Microdep and perfSONAR

End-to-end monitoring to impove interdomain routing

3rd European perfSONAR Users Workshop
May 24-25, 2022

Otto J Wittner and Olav Kvittem - Sikt (Uninett)

Sikt

Relevance of end-to-end monitoring

- **Continuous end-to-end measurements** have significant importance
  - May compensate for "*end-to-end blindness*" due to only (traditional) per-device monitoring
- Enable NOCs to
  - Better understand how customers experienced delivered networking services
    - **Also, interdomain QoS**
    - **Early problem-awareness**, e.g. always before customer calls service centre
  - Evaluate and improve routing and forwarding
  - Faster discover and "debug" interdomain issues
- Enable customers and users to
  - Monitor network QoS towards critical application service providers
  - Easier differentiate between external- and internal-network issues

Sikt

Examples of end-to-end observations and resulting improvments

- **Periodic 2 min outage** in NORDUnet
  - MPLS-transporter in USA required to optimize configurations
- **Routing stopped for 30 min** in Geant network
  - Cause by upgrade failure
- **2 min BGP failover time** between customer's primary and secondary connections in Uninett/Sikt
  - Optimization in BGP and IS-IS configurations required
- **Down-time due to planned maintenance** in Uninett/Sikt
  - Routines for route deflection updated
- **Fine-grained understanding of load and queues** on customer access links in Uninett/Sikt
  - Enabled timely and well documented capacity upgrade warnings to customers (no longer "gut feeling based").
- … and "die hard" packets
  - 2 week old packets traversing the Geant network, 2 hour old packets in the Uninett/Sikt network

Sikt

Microdep funamentals

- Initially a measurement project (since 2010)
- Today a measurement system and a project
- Objectives
  - Reveal network **dependability issues** at **fine grained** level by **end-to-end measurements**
  - **Improve routing** in NRENs and **the global Internet**
- *… i.e. a parallel initiative to perfSONAR*

Sikt

Topology view with **event status**

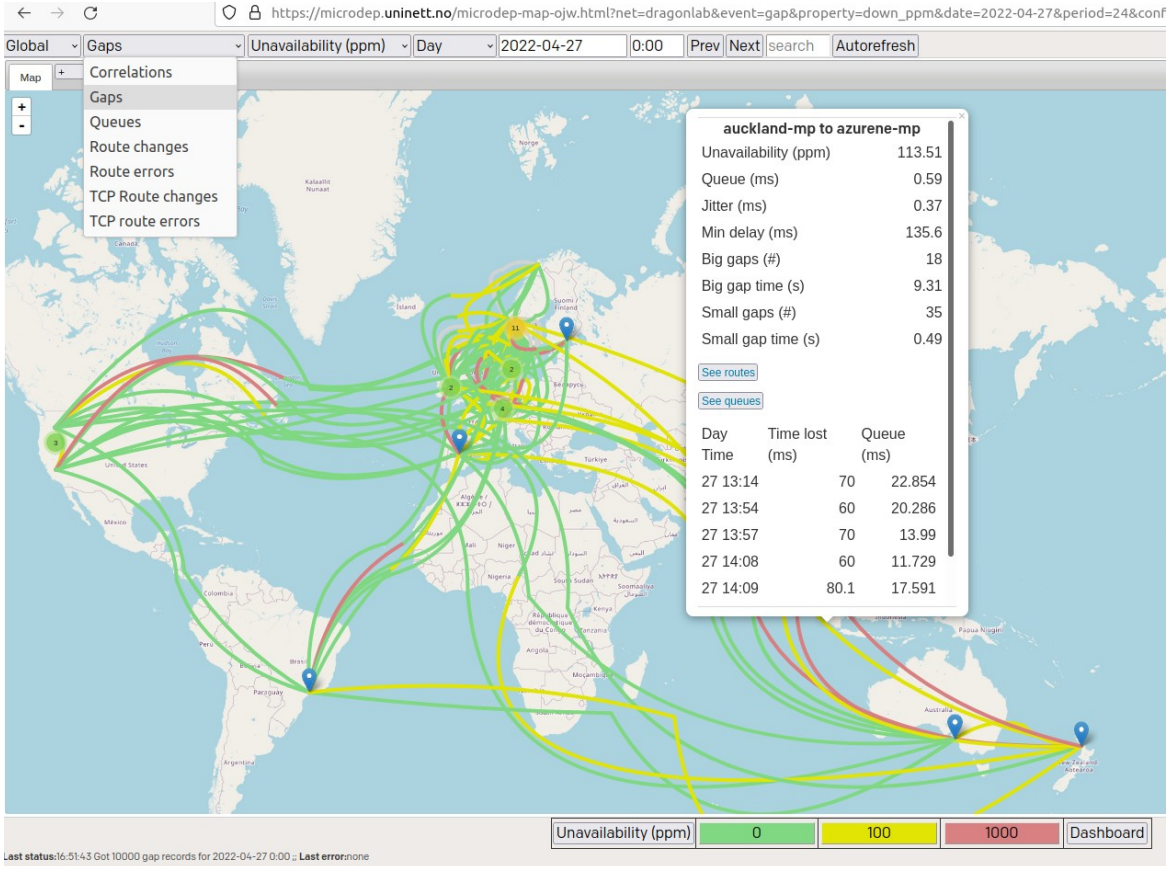

Sikt

https://microdep.uninett.no

Other views

Traceroute charts from stockholm-mp to madrid-mp(192.148.201.15) on 2022-04-05

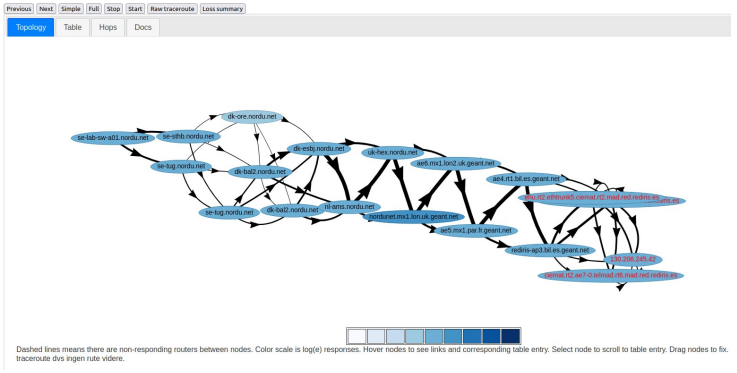

Traceroute charts from stockholm-mp to madrid-mp(192.148.201.15

Previous | Next | Simple | Full | Stop | Start | Raw traceroute | Loss summary

Topology | Table | Hops | Docs

Hop	Router	Avg ms	Min	Max	Sdv	Loss%	Seen	Address	Start	End	Error
1	se-lab-sw-a01.nordu.net	0.2	0.1	1.6	0.1	0.00%	8196	194.68.13.66	05 00:00:05	05 22:34:44	
2	se-sthb.nordu.net	1.0	0.4	29.8	1.9	4.56%	3661	109.105.97.182	05 00:00:05	05 22:34:44	
	se-tug.nordu.net	0.7	0.2	24.4	1.3	9.59%	3940	109.105.97.180	05 00:00:05	05 22:34:44	
3	dk-bal2.nordu.net	11.6	10.0	29.2	3.4	0.33%	1202	109.105.97.10	05 00:00:05	05 22:34:44	
	dk-ore.nordu.net	9.6	8.6	26.0	2.6	2.40%	609	109.105.97.130	05 00:00:05	05 22:34:44	
	se-tug.nordu.net	1.2	0.3	32.8	2.8	67.58%	2062	109.105.97.245	05 00:00:05	05 22:33:41	
4	dk-bal2.nordu.net	11.1	9.5	46.3	4.1	2.37%	2058	109.105.97.249	05 00:01:00	05 22:33:41	
	dk-esbj.nordu.net	15.9	14.6	41.2	3.1	0.25%	3982	109.105.97.3	05 00:00:05	05 22:34:44	
	dk-bal2.nordu.net	11.8	10.1	46.9	4.4	2.26%	2037	109.105.97.10	05 00:00:05	05 22:34:44	
5	nl-ams.nordu.net	22.8	20.7	87.9	5.0	0.12%	4093	109.105.97.75	05 00:00:05	05 22:34:44	
	dk-esbj.nordu.net	16.2	14.1	62.0	4.9	0.90%	4059	109.105.97.3	05 00:00:05	05 22:34:44	
6	uk-hex.nordu.net	26.2	25.4	71.3	2.8	0.02%	4040	109.105.97.125	05 00:00:05	05 22:34:44	
	nl-ams.nordu.net	23.0	20.2	77.9	6.0	0.14%	4149	109.105.97.75	05 00:00:05	05 22:34:44	
7	nordunet.mx1.lon.uk.geant.net	26.4	25.5	71.8	3.6	0.10%	4080	62.40.124.129	05 00:00:05	05 22:34:44	
	uk-hex.nordu.net	26.0	24.9	65.8	3.0	0.00%	4112	109.105.97.125	05 00:00:05	05 22:34:44	
8	ae6.mx1.lon2.uk.geant.net	27.1	26.3	76.3	2.8	0.00%	3998	62.40.98.37	05 00:00:05	05 22:34:44	
	nordunet.mx1.lon.uk.geant.net	26.4	25.0	78.1	4.0	0.00%	4197	62.40.124.129	05 00:00:05	05 22:34:44	
9	ae5.mx1.par.fr.geant.net	33.7	32.7	79.3	3.4	0.00%	4070	62.40.98.179	05 00:00:05	05 22:34:44	
	ae6.mx1.lon2.uk.geant.net	27.0	25.8	69.8	3.2	0.00%	4126	62.40.98.37	05 00:00:05	05 22:34:44	
10	ae4.rt1.bil.es.geant.net	44.7	43.8	87.4	3.3	0.00%	4111	62.40.98.222	05 00:00:05	05 22:34:44	
	ae5.mx1.par.fr.geant.net	33.5	32.2	78.5	3.6	0.00%	4085	62.40.98.179	05 00:00:05	05 22:34:44	
11	ae4.rt1.bil.es.geant.net	44.4	43.2	83.1	3.1	0.00%	4073	62.40.98.222	05 00:00:05	05 22:34:44	
	rediris-ap3.bil.es.geant.net	44.5	44.0	63.6	1.8	0.00%	4123	62.40.127.183	05 00:00:05	05 22:34:44	
12	ehu.rt2.ethtrunk5.ciemat.rt2.mad.red.rediris.es	53.5	52.9	72.9	1.7	0.00%	2053	130.206.245.5	05 00:00:05	05 22:34:44	

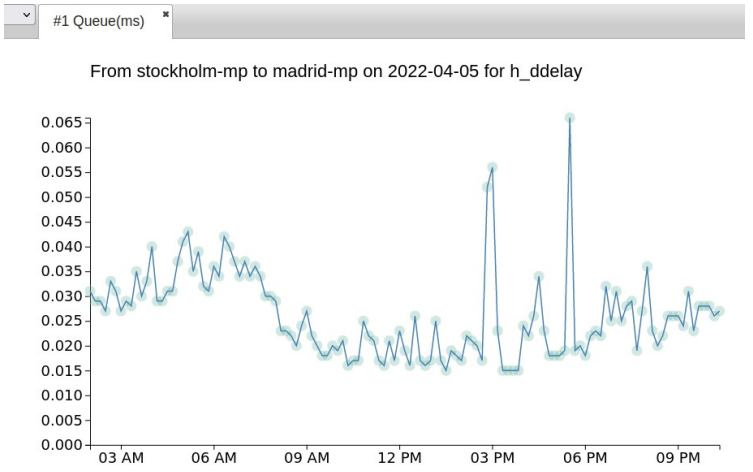

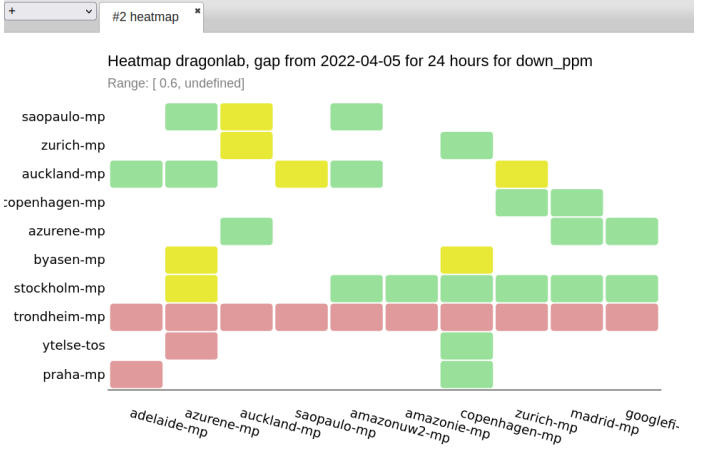

Microdep system details

- End-to-end measurements 24/7
  - 100 packets/s probe traffic
  - 60 per hour traceroutes
  - ICMP response monitoring
- 51 nodes, 212 flows in Norway
- 24 nodes, 238 flows globally
  - 8 DC-nodes (amazon, azure, google)

- Realtime event analysis
  - Packet-loss (gaps)
  - Queues (jitter)
  - Route failures and changes (traceroute)
  - Correlated events
- ML based joint event anomality under investigation.
- perfSONAR integration in progress

Sikt

Microdep and PerfSONAR today

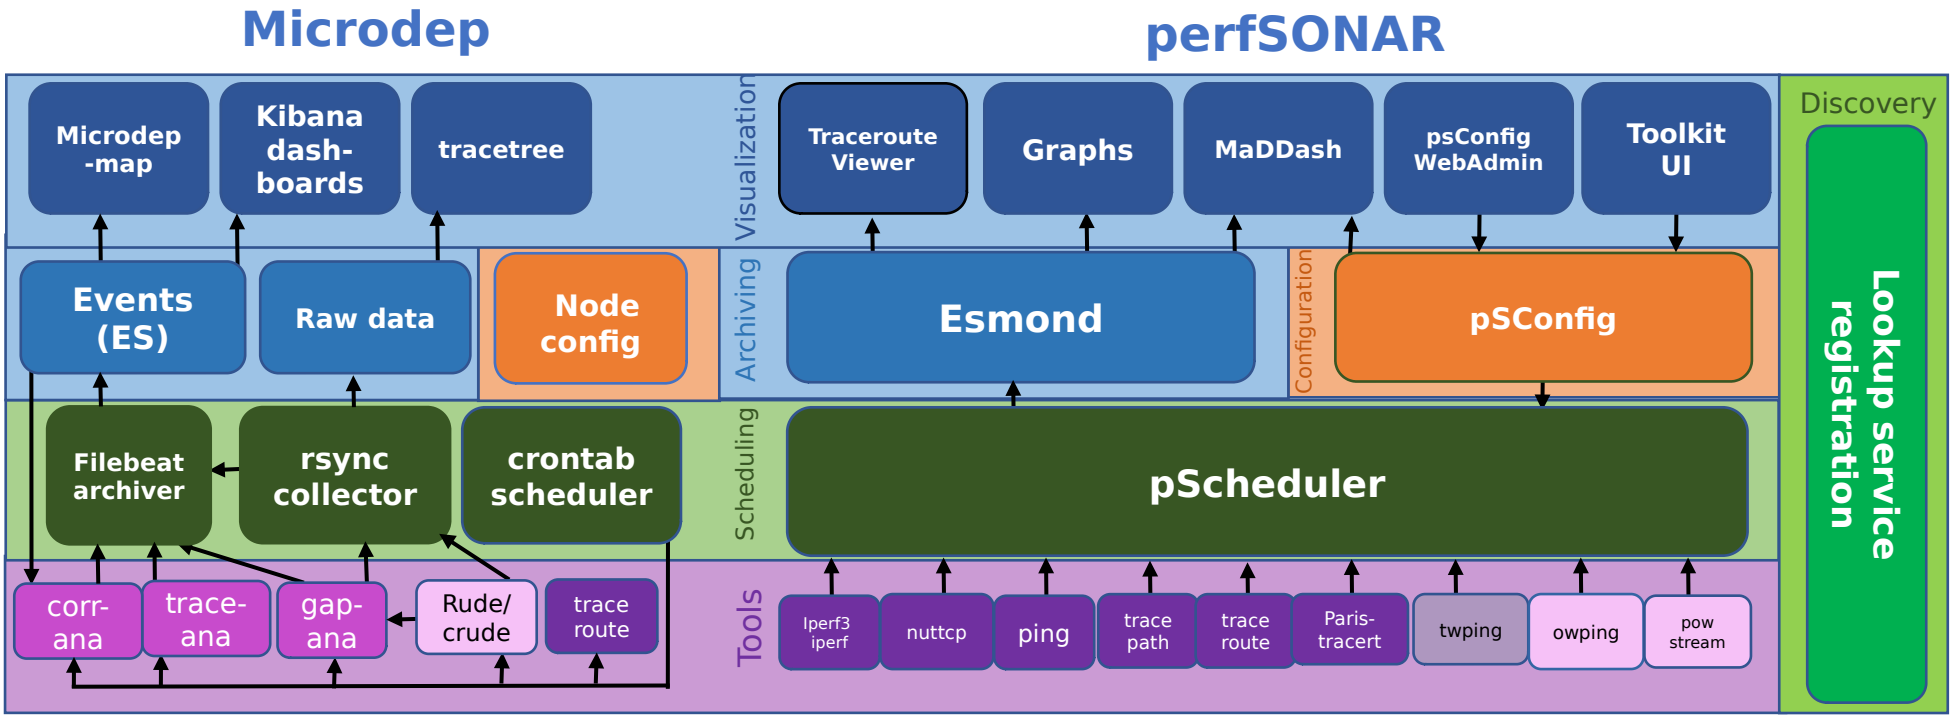

Analyser tools / plugins

Sikt

Microdep integrated in perfSONAR

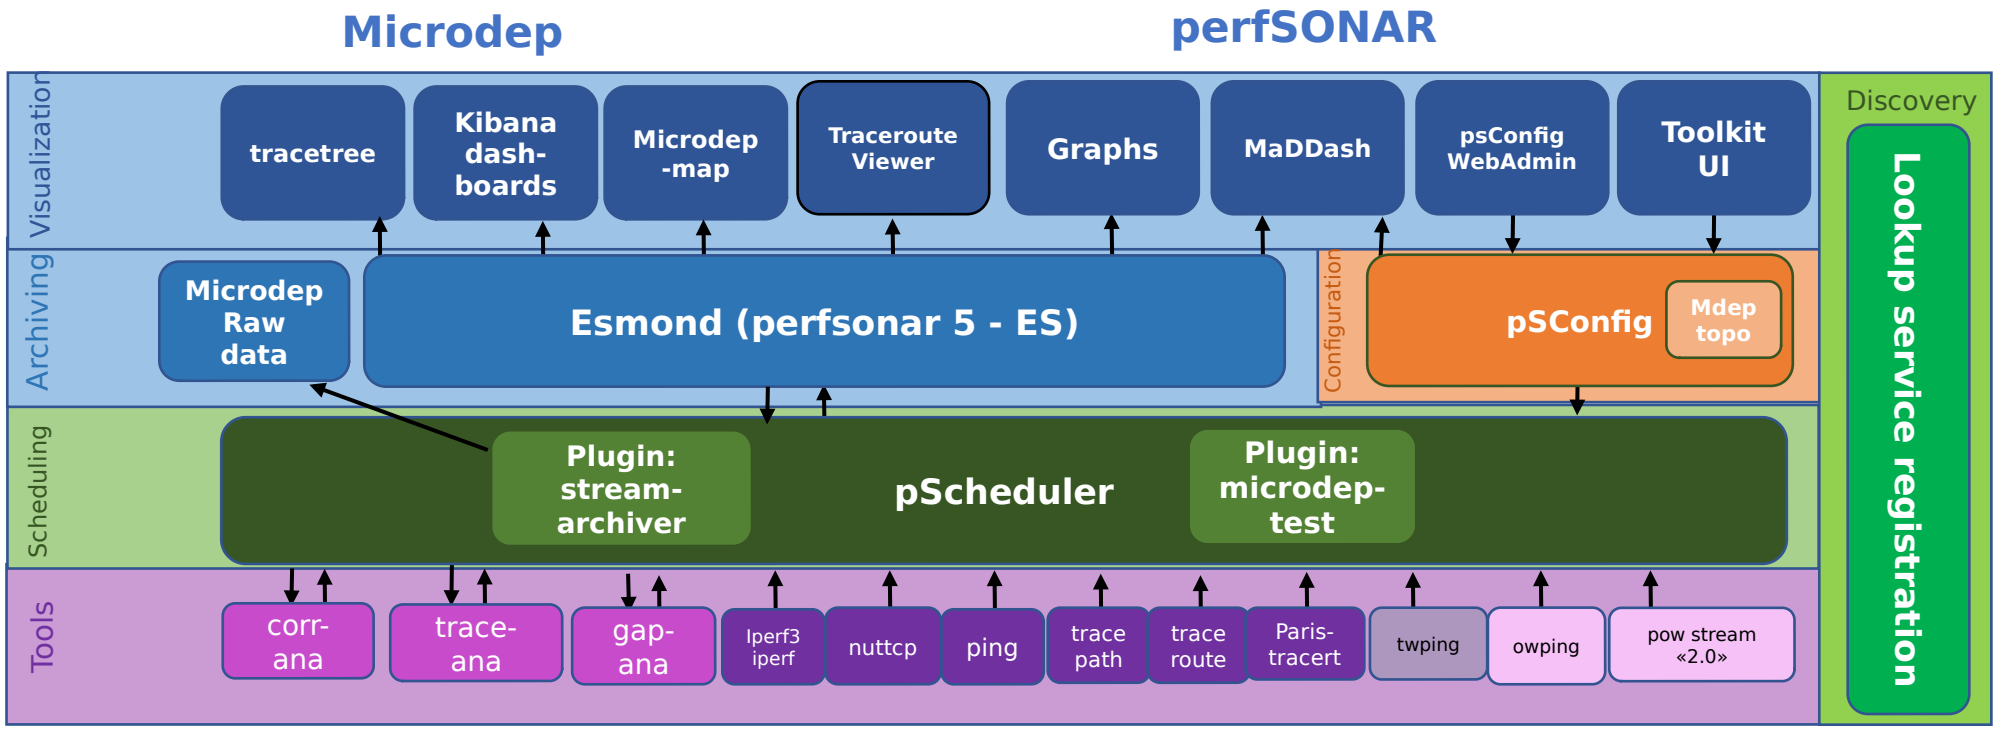

Sikt

A quick demo...

How to join in as analyst

- Access the Microdep online tool via **https://microdep.uninett.no**
- Add a node to the topology
  - Prepare a Debian or Ubuntu system (VM, container, physical)
  - Open some ports:
    - UDP 10001 and 34464-34564
    - TCP 22 and 80.
  - Run

```
wget -O- http://apt.uninett.no/uninett_apt.gpg | apt-key add -
apt-add-repository 'deb [arch=amd64] http://apt.uninett.no/debian buster main'
apt update && apt install mp-dragonlab
```

  - Email IP-address of node to *otto.wittner@sikt.no* or *olav.kvittem@sikt.no*

Sikt

Extra slides on analysis details:

Gaps / packet loss events

- Windows of 2000 pkts -> min one-way delay
- Gap event = 5 or more pkts lost, i.e. 50 ms downtime
  - 5 successfull pkts ends gap
- Stats on head and tail of gaps (50 pkts)
- Smaller gaps + other stats in daily summary reports

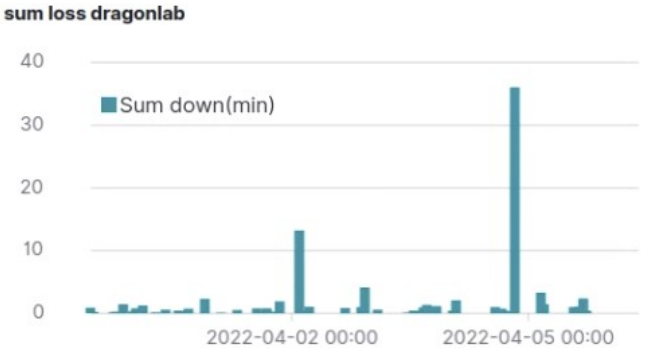

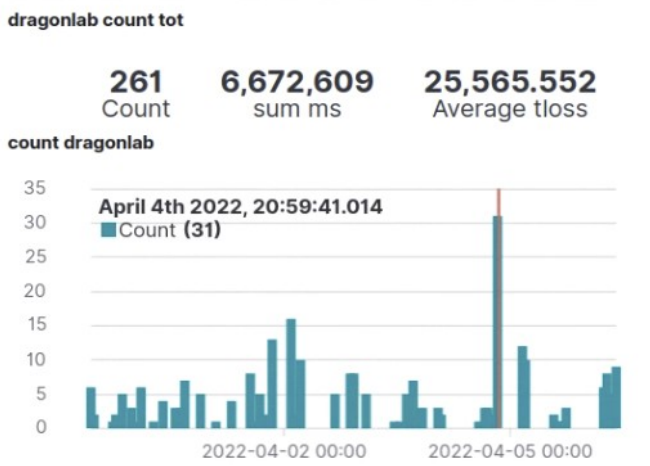

Sikt

Queues / Jitter events

- Jitter definition from RTCP (rfc3550)
  - ... but show only minor variances
  - Order of few ms
- Queue-buildup events by change in differential one-way delay
  - (delayB – delayA) – mindelay
  - Order of 10-100 ms

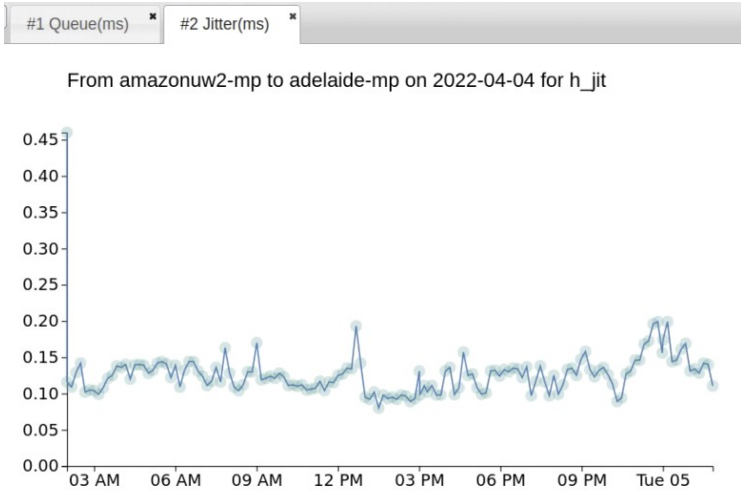

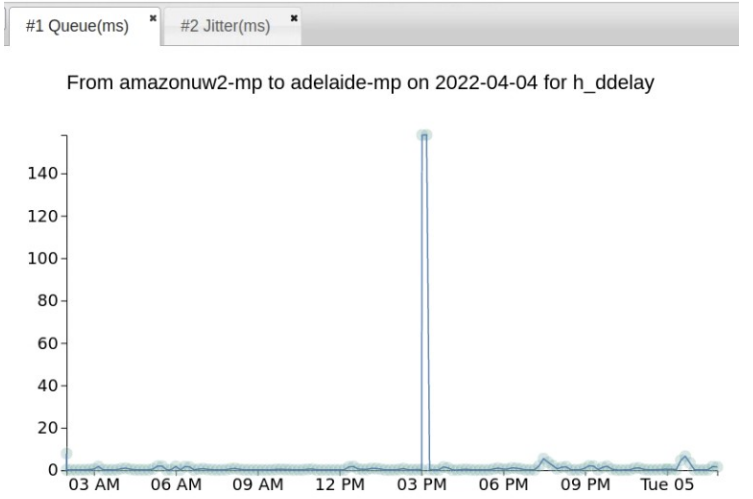

Sikt

Route failure events

- Route failure = «never ending» traceroute
- Detect periodes with route failures
  - Find «* * * * * *» at max-hops
- Report ICMP errors
  - Network unreachable (N!)
  - ...

Sikt

Route change events

- Route change = **significant** new route
- Detects **change in distribution of seen ip** addresses for each traceroute hop
  - Differential cross entropy
- «Learns» which route changes are normal

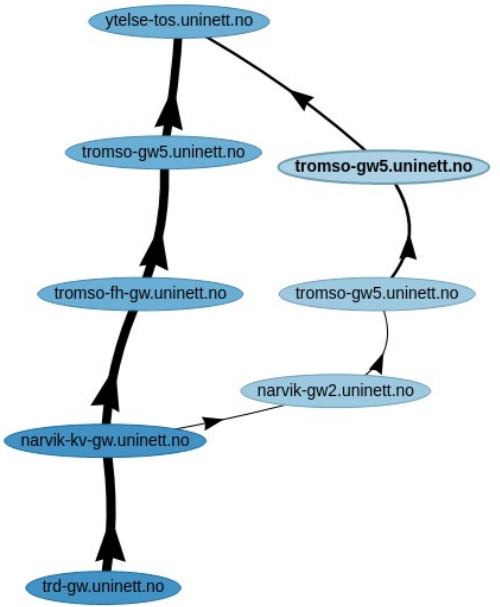

Sikt

Correlated events

- Gap and routechange in same time window
- Downtime + path anomality
- ASN to be added...

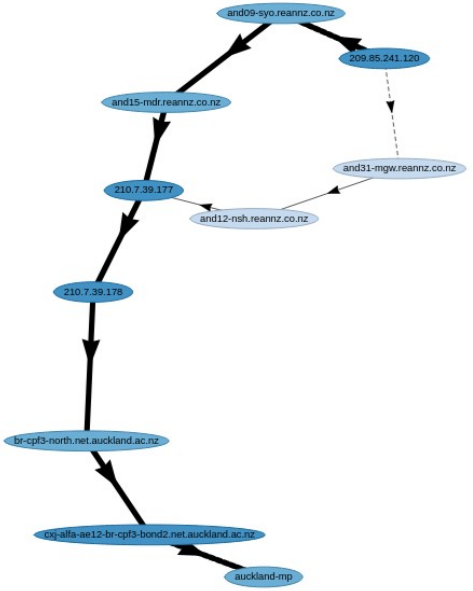

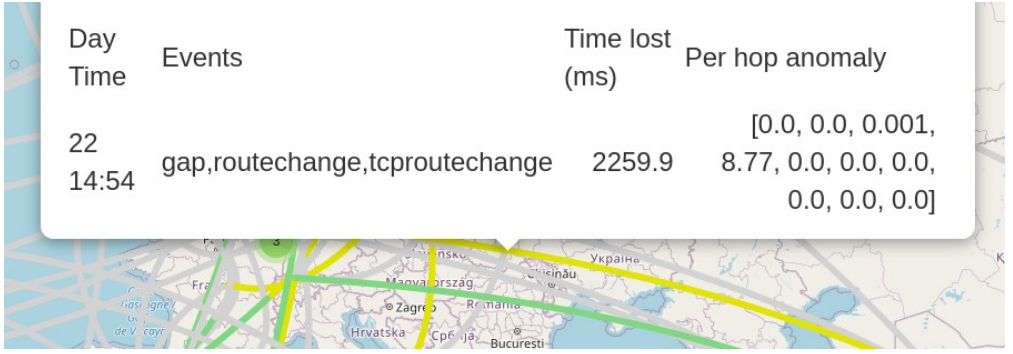

Sikt