



LabIX

Creating an Internet Exchange in Your Local Hackerspace

Emil Petersen (AS211153)


2023-08-17

Chaos Communication Camp 2023

ix.labitat.dk

Why am I making this talk?

- Difficult to find information on IXP operation
 - Everything is an acronym soup
- Difficult to understand the administration of IP allocation
 - What policies is needed?




What is Labitat Internet Exchange (LabIX)?



- A community run internet exchange in Labitat, Copenhagen
- A place to learn network routing, including:
 - BGP
 - Route Server ([AS60247](#))
 - ASN and IP allocation
 - BGP security
 - RPKI, filtering, GTSM
 - 9000 MTU (very soon)
 - Expired RFC draft: [Jumbo Frame Deployment at IXPs](#)

What does your hackerspace need for an IX?

- Friendly uplink connection
 - Peering policies and filters needs to be updated
- Server housing
 - A layer 2 switch
 - Pre-used rack or similar
- At least three peering clients
 - Requirement from RIPE, PeeringDB and IXPDB
- Interested parties
 - Hackerspace projects tend to die over time



What is an Internet Exchange?



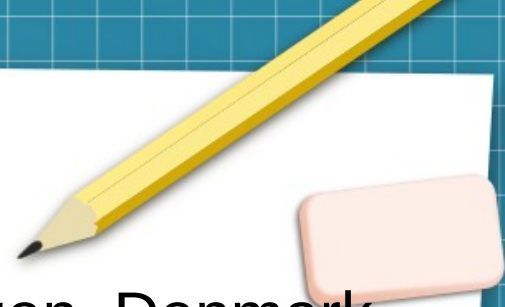
- Very simplified:
 - A VLAN on a layer 2 switch, with a bunch of connected network operators
- Its more than that, so lets first explain the internet

How do you do internet routing?



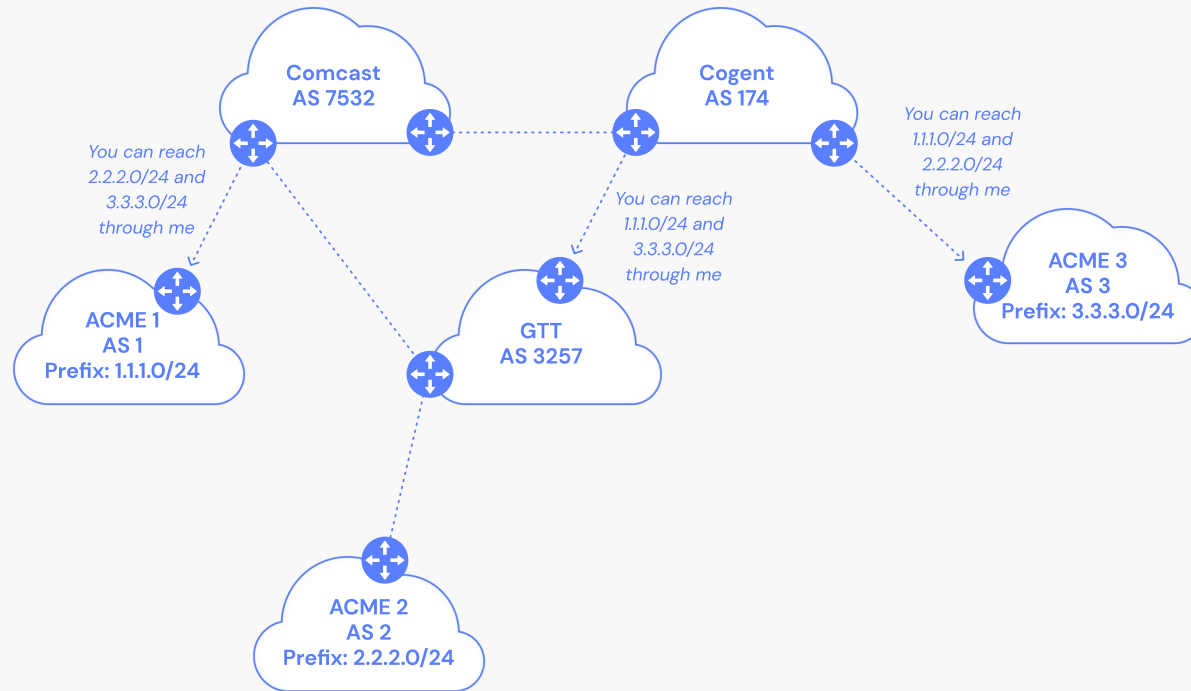
- 1) Get IP prefixes allocated to you
- 2) Find hosting facility
- 3) Install machine
- 4) Setup an eBGP routing daemon
- 5) Find peers and peer with them
- 6) Find a transit uplink for connecting to everyone else
- 7) Profit?

Let us set a senario




- You are a (eyeball) network provider in Copenhagen, Denmark
 - Meaning: Customers consuming content with their eyeballs
- You need to get content from content providers
- You need IP addresses for your customers and servers
- You need server housing and physical connections (fiber).

Networks connect to networks



Autonomous System



- RFC 1930 describes an AS as
 - “a connected group of one or more IP prefixes run by one or more network operators which has a SINGLE and CLEARLY DEFINED routing policy”.
- Each AS is assigned an unique 16- or 32-bit Autonomous System Number (ASN)
 - Used to identify organizations and their prefixes

So what is a routing policy?




- These are called Internet Routing Registry (IRR) objects
 - Found in the RIR databases like RIPE DB
- Example:
 - These IP prefixes should be destined to me
 - AND I accept prefixes destined to- and from my business customers AS 65536 and AS 65537.

Read more:

- RFC 2622: Routing Policy Specification Language (RPSL)
- RIPE NCC. [“BGP Security: IRR and Filtering”](#). Feb. 15, 2021

Peering and transit



- Peering
 - Free (but includes a setup cost for things like cross-connects and optics)
 - “We only send traffic destined to each other”
- Transit
 - €-cost / Mbit
 - “I will transport traffic from you to other networks”

The Internet Exchange

- Keep local traffic local
 - Avoid the trambone effect
- Lower latency
- Easier to interconnect with other networks
 - Recuring cost for data-centre cross-connects

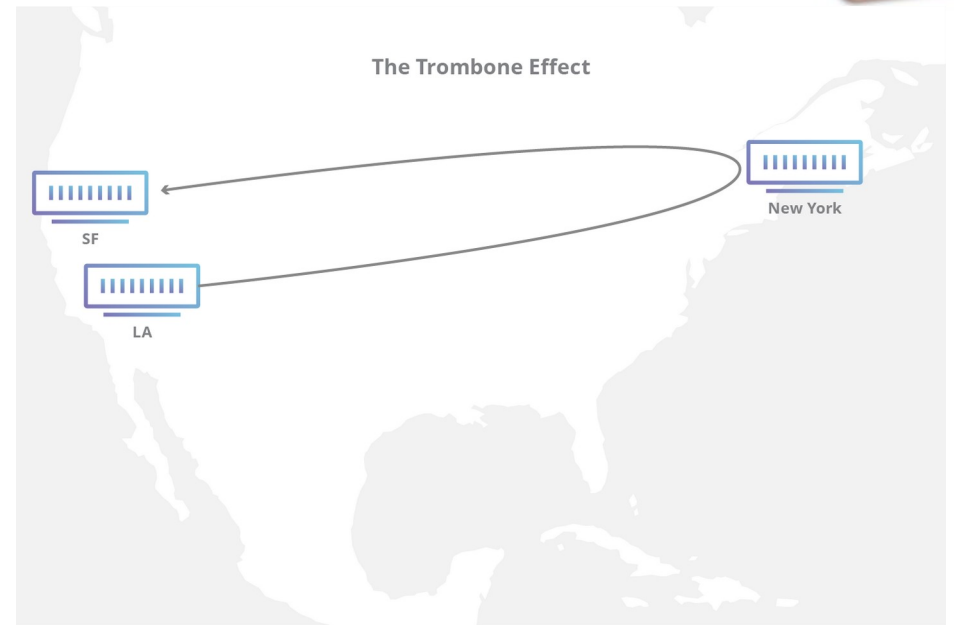
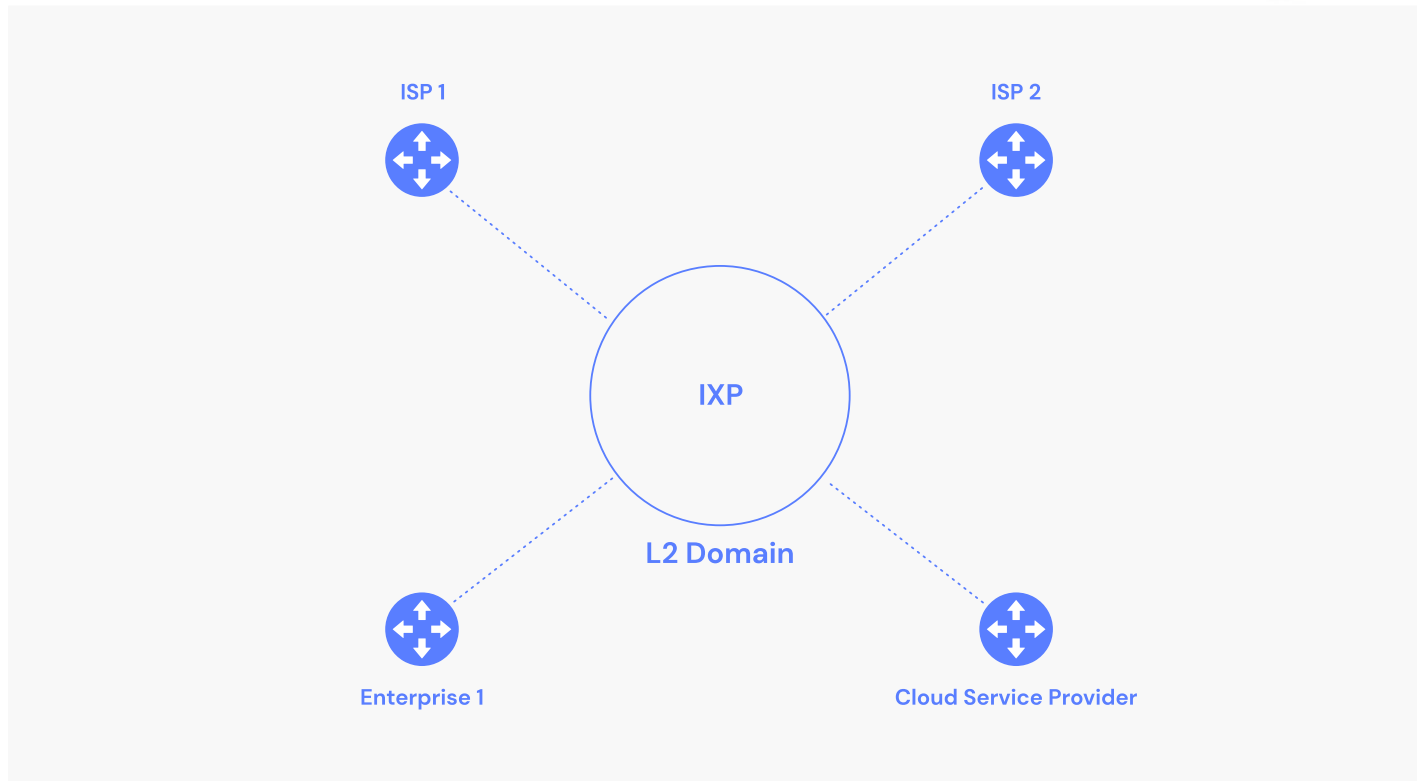


Image source: [Cloudflare](#)

Simplified IXP topology



What is a Route Server?

- Peering interconnections changes
 - From: Many-to-Many (Bi-lateral)
 - To: One-to-Many (Multi-Lateral)
- Control-plane traffic is aggregated by the route server





Figure 1: IXP peering options.
Richter, et al, 2014

Advantages of route servers

- Low maintenance aggregation point sessions
- Immediate value for newcomers.
- Debugging tools to have a sense what's going on at the IX




Job Snijders, [LACNIC29](#), image by [bknix.co.th](#)

Our route server

- Hardware: Wyse 3030 LT
 - Power usage: 6 Watts
 - Free and available at hand
- Software
 - Debian as operating system
 - [Arouteserver](#) as config parser
 - [BIRD](#) as BGP speaker



Server Hosting




- Machine can be colocated in LabiColo, in Labitat
 - [Internal wikipage \(PeeringDB facility\)](#)
- Intented for small embedded devices such as Raspberry Pi and APU2 ([EOL](#))

Short server housing introduction: [IBM.com](#)

Our switch and rack

- Pre-used Cisco SG300
 - Was already in use at Labitat
 - Can do MAC port filtering




IX Peering Policy



- Allowed Ethernet frames:
 - 0x0800 - IPv4
 - 0x0806 - ARP
 - 0x86dd - IPv6
- Proxy ARP, OSPF, DHCP, STP/RSTP/VSTP with any kind of CDP/ZEROTIER are not allowed
- MAC-addresses are filtered per port
- Routes learnt through peering network are not allowed to be announced by BGP
- Connecting via any type of tunnel or VPN/Defragmented MTU is prohibited.

Based on [LocIX peering policy](#)

IX Client Peering Policy



- Open
 - “We peer with everyone”
 - Peering with route server, including newcomers
- Selective
 - “We peer with most networks, but contact us first”
 - Direct peering connections
- Restricted
 - “You must meet our conditions before we peer”
 - Ex. transit and large networks

Source: nsrc.org

Getting internet resources allocated


- IANA -> RIR -> LIR -> End User
 - Network operator: Eyeball and/or business customers, servers and routers
 - Globally routable
 - IXP: For uniquely identify routers of IX members
 - Not globally routed-ish



■ IANA ■ RIPE ■ LIR ■ End User

Illustrative purpose only

RIPE Database



- Create objects:
 - org: create new, or use existing org
 - mnt
 - person
- Optionally, but required for non-IXPs:
 - AS-Set

Course: [RIPE Database](#) from RIPE academy

Miscellaneous info when talking to LIR



- IX name
- Want route server? YES, then:
 - ASN required
 - Request 16-bit ASN as client can potentially bring old equipment
 - IXP prefix?
 - RIPE has a prefix dedicated for IXPs
 - Specify to only announce withing peering LAN
 - Size: /24 Legacy IP and /48 IPv6

Miscellaneous info when talking to LIR



- Get hold of a receipt or contract confirming connectivity to IXP colo site
- Have at least three customers that is willing to peer at IXP.
- Facility postal address that is able to receive physical mail
- Contact person name, organisation and phone number

Finding Local Internet Registry (LIR)




- Ask your local network wizard
- We asked the [Freetransit.ch](https://freetransit.ch) project from Openfactory
 - Bias: We personally know some of the nice people behind it
 - Run by people attending this camp
- Search “RIPE LIR services” in a web search engine

Similar projects




- The many community run internet exchanges around the globe
- Hosting
 - Coloclue in Amsterdam ([PeeringDB](#))
 - [K-space](#) in Tallinn
 - communityrack.org in Zürich ([PeeringDB](#))
- bgp.wtf in Warsaw Hackerspace ([PeeringDB](#))
- Individual Network Berlin (in-berlin.de), in Berlin ([PeeringDB](#))

Thanks to...



- Project collaborators:
 - Olivia Wenya ([AS199750](#))
 - Daniel Brasholt ([AS198886](#))
- Jori from the [Freetransit.ch](#) project for LIR and RIPE matters
- Asbjorn Tønnesen ([AS207727](#)) for sharing networking knowledge


Good resources



- [PeeringDB.com](https://www.peeringdb.com)
 - The interconnection database
- [BGP.tools](https://www.bgp.tools)
 - Browse the Internet ecosystem
- Local Network Operator Meetups
 - DKNOG in Denmark
 - DENOG in Germany
- RIPE academy. academy.ripe.net
 - Free courses in IPv6, BGP security, RPKI and more
 - Certificate available for a fee

Good resources

- Network Startup Research Center. nsrc.org
 - BGP and IXP introduction
- NLNOGs [BGP Filter Guide](#) and NLNOG [ring](#)
 - BGP filtering and network reachability map
- Euro-IXs [Best Current Operational Practices](#) and the [IXP Wishlist](#)
 - Overview of infrastructure considerations for an IX
- Presentation: [The 1000\\$ Internet Exchange](#) at UKNOF32 by Remco van Mook




Find us



- Camps
 - **BGP-enabled hackerspaces meet-up at Chillfloor at 17:00, on day 3**
 - CCCamp at Bornhack village in Cold North cluster
 - BornHack 2024
- Every other day
 - Labitat in Copenhagen
- Website: ix.labitat.dk
- PeeringDB: peeringdb.com/ix/4193
- Plain old email: labix@labitat.dk


Sponsors?

- All hardware is pre-used
- We are looking for newer hardware, second uplink and stickers!
- LIR fees paid from out of own (student) pocket



Whats next?

- Looking Glass for easier debugging
- Limit ARP broadcast storm with ARP Sponge
- Birdwatch: Prometheus Exporter
- Much more documentation
- Setup RPKI relay
 - Labitats uplink already filters for RPKI invalids



Thank you for listening

- Download the presentation
 - URL: <https://labix-presentation-v2.hafnium.me>





This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International (CC BY-SA 4.0). It makes use of the works of Mateus Machado Luna.



Images credited to its own authors on respective slides

