



AS201701

build an internet backbone with pure 1he servers



Oliver Herms
@takt



Philip Berndroth
@pberndro



takt

- Oliver Herms - oliver@freifunk-rheinland.net
- Google Edge Network SRE - takt@google.com
- Autonomous System: AS199714
- AS_PATH: 15169, 201701, 8560, 199714, 12581, 20773, 8422
- Twitter: [@taktv6](https://twitter.com/taktv6)





pberndro

- Philip Berndroth - philip@freifunk-rheinland.net
- Member of the board Freifunk Rheinland e.V.
- Run my own company with software and datacenter stuff
- Fan of Irish Steak and Guinness
- Twitter: @pberndro - Facebook: facebook.com/philip.berndroth



What is Freifunk and what does the FFRL?

- Research in free wifi and mesh networks
- Empower people to build free and open infrastructure
- Organization of regional, national and international meetings and conferences
- provide ip transit for the most Freifunk communities in Germany
- **Access to the internet for socially disadvantaged people**



Bring the Internet to _all_ people



A United Nations report said:
***disconnecting people from
the internet is a human rights
violation and against
international law.***



Free wifi for refugees

- a government task but the Freifunk community fixed it
- Communication with family and friends
- Access to education and information in a foreign country
- Use your money for food and family and not for prepaid internet dealers
- There is no dsl/fiber in a refugee camp



Build internet into the wild is hard...

- There is no business model -> we need donations
- Bring volunteers and professionals together
- many technical challenges to build wifi with mesh networks for the mass
- Freifunk communities accepted the challenge
- Build free internet for many thousands of people in Germany













VPN providers don't scale

- only for a single user
- mostly OpenVPN -> userspace daemon
- you have to pay per user every month
- abuse is a problem, it's not your own network
- -> run your own fucking infrastructure (*is hard*)



DIY ISP

Problems:

- No Budget
- No Number Ressources
- No Hardware
- No Colo
- No Transit
- No Peering





How it began

First idea in November 2013

Freifunk Rheinland e.V. became a LIR in August 2014

First Deployments in September 2014



First deployment

Weekend in September 2014:

Frankfurt (Interxion):

2x Supermicro, Xeon 5405, 4GB RAM, 4x SAS HDD RAID-1

Arch Linux, Quagga

2x 1Gbit/s Transit from takt's former employer

Berlin (Speedbone):

2x HP DL320 Gen8, Xeon E3-1220, 8GB RAM, SD Card

Arch Linux, Quagga

2x 1Gbit/s Transit from IN-Berlin AS29670



Deployment #2


Düsseldorf (Interxion):

2x HP DL320 Gen8, Xeon E3-1220, 8GB RAM, SD Card

Arch Linux, Quagga

2x 1Gbit/s Transit from rrbone AS39138

End of 2014





Problems in Q1/2015

Inter-POP connectivity = GRE

1G circuits are not big pipes

Transit sucks, Peering rocks

XCONs limited or impossible in some of our locations

Quagga sucks

Arch sucks (at least for our usecase)



Solutions

Move FRA POP to another location

Move DUS POP to another location

BCIX sponsored 10G port in BER

ECIX sponsored 10G ports in FRA, DUS, BER, HAM (used by Freifunk Hamburg)

ECIX sponsored 1G L2VPN between all Sites

NLix sponsored 1G ports in FRA, DUS, BER

Somehow we are on DECIX FRA

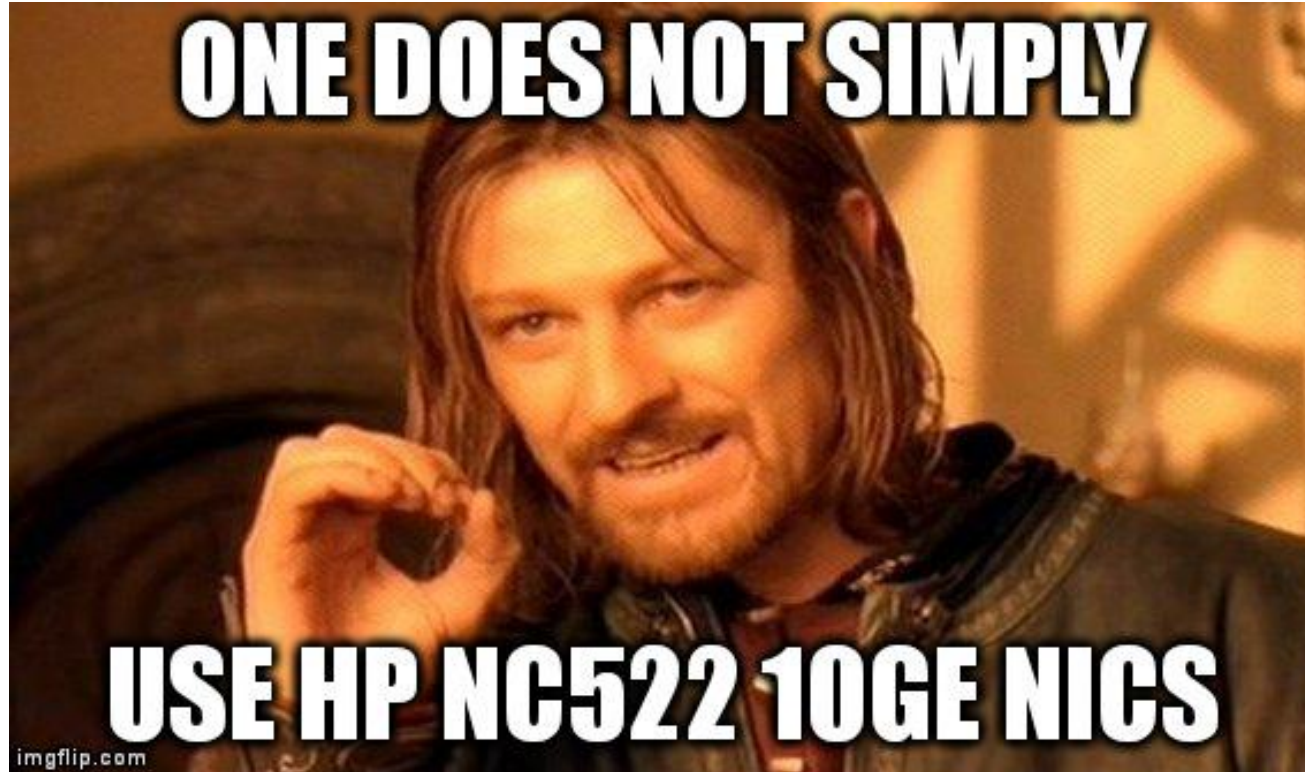
Replaced Quagga with BIRD

Replaced Arch with Debian

Replaced Debian Kernel with CentOS Kernel



First approach to 10GE (HP NetXen)





Mellanox ConnectX on Kernel 4.2

```
for interface in `cat /proc/net/dev | grep eth | awk -F: '{print $1}'`; do
```

```
    /sbin/ethtool -K $interface tx off
```

```
    /sbin/ethtool -K $interface gro off
```

```
    /sbin/ethtool -K $interface gso off
```

```
    /sbin/ethtool -K $interface tso off
```

```
    /sbin/ethtool -K $interface sg off
```

```
done;
```



Result

```

iLO Integrated Remote Console - Server: bb-a.ak.ber.de.ffrl.de | iLO: ilo-bb-a-ak-ber-de.ffrl.de
Power Switch Virtual Drives Keyboard
[ 1282.133125] [<ffffffff8147c351>] ? napi_gro_frags+0xa1/0x1e0
[ 1282.133152] [<ffffffffa023d5a7>] ? mlx4_en_process_rx_cq+0x827/0xc20 [mlx4_e
n]
[ 1282.133187] [<ffffffff8101d285>] ? read_tsc+0x5/0x10
[ 1282.133211] [<ffffffffa023da86>] ? mlx4_en_poll_rx_cq+0x86/0x150 [mlx4_en]
[ 1282.133235] [<ffffffff8147be22>] ? net_rx_action+0x212/0x340
[ 1282.133258] [<ffffffff810772fa>] ? __do_softirq+0x11a/0x290
[ 1282.133281] [<ffffffff810775e5>] ? irq_exit+0x95/0xa0
[ 1282.133303] [<ffffffff8157bf9f>] ? do_IRQ+0x4f/0xe0
[ 1282.133325] [<ffffffff81579e2e>] ? common_interrupt+0x6e/0x6e
[ 1282.133347] <EOI>
[ 1282.133353] [<ffffffff810d73a7>] ? get_next_timer_interrupt+0xd7/0x250
[ 1282.133392] [<ffffffff8144571d>] ? cpuidle_enter_state+0xad/0x200
[ 1282.133415] [<ffffffff814456ec>] ? cpuidle_enter_state+0x7c/0x200
[ 1282.133438] [<ffffffff810b262b>] ? cpu_startup_entry+0x3ab/0x430
[ 1282.133462] [<ffffffff81928f54>] ? start_kernel+0x479/0x484
[ 1282.133484] [<ffffffff81928120>] ? early_idt_handler_array+0x120/0x120
[ 1282.133507] [<ffffffff81928120>] ? early_idt_handler_array+0x120/0x120
[ 1282.133530] [<ffffffff81928605>] ? x86_64_start_kernel+0x148/0x157
[ 1282.133552] Code: c2 0f 83 91 01 00 00 44 8b 04 24 89 2c 24 8b 6c 24 18 4c 89
64 24 40 4c 8b 54 24 20 4c 8b 5c 24 70 44 8b 64 24 28 41 39 ef 7c 3c <41> 8b 82
84 00 00 00 41 39 82 80 00 00 00 0f 85 60 04 00 00 45
[ 1282.133833] RIP [<ffffffff8146d0f4>] skb_segment+0x394/0xa50
[ 1282.133858] RSP <ffff880202203748>
[ 1282.133877] CR2: 0000000000000084
[ 1282.134264] ---[ end trace ddf426a7f9c9889 ]---
[ 1282.134332] Kernel panic - not syncing: Fatal exception in interrupt
[ 1282.134442] Kernel Offset: disabled
[ 1282.134509] ---[ end Kernel panic - not syncing: Fatal exception in interrupt

```




Solution: Downgrade

```
takt@bb-a.ix.dus.de.ffrl.de ~ % uname -a
```

```
Linux bb-a.ix.dus.de.ffrl.de 3.10.0-229.14.1.el7.x86_64 #1 SMP Tue Sep 15 15:05:51 UTC 2015 x86_64 GNU/Linux
```

```
takt@bb-a.ix.dus.de.ffrl.de ~ % cat /etc/debian_version
```

```
8.5
```



AS201701 / BER



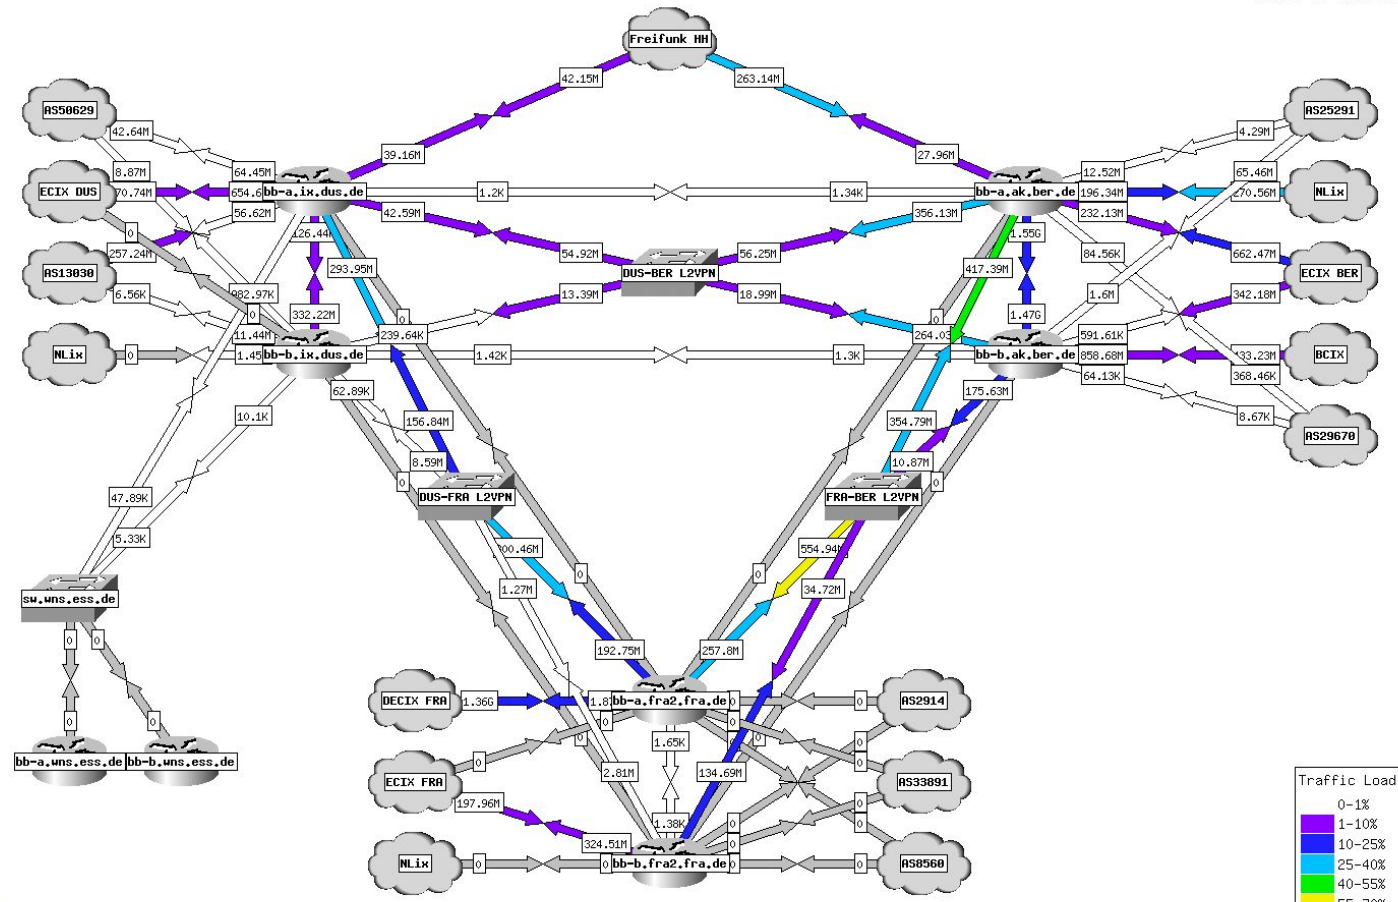
RUNNING FOR 10 DAYS



Quagga at work



Replaced with BIRD. No pain.





fra2.fra.de.ffrl.de





ix.dus.de.ffrl.de (and some more)





Todo Q3/2016





Traffic Level

AS201701 Overall Traffic





Many thanks to Sponsors and Supporters!

