

A Network Of Networks

Ben Cartwright-Cox - RIPE 88 Student Event

I find infrastructure
utterly fascinating



```
[11:30:24] ben@metropolis:~$ ping 163.237.247.17
```

```
PING 163.237.247.17 (163.237.247.17) 56(84) bytes of data.
```

```
64 bytes from 163.237.247.17: icmp_seq=1 ttl=52 time=90.3 ms
```

```
^C
```

Best departing flights

Ranked based on price and convenience ⓘ Prices include required taxes + fees for 1 adult. Optional charge



08:10 – 10:45

American · British Airways, Iberia, Finnair

8 hrs 35 min

LHR-ORD

Non-stop

Real World (One way)

Best departing flights

Ranked based on price and convenience ⓘ Prices include required taxes + fees for 1 adult. Optional charge



08:10 – 10:45

American · British Airways, Iberia, Finnair

8 hrs 35 min

LHR-ORD

Non-stop

Host	Loss%	Snt	Last
1. _gateway	0.0%	23	0.9
2. cyclone-flat.p2p.b621.net	0.0%	23	6.0
3. lo.router-cam-a.mythic-beasts.com	0.0%	23	5.9
4. lo.router-sov-d.mythic-beasts.com	0.0%	23	6.9
5. ldn-b4-link.ip.twelve99.net	23.8%	22	5.5
6. ldn-bb2-link.ip.twelve99.net	68.2%	22	6.4
7. ldn-b3-link.ip.twelve99.net	0.0%	22	6.2
8. (waiting for reply)			
9. be2871.ccr42.lon13.atlas.cogentco.com	0.0%	22	5.6
10. be2101.ccr32.bos01.atlas.cogentco.com	0.0%	22	74.6
11. be3600.ccr22.alb02.atlas.cogentco.com	0.0%	22	78.0
12. be2879.ccr22.cle04.atlas.cogentco.com	0.0%	22	84.8
13. be2718.ccr42.ord01.atlas.cogentco.com	0.0%	22	93.1
14. be2513.rcr51.b025243-2.ord01.atlas.cogentco.com	0.0%	22	95.0
15. 38.104.103.70	0.0%	22	91.4
16. 163.237.254.128	0.0%	22	92.2
17. yourorg01.ring.nlnog.net	0.0%	22	91.0

Best departing flights

Ranked based on price and convenience ⓘ Prices include required taxes + fees for 1 adult. Optional charge



08:10 – 10:45

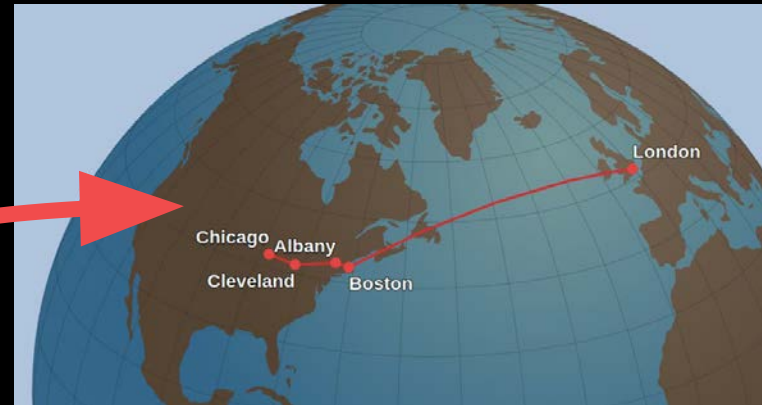
American · British Airways, Iberia, Finnair

8 hrs 35 min

Non-stop

LHR-ORD

Host	Loss%	Snt	Last
1. _gateway	0.0%	23	0.9
2. cyclone-flat.p2p.b621.net	0.0%	23	6.0
3. lo.router-cam-a.mythic-beasts.com	0.0%	23	5.9
4. lo.router-sov-d.mythic-beasts.com	0.0%	23	6.9
5. ldn-b4-link.ip.twelve99.net	23.8%	22	5.5
6. ldn-bb2-link.ip.twelve99.net	68.2%	22	6.4
7. ldn-b3-link.ip.twelve99.net	0.0%	22	6.2
8. (waiting for reply)			
9. be2871.ccr42.lon13.atlas.cogentco.com	0.0%	22	5.6
10. be2101.ccr32.bos01.atlas.cogentco.com	0.0%	22	74.6
11. be3600.ccr22.alb02.atlas.cogentco.com	0.0%	22	78.0
12. be2879.ccr22.cle04.atlas.cogentco.com	0.0%	22	84.8
13. be2718.ccr42.ord01.atlas.cogentco.com	0.0%	22	93.1
14. be2513.rcr51.bu25z43-2.ord01.atlas.cogentco.com	0.0%	22	95.0
15. 38.104.103.70	0.0%	22	91.4
16. 163.237.254.128	0.0%	22	92.2
17. yourorg01.ring.nlnog.net	0.0%	22	91.0



6600KM+ for practically zero cost


The machine is huge

- Despite gloom and doom, we still have one internet for the most part
 - It is economically important for countries to remain connected, **even if it costs them domestic control**
 - The propagation of information will have changed a unthinkable amount of things, including wars.
- While some parts of the machine have centralized bodies for tasks like number allocation, the rest of the internet is constructed from a large mutual understanding of how things should work


Inside the machine

- Networks (Companies) link up to other networks
 - Sometimes far away, sometimes in the same building
- Their computer tells the other computers what addresses they want traffic for

- Some network provided "policy" (aka code/logic) is run to decide if they want to accept that information, and if it should be used right now



This scales up!



In reality this is what BGP updates look like:

```
+|62.193.131.0/24|28917 13004 6700|195.208.208.147|i|||28917:4000 28917:4100 28917:4103|195.208.208.147 28917|1667136299|1
+|84.237.71.0/24|28917 9049 25549 25549 30797|195.208.208.147|i|||28917:3000|195.208.208.147 28917|1667136299|1
-|213.145.128.0/24| | | | |195.208.210.40 6939|1667136297|1
```

In reality this is what BGP updates look like:

Prefix	AS Path	Next Hop	BGP Communities	Data Source / ASN
+ 62.193.131.0/24	28917 13004 6700	195.208.208.147 i	28917:4000 28917:4100 28917:4103	195.208.208.147 28917 1667136299 1
+ 84.237.71.0/24	28917 9049 25549 25549 30797	195.208.208.147 i	28917:3000	195.208.208.147 28917 1667136299 1
- 213.145.128.0/24				195.208.210.40 6939 1667136297 1

Prefix	AS Path	Next Hop	BGP Communities	Data Source / ASN
--------	---------	----------	-----------------	-------------------

```
route-views.amsix.routeviews.org> show bgp ipv4 185.230.223.0/24
```

```
BGP routing table entry for 185.230.223.0/24
```

```
Paths: (27 available, best #13, table default)
```

```
Not advertised to any peer
```

```
211398 34854 44684 206924
```

```
185.1.167.45 from 185.1.167.45 (31.24.249.245)
```

```
Origin IGP, valid, external
```

```
Community: 34854:1001 65500:10000 65500:10100 65532:1000
```

```
Large Community: 44684:0:900 44684:1:2 44684:2:206924 44684:3000:1299 211398:245:100 211398:2760245:34854
```

```
Last update: Sun Oct 30 01:17:05 2022
```

```
39120 3356 3170 206924
```

```
80.249.210.28 from 80.249.210.28 (195.60.190.29)
```

```
Origin IGP, valid, external, atomic-aggregate
```

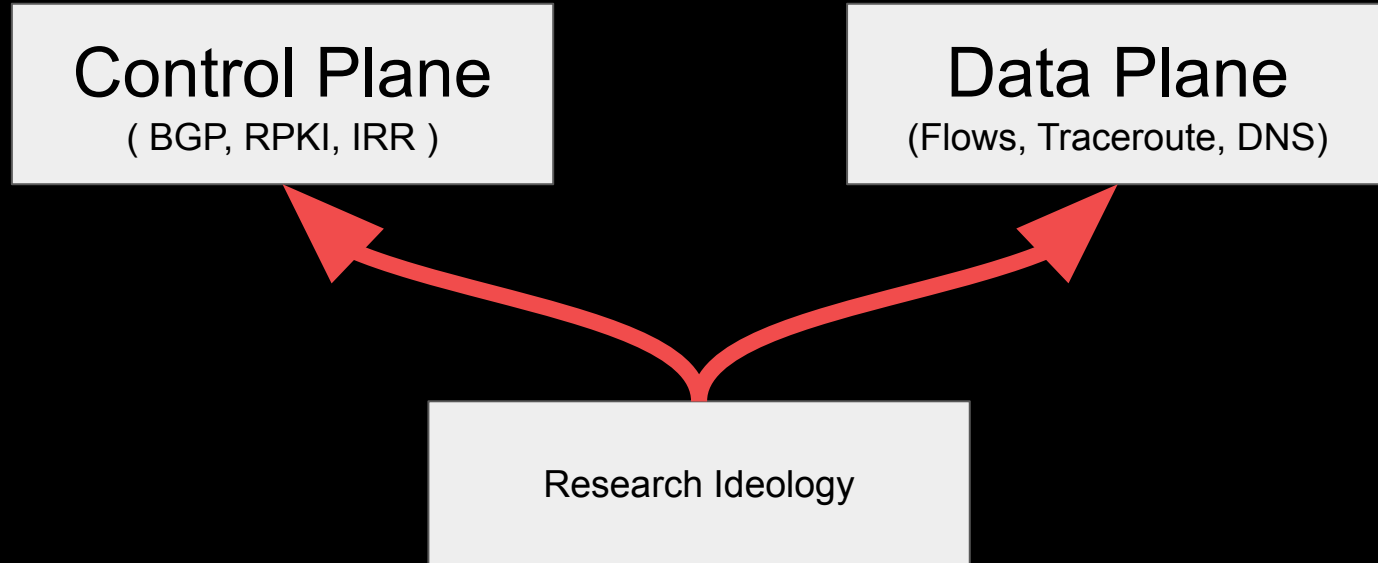
```
Community: 3356:2 3356:22 3356:100 3356:123 3356:500 3356:901 3356:2064 60945:0 65002:6830 65532:1000
```

```
Last update: Sat Oct 29 08:33:29 2022
```

BGP Implicitly encodes so much extra information

- Routing around expensive providers
 - You can observe when some carriers become uncompetitive over time
- Routing around physical issues
 - Recent de-prefing of some ISPs during the recent APAC undersea cuts
- Routing around with politics in mind
 - North Korea brought up a 2nd upstream (Russia) after suspected tensions(?) with China
- Literally routing stuff into blackholes with *Politics* in mind

Two ways to look at the internet



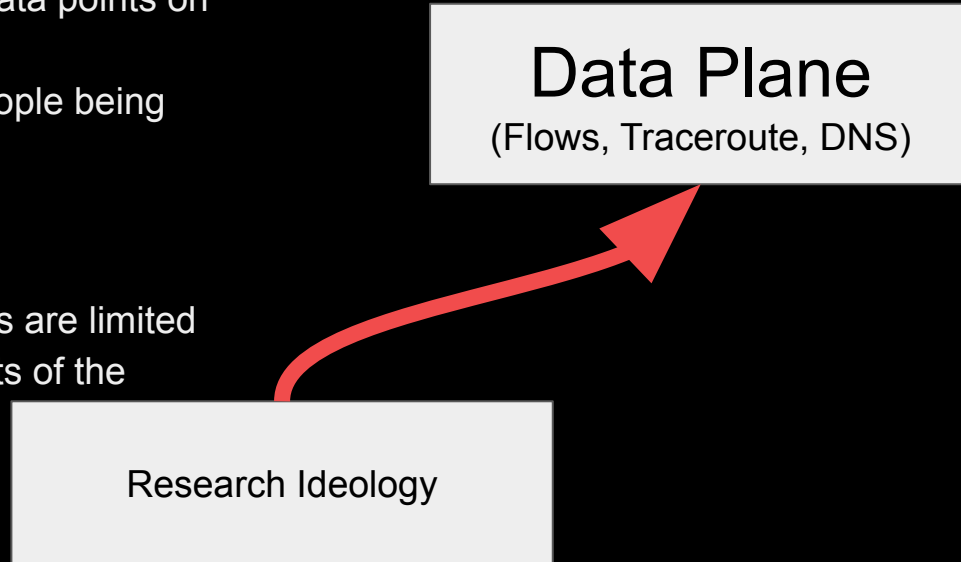
Two ways to look at the internet

Pros:

- Provides a set of more real data points on where traffic goes
- Easier to do without other people being involved

Cons:


- Coverage for testing networks are limited
- Hard to take atomic snapshots of the internet with this
- Firewalls can get in the way



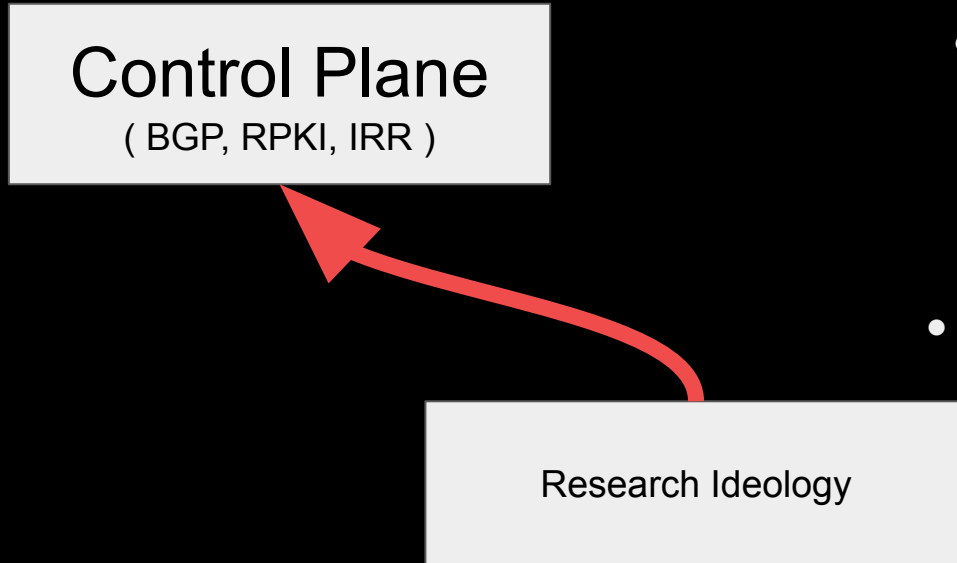
Data plane "cool stuff" - According to Ben

- Internet mapping to see how the internet is being used over time
- Automatic interpretation of traceroute data
- BGP anycast detection
- Network to Network latency breakdowns

Internet mapping!



Two ways to look at the internet



Pros:

- Lots of downloadable data
- Data generally covers the whole internet in one go

Cons:

- Control Plane does not always align with the Data Plane
 - Lots of stale data around
- Some of the "good" data needs human interaction

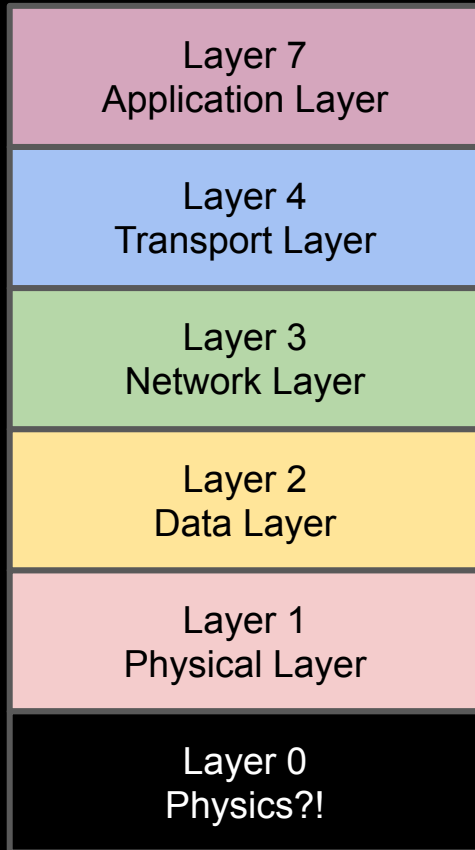
Control plane "cool stuff" - According to Ben

- Better BGP Software
 - The industry as a whole has mediocre software quality
- RPKI Adoption tracking
 - We can only improve internet routing security by checking if we are making progress or not
- Network Ranking
 - How do you rank a network as the best? Many ways to do it, many real world networks would love to argue it with you!

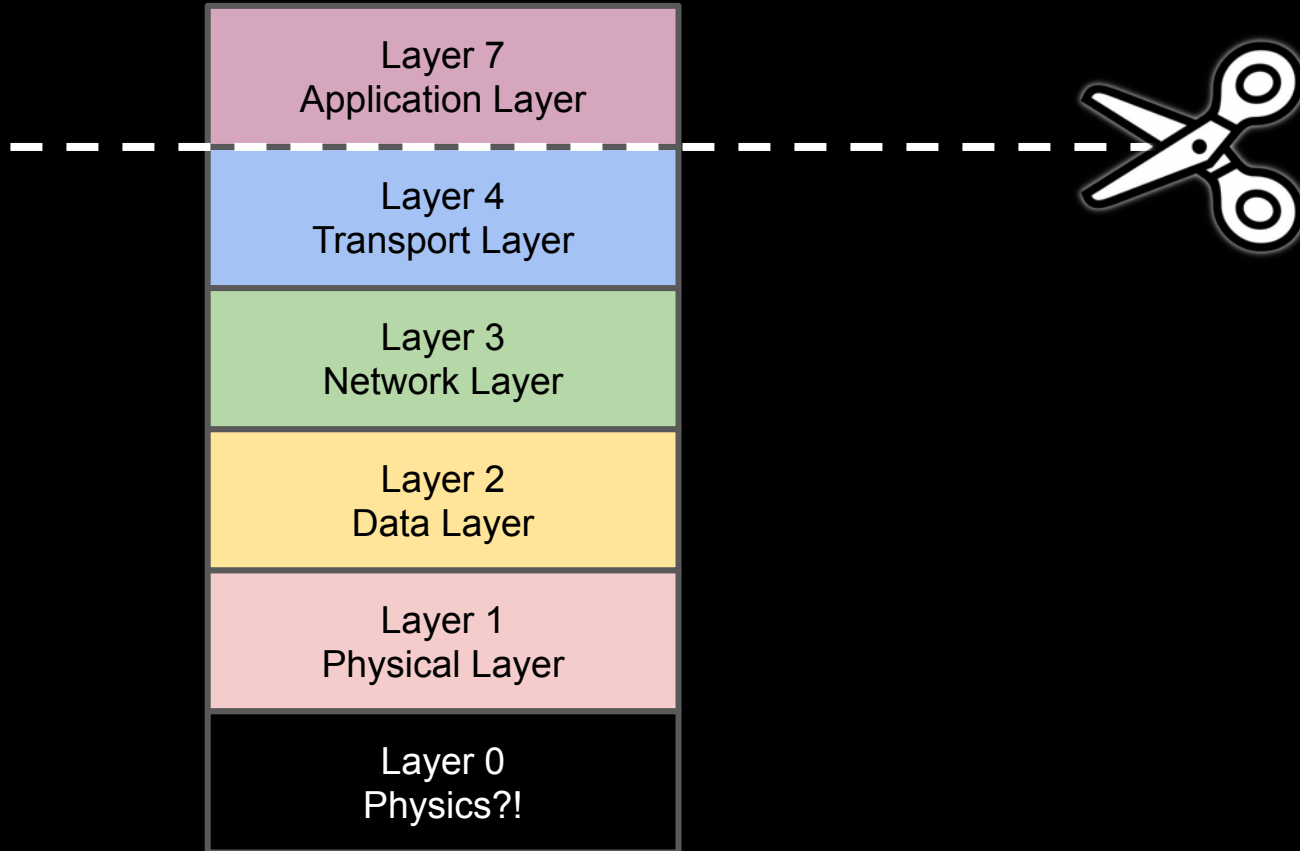
The internet is remarkably flexible

Because it's built out of the wills of the people who run their
own section of the network

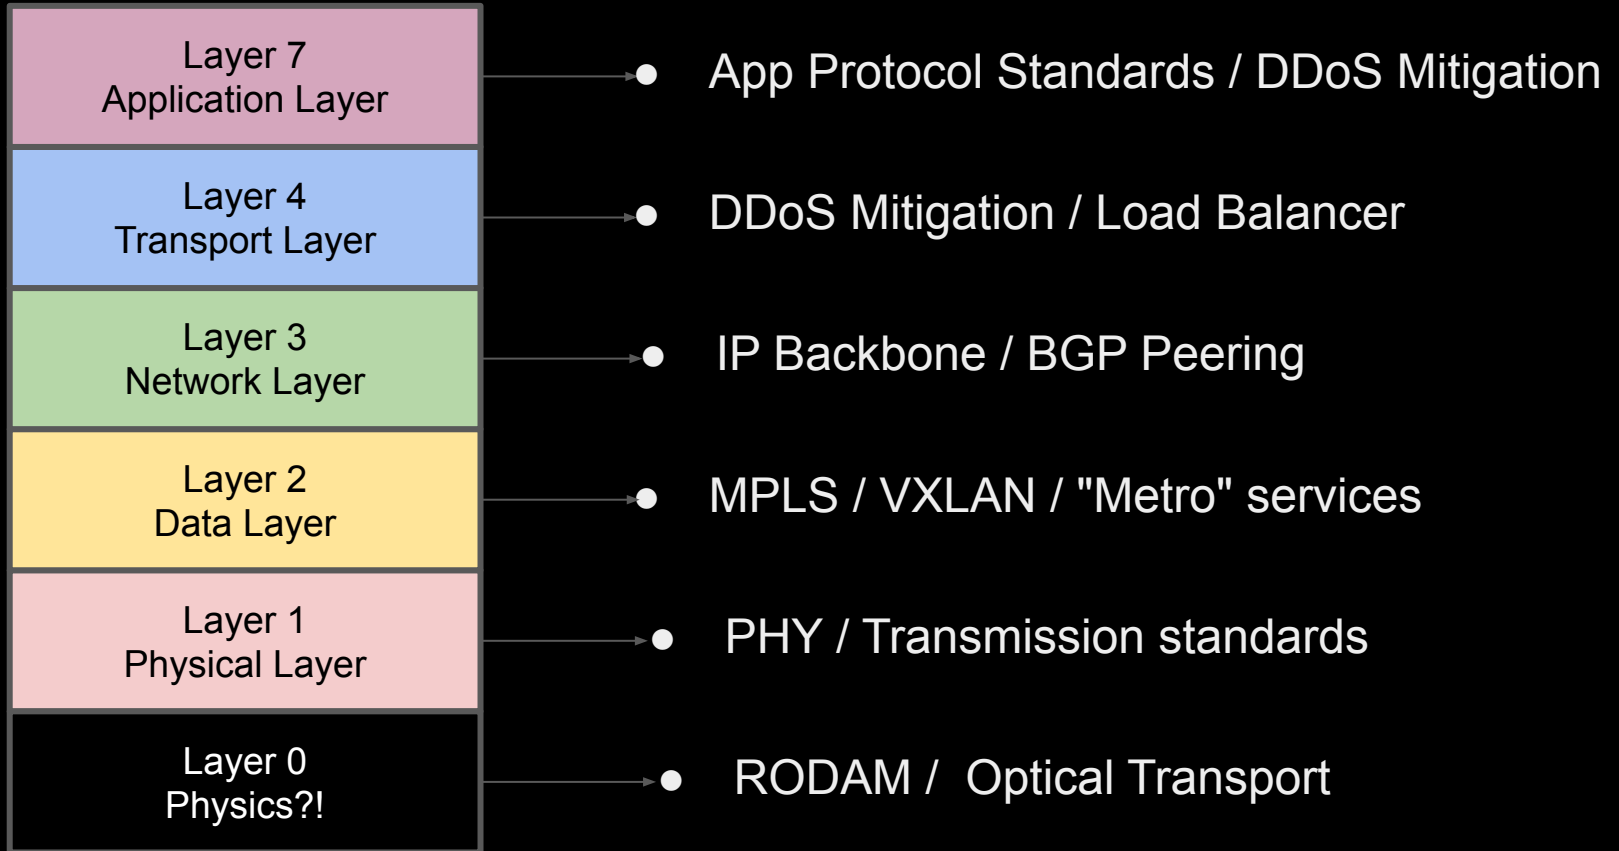
Meet the people running the networks



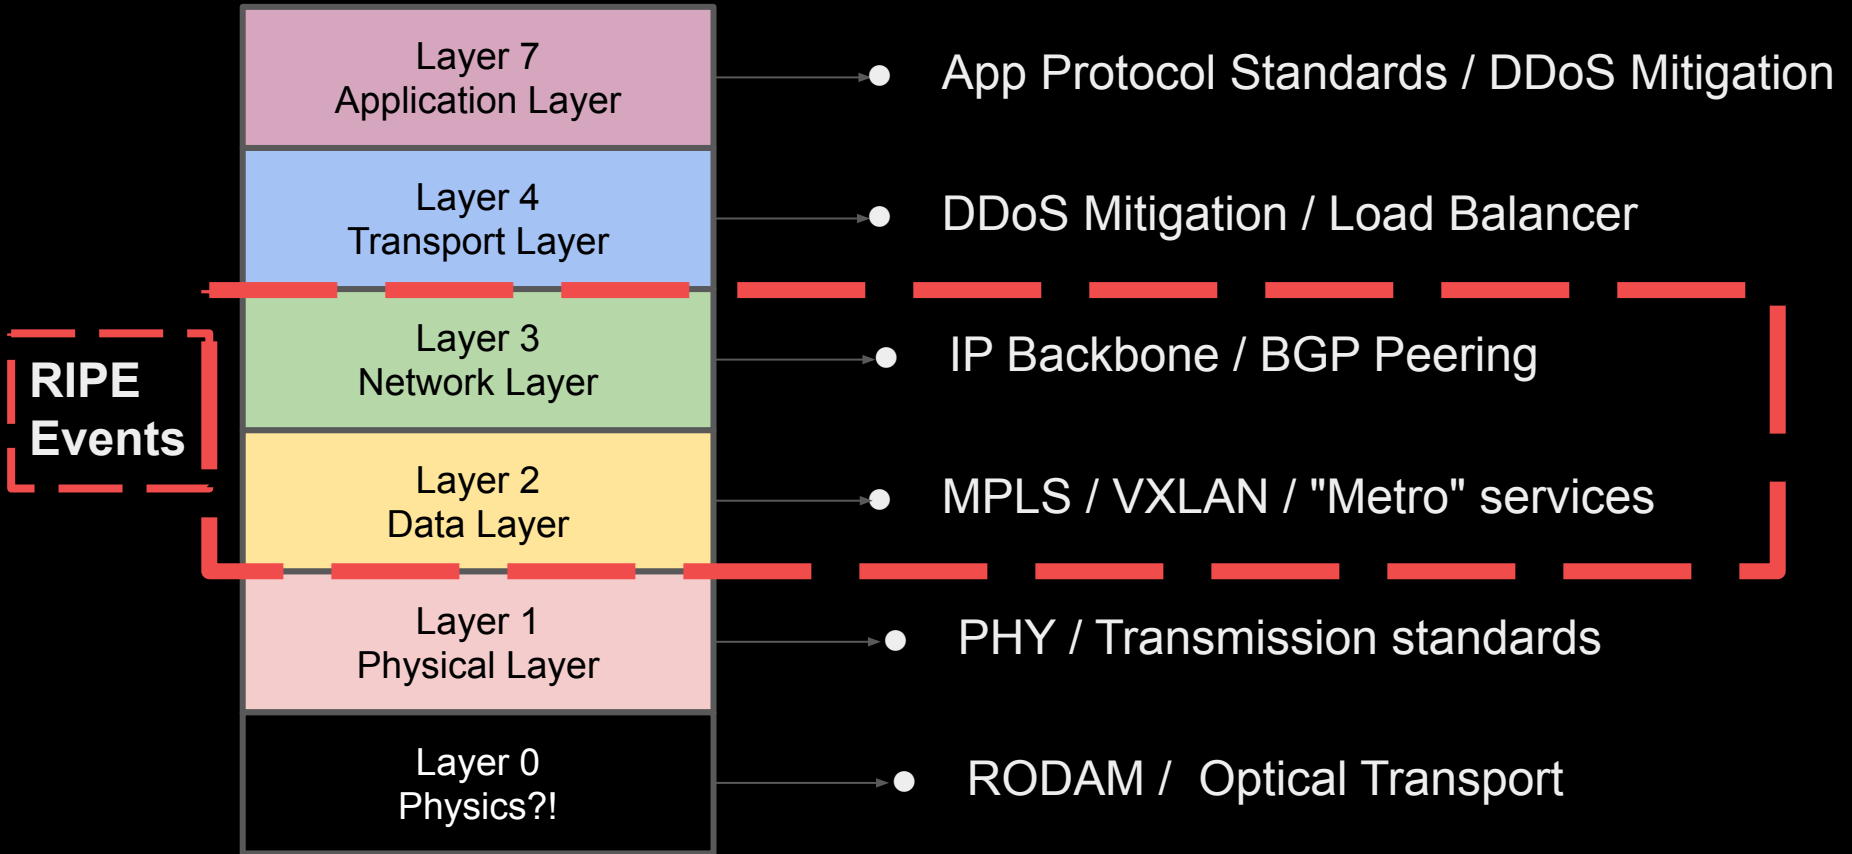
Meet the people running the networks



Meet the people running the networks



Meet the people running the networks




Here is what inspires me

- I worked in software development doing various things, but always have had a lingering interest in the networking that makes the internet work
- In 2017 I got a RIPE autonomous system number (ASN) and ran my own network for a while
- I grew frustrated with the tooling that the industry had/has, so much so that I started to build my own
- What started as a hobby, turned into a business

Here is what inspires me

- I worked in software development doing various things, but always have had a lingering interest in the networking that makes the internet work
- In 2017 I got a RIPE autonomous system number (ASN) and ran my own network for a while
- I grew frustrated with the tooling that the industry had/has, so much so that I started to build my own
- What started as a hobby, turned into a business

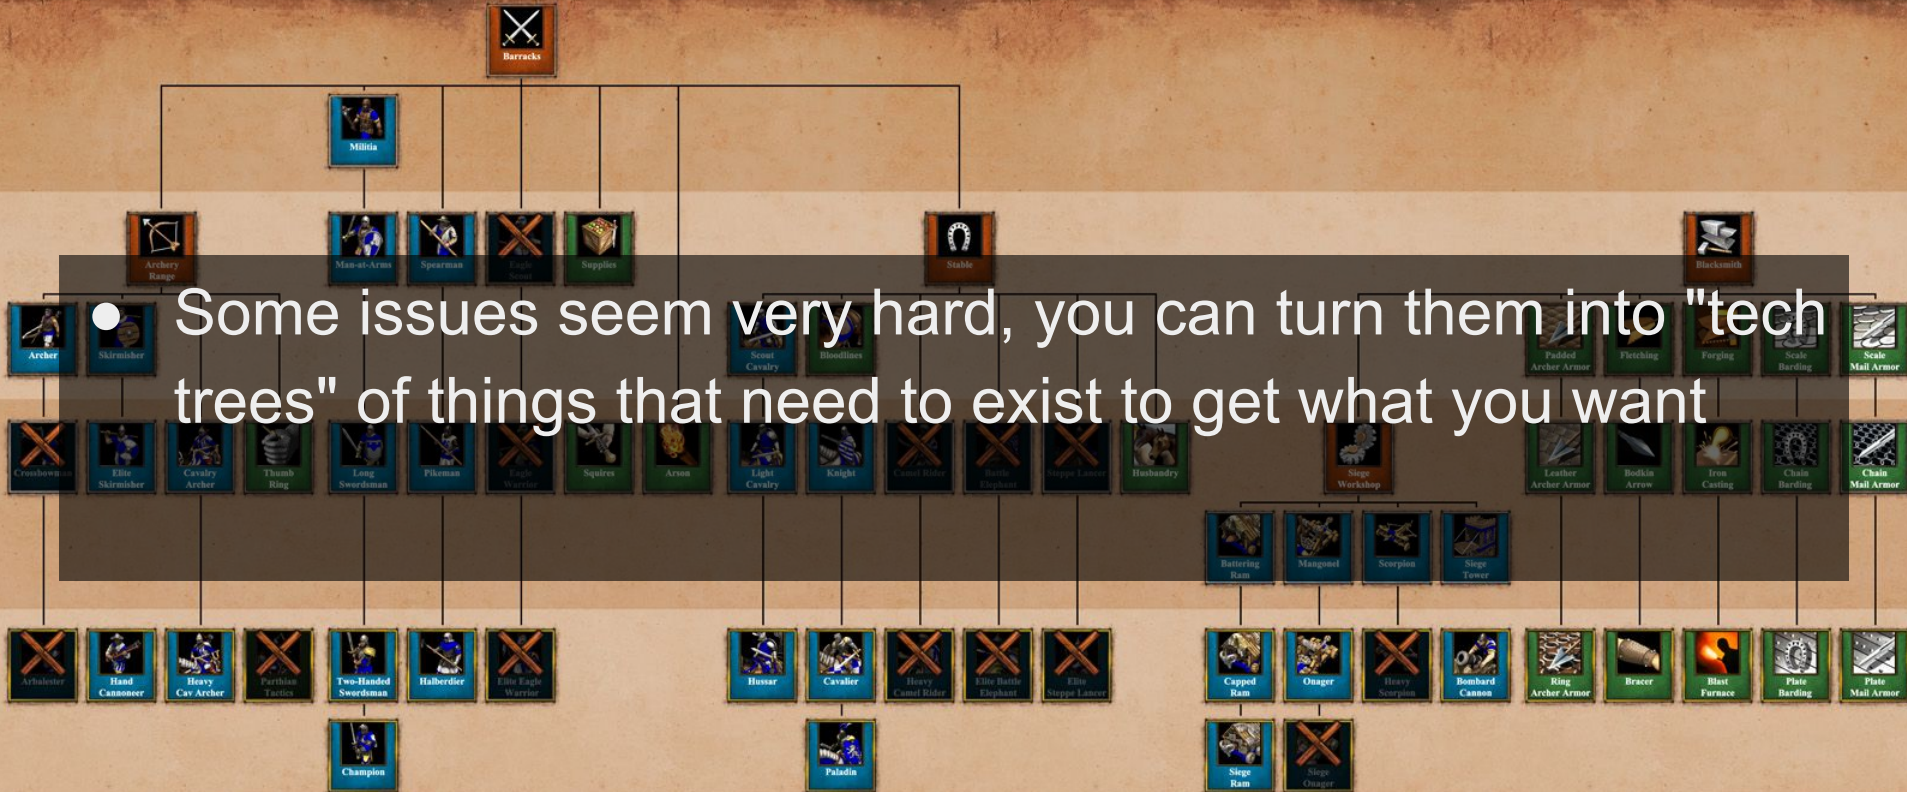


bgp.tools

Problems can be fixed!

- Some issues seem very hard, you can turn them into "tech trees" of things that need to exist to get what you want

Problems can be fixed!

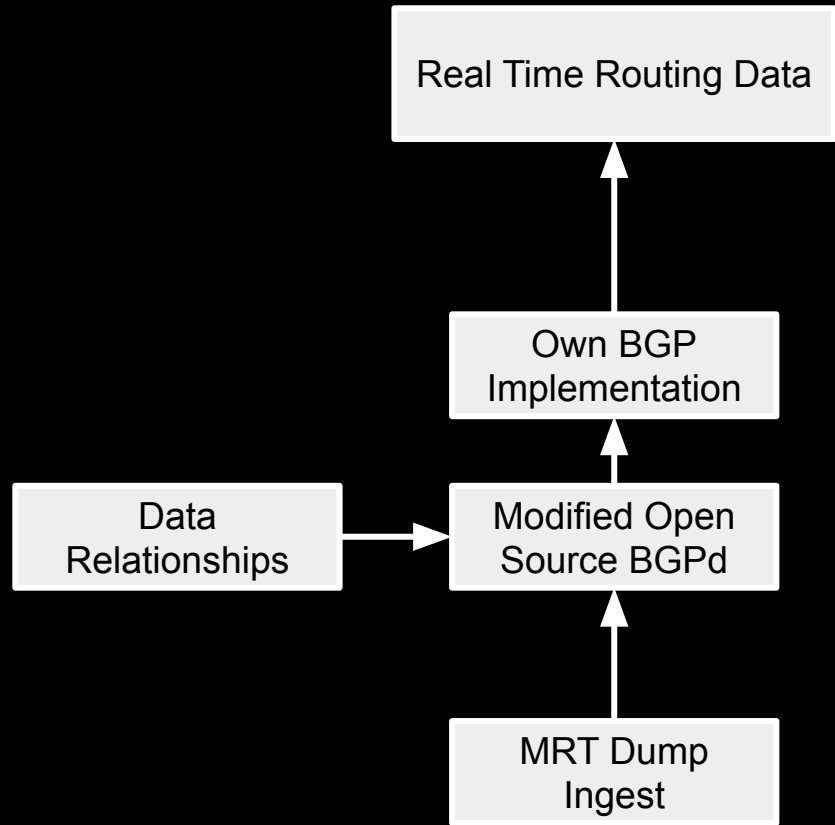


Ideal: I want the freshest data for my tool

- You can freely get BGP data from:
 - RIPE RIS: <https://ris.ripe.net/docs/mrt/>
 - RouteViews: <https://routeviews.org/>
- You can then use MRT files to build the final data you want
- Problem:
 - The "full" dump files are only made every 4 to 8 hours
 - This is not fresh enough to be useful for real time debugging

Tech Tree to success

- At first you have to build what you are able to
- At some point I was using modded versions of BIRD BGP
- Developing relationships with networks for data
- Outscaled BIRD, wrote my own BGPd for my own needs
- Scale scale scale!



I don't think there is a "reason" the internet works

- You can point to the standards
 - But ultimately there are 90,000+ different entities choosing to do them
- You can point to the allocation bodies
 - But ultimately most 90,000+ networks are choosing to believe them

The internet is a machine with no defined purpose,
"We" found a purpose for it after it was designed, and that purpose changes over time

The real "reason" the internet works

- We want it to work

The real "reason" the internet works

- We want it to work
- However, there is still a lot of things that don't work, and for that we need researchers and operators to keep the machine understood and working
- The machine that supports one of the greatest (or worst) things humanity has done so far

On that note

Questions?

Shy? Email me on riperaci@benjojo.co.uk