ffmuc presents #FFMEET

Technical challenges as well as social challenges during COVID-19 when running an open source conference system

Tech talk at UKNOF 2020

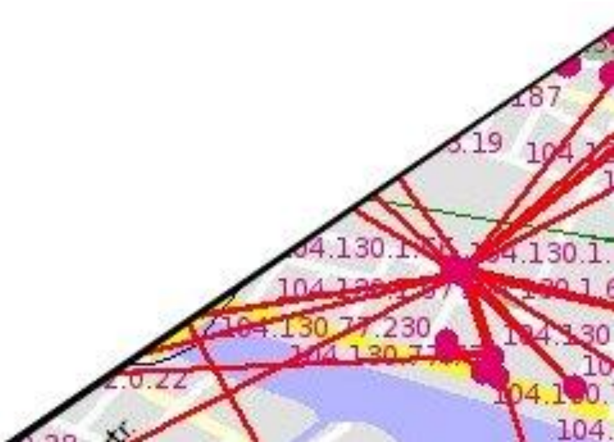

Who are we?

awlnx

- Annika Wickert
- Senior Network Engineer @AS51324
- Twitter @awlnx

krombel

- Matthias Kesler
- IT Consultant
- Twitter @kr0mbel

Why did we set up jitsi?

- An upcoming Freifunk Meeting was about to get canceled because of Corona
- The architecture of jitsi looked straight forward and easy to scale
- In the past we used NextCloud talk but we had issues (max 6. users per conference)

– We needed something different for at least 20 - 30 people

– jitsi seemed easy and straightforward to install

  - “apt install jitsi-meet” => Done ¯_(ツ)_/¯

First test - looks promising

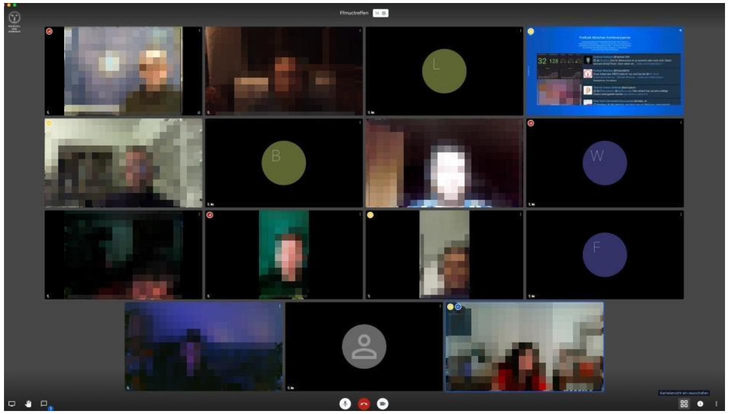

- Testing with some people of Freifunk Munich
- Worked like a charm, with a handful of people
- The idea was born:
  - Maybe other people have the same problem?
  - Why not open it for the public?
  - Maybe teachers, healthcare workers, etc. need that, too?

We need insights! Let’s monitor!

- How many users are on the platform?
- What’s the impact on the Hardware/VMs?
- Do we have bottlenecks? If so: Where?

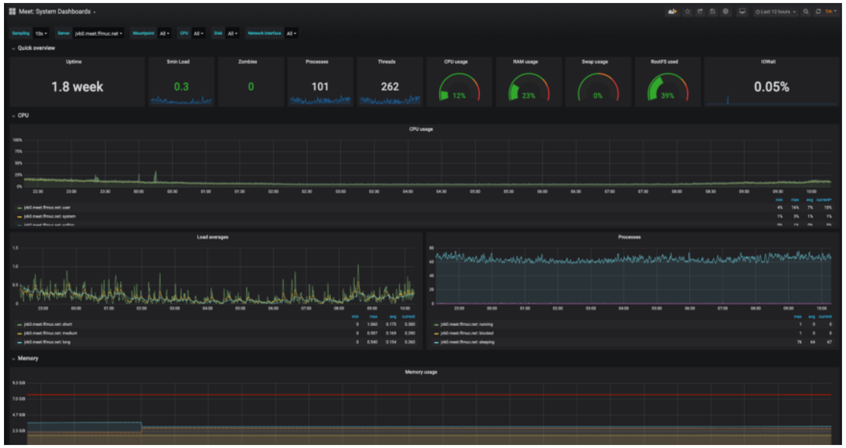

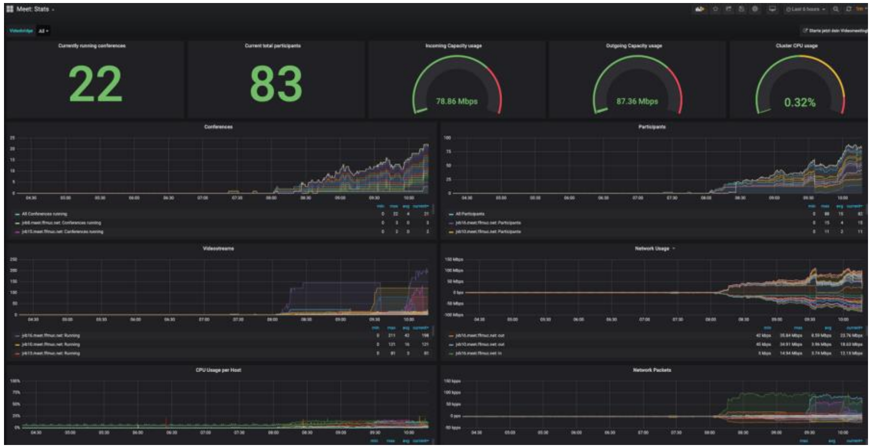

muenchen.freifunk.net

The project gains attention

- Interviews
- BR, DigitalCourage, Focus Online, Stimme, SZ, …
- appearing on several public server lists

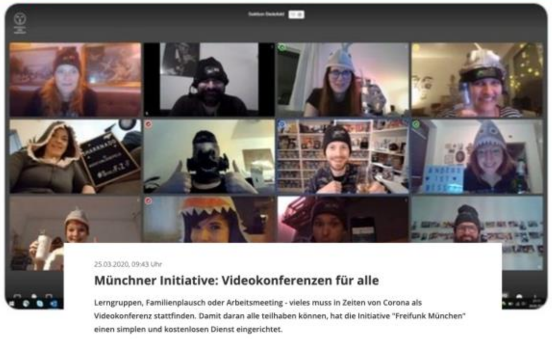

25.03.2020, 09:43 Uhr

Münchner Initiative: Videokonferenzen für alle

Lerngruppen, Familienplausch oder Arbeitsmeeting - vieles muss in Zeiten von Corona als Videokonferenz stattfinden. Damit daran alle teilhaben können, hat die Initiative "Freifunk München" einen simplen und kostenlosen Dienst eingerichtet.

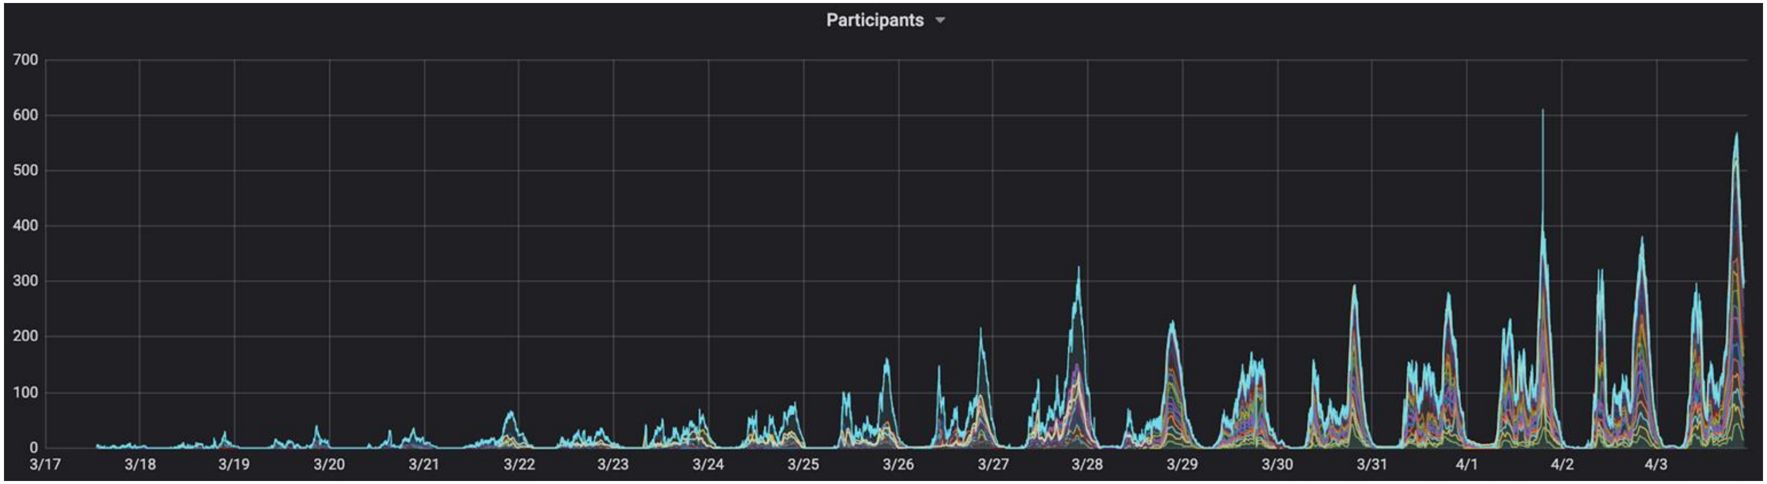

Holidays are over … school starts - 20.04.2020

muenchen.freifunk.net

- All time record before was around 600 Users

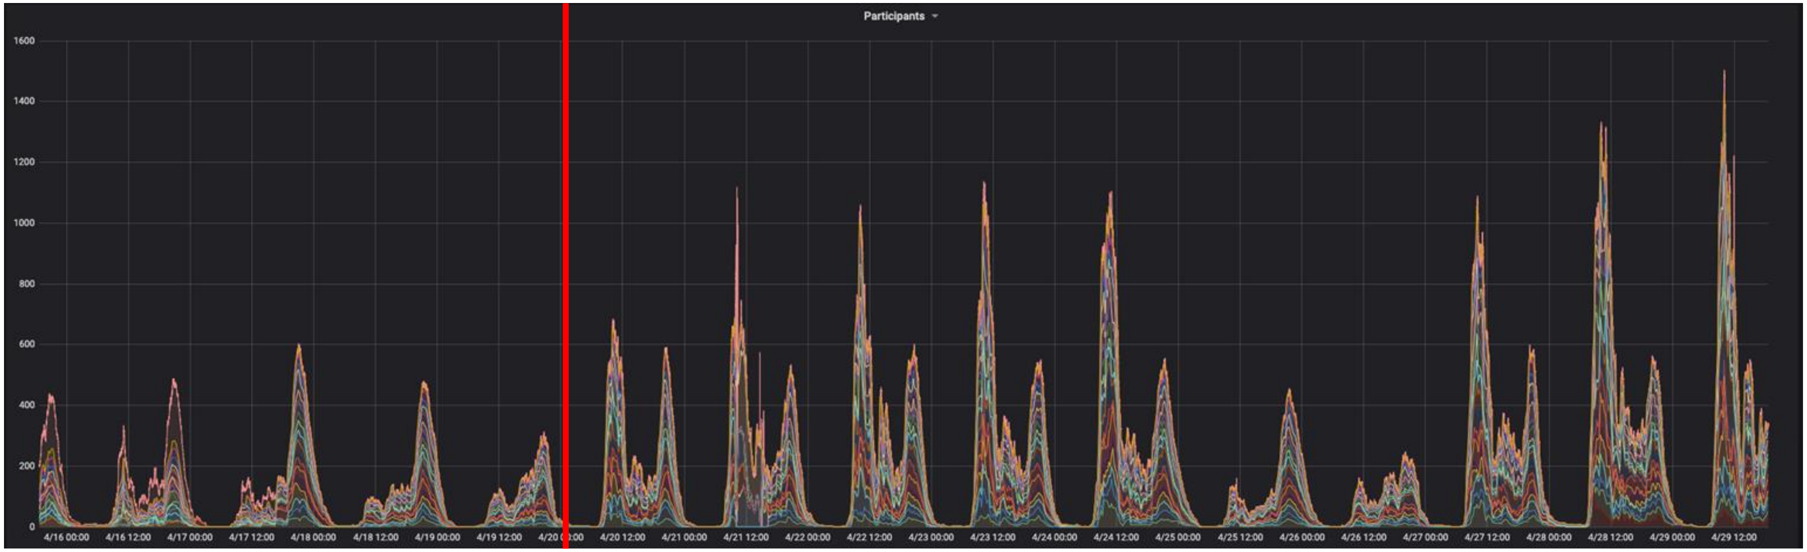

Some unexpected things happen

People started depending on our system

- Schools & universities (not only DE)
- Small companies
- Open-Source-Communities (repeatedly) meet on our platform
  - Typo3, BitCoin, dotfmp
- In Peru and they use #FFMEET as their primary conf system

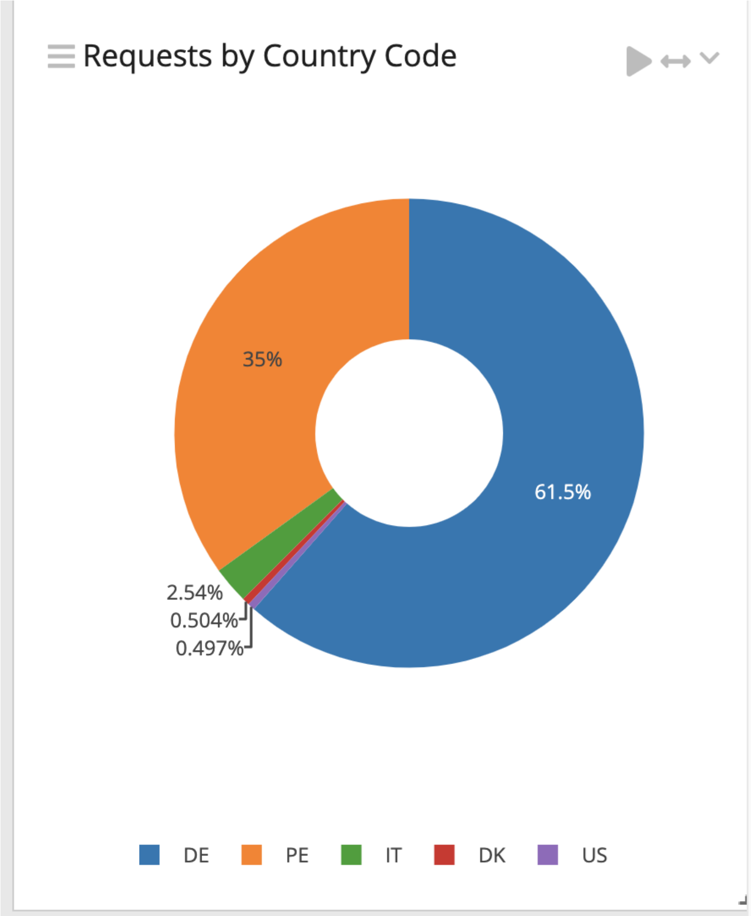

People tend to be in firewalled networks => TURN

- Especially after school started again we got many people complaining about video blackouts etc.
- Reason: Firewalls/NAT/Proxy so no UDP nor direct connection possible
- sysctl -w net.ipv4.tcp_congestion_control=bbr
- Solution: TURN (Traversal Using Relays around NAT) RFC 5766 and 6156
  - Acts as man in the middle to forward packets back and forth

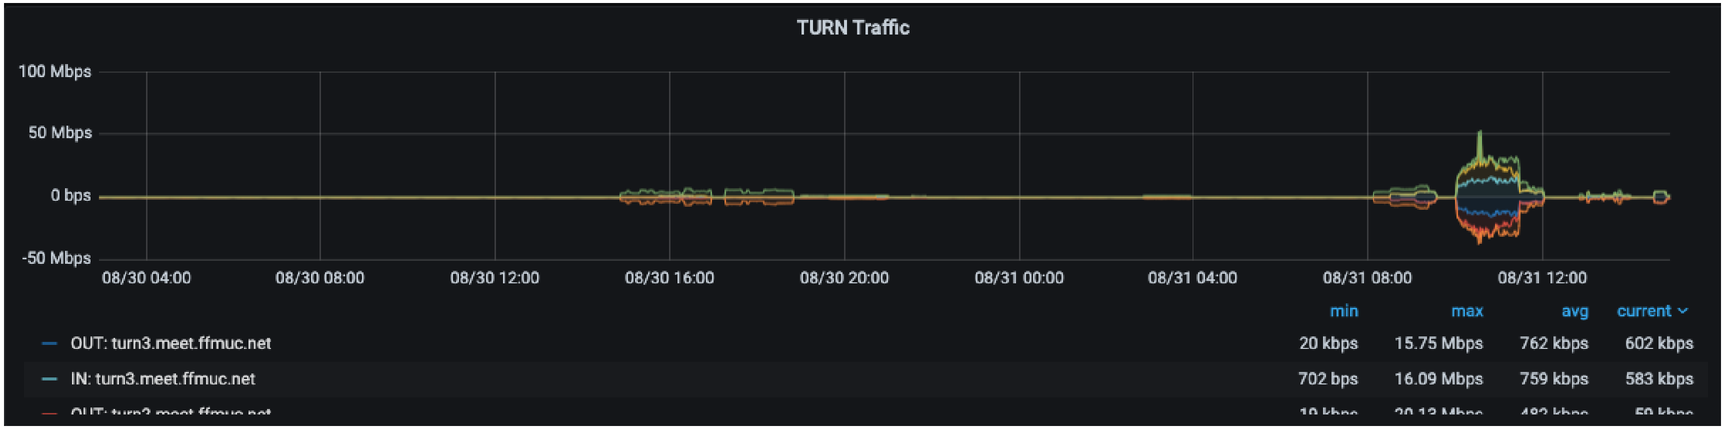

	min	max	avg	current
OUT: turn3.meet.ffmuc.net	20 kbps	15.75 Mbps	762 kbps	602 kbps
IN: turn3.meet.ffmuc.net	702 bps	16.09 Mbps	759 kbps	583 kbps
OUT: turn2.meet.ffmuc.net	10 kbps	20.13 Mbps	483 kbps	50 kbps

Of course: It did not work as we intended it to do...

- Ugly “Feature” in browsers
  => They prefer the first TURN server in the list …
  => Patch Jitsi to randomize the sequence

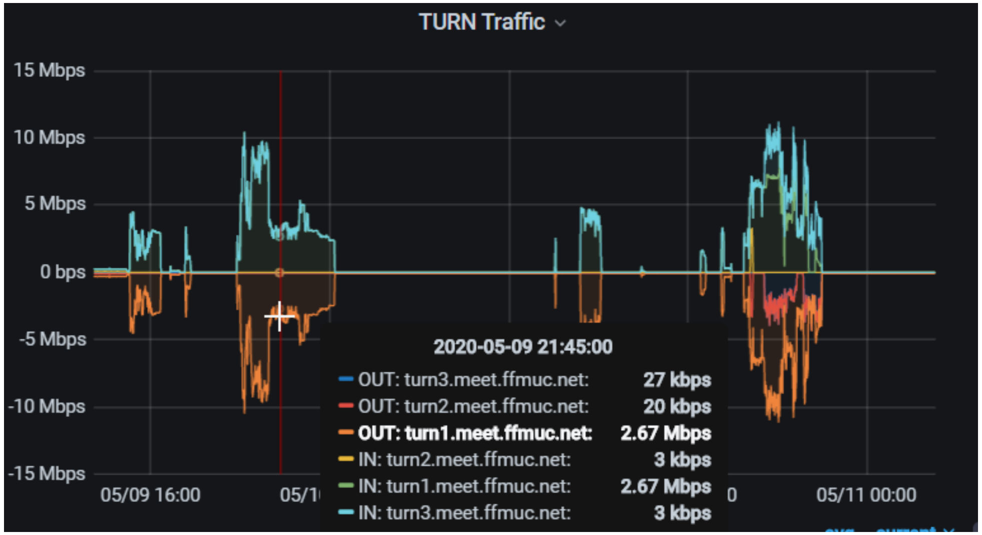

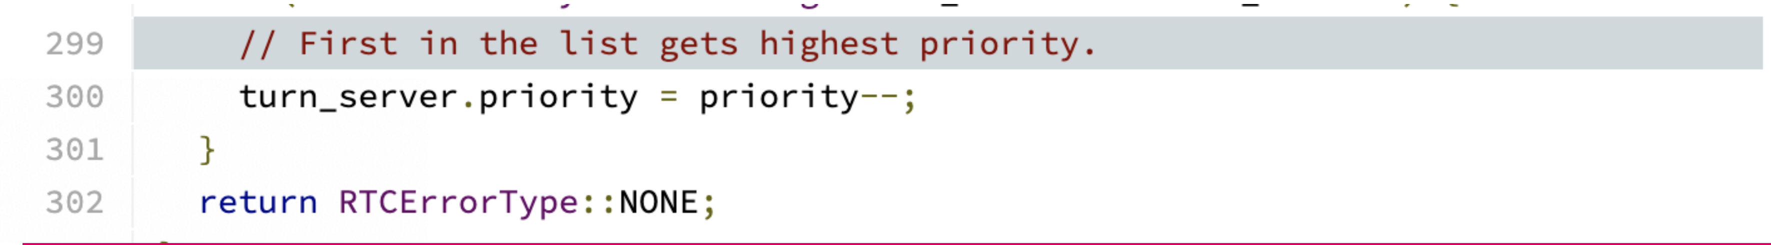

```
    // First in the list gets highest priority.
    turn_server.priority = priority--;
  }
  return RTCErrorType::NONE;
```

21.04.2020 - Another record! Sadly also another bug

muenchen.freifunk.net

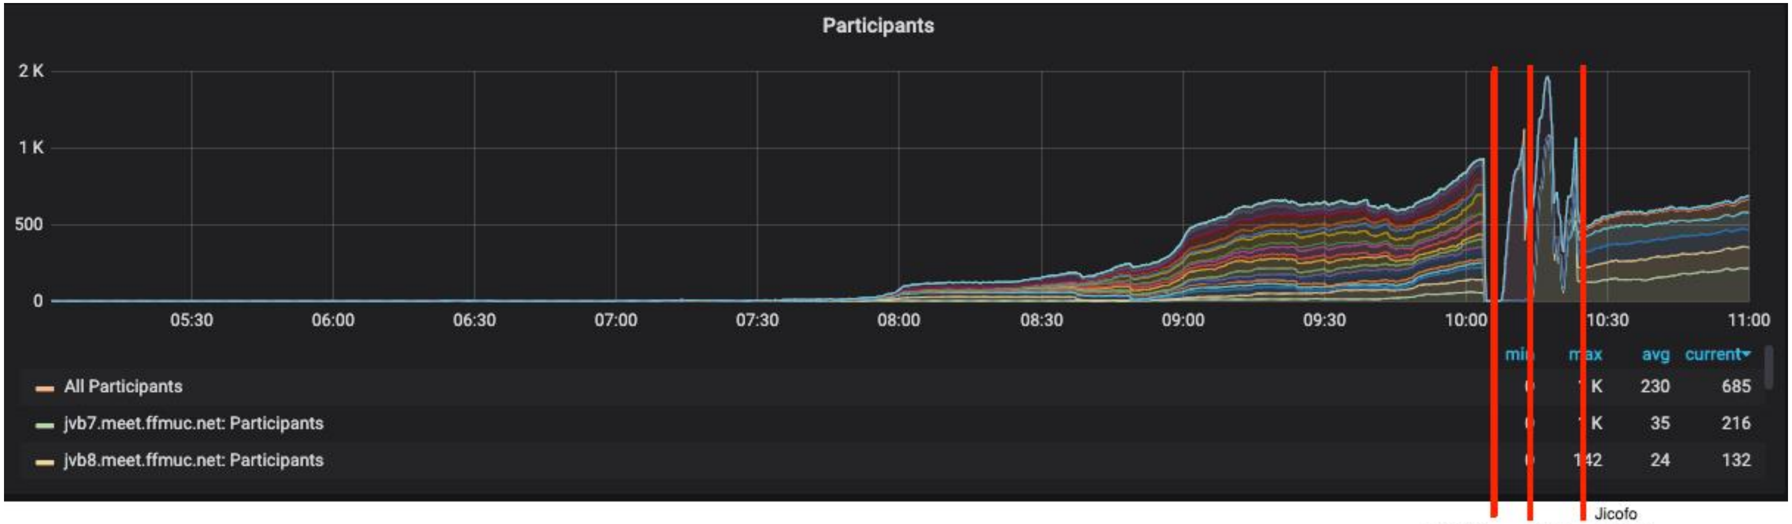

	min	max	avg	current▾
▬ All Participants	0	1 K	230	685
▬ jvb7.meet.ffmuc.net: Participants	0	1 K	35	216
▬ jvb8.meet.ffmuc.net: Participants	0	142	24	132

Prosody too many connections drops users

Jicofo Starts to overload videobridges

Jicofo Update Service back Online

- 20 min. to find the issue, isolate and fix it!
- Unfortunately there was an other bug hiding after the crash
- We run at reduced capacity for a few hours but as we have so many servers that was not an issue.
- Crash of prosody crashed some important DBs :/ ⇒ Loadbalancing failed

We run on VMs with high PPS and many interrupts

- Oversubscribed hosts are an issue for us => pure hardware performs best
- We do many PPS and many interrupts on a single videobridge
- Tuning buffers and memory of the kernel

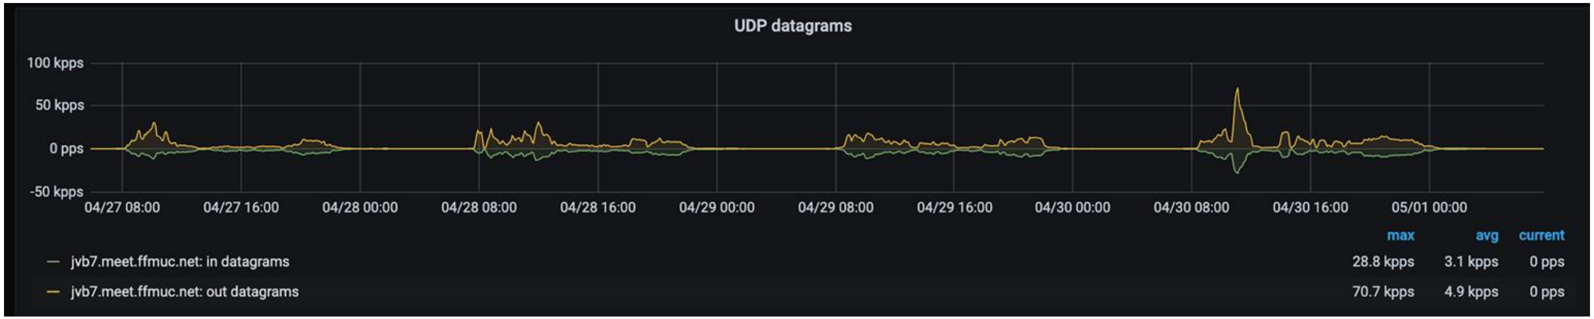

UDP datagrams

	max	avg	current
— jvb7.meet.ffmuc.net: in datagrams	28.8 kpps	3.1 kpps	0 pps
— jvb7.meet.ffmuc.net: out datagrams	70.7 kpps	4.9 kpps	0 pps

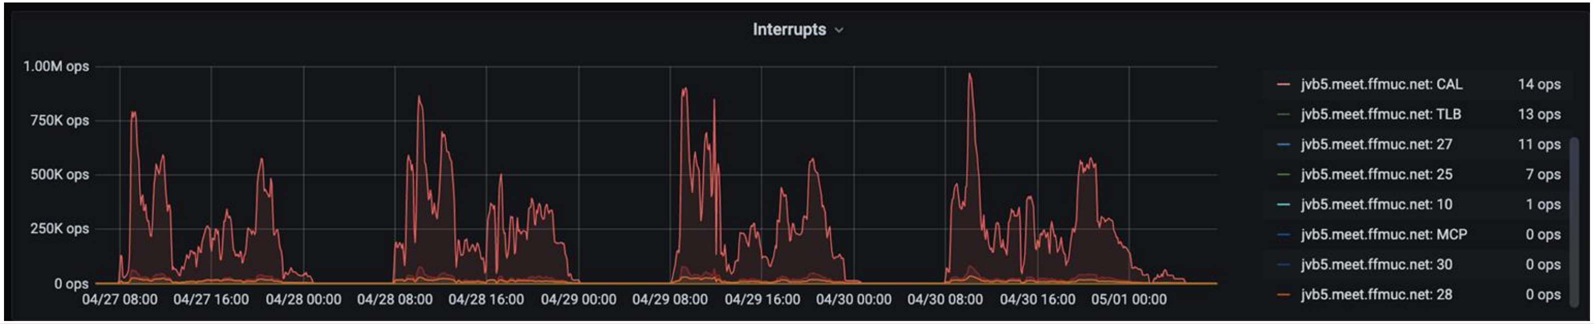

Interrupts

— jvb5.meet.ffmuc.net: CAL	14 ops
— jvb5.meet.ffmuc.net: TLB	13 ops
— jvb5.meet.ffmuc.net: 27	11 ops
— jvb5.meet.ffmuc.net: 25	7 ops
— jvb5.meet.ffmuc.net: 10	1 ops
— jvb5.meet.ffmuc.net: MCP	0 ops
— jvb5.meet.ffmuc.net: 30	0 ops
— jvb5.meet.ffmuc.net: 28	0 ops

Scaling the infrastructure

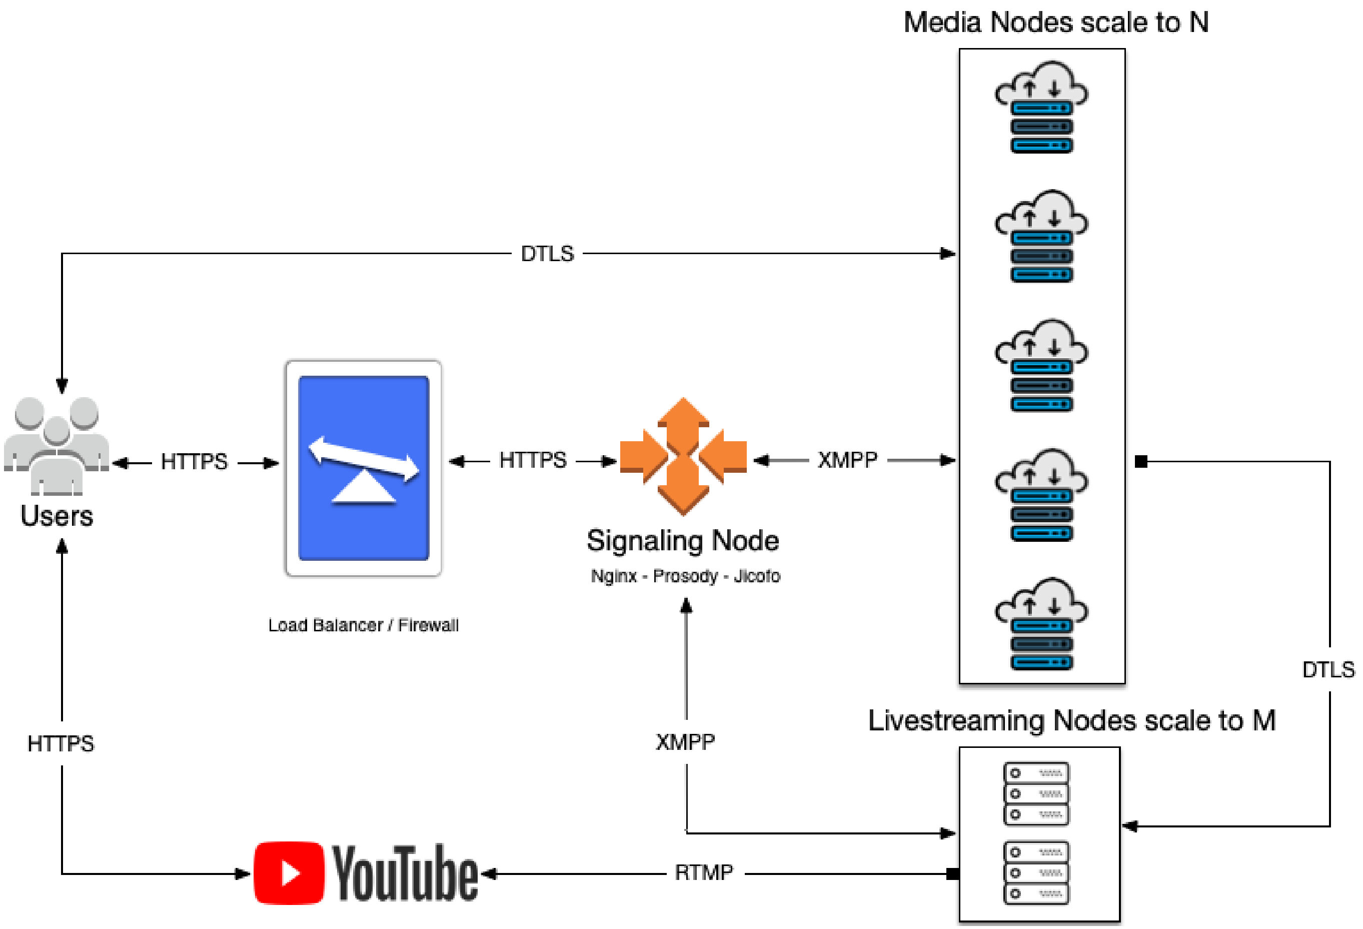

People ask for bigger meetings

- Octo solves the problem of only one videobridge per conference (server-side is no longer an issue)
- Sadly video traffic is unencrypted between videobridges => we need a solution for that
- As we do high packet rates we need a full-mesh VPN without a central point where all traffic has to pass

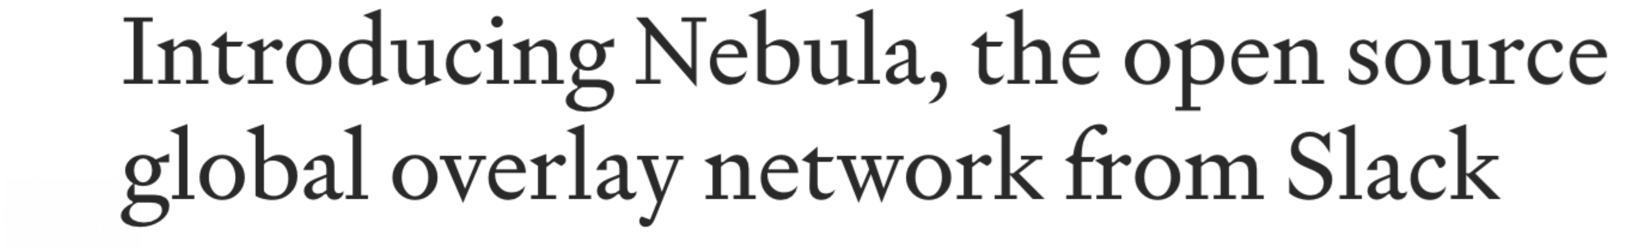

Introducing Nebula, the open source global overlay network from Slack

Octo on all bridges

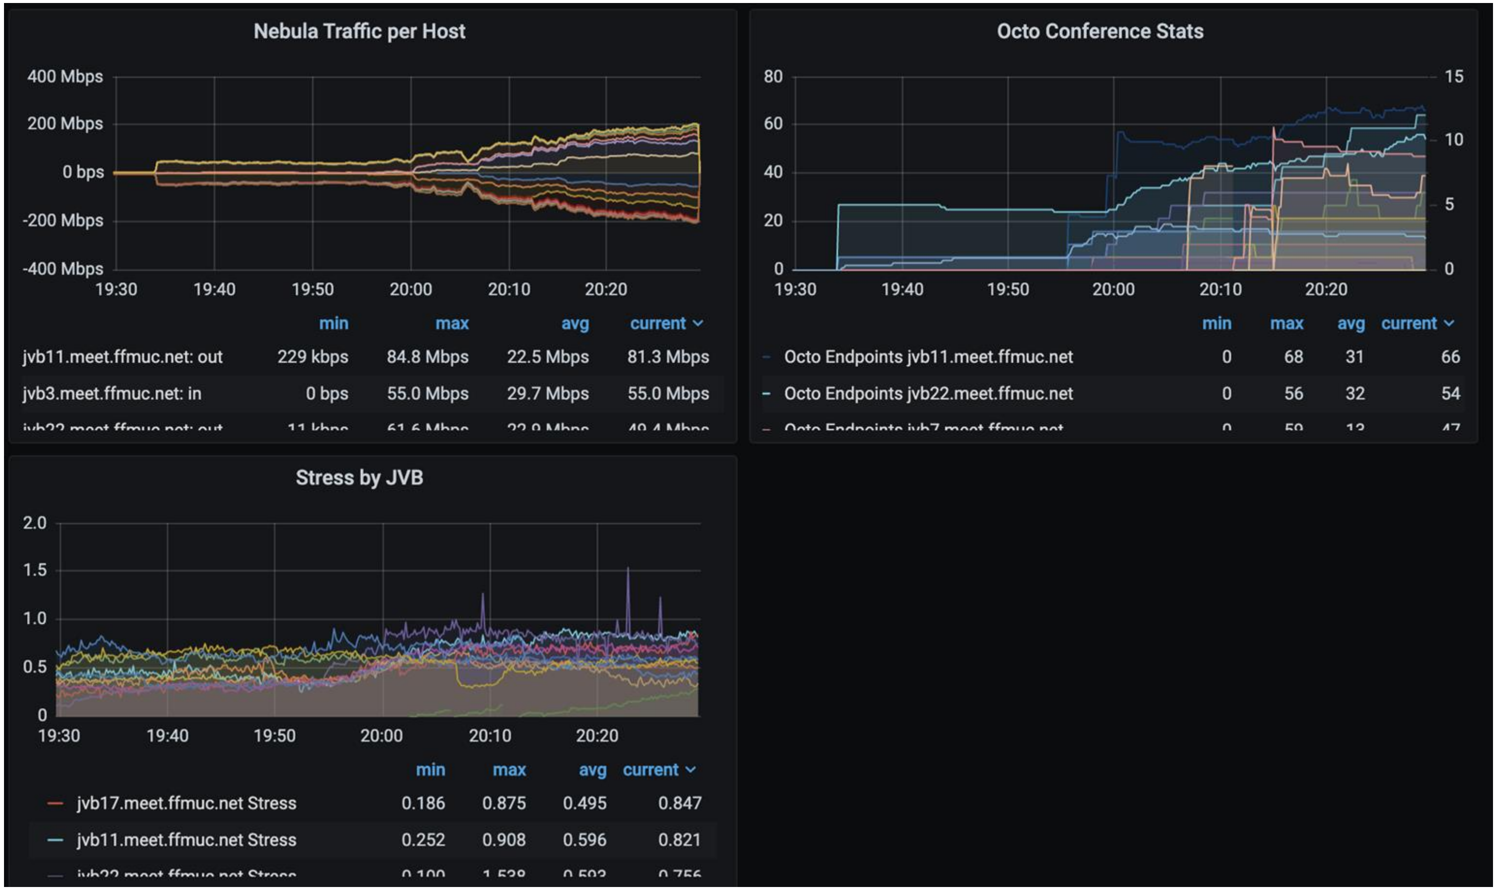

Nebula Traffic per Host

	min	max	avg	current ⌄
jvb11.meet.ffmuc.net: out	229 kbps	84.8 Mbps	22.5 Mbps	81.3 Mbps
jvb3.meet.ffmuc.net: in	0 bps	55.0 Mbps	29.7 Mbps	55.0 Mbps
jvb22.meet.ffmuc.net: out	11 kbps	61.6 Mbps	22.9 Mbps	49.4 Mbps

Octo Conference Stats

	min	max	avg	current ⌄
— Octo Endpoints jvb11.meet.ffmuc.net	0	68	31	66
— Octo Endpoints jvb22.meet.ffmuc.net	0	56	32	54
— Octo Endpoints jvb7.meet.ffmuc.net	0	59	13	47

Stress by JVB

	min	max	avg	current ⌄
— jvb17.meet.ffmuc.net Stress	0.186	0.875	0.495	0.847
— jvb11.meet.ffmuc.net Stress	0.252	0.908	0.596	0.821
— jvb22.meet.ffmuc.net Stress	0.100	1.538	0.593	0.756

How it evolved

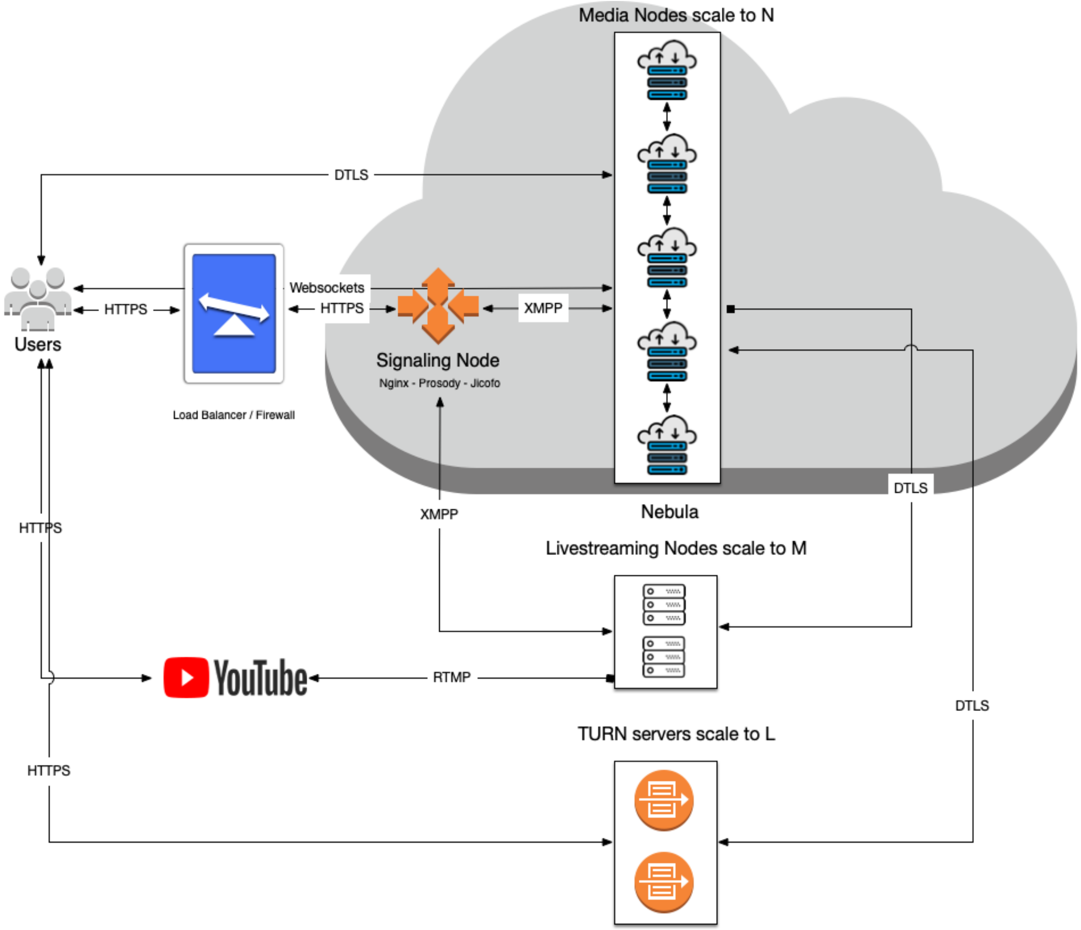

Some figures? Since we started #FFMEET

- ~ 1.109 TB videotraffic
- ~ 1.407.653.191 HTTPS Requests
- ~ 360 CPU Cores / 720GB RAM
- ~ 1.100.000 Unique Users
- ~ 700 tickets solved (schools, charity projects, individuals etc.)
- 2 Admins and at least 15 people in the background for translations, press, wiki etc.

Thank you very much!

A special thanks to everyone who helped writing/translating the FAQs.

Also thanks to all the users and of course our supporters:

Alibaba Cloud

CLOUDFLARE

ESPLENDOR GROUP

level66.network

plusserver

freifunk-franken.de

karlsruhe.freifunk.net

neander funk

Open Telekom Cloud

Backup Slides

muenchen.freifunk.net

We are running on the internet aren’t we?

- From time to time there are issues in the internet beyond our control (route leaks, connectivity issues between providers, maintenance)
- We have many servers and providers so we can easily mitigate connectivity issues

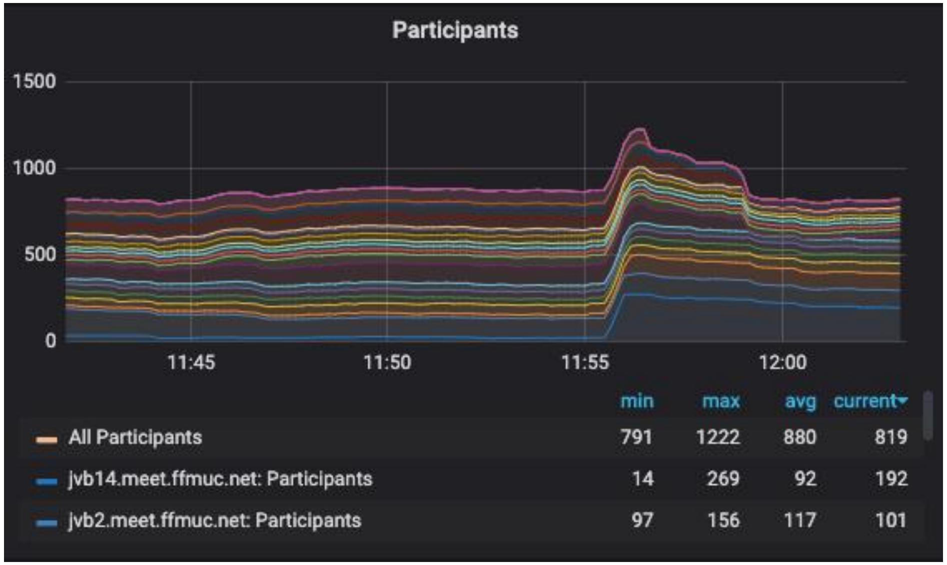

	min	max	avg	current▾
▬ All Participants	791	1222	880	819
▬ jvb14.meet.ffmuc.net: Participants	14	269	92	192
▬ jvb2.meet.ffmuc.net: Participants	97	156	117	101

We got more popular so DDOS attacks happen

- People attacked our infrastructure and didn’t stop when they couldn’t bring down #FFMEET
- They started to attack other parts of our infrastructure

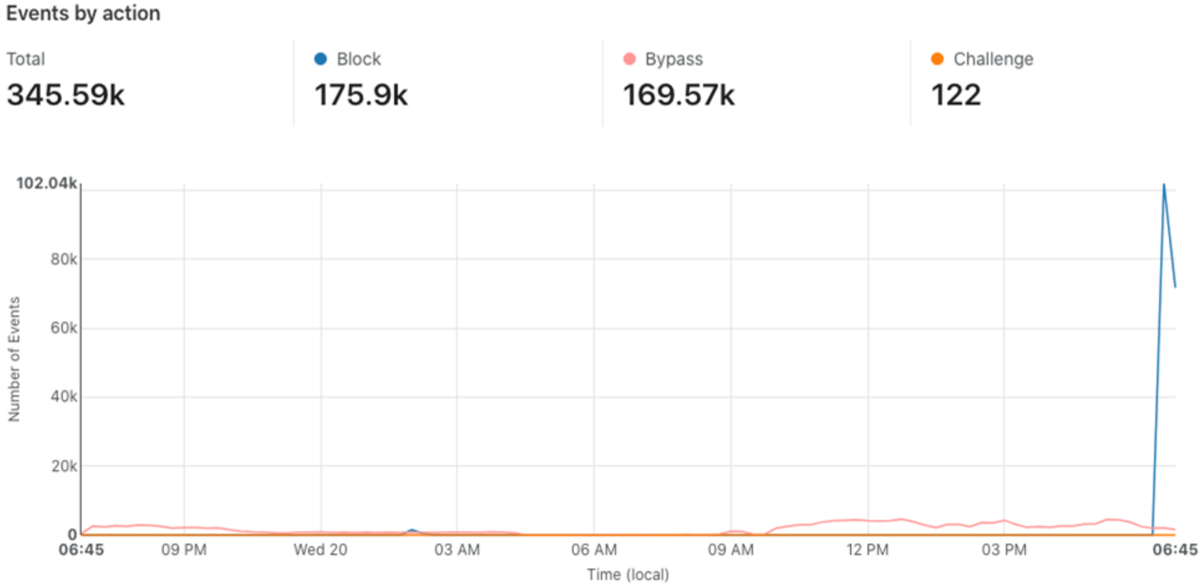

Events by action

Total	● Block	● Bypass	● Challenge
345.59k	175.9k	169.57k	122

Problems start - Prosody

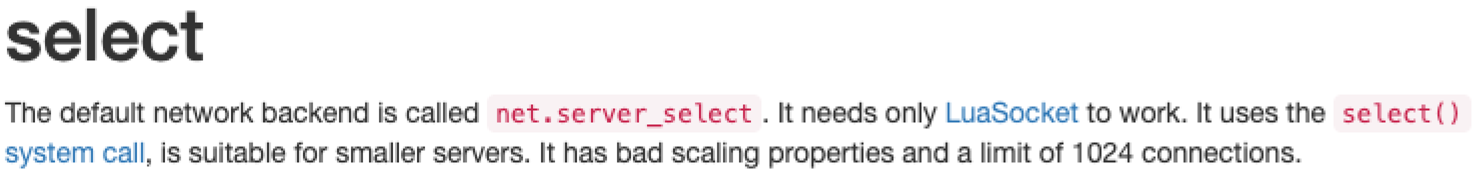

select

The default network backend is called `net.server_select`. It needs only LuaSocket to work. It uses the `select()` system call, is suitable for smaller servers. It has bad scaling properties and a limit of 1024 connections.

- Change network backend to “epoll” - No more 1024 connections as a limit

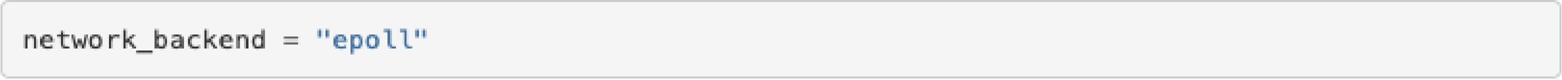

```
network_backend = "epoll"
```

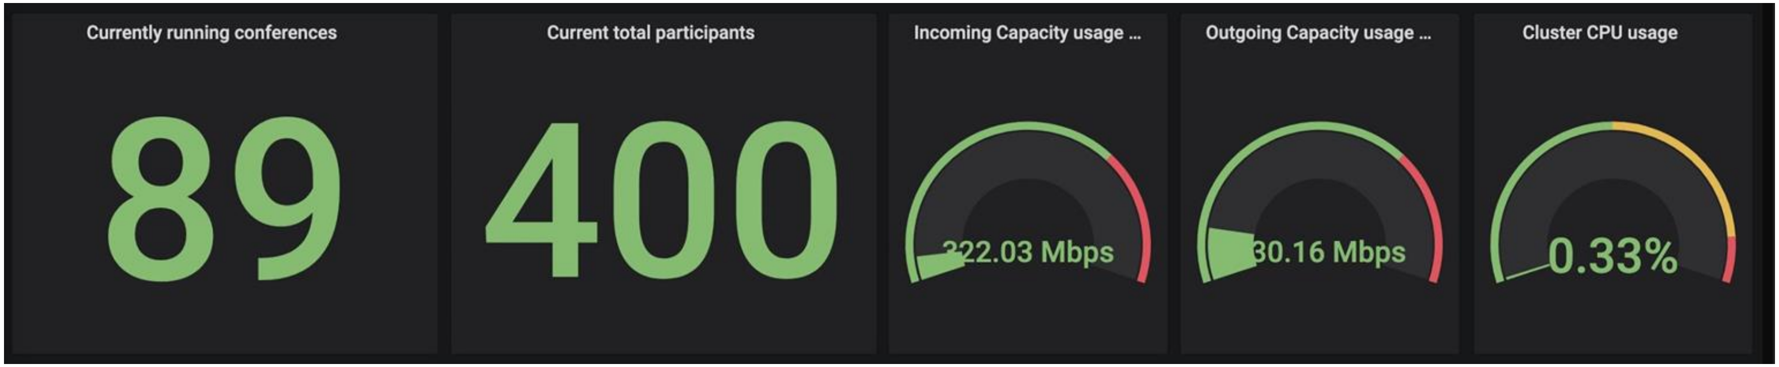

Currently running conferences
89
Current total participants
400
Incoming Capacity usage ...
222.03 Mbps
Outgoing Capacity usage ...
30.16 Mbps
Cluster CPU usage
0.33%

Don’t tune only one part of the stack - nginx

- 502 all over the place
- Raise the number of nginx workers
- Raise max open files limit

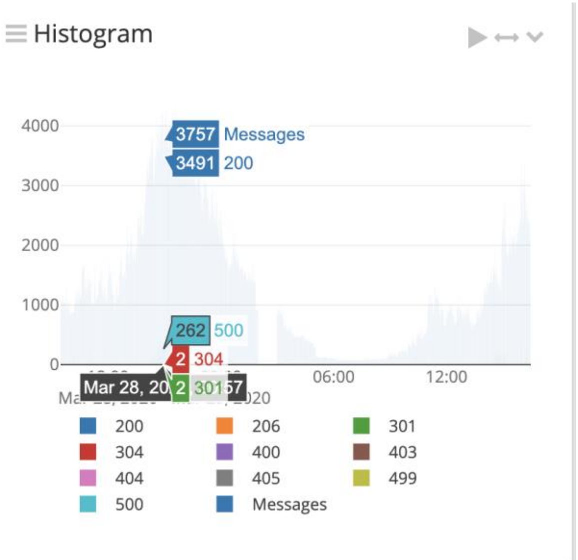

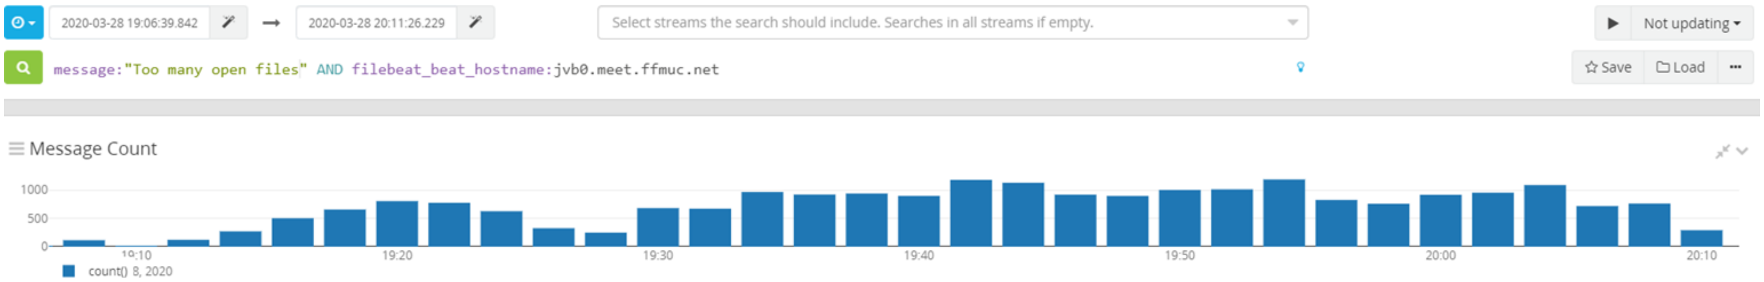
2020-03-28 19:06:39.842 → 2020-03-28 20:11:26.229

Select streams the search should include. Searches in all streams if empty.

```
message:"Too many open files" AND filebeat_beat_hostname:jvb0.meet.ffmuc.net
```

How do we update without user impact?

- We can upgrade the videobridges one by one
- We query the videobridge API and wait for 0 users then run our custom upgrade script
- Upgrades of videobridges just happen in the background during the day/night
- Jicofo/Prosody upgrades … sadly we have to do update them manually but we don’t need to update them very often