Today's Agenda

- BGP Overview
  - Note: it takes years to really master BGP
  - Many slides stolen from Prof. Zhi-Li Zhang at Minnesota and from Avi Freedman's slides
- AS Relationship Inference
  - There'll be some openresearch topics here

Inter-Domain Routing: BGP

- Internet connectivity and BGP
  - connectivity services, AS relationships
- BGP Basics
  - BGP sessions, BGP messages, BGP attributes
- BGP Policy Control: Examples
  - Cisco filtering mechanisms
- Others: scalability and stability

Dynamic Routing: Intra- vs. Inter-AS

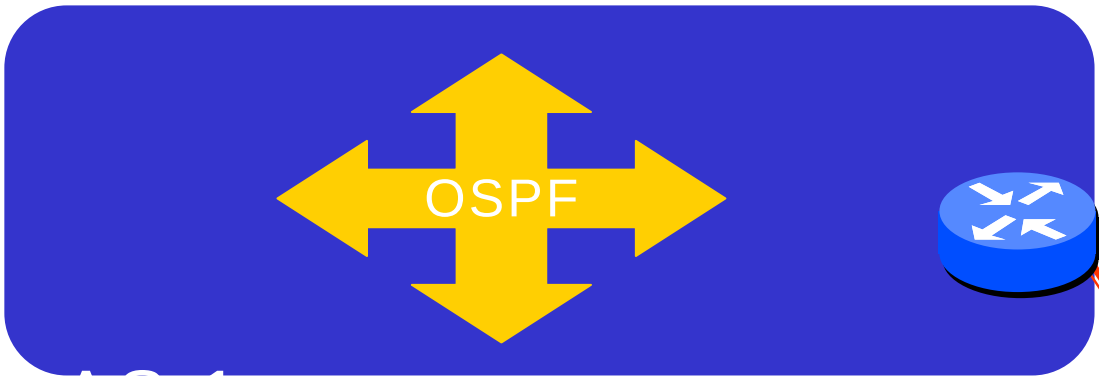

BGP

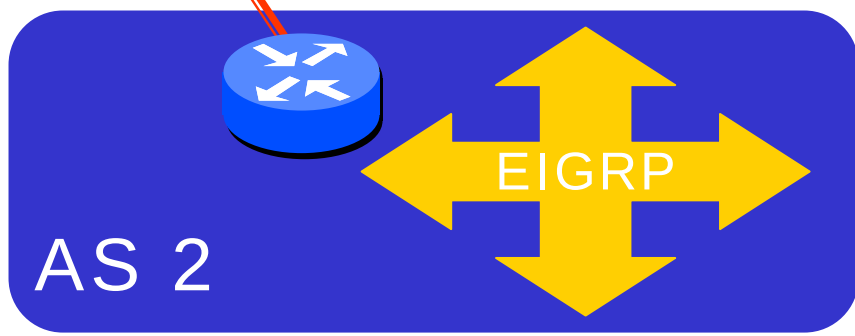

IGP = Interior Gateway Protocol

Metric based: OSPF, IS-IS, RIP, EIGRP (cisco)

EGP = Exterior Gateway Protocol

Policy based: BGP

The Routing Domain of BGP is the entire Internet

Where Does Forwarding Table Come From?

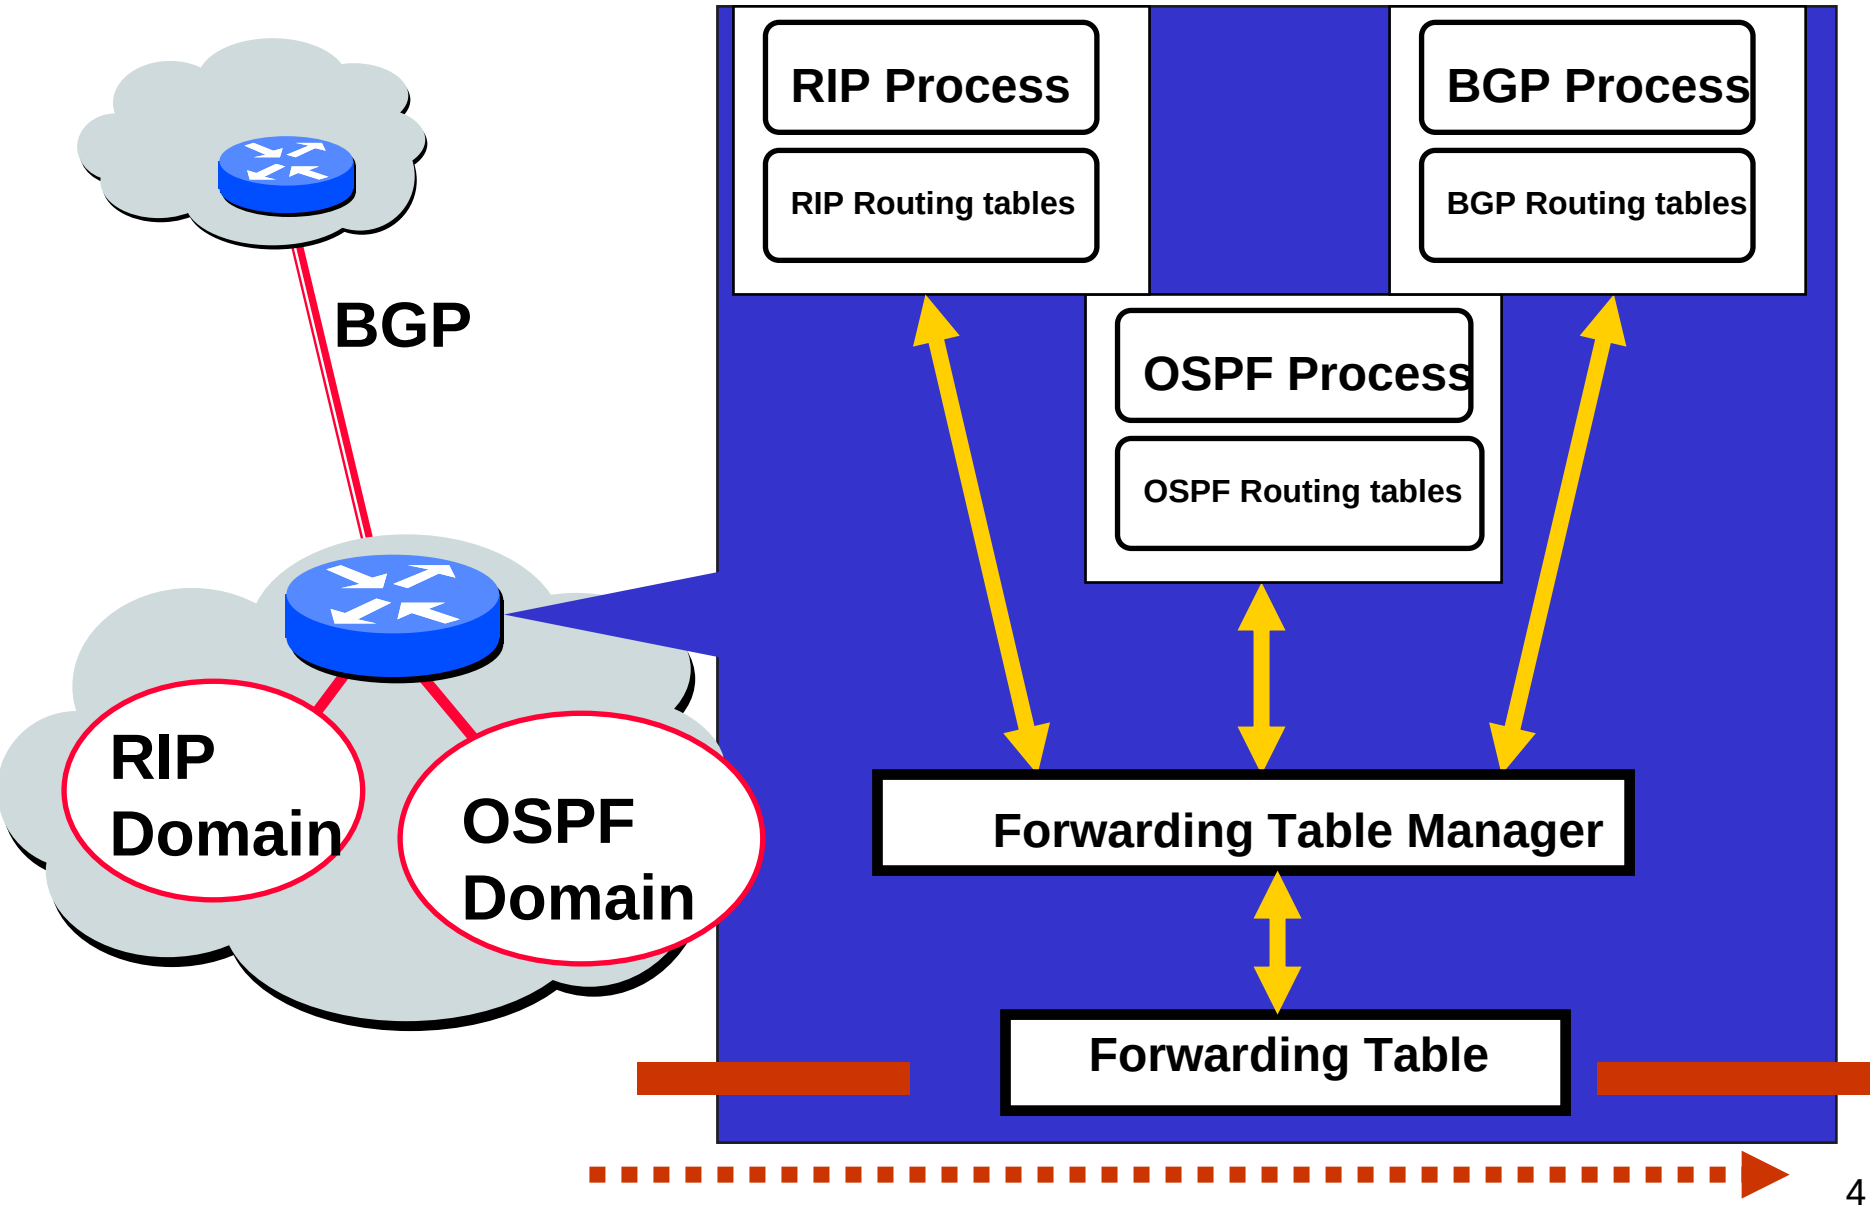

Internet Architecture

Internet: “network of networks”

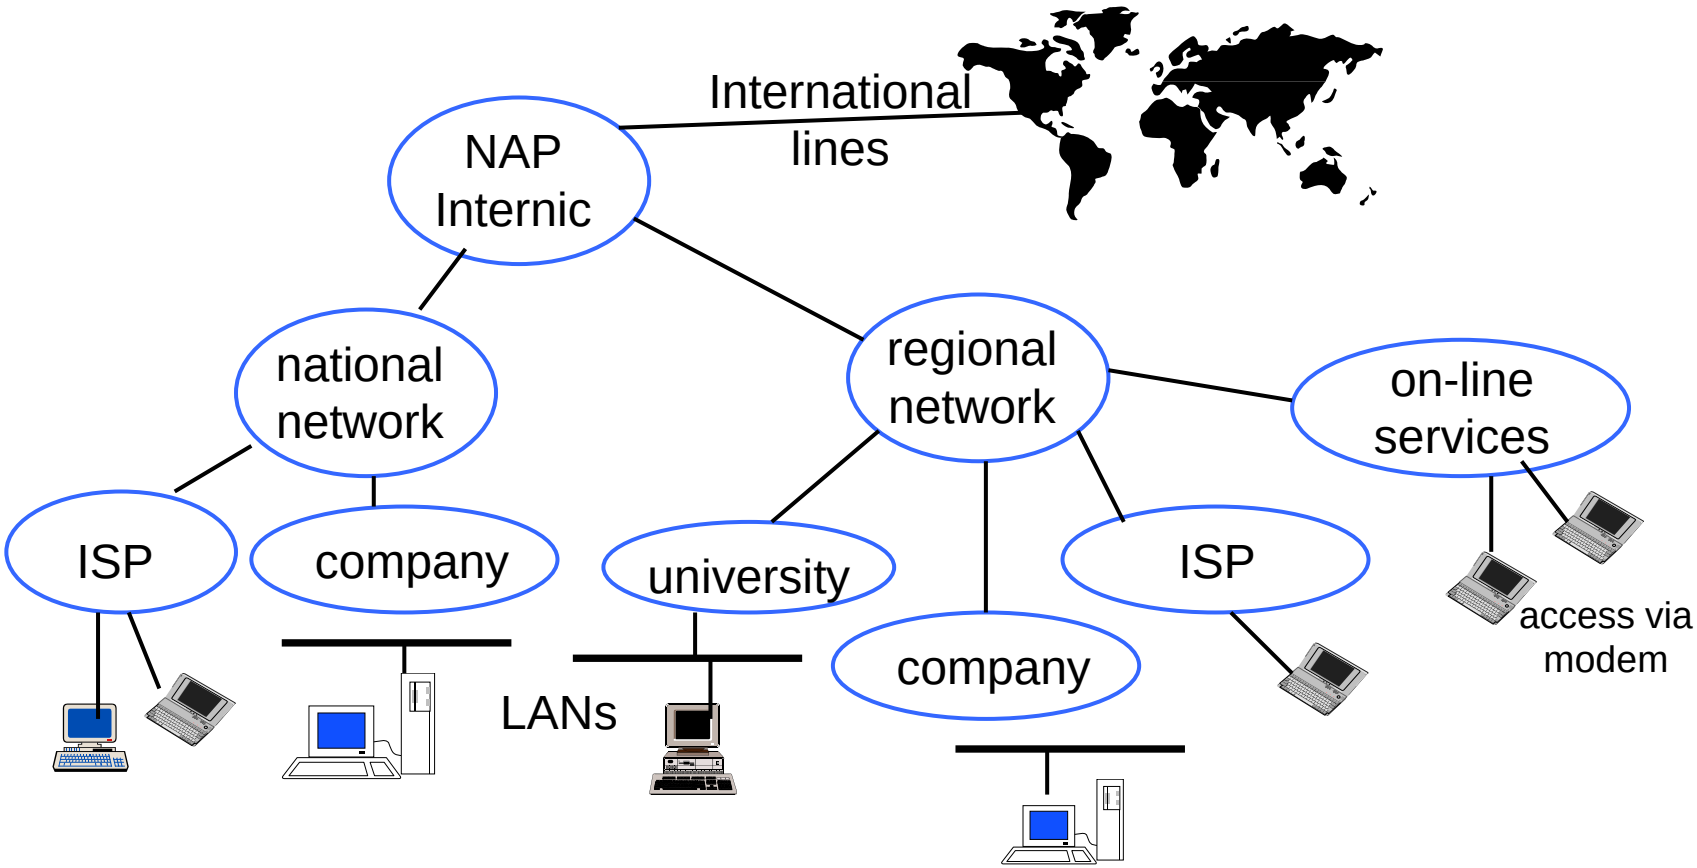

AS and AS Numbers (1)

- Autonomous Systems (each with an ASN)
  - Multihomed AS: connections to > 1 ISP (no transit traffic)
  - Stub AS: connection to 1 ISP (waste of AS number)
  - Transit AS

- ASNs assigned by IANA (Internet Assigned Number Authority)

- ASNs are 16-bit integers (running out!)
  - Public ASNs: 1 – 64511
  - Private ASNs: 64512 – 65536 (used internally in an AS)

AS and AS Numbers (2)

Currently over 20,000 in use.

- Genuity: 1
- MIT: 3
- Harvard: 11
- AT&T: 7018, 5075, ..., 6341, ...
- UUNET: 701, 702, 284, 12199, ...
- Sprint: 1239, 1240, 6211, 6242, ...
- University at Buffalo: **3685** (since 1994)
- ...

ASNs represent units of routing policy

Number of Used ASNs

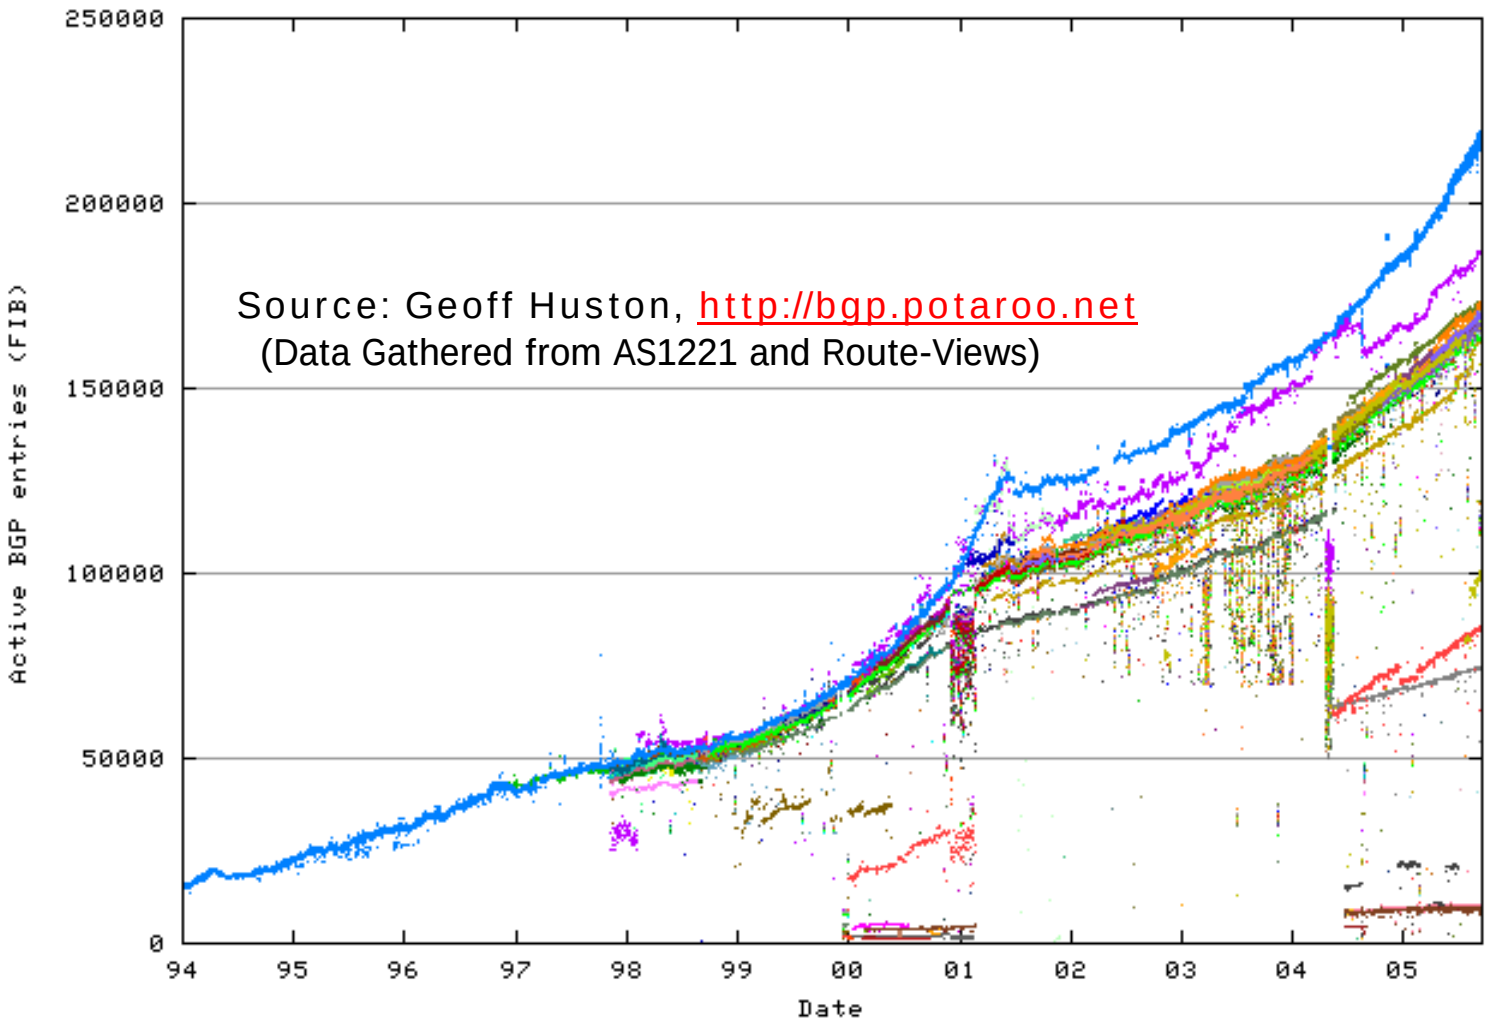

University at Buffalo Neighborhood (Data from 2002)

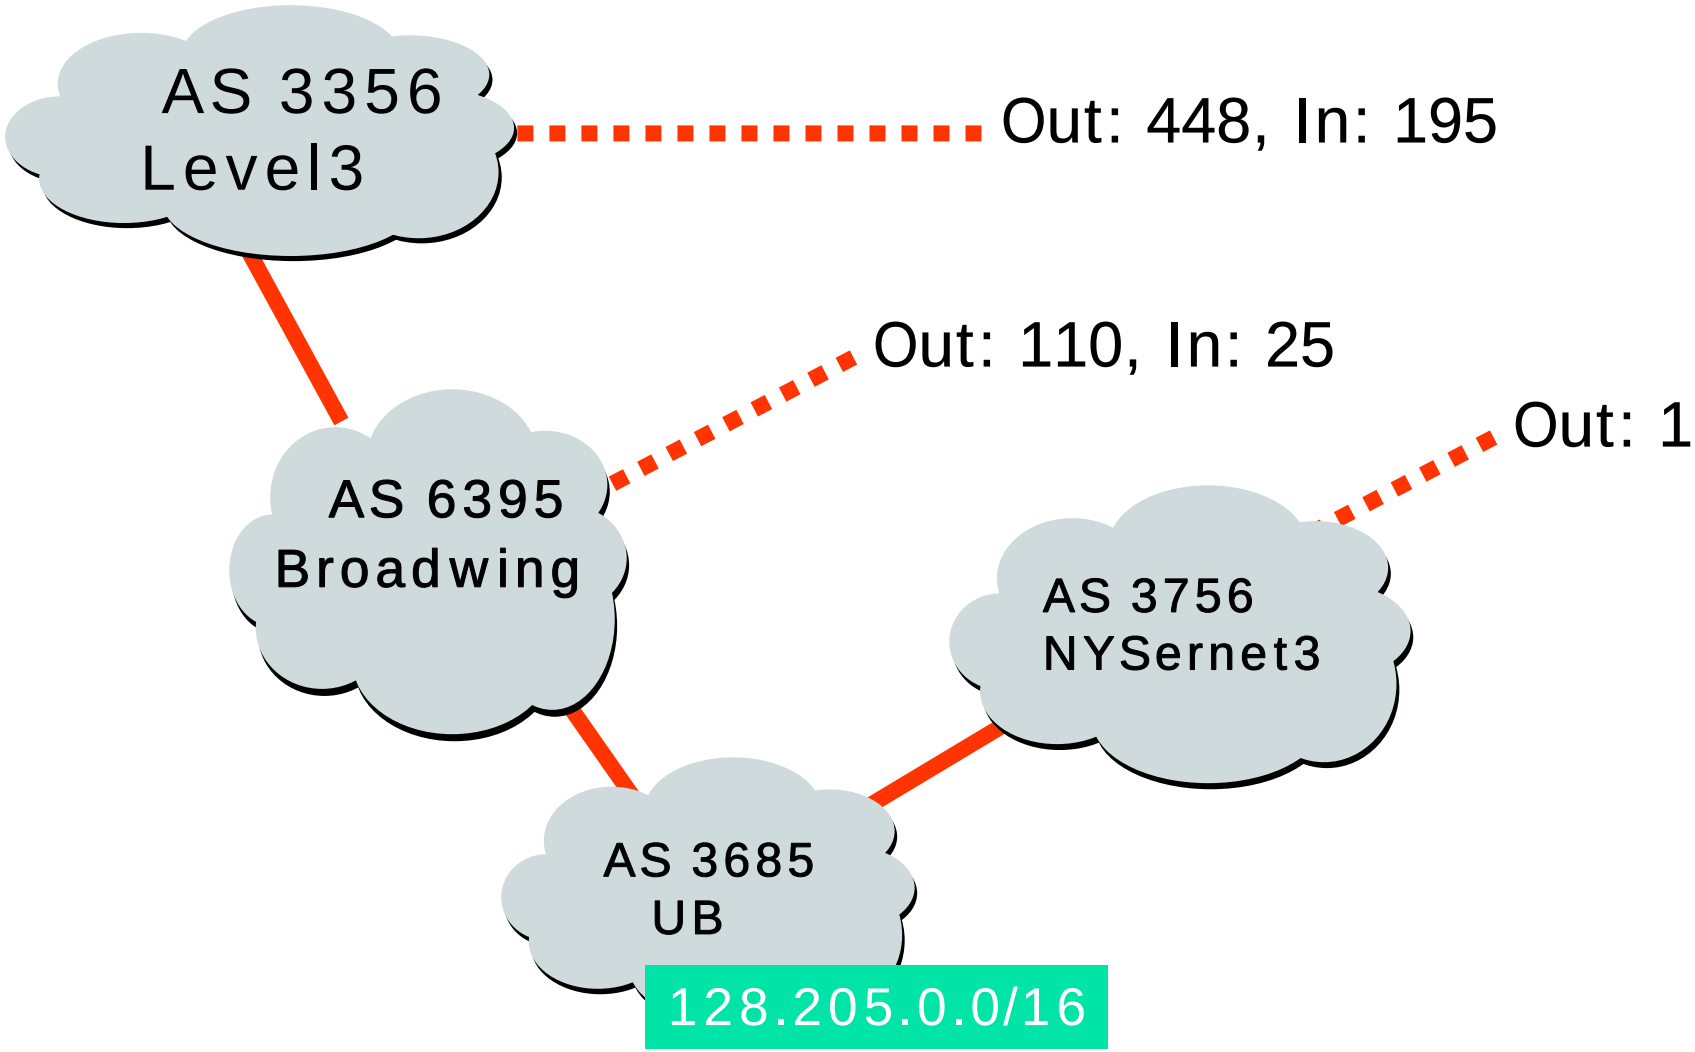

Inter-domain Links

- Exchange Points
  - Layer 2 or Layer 3
  - Usually Gigabit (or more) Ethernet switch

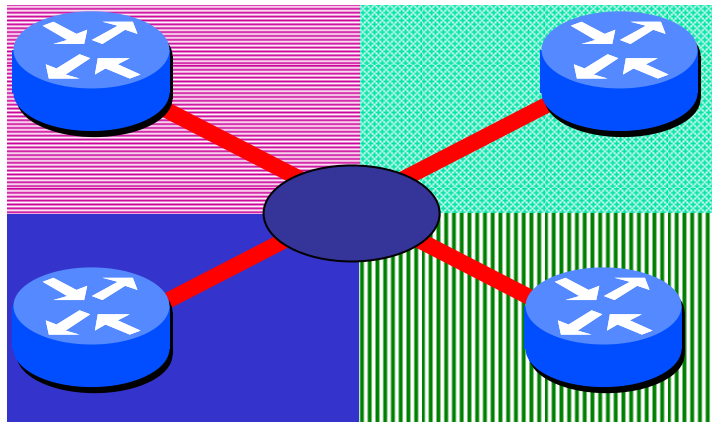

- Private Circuit
  - May be provided by a third party

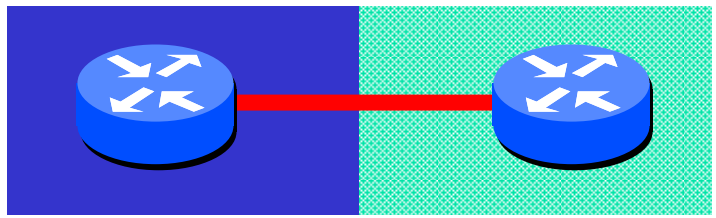

BGP: The Glue of Internet

- Allow ASs to tell other ASs about “routes” that they are “responsible” for and how to reach them
  - Using “route advertisements” - also called **NLRI** (network-layer reachability information)
  - *Path-vector* routing protocol
- Allow setting routing policies
  - Selecting outbound paths
  - Announcing internal routes
- Relatively “simple” protocol
  - Configuration is complex
  - Entire world can see, and be impacted by, your mistakes

Growth of BGP Routes

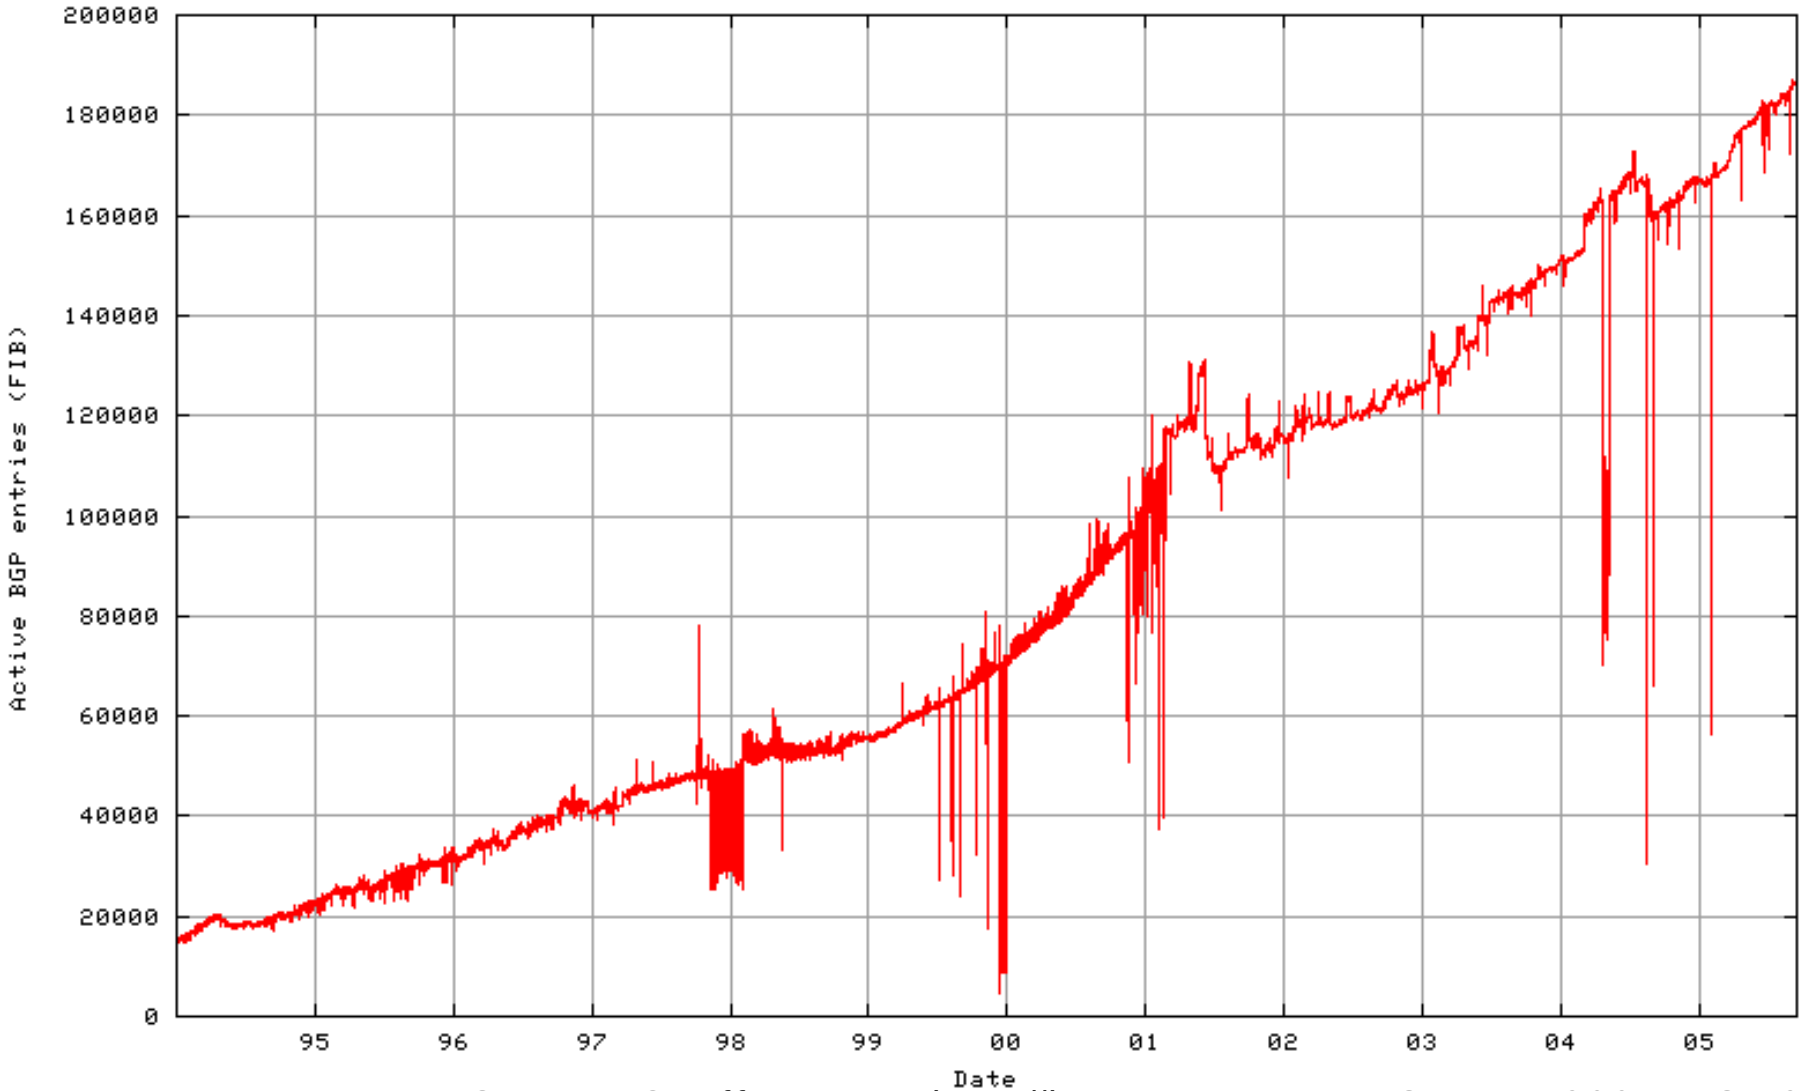

Source: Geoff Huston, http://bgp.potaroo.net, Sep 14, 2005, AS 1221

Disadvantages of not using BGP

- You gain a bit more control of your destiny when you speak BGP yourself
  - Break up your routes in an emergency
  - Tune traffic
- You can “pad” your announcements to de-prefer one or more up-streams
- Also, you lose the ability to fine-tune outbound traffic flow to the “best” upstream

Tips and Tricks 9

Hot-potato routing

BGP: Some Basics

- BGP exchanges routes between ASs.
- ASNs are stamped on the routes "on the way out"
  - adding one "AS hop" per network traversed → AS path
  - no concept of pipe size, internal router hop-count, congestion → in some sense BGP treats all ASs the same
- Routes are exchanged over "peering sessions"
  - Sessions run on top of TCP, port 179
  - The routes are "objects", or "bags of attributes"
- BGP is actually two protocols
  - iBGP, designed for "internal" route exchange
  - eBGP, designed for "external" route exchange
- 1995: BGP-4 [RFC 1771]
  - Support for Classless Inter-domain Routing (CIDR)

BGP: Net Prefixes, ASNs and Route Advertisements

BGP route advertisement:
Net prefix: 207.8.128.0/17
AS path: 4969 6461

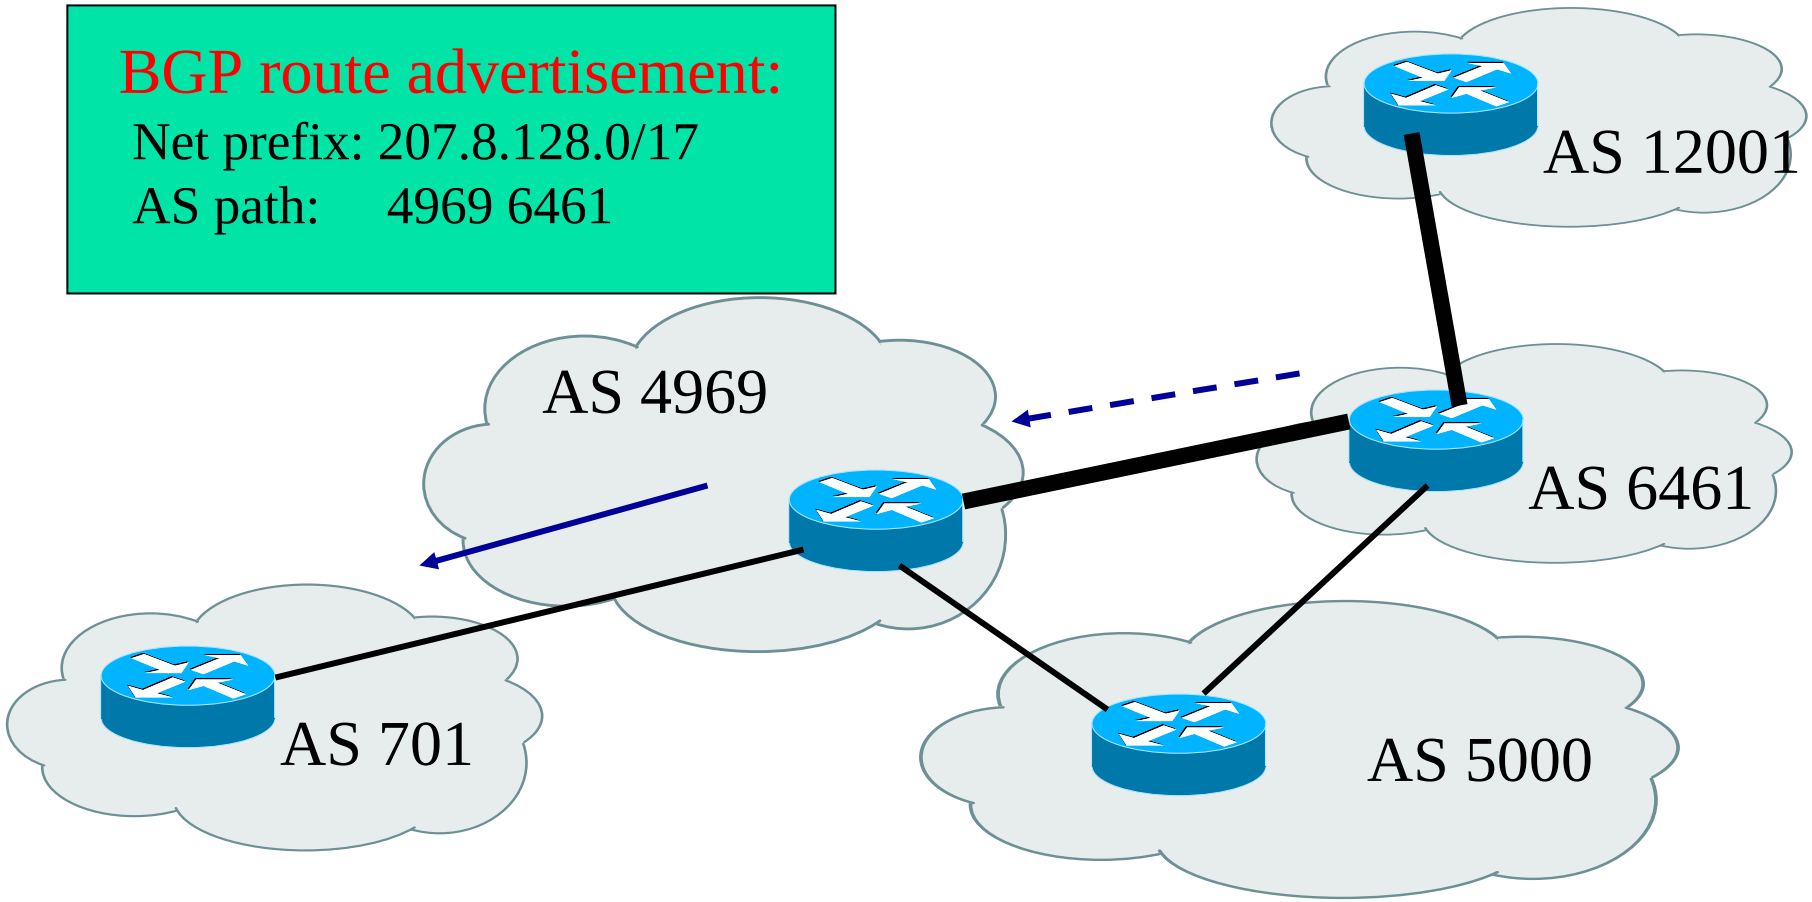

BGP Route Advertisement

- Think of a BGP route as a “promise”
  - If I advertise 207.8.128.0/17, I promise that if you deliver traffic destined to any IP address within 207.8.128.0/17 to me, I know how to deliver it (at least as well as anyone else)
- By making sure these routes are heard by all ASes, your provider ensures a return path for all of your packets
  - Sending packets *out* is easier than getting them back.
  - Sending routes *out* causes IP traffic to come *in*

BGP Route Advertisements and IP Address Space

- Route aggregation is used
- Recall “longest prefix matching” to look up forwarding table
  - If one of my customers’ ISPs is advertising 207.8.240.0/24, all incoming traffic from other networks will start flowing in that pipe.
  - So I must “punch a hole” in my aggregate announcement and advertise 207.8.128.0/17 and 207.8.240.0/24

BGP Operations (Simplified)

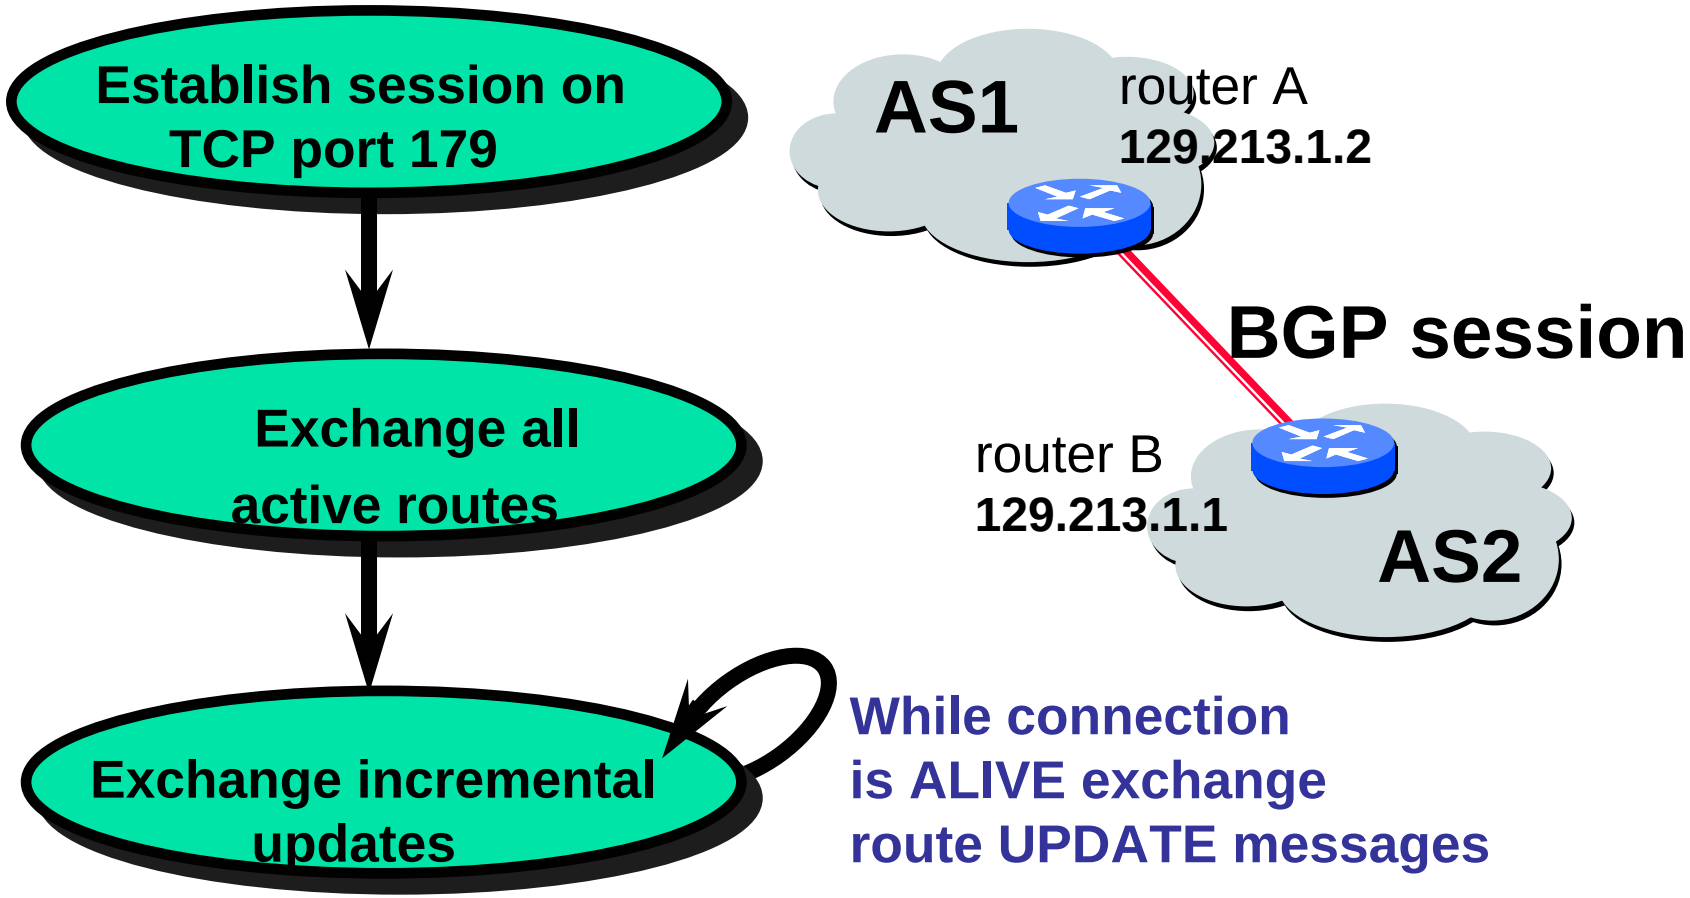

BGP Peering Sessions

- BGP session set up over TCP
  - When session set up, both sides flood the other end with all of their best BGP routes
  - Over time, only incremental updates are exchanged
  - If session dies, all associated routes must be withdrawn
- BGP peers (neighbors) must be specified explicitly
- BGP session set-up: Cisco Example

Router A in AS 1

```
router bgp 1
neighbor 129.213.1.1 remote-as 2
```

Router B in AS 2

```
router bgp 1
neighbor 129.213.1.2 remote-as 1
```

Customers and Providers

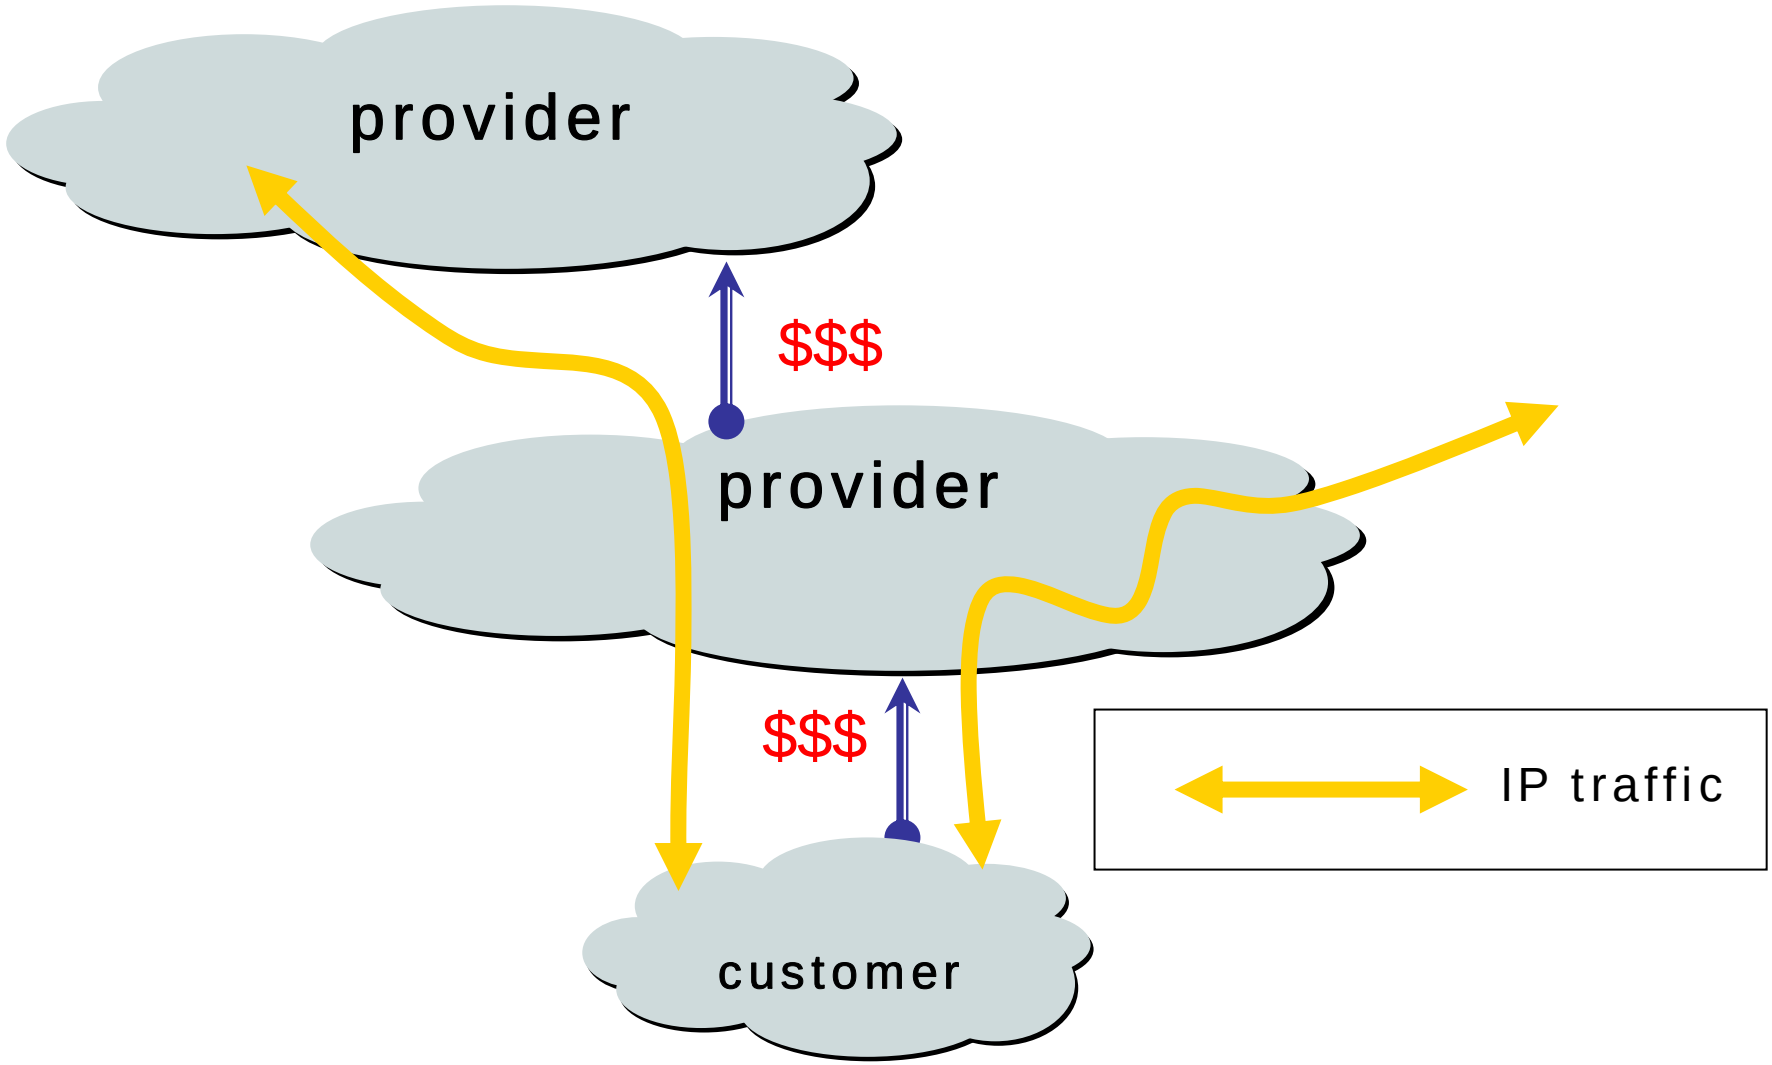

The Peering Relationship

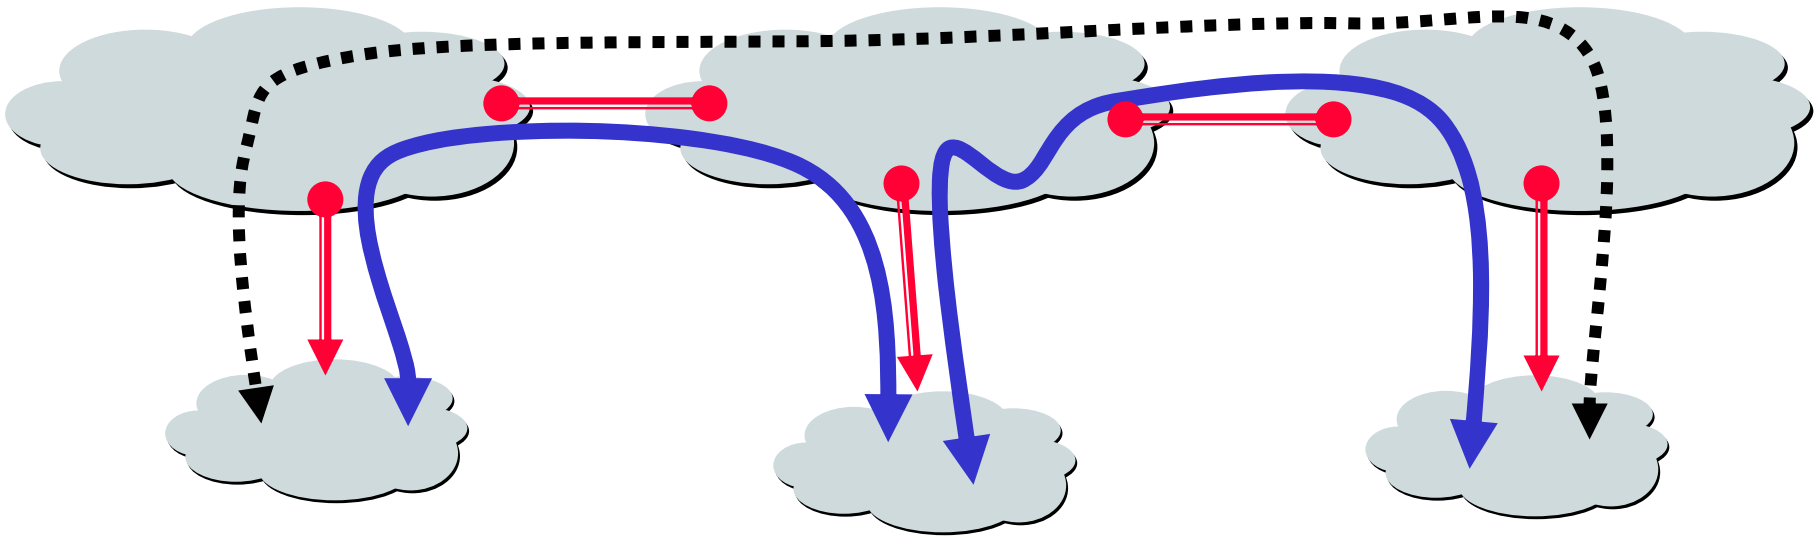

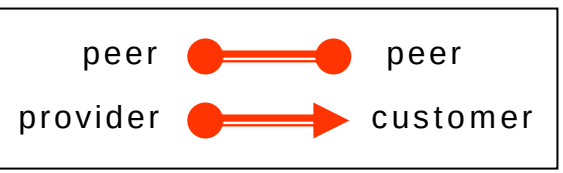

traffic allowed

traffic NOT allowed

Peers provide transit between their respective customers

Peers (often) do not provide transit between peers

Peers (often) do not exchange $$$

Peering Provides Shortcuts

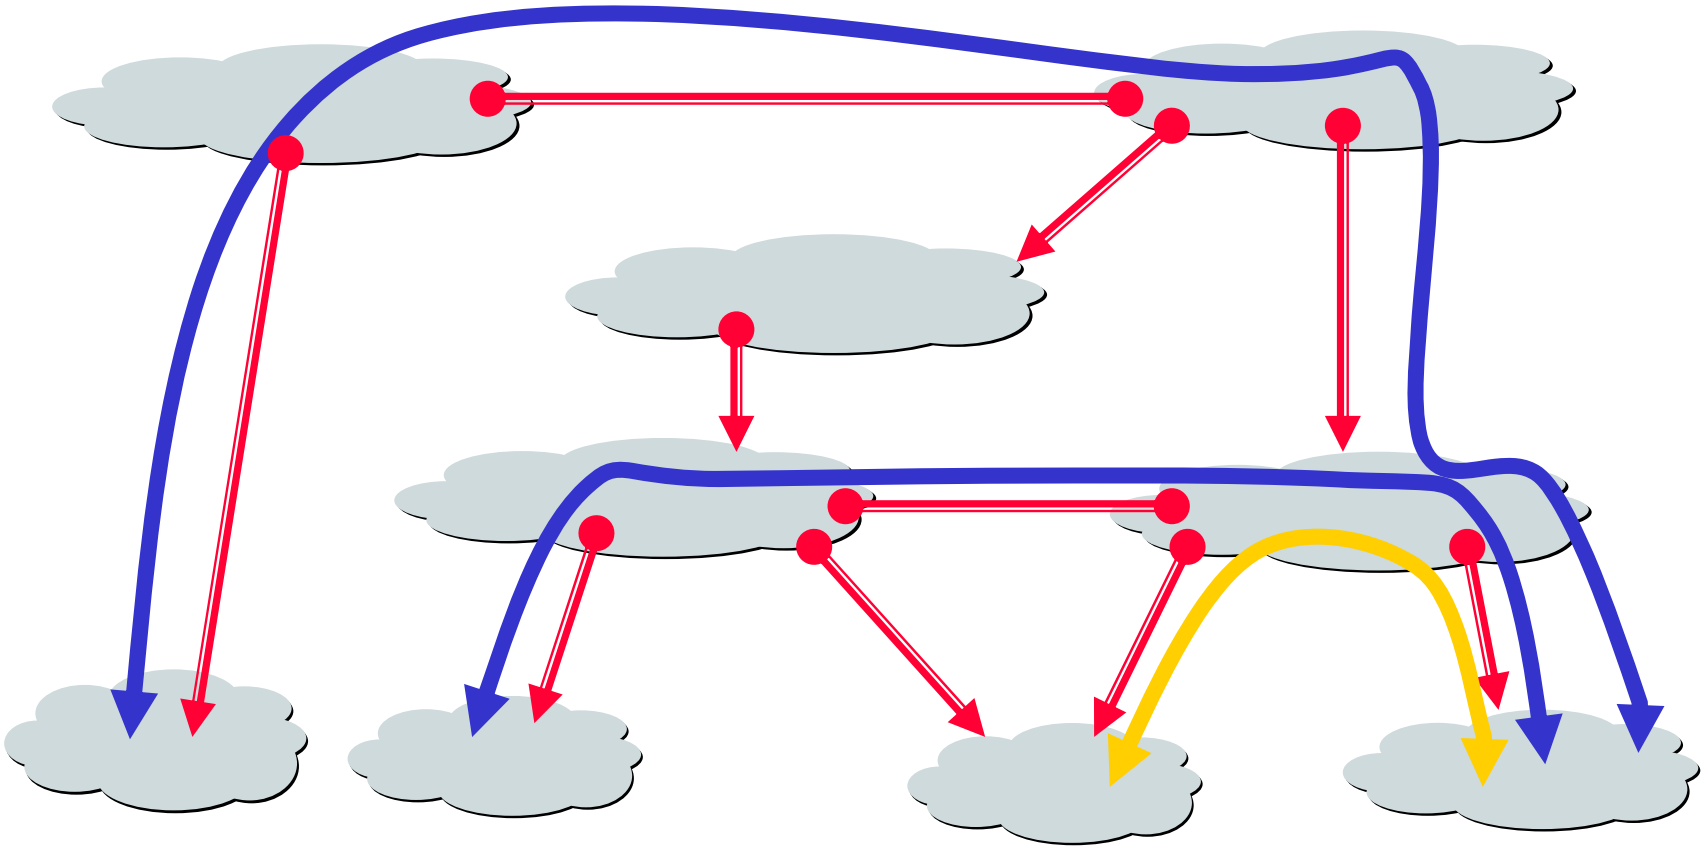

Peering also allows connectivity between the customers of “Tier 1” providers.

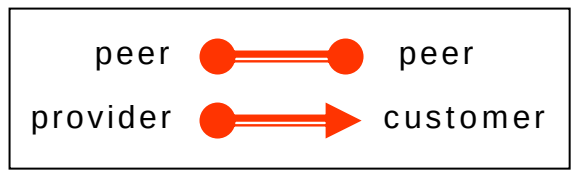

Don't Always Need BGP or an ASN

Static routing is the most common way of connecting an autonomous routing domain to the Internet. This helps explain why BGP is a mystery to many …

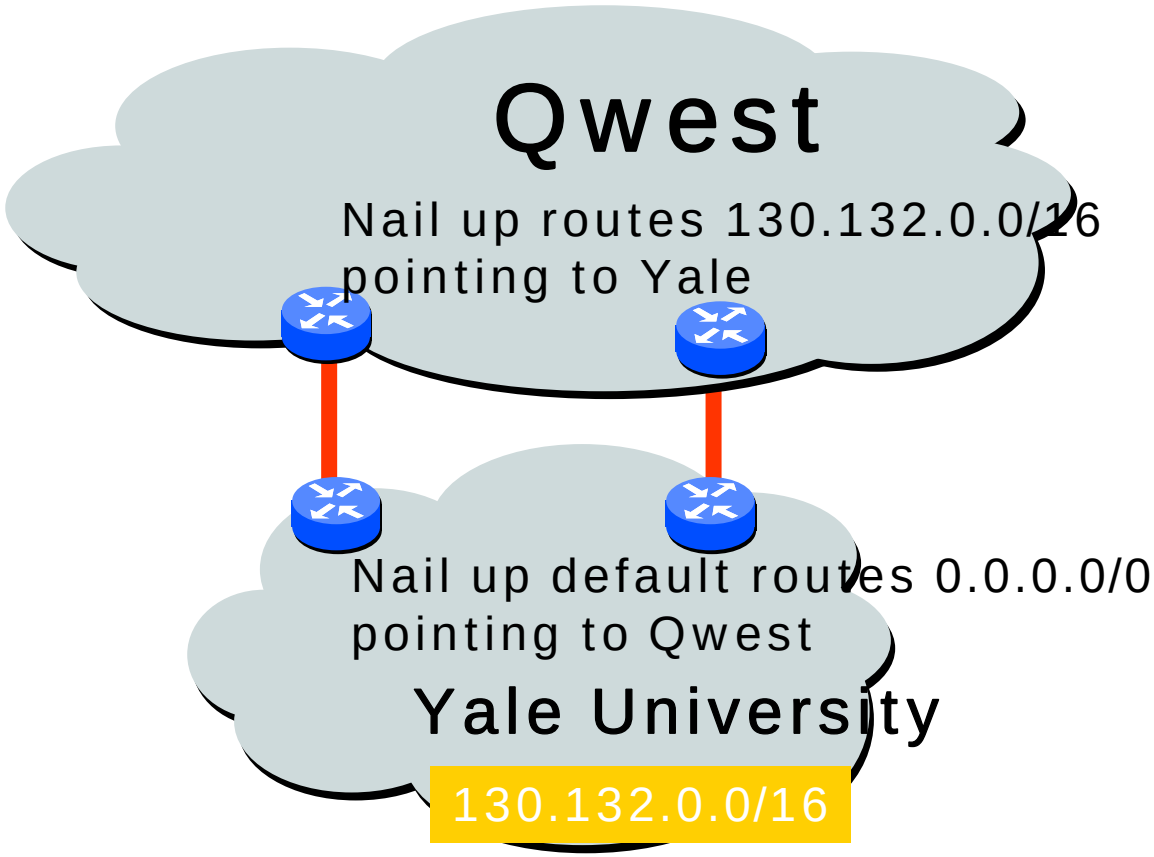

EBGP vs. IBGP Sessions

- EBGP: between (usually directly-connected) routers in different ASs
- IBGP: between (BGP-speaking) routers in same AS
- Different (operational) rules and polices apply!

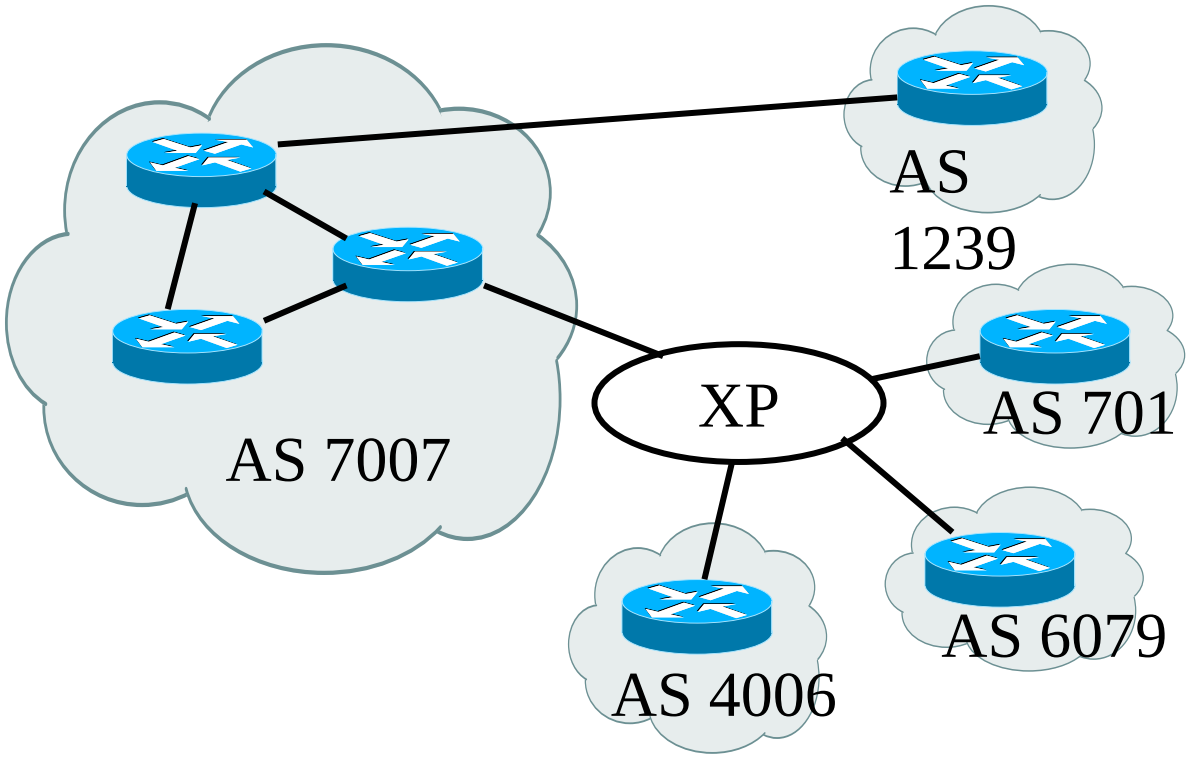

iBGP

- IBGP speakers are (usually) fully meshed: why?
- IBGP session set up:

Router A in AS 3847

router bgp 3847
neighbor 129.213.1.1 remote-as 3847
neighbor 128.28.10.2 remote-as 3847

Router B in AS 3947

router bgp 3847
neigbhor 129.213.1.2 remote-as 3847
neighbor 127.101.1.1 remote-as 3847

Router C in AS 3947

router bgp 3847
neigbhor 128.28.10.1 remote-as 3847
neigbhor 127.101.1.2 remote-as 3847

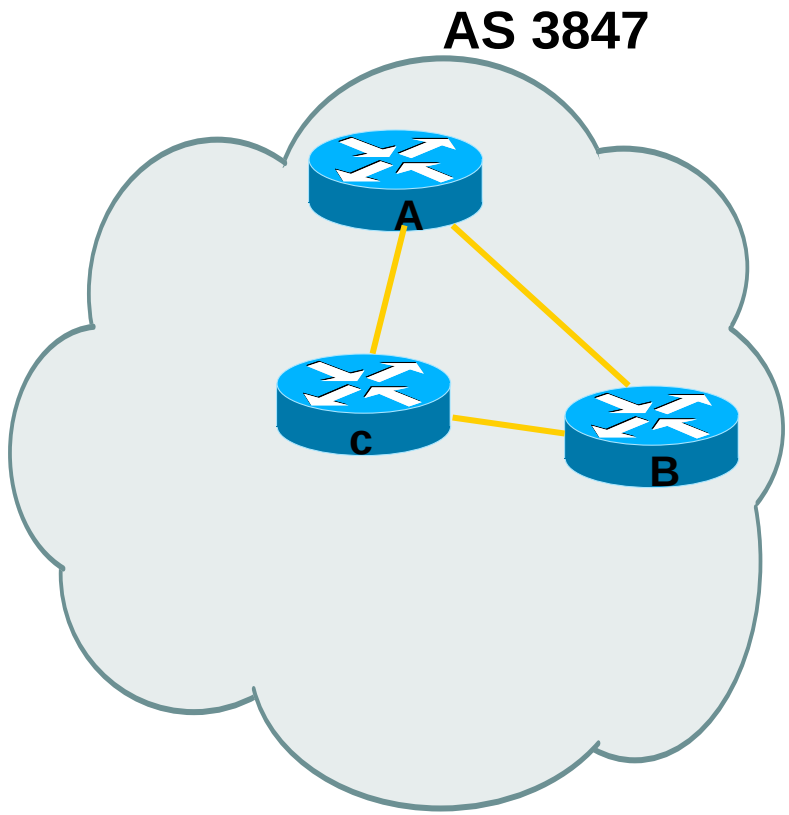

BGP Messages: Four Types

- **Open** : Establish a peering session.
- **Keep Alive** : Handshake at regular intervals.
- **Notification** : Shuts down a peering session.
- **Update** : Announcing new routes or withdrawing previously announced routes.

route announcement
=
prefix + attributes values

What is an Attribute?

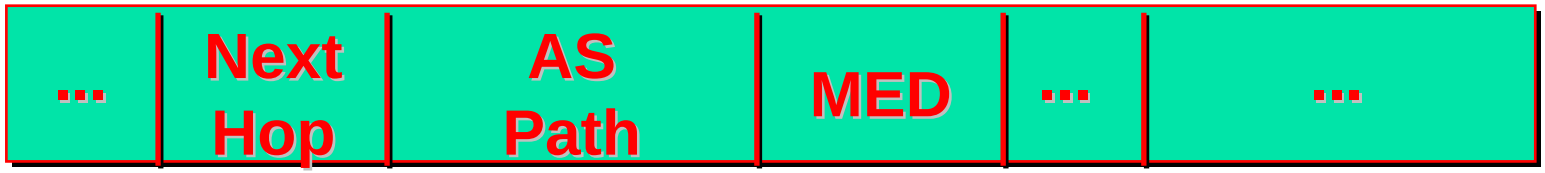

- A BGP message consists of a prefix and information about that prefix (i.e., local-pref, med, next-hop, originator, etc...)
- Attribute encoded in a TLV (type-length-value) format.
- Attribute length is 4 bytes long
- Attributes can be transitive (across ASs) or non-transitive (between AS neighbors only)
- Some are mandatory: e.g., AS Path, Next-Hop, etc.

BGP Attributes

Value	Code	Reference
1	ORIGIN	[RFC1771]
2	AS_PATH	[RFC1771]
3	NEXT_HOP	[RFC1771]
4	MULTI_EXIT_DISC	[RFC1771]
5	LOCAL_PREF	[RFC1771]
6	ATOMIC_AGGREGATE	[RFC1771]
7	AGGREGATOR	[RFC1771]
8	COMMUNITY	[RFC1997]
9	ORIGINATOR_ID	[RFC2796]
10	CLUSTER_LIST	[RFC2796]
11	DPA	[Chen]
12	ADVERTISER	[RFC1863]
13	RCID_PATH / CLUSTER_ID	[RFC1863]
14	MP_REACH_NLRI	[RFC2283]
15	MP_UNREACH_NLRI	[RFC2283]
16	EXTENDED COMMUNITIES	[Rosen]
...		
255	reserved for development	

From IANA: http://www.iana.org/assignments/bgp-parameters

Most important attributes

Not all attributes need to be present in every announcement

Next Hop Attribute

- Next-hop IP address to reach a network.
- Router A will advertise 198.3.97.0/24 to router B with a next-hop of 207.240.24.202.
- With iBGP, the next-hop does not change.
- IGPs should carry route to next-hops, using intelligent forwarding decision (i.e., via IGP).

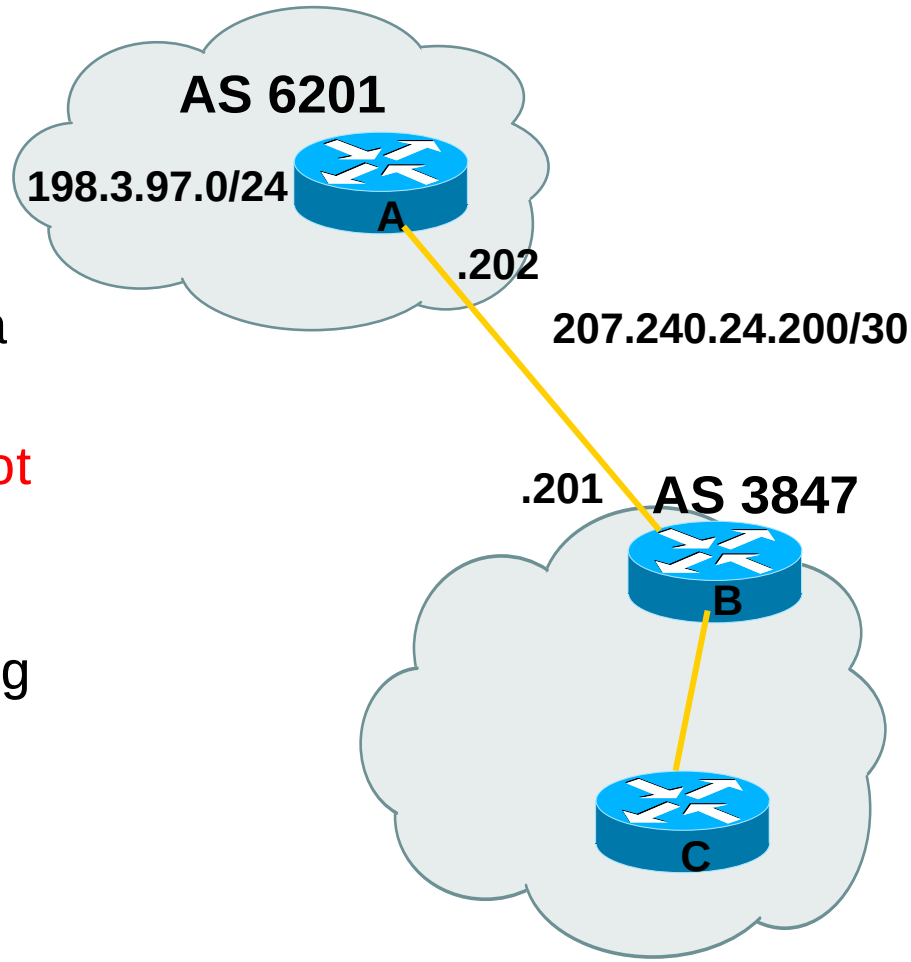

AS Path Attribute

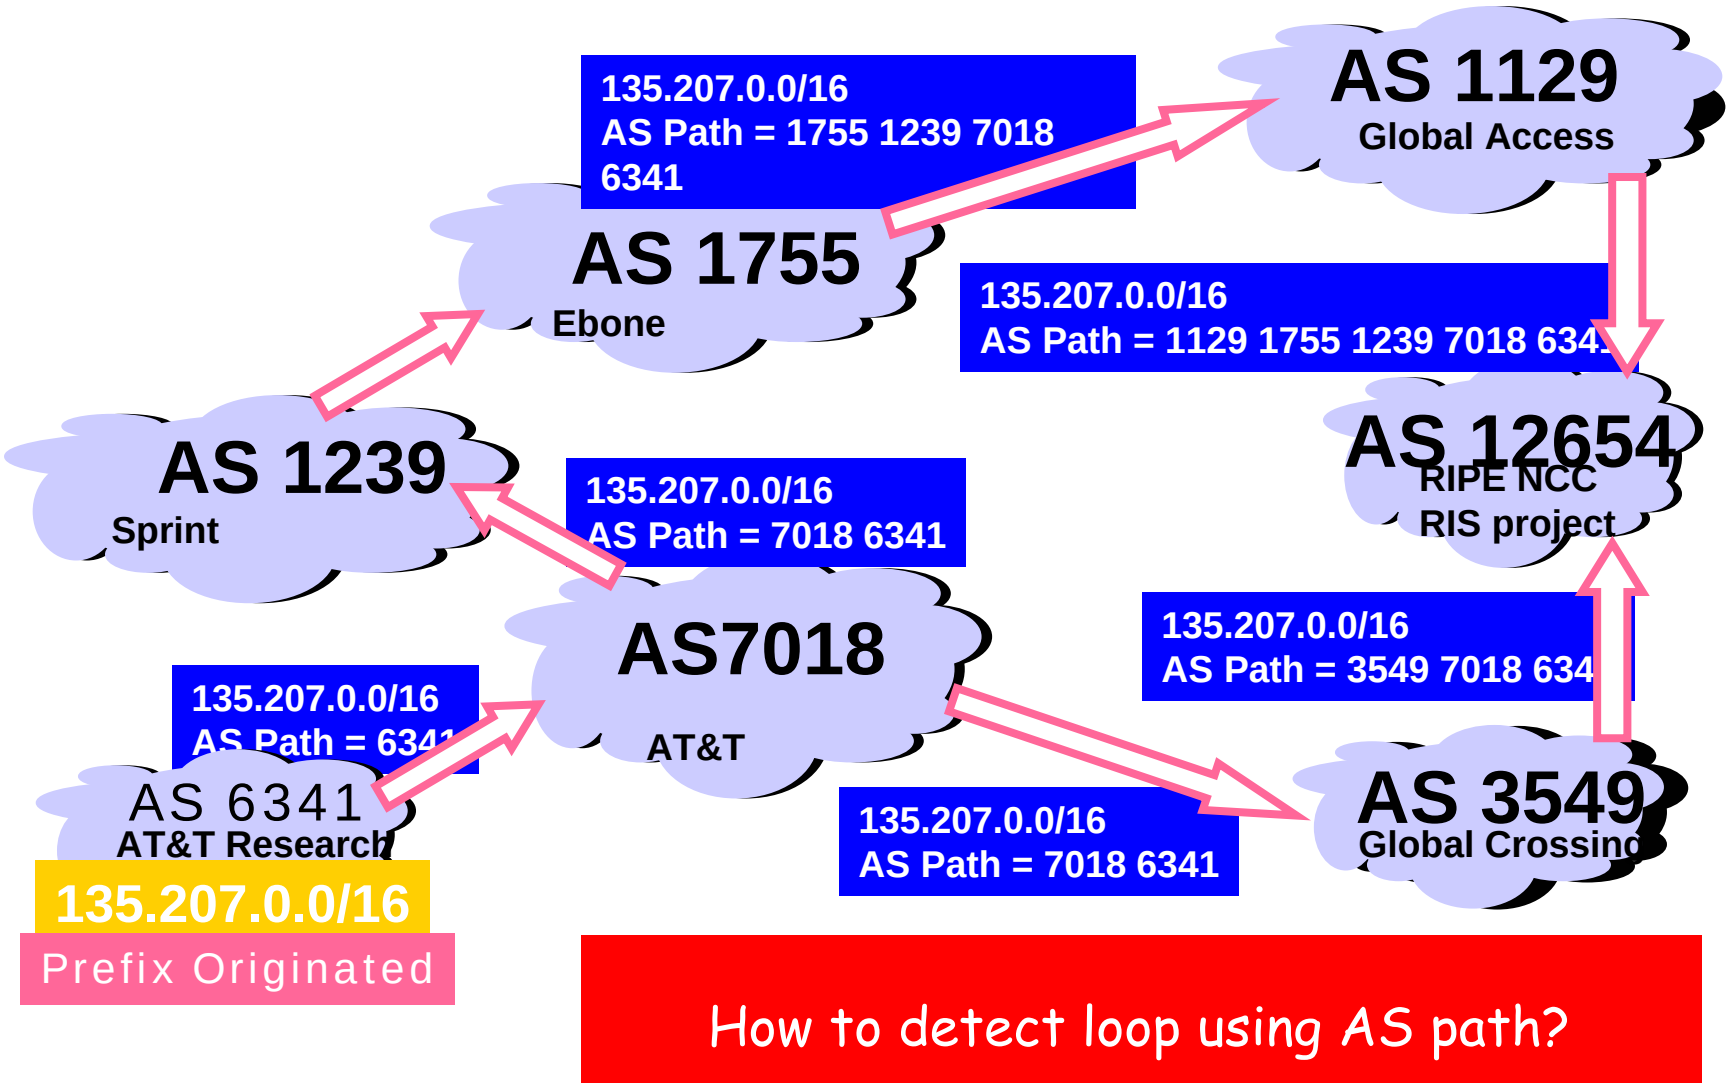

Local Preference Attributes

- Local to AS
  - transitive throughout your network. It is never advertised to an eBGP peer.
- Used to influence BGP path selection
- Default 100
  - Highest local-pref preferred
- For example, you can express the policy “prefer private connect” by making the “local_pref” be 150 and leaving all other peers at 100.

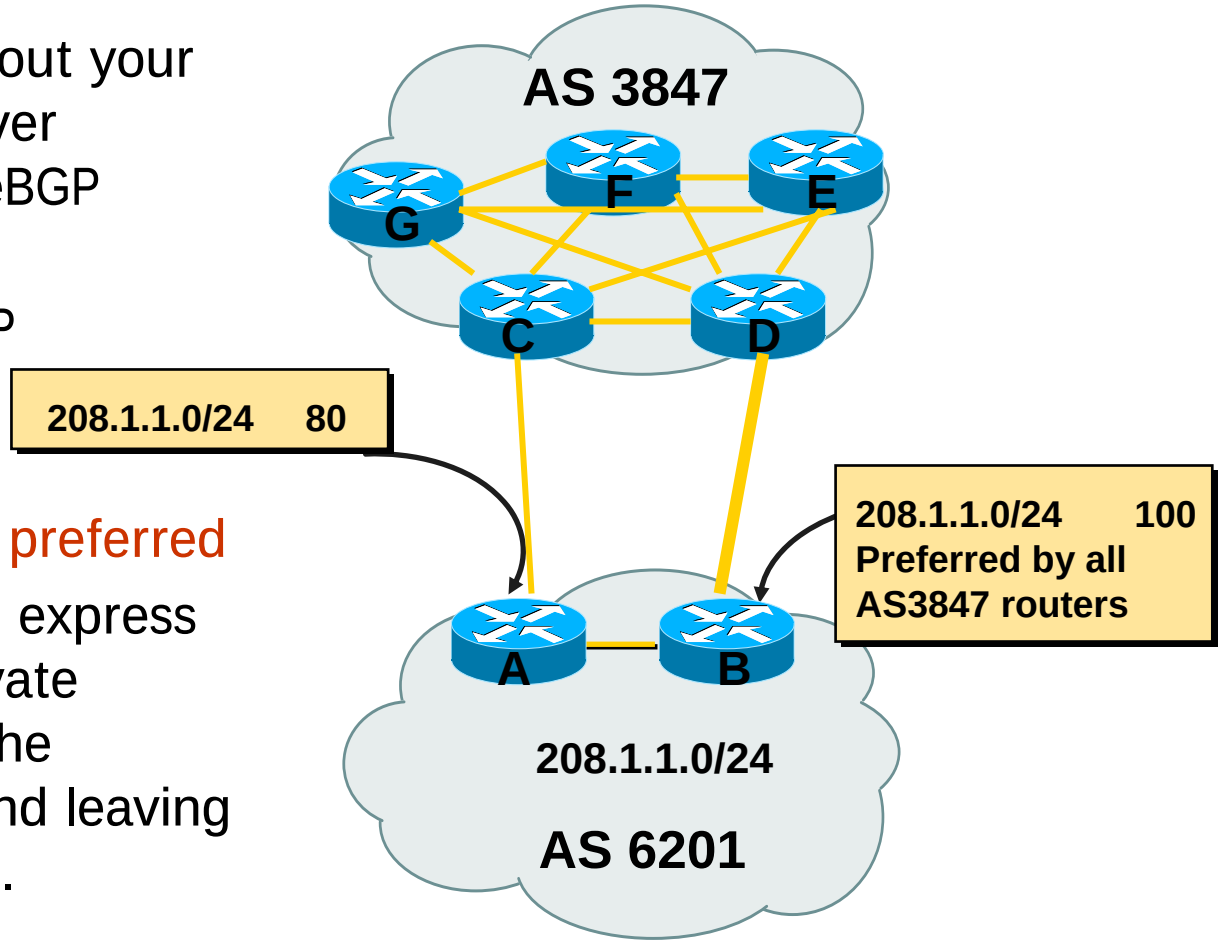

Multi-Exit Discriminator (MED) Attribute

- Indication to *external* peers of preferred path into an AS
  - Advertised to external neighbors
  - Neighbors are *not* obliged to heed it
- Affects routes with *same AS path*
* Lowest MED preferred
- A commonly used attribute by ISPs
  - Usually based on IGP metric
  - For example, big ISPs with multiple connections with each other use MED to indicate which PoP is “closest” to an advertised route, thus more preferred
- It comes after AS_PATH in evaluation, and thus isn’t quite as much of a “hammer” as local-pref

MED: Example

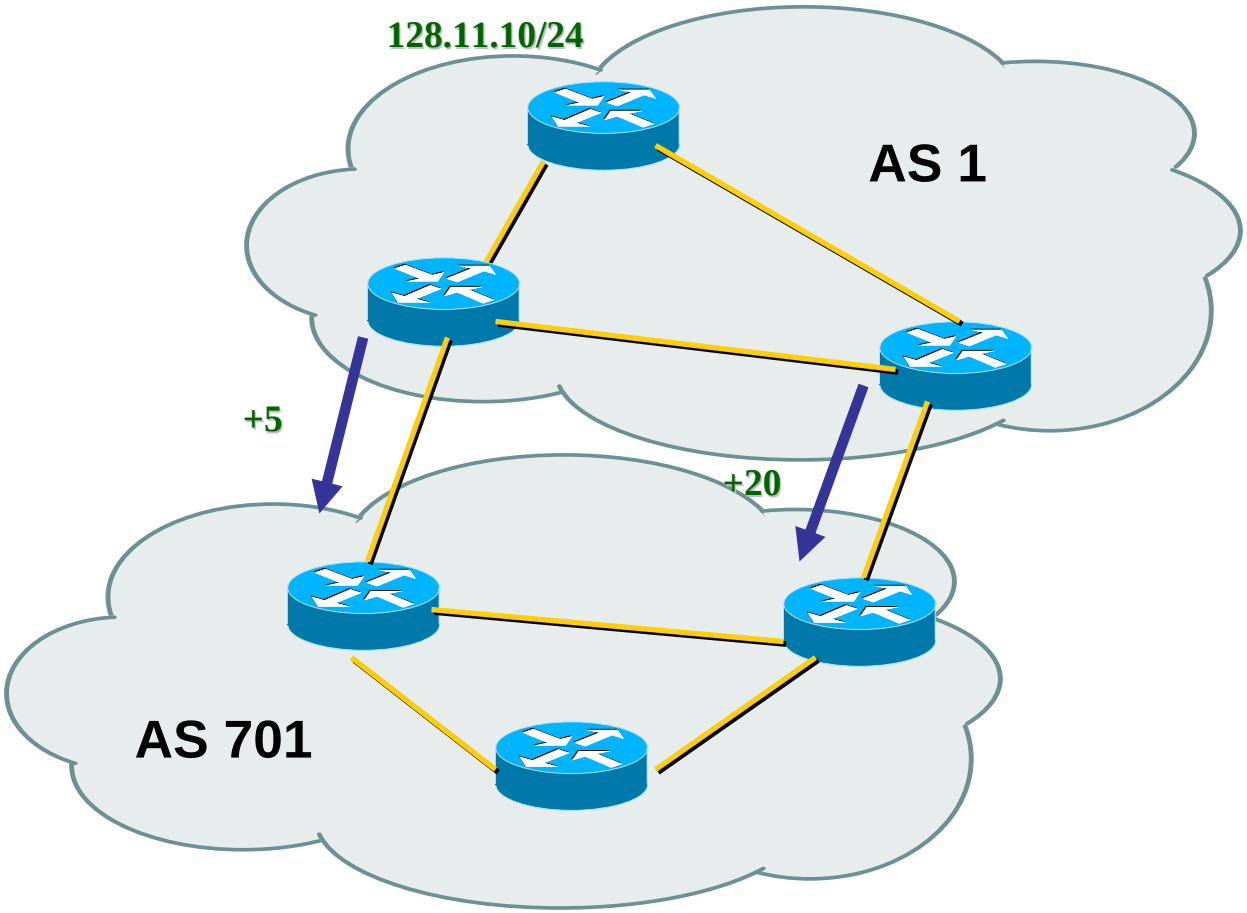

Weight Attribute

- Cisco proprietary, not part of any spec.
- Local to router.
- Value 0-65535 (default if originated by router - 32768, other - 0)
* Highest weight preferred
- Weight is rarely used. It overrides almost all other attributes in the decision path, and *is local to a specific router* - it is never sent to other routers, even ones inside your ASN.
- Usually used for temporary “I-don’t-have-time-to-think-about-it” fixes.

Origin Attribute

- One of the mandatory, but minor, attributes of a BGP route is the origin. It is one of (in order of preference):
  - IGP (i) (from a network statement)
  - EGP (e) (from an external peer)
  - Unknown (?) (from IGP redistribution)
- It can be re-set, but that is not often done.
- It is almost-last in the path selection algorithm.

Community Attribute

- Defined in RFC 1997
- 32-bit integer:
  - represented as two 16-bit integer x:y
- Used to group routes (“net prefixes”)
  - Each route could be member of multiple communities
- Transitive: carried across ASs
- Very useful in applying policies
- Well-known communities
  - No-Export: do not advertise to eBGP peers
  - No-advertise: do not advertise to any peer
  - Local-AS: do not advertise outside local AS (only used with “confederations”)

Community Attribute: Example

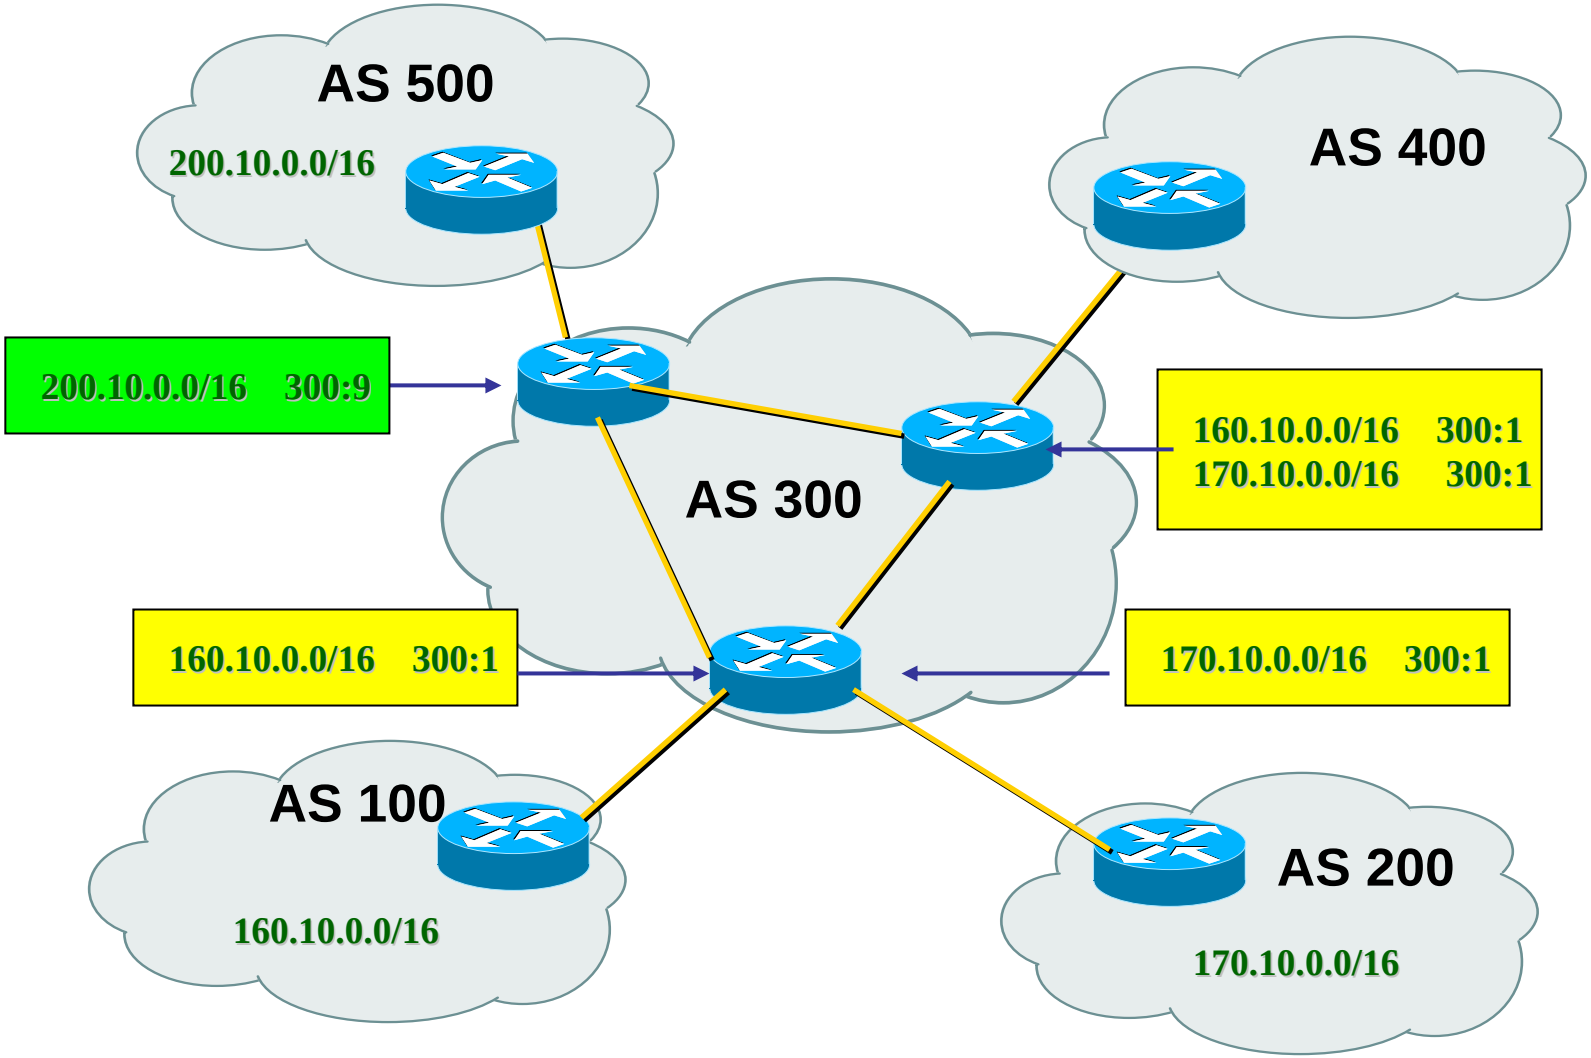

Attributes are Used to Select Best Routes

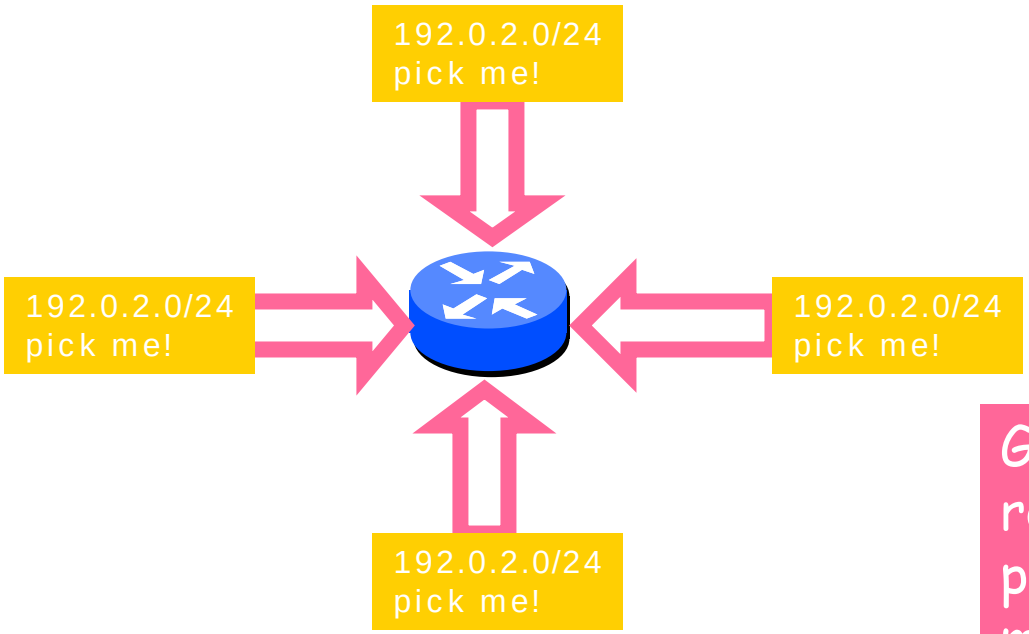

Given multiple routes to the same prefix, a BGP speaker must pick at most one best route

(Note: it could reject them all!)

Route Selection Priorities

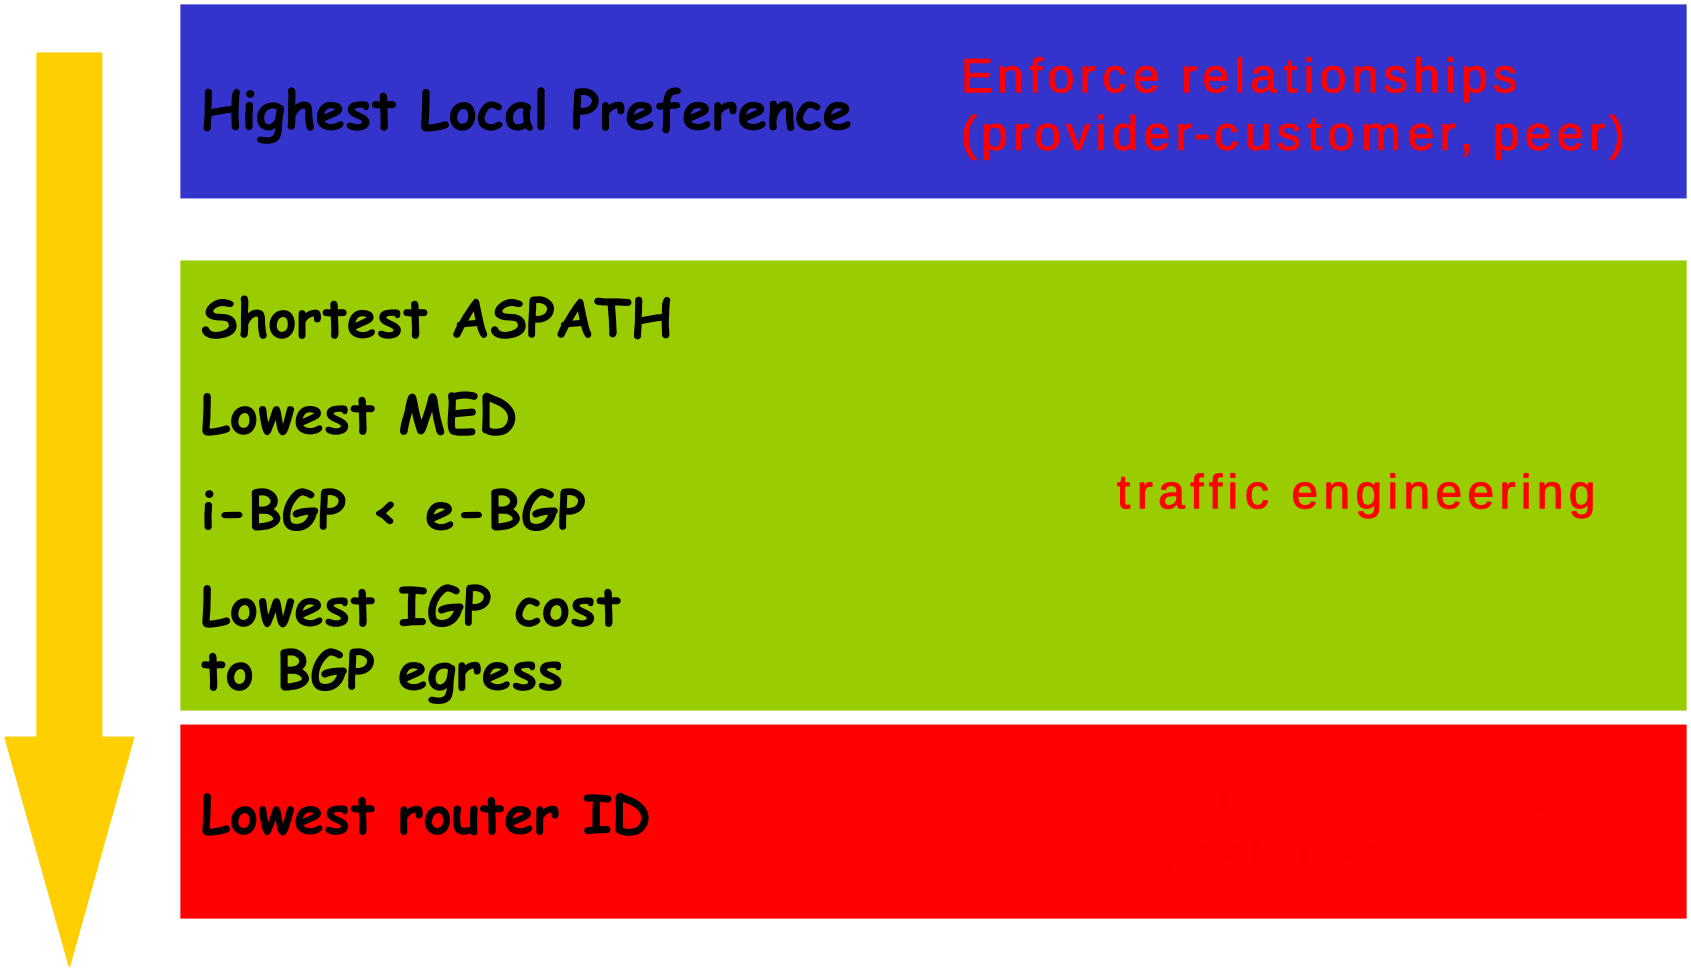

BGP Path Selection Algorithm

1. If next hop inaccessible, drop the update.
2. Largest weight (Cisco)
3. Largest local preference.
4. Paths originated by BGP running on this router.
5. Shortest AS_path.
6. Lowest origin type (IGP < EGP, EGP < incomplete).
7. Lowest MED attribute.
8. External path over the internal path.
9. Closest IGP neighbor.
10. Lowest IP address (BGP router ID).

Shorter Doesn't Always Mean Shorter

In fairness: could you do this "right" and still scale?

Exporting internal state would dramatically increase global instability and amount of routing state

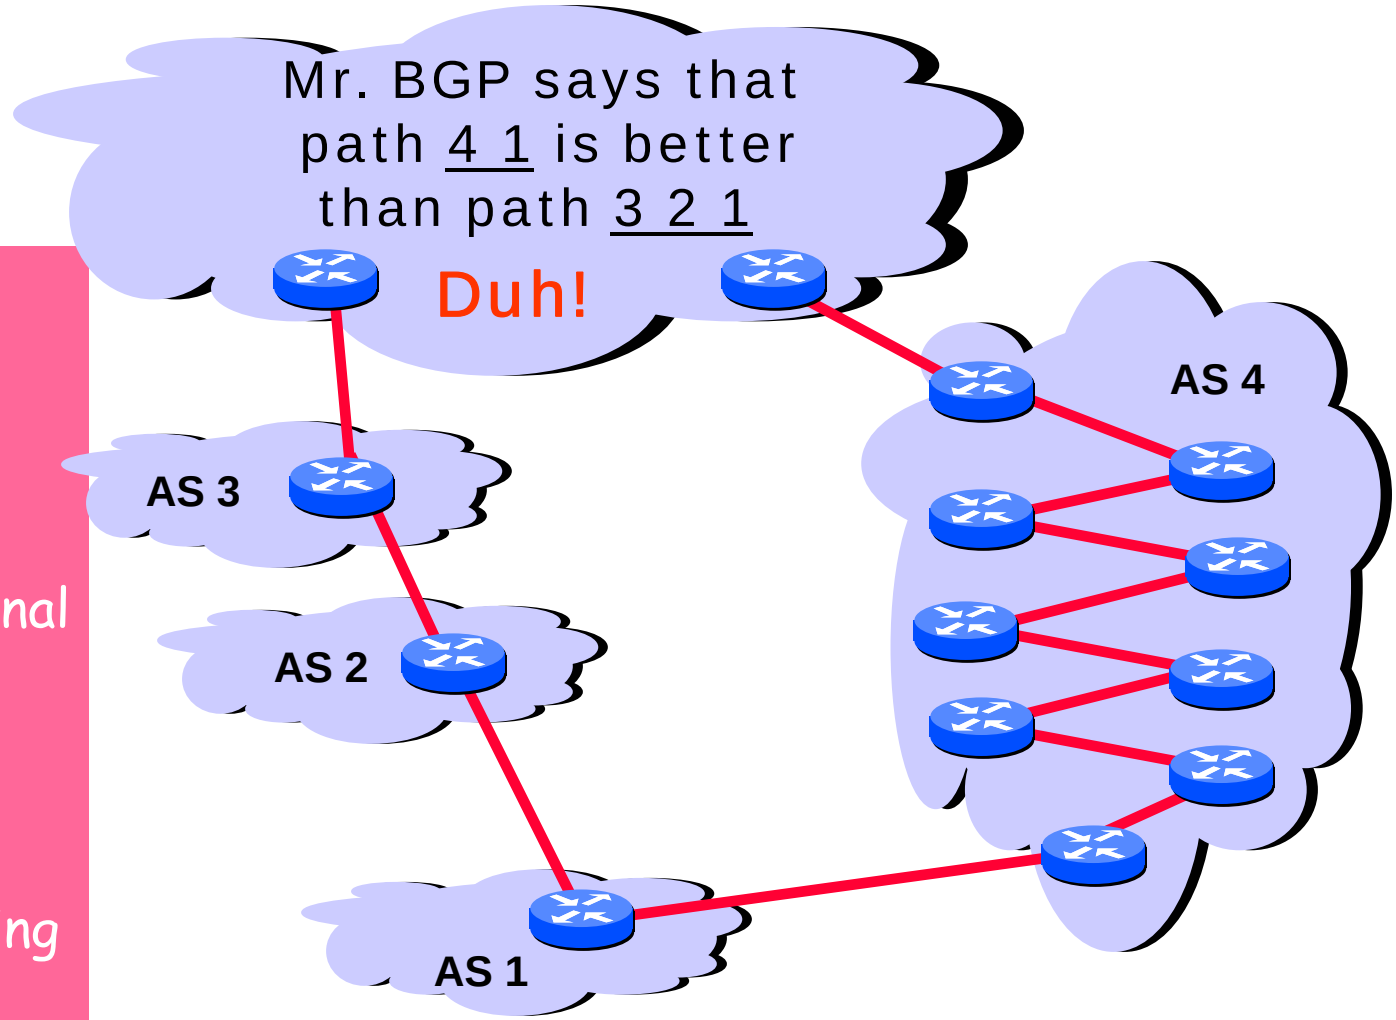

Tips and Tricks 11

The AS 7007 Incident

- Apr 1997, AS 7007 two BGP routers announced routes to most Internet, disrupt communications for 2 hours

The AS 3561 Incident

- Apr 2001, AS 3561 propagated > 5000 improper routes from a customer

eBGP vs. iBGP Revisited (1)

eBGP Rules:

- By default, only talk to directly-connected router.
- Sends the one best BGP route for each destination.
- Sends all of the important “attributes”; omits the “local preference” attribute.
- Adds (prepends) the speaker’s ASN to the “AS-Path” attribute.
- Usually rewrites the “next-hop” attribute.

eBGP vs. iBGP Revisited (2)

iBGP Rules:

- Can talk to routers many hops away by default.
- Can only send routes it "injects", or routes heard *directly* from an external peer.
- Thus, requires a *full* mesh.
- Sends all attributes.
- Leaves the "as-path" attribute alone.
- Doesn't touch the "next hop" attribute.
- With iBGP, next-hop is not a router *directly* connected.
  - So a "recursive lookup" is needed.
  - After the next-hop is found, a second lookup is made to figure out how to send the packet "in the direction" of the next-hop.

iBGP and Next-Hop: Example

In AS 2828:

- Router A: “next hop” for 170.10.0.0/16 will be the serial interface on Router D in AS1239 router
- This is true even in Router B’s and Router C’s forwarding table.

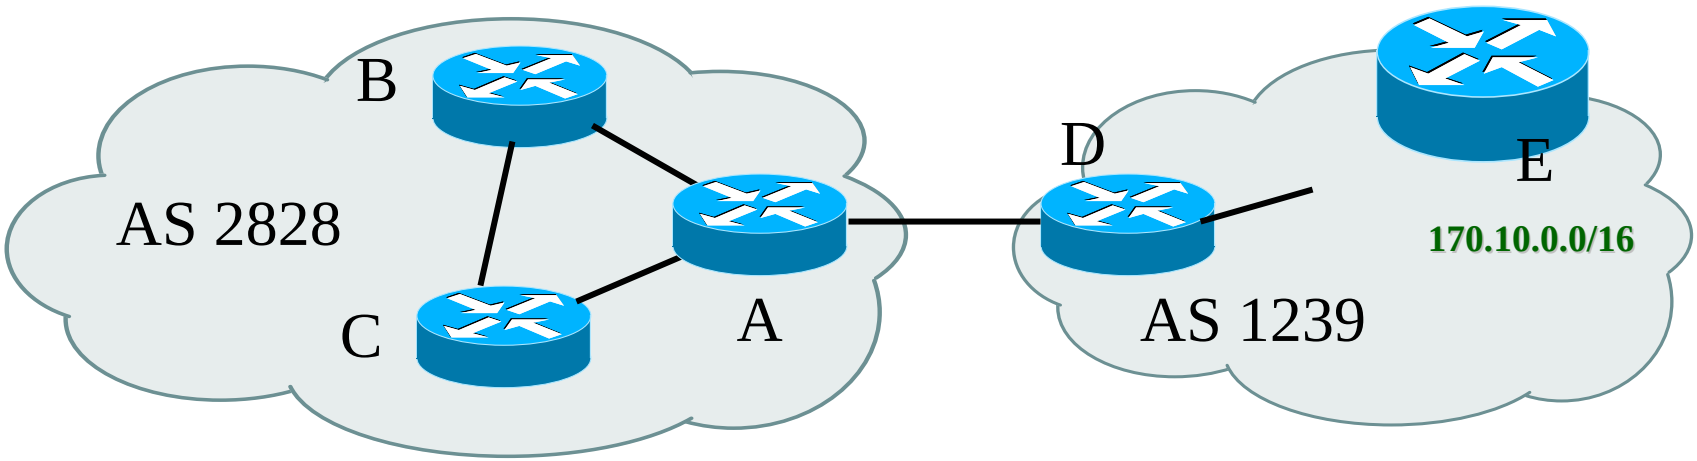

iBGP Route Distribution Restriction

- Assume AS1239 sends route 170.10.0.0/16 to AS2828. Router A will send that route to Routers B and C
- When Router B receives 170.10.0.0/16, it will not propagate that route to Router C because it was learned from an iBGP neighbor. Router C will behave similarly

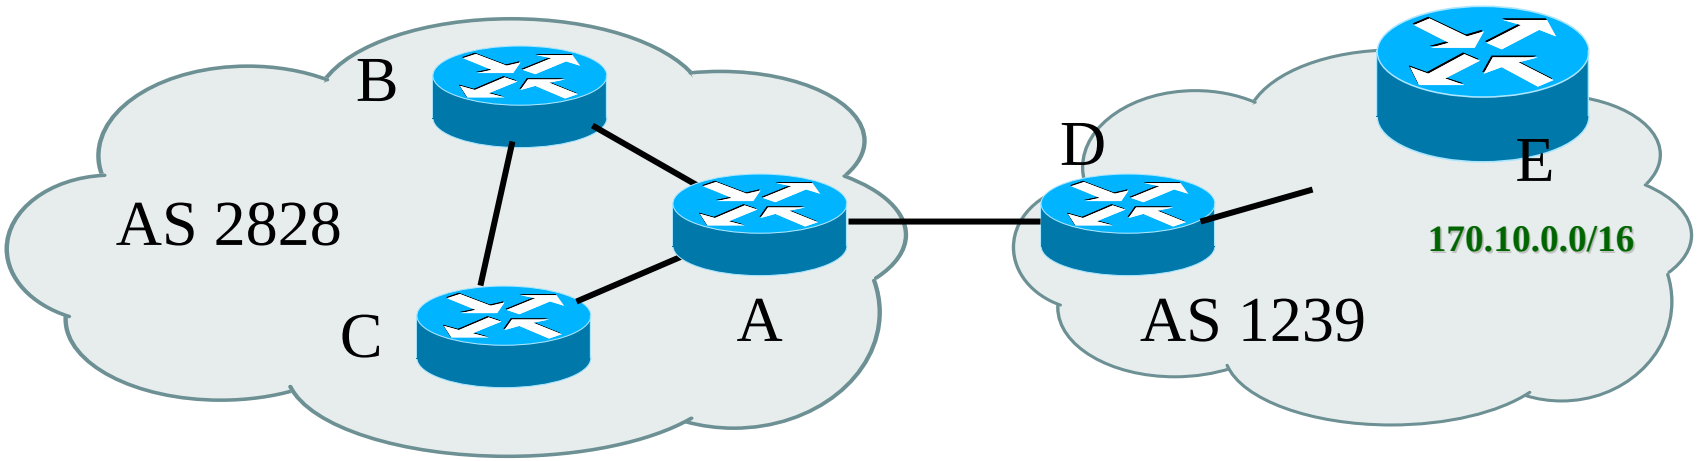

Synchronizaton

- In a transit AS, eBGP should not advertise a route before all routers in the AS learned about the route via IGP

- eBGP should wait until all the routes are propagated

Synchronization Example

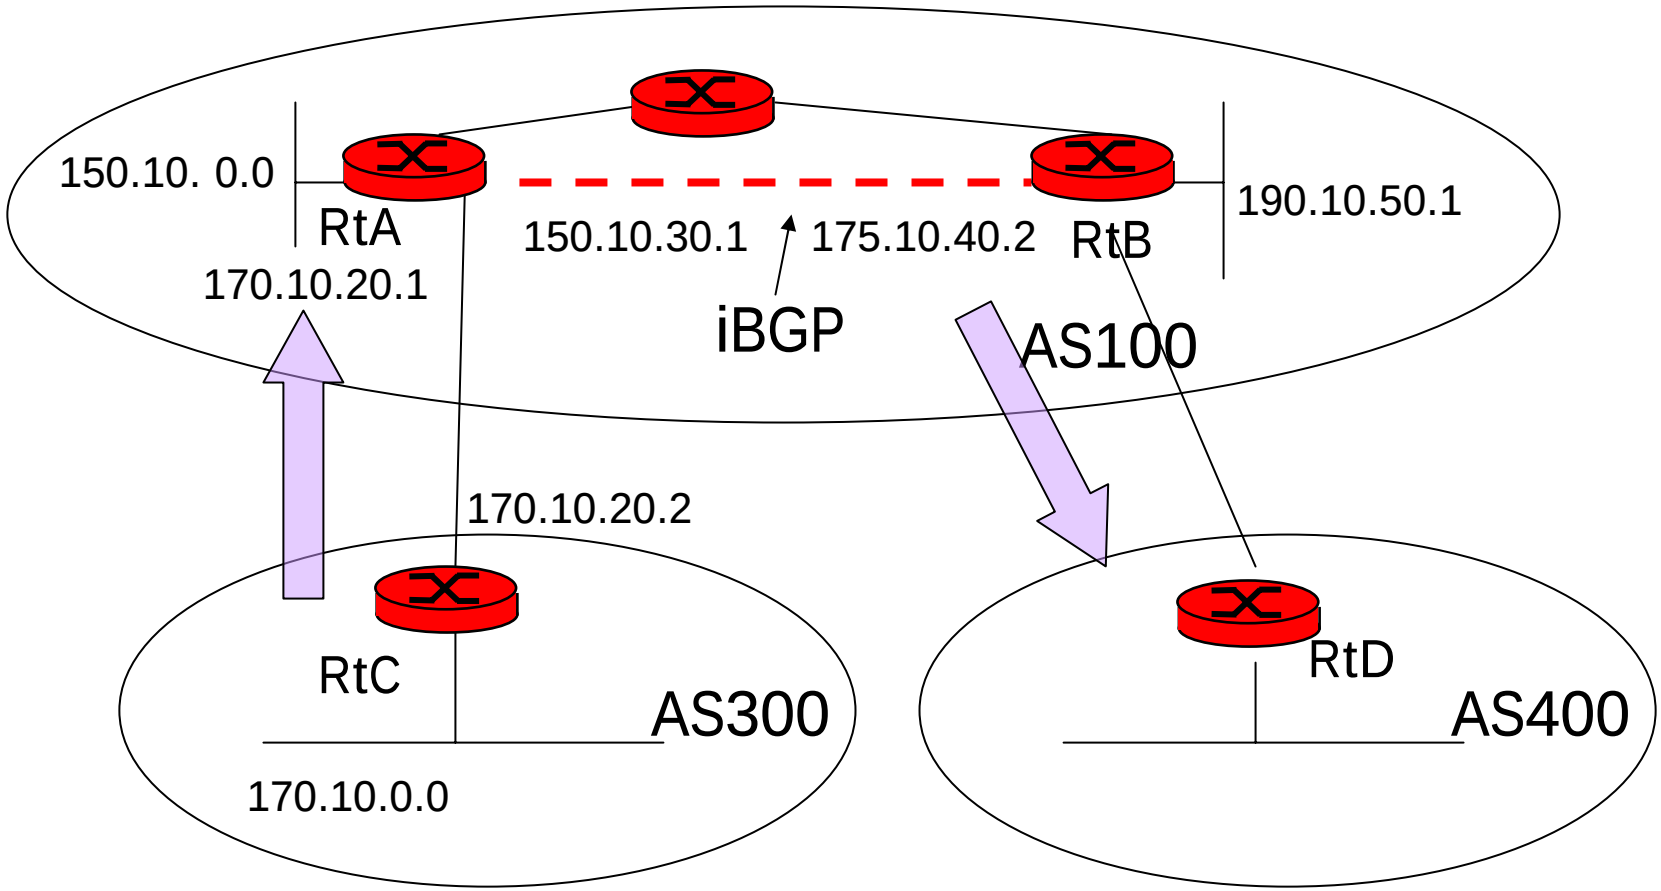

BGP Route Processing

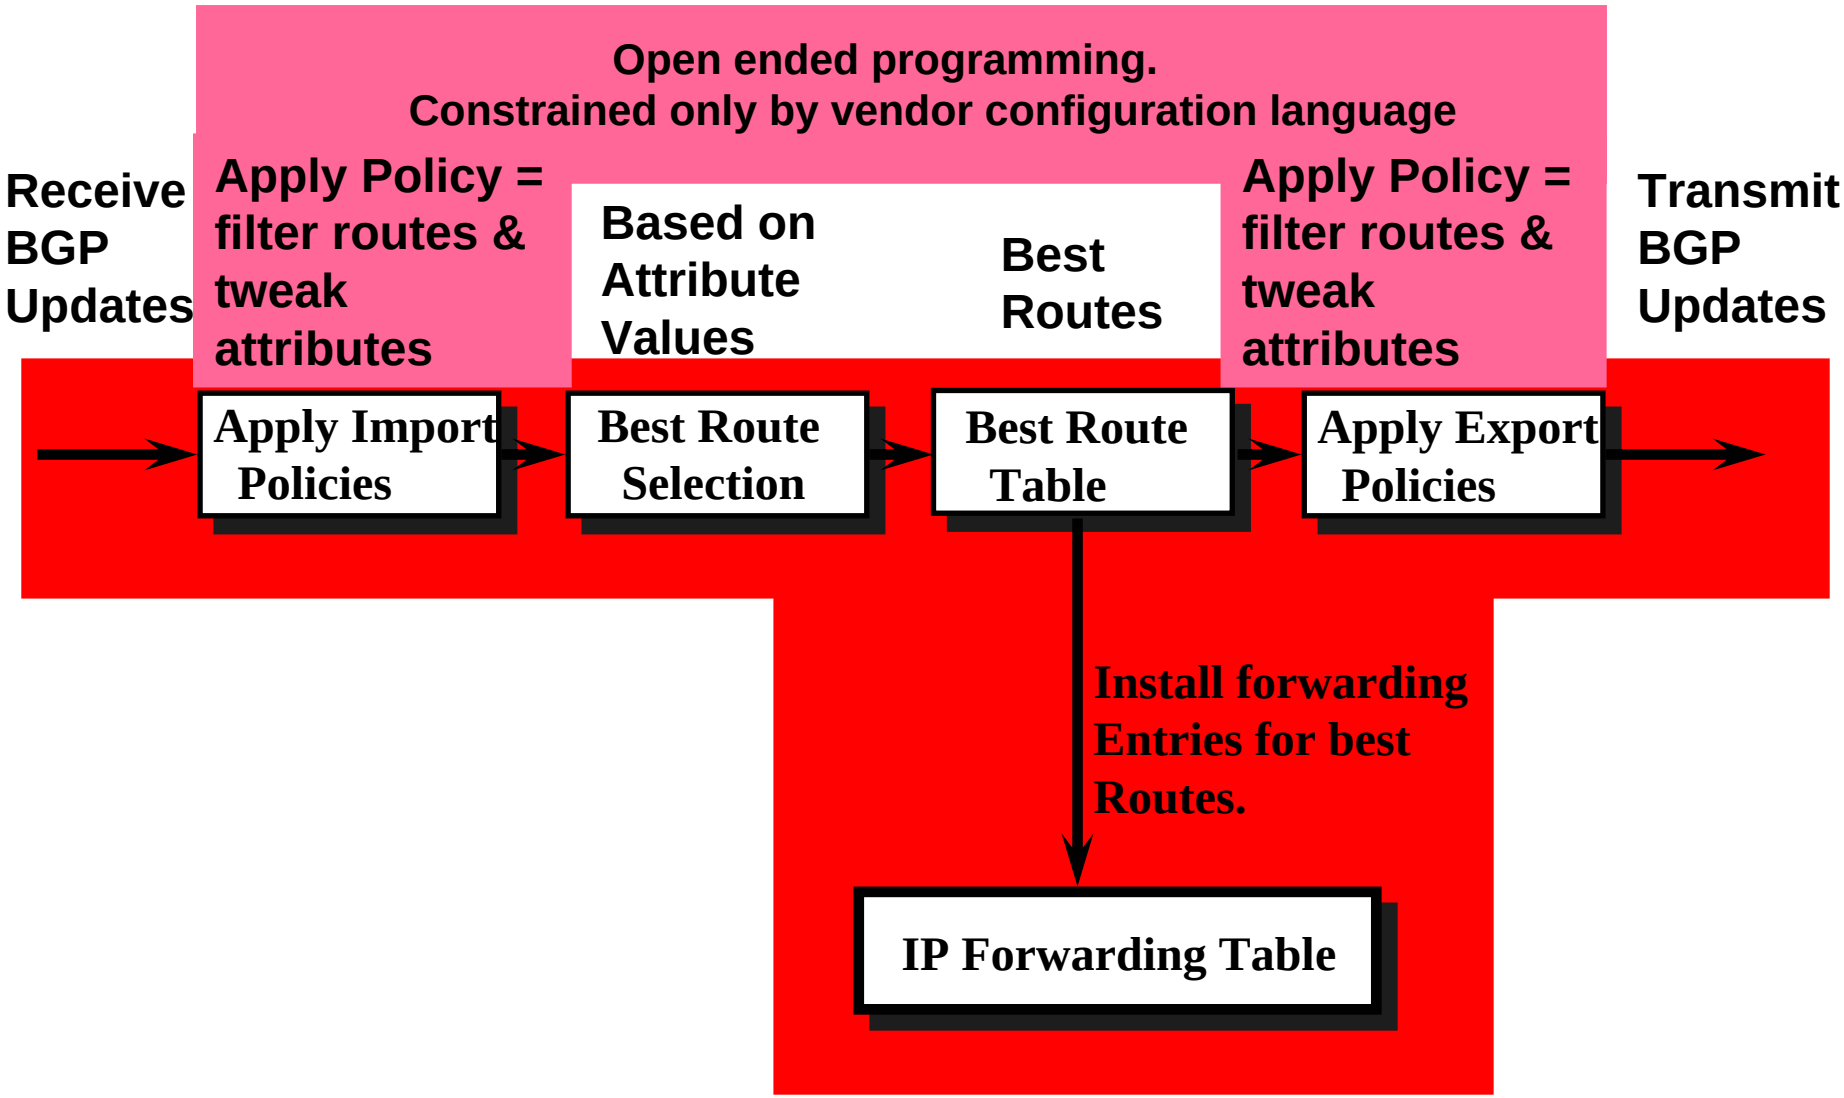

BGP Router - Processing Routes

- For each route received:
  - If it's a valid route AND passes any filters, it must be put into the BGP routing table.
  - Then, unless it is replacing a duplicate, a best-path computation must be run on all candidate BGP routes of the same prefix.
  - Then, if the best route changed, the RIB and/or FIB must be updated.
  - This process is done for ALL incoming BGP routes.

BGP Policy Control

- Filtering BGP routes
- Manipulate BGP route attributes
- Objective: control inbound/outbound traffic
- Some Cisco BGP filtering mechanisms:
  - To decide what routes can and can't go to various other routers, you can "filter" using:
    - "distribute lists" ("prefix filters") - lists of routes
    - "filter lists" ("as-path filters") - lists of regular expressions matching or denying ASs
    - "route maps" ("BGP Basic programs") that allow you to match and change most BGP attributes

Tweak Tweak Tweak

- For inbound traffic
  - Filter outbound routes
  - Tweak attributes on outbound routes in the hope of influencing your neighbor's best route selection
- For outbound traffic
  - Filter inbound routes
  - Tweak attributes on inbound routes to influence best route selection

In general, an AS has more control over outbound traffic

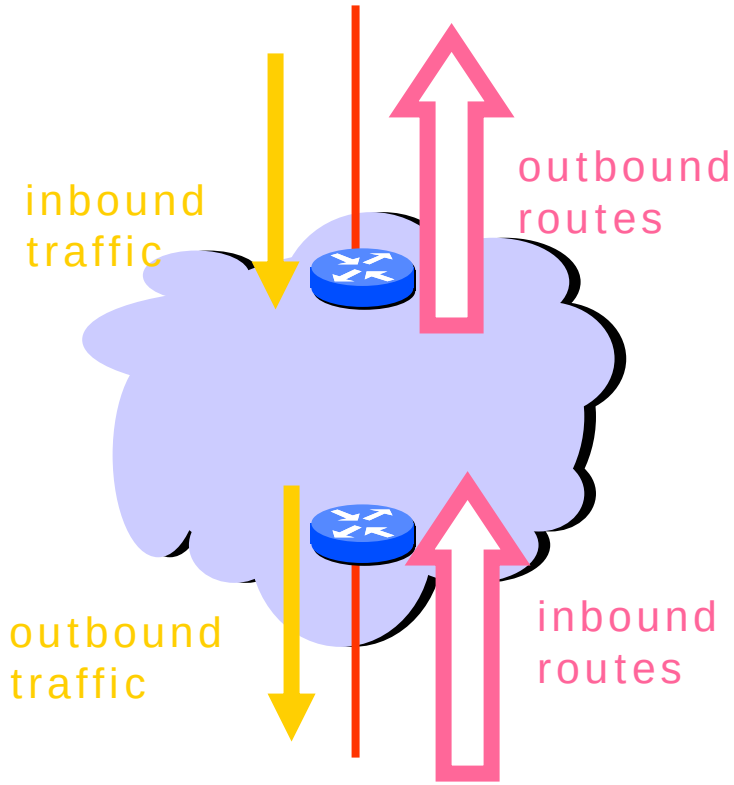

Tuning Inbound BGP Routes

- Inbound BGP routes make traffic go out
  - Having a route means that an outbound packet can use it as basis for a forwarding decision (well, the router can)
  - It is far easier to control outbound traffic than inbound
- Goal is generally to provide fastest, lowest-loss, path for all destinations,

  i.e., to optimize connectivity “quality”, whatever that is
  - E.g., to optimize throughput and latency
  - to reduce transit costs, say,
    - squash traffic via a certain provider,
    - prefer customer than peer/provider paths, and prefer peer than provider paths
  - to load balance, or to ensure reliability with back-up routes

Controlling Outbound BGP Routes

- Outbound BGP routes make traffic come in
  - It's a lot harder to control inbound traffic as other ASs' policies complicate your life!
- If you are a stub AS with a single connection
  - Not much you need to do except to filter out routes not in your AS
- If you are a multi-homed stub AS,
  - You may want to control through which link/provider that traffic to certain destinations in your AS may take, to load balance or for back-up
- If you are an ISP, you want to minimize transit cost,
  - carry transit traffic from customers only !
  - use "hot-potato" routing to hand off traffic to peers/providers as soon as possible
  - to load balance, or to ensure reliability with back-up routes

Shedding Inbound Traffic with ASPATH Padding Hack

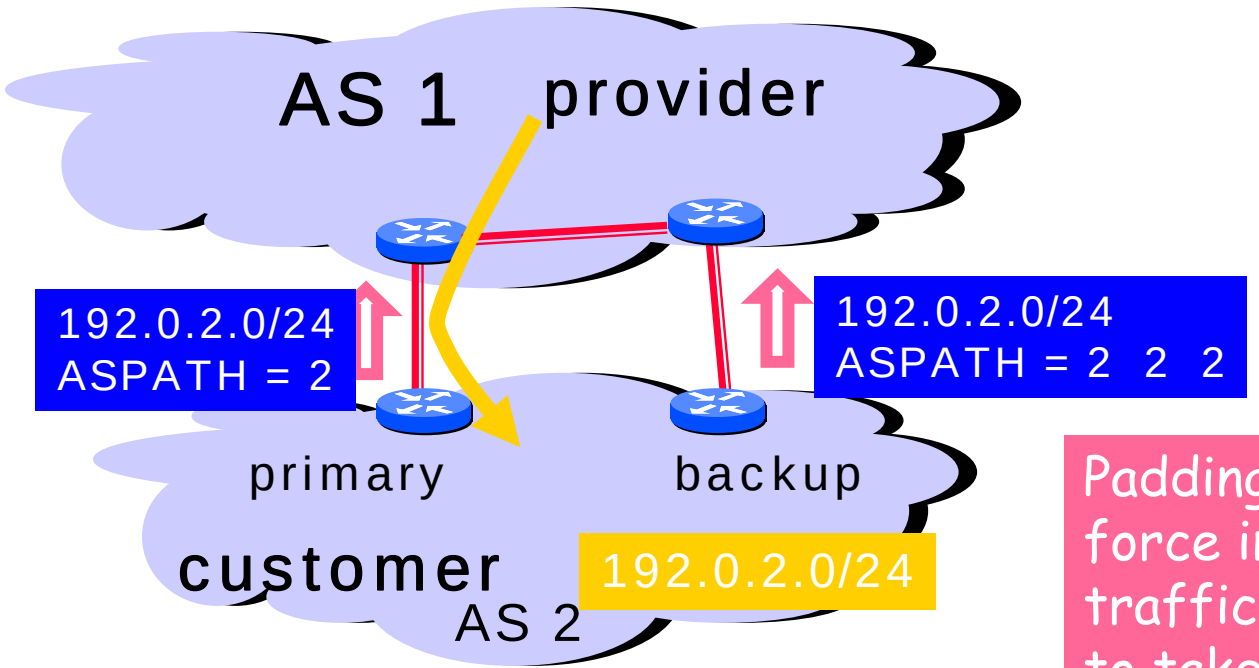

Padding will (usually) force inbound traffic from AS 1 to take primary link

Padding May Not Shut Off All Traffic

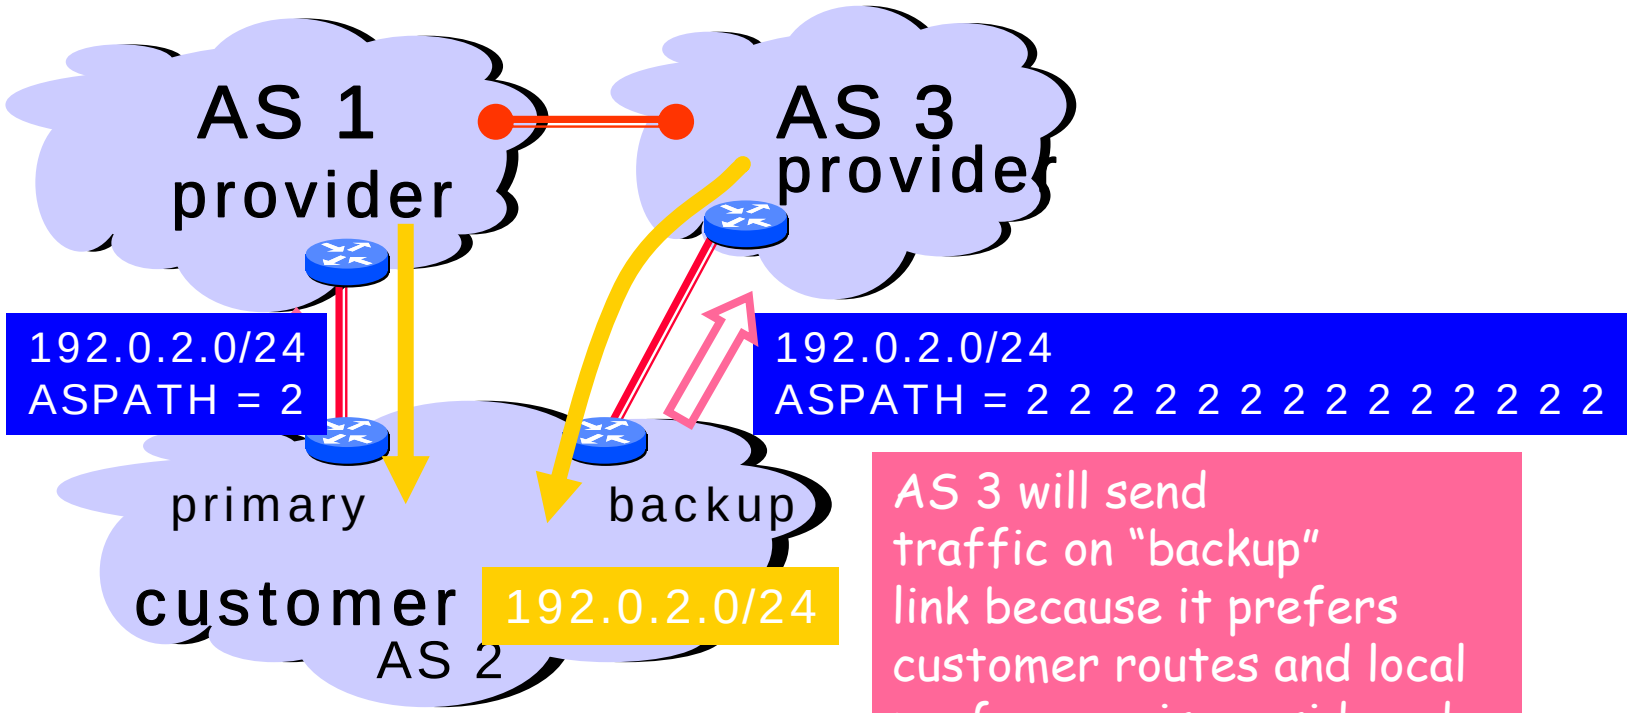

AS 3 will send traffic on "backup" link because it prefers customer routes and local preference is considered before ASPATH length!

Padding in this way is often used as a form of load balancing

Hot Potato Routing: Go for the Closest Egress Point

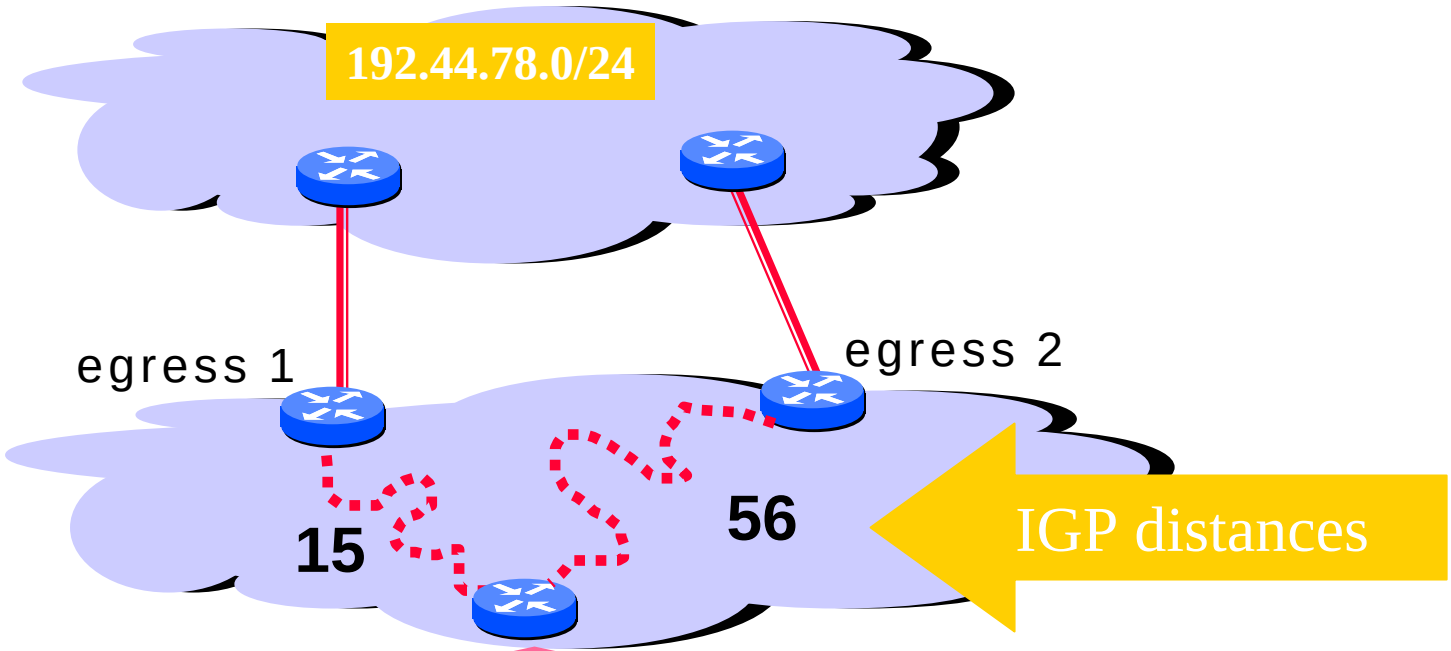

This Router has two BGP routes to 192.44.78.0/24.

Hot potato: get traffic off of your network as Soon as possible. Go for egress 1!

Getting Burned by the Hot Potato

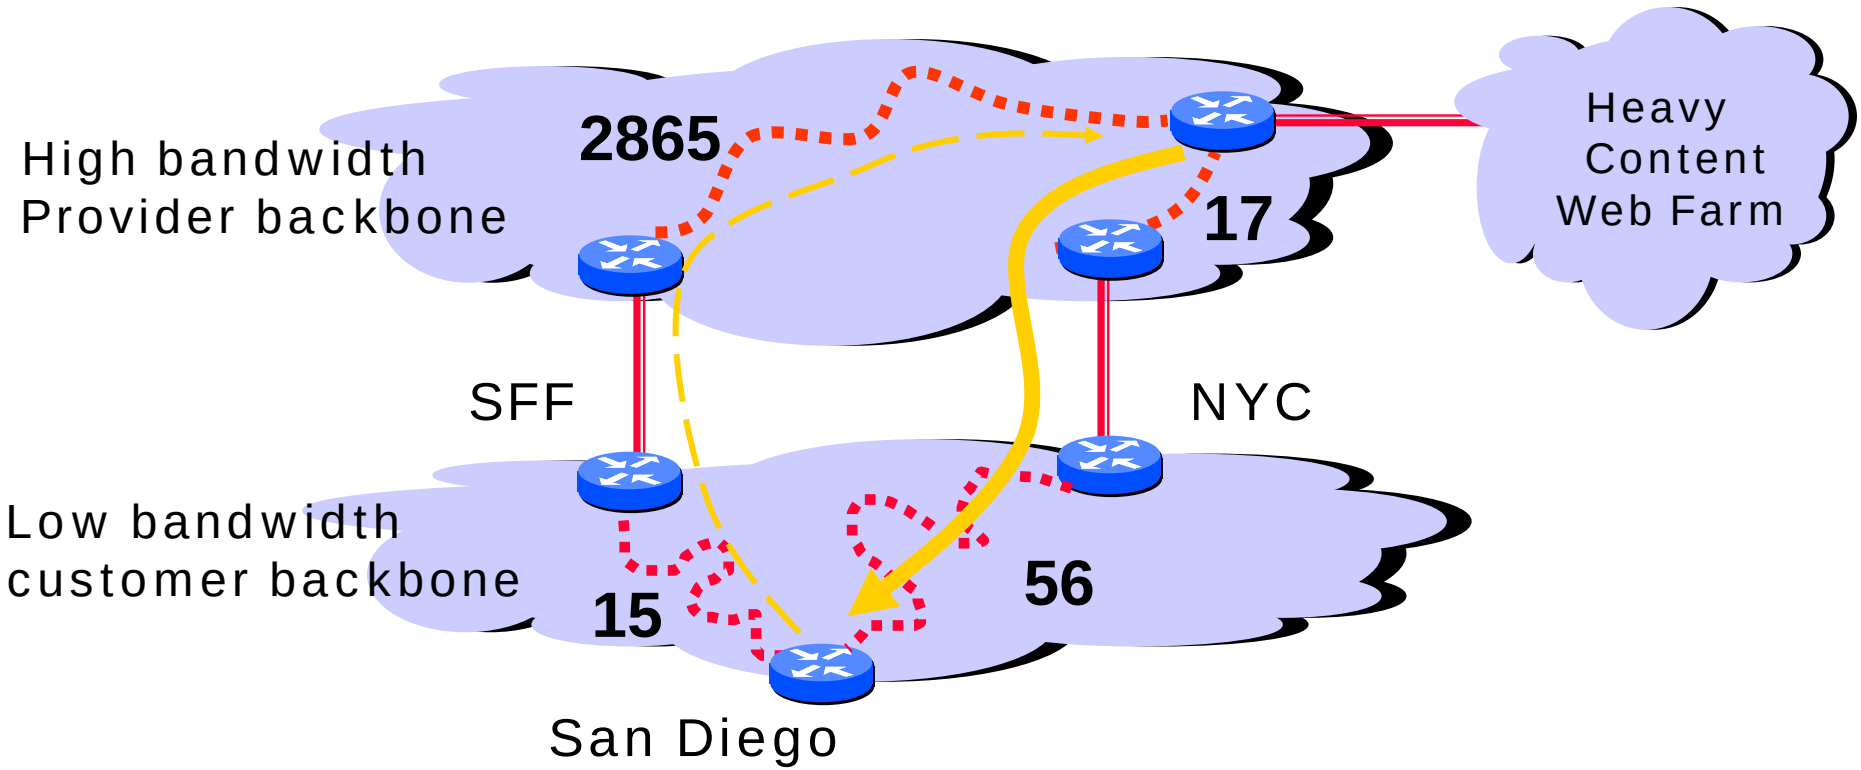

Many customers want their provider to carry the bits!

tiny http request

huge http reply

Route Maps (1)

- Route-maps are cisco's mechanism to select and modify routes with if/then style algorithms.
  - Route-maps are used to match and set attributes of routes.
- They are a little logic flow of ANDs and NOT ANDs.
  - Like a little basic program; evaluated in order of the sequence number.
  - At the end of evaluation, if a route has been permitted at some point, it passed.
- A route-map is ADDITIVE to other filters
- Route-maps follow this format:

```
route-map <name> <per|deny> <#>
  [match statements]
  [set satements]

 [repeat with unique sequence numbers as needed]
```

Route Maps (2)

- For route-maps with the keyword “permit”,
  - If the prefix being examined passes the match statement, the set commands are executed and the route-map is exited.
  - If the match statement is not passed, the next sequence number is executed.
  - If there are no more sequence numbers, the prefix is filtered/dropped.
- For route-maps with the keyword “deny”,
  - if the prefix being examined passes the match statement, the prefix in question is filtered and no more sequence numbers are executed.
  - If the prefix does not pass the match statements, the next sequence number is executed.

Distribute List

- Per neighbor access list applied to BGP routes
- Inbound or outbound
- Based upon network numbers

Distribute List: Example 1

```
router bgp 3847
neighbor 207.240.8.246 remote-as 8130
neighbor 207.240.8.246 distribute-list 127 in
neighbor 207.240.8.246 distribute-list 101 out

access-list 127 permit ip host 207.19.74.0 host 255.255.255.0
access-list 127 permit ip host 208.198.100.0 host 255.255.252.0
access-list 127 permit ip host 208.204.80.0 host 255.255.252.0
access-list 127 permit ip host 208.212.249.0 host 255.255.255.0
access-list 127 permit ip host 207.240.120.0 host 255.255.255.0
access-list 127 permit ip host 208.220.144.0 host 255.255.248.0
access-list 127 permit ip host 208.225.192.0 host 255.255.240.0
access-list 127 deny   ip any any
! explicit deny if not specified
```

Distribute List: Example 2

```
access-list 10 deny   ip 10.0.0.0 0.255.255.255
access-list 10 deny   ip 127.0.0.0 0.255.255.255
access-list 10 deny   ip 128.0.0.0 0.0.255.255
access-list 10 deny   ip 172.16.0.0 0.15.255.255
access-list 10 deny   ip 191.255.0.0 0.0.255.255
access-list 10 deny   ip 192.0.2.0 0.0.0.255
access-list 10 deny   ip 192.168.0.0 0.0.255.255
access-list 10 deny   ip 223.255.255.0 0.0.0.255
access-list 10 deny   ip 224.0.0.0 31.255.255.255
access-list 10 deny   ip 207.240.0.0 0.0.3.255
access-list 10 permit ip any
```

A sanity filter like this keeps your table neat and prevents you from advertising crud to your peers.

Filter List

- Filter routes both inbound and outbound based on value of AS path attribute.
- Called “as-path” access, or filter, lists.
- Configuration

```
router bgp 3847
neighbor 207.240.10.100 remote-as 2900
neighbor 207.240.10.100 distribute-list 100 in
neighbor 207.240.10.100 distribute-list 101 out
neighbor 207.240.10.100 filter-list 10 in

ip as-path access-list 10 permit ^2900$
ip as-path access-list 10 deny .*
```

Cisco Regular Expressions

Much like standard vi, Perl regular expressions:

. Period matches any single character, including white space.
***** Asterisk matches 0 or more sequences of the pattern.
+ Plus sign matches 1 or more sequences of the pattern.
? Question mark matches 0 or 1 occurrences of the pattern
^ Caret matches the beginning of the input string.
$ Dollar sign matches the end of the input string.
_ Underscore matches a comma (,), left brace ({), right brace (}) left parenthesis, right parenthesis, the beginning or end of the input string, or a space.
[] Square brackets designate a range of single character patterns.
- Hyphen separates the endpoints of a range.

Applying AS Path Filtering

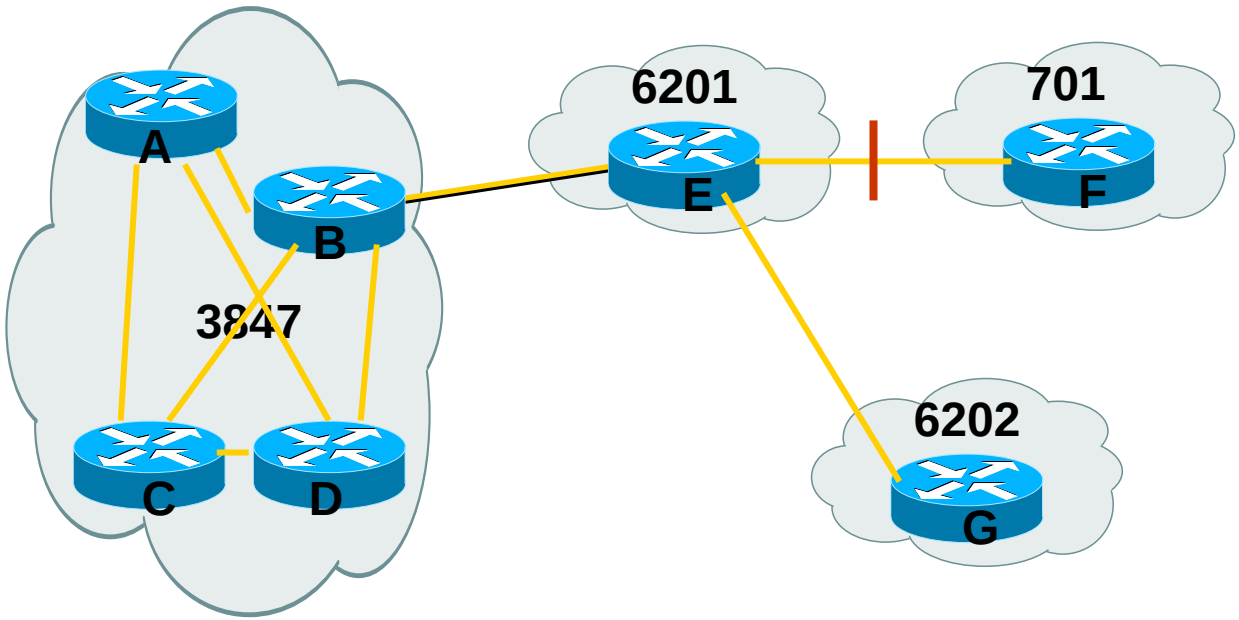

The following configuration could be used on router B to accept routes from AS6201 & 6202 and deny all others.

```
ip as-path access-list 10 permit ^6201$
ip as-path access-list 10 permit ^6201_6202$
ip as-path access-list 10 deny .*
```

AS-Path Filtering: Default Access Control Lists

- 3 default lists
- (Permit all; Deny all; Permit only our routes)

```
ip as-path access-list 1 permit .*
ip as-path access-list 2 deny .*
ip as-path access-list 3 permit ^$
```

Tuning Inbound BGP Routes: Mechanisms

Once you identify better paths,

- use Local-Pref
- use AS_PATH padding

- Identify the providers in question.
- Pick out the relevant AS_PATH regexp.
- Build a route-map to apply inbound

Tuning Inbound BGP Routes: Exp

- Simple route-map

```
ip as acc 20 permit ^701 1673_
route-map inbound-uu permit 10
   match as 20
   set as pre 701 701
route-map inbound-uu permit 20
   match as 1
```

- Always best to leave a specific match all at the end.

Controlling Outbound BGP Routes: Mechanisms

- Your two main tools are:
  - Padding your outbound AS_PATHs
  - De-aggregating announcements
- And:
  - With a cooperative provider, using communities
  - MED, but can be ignored

Tuning Outbound - Padding

- When your router announces iBGP routes, it normally creates a 1-entry AS_PATH with your ASN. So, by adding one or more copy of your own ASN, you cause the providers who listen to that route to de-prefer it a bit (since the AS_PATH is now 1 longer, thus making it win less often).
- Example:

```
route-map pad-me-once
  match as 1
  set as prepend 22222

router bgp 22222
  neigh 207.106.2.45 route-map pad-me-once out
```

Tuning Outbound - Communities

- If your providers are good, they'll give you the ability to control your destiny with communities
- Implementing communities: example

```
route-map set-transit
  match ip address 40
  set comm 4969:1200 4969:666 additive

router bgp 22222
  neigh <custip> route-map set-transit in
```

Implementing Communities (cont'd)

```
ip comm 4 permit 4969:123
ip comm 4 permit 4959:1200

ip comm 20 permit 4969:0
ip comm 21 permit 1239:1
ip comm 22 permit 1239:2

route-map tosprint deny 20
   match comm 20
route-map tosprint permit 21
   match comm 21
   set as pre 4969
route-map tosprint permit 22
   match comm 22
   set as pre 4969 4969
route-map tosprint permit 30
   match comm 4
```

Tuning Outbound – De-Aggregation

- I have 207.106.128.0/17.
  I want to advertise 207.106.128.0/17 to Sprint and UUNET, and 207.106.128.0/18 to Sprint alone.

  access 56 deny 207.106.128.0 0.0.63.255

  access 56 <insert lines from access 55>

  neigh <uunetip> dist 56 out

Making BGP Scalable

- Address and route aggregation
- iBGP fully meshed, not scalable for large AS
- Two mechanisms:
  - BGP route reflector (RR)
    - *Client:* used to identify "client" of the RR(s).
    - *Non-client:* identifies standard BGP peers.
    - *Cluster:* a group of clients under same RR(s).
    - *Cluster-id:* unique identifier for a cluster.
    - *Originator-id:* router-id of the originator of the route.
  - BGP confederation, e.g.,
    - Fully-mesh all BGP speakers at a POP
    - Use fake ASNs at each POP
    - Between POPs, use eBGP rules (send everything)
    - Within POPs, use iBGP rules
    - Preserve local_prefs between POPs

BGP Route Reflector: Illustration

RR router defines:
neighbor <ip-addr> route reflector-client

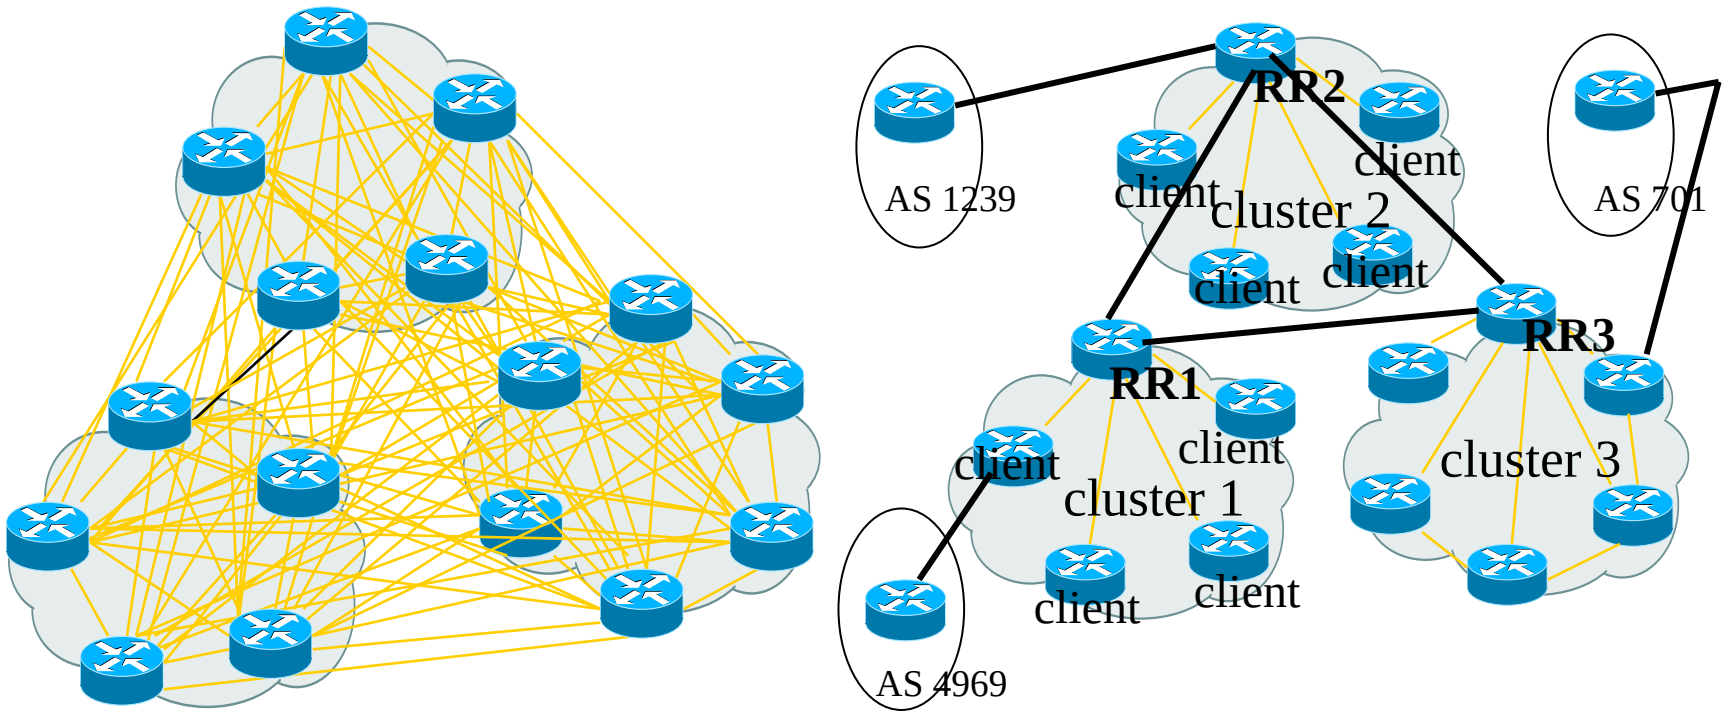

BGP Confederation: Illustration

‘bgp confed identifier 4969’

‘bgp confed peers 64512 64513 64514 64515’

put in extra confed peers up-front

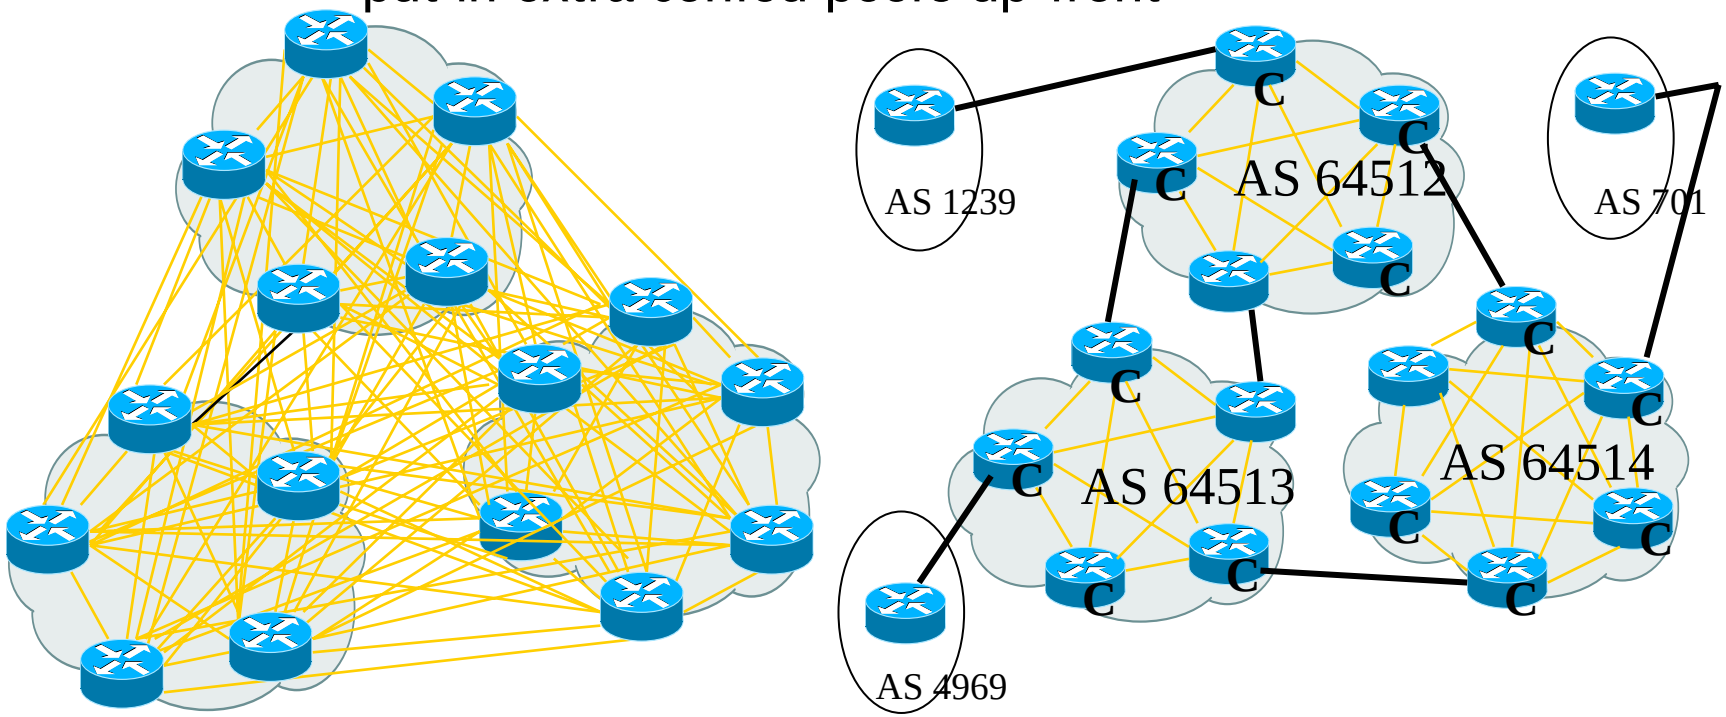

Making BGP Stable

Some Mechanisms:

- Nail routes to loopback
  - Peering between loopbacks enhances stability, since loopbacks don't go down.
  - Also, good for load-balancing

    set up lo0, then

    "neigh x.y.z.q update-source looback0"
- Cisco soft reconfiguration and route refresh
- Route dampening: watch out for flapping routes
  - if a route shows instability, it may be "blackholed" for some time (30-90 minutes) until it stabilizes.