

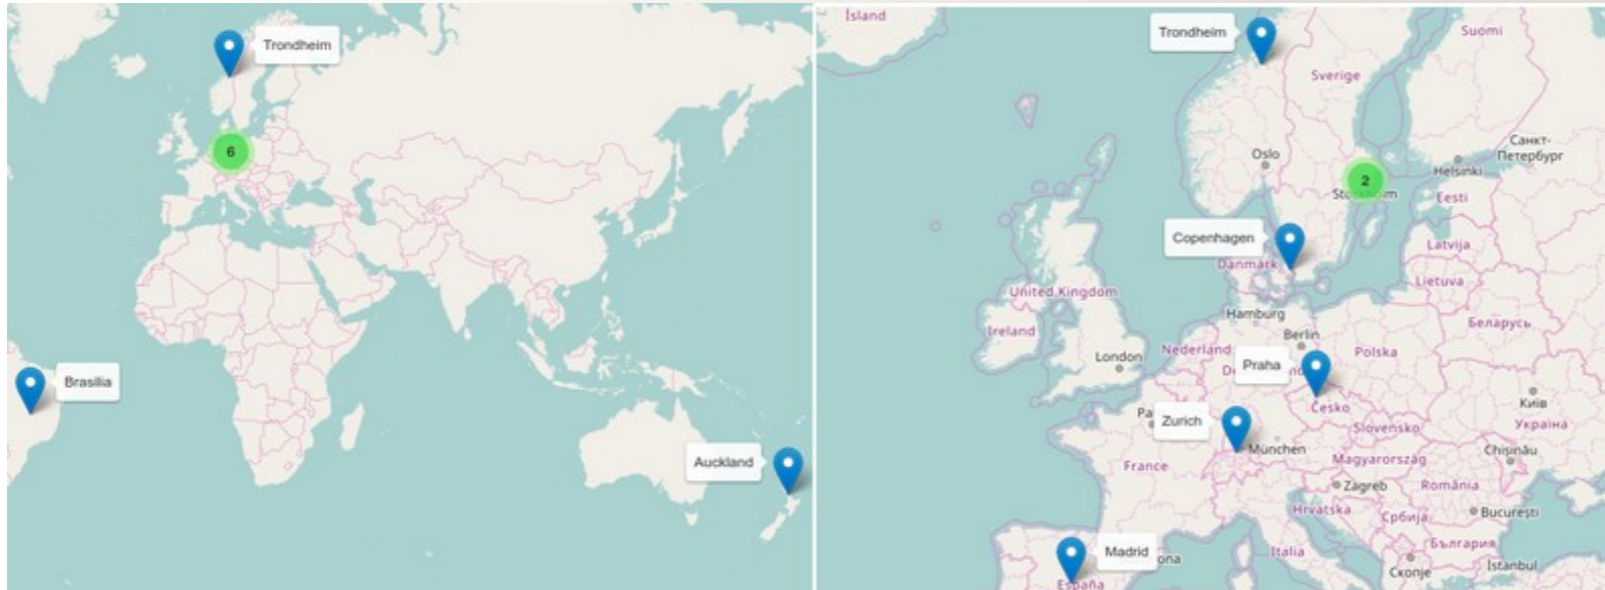
# Micro dependability measurements

## RIPE MAT WG

Olav Kvittem  
UNINETT  
2019-10-17

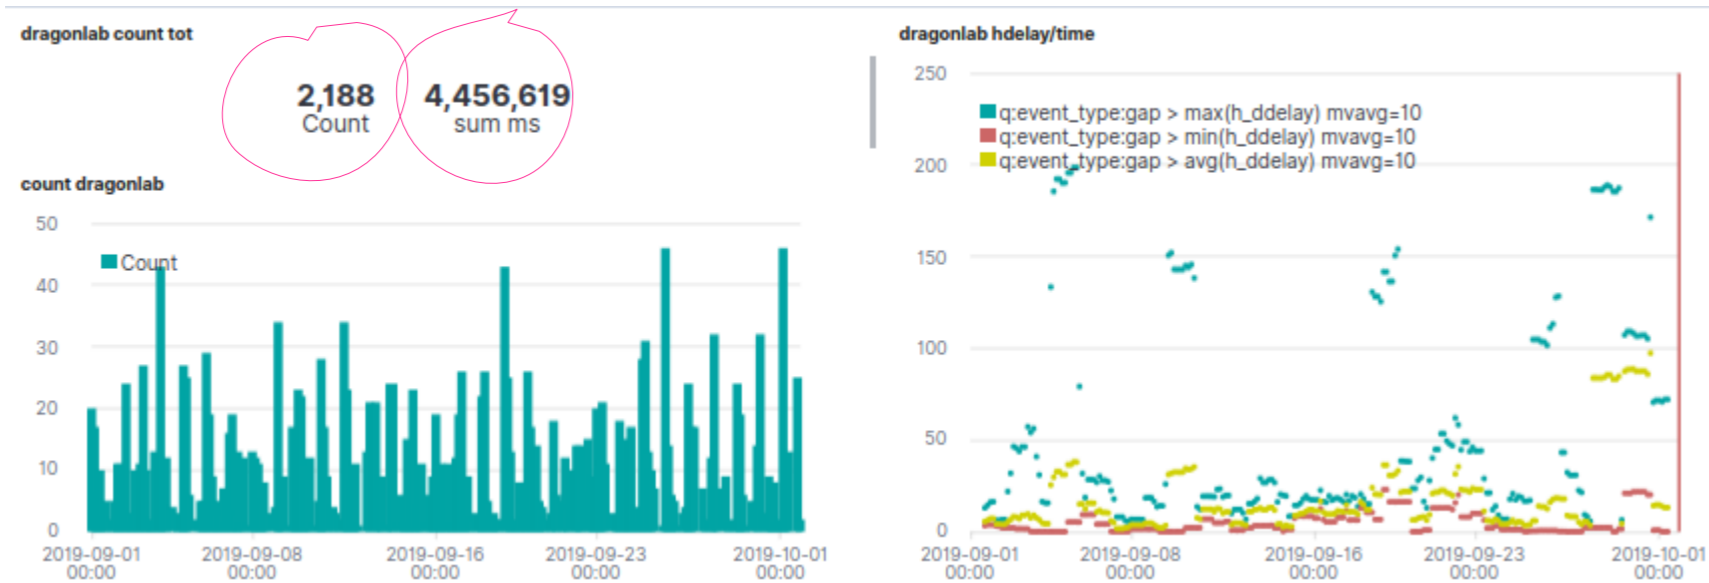


# The Dragonlab participants





- Participants
  - U.Auckland, RNP(Brasil), Cernet(out but co-initiator)
  - Nordunet, Sunet, Uninett, NTE(commercial)
  - Switch, Cesnet, RedIris, Aarnet
  - Cloud : Amazon, Azure, Google
- NTNU university : PHD's and MSC's
- Measurement setup
  - 100 packets per second with timestamps
  - Traceroute
  - analyze packet behaviour around loss events into ELK
    - 50 ms ITU failover time

# Outage stats



- Sep 2019 – lost 1.5 hours in 2000 gaps
  -
- Might be the same actual outage – hitting many
  - Site power loss – 12 minutes
  - Router upgrade 2 minutes
  - Fiber instability (10\*1 min)
  - Dual router reboot – 20 minutes..
- Oddities
  - Old packets sent on when link comes up – hours afterwards
  - Tails of routing loops – low TTLs

# September outages > 500ms

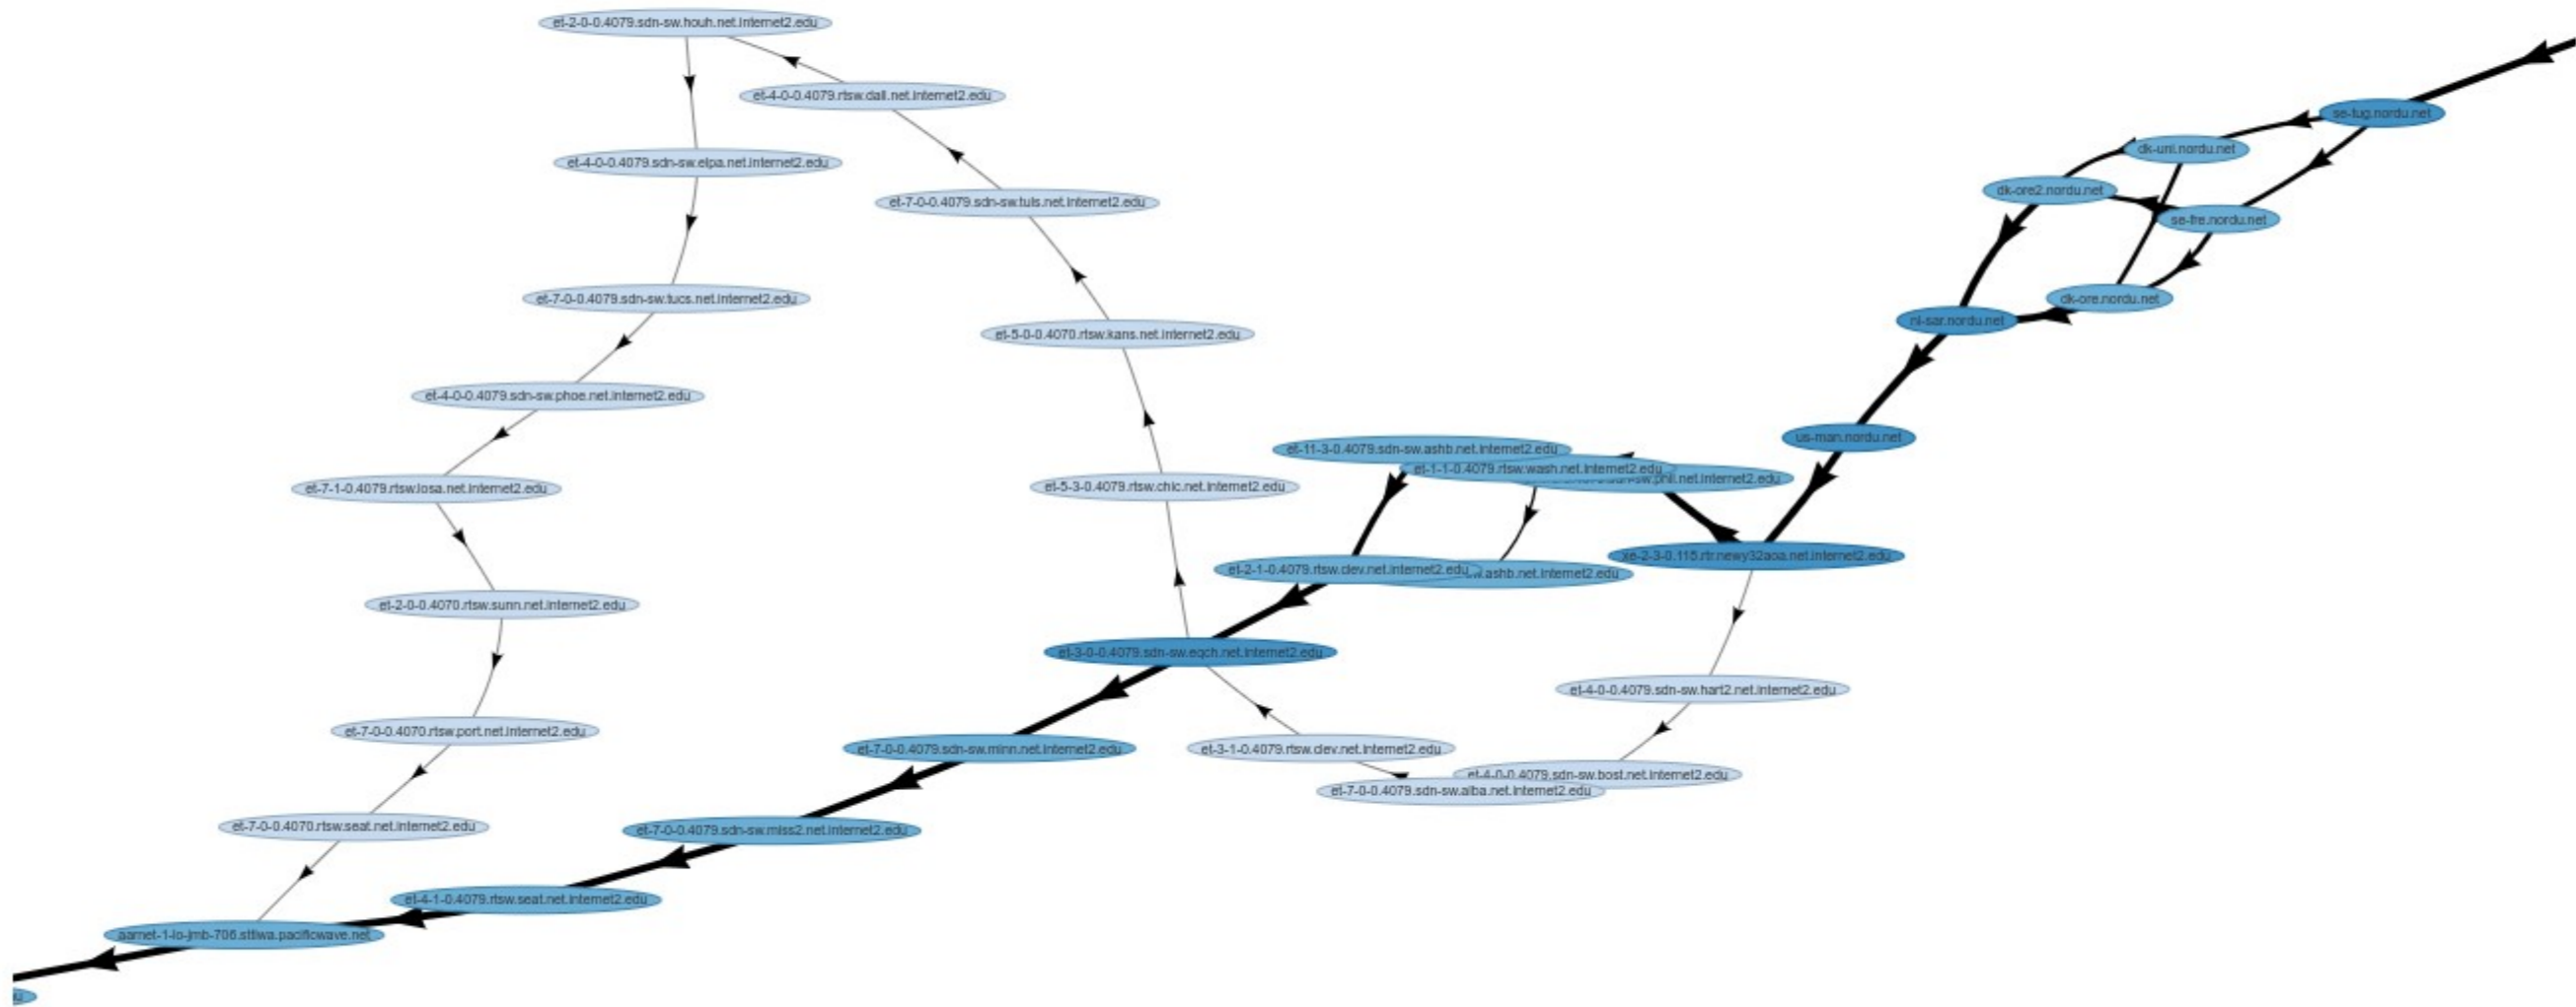


# Tracing rerouting

Traceroute charts /dragonlab/mp/trondheim-mp/20170626/trace/


Previous Next Up Raw traceroute Loss summary

Topology Hops Docs




# Queueing time


Jitter queue/time dragonlab



Jitter sum distribution dragonlab




Jitter heatmap dragonlab



Jitter count distribution dragonlab



# Queue buildup to congestion loss



Packet delay increases with 400 ms in 150 ms

# Unreliability - max 711 ppm, average 86 ppm

from.keyword: Descending ↕	to.keyword: Descending ↕	Average down_ppm ↕
saopaulo-mp	amazon-mp	711.253
saopaulo-mp	trondheim-mp	573.61
saopaulo-mp	adelaide-mp	537.161
amazon-mp	saopaulo-mp	509.679
saopaulo-mp	amazonff-mp	433.633
saopaulo-mp	googleeu-mp	407.333
madrid-mp	amazon-mp	389.398
saopaulo-mp	madrid-mp	312.583
saopaulo-mp	zurich-mp	271.626
praha-mp	stockholm-mp	177.997
amazon-mp	stockholm-mp	176.32
praha-mp	madrid-mp	173.292
madrid-mp	saopaulo-mp	170.666
praha-mp	trondheim-mp	163.699
praha-mp	adelaide-mp	162.901
copenhagen-mp	sunet-mp	147.176
zurich-mp	saopaulo-mp	134.68
amazon-mp	googleeu-mp	125.361
trondheim-mp	saopaulo-mp	117.765
adelaide-mp	googleeu-mp	115.381
		<b>86.826</b>



# Lessons learned

- Rerouting takes time..
- BGP passive mode – took 2 minutes to learn new route when primary router rebooted
- Writing full routing table contributed probably to 79 secs on a fibre-instability – 10 outages one day = 790 seconds
- PIC – Prefix Independent Convergence
  - Tree-shaped forwarding table
  - Only one alternative ..
- Deflect routing before maintaining routers
  - =>Not detected on a 10ms scale
- IGP (IS-IS) optimization
  - «loop-free alternate»
  - «fast re-route» via MPLS
- Still a lot to be learned
  - About BGP ??

# Further work

- Into production – with basic analysis
- Open issues –
  - Automated classification of outages ..
  - Combine gaps with :
    - BGP announcements – RIPE BGP logs
    - ISIS state via I-BGP
    - SNMP Traps
    - ICMP back scatter (network unreachables)
    - Traceroute hickups
    - .. ?
- Data are available 10 years of data
  - 300TB raw logs
  - Reduced data – one record -gigabytes

# Better coverage

- Own scripts
  - 17 nodes Dragonlab
- Custom debian packages
  - 50 nodes in uninett
- Perfsonar – research networks (CERN, Internet2, EU, Esnet)
- RIPE Atlas
- .. ?

# References

- Project : <https://in.uninett.no/dragonlab/>
- Talks :
  - Nordunet technical workshop – 2016
  - GEANT – SIG-NGN – Next generation Networking
  - GEANT – SIG-PMV – Performance Measurement and Verification
  - Terena Networking Conference – 2017
  - Perfsonar project
    - Rutevisualisering implementert
  - IRTF – Measurement of Protocols WG (Maprg) – 2019
  - Simula workshop on BGP security
  - Møte med Broadnet