The Internet Ecosystem and Evolution

Contents

- AS-AS business relationships
- Transit
  - the customer cone
  - transit pricing
- Peering:
  - definition and notation
  - peering policies, IXPs
- The AS-level hierarchy of the Internet

AS-AS business relationships

Autonomous Systems

- **Autonomous System (AS):** a set of hosts, routers, and networks that **an organization owns and administers as a unit** and **exposes a unified routing policy** to the Internet

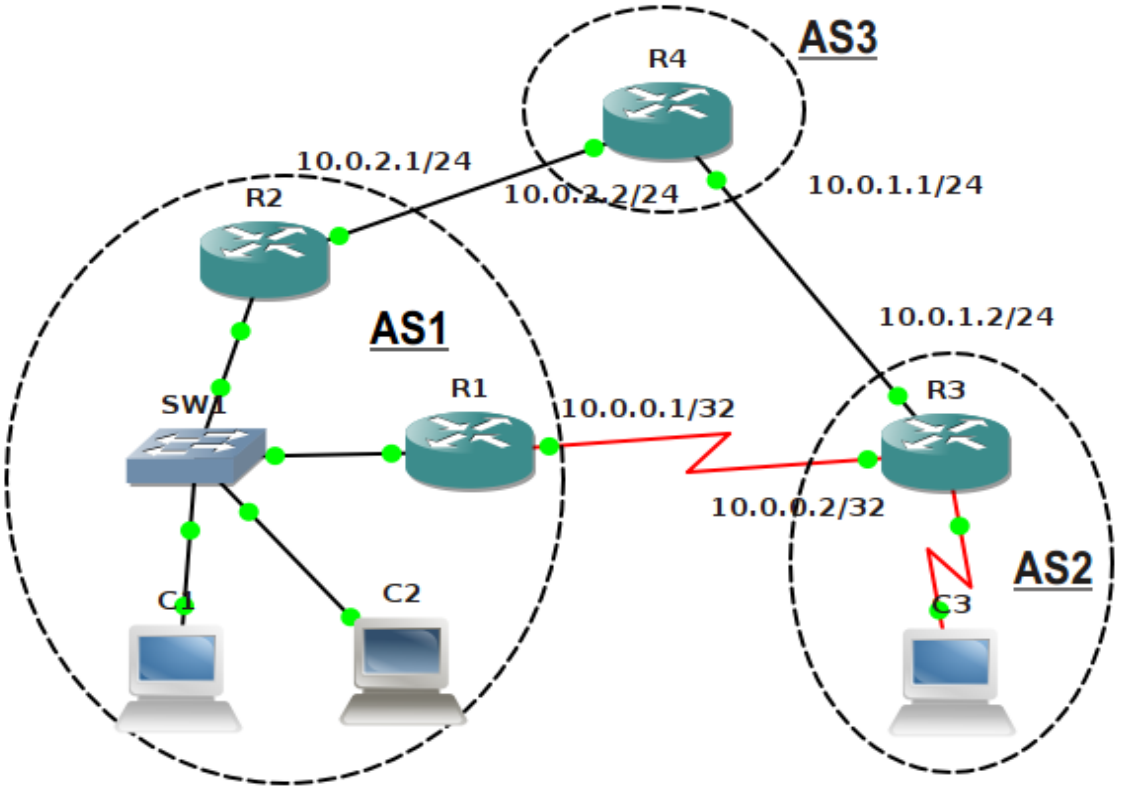

Internet inter-domain routing

- **Service model:** how an AS decides whether to forward or block other ASes' traffic
- Reflects the business/political/security/etc. interests of ASes: **policy routing**

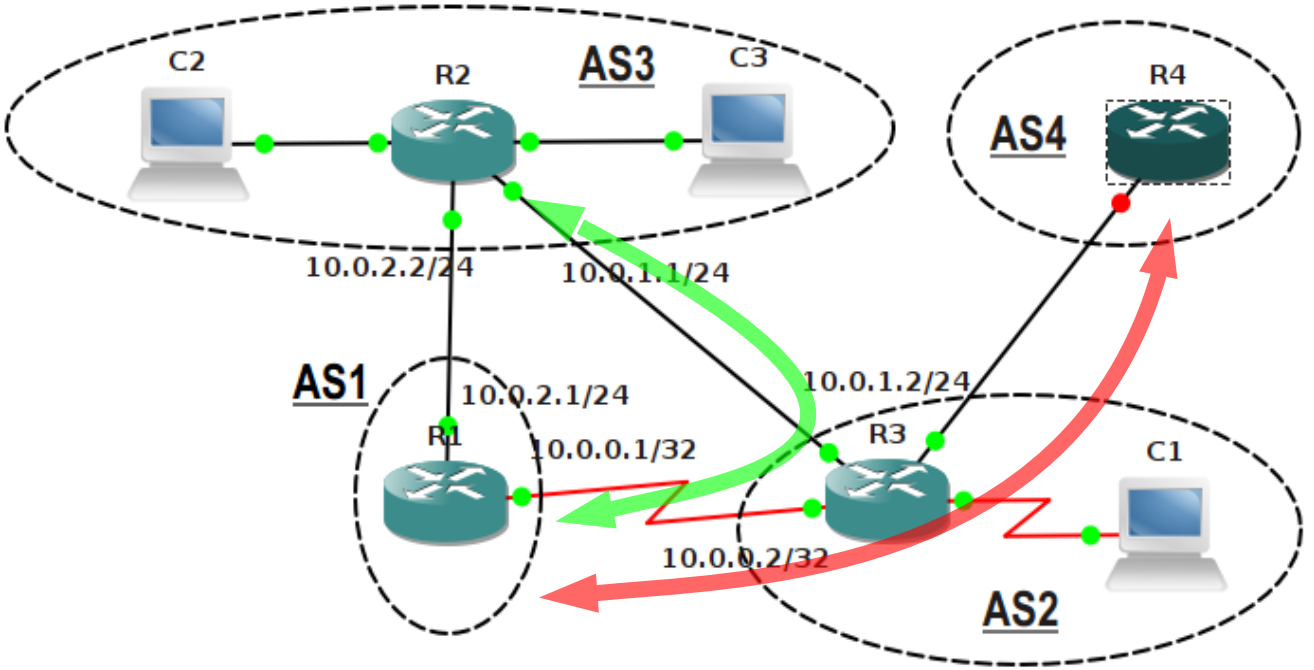

The transit service

- A **customer** AS contracts a service **provider** AS to deliver its egress traffic to any host connected to the Internet and ingress traffic from the Internet back to the customer
- The customer is **charged** by the traffic rate

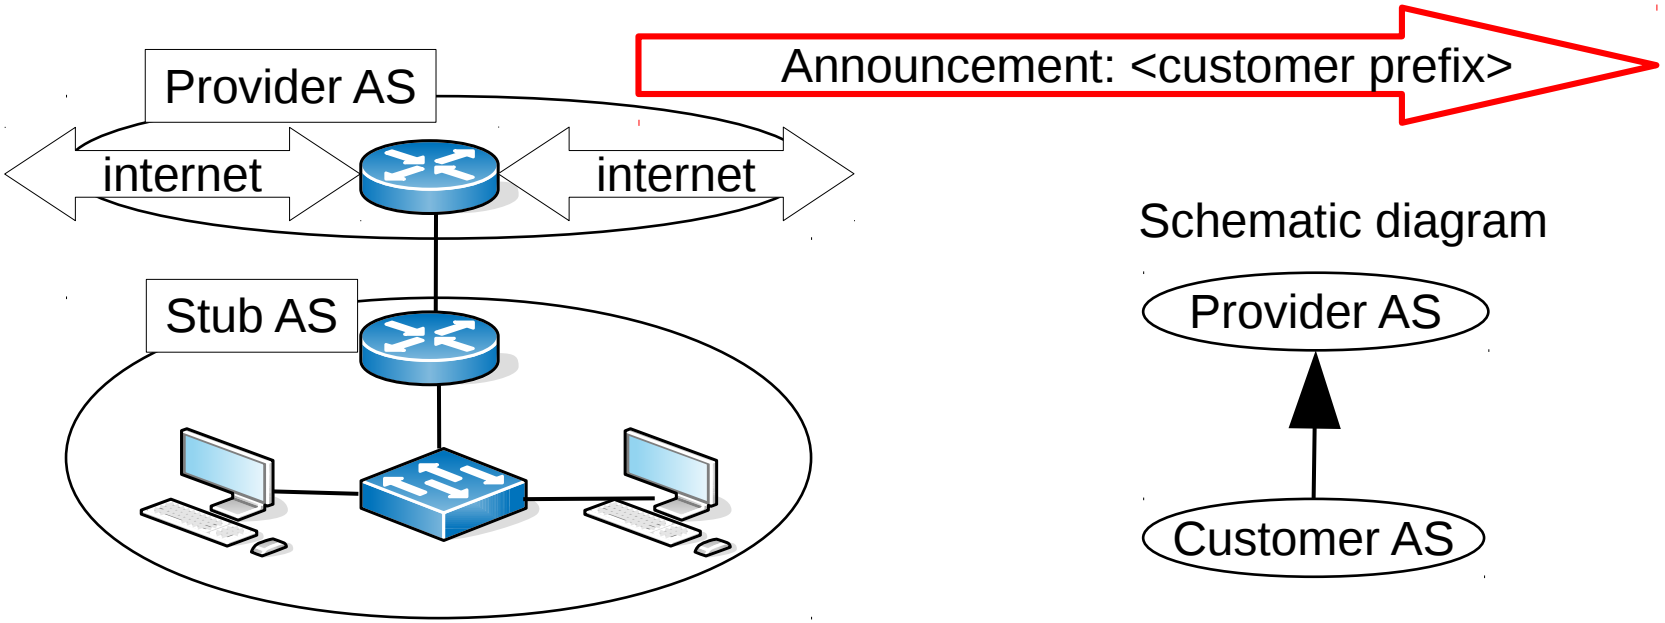

Size of ASes

- The importance of an AS is proportional to the number of customer ASes to which the AS provides transit service (**customer cone**)
- A more characteristic metric may be **the size of the address range announced by an AS** into the inter-domain routing system
  - an AS must announce each prefix of each of its customers to provide access to them
  - so the extent of the address range announced (equivalent /8s) will be a good size metric
- This is called the **size of an AS**

Transit pricing

- The quality of the transit service is specified in the **SLA (Service Level Agreement)**: a formal contract between the customer and the provider
  - availability: for instance, the transit service is available in 99.99% of the time
  - rate, packet loss, delay, and delay variation
  - monitoring: the signers agree how and where to measure the quality of the service
- **Progressive pricing** according to the measured traffic rate on the provider–customer AS link

Transit pricing

- **Commit:** customer's expected traffic rate

Commit rate	Price per unit	Minimum price
10 Mbps	12$ / Mbps	120 $/month
100 Mbps	5$ / Mbps	500 $/month
1 Gbps	3.5$ / Mbps	3 500 $/month
10 Gbps	1.2$ / Mbps	12 000 $/month
100 Gbps	0.7$ / Mbps	70 000 $/month

- E.g., choosing commit „10 Mbps/12$” means
  - customer pays 12 USD per Mbps traffic unit
  - but customer pays the minimum price (120$) even if the rate remains below 10 Mbps
- The larger the commit the smaller the price per unit!

Measuring the transit rate

- Internet traffic varies on a wide scale on an hourly, daily, and weekly basis

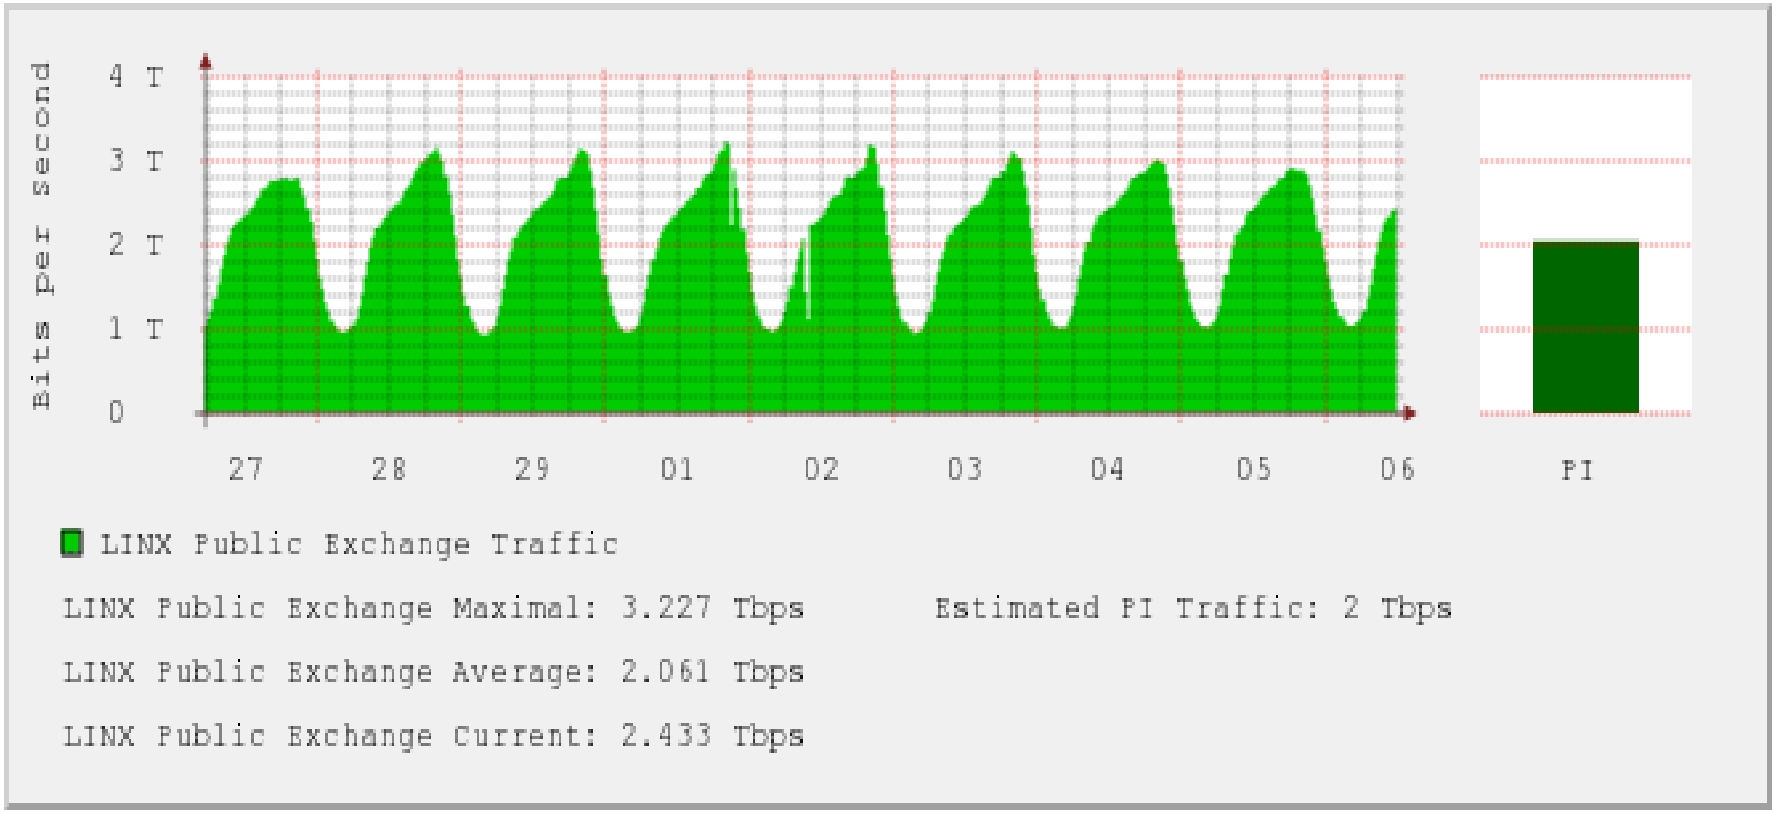

https://www.linx.net/tech-info-help/traffic-stats, 2016. 03. 06.

Measuring the transit rate

- **95th percentile** (**95/5**) methodology
  - sample traffic rate on the customer-provider AS-AS link every 5 minutes
  - at the end of the measuring period, sort measured samples in ascending order
  - drop erroneous samples (overflow, reboot)
  - calculate the 95th percentile: the smallest sample that is larger than, or equal to, the 95% of the samples
- In other words, the customer's rate was below the 95th percentile in 95% of the time

Transit pricing: Example

- 95th percentile for the sequence `[1,...,1, 100]` of 100 samples is 1 Mbps (the 95-th entry in the list sorted in ascending order)
- But the sequence `[1,...,1,15,16,16,17, 21,40]` (again of length 100) the 95th percentile is 15 Mbps
- For this, the transit price at commit 10Mbps is

12 $/Mbps * 15 Mbps = 180$

- But at commit 100 Mbps the price is 500$, as the minimum price must be payed:

commit rate (100$) * unit price (5$/Mbps)=500$

A peer AS-AS relationship

The peer AS-AS relationship

- Suppose that both `AS1` and `AS2` are customers of `AS3` and they are of roughly the same size
- It's cheaper for `AS1` and `AS2` to exchange traffic directly, eliminating `AS3` from the loop
- They enter into a so called settlement-free **peering** relationship

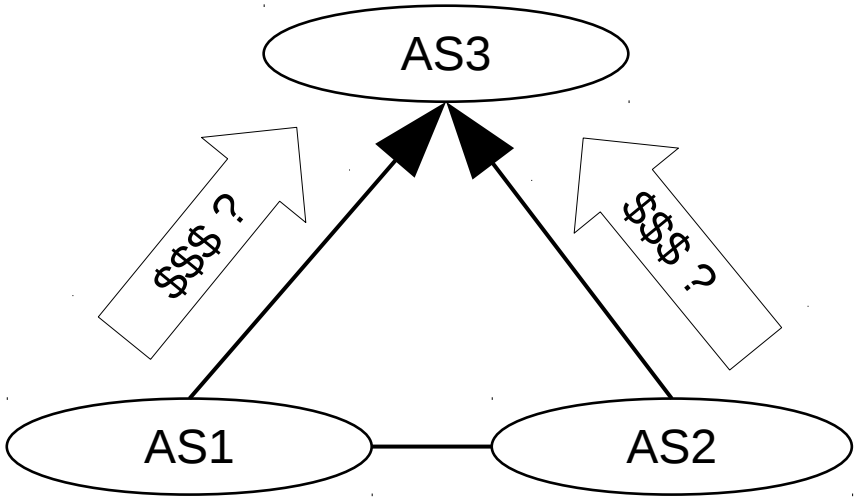

- The peer relationship is marked by an **arrowless edge**: no cash-flow!

Schematic diagram

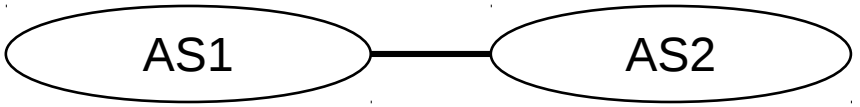

The peer AS-AS relationship

Internet Peering is the business relationship whereby companies reciprocally provide access to each others' customers.

(DrPeering)

The peer AS-AS relationship

- **Bilateral agreement** between two ASes to forward traffic *(i)* between each other and *(ii)* between all their customers

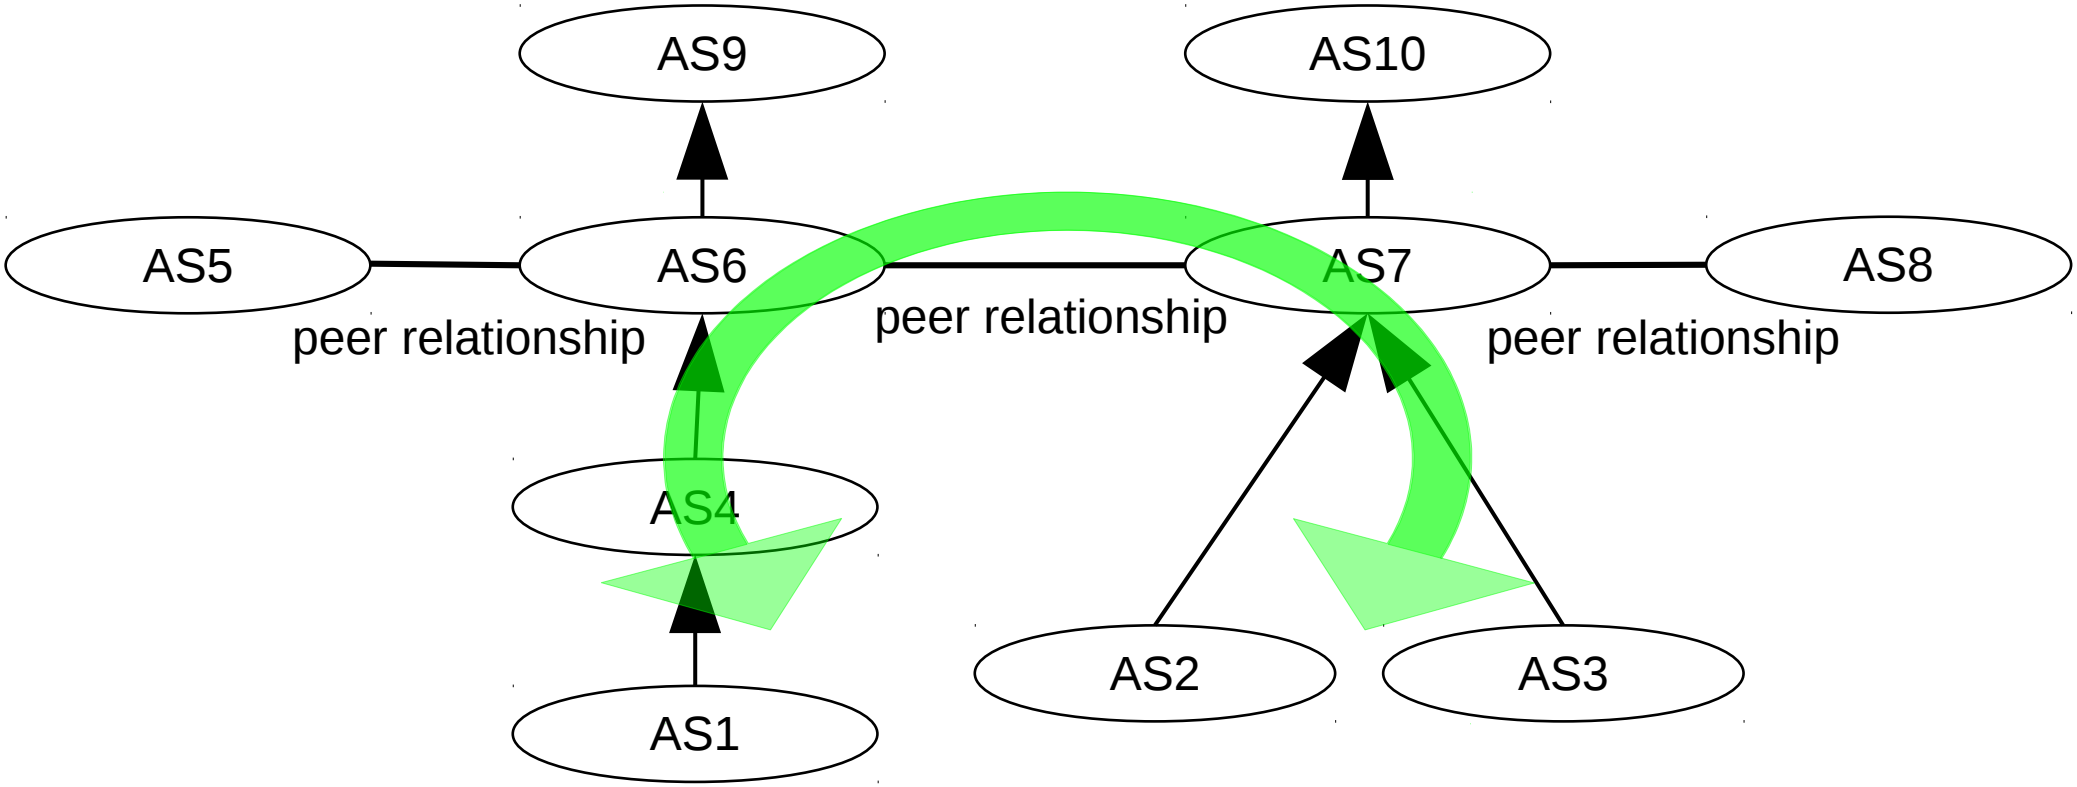

Peering: Considerations

- A peer link essentially comes for free
- Of course, operating the direct link costs money
- Still, not all ASes form a peer link between one another
- A larger AS will typically not peer with a smaller one
  - since this would allow access to its customers
  - instead, it could rather charge the other AS by providing transit instead of a peering
  - so such a peering link would bring profit loss

Peering policy

- Every AS sets the conditions under which it would go into a peering relationship
- **Selective peering:** strict peering policy
  - customer cones of roughly the same size
  - symmetric traffic demands
  - multiple, geographically diverse POPs
  - 24x7 support
- **Open peering:** may peer with any AS
- See also: PeeringDB

Peering: selective vs. open

Facebook Gold Sponsor

Organization	Facebook
Also Known As	Facebook, Instagram, WhatsApp
Company Website	https://www.facebook.com/
Primary ASN	32934
IRR Record	AS-FACEBOOK
Route Server URL	
Looking Glass URL	
Network Type	Content
IPv4 Prefixes	100
IPv6 Prefixes	100
Traffic Levels	1 Tbps+
Traffic Ratios	Heavy Outbound
Geographic Scope	Global
Protocols Supported	Unicast IPv4 ○ Multicast ✓ IPv6
Last Updated	2016-07-06T00:29:07Z
Notes	We have a selective peering policy requiring a minimum of 50 Mbps of in-continent traffic destined to or through your network. We welcome the opportunity to engage in peering with responsible BGP speakers in an effort to improve the experience of our millions of users throughout the globe. We require an up-to-date peeringdb entry for all public peering requests, including exchange information with properly formatted public fabric addresses, asns, and noc/peering contact information. We ask that peers also maintain their private peering facilities, as we use this information for private peering (PNI) targeting.

Peering Policy Information

Peering Policy	https://www.facebook.com/peering/
General Policy	Selective
Multiple Locations	Not Required
Ratio Requirement	No
Contract Requirement	Not Required

Google Inc. Platinum Sponsor

Organization	Google Inc.
Also Known As	Google, YouTube (for Google Fiber see AS16591 record)
Company Website	https://www.google.com/
Primary ASN	15169
IRR Record	AS-GOOGLE
Route Server URL	
Looking Glass URL	
Network Type	Content
IPv4 Prefixes	15000
IPv6 Prefixes	750
Traffic Levels	Not Disclosed
Traffic Ratios	Mostly Outbound
Geographic Scope	Global
Protocols Supported	✓ Unicast IPv4 ○ Multicast ✓ IPv6
Last Updated	2017-01-05T14:08:35Z
Notes	Peering Operational Issues: Contact noc@google.com 24x7 Peering Requests: https://isp.google.com/iwantpeering We have a generally open peering policy. Please visit the following link: https://peering.google.com/#/options/peering This link also has information about our traffic delivery and management practices. Please note, not all Google content and services may be available at each PoP or Exchange. Related ASNs Google also manages the following ASNs: AS36040, AS43515, AS36561

Peering Policy Information

Peering Policy	https://peering.google.com/#/options/peering
General Policy	Open
Multiple Locations	Preferred
Ratio Requirement	No
Contract Requirement	Not Required

Peering wars

- Peering wars: to peer or not to peer with an AS

Peering	No peering
may decrease transit costs (no need to pay transit price to traffic that flows on the peer link)	potential of profit loss (compared to if the other AS would be charged for a transit service)
may decrease latency (the peer link allows direct traffic exchange, saving the round-trip to the transit)	may also increase latency (transit may provide better network: faster connectivity)

- By far the most contentious policy tussle in today's Internet

IXP

- Most peer relationships are created at IXPs
- **Internet eXchange Point (IX/IXP):** special network infrastructure dedicated to allow ASes to enter into peering relationships easily
- Typically a well-connected data center where ISPs can co-locate their POPs
- The IXP ensures that any two member ASes can connect via their POPs
- Mass effect: if an ISP appears at an IXP, then it's very cheap to establish new peering relationships

IXP

- An IXP can be for-profit (USA) or non-profit (Europe)
- **An IXP's goal is to attract as many ISPs as possible into its data center**
- Member ISPs are charged on a per-port basis

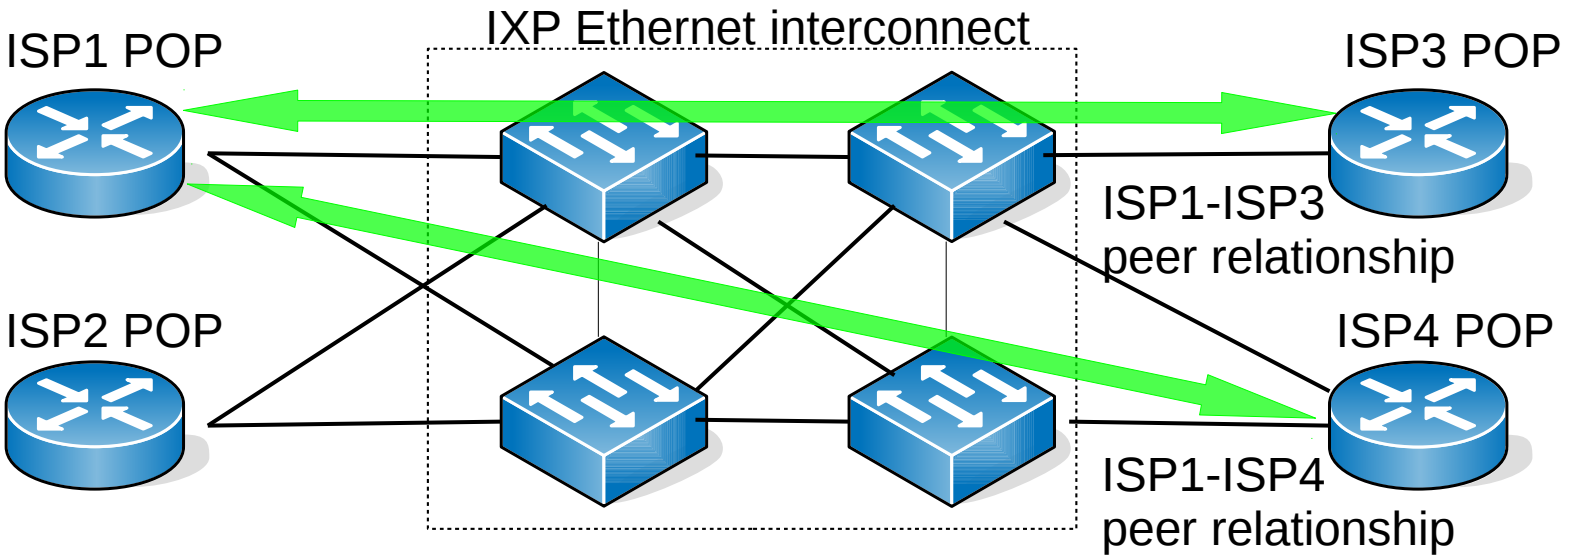

IXP

- An IXP may connect hundreds of ASes, the transfer rate can match that of largest ISPs
  - DE-CIX (Frankfurt, Hamburg, Munich): 600+ ISP, 5 Tbps average rate
  - AMS-IX (Amsterdam Internet Exchange), LINX (London Internet Exchange), Equinix
  - BIX: Budapest Internet eXchange

Stub versus transit ASes

- **Stub AS:** does not provide transit to any AS
- Roughly half/two-thirds of the 50 thousand ASes on the Internet today are stubs
- **Tranzit AS:** non-stub AS

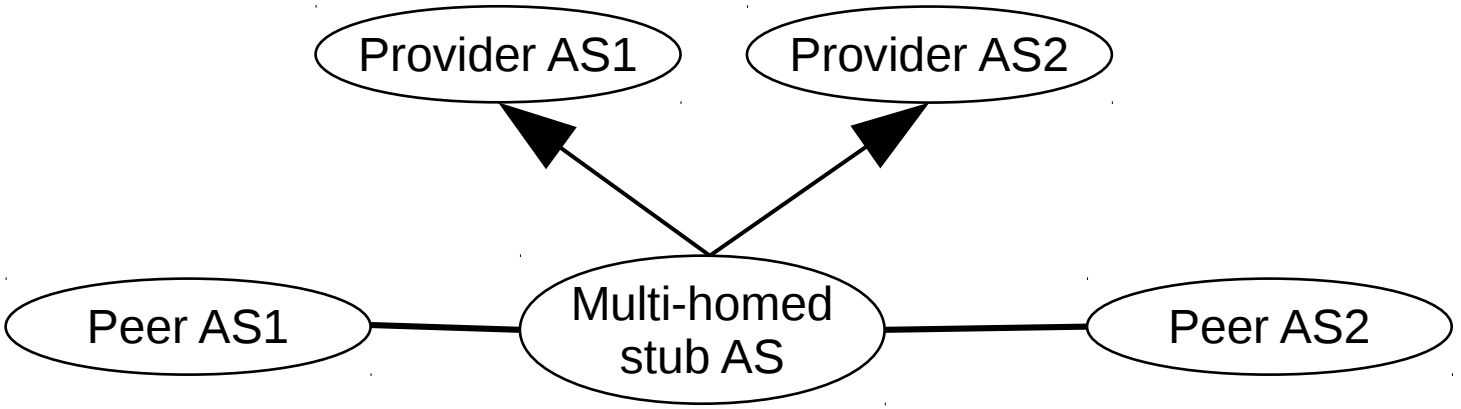

HBONE (AS1955)

- Hungarian academic backbone: education, R&D, libraries, government, etc.

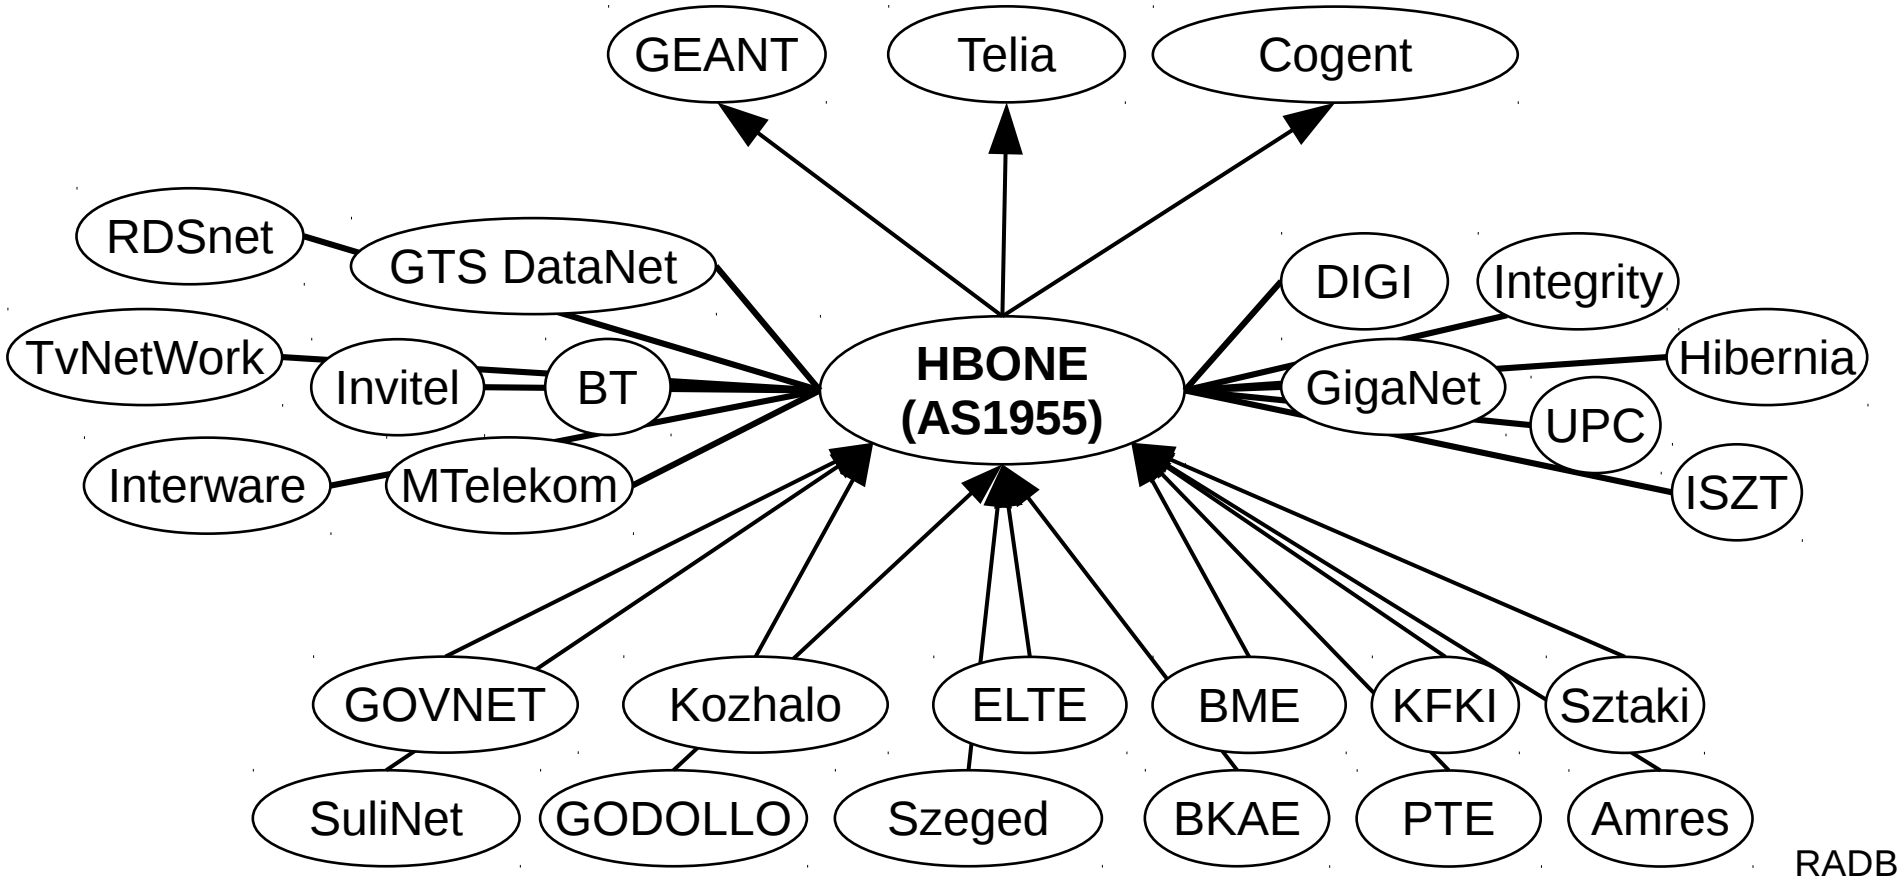

The AS-level structure of the Internet

Transit hierarchy: Tiers

- A customer can also provide transit to other customer ASes, and so on
- **Tier 1:** AS with no **upstream AS** (no provider)
  - **global:** can connect any two ASes regardless of geographic distance
  - Tier1 ASes form a **full peer-mesh** between each other
  - 12 Tier1 ASes: AT&T, CenturyLink, Cogent, GTT, Deutsche Telekom, Level3, NTT, Sprint, Tata, Seabone, TeliaSonera, Verizon, XO

Transit hierarchy: Tiers

- **Tier 2:** customer of a Tier 1 AS (often a global network on its own right)
- **Regional provider:** providing Internet access throughout some geographic region (e.g., Comcast: USA, Orange: Europe)
- **National provider:** country-wide Internet service provisioning
- At the bottom of the transit hierarchy: **single- and multi-homed customers and stub ASes**

Transit hierarchy: Tiers

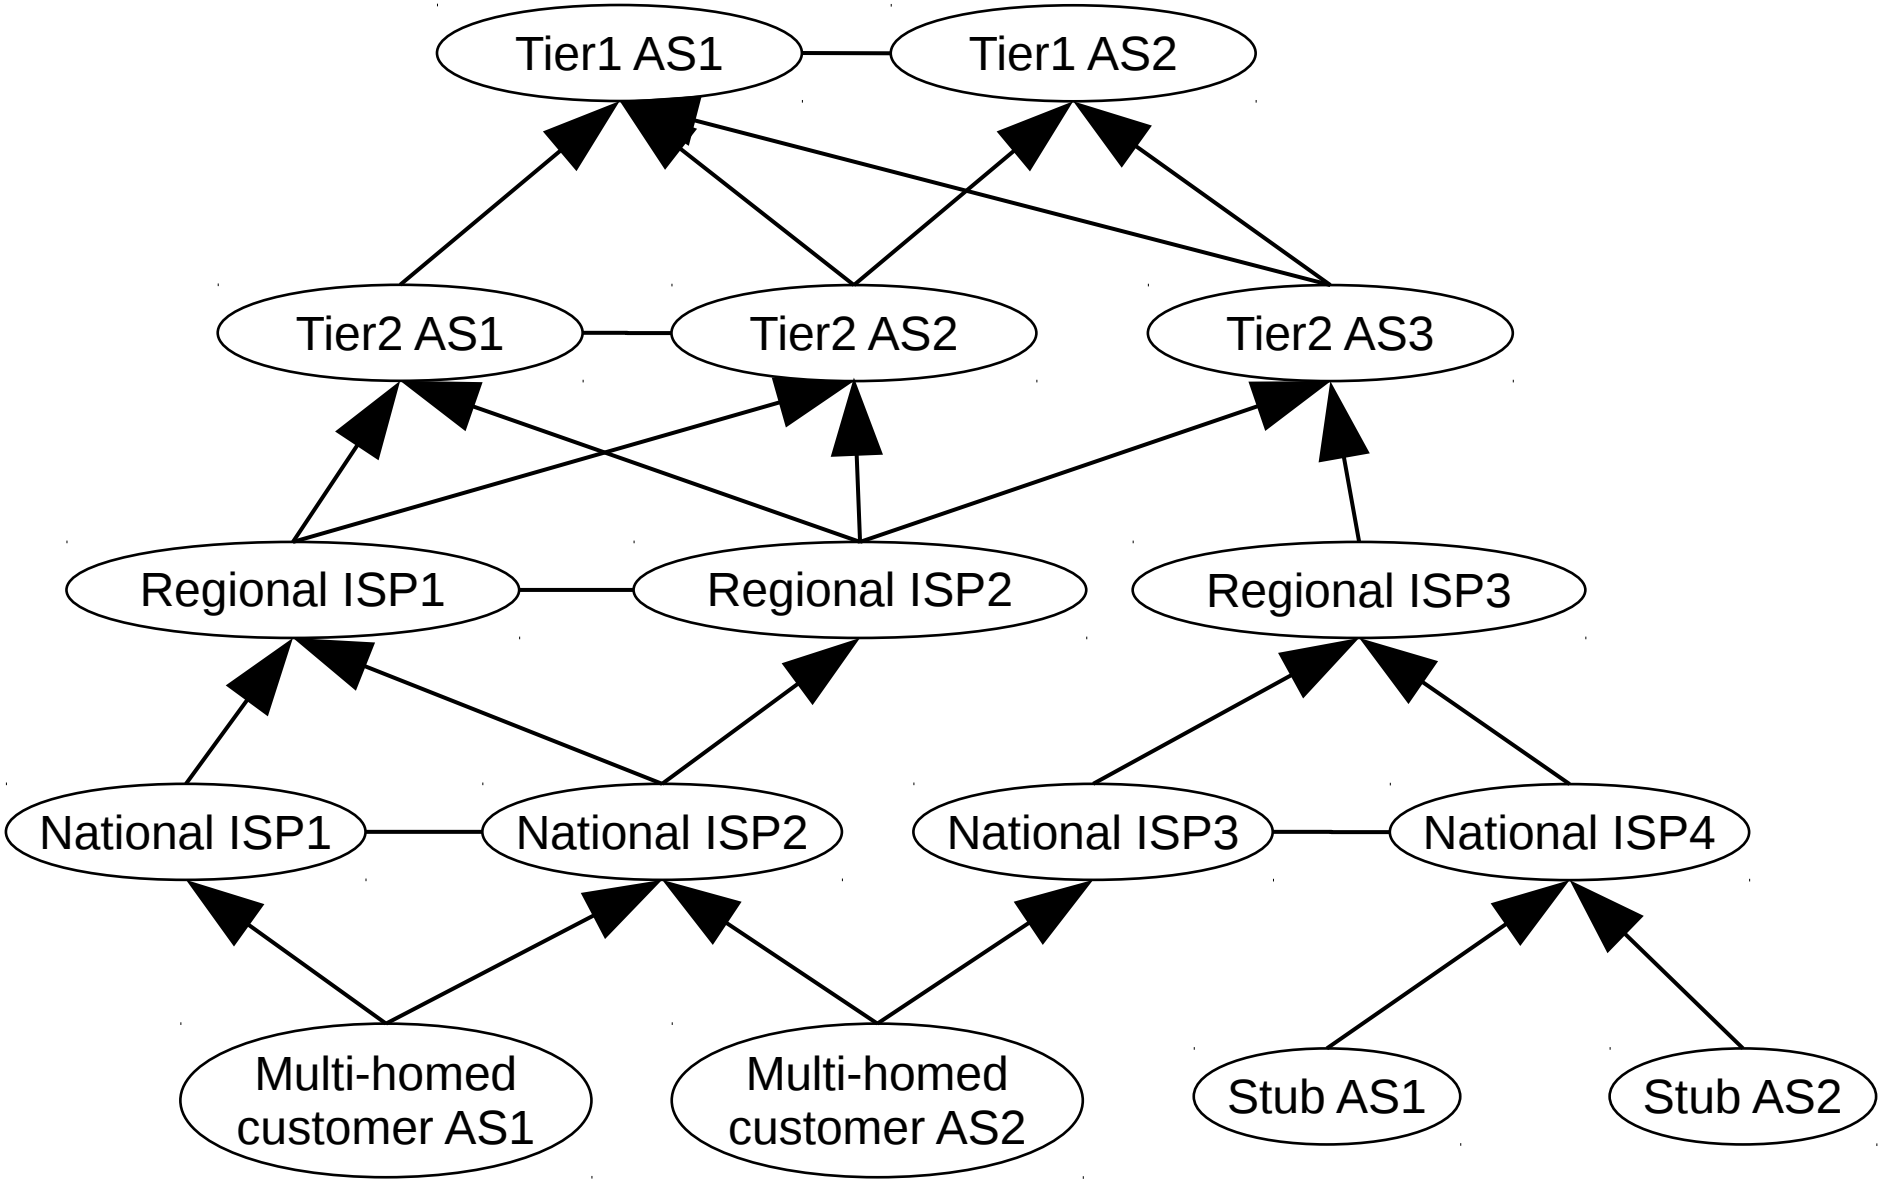

Internet: Terra incognita

- The transit-peer taxonomy covers only about 70% of real AS-AS business relationships
  - **paid peering:** peering at an IXP for a fee
  - **sibling:** mutual transit between two ASes
  - many other unclassified policies
- AS-AS business relationships are secret!
  - knowing an AS's business strategy is a competitive advantage
  - still, some ASes public their relationships: IRR
  - or we can infer from `traceroute` measurements

Internet: Terra incognita

- The tier-classification is only a guess
  - many Tier1 connect directly to national ISPs
  - not only ASes at the same level peer
  - sometimes Tier1s also terminate peering (**depeering**) → peering wars
- **Flattening:** the Internet slowly transitions from the initial strict transit AS hierarchy to a fully decentralized peer full-mesh
- But the precise AS-level hierarchy is not known!!!