

# HOW TO NOT LOSE ALL YOUR MESH NODES

## UPDATING INCOMPATIBLE ROUTING PROTOCOLS

PHILIPP

PAUL

FREIFUNK LÜBECK

NOVEMBER 17, 2020

1 Where we are

2 Why upgrading?

3 Planing-Stage

- Infrastructure
- Firmware

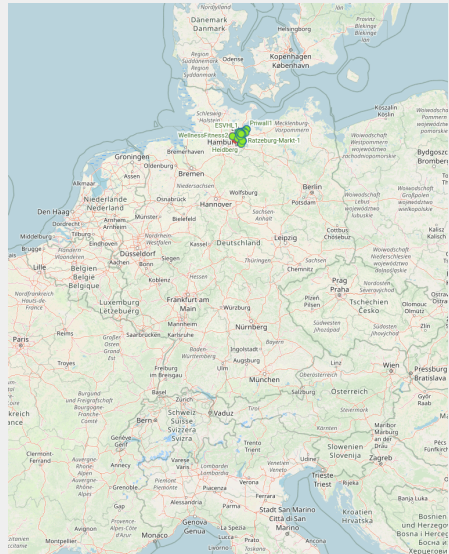
4 How we did it

- Infrastructure Side
- Firmware Side
  - Autoupdater
  - scheduled-domain-switcher

**WHERE WE ARE**

# Where we are

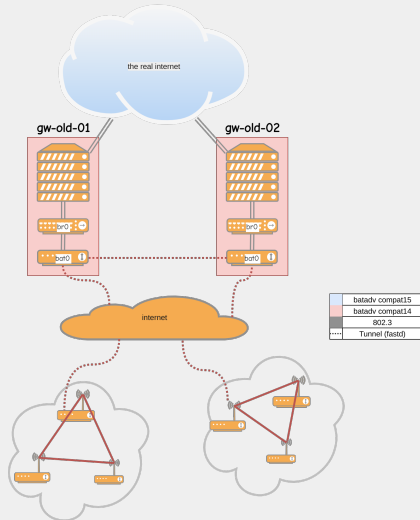
- Medium sized Freifunk Community
- Northern Germany
- ~250 Nodes
- ~400 simultaneous users at peak
- ~10.000 unique devices per day
- We are still quite new to Freifunk



# Current Infrastructure

- Gluon based
- 802.11s for meshing
- B.A.T.M.A.N. advanced for routing
- Gateways for internet exit and interconnecting meshclouds
- fastd for tunneling
- Autoupdates

# Network Architecture



# WHY UPGRADING?

# Advantages on upgrading batadv

- batman compat 14 and compat 15 are incompatible
- Batadv14 no longer supported by Gluon v2019.1 and newer
- new features in batadv 2013+
  - ▶ TVLV support => no more incompatibilitys
  - ▶ distributed ARP-cache optimizations
  - ▶ improved roaming
  - ▶ (soon) multicast optimizations
- Working commandline tools



# PLANING-STAGE

# Challenges

## ■ Goals

- ▶ No to minimal downtime
- ▶ Automated
- ▶ Testing should be easy
- ▶ Networks can run simultaneously
- ▶ clients from both networks can reach each other

## ■ Challenges:

- ▶ avoid unnecessary redundancy in infrastructure
  - Firmware-Server
  - Nodemap
  - Statistics
- ▶ keep our AS reachable
- ▶ Not cutting of nodes

## ■ Missions

- ⇒ Interconnect both networks (serverside)
- ⇒ Find a safe update procedure (nodeside)

# PLANING-STAGE


## INFRASTRUCTURE

- Can't load batadv-compat14 and batadv-compat15 at the same time  
⇒ additional gateway needed
- Strategies
  - ▶ Layer 3: IP Subnets for each compat
  - ▶ Layer 2: Bridge it!


# PLANING-STAGE

**FIRMWARE**


# Safe Update: How to cut off a node



1.



2.



3.

## ■ Challenges

- ▶ Timing (when is time?)
- ▶ Bad link quality
- ▶ Nodes with deactivated autoupdater

## ■ Strategies

- ▶ Systematic update: Update outside-in, server-side controlled (MIAU, FFUA)<sup>1</sup>
- ▶ Scheduled-Switch: all nodes switch config simultaneously at a specific time
- ▶ Node fallback in clientmode (never implemented?)
- ▶ non-gluon: Weimar Freifunk Community <sup>2</sup>

---

<sup>1</sup><https://freifunk.in-kiel.de/blog/2019/04/13/BATMAN-migration.html>

<sup>2</sup><https://www.youtube.com/watch?v=IC-etQwLYAE&list=PL3bvPCw5QCLJ-VJPamVeQx-UPNBVyaopj&index=6>

# HOW WE DID IT



- infrastructure-side: Layer2 Bridge
  - ▶ very easy to set up
  - ▶ avoids redundant infrastructure
  - ▶ Nodes/server reachable from each batadv network
  - ▶ allows easy testing
- Firmware: domains + scheduled-domain-switcher
  - ▶ worked nice with ibss -> 802.11s switch
  - ▶ has fallback mechanism
  - ▶ requires good planing
  - ▶ requires accurate date

# HOW WE DID IT

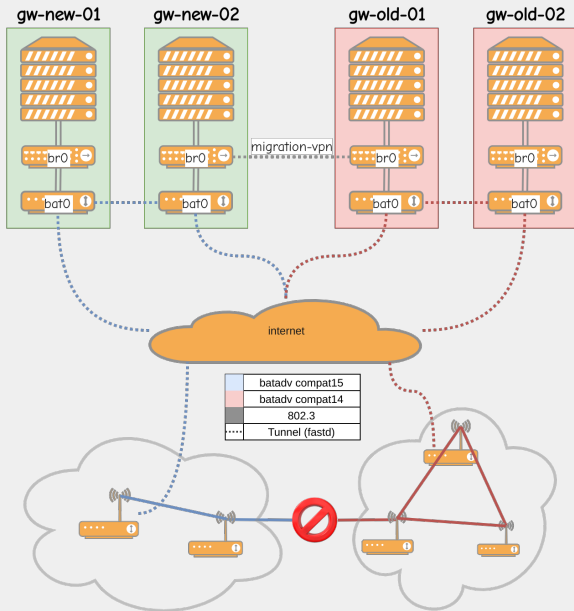
## INFRASTRUCTURE SIDE

# First attempt

- "just connect two routers, lol"
- worked surprisingly well
- Map, Statistics, DHCP, internet access just worked (kinda)



■ 2 new Gateways (thx FFHH and FFOH)



- Layer8
  - ▶ No documentation for the old gateway setups
  - ▶ Who has access to \$account?
- MTU migration vpn
- DHCP
- NDP + Gluon + multicast

- **DHCP requests leaking to different gateways**
- batman's *gateway-feature*
  - ▶ Nodes can be configured as a gateway
  - ▶ batman routes specific broadcast packets to the nearest gateway e.g. DHCP
- DHCP requests leave the mesh at bridge
- gets broadcasted to another server via migration-vpn
- ⇒ block DHCP

- **Incoming Packets not routed into the mesh**
- `ip neigh` shows a lot addresses as INCOMPLETE
- NDP
  - ▶ NDP works with multicast
  - ▶ Gluon blocks some multicast traffic
  - ▶ batman intercepts MLD for multicast optimizations
- disabling multicast-snooping on gateways fixed it

# Blocked traffic on migration-vpn

- DHCP
- some ICMPV6 types
  - ▶ ICMP redirect (annoying for testing)
  - ▶ Router Advertisement



# HOW WE DID IT

## FIRMWARE SIDE

1. Get new firmware to the nodes  
⇒ Autoupdater
2. Switch from batadv-legacy to batadv simultaneously  
⇒ Scheduled-Domain-Switch

- Central update Server
- Manifest file with signatures
- Incremental rollout

# What are "domains" exactly?

- `site.conf`
- default place for most config variables

```
1 default_domain = 'default',
2 site_name = 'Freifunk Lübeck',
3 site_code = 'ffhl',
4
5 wifi24 = {
6     channel = 1,
7 },
8 wifi5 = {
9     channel = 44,
10    outdoor_chanlist = '96-140',
11 }
```

# What are "domains" exactly?

- domains/default.conf
- Often used to split a large network into multiple smaller
- Network specific config variables
- Switched with uci, webinterface or scheduled

```
1 mesh = {
2     vxlan = false,
3     batman_adv = {
4         routing_algo = 'BATMAN_IV_LEGACY',
5     },
6 },
7 prefix4 = '10.130.0.0/20',
8 prefix6 = 'fdef:ffc0:3dd7::/64',
9 extra_prefixes6 = { '2001:67c:2d50::/48' },
```

# What are "domains" exactly?

```
1 mesh = {
2     vxlan = false,
3     batman_adv = {
4         routing_algo = 'BATMAN_IV_LEGACY',
5     },
6 },
7 prefix4 = '10.130.0.0/20',
8 prefix6 = 'fdef:ffc0:3dd7::/64',
9 extra_prefixes6 = { '2001:67c:2d50::/48' },
```

Do 03 Sept

22:00:00

=>

```
1 mesh = {
2     vxlan = false,
3     batman_adv = {
4         routing_algo = 'BATMAN_IV',
5     },
6 },
7 prefix4 = '10.130.0.0/16',
8 prefix6 = 'fdef:ffc0:3dd7::/64',
9 extra_prefixes6 = { '2001:67c:2d50::/48' },
```

# Scheduled-Domain-Switch

- Gluon-Package
- Switches to `target_domain` if the `switch_time` has passed
- Fallback: regulary pings `connections_check_targets` and switches configuration if offline for more than `switch_after_offline_mins`



```
1 domain_switch = {
2   target_domain = 'bat15_migration',
3
4   switch_after_offline_mins = 120,
5   switch_time = 1599163200, -- Do 03 Sept 22:00:00 CEST 2020
6
7   connection_check_targets = {
8     '2001:4860:4860::8888',
9     '2001:4860:4860::8844',
10  },
11 },
```

# Timeline

## ■ Normal Update Cycle with Beta phase not possible

Day

- n Build Firmware and test it
- 10 Sign firmware and upload it  
Beta release
- 9 All beta nodes still alive?  
Stable release, incremental rollout
- 3 wait more days for stragglers
- 0 Domain-Switch, pray to your gods



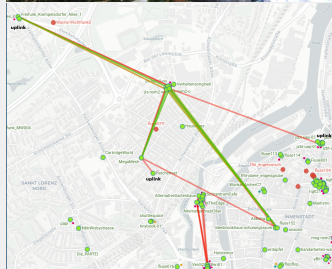


# Issues - Example JLS

- Link TQ really bad
- some nodes upgraded some didn't
- Link was dead for 120mins (not even one ping)
- some Nodes switched domain, others couldn't download new firmware

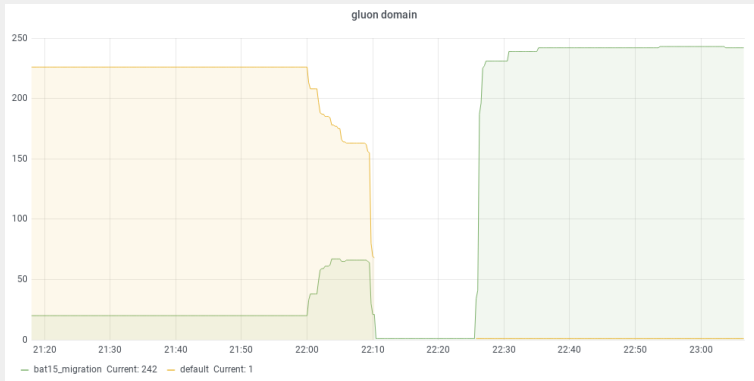
⇒ set `switch_after_offline_mins` to a longer interval

⇒ less days of incremental rollout



# What now?

■ Wait.



# Thx, Questions?

- [luebeck.freifunk.net](http://luebeck.freifunk.net)
- [@freifunkluebeck](https://twitter.com/freifunkluebeck)
- [#ffhl:matrix.org](https://matrix.org/#ffhl) or [#ffhl](https://freenode.net) on freenode