How To Internet

The basics of subnetting routing and the building blocks of the Internet

$ whoami

- famfo
  > https://famfo.xyz
- networking, distributed systems and some embedded things
- you may know me from authoritative DNS shenanigans
  > https://beta.servfail.network
- @famfo@chaos.social
- @famfo@1.6.0.0.8.0.0.b.e.d.0.a.2.ip6.arpa

Some Notes

- examples for iproute2, busybox (Linux utils)
- some binary representations of numbers
- due to IPv4 exhaustion, I only found IPv6 (*new in 1998!*) addresses as examples

What Are We Looking At?

- TCP/IP reference model
- Everything on the IP layer

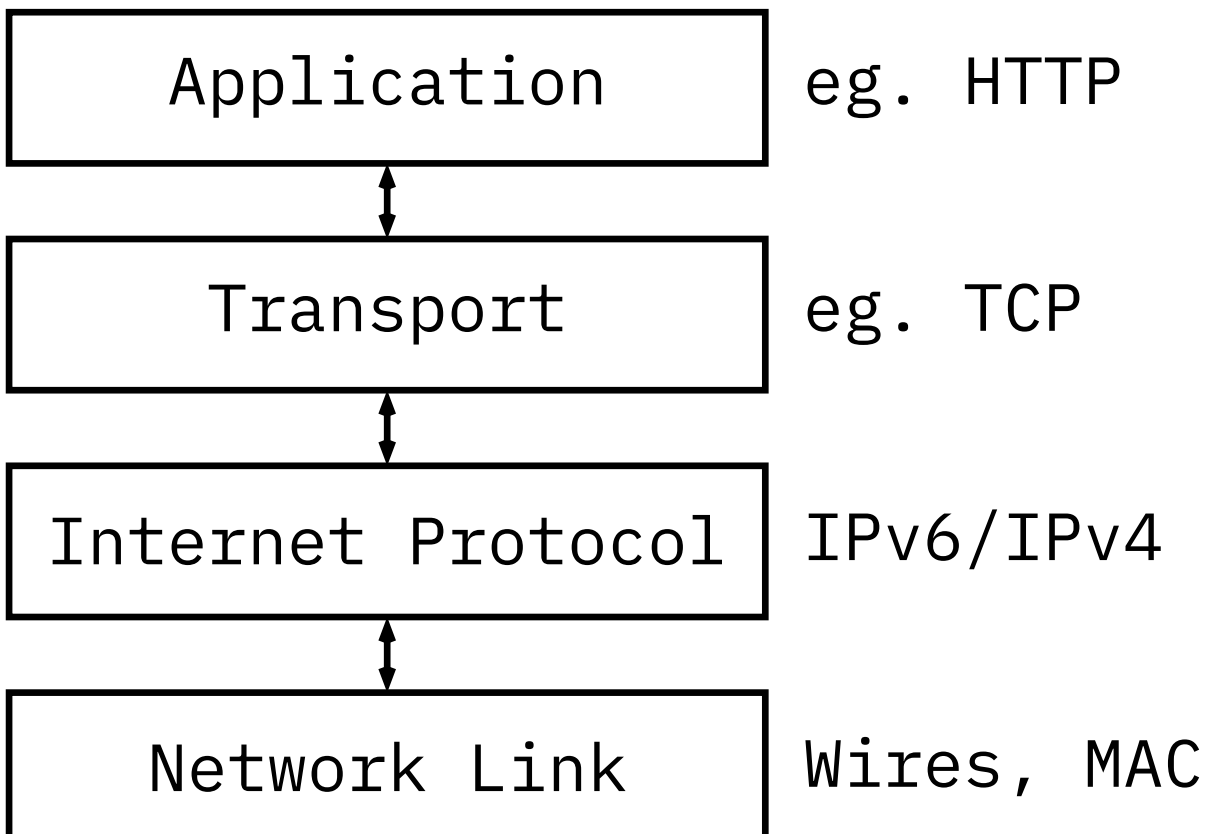

IP Addresses (1)

- address hosts on the internet
- IPv4 (32 bit) and IPv6 (128 bit)
  - 1.0.0.1
  - 2620:fe::fe

    Compressed IPv6 address, expands to:

    2620:00fe:0000:0000:0000:0000:0000:00fe

IP Addresses (2)

How do they arrive on your device?

- IPv6 autoconfiguration
  > IPv6 Router Advertisements
- IPv4 DHCP

IP Addresses (3)

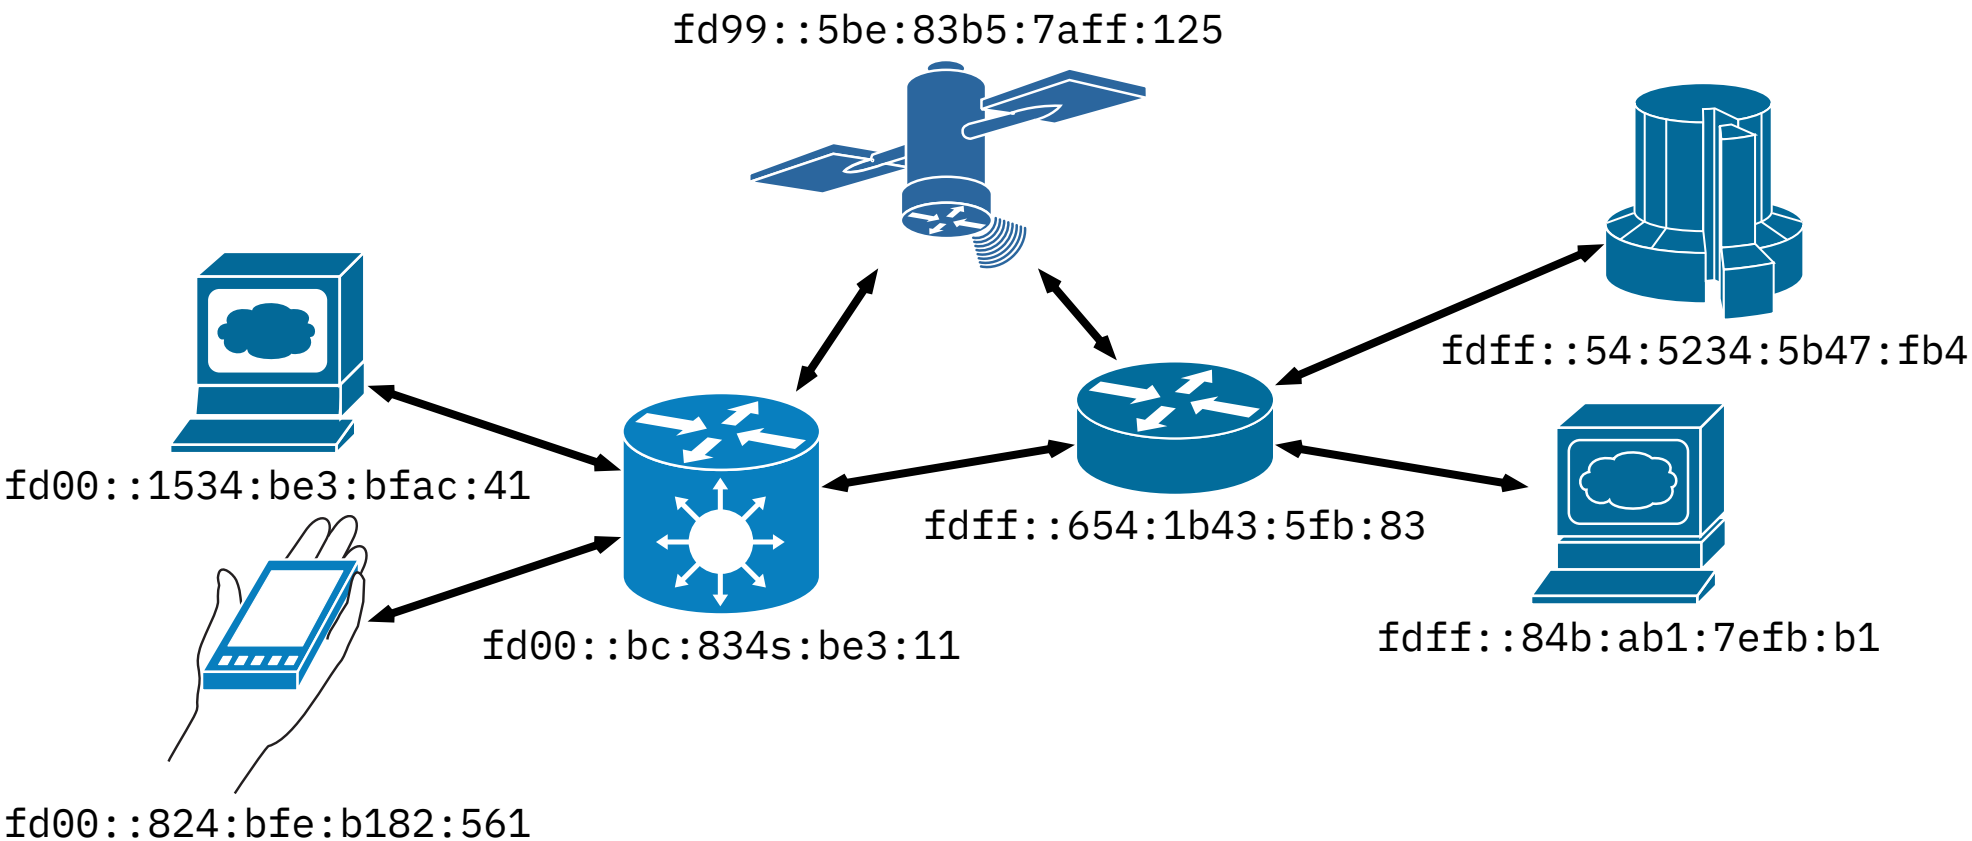

Routes

```
$ ip -6 route

default via fd00::1 dev eth0
fd00::/64 dev eth0
```

The Default Route (1)

- packets go there
- “internet behind this interface”
- if no route is found, packets get dropped

The Default Route (2)

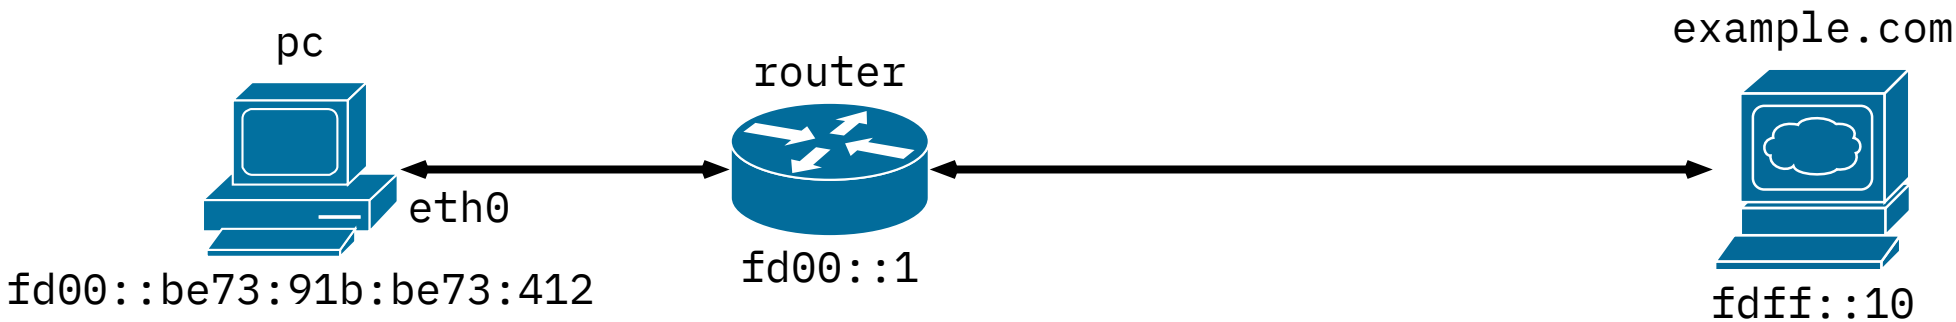

> HTTP request to fdff::10

```
pc$ ip -6 route

default via fd00::1 dev eth0
```

“Behind fd00::1 on link eth0 is the internet”

Subnetting (1)

- *Sub* divide a range of IP addresses
- Classless Inter-Domain Routing (CIDR)

> Historically used: network classes, obsolete since 1993! Still taught in your German school of choosing.

Subnetting (2)

```
$ ip address

fd23:ccc:ac00:1::1/64
------------------ --
|                  |
+-- IP Address     +-- Size of the subnet
```

Subnetting (3)

- size says *how many bits are set*
  - a *mask* can be created from that
- if a bit is set, it can’t change in the address
- last IPv4 address is the *broadcast address*

Fictional 8bit Address Type (1)

Address: 01010000/8
Mask: 11111111
> The eight most significant bits can’t change
> One address in the subnet

Fictional 8bit Address Type (2)

Address: 01010011/4
Mask: 11110000
> The four most significant bits can’t change
> Sixteen addresses in the subnet
> 01010000 - 01011111

Subnetting (4)

fd23:ccc:ac00:1::1/64

> 64 bit are set

fd23:0ccc:ac00:0001:0000:0000:0000:0001
ffff:ffff:ffff:ffff:0000:0000:0000:0000

Subnetting (5)

192.168.0.0/24

> 24 bits are set

> 24 + 8 = 32

Subnetting (5)

192.168.0.0/24

> 24 bits are set

> 24 + 8 = 32

192.168.000.000

255.255.255.000

The Other Thing In The Routing Table (1)

```
$ ip -6 route

default via fd00::1 dev eth0
fd00::/64 dev eth0
```

- `fd00::/64` a so-called *prefix*
- first address of the subnet
- everything to the subnet, sent over `eth0`

The Other Thing In The Routing Table (2)

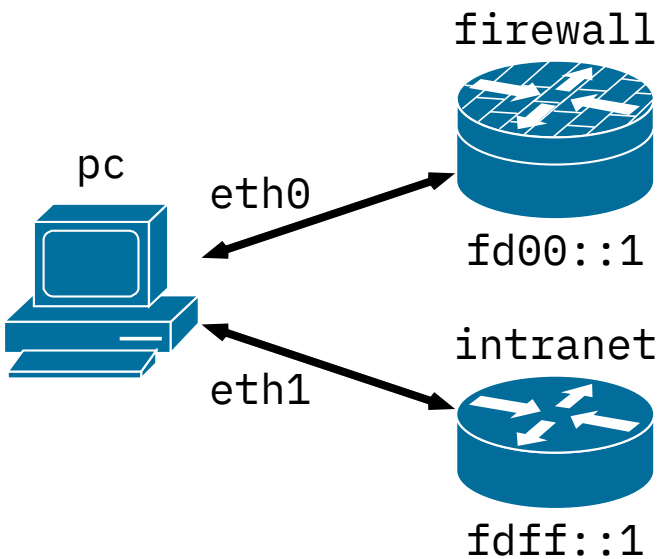


```
pc$ ip -6 route
default via fd00::1 dev eth0
fdff::/64 dev eth1
```

A Small Routing Table

```
$ ip -6 route

default via fd00::1 dev eth0
fd00::/64 dev eth0
fdff:0:0::/48 dev eth1
fdff:0:0:aa::/56 dev eth0
```

- the most specific route will be used

After The Default Route

- default-free zone

What Is An ISP?

- *Autonomous System* (AS)
  > identifiers for owners of IP prefixes
- Global authority *IANA*
- Delegated resources to *Regional Internet Registries* (RIR), for Europe *RIPE*
- exchange *reachability information*
  > *Border Gateway Protocol* (BGP)

Exchanging Reachability Information (1)

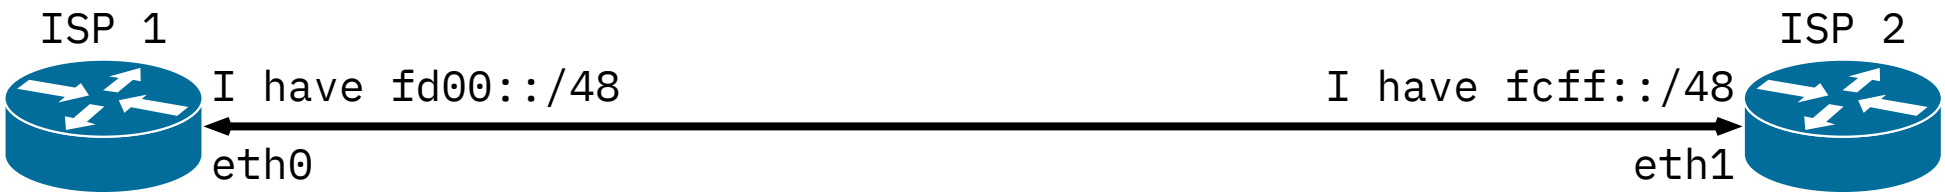

Routing Tables

ISP 1

Prefix	Interface
fd00::/48	lo
fcff::/48	eth0

ISP 2

Prefix	Interface
fcff::/48	lo
fd00::/48	eth1

Exchanging Reachability Information (2)

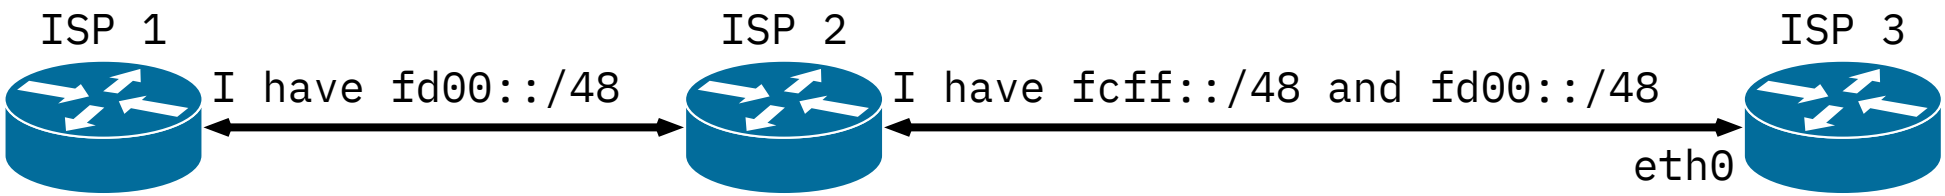

ISP3 Routing Table

Prefix	Interface
fd99::/48	lo
fd00::/48	eth0
fcff::/48	eth0

Internet routing table size

- IPv6: ~209653 prefixes
- IPv4: ~998127 prefixes

Counts from he.net's looking glass

Exchanging Reachability Information (3)

- Initially a daemon sends entire routing table
- Incremental update
  - New router joining the network
  - Existing router exiting (withdraw, timeout, …)

Exchanging Reachability Information (4)

- A BGP announcement contains more information
- *AS Path*

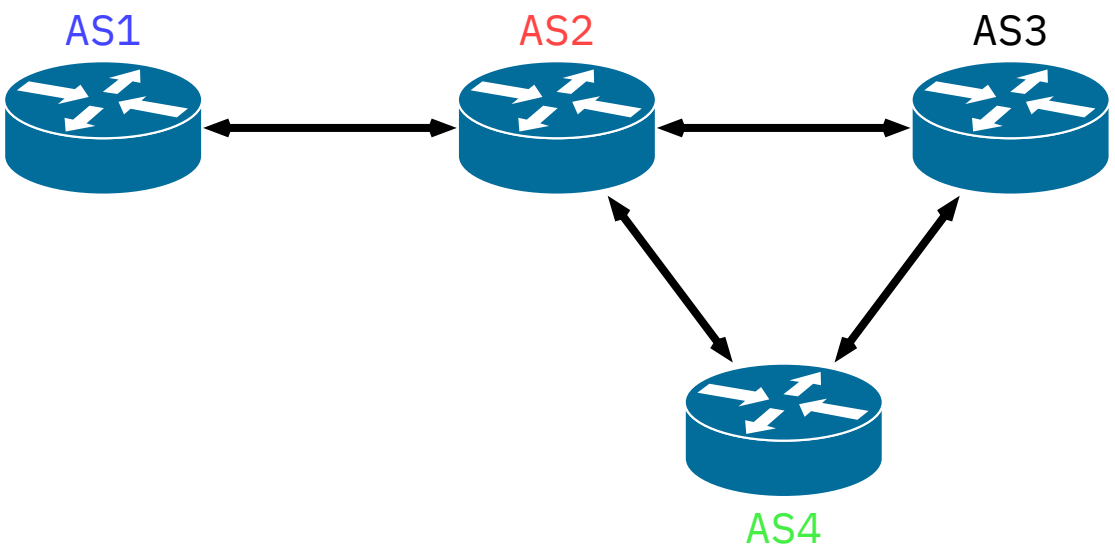

View from AS4

From	AS Path
AS2	AS2, AS1
AS3	AS3, AS2, AS1

> The shortest *AS Path* will be chosen

Further Reading And Experimenting

- dn42 - pretend internet
  > Private AS and IP range
  > https://dn42.cc
- there are more routing protocols!
  > BGP is only for inter AS communication
  > ASs also do internal dynamic routing
  > OSPF, IS-IS, RIP, Babel, …
- there are other forms of making a routing decision (eg. MPLS)

Summary

- IP addresses (IPv4 and IPv6)
- default route
- subnetting
- how a routing table looks
- default-free zone
- BGP basics