Making Network Education More Accessible

Practical learning with the mini-Internet project

2024-06-24

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
- Running your own mini-Internet
- Future of the mini-Internet

How do we traditionally teach how the Internet works?

How do we traditionally teach how the Internet works?

theory

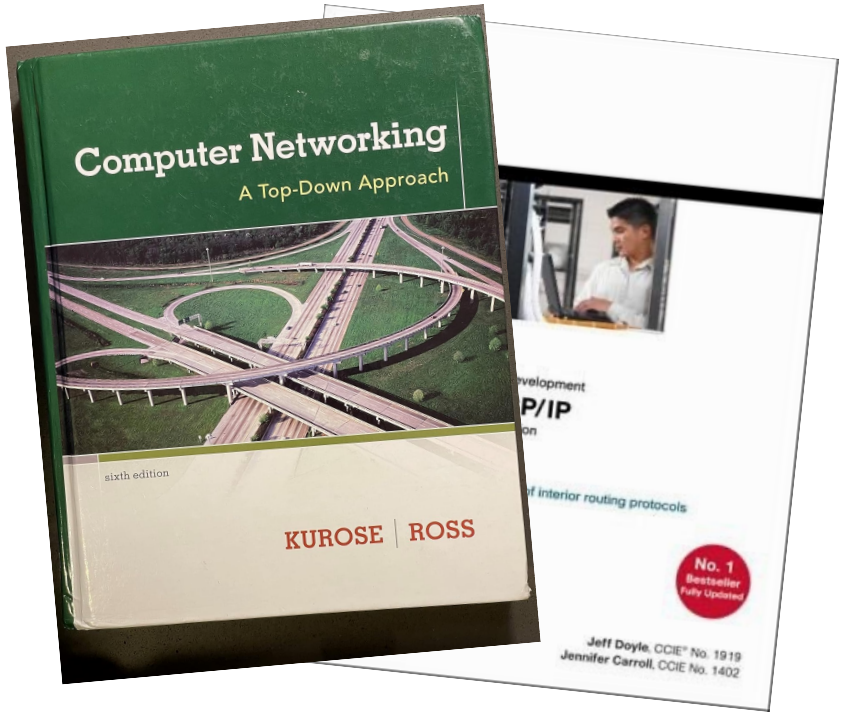

assignments

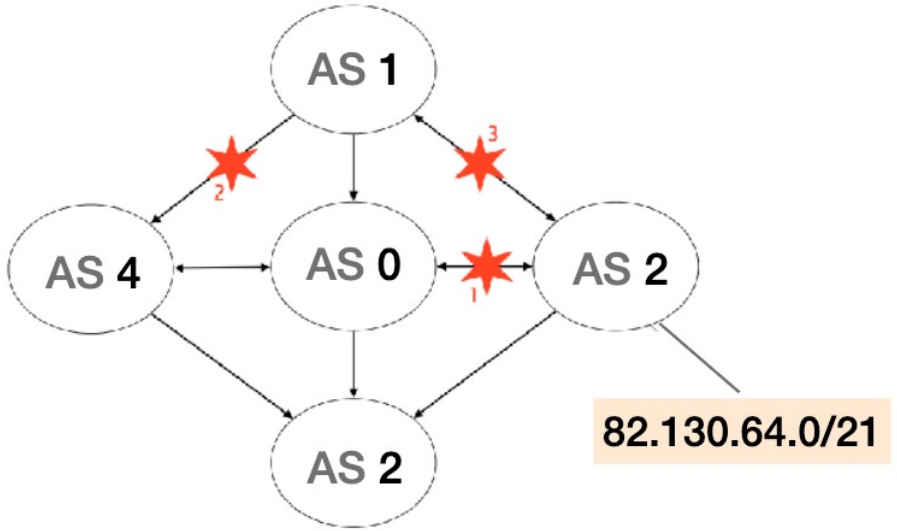

Which messages are exchanged?

Labs

Hosts

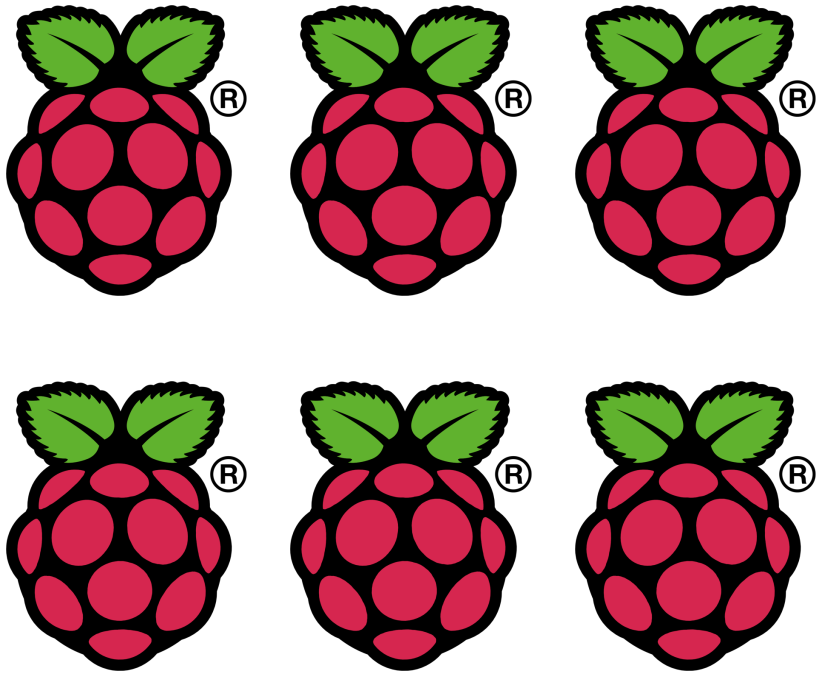

Labs

Hosts

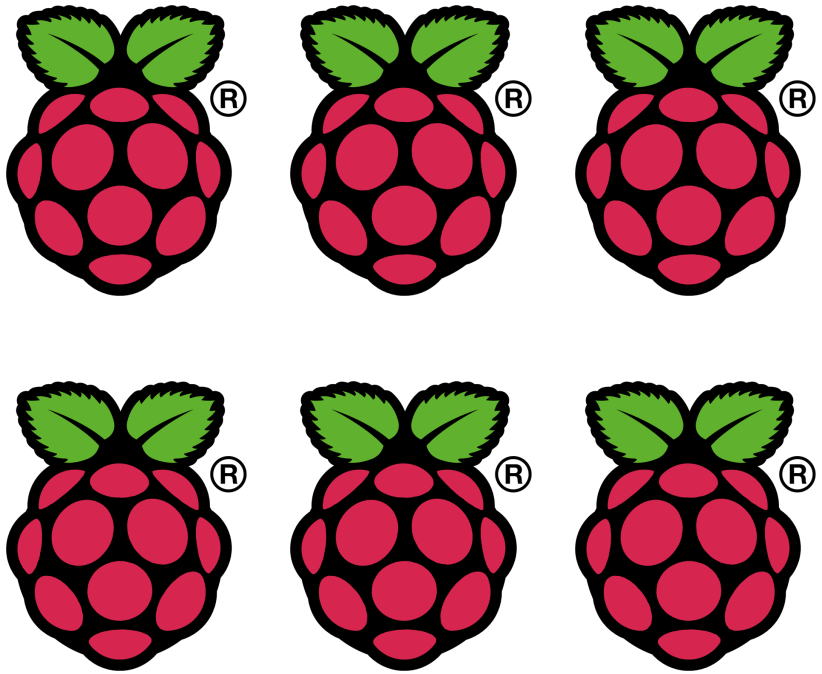

Switches

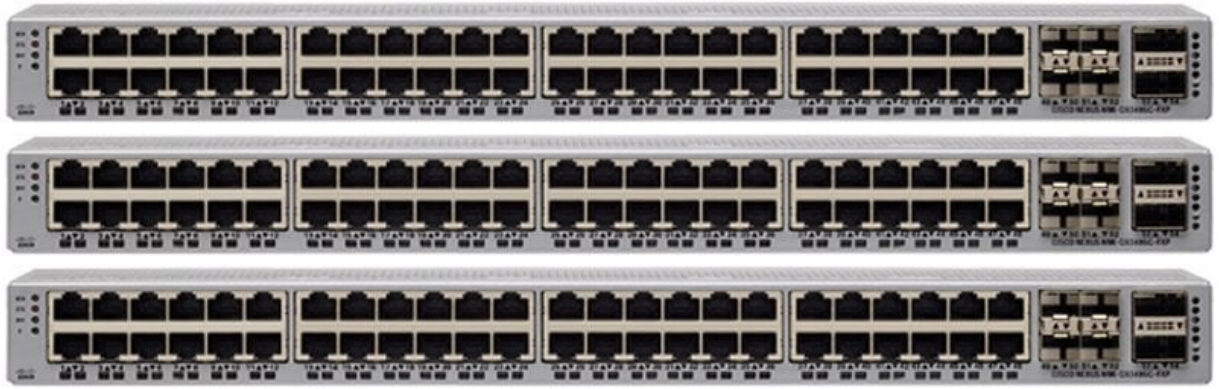

Labs

Routers

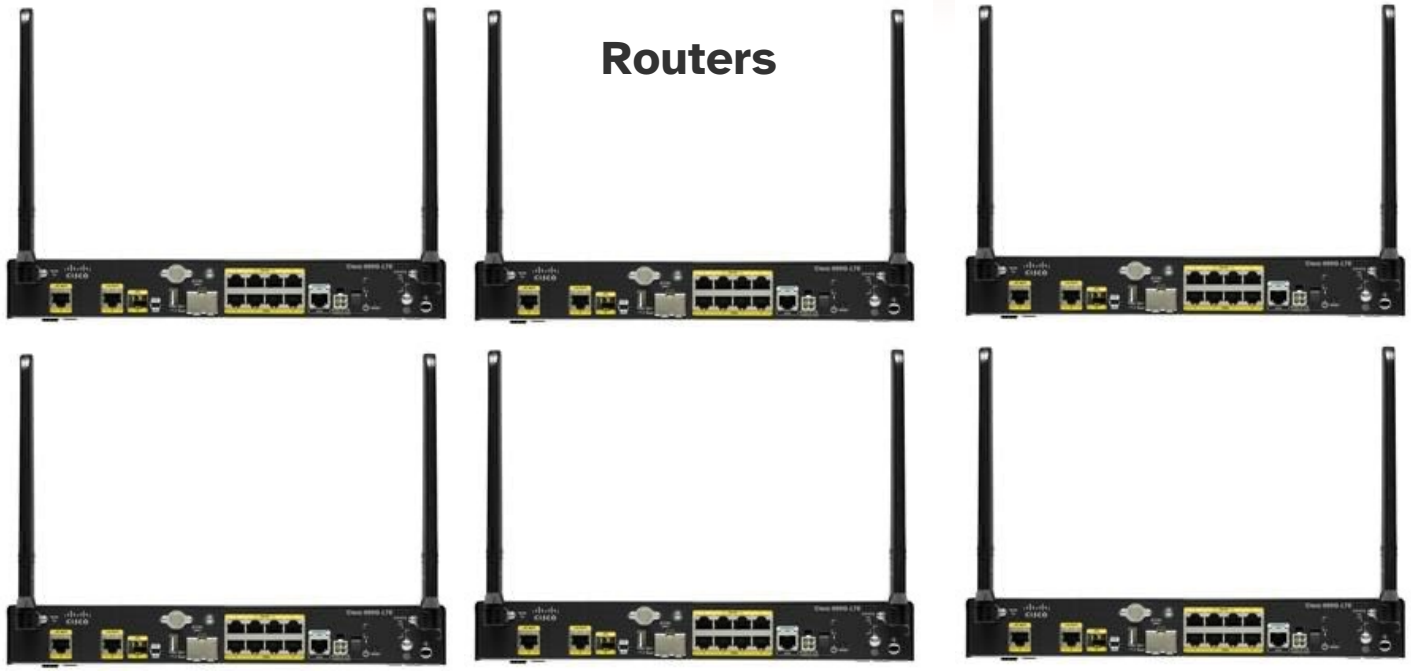

Hosts

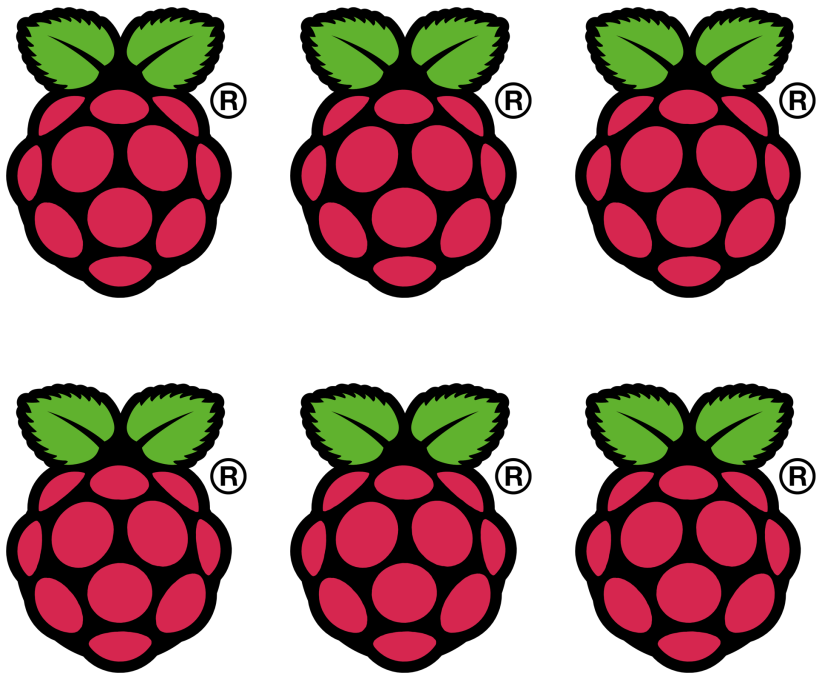

Switches

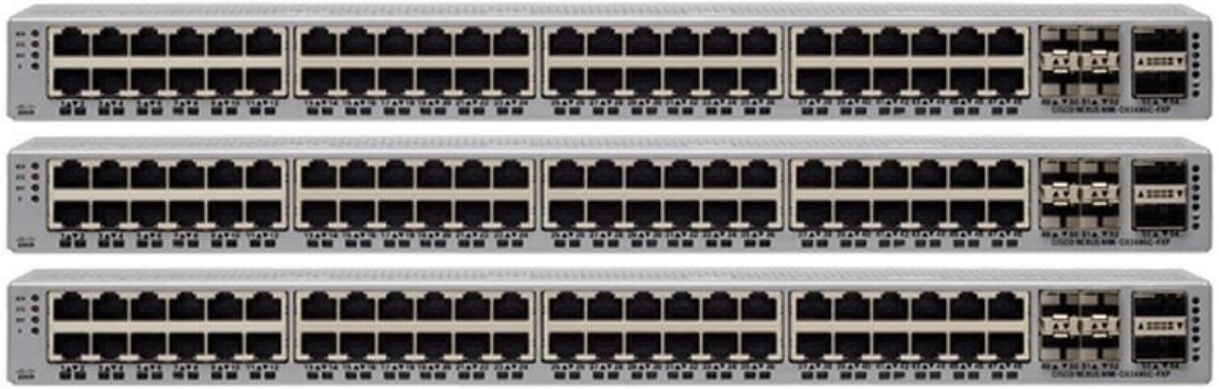

How do we traditionally teach how the Internet works?

Labs

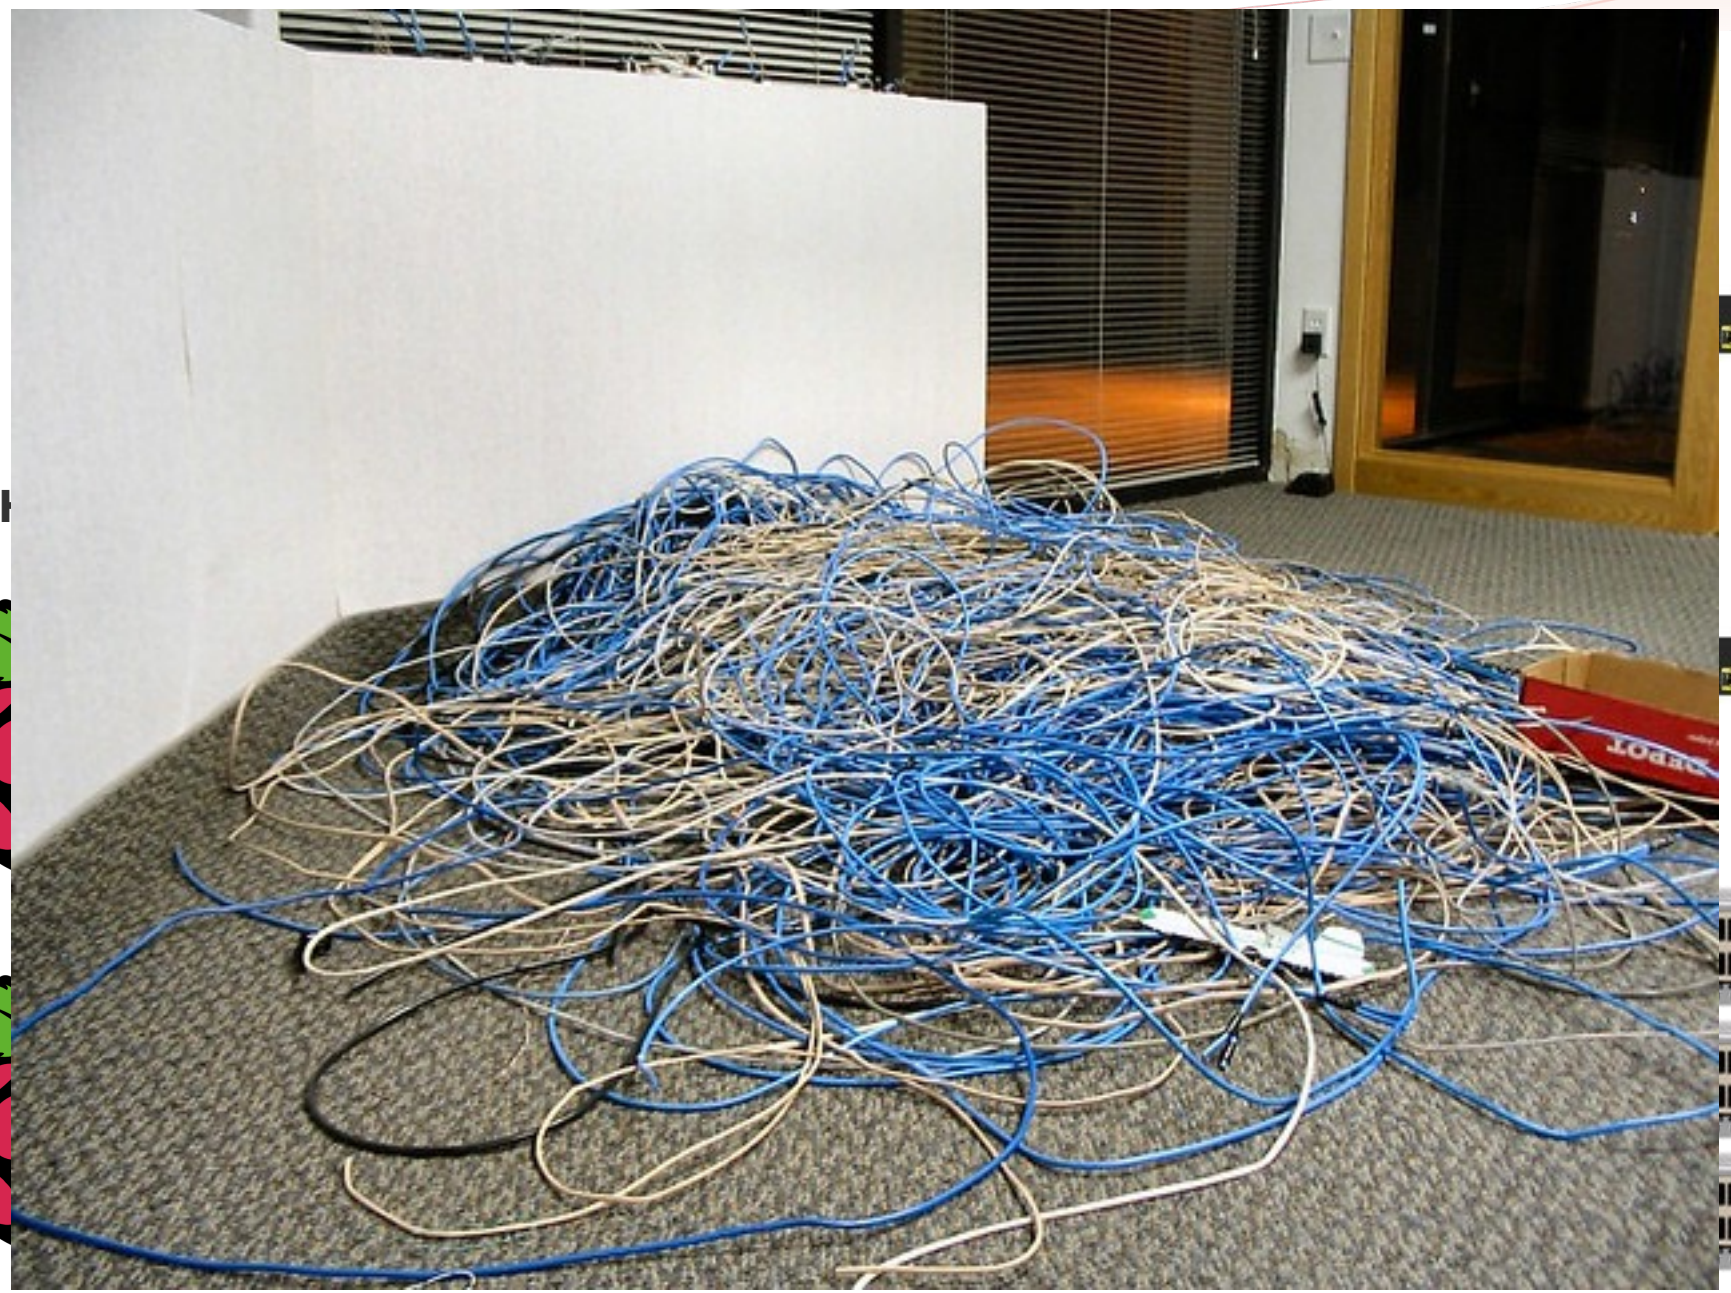

Networking is not just about the technical aspects

- 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
- 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?

Hosts

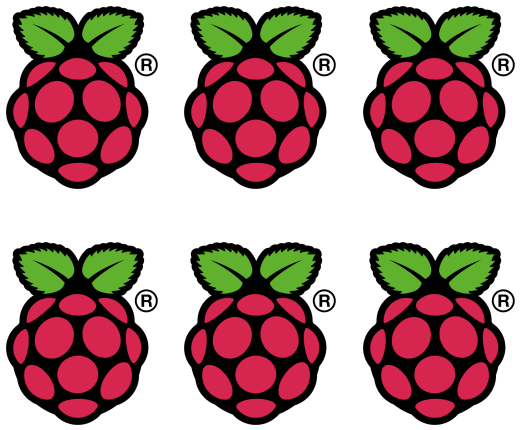

Routers

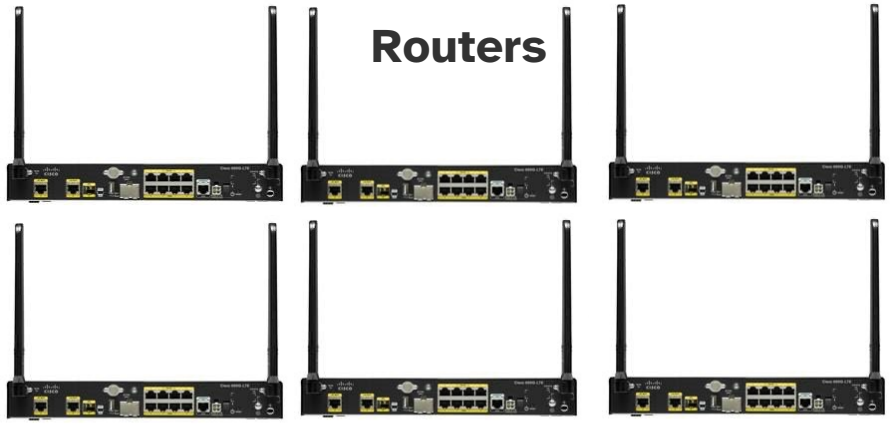

Switches

x 2? 5? 10?

- 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

- 55 AS
  - 5 regions
  - 5 IXPs
  - 6 Student AS per region
  - 4 TA run AS per region

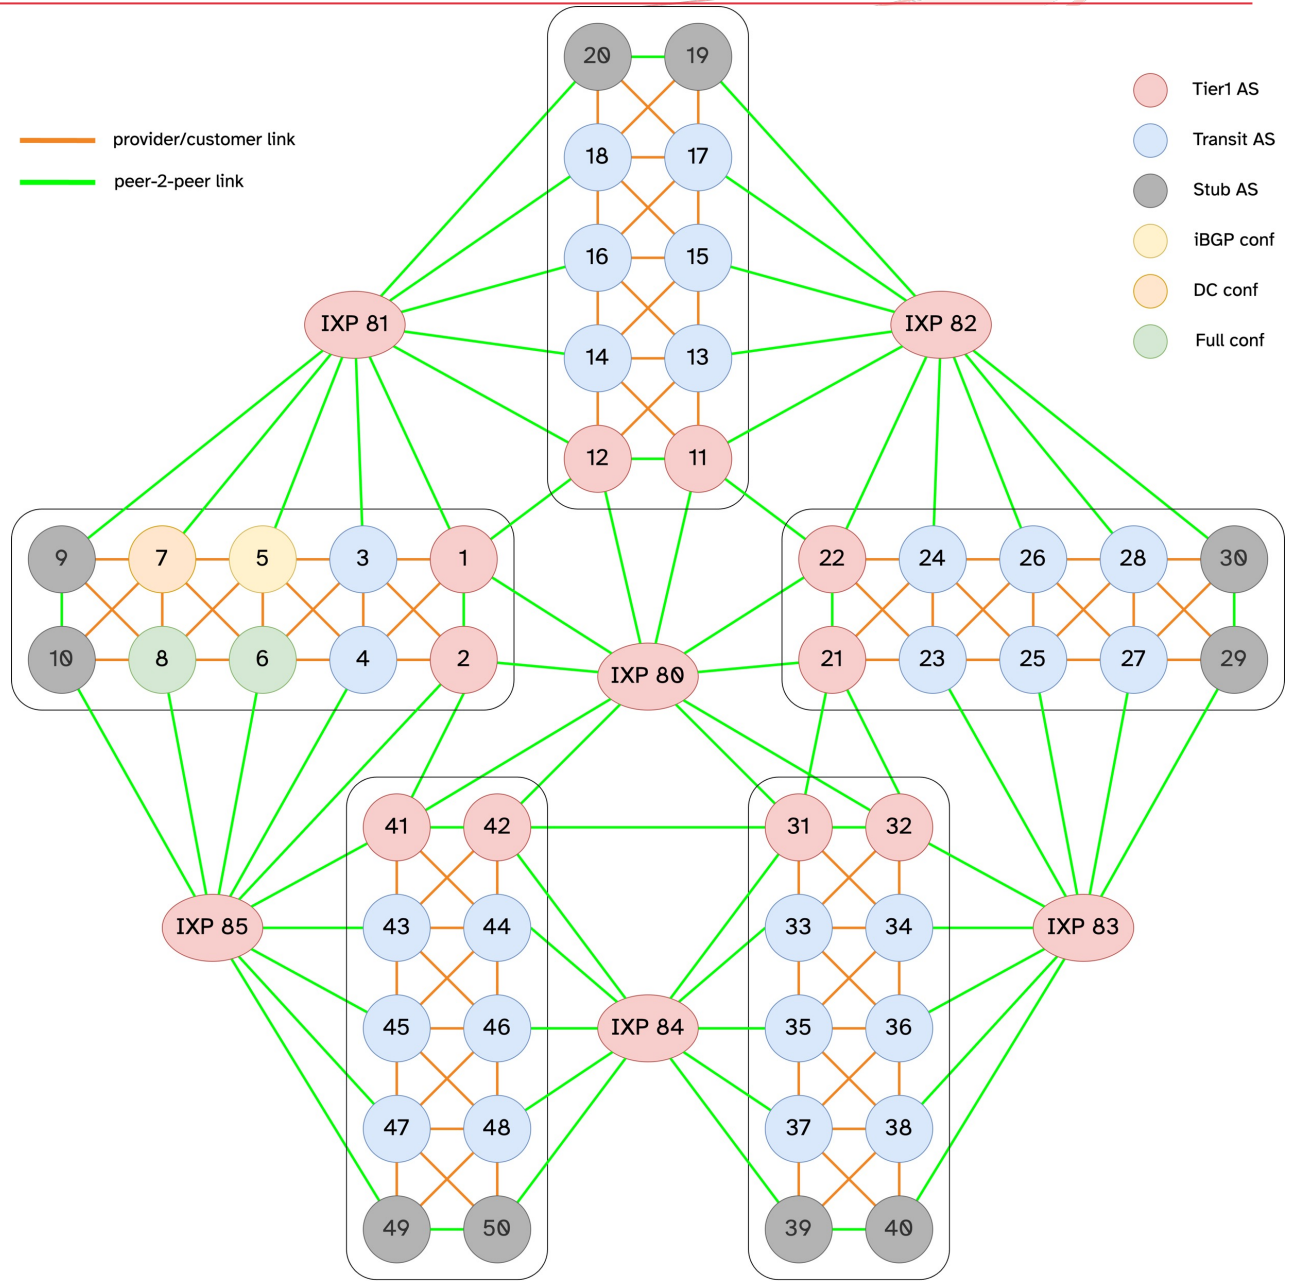

- 2 Data centres
  - North DC
  - South DC
- 8 Routers
- 1 Measurement host
- Shared DNS host

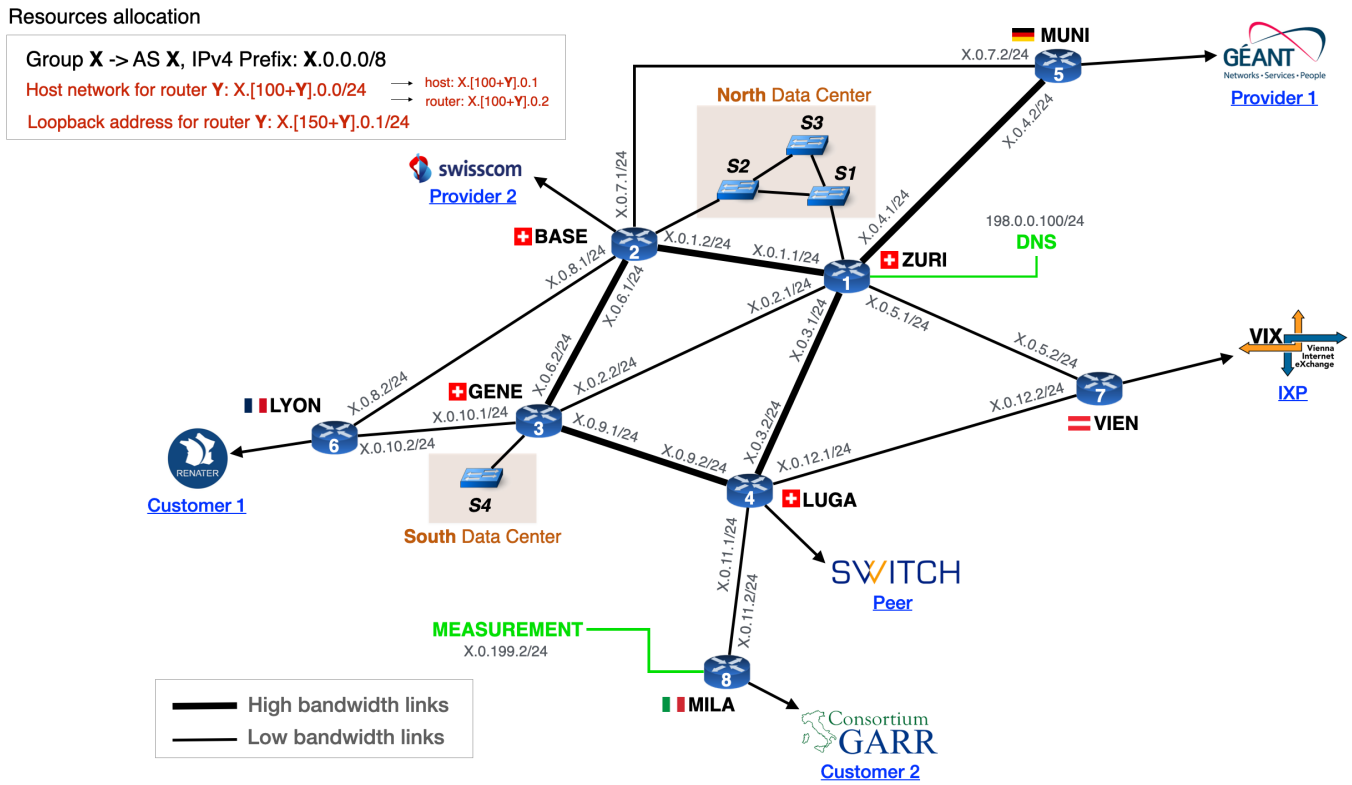

- 2 Hosts per switch
  - FIFA – VLAN 10
  - UEFA – VLAN 20
- South DC v6 only

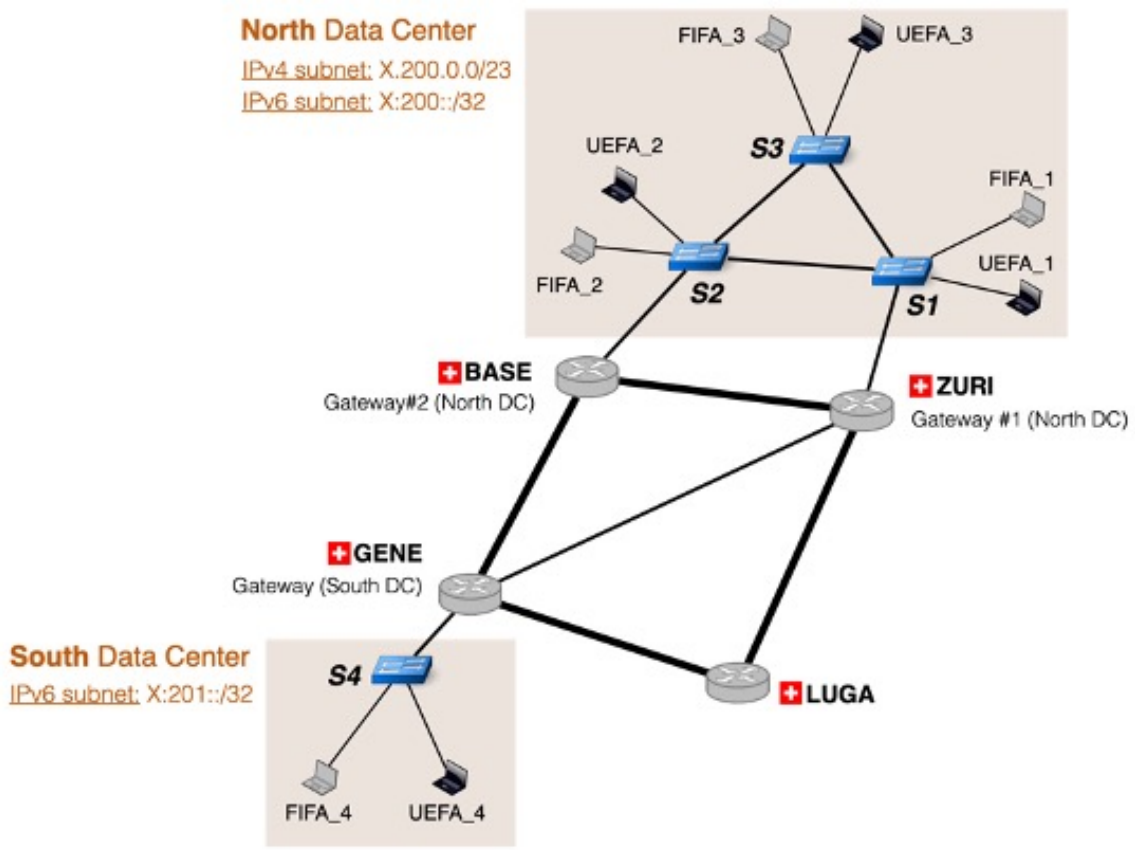

- Runs on a single machine
- Each component is a docker container
- Use common network tools

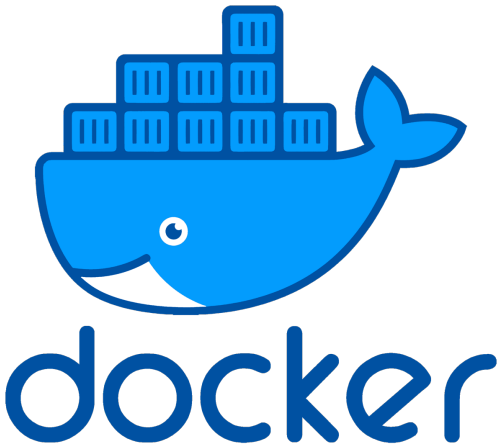

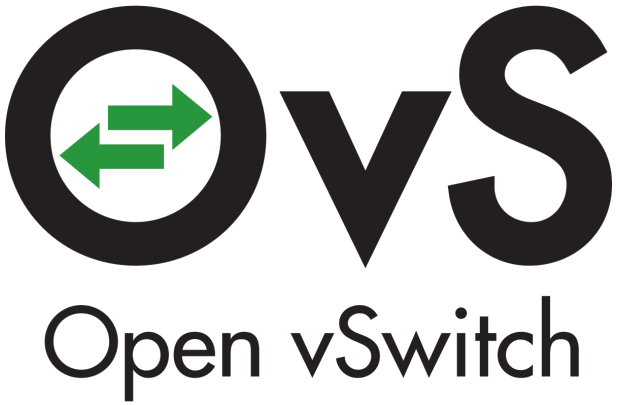

The mini-Internet

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

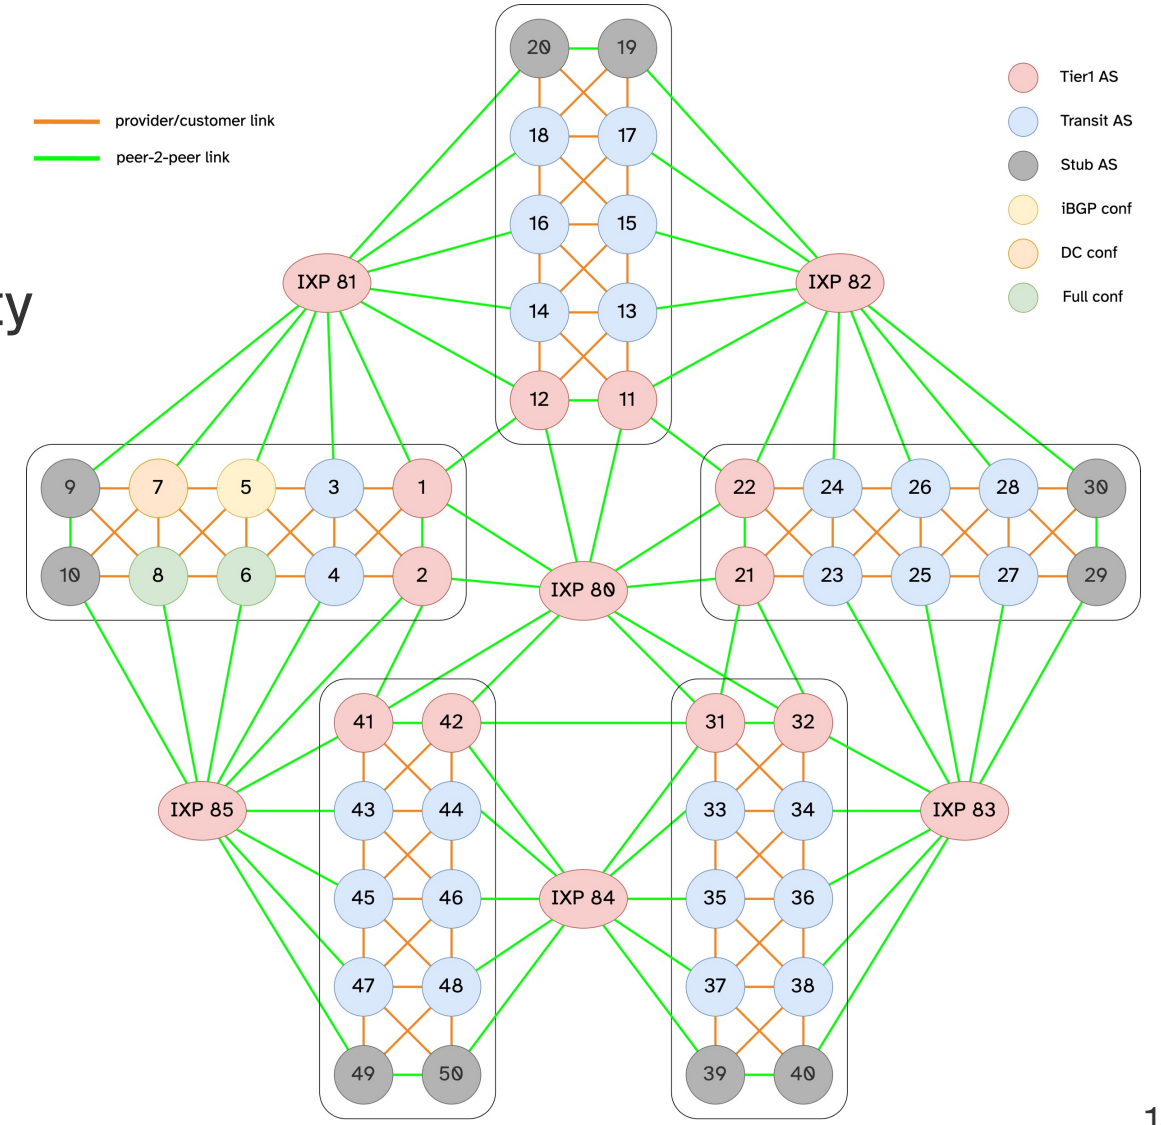

Enable internal connectivity

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

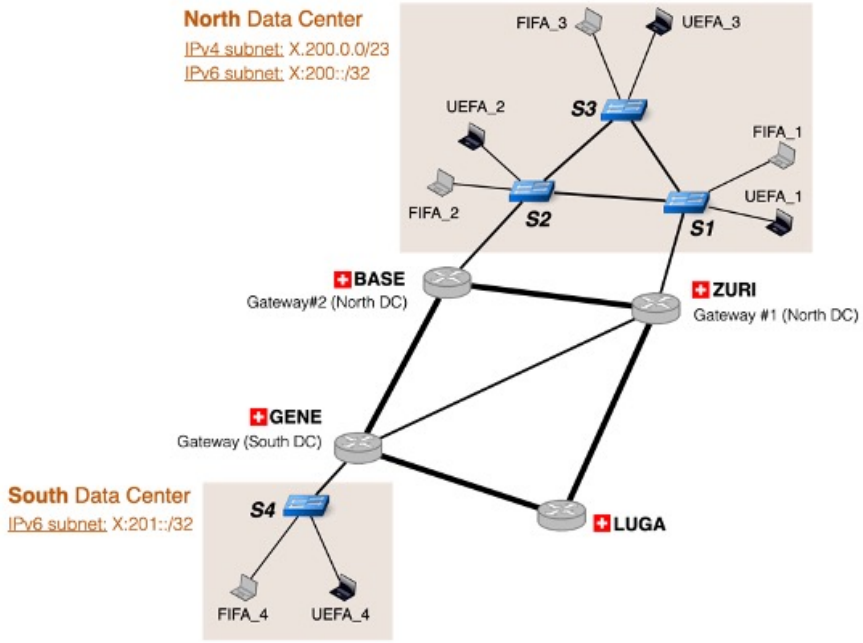

Enable internal connectivity

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

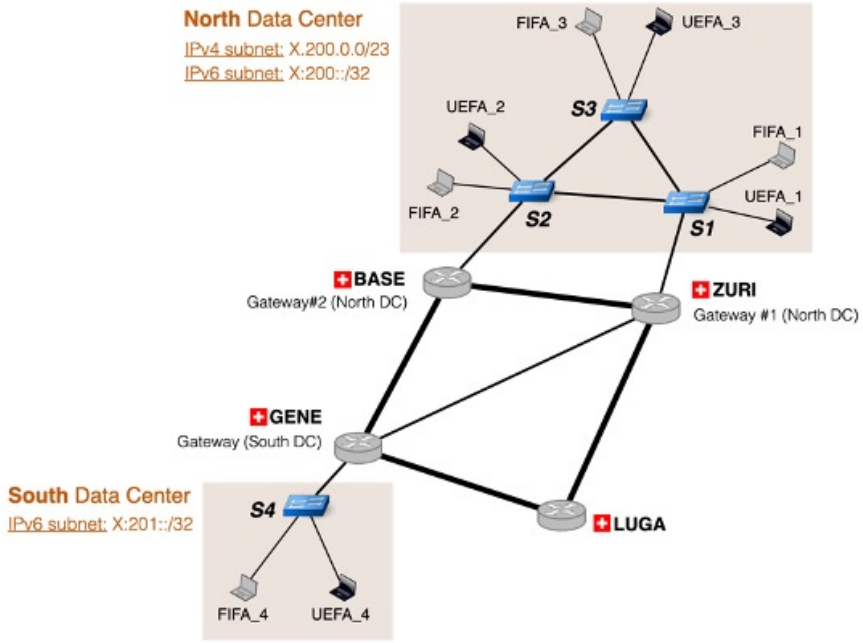

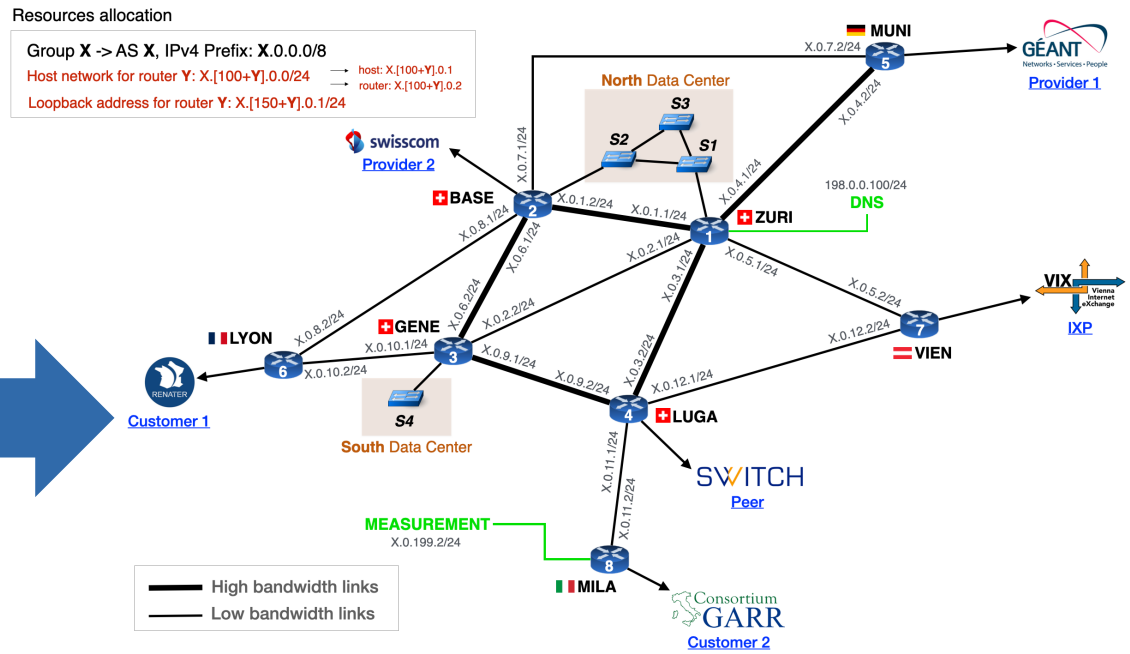

“Peering day”

- 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

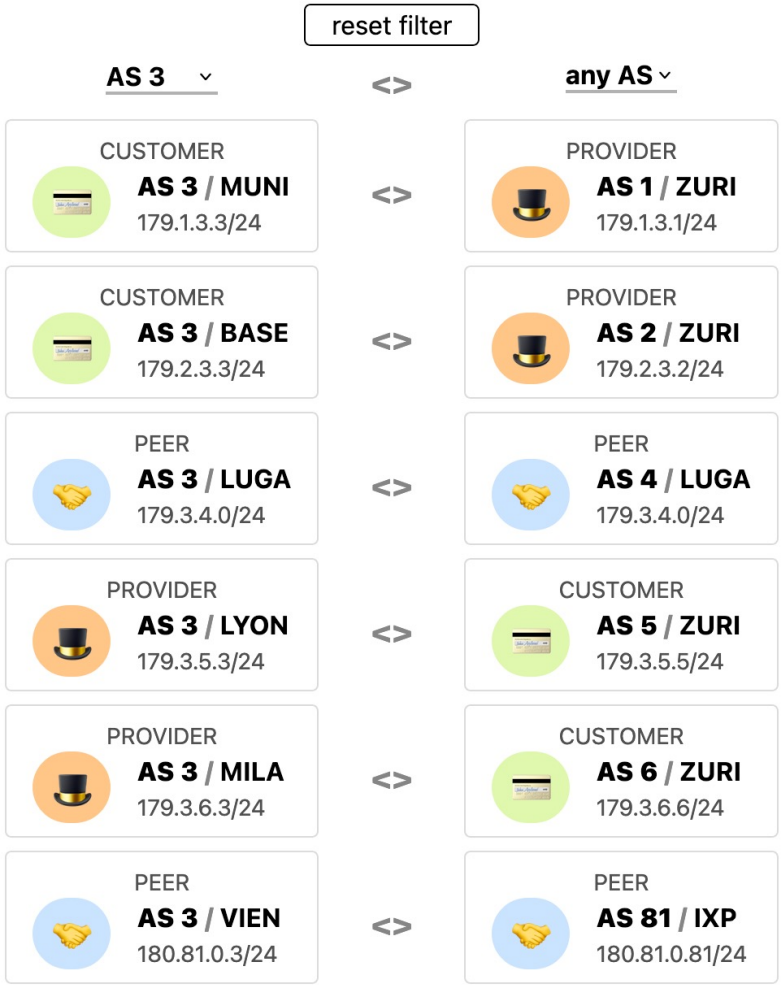

MANRS

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

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 + 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

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
```

reset filter

AS 3	<>	any AS
CUSTOMER **AS 3 / MUNI** 179.1.3.3/24	<>	PROVIDER **AS 1 / ZURI** 179.1.3.1/24
CUSTOMER **AS 3 / BASE** 179.2.3.3/24	<>	PROVIDER **AS 2 / ZURI** 179.2.3.2/24
PEER **AS 3 / LUGA** 179.3.4.0/24	<>	PEER **AS 4 / LUGA** 179.3.4.0/24
PROVIDER **AS 3 / LYON** 179.3.5.3/24	<>	CUSTOMER **AS 5 / ZURI** 179.3.5.5/24
PROVIDER **AS 3 / MILA** 179.3.6.3/24	<>	CUSTOMER **AS 6 / ZURI** 179.3.6.6/24
PEER **AS 3 / VIEN** 180.81.0.3/24	<>	PEER **AS 81 / IXP** 180.81.0.81/24

routing project [*] matrix [ ] looking glass [ ] connections [ ] krill

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 ( );
- can reach with an invalid AS-level path ( );
- cannot reach ( ).

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.

96% 4% 0%

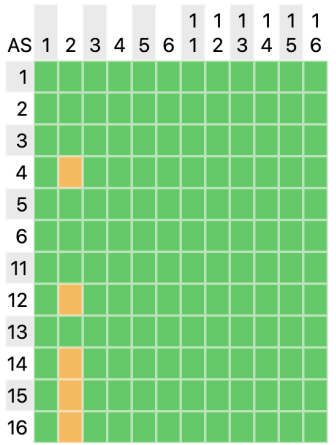

The Mini-Internet use cases

- 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

The end goal: A mini-Internet for all

The end goal: A mini-Internet for all

Content

The end goal: A mini-Internet for all

Content

Platform

The end goal: A mini-Internet for all

Content

- More content
  - IPv6 transition
  - RPKI deployment
  - Network management
  - Setting up an IXP
  - Etc.
- Increasing accessibility
  - Applicable to local community
  - Translation

Platform

The end goal: A mini-Internet for all

Content

- More content
  - IPv6 transition
  - RPKI deployment
  - Network management
  - Setting up an IXP
  - Etc.
- Increasing accessibility
  - Applicable to local community
  - Translation

Platform

- More topologies
- More specialized nodes
  - They are all docker nodes
  - Kubernetes?
- Containerlab
- Physical hardware
- Etc.

Train Trainers That Train Trainers to Train…

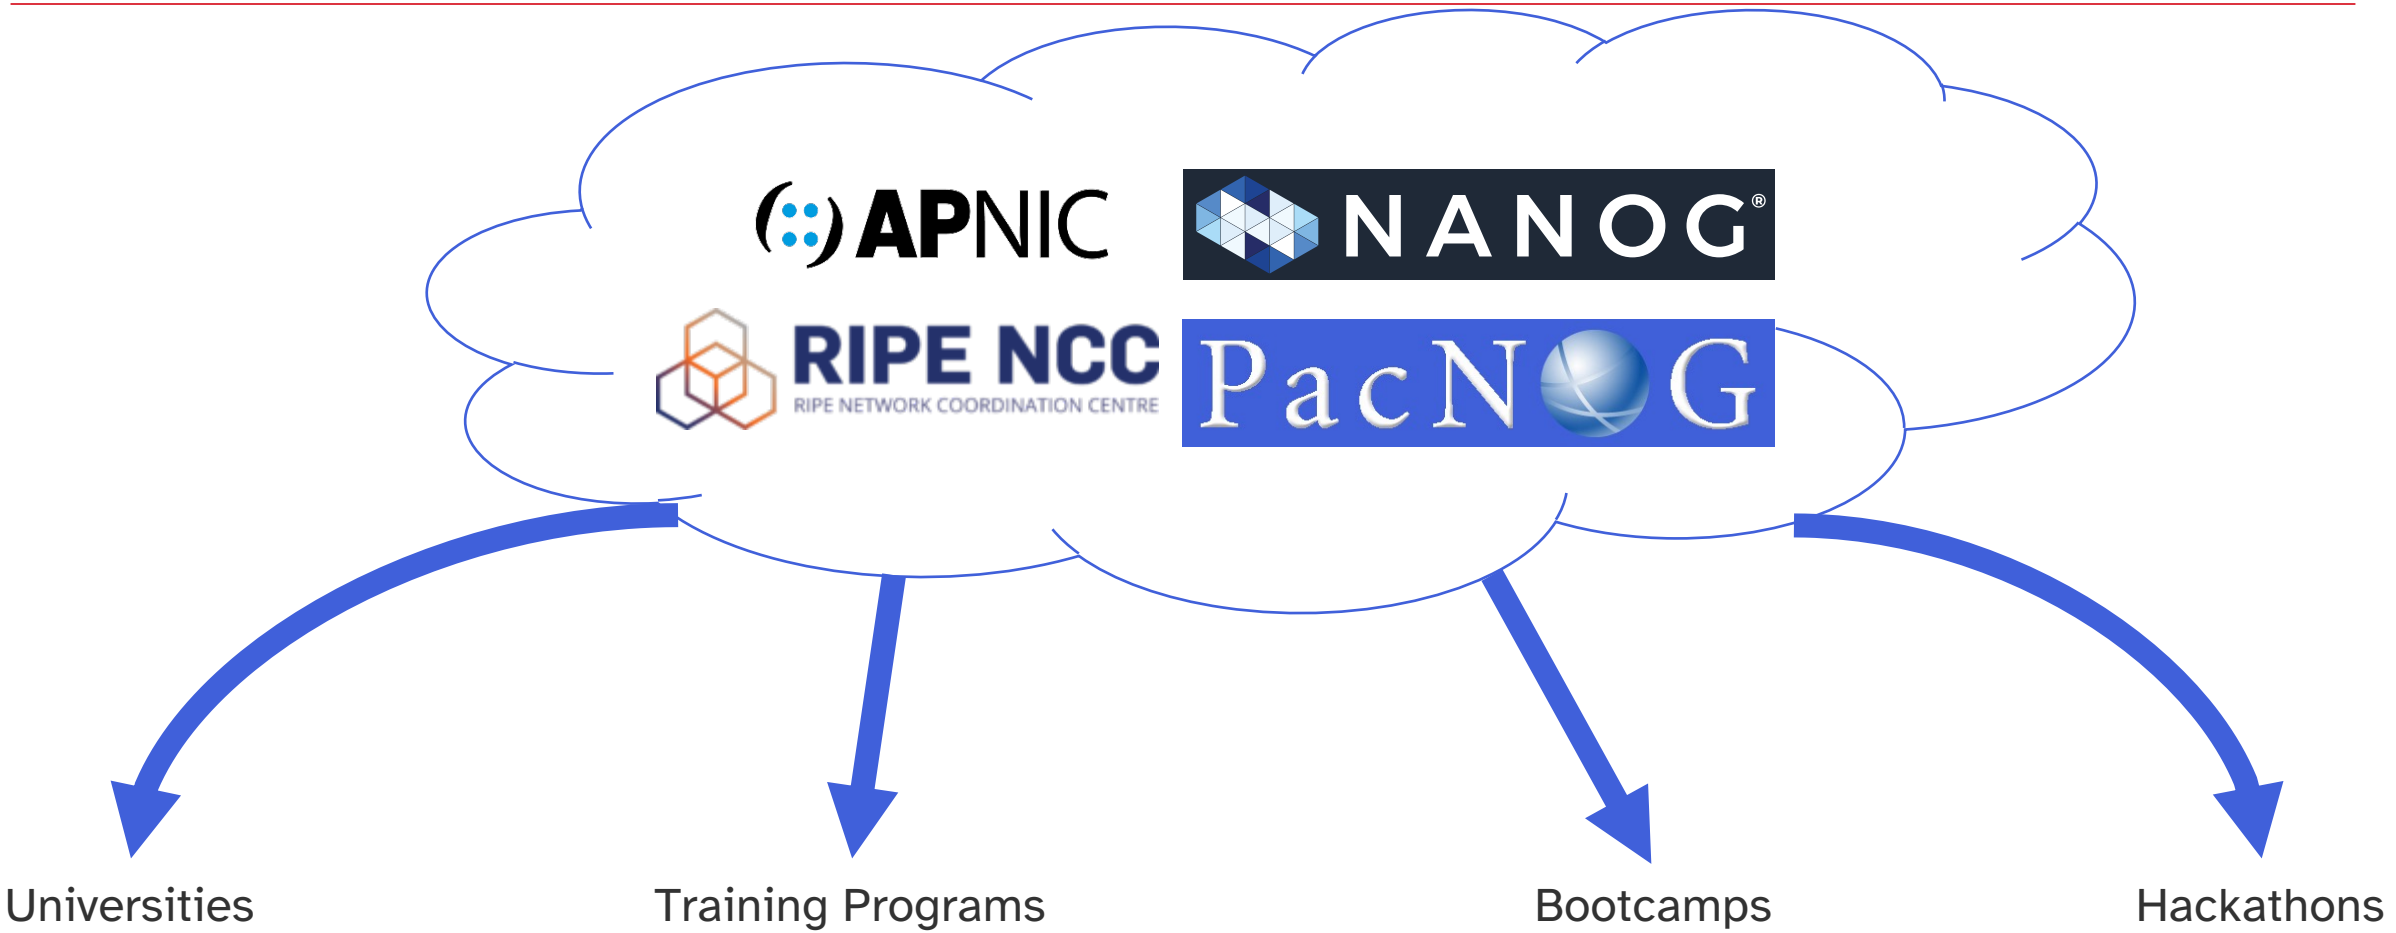

Running your own Mini-Internet

Run your own

- Clone the repo
  https://github.com/nsg-ethz/mini_internet_project

- Follow the wiki
  https://github.com/nsg-ethz/mini_internet_project/wiki
  - Has an example topology to play with

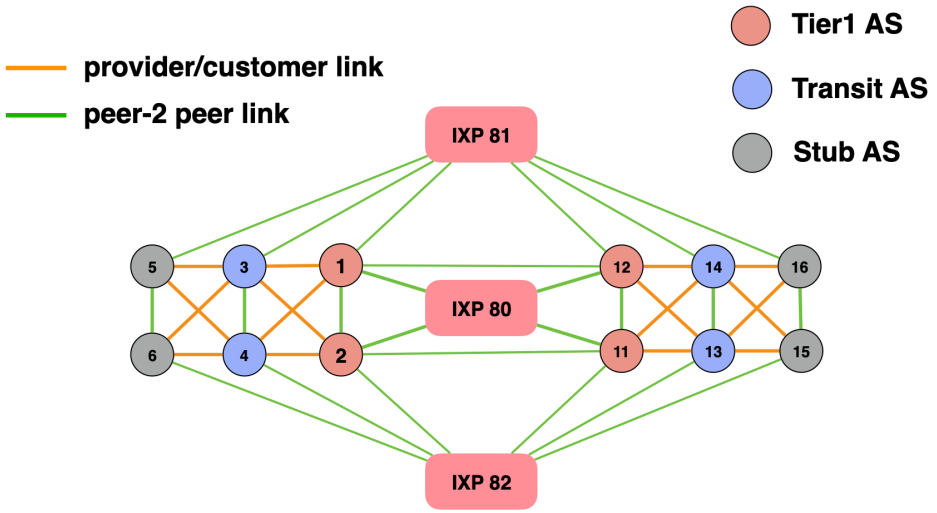

Run your own

- Clone the repo
  https://github.com/nsg-ethz/mini_internet_project

- Follow the wiki
  https://github.com/nsg-ethz/mini_internet_project/wiki
  - Has an example topology to play with
  - Or define your own
  - Make sure to follow the prerequisites

Run your own

- Clone the repo
  https://github.com/nsg-ethz/mini_internet_project

- Follow the wiki
  https://github.com/nsg-ethz/mini_internet_project/wiki
  - Has an example topology to play with
  - Or define your own
  - Make sure to follow the prerequisites

- Sample course to follow:

https://pacnog33.iijlab.net/

Here is one we we prepared earlier…

Interactive demo

Topology

- Tier 1 + Stub as
  - Preconfigured/unmanaged
- AS 7
  - Configured on a DC level
- AS 5
  - Configured up to iBGP stage
- AS 8 + 6
  - Fully configured
- AS 3 – Demo (me)
- Unconfigured
  - Y3, Y4, Y5, Y6, Y7, Y8
  - Replace Y with region #
- 2~3 per AS

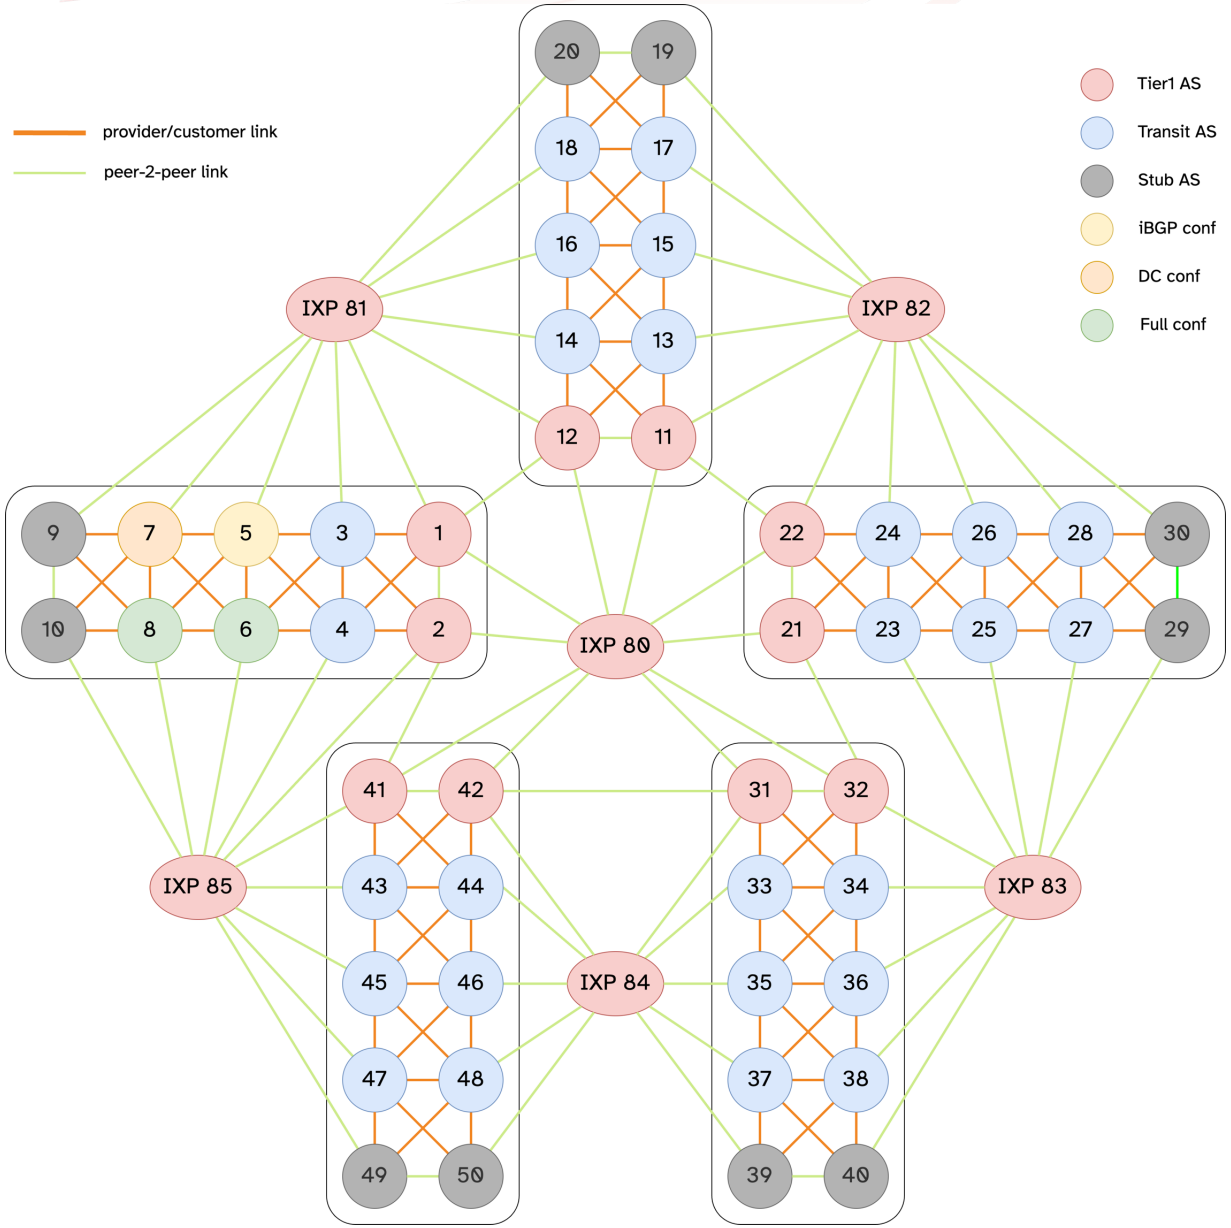

Access

- https://pacnog33.iijlab.net
- Tutorial/Instructions
- Monitor
- SSH
  - Select your AS
  - User: root
  - Pass: pacnog-ASX
  - Replace X with AS #
- Use the helper script e.g.:
  - Router - ./goto.sh LUGA router
  - Switch - ./goto.sh DCN S2
  - Host - ./goto.sh DCN FIFA_1

Step 1: DC connectivity

- Configure layer-2 network
  - X.200.0.0/23
- Ensure VLAN separation between FIFA and UEFA hosts
- Note:
  - Switches are OvS switches

Step 2: OSPF and iBGP

- Configure router IP addresses
- Configure iBGP

Step 3: Routing

- Configure eBGP sessions
- Connect to an IXP

The future of the Mini-Internet

And my work interest

- Mimics the “real” Internet
- Safe environment for experimentation
- Ideal for:
  - Testing
  - Hackathons
  - Exploration

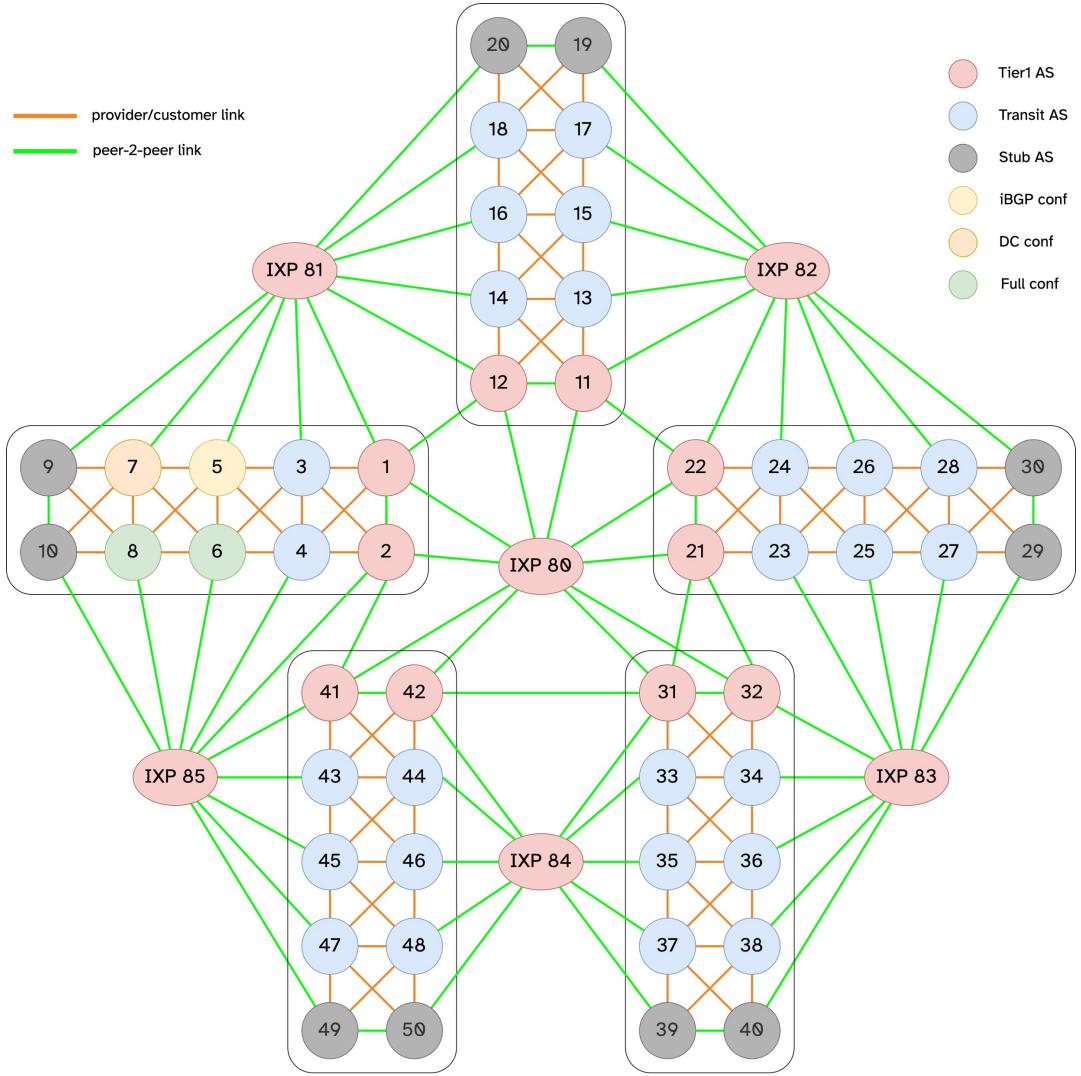

- Split into smaller modules
- More specialised
  - Courses
  - Experiments and testing
- Current design is very monolithic

Simplify:

- Setting up and configuring test bed
- Digital twinning

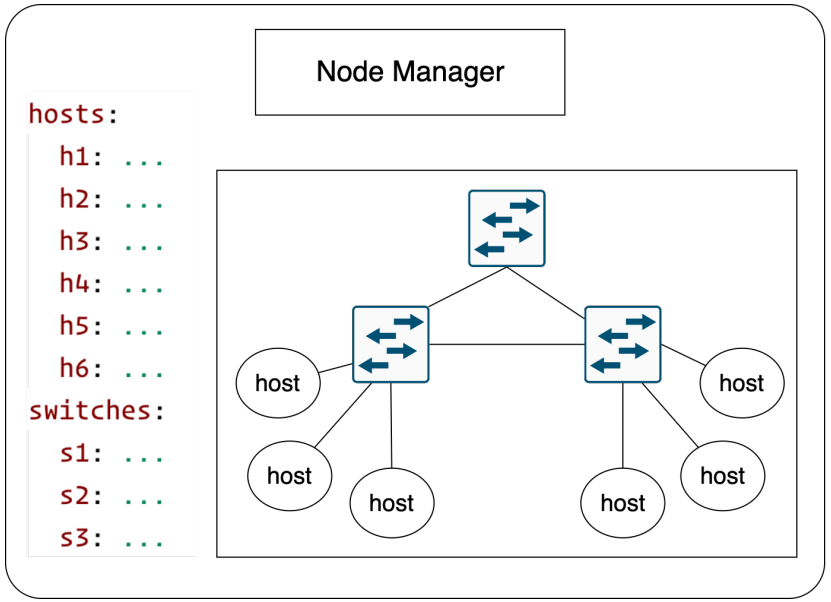

Simplify:

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

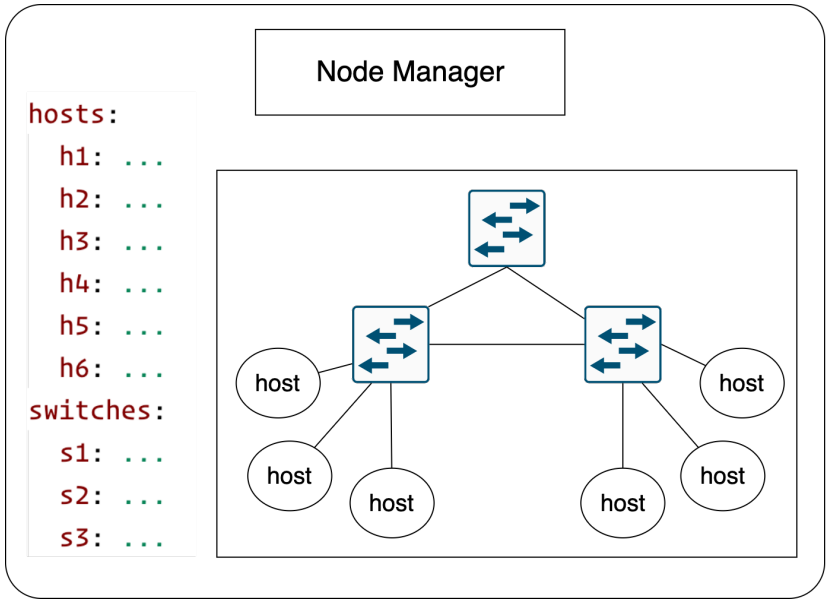

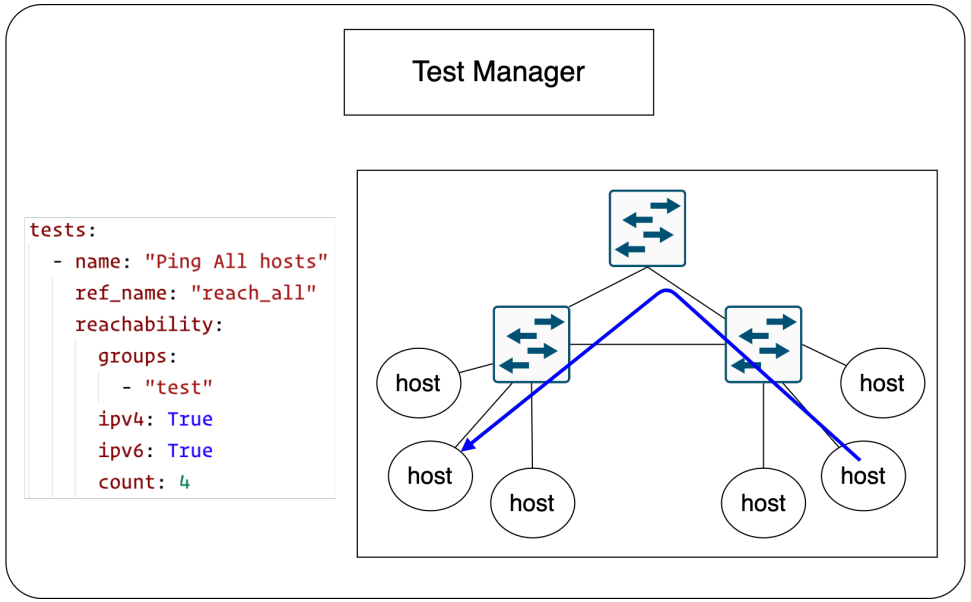

Simplify:

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

Node Manager

```
hosts:
  h1: ...
  h2: ...
  h3: ...
  h4: ...
  h5: ...
  h6: ...
switches:
  s1: ...
  s2: ...
  s3: ...
```

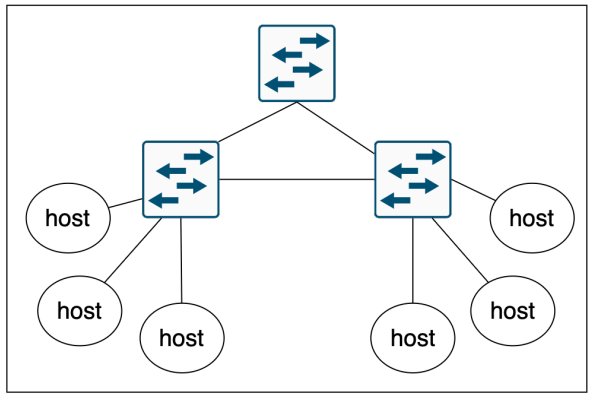

Test Manager

```
tests:
  - name: "Ping All hosts"
    ref_name: "reach_all"
    reachability:
      groups:
        - "test"
      ipv4: True
      ipv6: True
      count: 4
```

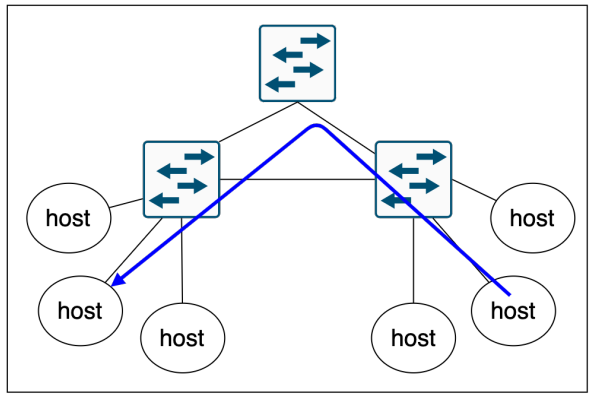

Verification Manager

```
verify:
  - name: "No packets lost"
    reachability:
      packets_lost: 0
```

No packets lost: ✅

- 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
- How to run your own mini-Internet
- Network Testing Framework
  - Simplify configuring, testing and verification

Questions?

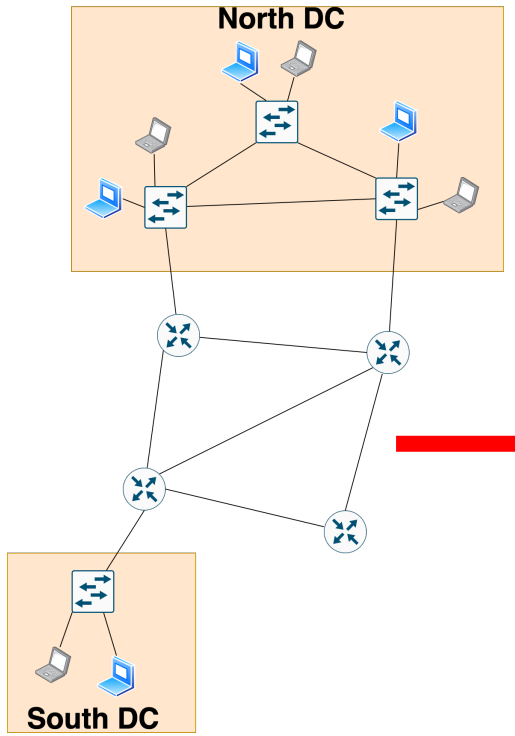

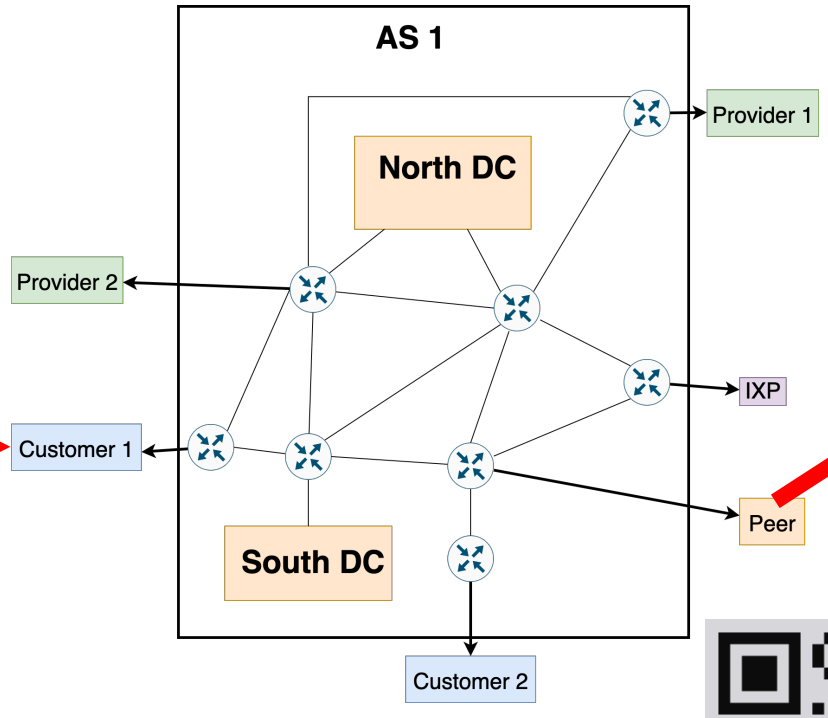

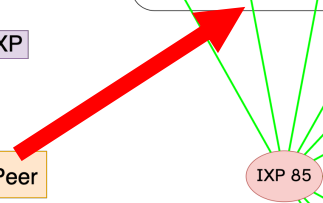

e-mail: christoff@iij.ad.jp

日本のインターネットは1992年、IIJとともにはじまりました。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。