Connectivity Fäscht 2026

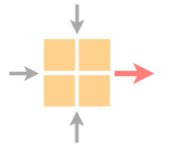

Networked Systems
ETH Zürich — seit 2015

Grab a piece of paper and write your AS name on it!

Try to sit together with regards to your area in the mini-internet

Lukas Röllin
roellinl@ethz.ch

Slides inspired by
Alexander Dietmüller

Please Sign up for a Discussion Slot (oral part of the project)

You can sign up by using the following link:

https://docs.google.com/spreadsheets/d/1ZLmFonBnD5F3TjY_-6l5my3uhW39Fp70PBvBekGnrCE/edit?gid=0#gid=0

Goals for today

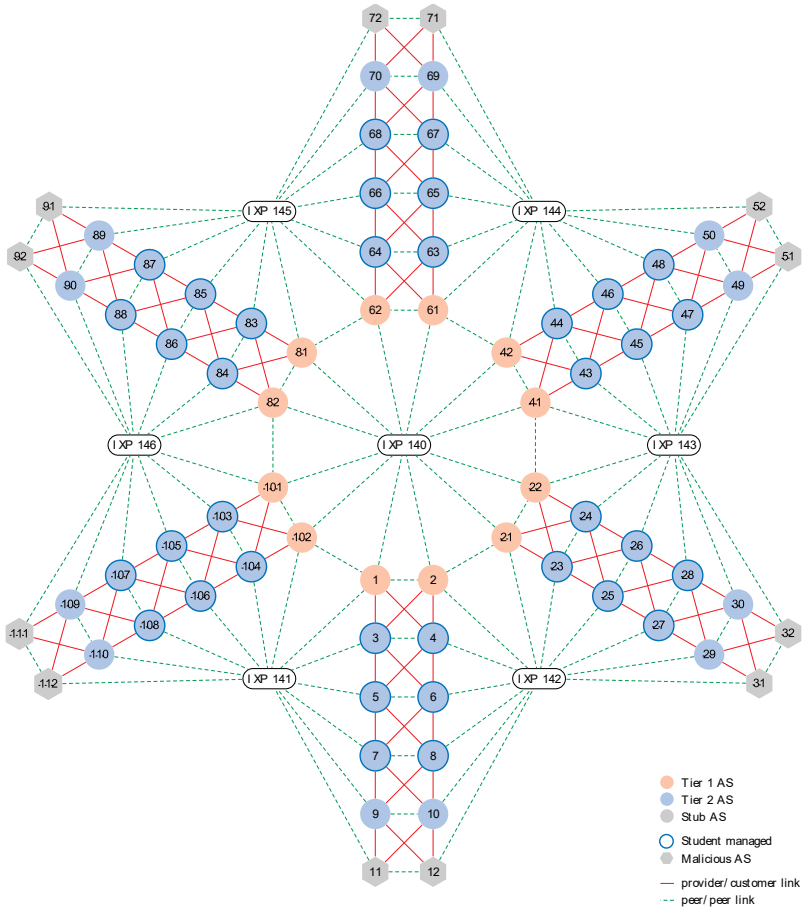

- Connectivity across ASes
- Learn how to debug common problems, especially with BGP
- Have fun and connect with other people from the class

What you should have done already

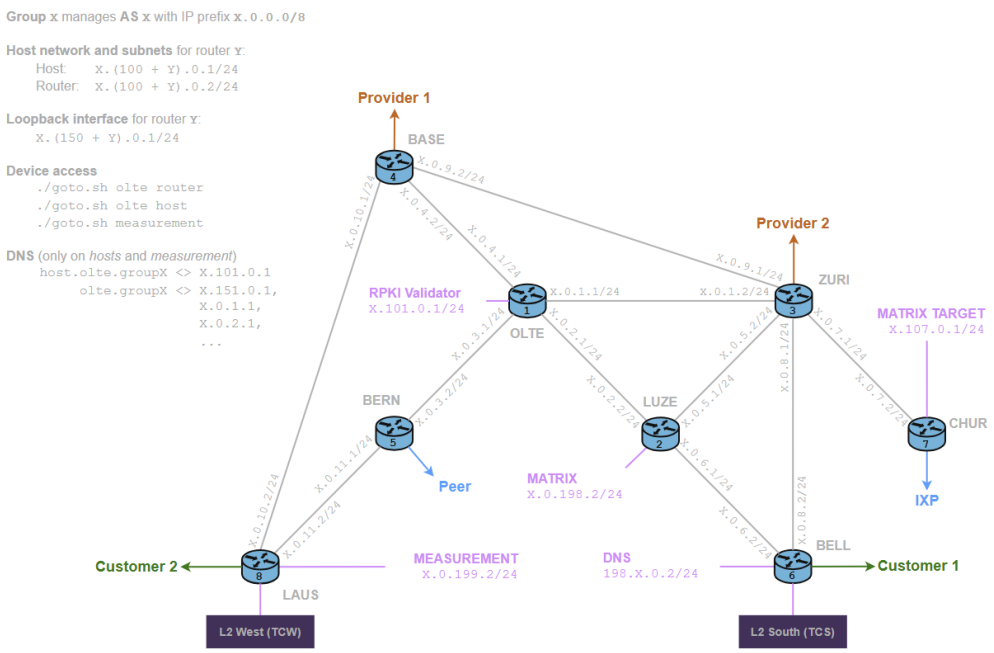

- Connectivity inside your own AS
  - Question 1.2: Routers need IP addresses and OSPF needs to be configured
  - **Test: ping between your routers and hosts**
- Setup full iBGP mesh
  - Question 2.1: All routers need iBGP sessions with each other -> full mesh
  - **Test: show ip bgp summary**

You need a total of 56 iBGP sessions

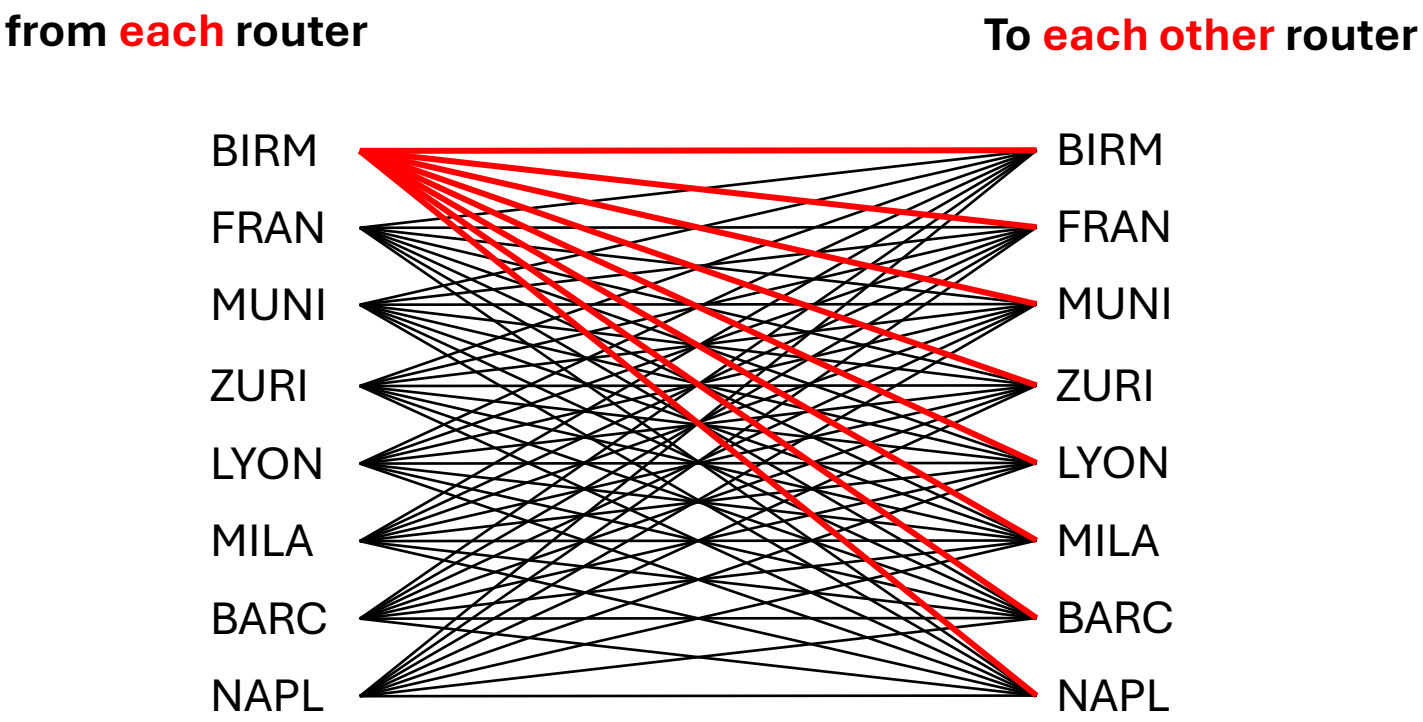

Today

- You will connect your own networks with your neighbors' networks to form the mini-internet
- Internet -> interconnected networks

duvel.ethz.ch/matrix will show the current mini-Internet connectivity

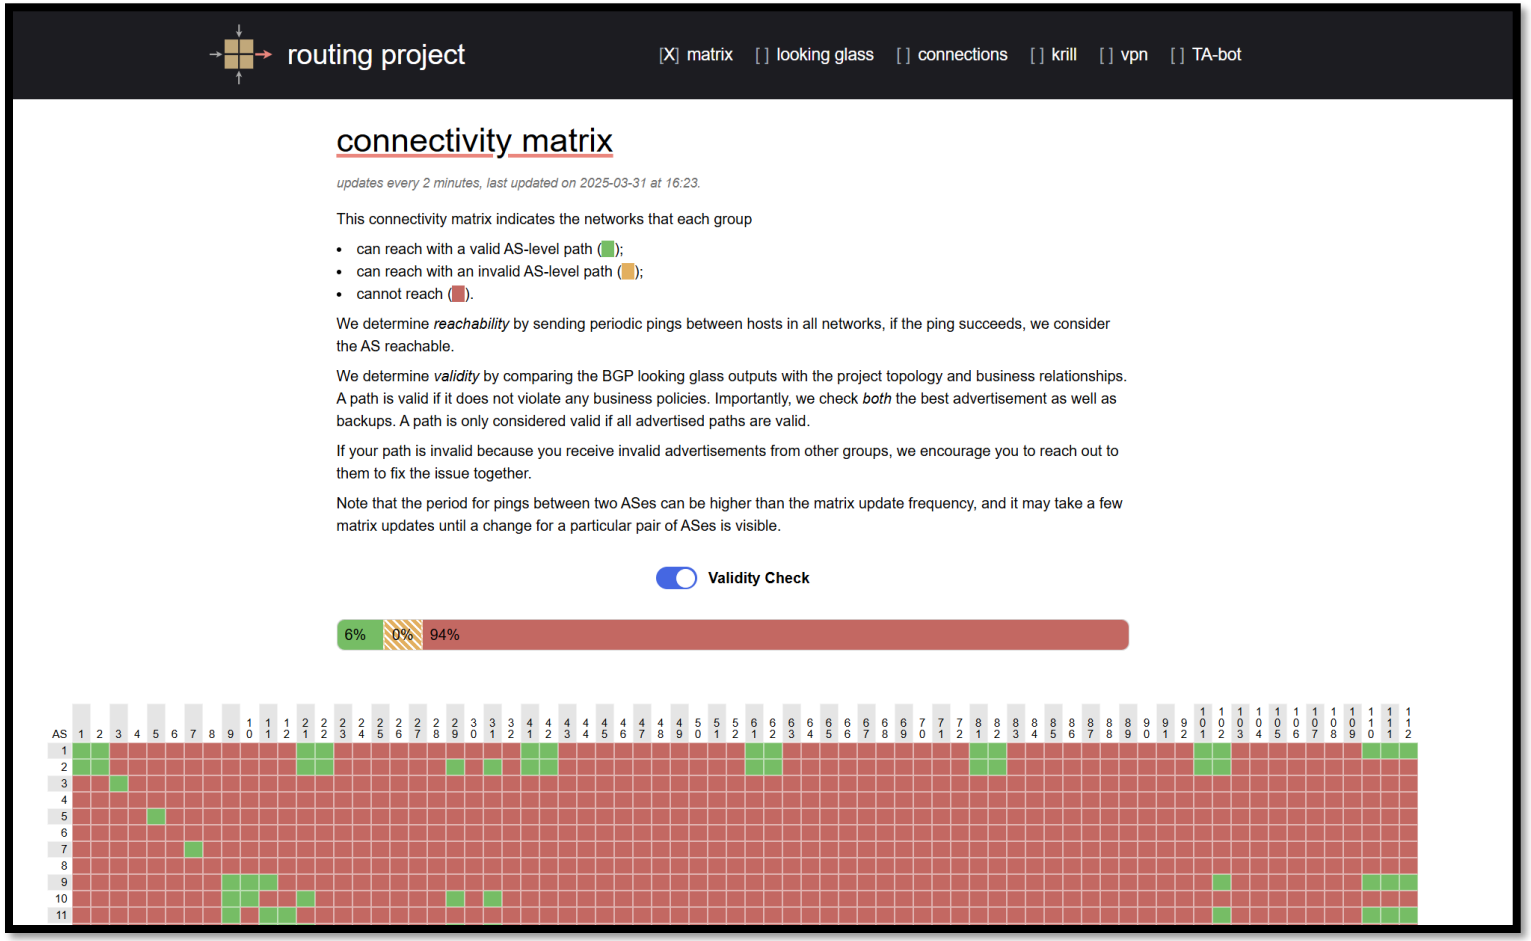

duvel.ethz.ch/matrix will show the current mini-Internet connectivity

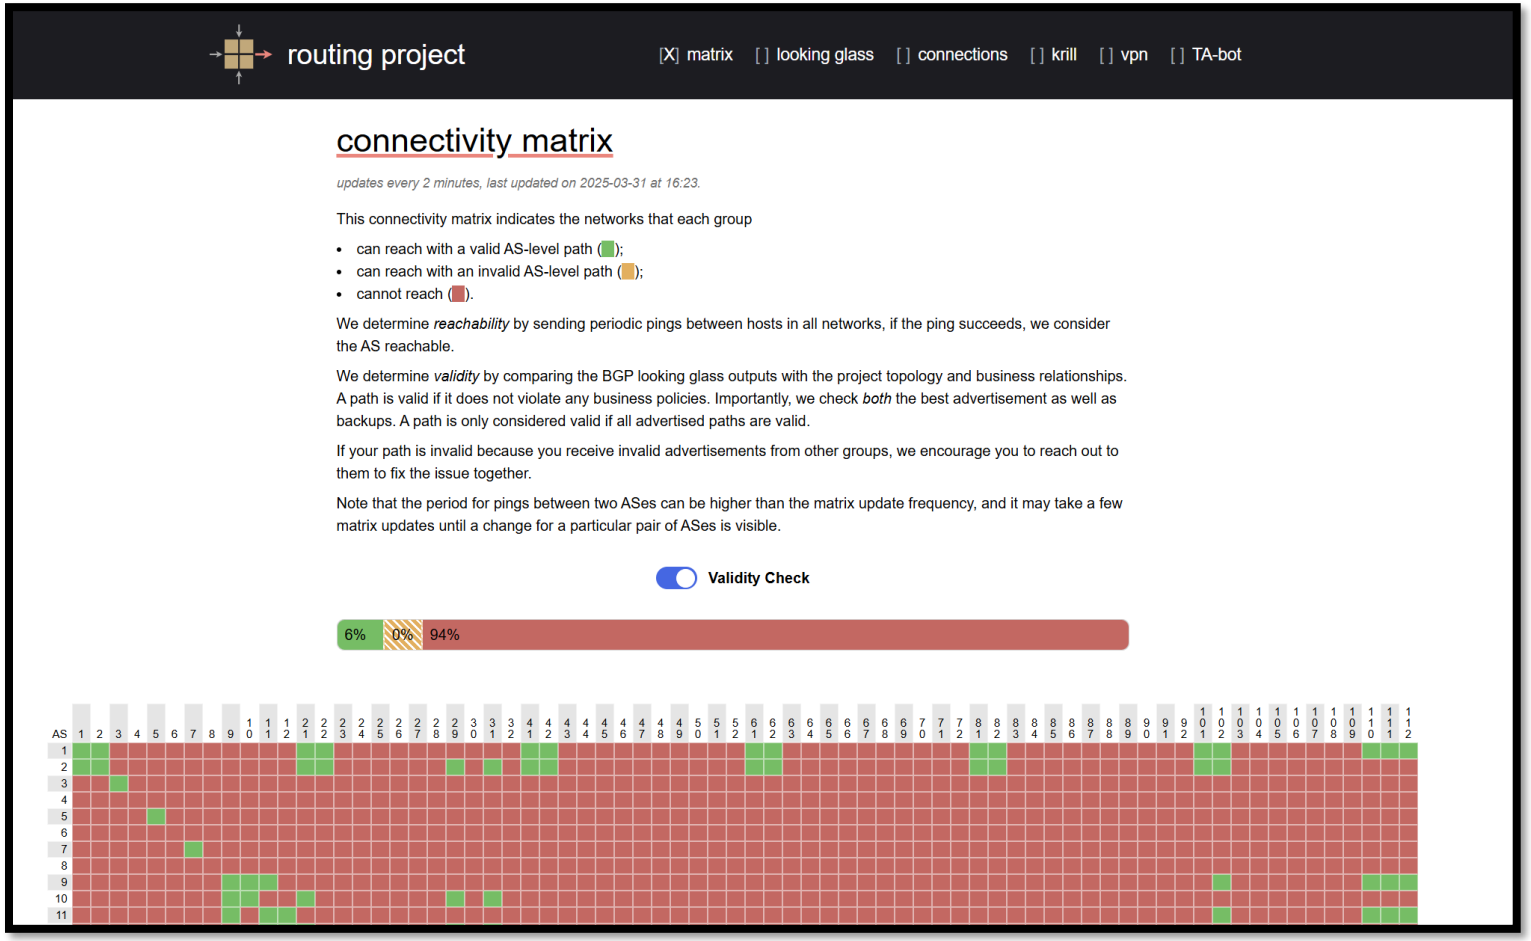

RED
Not reachable

YELLOW
Reachable but not following business relationships

GREEN
Reachable using a «valid» path

What you will be doing today

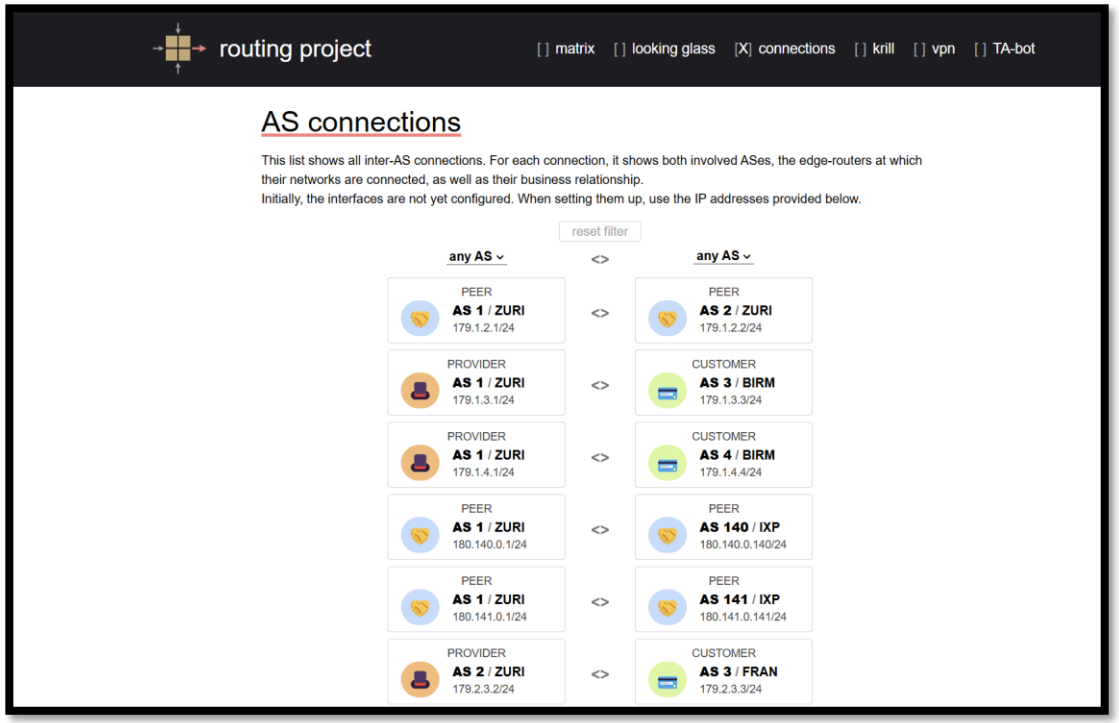

- Have a look at duvel.ethz.ch/as-connections and enter your AS number in the filer
- This webpage will show you the business relations as well as the subnets to use

What you will be doing today

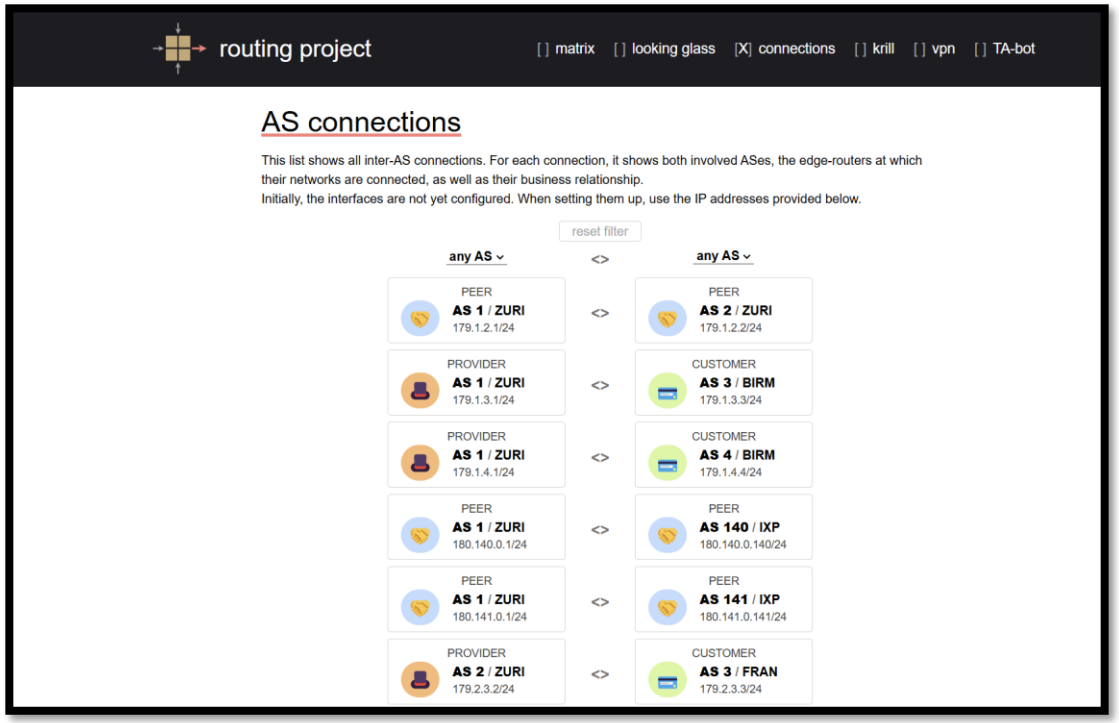

- **You will need to find the students that manage a neighboring AS**
  - Discuss which IP address each of you will use within the subnet provided such that you can establish a session
- For TA-managed ASes, use the provided IP address

What you will be doing today

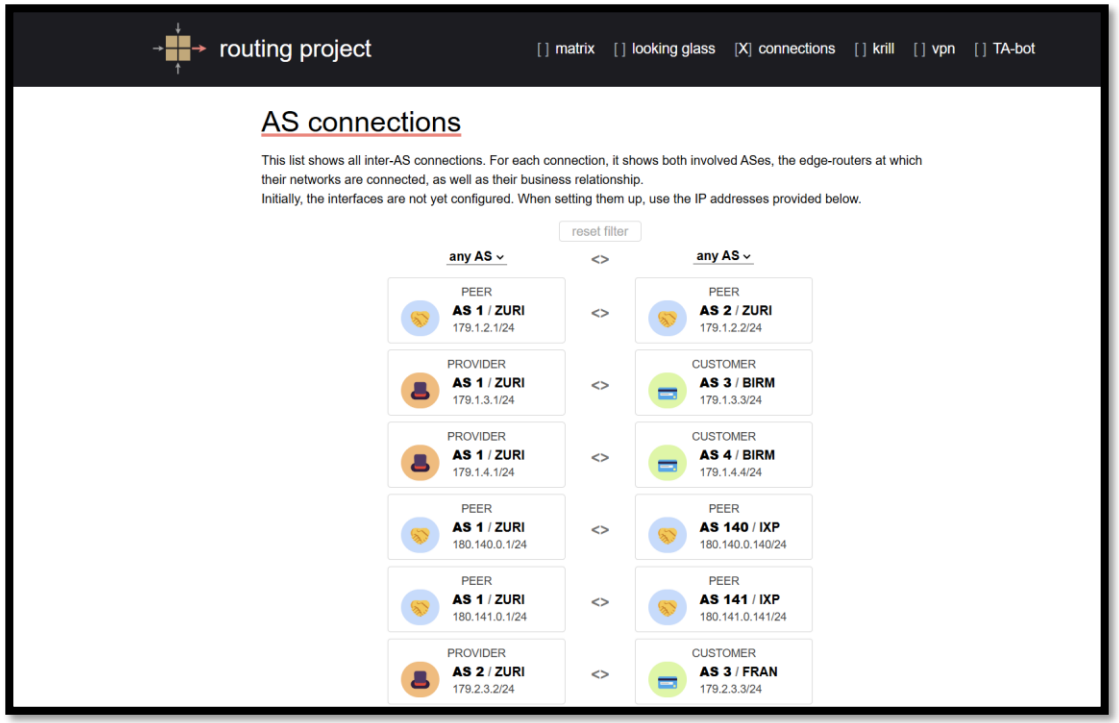

- If your neighbors aren’t here reach out to them on the routing project matrix channel later
- If this doesn’t work let us know

There are several useful tools for debugging BGP

- **show ip bgp summary**
  BGP neighbors and session status
- **show ip bgp**
  Routes learned with BGP with the AS path
- **show ip bgp neighbor x.x.x.x routes**
  Routes received from x.x.x.x
- **show ip bgp neighbor x.x.x.x advertised-routes**
  Routes advertised to x.x.x.x

Debug output from other ASes

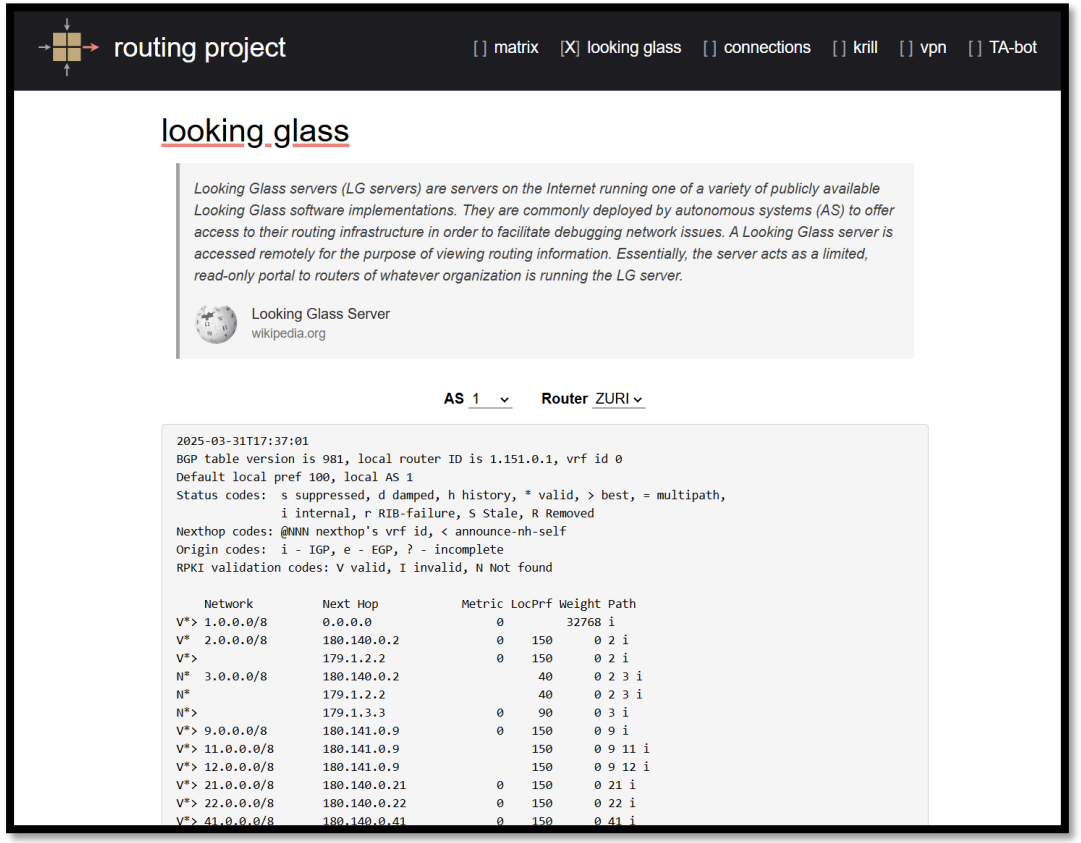

- Look at duvel.ethz.ch/looking-glass
- Shows the output of *show ip bgp* for a given AS and router
- Can help you understand how other ASes reach you
- Police Analyzer to help with some common policy mistakes

Common mistakes

- iBGP & update-source
- eBGP & next-hop-self
- Business relationship

iBGP sessions should use the loopback interface

- Each router is available over multiple paths, if one link fails the iBGP session should still work
- For that reason use the loopback interface which is reachable as long as the router still has one link that works
- When configuring the iBGP session, use:
  neighbor X.X.X.X update-source lo
- The loopback address should only be reachable internally

border routers must use next-hop-self for their iBGP sessions

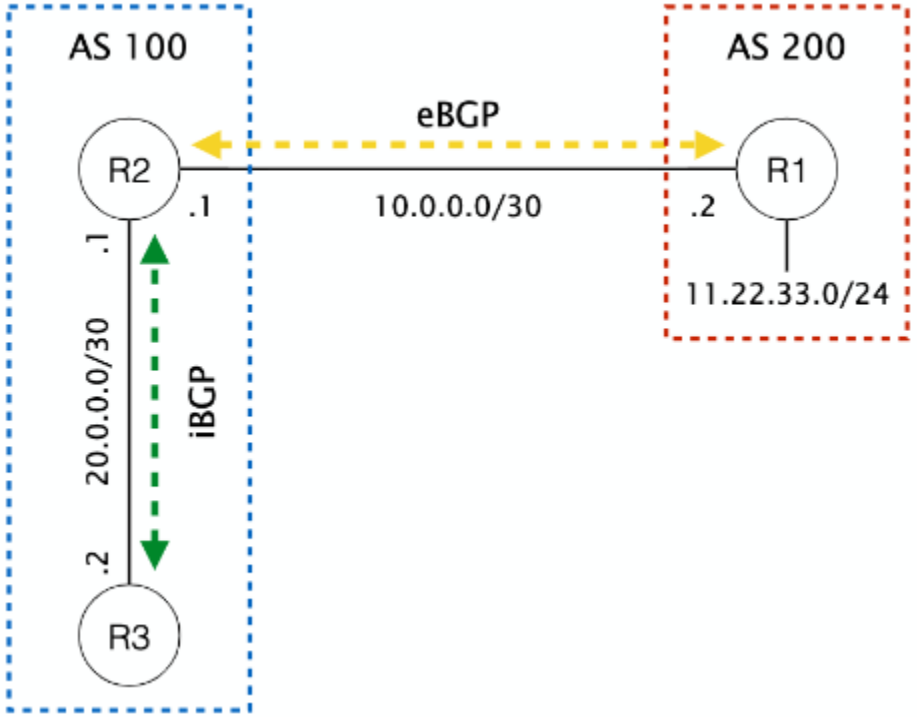

- R1 announces prefix for AS 200 to R2: *You can reach AS 200 with next hop 10.0.0.2*
- R2 propagates the announcement to R3 via iBGP

- **Problem: R3 does not know how to reach 10.0.0.2 only R2 does**

border routers must use next-hop-self for their iBGP sessions

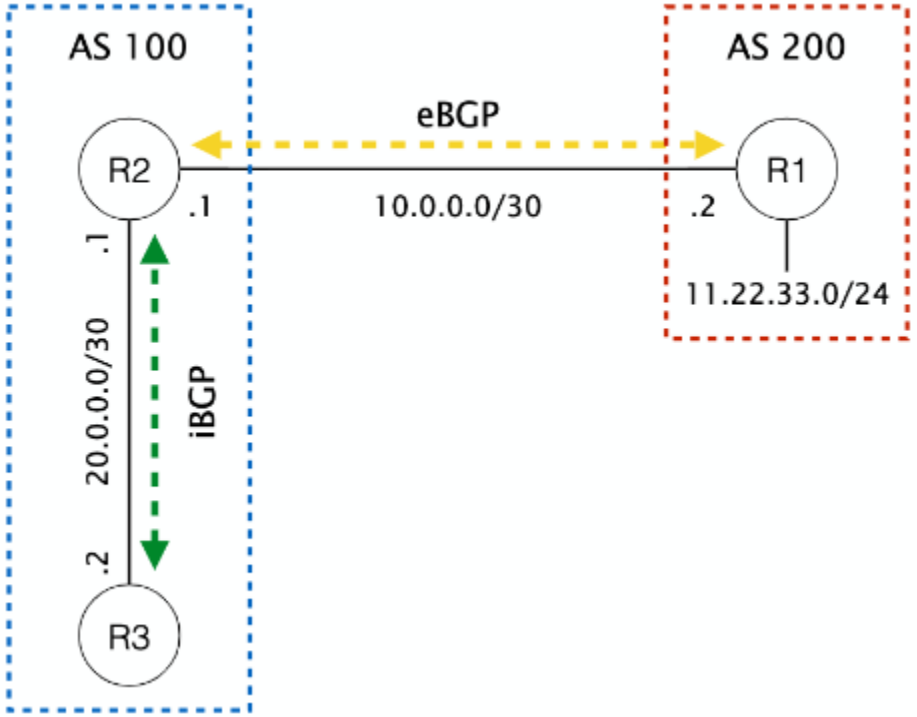

- **Solution:** When configuring the iBGP session on R2, use:
  neighbor x.x.x.x next-hop-self

- **Result:** R2 send the following to R3:
  «You can reach AS 200 with next hop <loopback address of R2>»

BGP business relationships

- You might have heard about local-pref which can be used to prioritize certain announcements

- But why would you prefer one advertisement over another?

BGP business relationships

- Of course it is **money**

- More details in the lecture

BGP business relationships

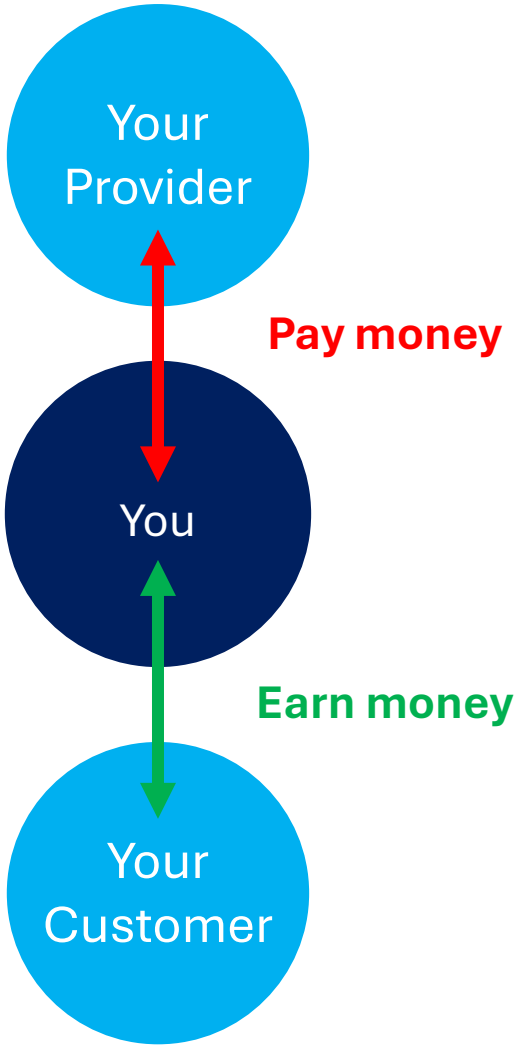

- In the real world you pay your provider for access
- Similarly your customers will pay you for access

- So if at all possible you want to send your traffic to your customer not to your provider, because that generates you money

- How?
  Filtering advertisements and using local-pref

BGP business relationships

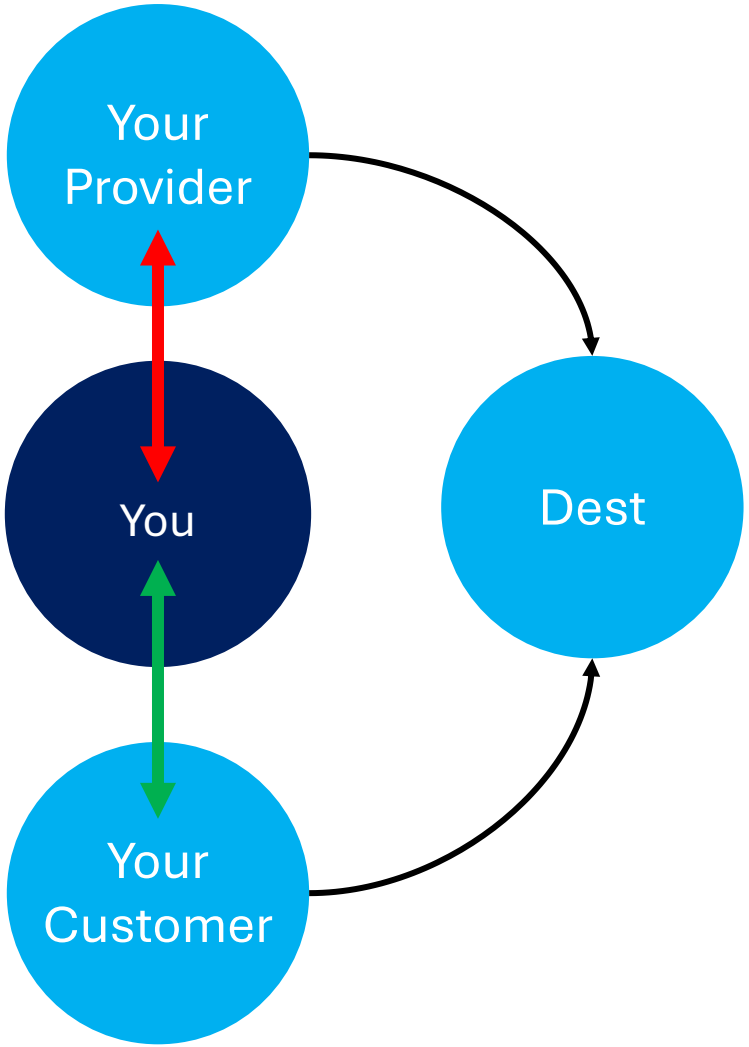

- Destination is reachable both through your provider as well as your customer
- Prefer the one that goes over your customer

Peers exchange traffic for free out of common interest.

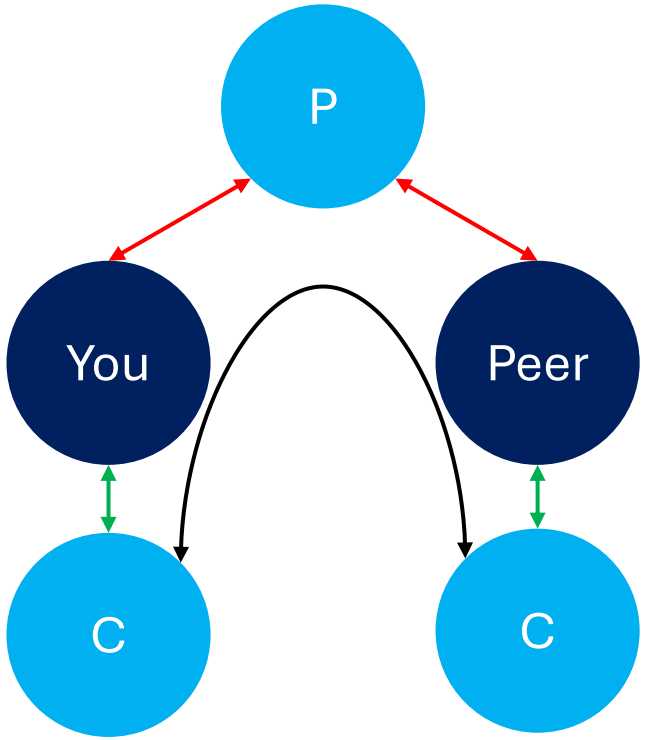

No Peering: You get paid by your customer but you also have to pay your provider

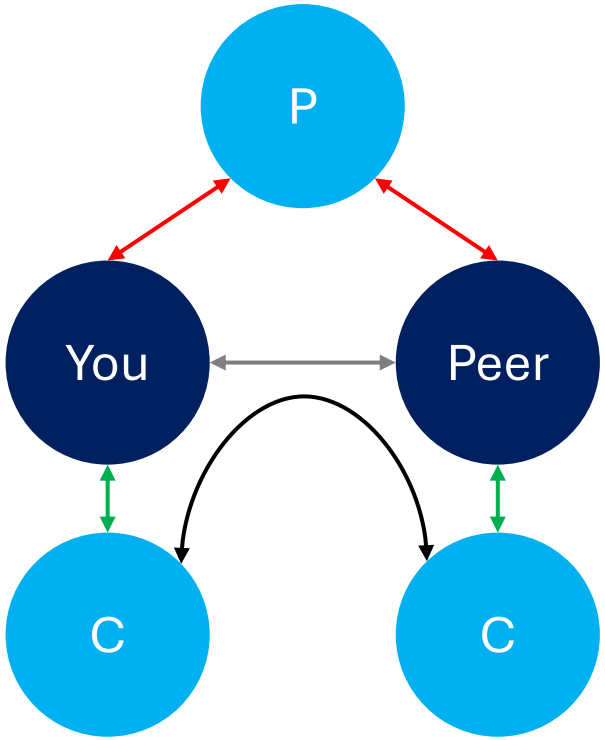

Good Peering: Your customer still pays, but you don't have to pay your provider

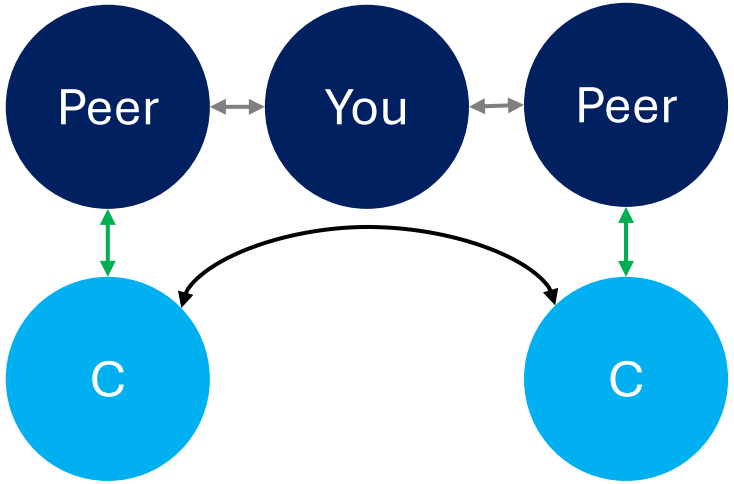

Bad Peering: You need to do work but don't get paid by anyone

How to enforce those rules

- **Outgoing traffic is determined by which incoming advertisements you select (route-map in)**
  If you deny an advertisement, you cannot use that path.
  For multiple paths, use local-pref to select the best.

- **Incoming traffic is determined by which selected advertisements you send out (route-map out)**
  If you deny, others cannot use this path via your AS.
  Remember that only the selected advertisements are considered.

Cheat sheet

import	**local-pref**
customer	high
peer	medium
provider	low

		send to		
	export	customer	peer	provider
from	customer	✅	✅	✅
	peer	✅	🛑	🛑
	provider	✅	🛑	🛑