

# **Making Network Education More Accessible**

## **Practical learning with the mini-Internet project**



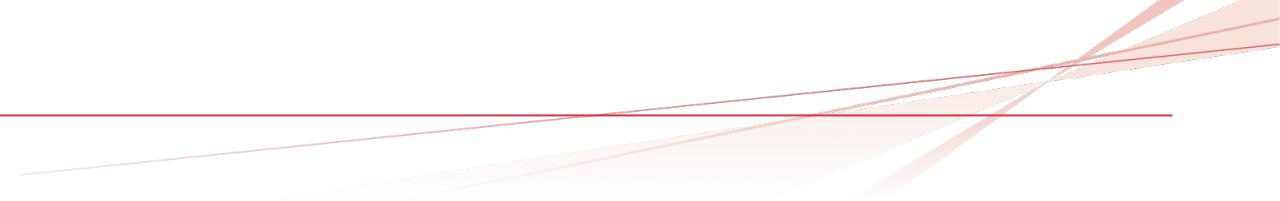
2024年4月18日

株式会社インターネットイニシアティブ  
Christoff Visser  
IIJ - Research Lab

- How do we traditionally teach how the Internet works?
  - Introducing the mini-Internet:
    - Mimic the real one and is virtual
    - Turn students into network operators
    - Provide tools to ease operations
  - The mini-Internet use cases
  - The mini-Internet in research
- これまでインターネットの仕組みをどのように教えてきたか？
  - ミニインターネットの紹介：
    - 仮想環境で実際のインターネットを模倣
    - 学生をネットワークオペレータに育てる
    - オペレーションを容易にするツールを提供
  - ミニインターネットの使用例
  - 研究におけるミニインターネット

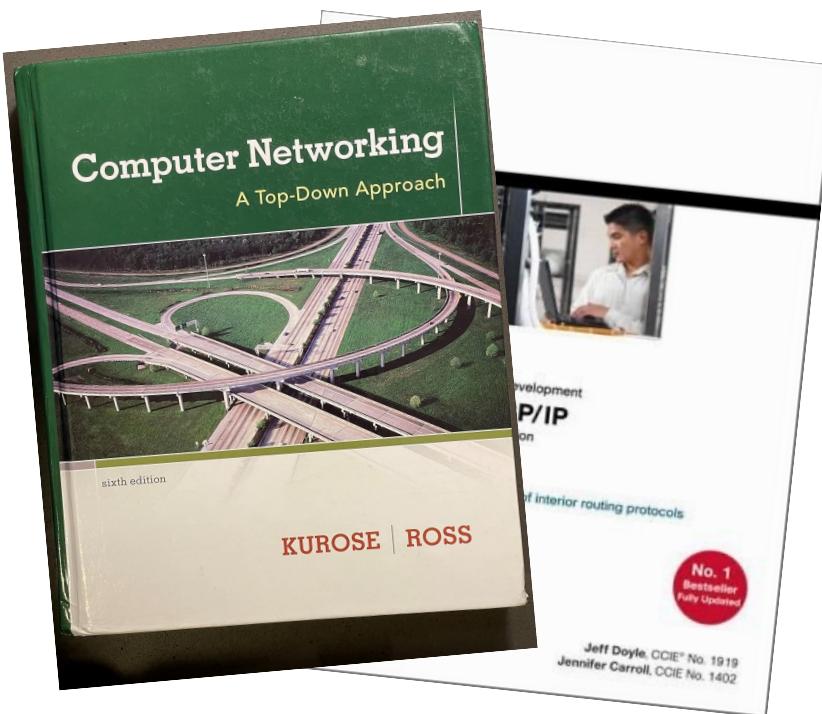
## How do we traditionally teach how the Internet works?

---

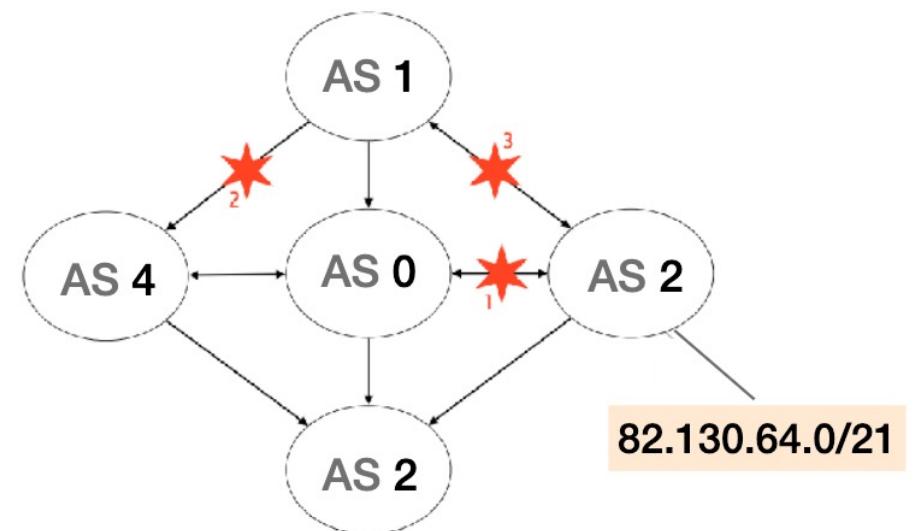


## How do we traditionally teach how the Internet works?

theory



assignments



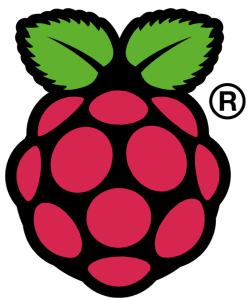
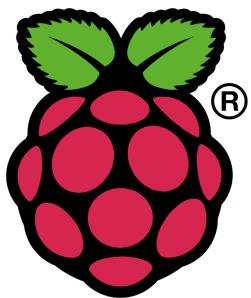
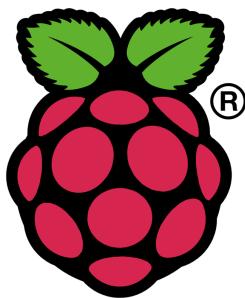
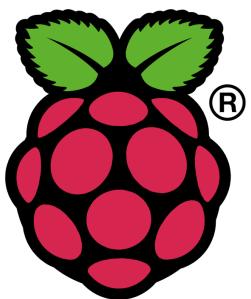
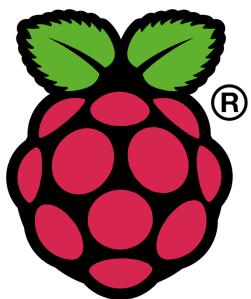
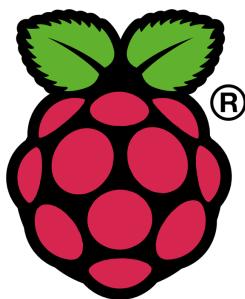
Which messages are exchanged?

## How do we traditionally teach how the Internet works?

---

Labs

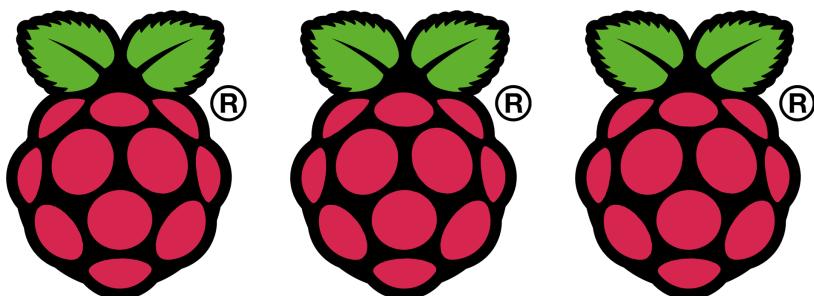
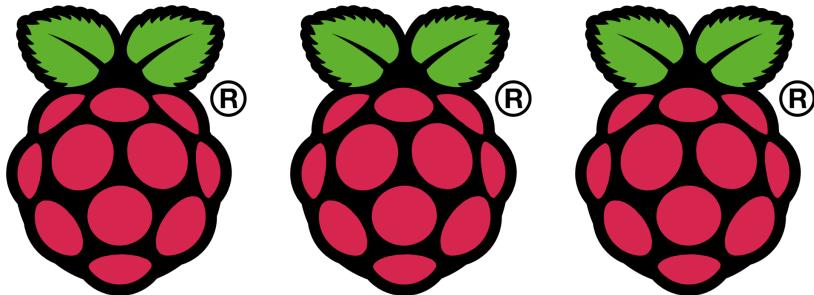
Hosts



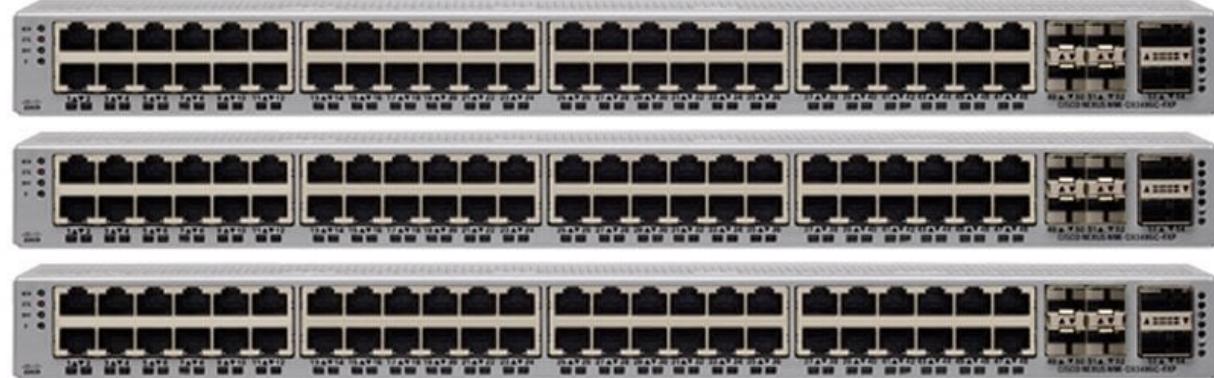
## How do we traditionally teach how the Internet works?

### Labs

**Hosts**

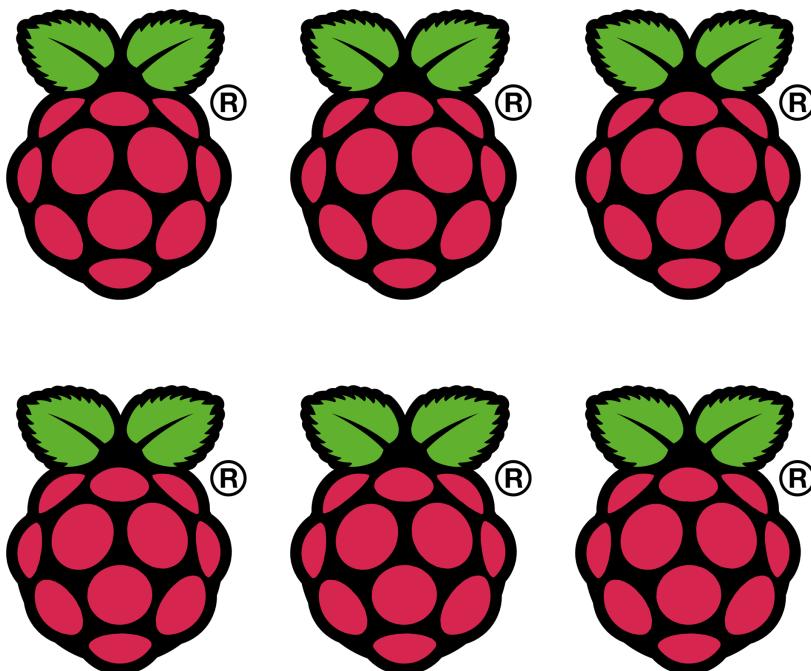


**Switches**



## How do we traditionally teach how the Internet works?

Labs



Hosts



Routers

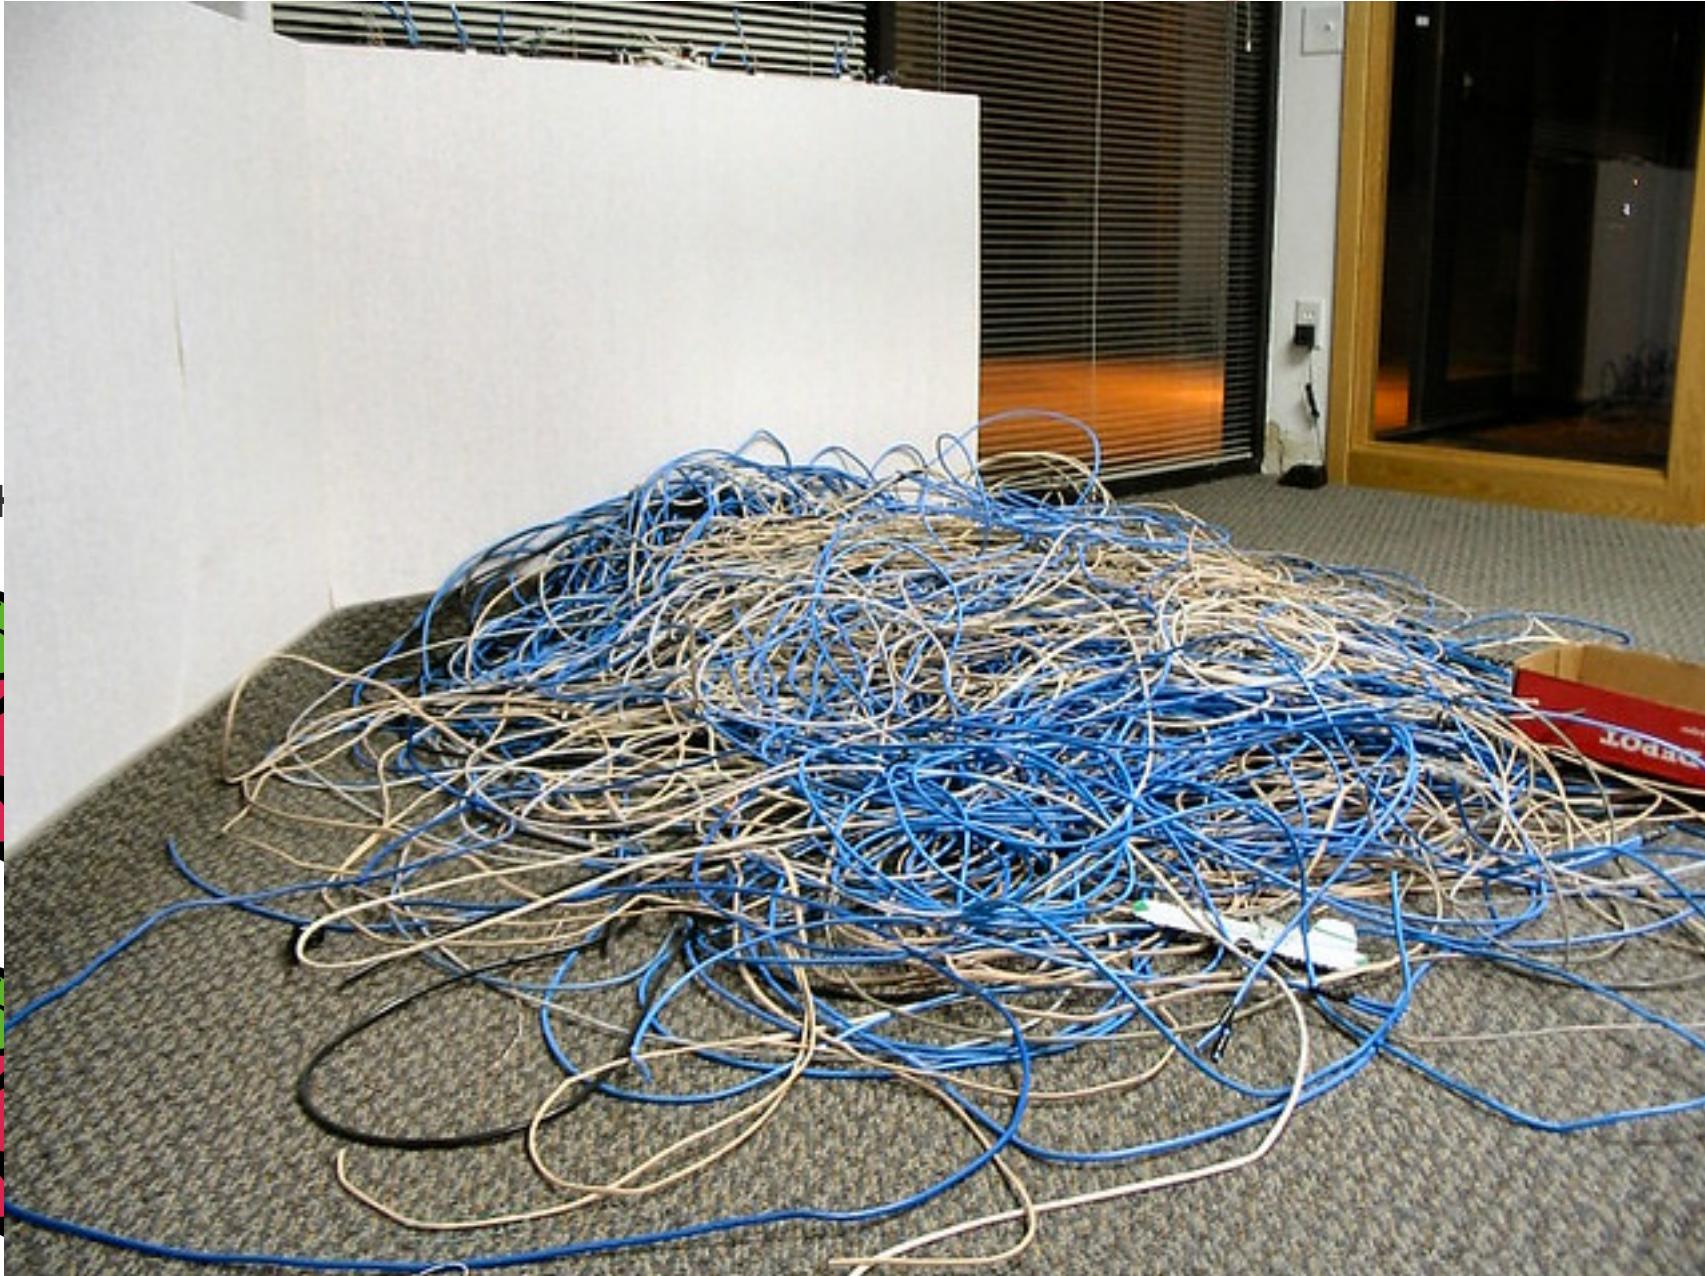
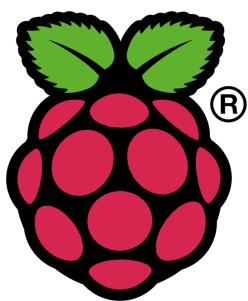
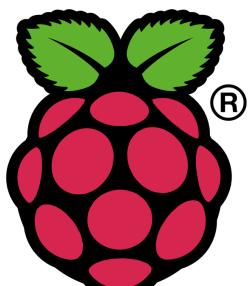


Switches



## How do we traditionally teach how the Internet works?

Labs



## Networking is not just about the technical aspects

- Peering agreements, peering events
- Mailing lists, slack, etc.

**Starting:** Fri Sep 1 01:40:32 UTC 2023

**Ending:** Sat Sep 30 22:55:22 UTC 2023

**Messages:** 413

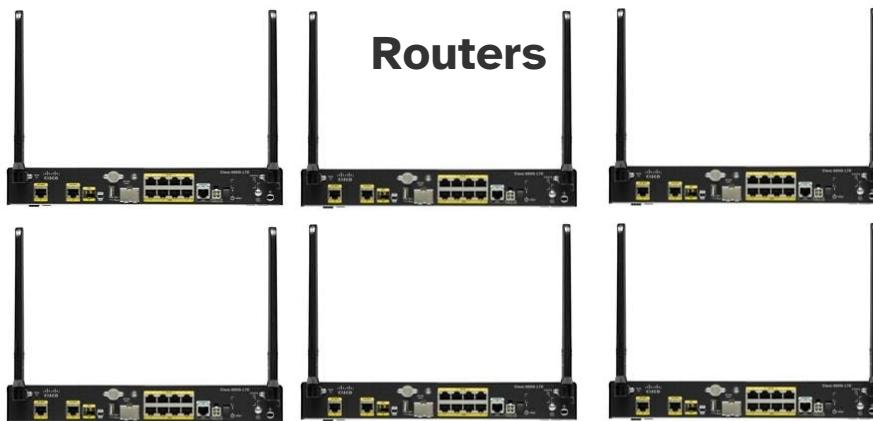
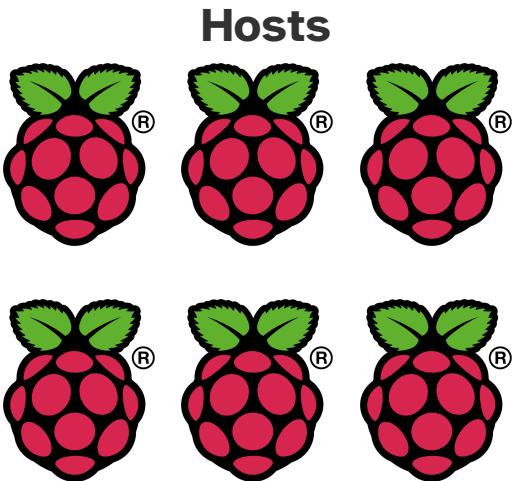
- [JunOS/FRR/Nokia et al BGP critical issue](#)
  - [JunOS/FRR/Nokia et al BGP critical issue](#)
    - [JunOS/FRR/Nokia et al BGP critical issue](#)
      - [JunOS/FRR/Nokia et al BGP critical issue](#)
      - [JunOS/FRR/Nokia et al BGP critical issue](#)
      - [JunOS/FRR/Nokia et al BGP critical issue](#)
      - [JunOS/FRR/Nokia et al BGP critical issue](#)
      - [JunOS/FRR/Nokia et al BGP critical issue](#)
      - [JunOS/FRR/Nokia et al BGP critical issue](#)
      - [JunOS/FRR/Nokia et al BGP critical issue](#)
  - [Lossy cogent p2p experiences?](#)
    - [Lossy cogent p2p experiences?](#)
      - [Lossy cogent p2p experiences?](#)
        - [Lossy cogent p2p experiences?](#)
        - [Lossy cogent p2p experiences?](#)
        - [Lossy cogent p2p experiences?](#)
        - [Lossy cogent p2p experiences?](#)
        - [Lossy cogent p2p experiences?](#)
        - [Lossy cogent p2p experiences?](#)
        - [Lossy cogent p2p experiences?](#)
        - [Lossy cogent p2p experiences?](#)
        - [Lossy cogent p2p experiences?](#)
        - [Lossy cogent p2p experiences?](#)



APRICOT 2024 Peering Social

How does this scale?

---



x 2? 5? 10?



## Introducing the mini-Internet project

---

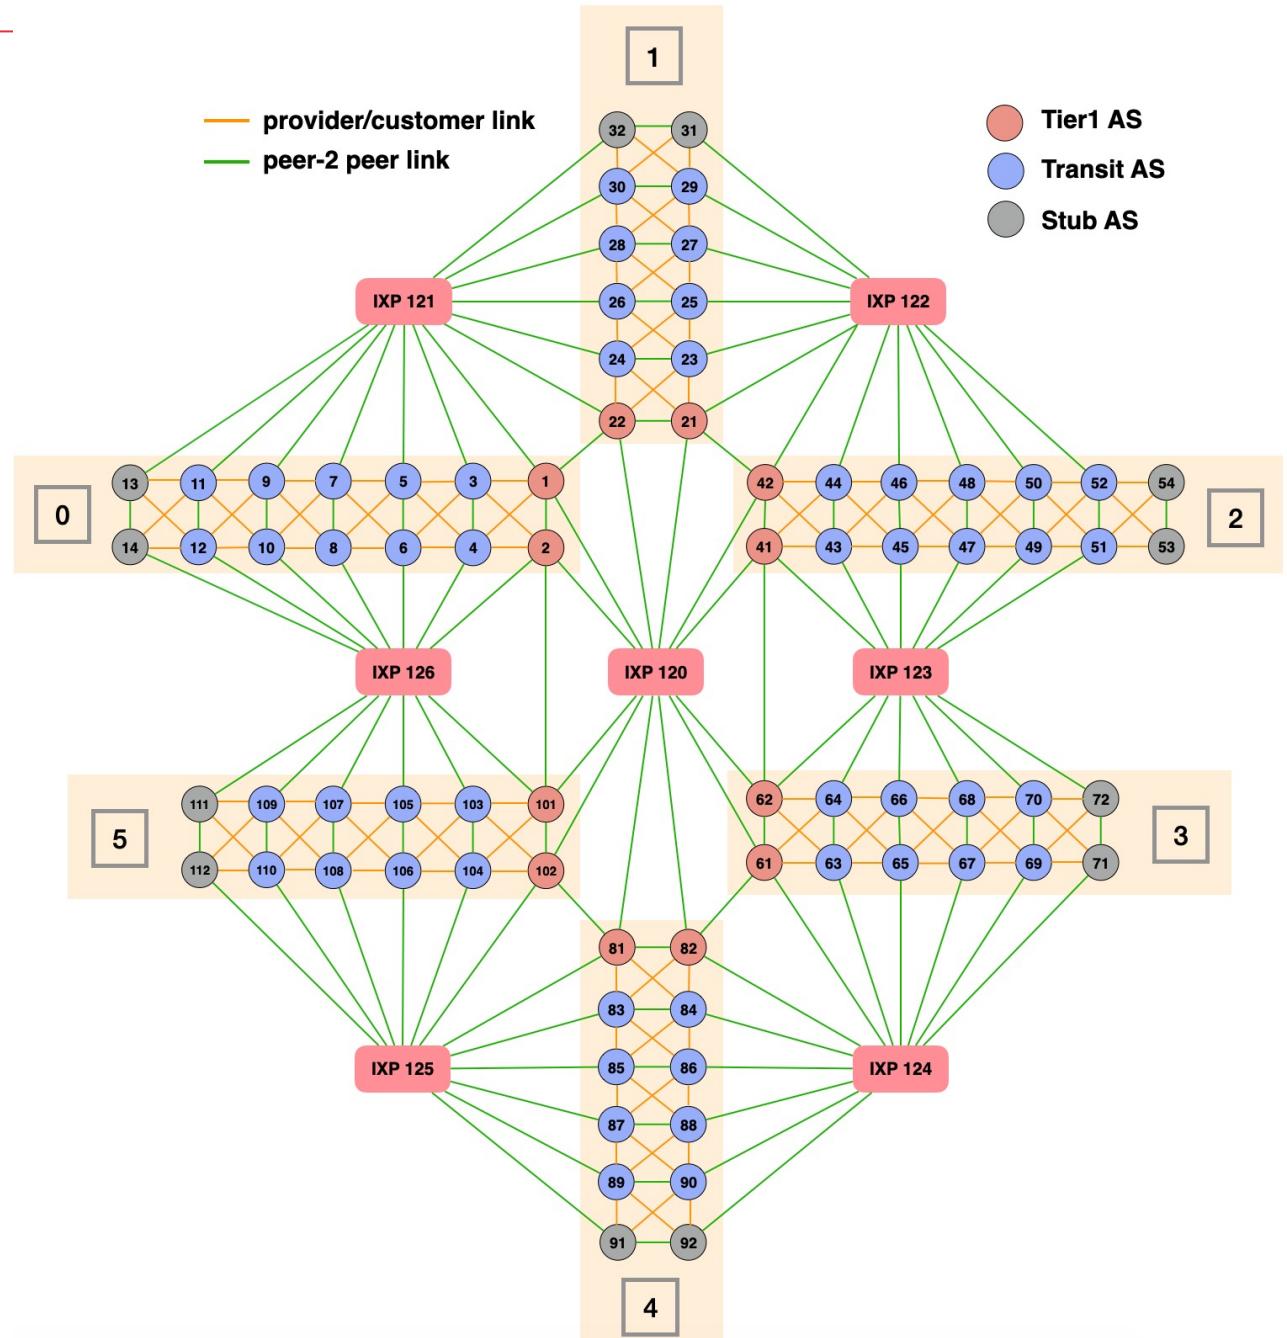
- Developed by ETH Zurich
- Goal:
  - Mimic the real Internet but is virtual
  - Turn students into network operators
  - Provide tools to ease operations

Holterbach T, et al. 2020. An Open Platform to Teach How the Internet Practically Works.  
SIGCOMM Comput. Commun. Rev. 50, 2 (April 2020), 45–52.  
<https://doi.org/10.1145/3402413.3402420>

# The mini-Internet

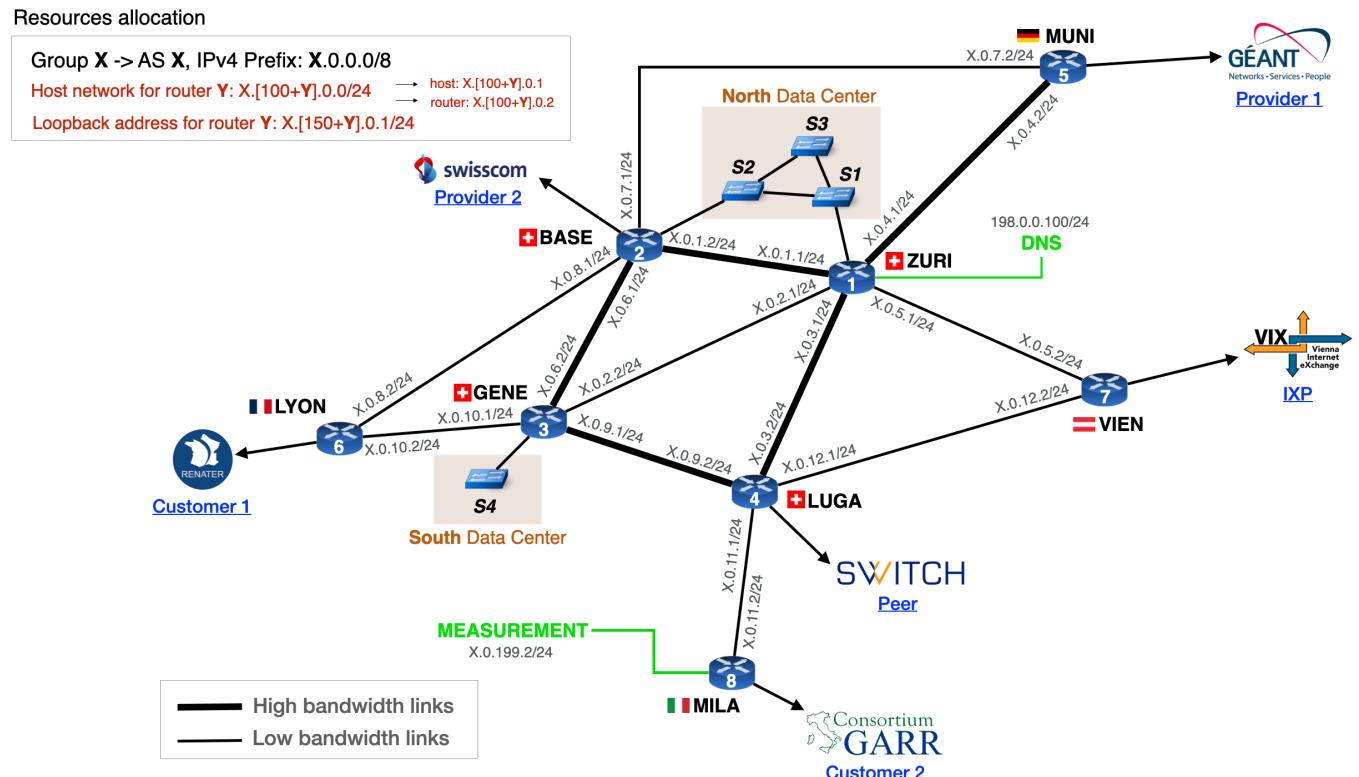
Mimics the real one and is entirely virtual

- 91 AS
- 6 regions
- 7 IXPs
- 8 Student AS per region
- 4 TA run AS per region



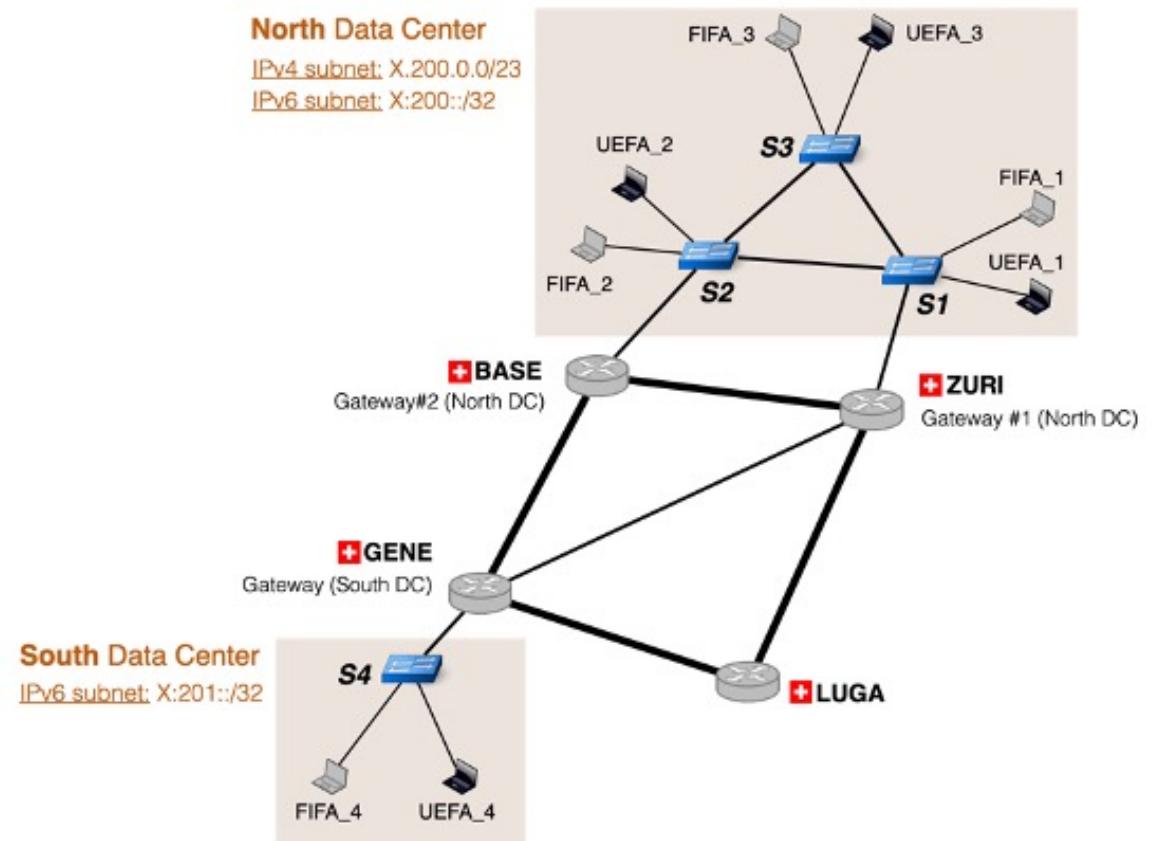
## Per AS Level Topology

- 2 Data centres
  - North DC
  - South DC
- 8 Routers
- 1 Measurement host
- Shared DNS host

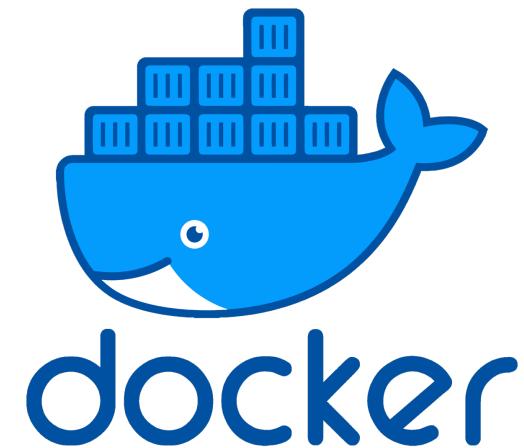


## DC Level Topology

- 2 Hosts per switch
  - FIFA – VLAN 10
  - UEFA – VLAN 20
- South DC v6 only



- Runs on a single machine
- Each component is a docker container
- Use common network tools

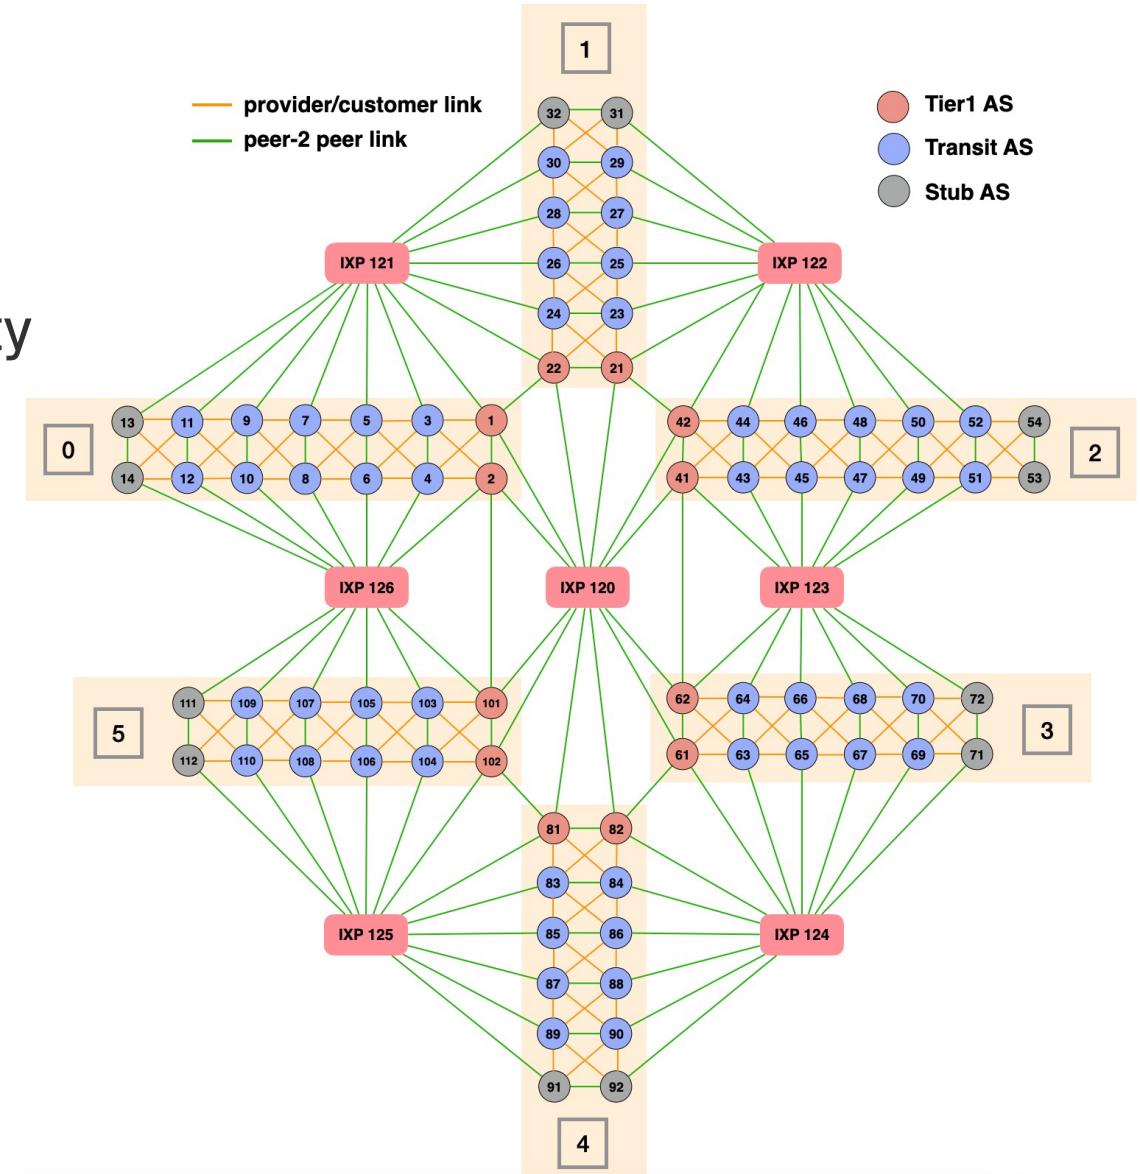


# The mini-Internet

Turn students into network operators

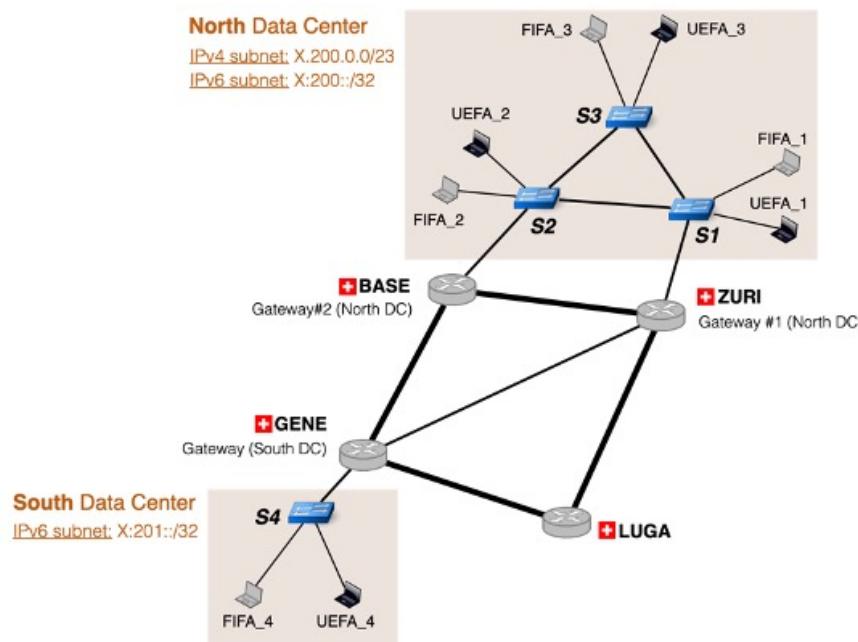
Each group has their own transit AS and one IP prefix

- Goal:
  - Enable ‘mini’ Internet-wide connectivity
- Intra-domain connectivity
- Inter-domain connectivity
- Security



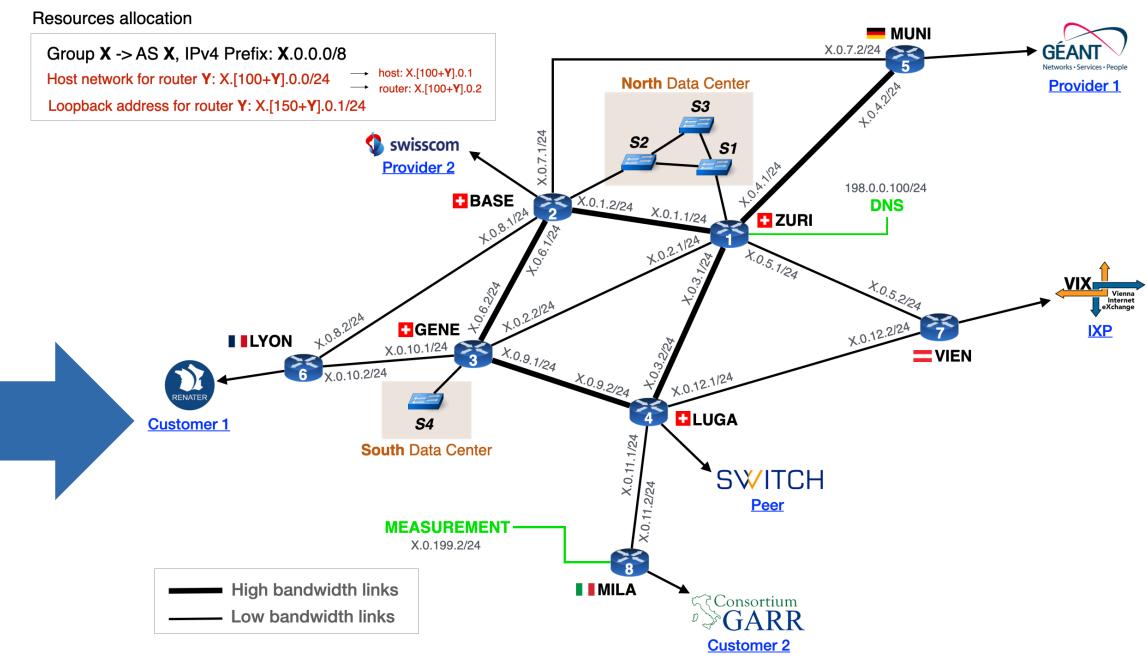
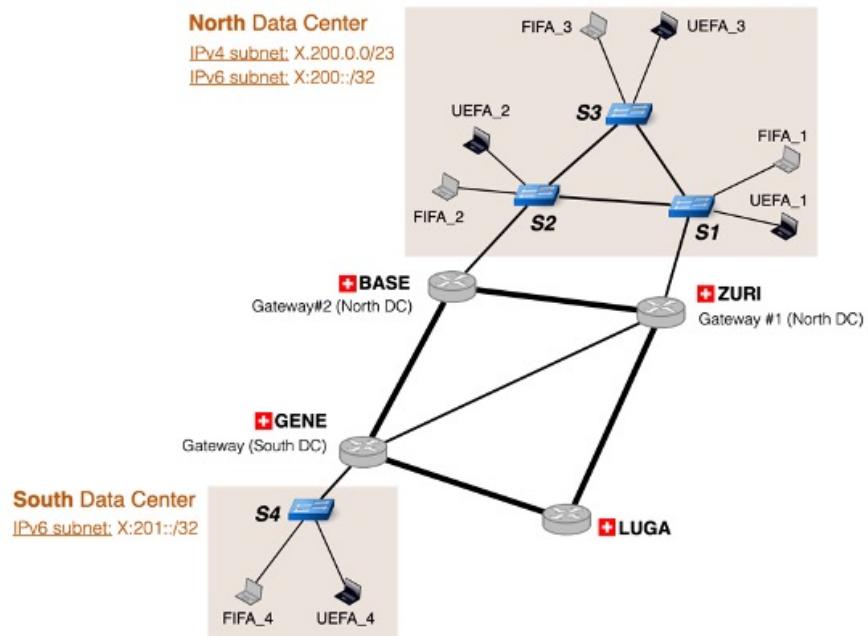
Enable internal connectivity

- Host configuration
- VLANs
- Gateways, static routes, etc.



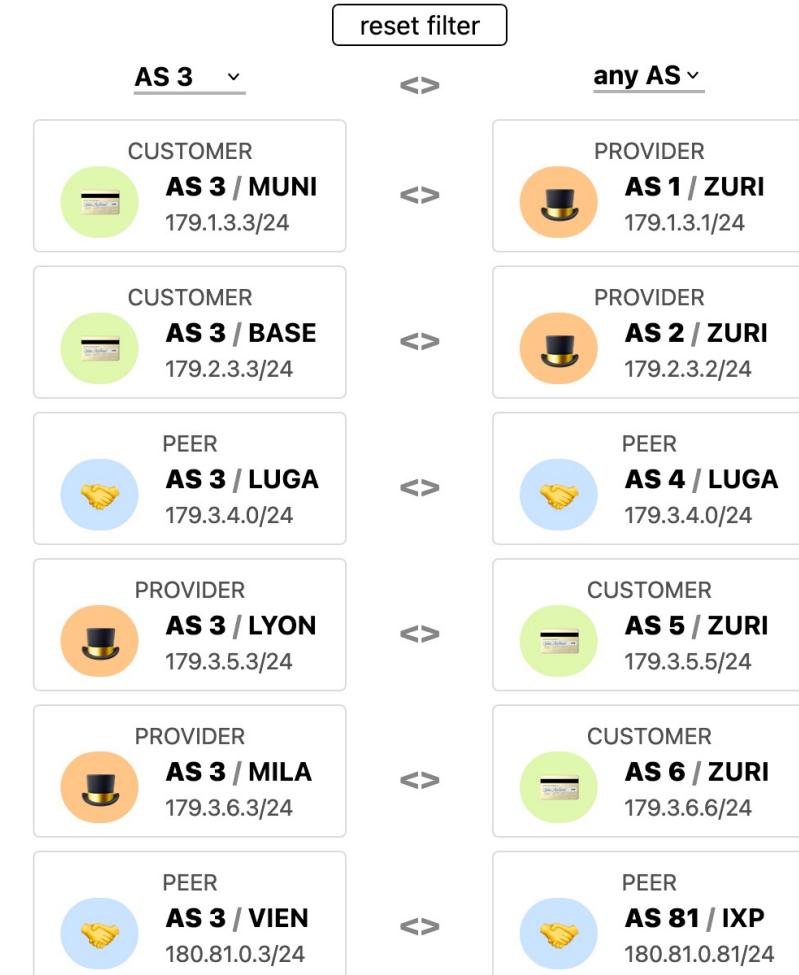
### Enable internal connectivity

- Host configuration
- OSPF + load balancing
- VLANs
- IPv6 + 6in4 Tunnel
- Gateways, static routes, etc.
- iBGP



### “Peering day”

- Networking is not just technical
  - There is also the social aspect
- Get together to configure BGP sessions
- Routing policies
  - Local-preference and export rules
- Follow peering preferences
  - e.g., one provider is preferred
- Connecting IXPs with communities



- Route hijacking
- Issue ROA
- Implement routing policies based on RPKI status
- Reports on why RPKI is important and implementation challenges



# MANRS



# The mini-Internet

Provide students with tools to ease operations

### Access

- Students are not expected to have any networking knowledge
- Helper scripts for access and working on their network
- SSH ports based on group number:  $X = 2000 + \text{group\_number}$ 
  - Group 7 use port 2007
  - Access proxy to access rest of groups network
- Helper script to help students that might not be used to SSH:
  - `./goto.sh` LUGA router
  - `./goto.sh` DCN S2
  - `./goto.sh` DCN FIFA\_3

### Student wiki

- Split into Questions and Tutorial
- Can be self-run

0. Routing Project Overview

1. Assignment

- 1.1 General Instructions
- 1.2 Your mini Internet
- 1.3 Questions
- 1.4 Tools to help you

2. Tutorial

- 2.1 Accessing your devices
- 2.2 Configuring a host
- 2.3 Configuring Open vSwitch
- 2.4 Configuring 6in4 tunnels
- 2.5 Configuring IP routers
  - 2.5.1 The FRRouting CLI
  - 2.5.2 Router interfaces
  - 2.5.3 Static routes
  - 2.5.4 OSPF
  - 2.5.5 BGP
  - 2.5.6 BGP policies
- 2.6 VPN Configuration
- 2.7 RPKI Configuration

3. FAQ

## Course Question

### Question 1.1 (1.25 point)

Your goal for this question is to enable direct layer-2 connectivity in the North Data Center between the `FIFA` hosts, between the `UEFA` hosts, but not in between them. Yet, `FIFA` hosts and `UEFA` hosts in the North Data Center should still be able to communicate between themselves, but via a layer-3 router. This will prevent typical layer-2 attacks such as MAC spoofing used to impersonate a type of user and get access to sensitive data. To do that, you will have to configure, in each Data Center, the hosts, the switches, and the routers directly connected to the Data Center.

*Configure the IP addresses and default gateway:* To enable end-to-end connectivity in the North Data Center, you will need to configure an IPv4 address as well as a default gateway on each host of each company, and on the interfaces of the routers `ZURI` and `BASE` connected to the Data Center. For this question, you must use the IPv4 subnet `X.200.0.0/23` where X is your group number (you do not have to configure IPv6 addresses for this question). You are free to use any IP address as long as it is in that subnet. To test connectivity, you can use `ping`.

In addition, every host needs to have an IPv4 default gateway to be able to reach external destinations. Configure it such that all hosts in the North Data Center use `ZURI` as standard gateway.

*Configure the VLANs:* You also have to configure VLANs: use VLAN 10 for `FIFA` and VLAN 20 for `UEFA`. VLAN 30 is reserved for later use. To configure VLANs, you will have to configure tagged and trunk ports on the switches. We explain how to do that in our tutorial.

👉 The interface of `ZURI` connected to North Data Center in VLAN 10 is named `ZURI-L2.10`, and the one in VLAN 20 is named `ZURI-L2.20` (you can see them with a `show interface brief` in the FRRouting CLI). The same convention is used for the interfaces in `BASE`. Do not use the interfaces `ZURI-L2` in the router `ZURI` and the interface `BASE-L2` in the router `BASE`.

📝 **To include in your report:** Explain what IP addresses you assigned to the different hosts in the North Data Center. Finally, show the output for one `traceroute` from `FIFA_1` to `FIFA_3` and one from `FIFA_1` to `UEFA_3`. In a few sentences, explain what you observe.

*Hint:* Use `traceroute` with the option `-n` to prevent traceroute from trying to translate the IP addresses to the respective domain names. This does not work because DNS is not yet configured.

## 2.2 Configuring a host

This section explains how you can configure the various hosts (connected to switches and routers) in your mini-Internet.

### Naming conventions

The interface to the router is called `<router-name>router`, e.g. `GENErouter` for `GENE`. The loopback interface has the name `lo`.

### Reading the configuration

You can see the interfaces with `ip address show`.

### Changing the configuration

The host configuration is rather simple. First, you have to assign an IP address and subnet to one of the interfaces available on the host. For that, you can use the following command:

```
ip address add IP/SUBNET_SIZE dev INTERFACENAME
```

For example, with

```
ip address add 111.0.222.3/24 dev BASErouter
```

you assign the IP address `111.0.222.3` inside the corresponding /24 subnet to the interface called `BASErouter`. Note that we configure the IP address and the corresponding subnet all at once. To remove an IP address, you can use the following command:

```
ip address del 111.0.222.3/24 dev BASErouter
```

At this point, the host knows how it can reach all the IP addresses inside the subnet `111.0.222.0/24`. Unfortunately, we cannot reach any other IP. (The subnet declaration tells the host in what subnet it is; the IP address in turn is the address by which other devices can reach this host). Hence, you normally add a route towards a default gateway that the host uses to reach all unknown IP addresses. To do that, use the command:

```
ip route add default via IP_ADDRESS
```

## 2.5.2 Configuring router interfaces

In order to explain how to configure interfaces, we will first describe their purpose, and then tell you about the naming conventions for this project. Next, we will show you how to read and change the interface configuration.

### The purpose of an interface

A router interconnects IP networks through several IP interfaces. When receiving a packet from one interface, it forwards it to another based on pre-computed forwarding decisions. Each IP interface must have an IP address configured and must be in a different subnet.

Try to think what would happen if that wasn't the case.

### Interface naming conventions

Each router has interfaces to its neighbouring routers whose names follow the pattern `port_<neighbor>`. For instance, the interface on `ZURI` connected to `LUGA` is named `port_LUGA`. Moreover, each router has an interface connected to the host named `host` and a loopback interface called `lo`. An interface connected to another AS is called `ext_<AS-number>_<router-name>`. For example, the interface on `BASE` in AS 86 connected to `LYON` in AS 84 has the name `ext_84_LYON`.

### Reading interface configurations

To get an overview of the interfaces, use the command

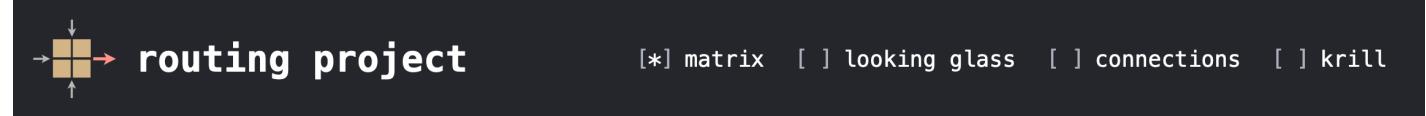
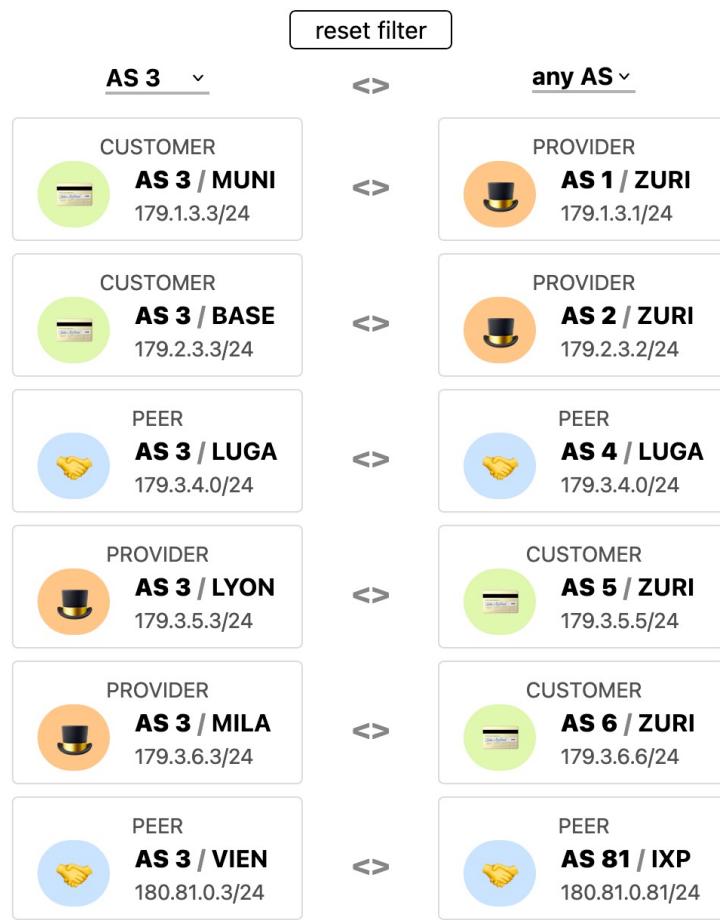
```
router# show interface
```

or for a briefer overview

```
router# show interface brief
```

You can get information for one specific interface with

```
router# show interface INTERFACENAME
```



## connectivity matrix

updates every 5 minutes, last updated on 2024-02-13 at 03:57.

This connectivity matrix indicates the networks that each group

- can reach with a valid AS-level path (green);
- can reach with an invalid AS-level path (orange);
- cannot reach (red).

We determine *reachability* by sending periodic pings between hosts in all networks, if the ping succeeds, we consider the AS reachable.

We determine *validity* by comparing the BGP looking glass outputs with the project topology and business relationships. A path is valid if it does not violate any business policies. Importantly, we check *both* the best advertisement as well as backups. A path is only considered valid if all advertised paths are valid.

If your path is invalid because you receive invalid advertisements from other groups, we encourage you to reach out to them to fix the issue together.

Note that the period for pings between two ASes can be higher than the matrix update frequency, and it may take a few matrix updates until a change for a particular pair of ASes is visible.



	1	2	3	4	5	6	1	2	3	4	5	6
1	green						green	green	green	green	green	green
2		green					green					
3			green									
4				orange								
5					green							
6						green						
11							green					
12							orange					
13								green				
14								orange				
15									green			
16									orange			

# The Mini-Internet use cases

- Education courses
  - APIE
  - APNIC, RIPE etc
  - University Courses



- Education courses
  - APIE
  - APNIC, RIPE etc
  - University Courses
- Workshops
  - NOG meetings
  - Hackathons

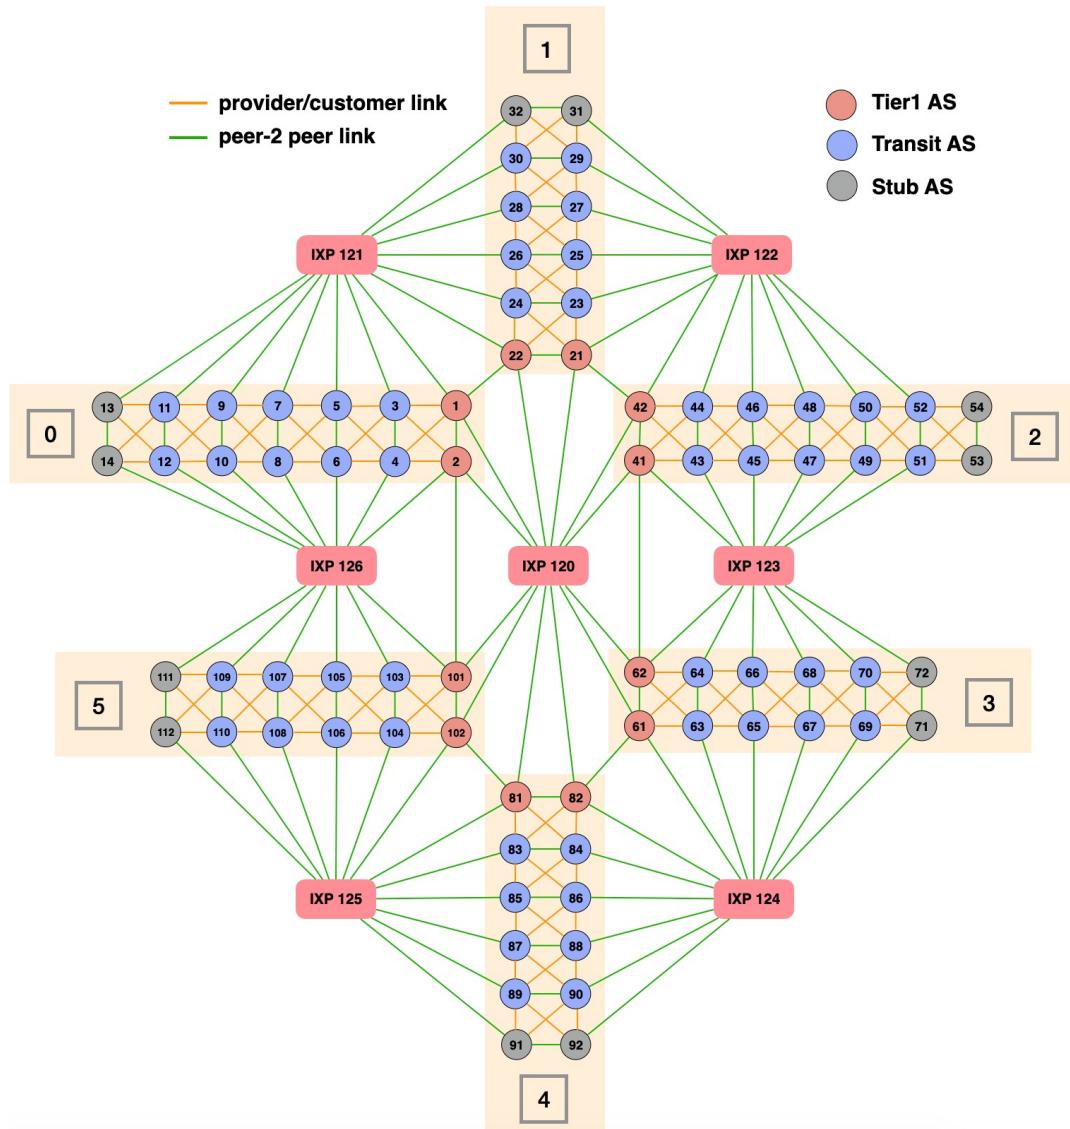


- Education courses
  - APIE
  - APNIC, RIPE etc
  - University Courses
- Workshops
  - NOG meetings
  - Hackathons
- Bootcamps
  - IIJ
  - RPKI
  - etc



# Mini-Internet in Research

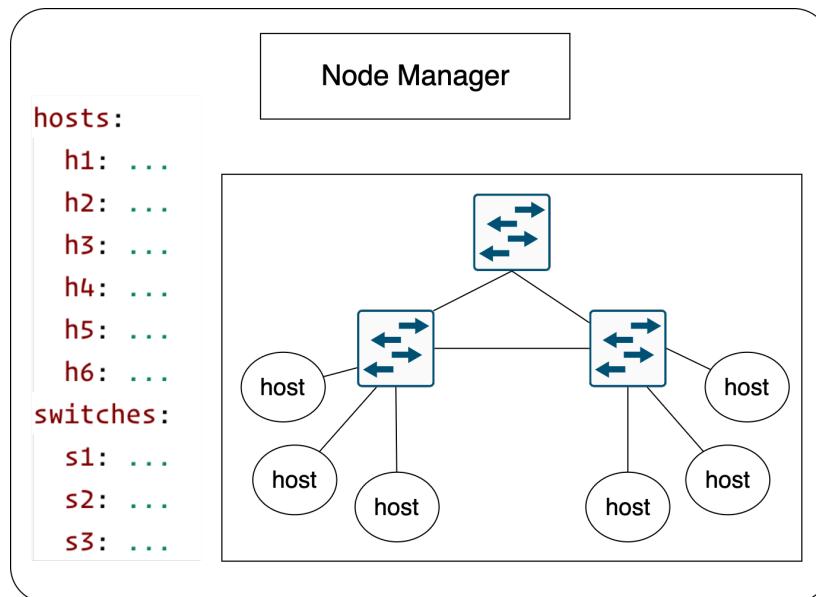
- Mimics the “real” Internet
- Safe environment for experimentation
- Ideal for testing



- Split into smaller modules
- More specialised
  - Courses
  - Experiments and testing
- Current design is very monolithic

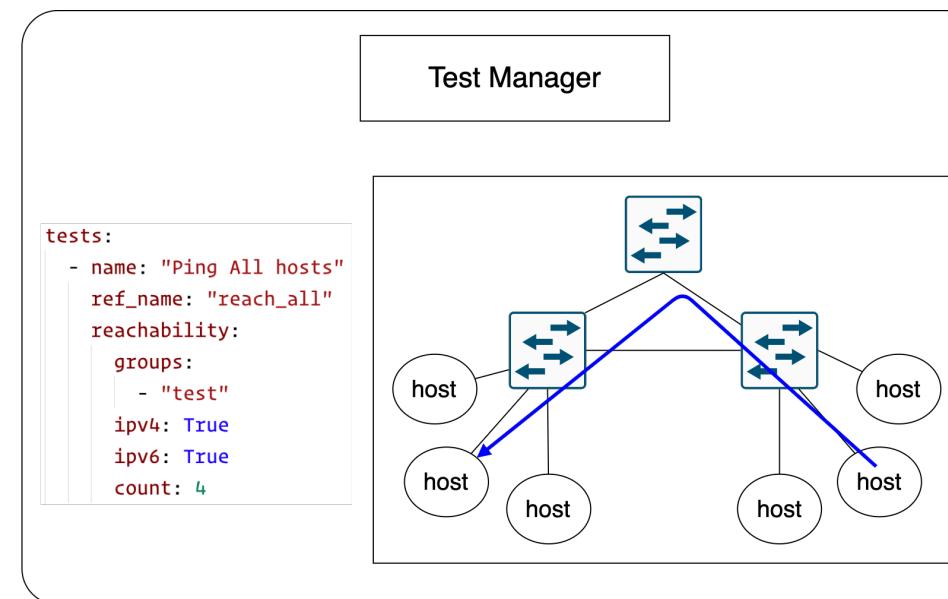
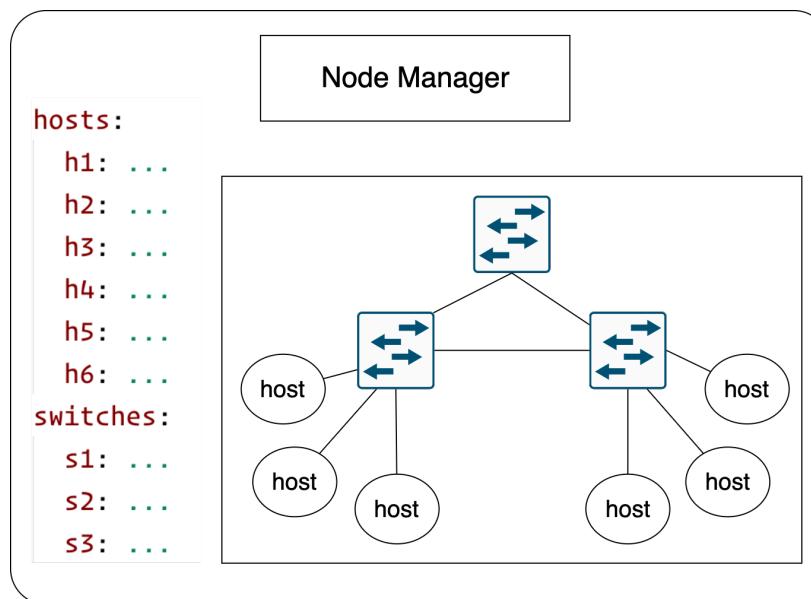
Simplify:

- Setting up and configuring test bed
- Digital twinning



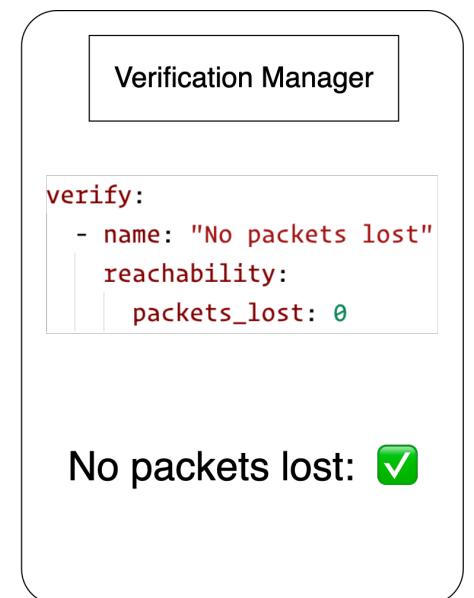
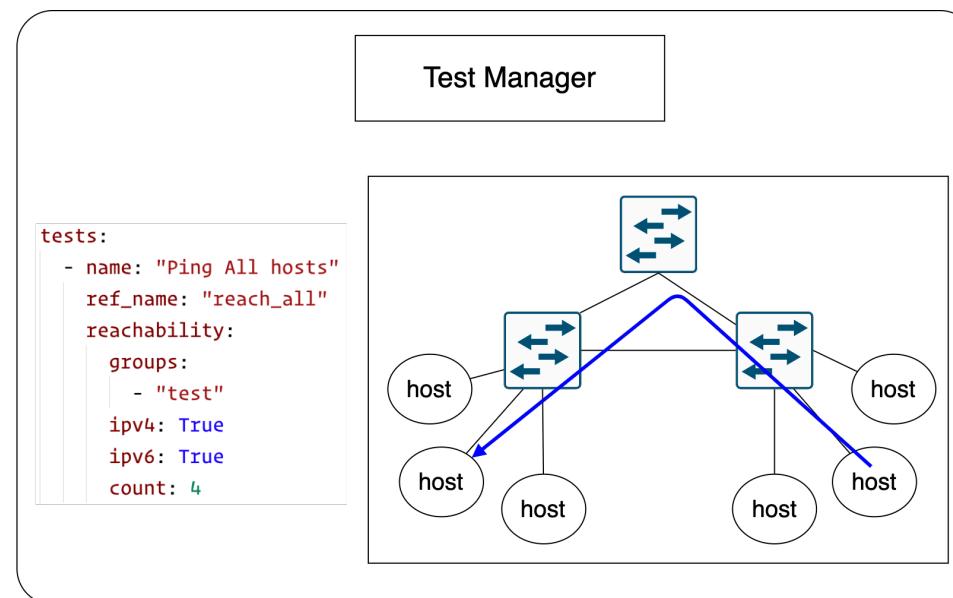
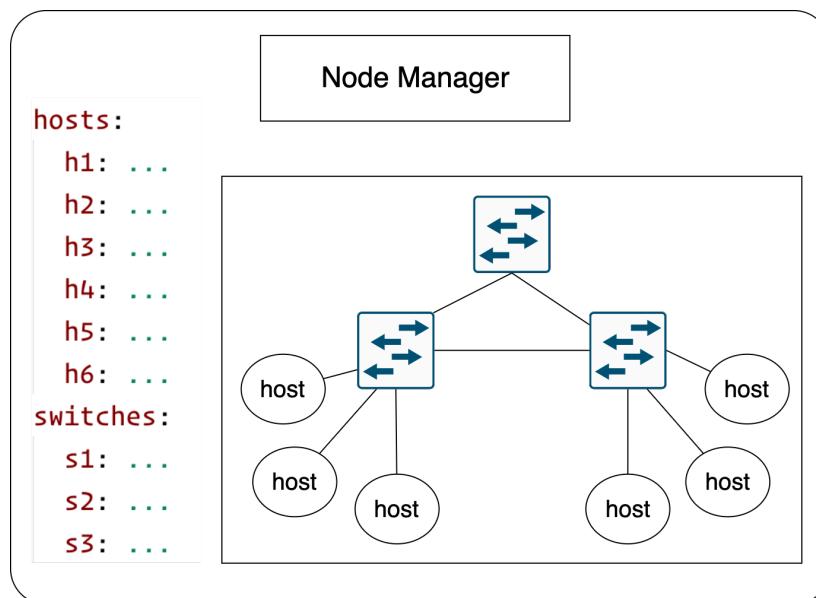
Simplify:

- Setting up and configuring test bed
- Digital twinning
- Network testing



Simplify:

- Setting up and configuring test bed
- Digital twinning
- Network testing
- Test validation



- Teaching
  - Reusable testbeds
  - Marking
- Research
  - Reproducibility
  - Build Datasets
- Industry
  - Test products
  - CI/CD



- The mini-Internet:
  - Mimics the real one and is virtual
  - Turns students into network operators
  - Provides tools to ease operations
- The mini-Internet use cases
  - Workshops, Courses, bootcamps
- The mini-Internet in research
- Network Testing Framework
  - Simplify configuring, testing and verification

- ミニインターネット：
  - 仮想環境で実際のインターネットを模倣
  - 学生をネットワークオペレータに育てる
  - オペレーションを容易にするツールを提供
- ミニインターネットの使用例
  - ワークショップ、コース、ブートキャンプ
- 研究におけるミニインターネット
- ネットワークテストフレームワーク
  - 設定、テスト、および検証の簡素化

- Mini-Internet
  - Run your own

[https://github.com/nsg-ethz/mini\\_internet\\_project](https://github.com/nsg-ethz/mini_internet_project)

- Or reach out:
  - Help hosting your own event
  - Network Testing Framework
  - Anything else:

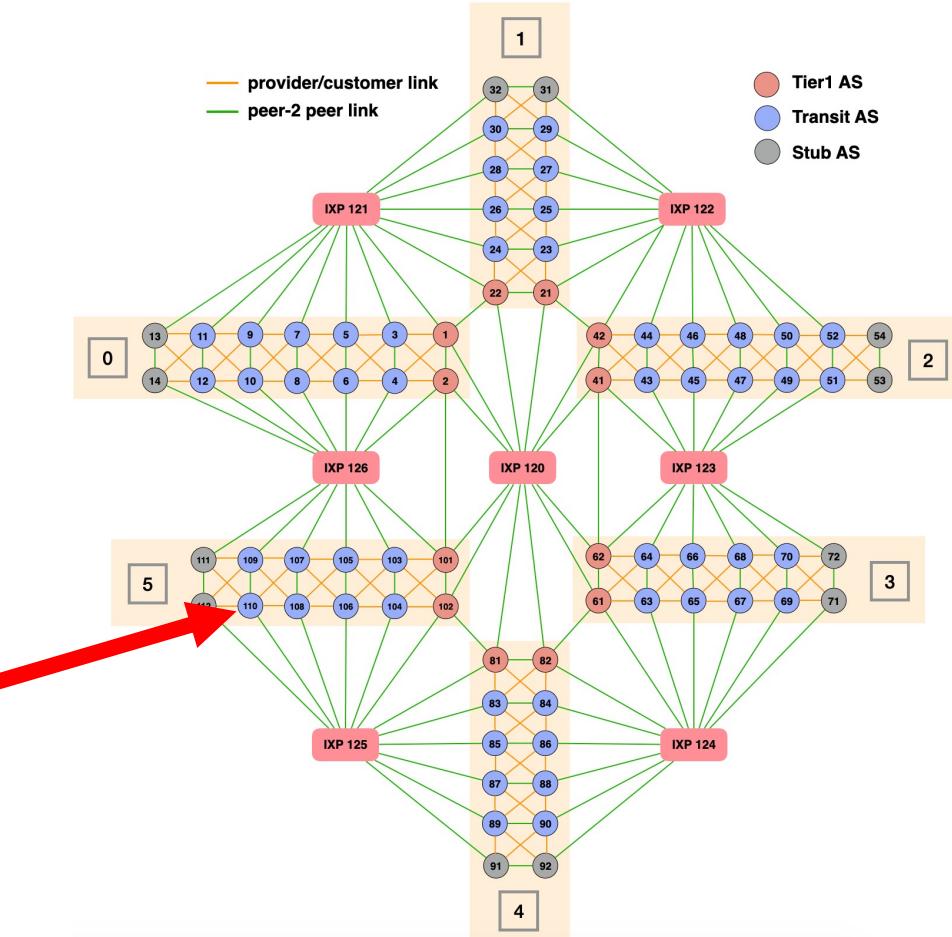
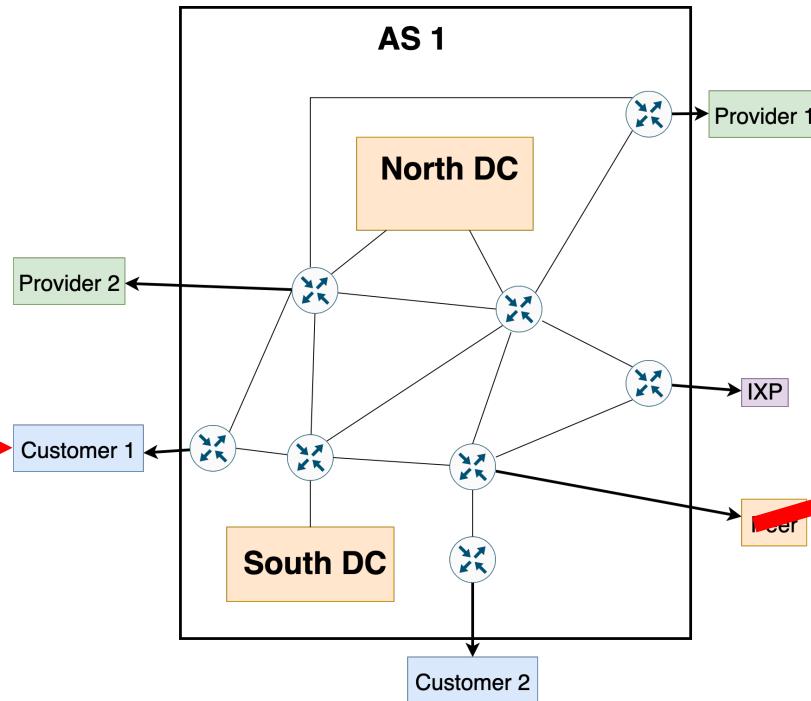
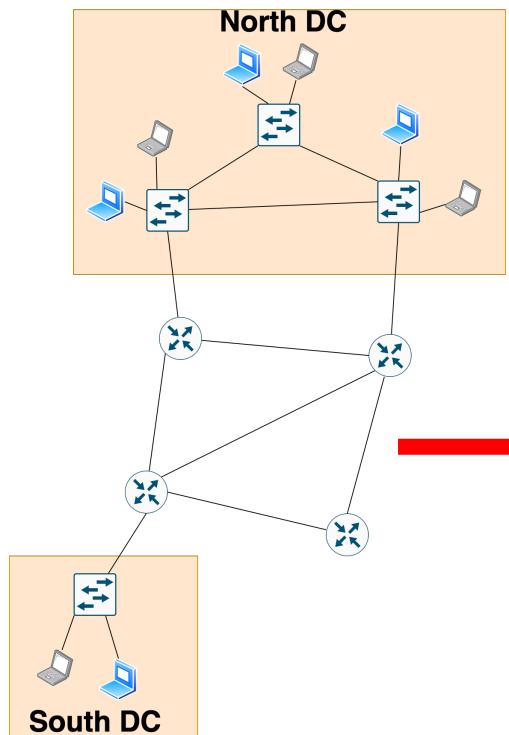
- ミニインターネット:
  - 自分で動かしてみる場合

[https://github.com/nsg-ethz/mini\\_internet\\_project](https://github.com/nsg-ethz/mini_internet_project)

- あるいは、ご相談ください:
  - ご自身のイベント開催をサポート
  - ネットワークテストフレームワーク
  - その他何か:

e-mail: [christoff@ij.ad.jp](mailto:christoff@ij.ad.jp)

# Questions?



e-mail: [christoff@ijj.ad.jp](mailto:christoff@ijj.ad.jp)



日本のインターネットは1992年、IIJとともにはじめました。以来、IIJグループはネットワーク社会の基盤をつくり、技術力でその発展を支えてきました。インターネットの未来を想い、新たなイノベーションに挑戦し続けていく。それは、つねに先駆者としてインターネットの可能性を切り拓いてきたIIJの、これからも変わることのない姿勢です。IIJの真ん中のIはイニシアティブ  
————— IIJはいつもはじまりであり、未来です。

本書には、株式会社インターネットイニシアティブに権利の帰属する秘密情報が含まれています。本書の著作権は、当社に帰属し、日本の著作権法及び国際条約により保護されており、著作権者の事前の書面による許諾がなければ、複製・翻案・公衆送信等できません。本書に掲載されている商品名、会社名等は各会社の商号、商標または登録商標です。文中では™、®マークは表示しておりません。本サービスの仕様、及び本書に記載されている事柄は、将来予告なしに変更することがあります。