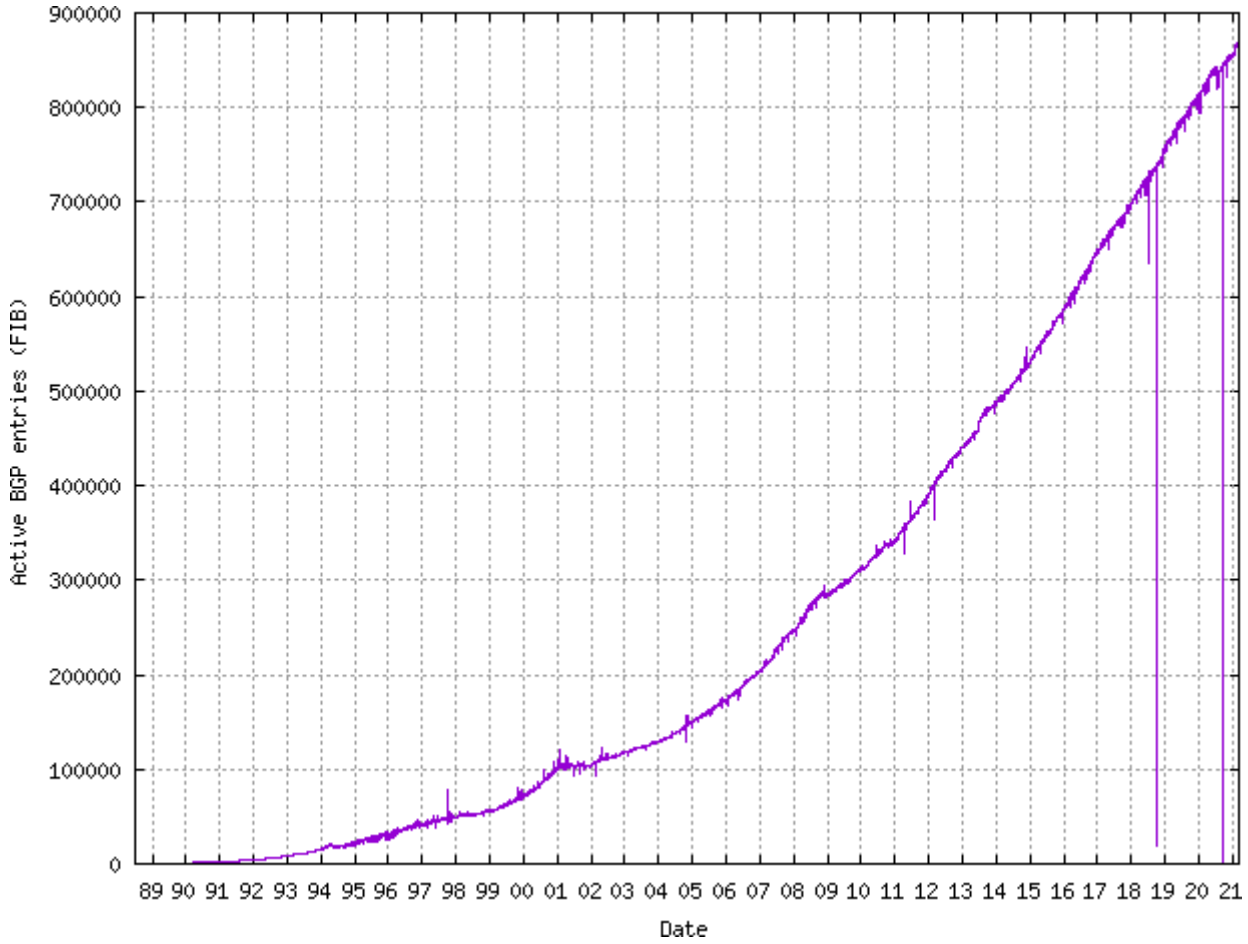

https://www.cidr-report.org/

100K in 2001 → 800K in 2019

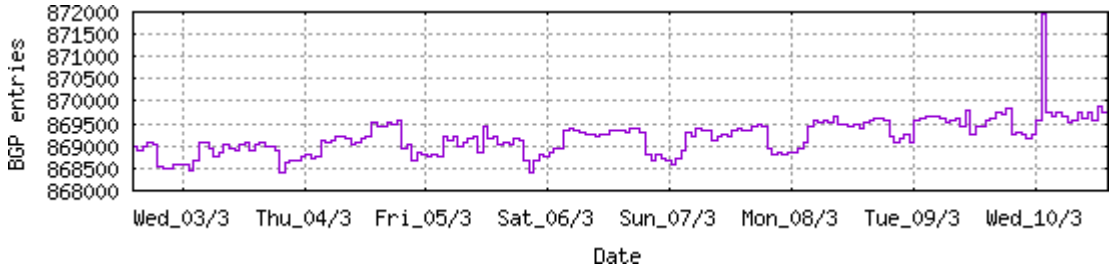

Bit of history – how the Internet evolved

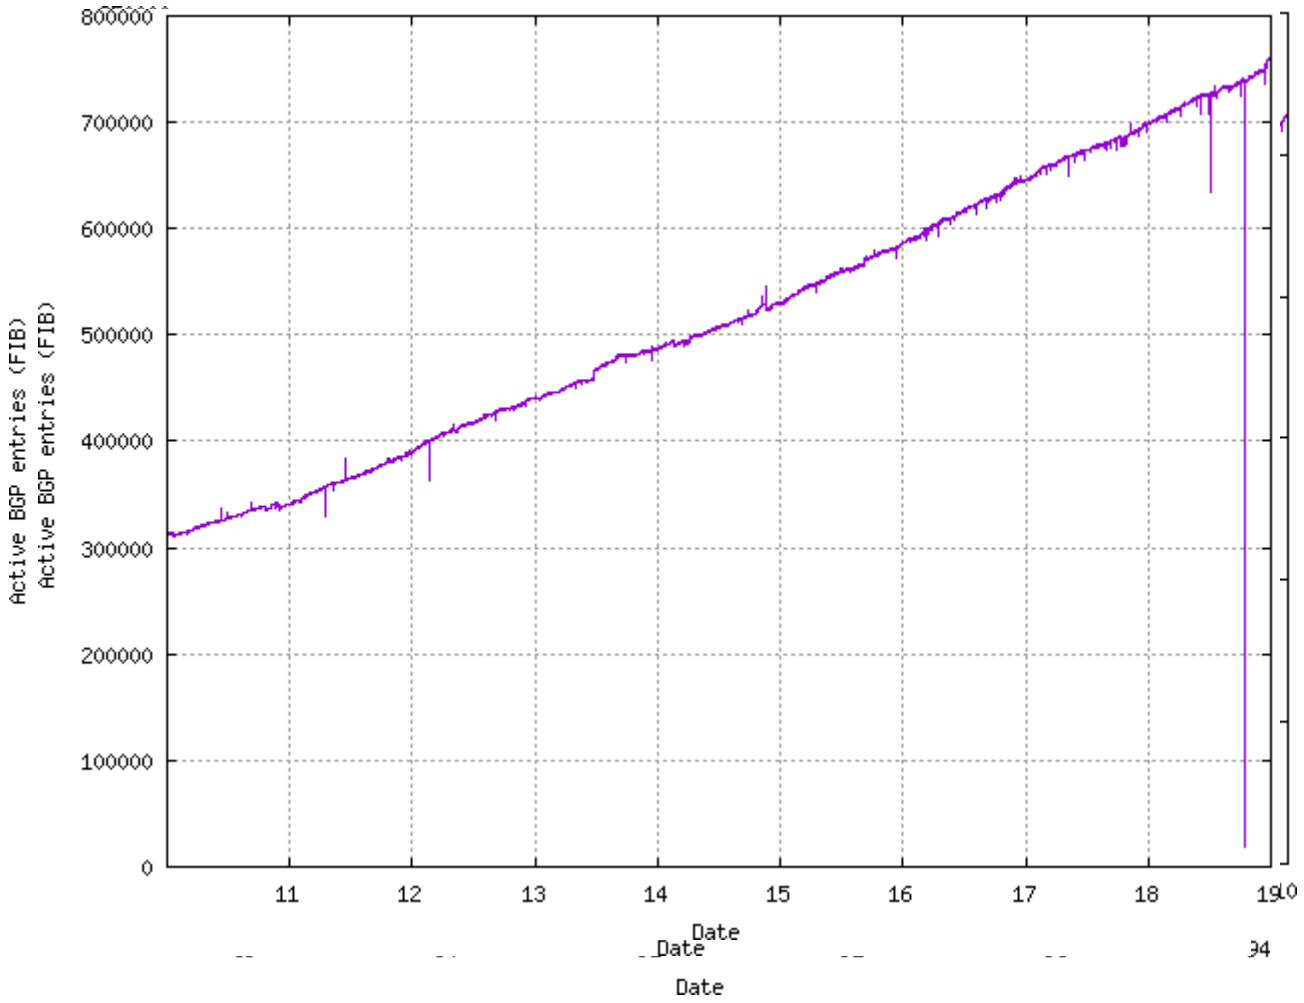

‘88-’94 - 0 → 14000
‘94-’00 – 90000 – Linear growth
’00-’10 – up to 300,000
’10-’19 – up to 800,000

Internet in the 1990s

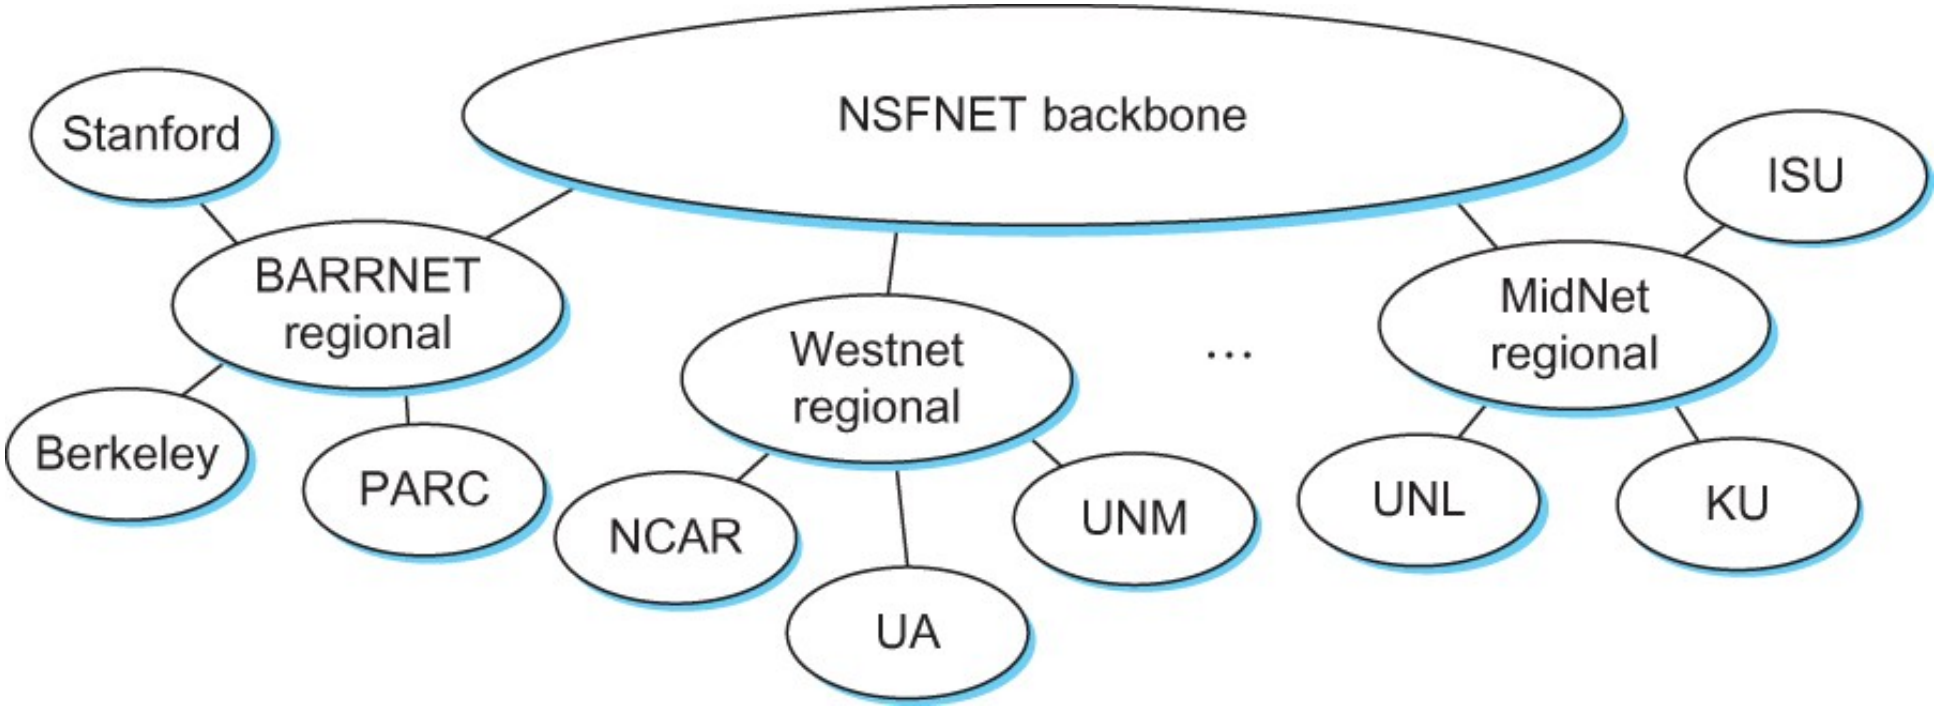

Internet now

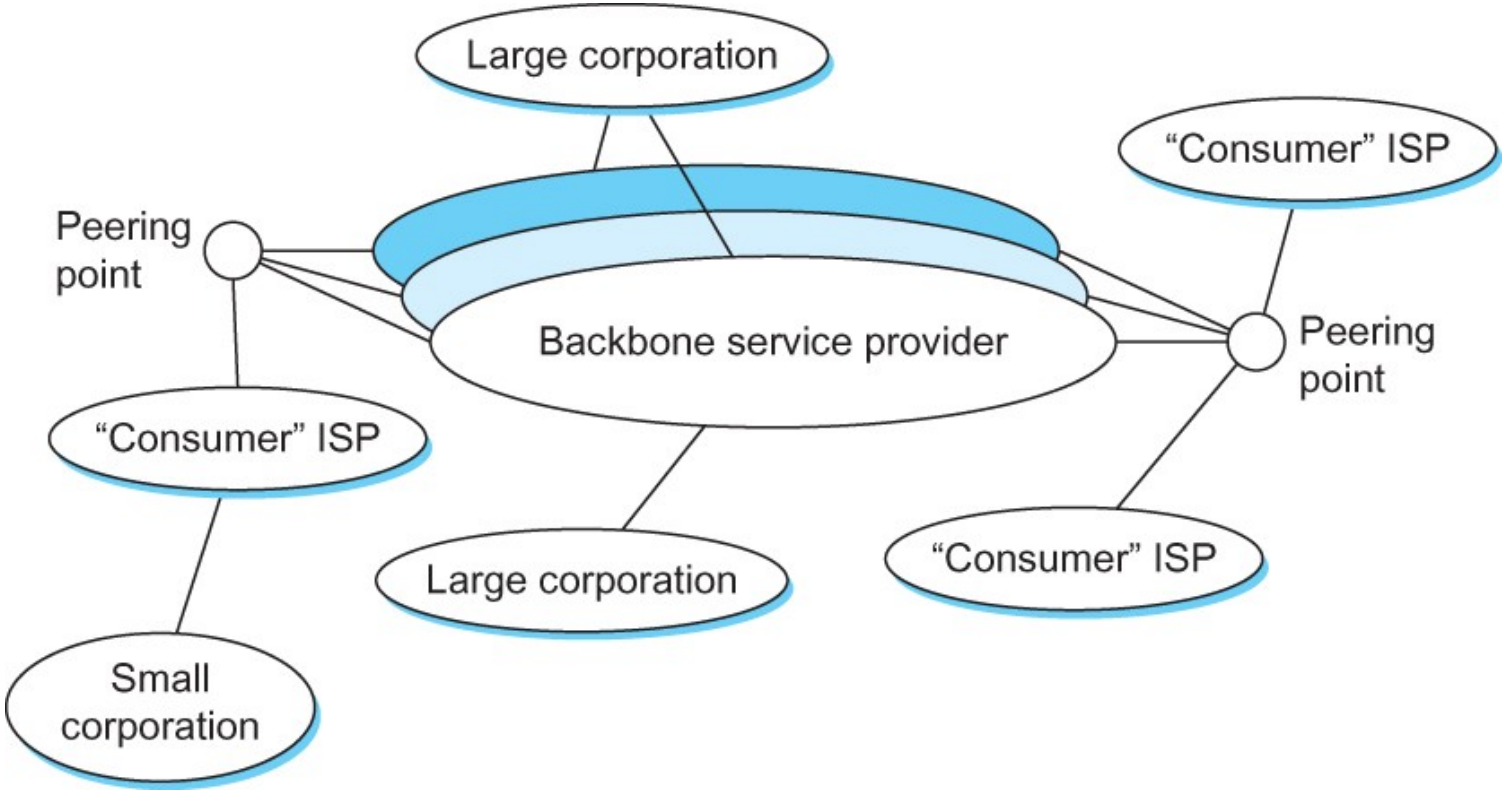

Hierarchical routing - Policy

scale: with 600 million destinations:

- can’t store all dest’s in routing tables!
- routing table exchange would swamp links!

administrative autonomy

- internet = network of networks
- each network admin may want to control routing in its own network

Autonomous systems (ASes)

- AS
  - A set of routers under a single technical administration
  - Uses IGP within the AS to route packets
  - Uses BGP between Ases to route packets
- What happens inside an AS stays within that AS!
  - That is, AS decides routing metrics internally

Interdomain Routing

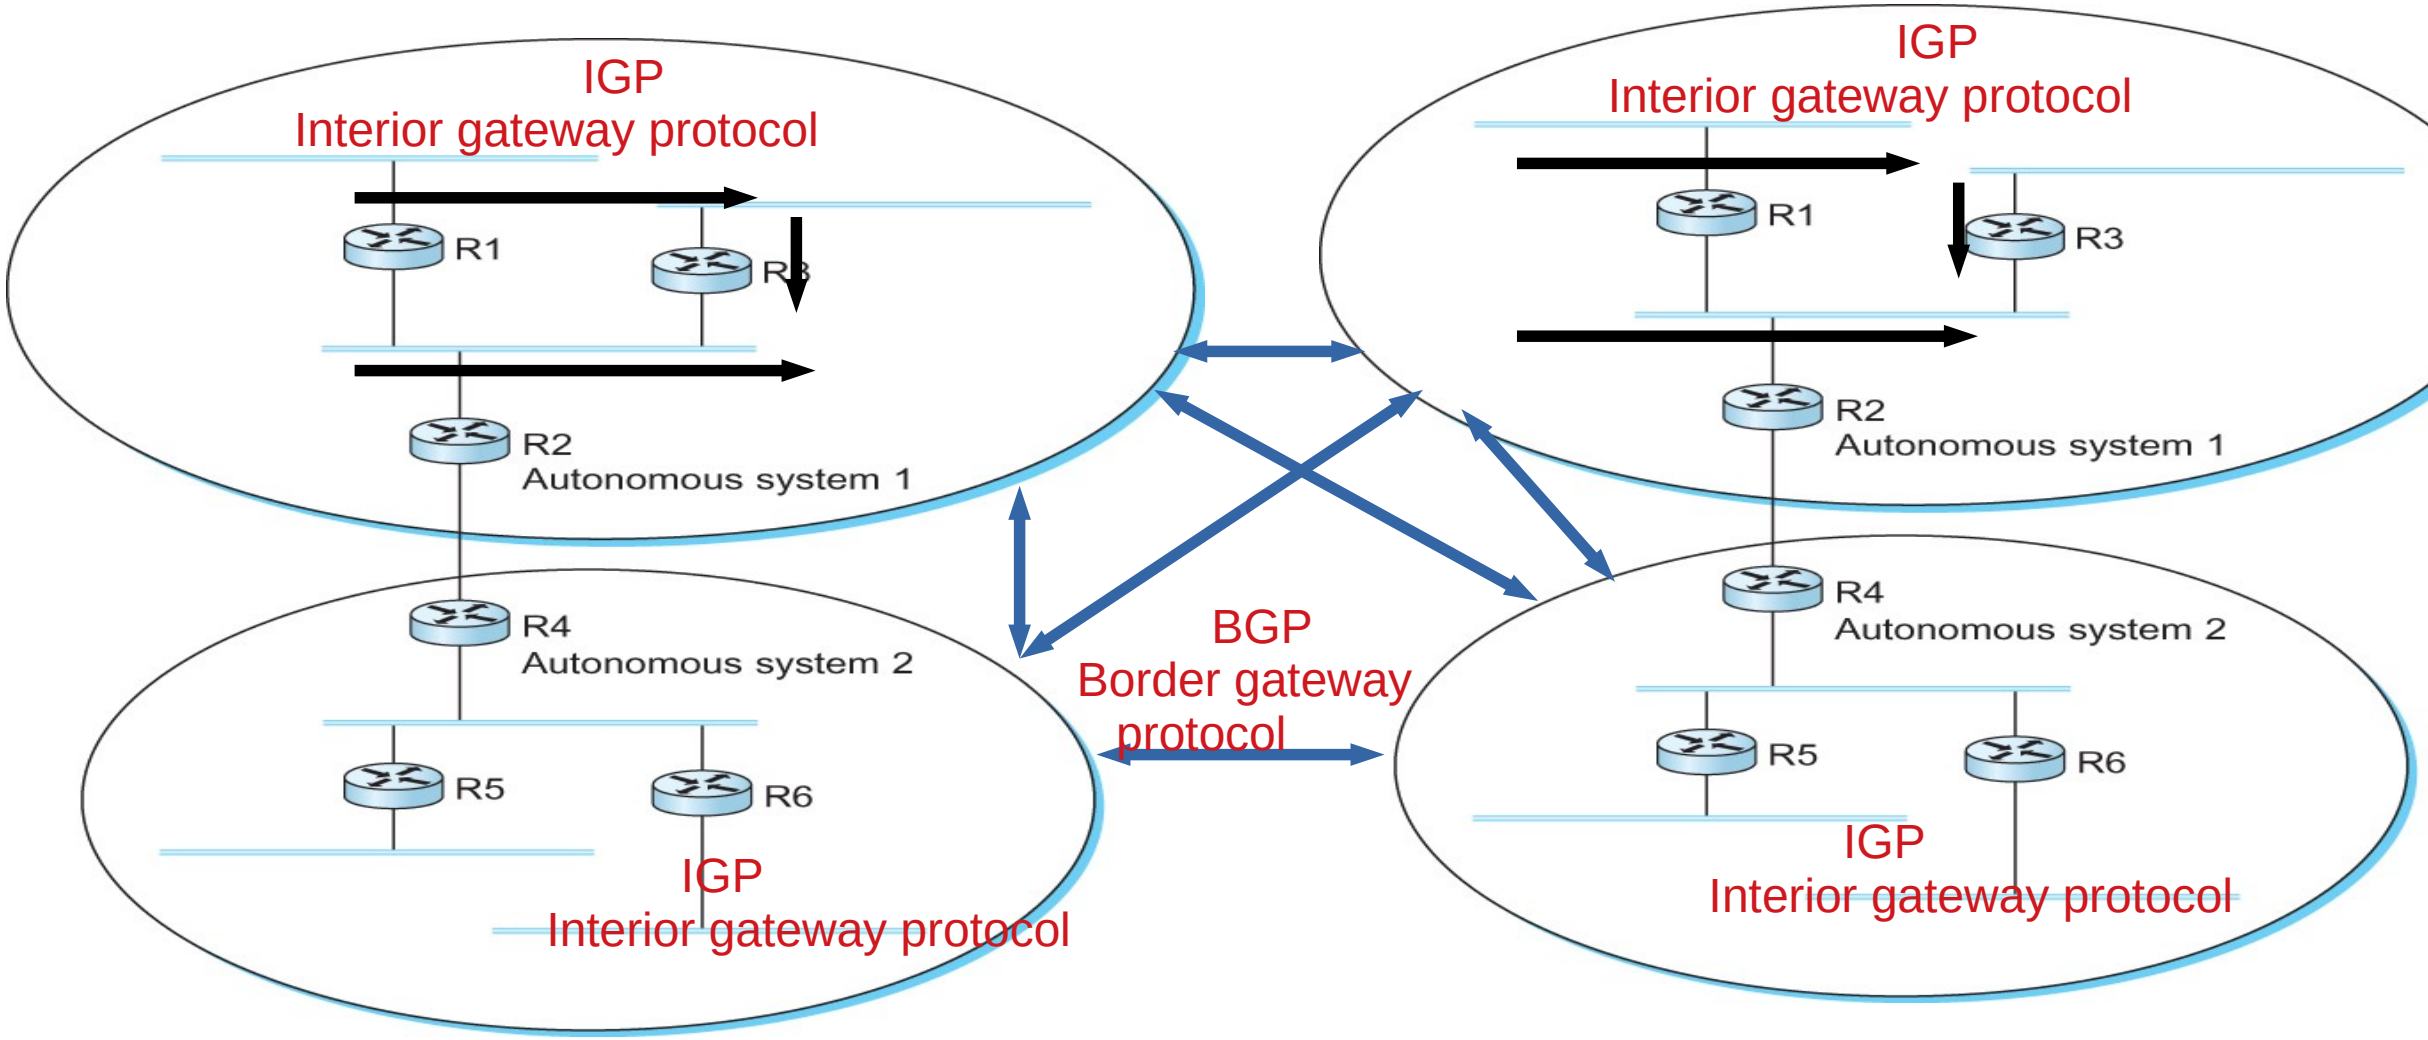

A network with four autonomous systems

BGP-4: Border Gateway Protocol

- Assumes the Internet is an arbitrarily interconnected set of AS's.

- Local traffic – within the AS
- Transit traffic – from AS1 to AS3 via AS2
- Three types of AS’s
  - *Stub AS*
  - *Multihomed AS*
  - *Transit AS*

BGP-4: Border Gateway Protocol

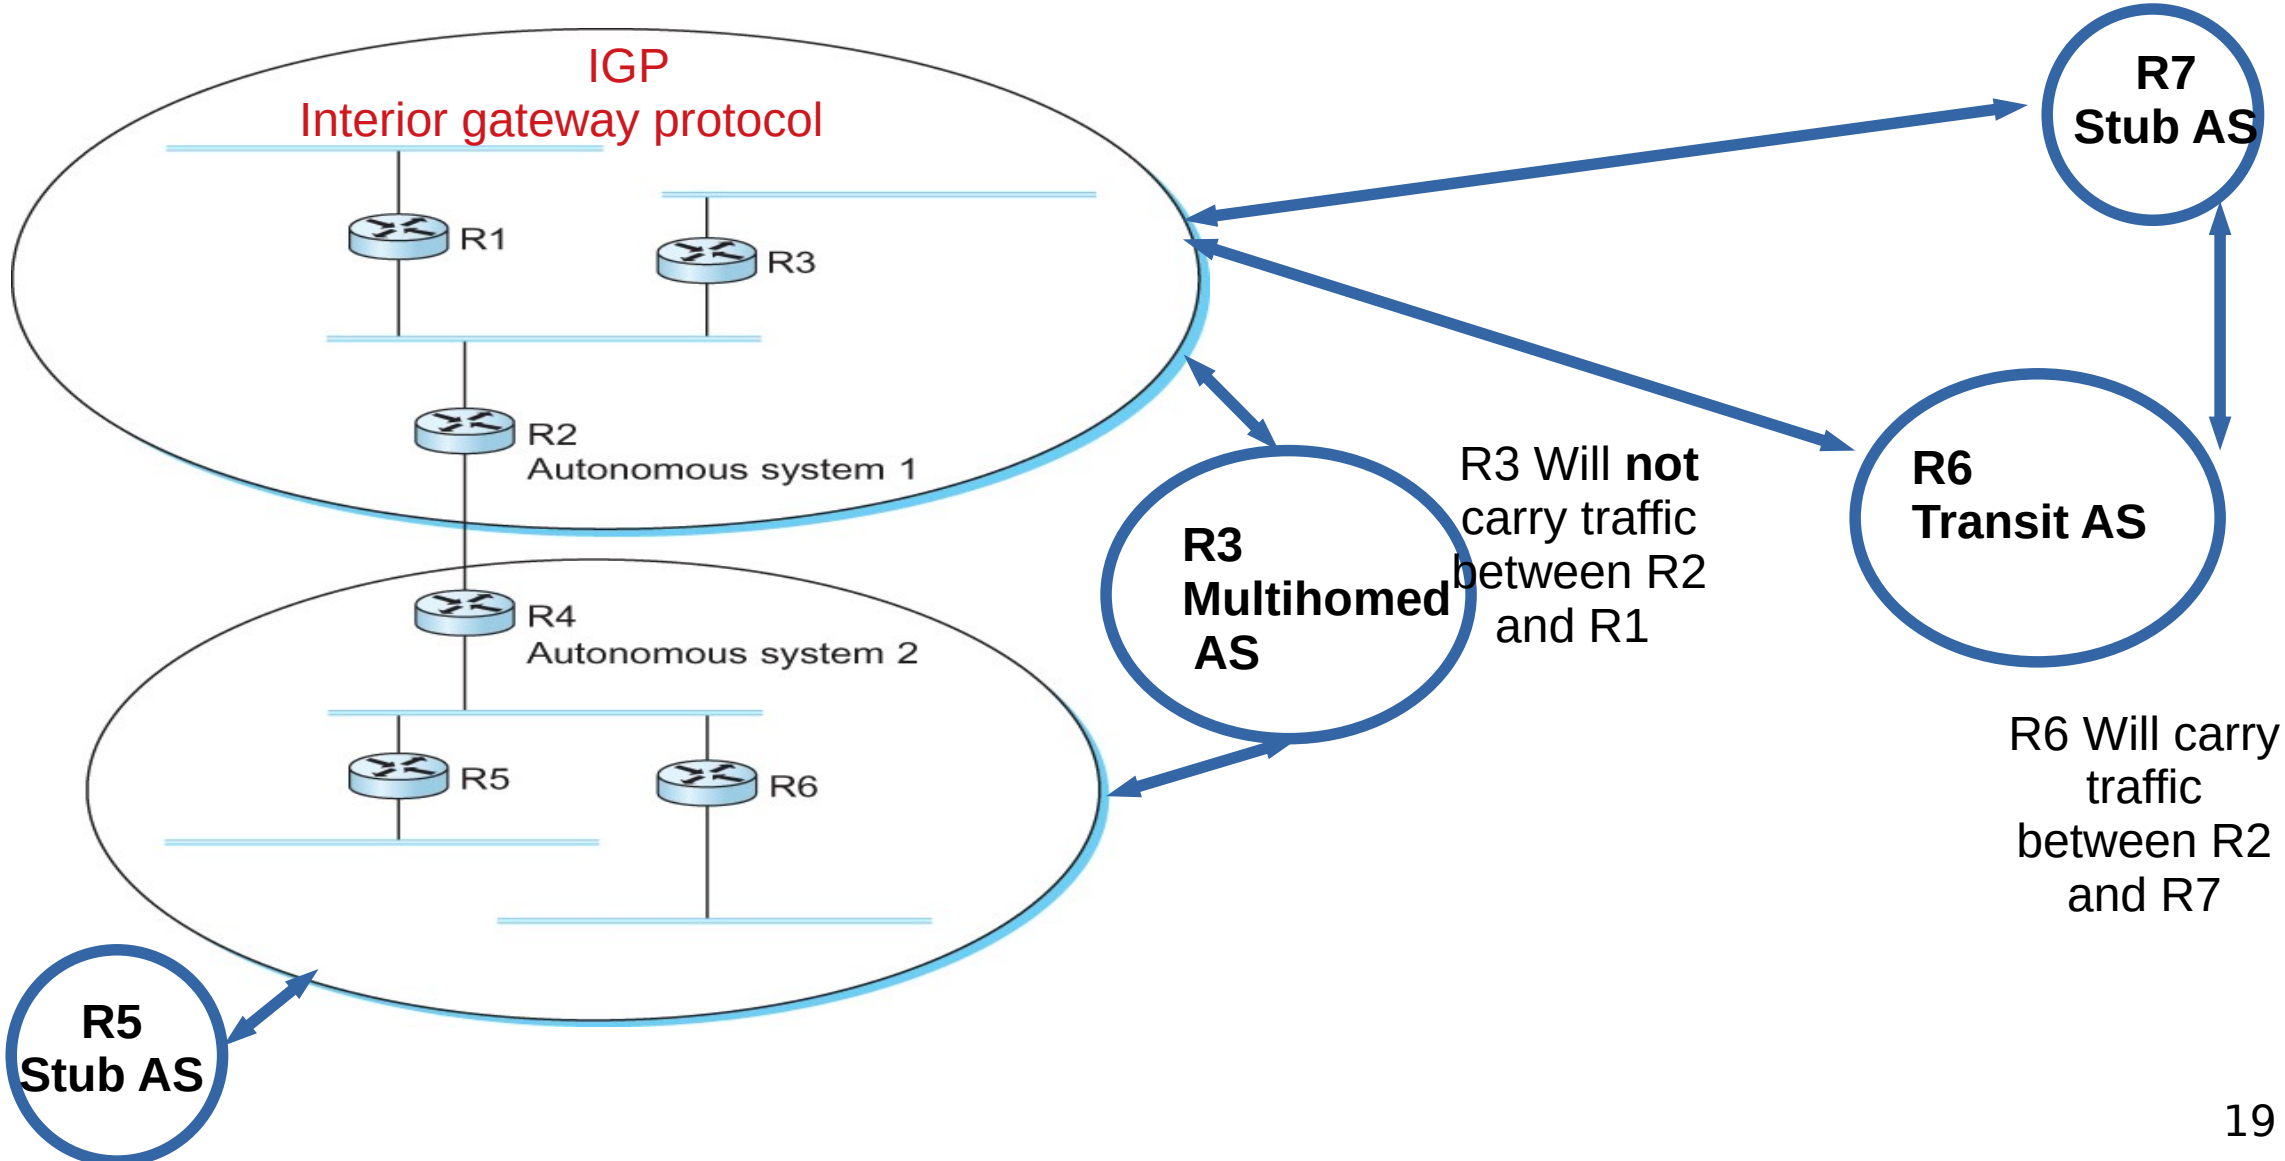

BGP: Which routing protocol?

Link state?

- Does not scale
- you can have loops
- exposes routing costs to others

Distance vector?

- Slow to converge, count-to-infinity
- No universal metrics

BGP - goals

- The goal of Inter-domain routing is to find **any path** to the intended destination that is **loop free**
  - **We are concerned with reachability than optimality**
  - Finding path anywhere close to optimal is considered to be a great achievement
- Why?

BGP - Goals

- Scalability: Forward any packet destined anywhere in the Internet
  - Having a routing table that will provide a match for any valid IP address
- Autonomous nature of the domains
  - impossible to calculate meaningful costs for a path crossing multiple ASs
  - A cost of 1000 is great at provider 1, terrible at provider 2
- Issues of trust
  - Provider A might be unwilling to believe certain advertisements from provider B

BGP: Path vector protocol

- Send the whole path with the routing update
- Loops are detected if an AS finds itself in the path
  - Reject if so
  - Accept otherwise
- Add self to the path and advertise to the neighbors
- Advantage: No loops, Local decision before advertizing

BGP: Path vector protocol

- Send the whole path with the routing update
- Loops are detected if an AS finds itself in the path
  - Reject if so
  - Accept otherwise
- Add self to the path and advertise to the neighbors
- Advantage: No loops, Local decision before advertising

BGP: Path vector protocol

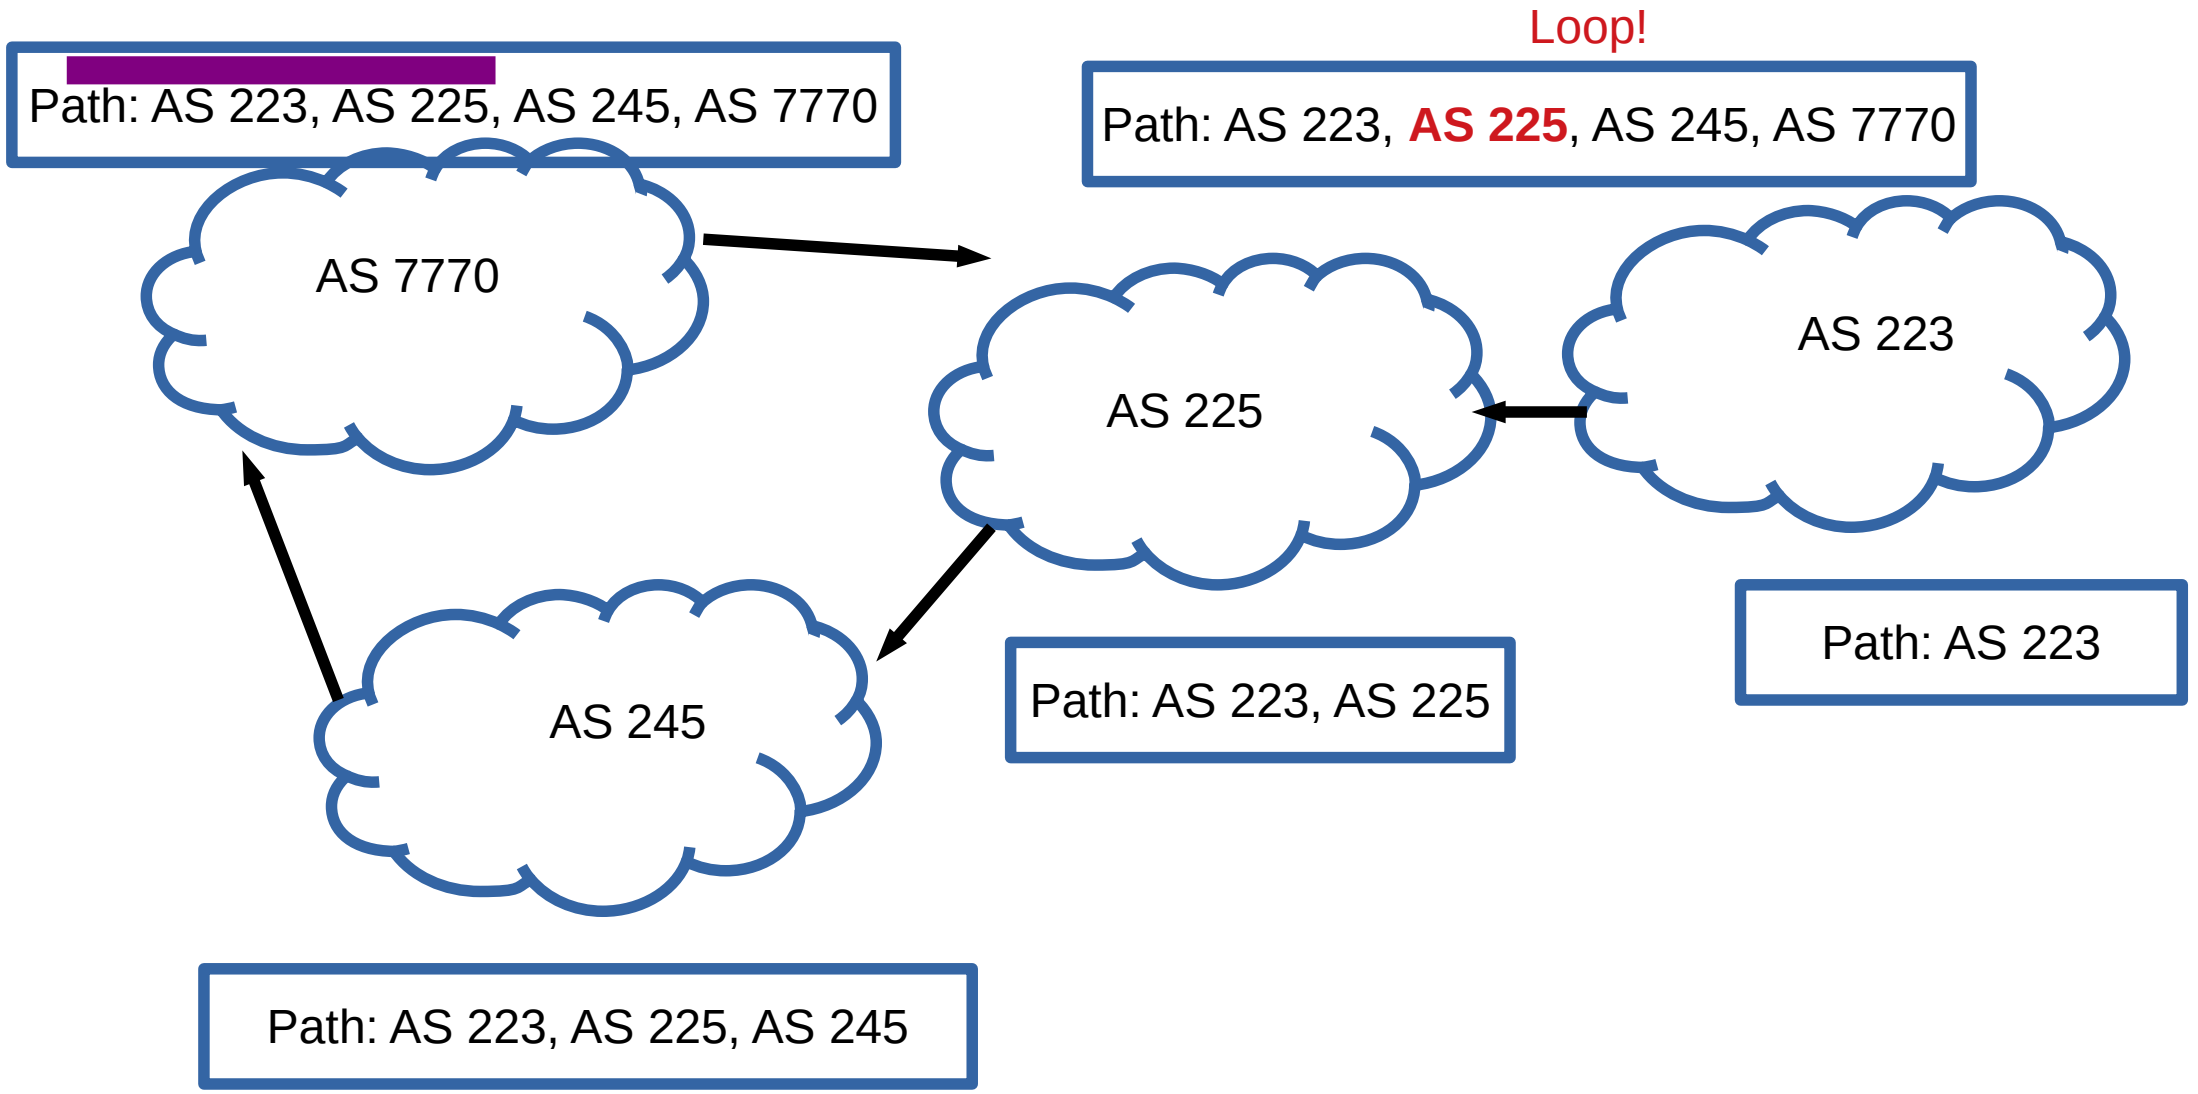

BGP: Interconnections

- Uses TCP port 179 to connect to **peers**
- Arbitrary connections between AS's
- Advantages:
  - Much simpler, no periodic update
  - Valid as long as TCP connection is valid (or withdrawn)
  - Incremental update (only a portion of the routing table)
- Disadvantages:
  - No security
  - Congestion control on routing messages

BGP: Security problems

Anyone can advertise anything!!!

c|net BEST PRODUCTS REVIEWS NEWS VIDEO HOW TO SMART HOME CARS DEALS DOWNLOAD JOIN / SIGN IN

CULTURE

How Pakistan knocked YouTube offline (and how to make sure it never happens again)

YouTube becoming unreachable isn't the first time that Internet addresses were hijacked. But if it spurs interest in better security, it may be the last.

BY DECLAN MCCULLAGH | FEBRUARY 25, 2008 4:28 PM PST

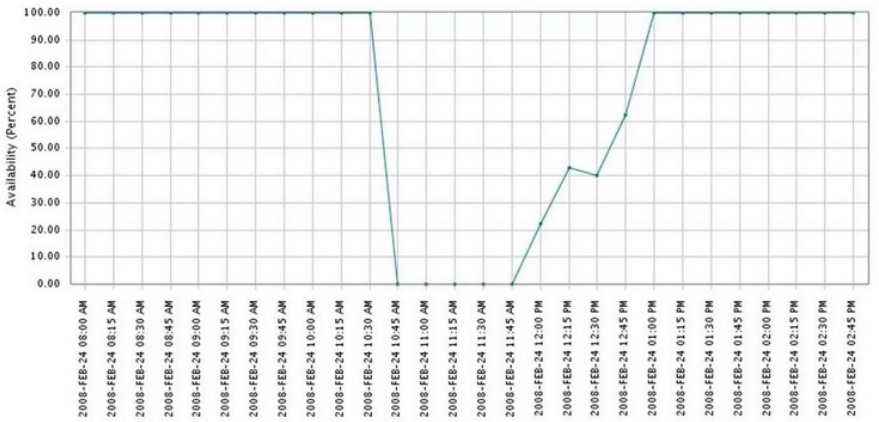

This graph that network-monitoring firm Keynote Systems provided to us shows the worldwide availability of YouTube.com dropping dramatically from 100 percent to 0 percent for over an hour. It

BGP: Hop by Hop model

- You can only tell others what you are using
  - But you control what you say

- BGP advertises only to peers
  - Tell them what you are using
  - Hop-by-hop model

BGP: Allows for policy

- Capable of enforcing various policies
  - AS2 → Don’t use AS1 to get to AS3
- Not part of BGP – configuration information that controls propagation of paths

Reading Assignments

- Scaling to billions:
  - https://book.systemsapproach.org/scaling/problem.html#problem-scaling-to-billions
    - ~2 minutes read

- Global Internet
  - https://book.systemsapproach.org/scaling/global.html#global-internet
  - Skip the routing areas section
  - Read until “Common AS Relationships and Policies”
  - ~40 minutes

Next steps