BGP monitoring

SIG-PMV Dublin 2019

Thomas Schmid
schmid@dfn.de

problem statement

- no history of BGP routing information
  - what? when? who?
- regular BGP session only annonces best routes
  - CLI, netconf?
- ris.ripe.net and BGPplay only good for an outside perspective
- iBGPplay dead
- other tools commercial or complex or too simple
  - BGPmon
  - BGPreader
  - BGPstream

BMP

- BGP monitoring protocol
  - „Telemetry for BGP" over TCP
  - RFC 7854, IETF GROW WG
  - streams BGP information per BGP session, not per router as a normal BGP session would
  - supports multiple AFs, e.g. BGP Linkstate
  - gives ADJ-RIB-In pre-policy and post-policy
  - unidirectional TCP session: nothing is sent to the router
  - no best-path selection, i.e. no RIB
    - drafts for ADJ-RIB-Out and Loc-RIB
  - more TLVs in the future

Cisco Config

```
bmp server 1
 host 2.3.4.5 port 5000
 update-source Loopback0
 initial-delay 60
 stats-reporting-period 300
 initial-refresh delay 60 spread 120
!
router bgp 680
  neighbor 1.2.3.4
    bmp-activate server 1
```

snas.io

- originally a Cisco-development, open-source
  - formerly known as openBMP using pnda.io
- https://github.com/SNAS/
  - chain of docker containers
  - python API for kafka
  - logstash integration etc.

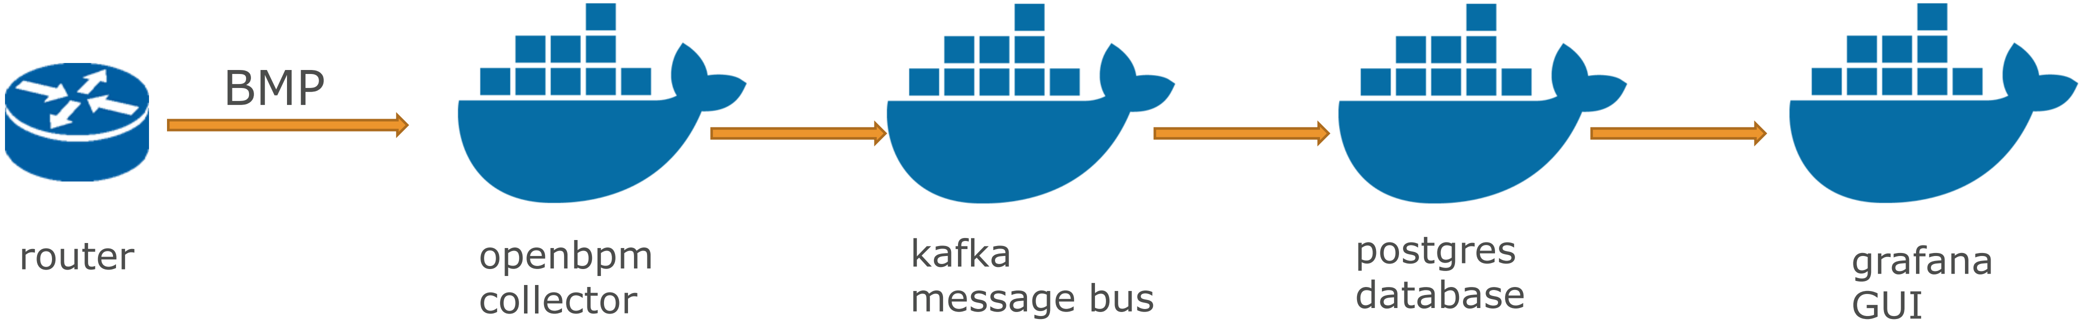

out-of-the-box GUI features

- 12 dashboards
  - routing history
    - per AS, per Peer, per Prefix, per Router
  - top-lists
  - IRR/RPKI dashboards
    - violations, per-AS view
  - inventory
  - looking-glass
  - live AS-view
  - ...

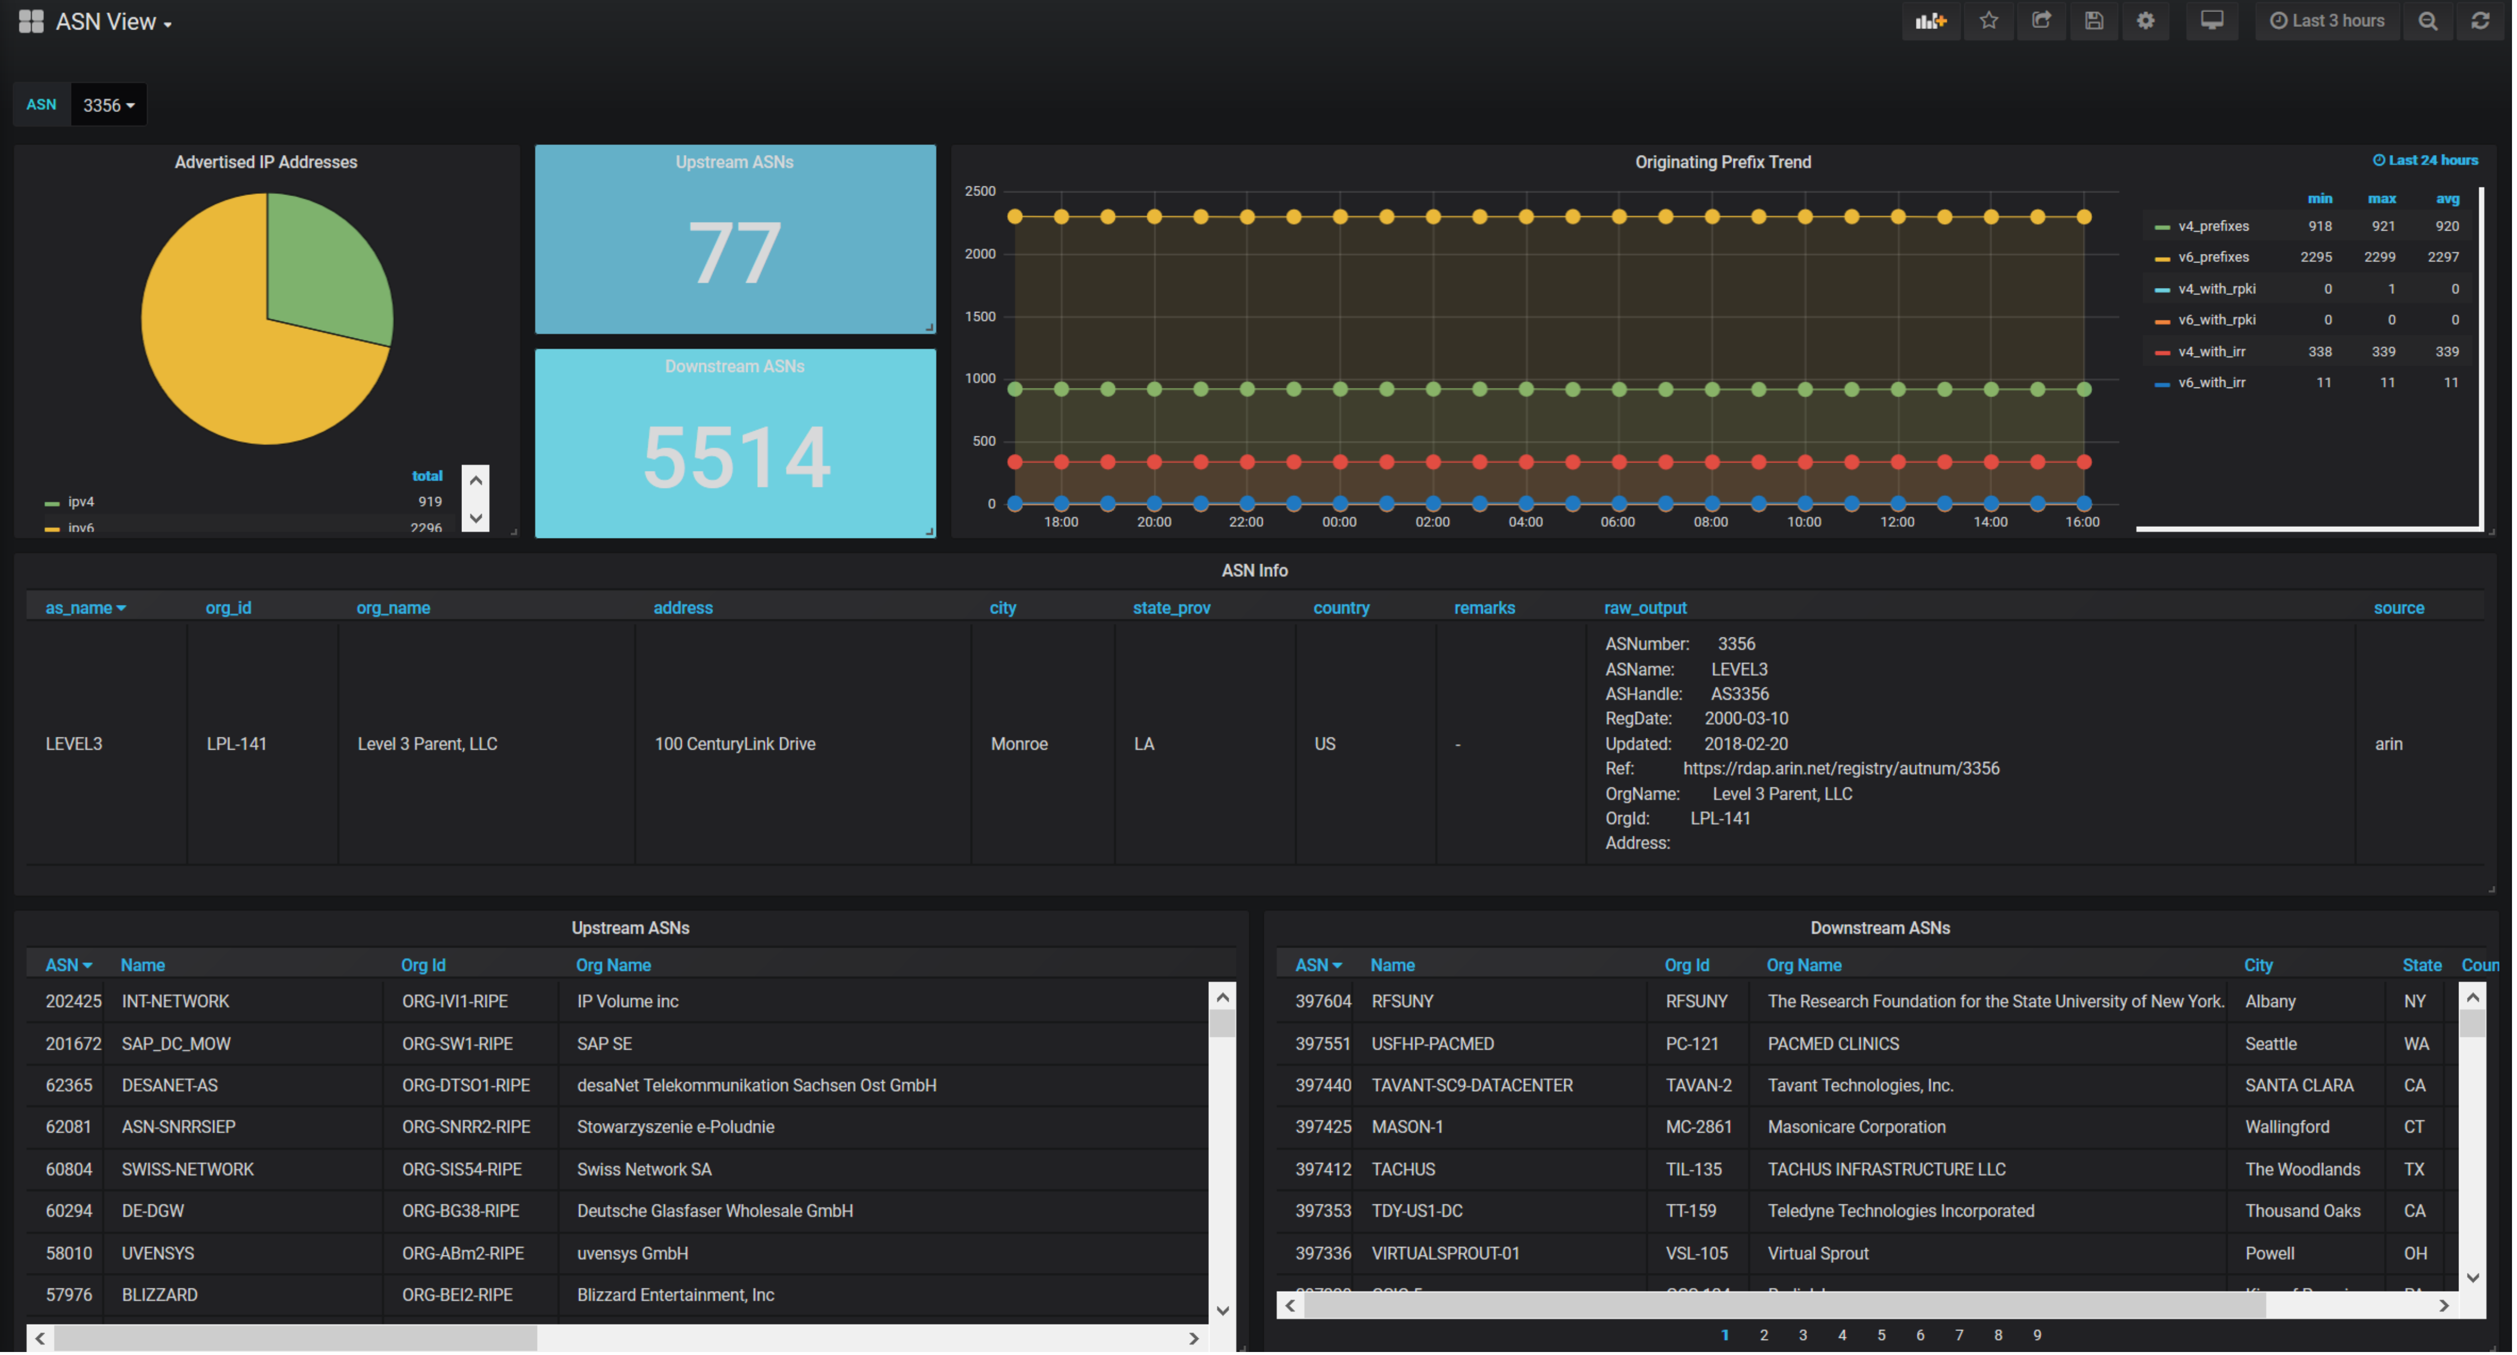

	total
ipv4	919
ipv6	2296

	min	max	avg
v4_prefixes	918	921	920
v6_prefixes	2295	2299	2297
v4_with_rpki	0	1	0
v6_with_rpki	0	0	0
v4_with_irr	338	339	339
v6_with_irr	11	11	11

ASN Info

as_name	org_id	org_name	address	city	state_prov	country	remarks	raw_output	source
LEVEL3	LPL-141	Level 3 Parent, LLC	100 CenturyLink Drive	Monroe	LA	US	-	ASNumber: 3356 ASName: LEVEL3 ASHandle: AS3356 RegDate: 2000-03-10 Updated: 2018-02-20 Ref: https://rdap.arin.net/registry/autnum/3356 OrgName: Level 3 Parent, LLC OrgId: LPL-141 Address:	arin

Upstream ASNs

ASN	Name	Org Id	Org Name
202425	INT-NETWORK	ORG-IVI1-RIPE	IP Volume inc
201672	SAP_DC_MOW	ORG-SW1-RIPE	SAP SE
62365	DESANET-AS	ORG-DTSO1-RIPE	desaNet Telekommunikation Sachsen Ost GmbH
62081	ASN-SNRRSIEP	ORG-SNRR2-RIPE	Stowarzyszenie e-Poludnie
60804	SWISS-NETWORK	ORG-SIS54-RIPE	Swiss Network SA
60294	DE-DGW	ORG-BG38-RIPE	Deutsche Glasfaser Wholesale GmbH
58010	UVENSYS	ORG-ABm2-RIPE	uvensys GmbH
57976	BLIZZARD	ORG-BEI2-RIPE	Blizzard Entertainment, Inc

Downstream ASNs

ASN	Name	Org Id	Org Name	City	State	Coun
397604	RFSUNY	RFSUNY	The Research Foundation for the State University of New York.	Albany	NY	
397551	USFHP-PACMED	PC-121	PACMED CLINICS	Seattle	WA	
397440	TAVANT-SC9-DATACENTER	TAVAN-2	Tavant Technologies, Inc.	SANTA CLARA	CA	
397425	MASON-1	MC-2861	Masonicare Corporation	Wallingford	CT	
397412	TACHUS	TIL-135	TACHUS INFRASTRUCTURE LLC	The Woodlands	TX	
397353	TDY-US1-DC	TT-159	Teledyne Technologies Incorporated	Thousand Oaks	CA	
397336	VIRTUALSPROUT-01	VSL-105	Virtual Sprout	Powell	OH	

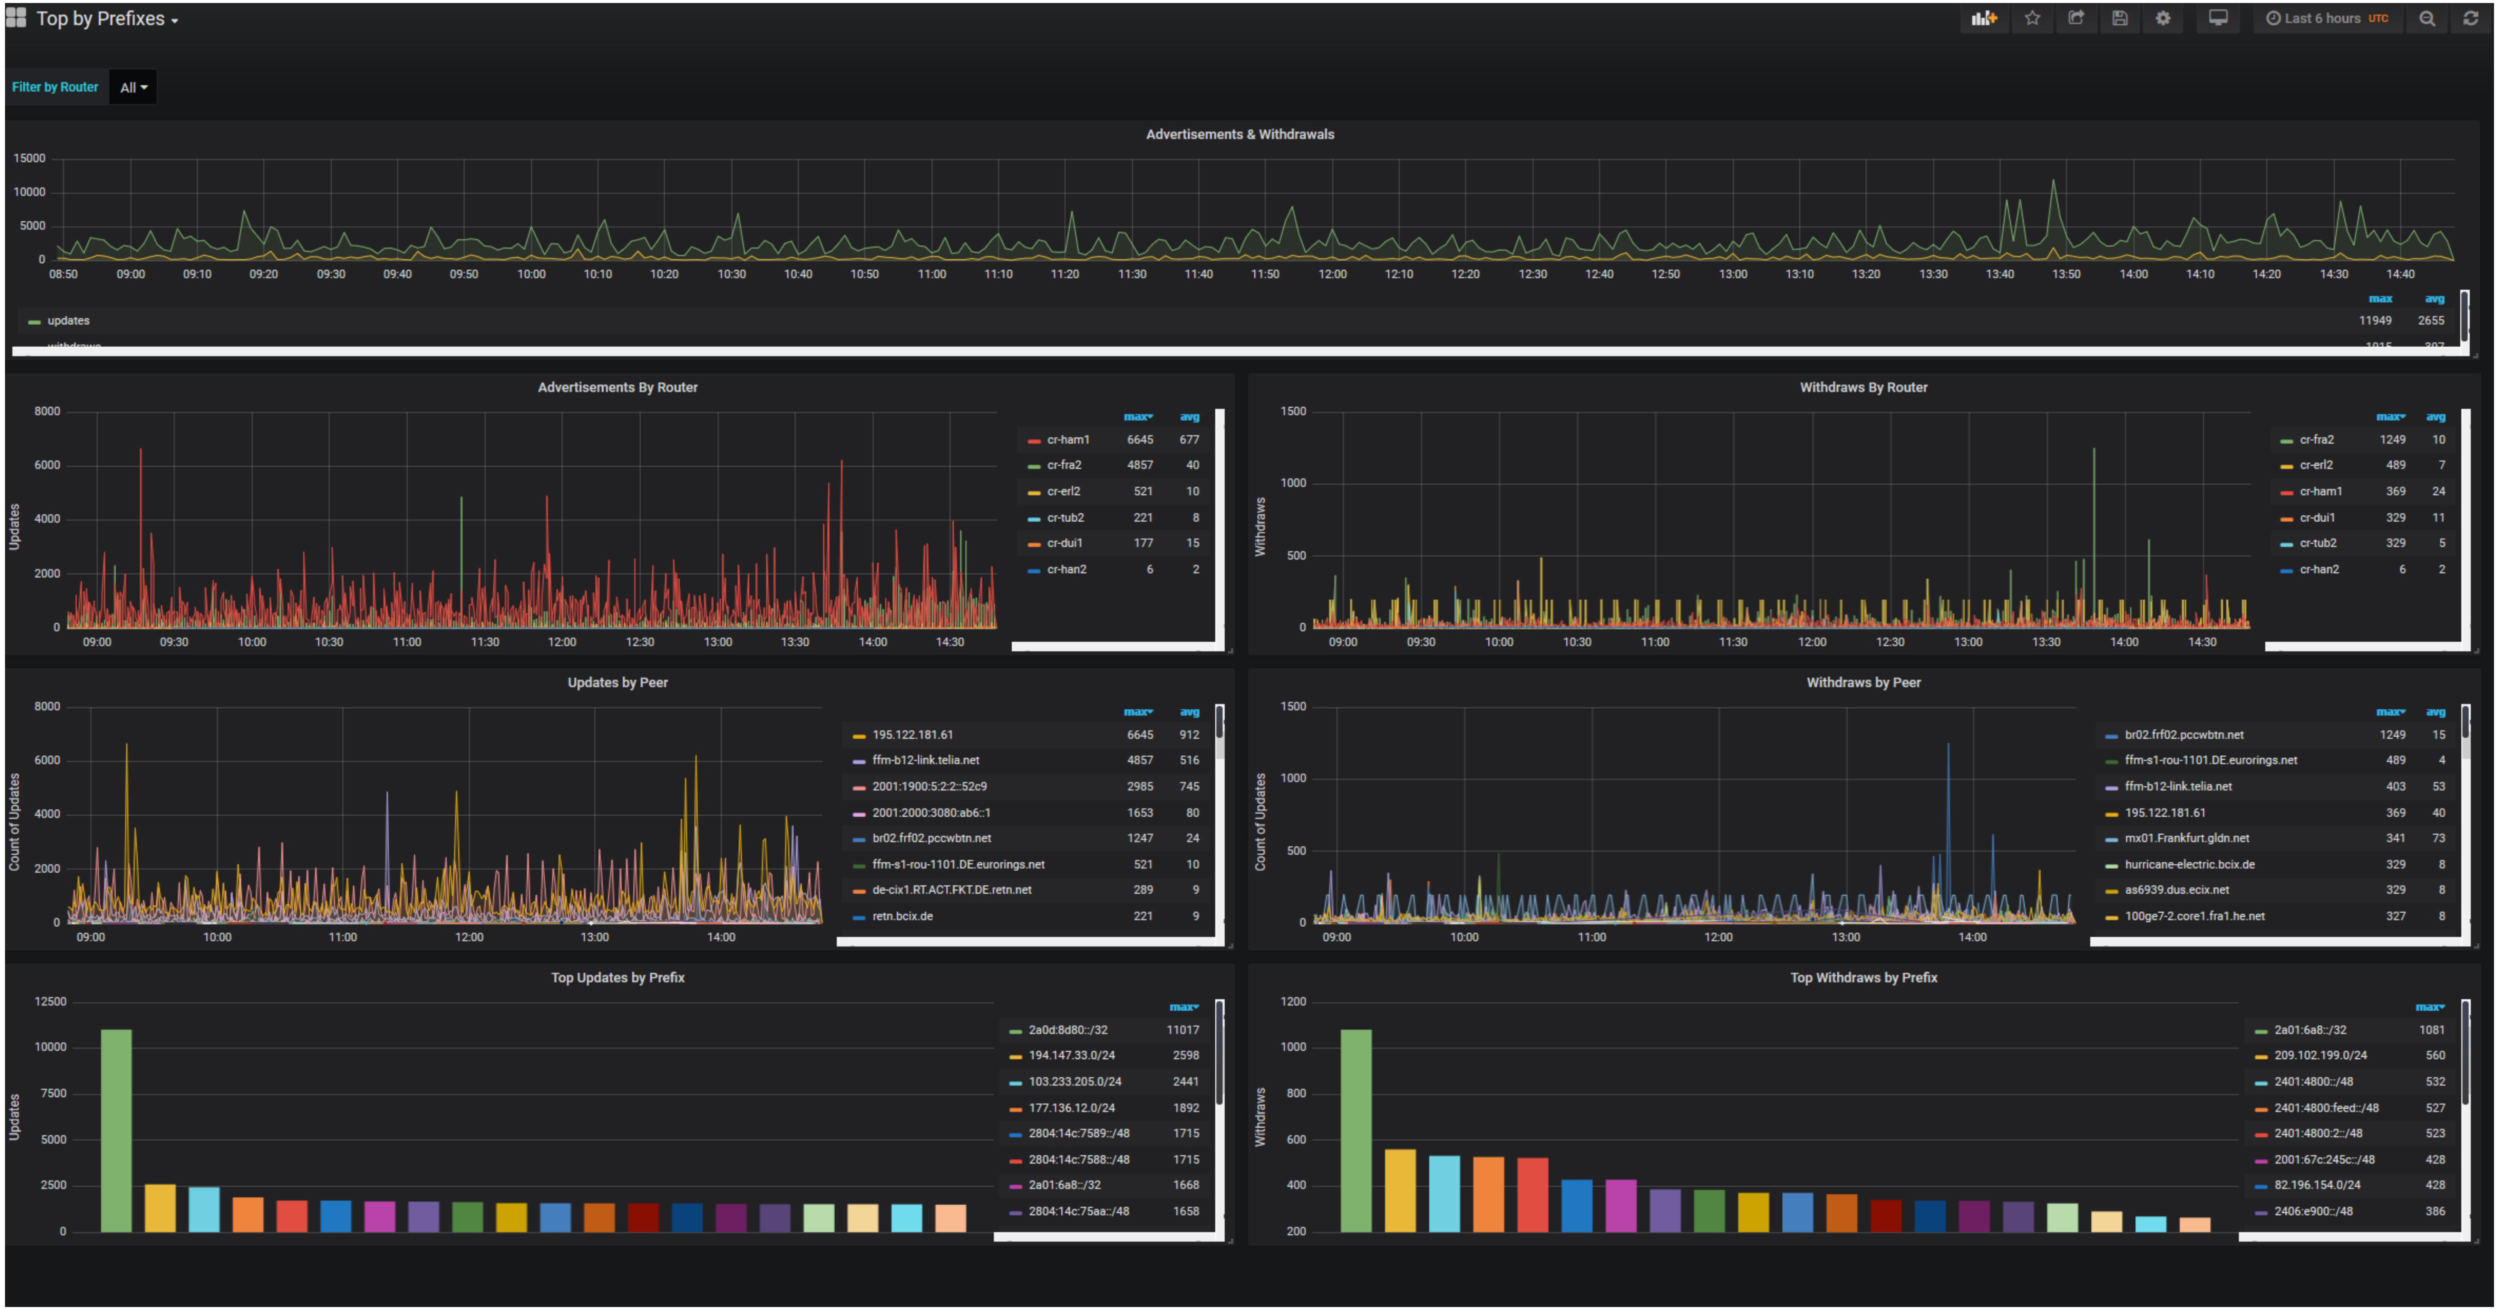

Filter by Router: All

Advertisements & Withdrawals

	max	avg
updates	11949	2655

Advertisements By Router

	max	avg
cr-ham1	6645	677
cr-fra2	4857	40
cr-erl2	521	10
cr-tub2	221	8
cr-dui1	177	15
cr-han2	6	2

Withdraws By Router

	max	avg
cr-fra2	1249	10
cr-erl2	489	7
cr-ham1	369	24
cr-dui1	329	11
cr-tub2	329	5
cr-han2	6	2

Updates by Peer

	max	avg
195.122.181.61	6645	912
ffm-b12-link.telia.net	4857	516
2001:1900:5:2:2::52c9	2985	745
2001:2000:3080:ab6::1	1653	80
br02.frf02.pccwbtn.net	1247	24
ffm-s1-rou-1101.DE.eurorings.net	521	10
de-cix1.RT.ACT.FKT.DE.retn.net	289	9
retn.bcix.de	221	9

Withdraws by Peer

	max	avg
br02.frf02.pccwbtn.net	1249	15
ffm-s1-rou-1101.DE.eurorings.net	489	4
ffm-b12-link.telia.net	403	53
195.122.181.61	369	40
mx01.Frankfurt.gldn.net	341	73
hurricane-electric.bcix.de	329	8
as6939.dus.ecix.net	329	8
100ge7-2.core1.fra1.he.net	327	8

Top Updates by Prefix

	max
2a0d:8d80::/32	11017
194.147.33.0/24	2598
103.233.205.0/24	2441
177.136.12.0/24	1892
2804:14c:7589::/48	1715
2804:14c:7588::/48	1715
2a01:6a8::/32	1668
2804:14c:75aa::/48	1658

Top Withdraws by Prefix

	max
2a01:6a8::/32	1081
209.102.199.0/24	560
2401:4800::/48	532
2401:4800:feed::/48	527
2401:4800:2::/48	523
2001:67c:245c::/48	428
82.196.154.0/24	428
2406:e900::/48	386

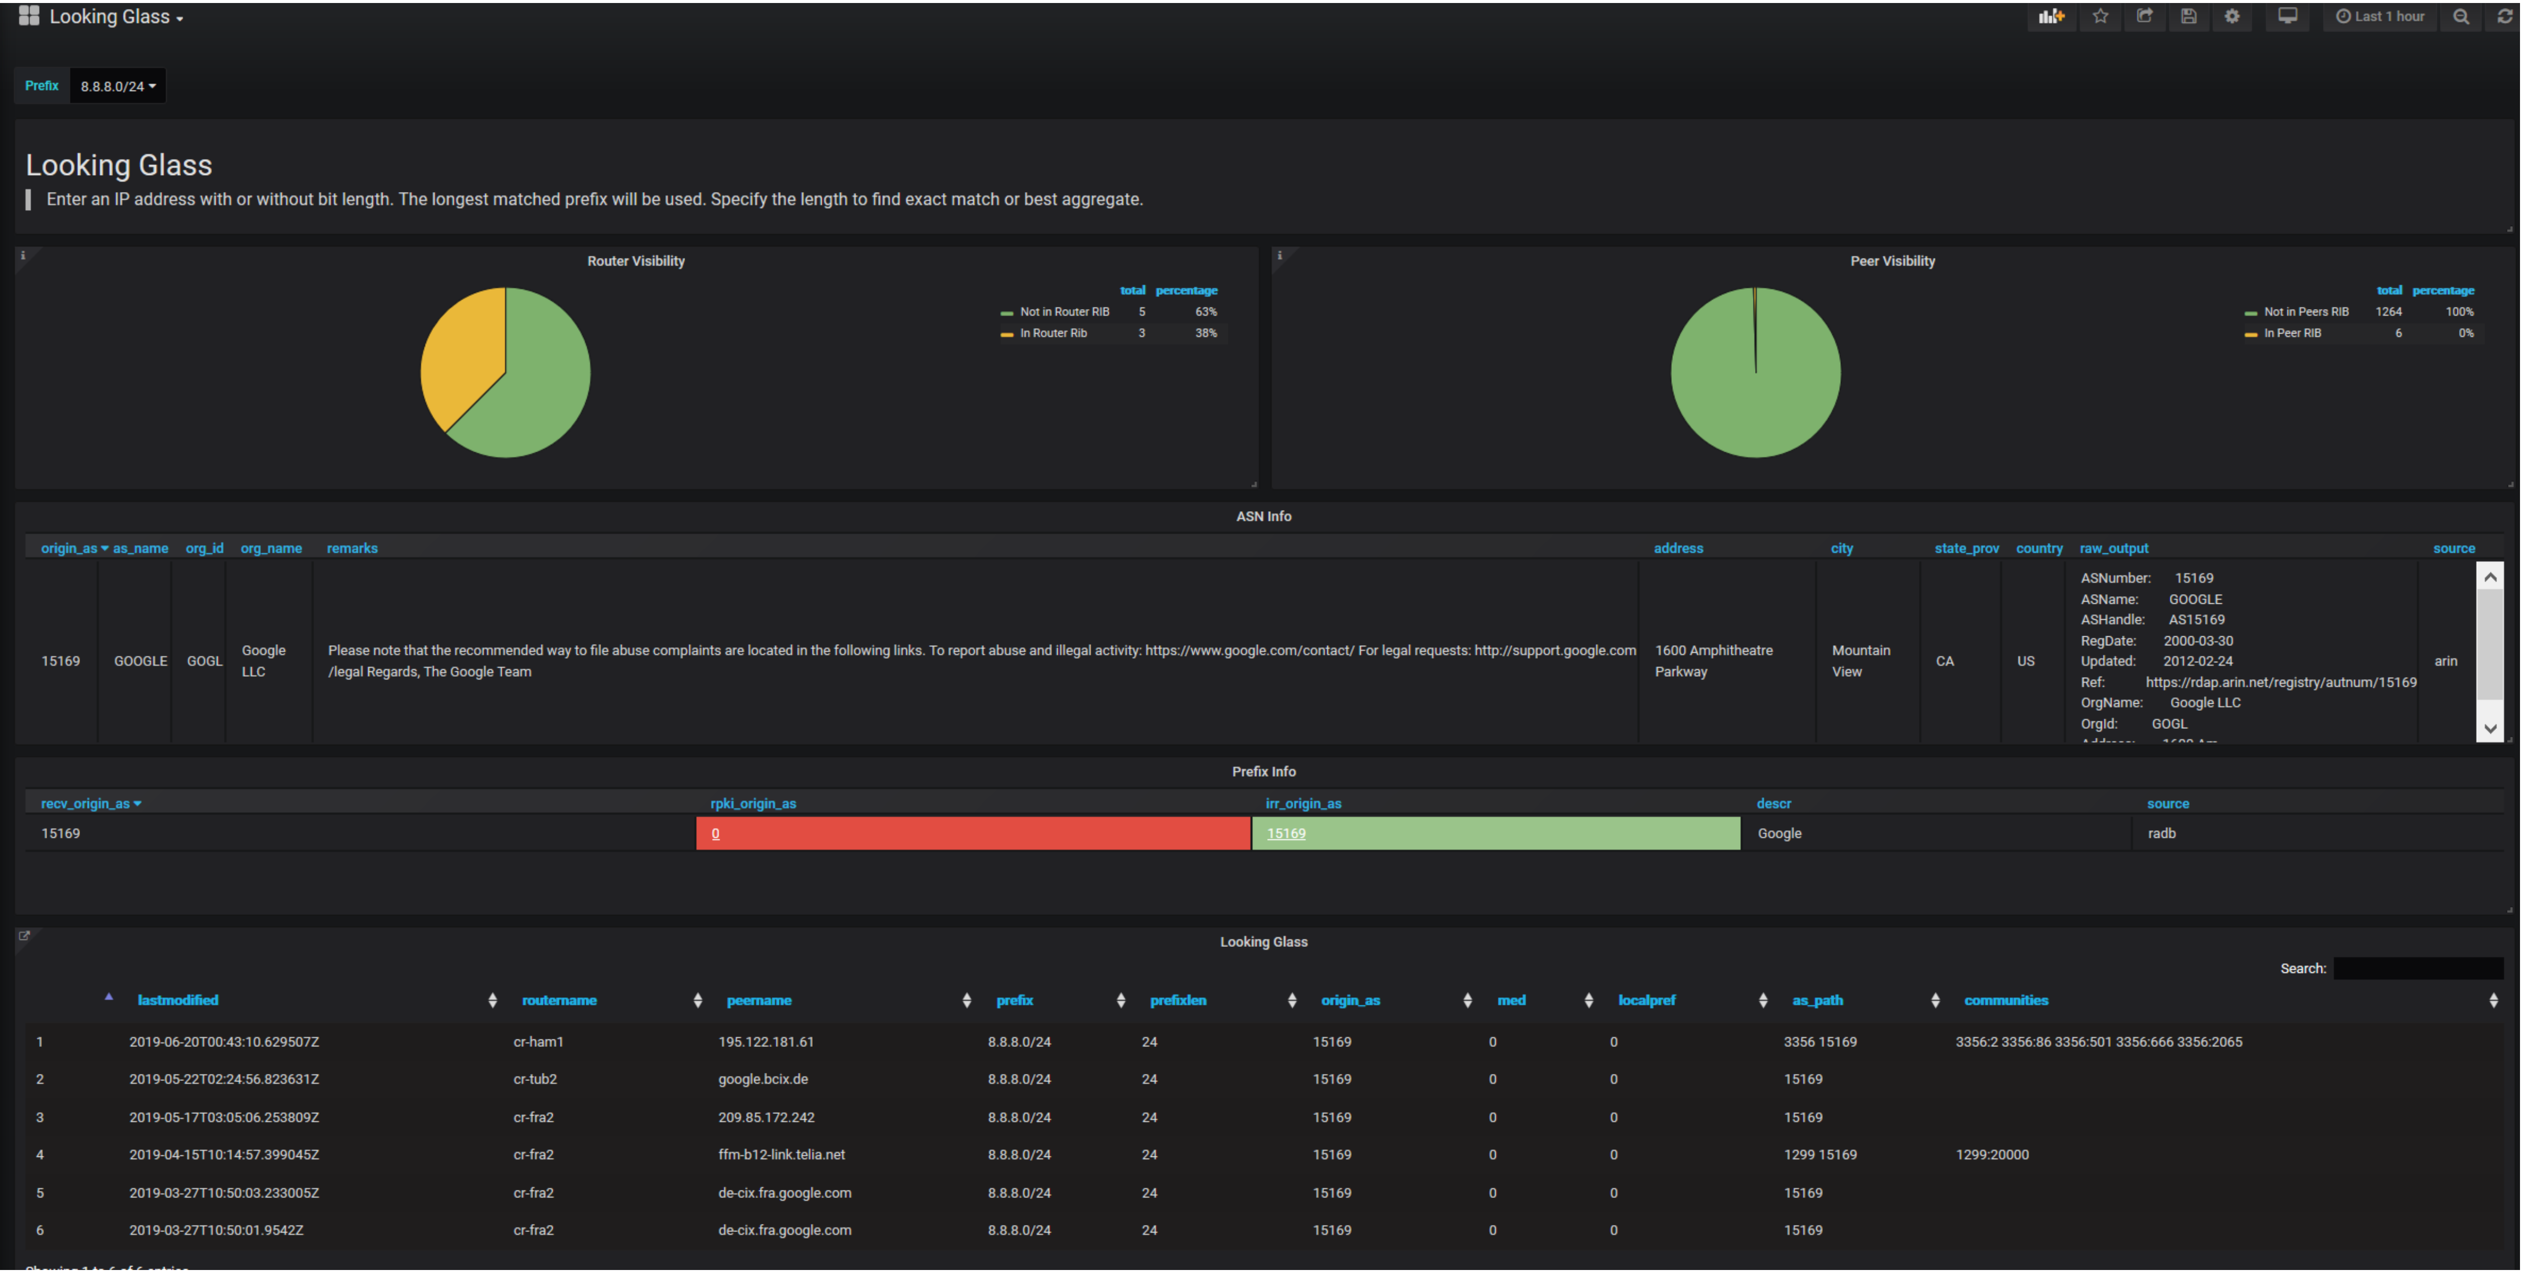

Prefix 8.8.8.0/24

Looking Glass

Enter an IP address with or without bit length. The longest matched prefix will be used. Specify the length to find exact match or best aggregate.

Router Visibility

	total	percentage
🟩 Not in Router RIB	5	63%
🟨 In Router Rib	3	38%

Peer Visibility

	total	percentage
🟩 Not in Peers RIB	1264	100%
🟨 In Peer RIB	6	0%

ASN Info

origin_as	as_name	org_id	org_name	remarks	address	city	state_prov	country	raw_output	source
15169	GOOGLE	GOGL	Google LLC	Please note that the recommended way to file abuse complaints are located in the following links. To report abuse and illegal activity: https://www.google.com/contact/ For legal requests: http://support.google.com /legal Regards, The Google Team	1600 Amphitheatre Parkway	Mountain View	CA	US	ASNumber: 15169 ASName: GOOGLE ASHandle: AS15169 RegDate: 2000-03-30 Updated: 2012-02-24 Ref: https://rdap.arin.net/registry/autnum/15169 OrgName: Google LLC OrgId: GOGL Address: 1600 Am	arin

Prefix Info

recv_origin_as	rpki_origin_as	irr_origin_as	descr	source
15169	0	15169	Google	radb

Looking Glass

Search:

	lastmodified	routername	peername	prefix	prefixlen	origin_as	med	localpref	as_path	communities
1	2019-06-20T00:43:10.629507Z	cr-ham1	195.122.181.61	8.8.8.0/24	24	15169	0	0	3356 15169	3356:2 3356:86 3356:501 3356:666 3356:2065
2	2019-05-22T02:24:56.823631Z	cr-tub2	google.bcix.de	8.8.8.0/24	24	15169	0	0	15169	
3	2019-05-17T03:05:06.253809Z	cr-fra2	209.85.172.242	8.8.8.0/24	24	15169	0	0	15169	
4	2019-04-15T10:14:57.399045Z	cr-fra2	ffm-b12-link.telia.net	8.8.8.0/24	24	15169	0	0	1299 15169	1299:20000
5	2019-03-27T10:50:03.233005Z	cr-fra2	de-cix.fra.google.com	8.8.8.0/24	24	15169	0	0	15169	
6	2019-03-27T10:50:01.9542Z	cr-fra2	de-cix.fra.google.com	8.8.8.0/24	24	15169	0	0	15169	

Showing 1 to 6 of 6 entries

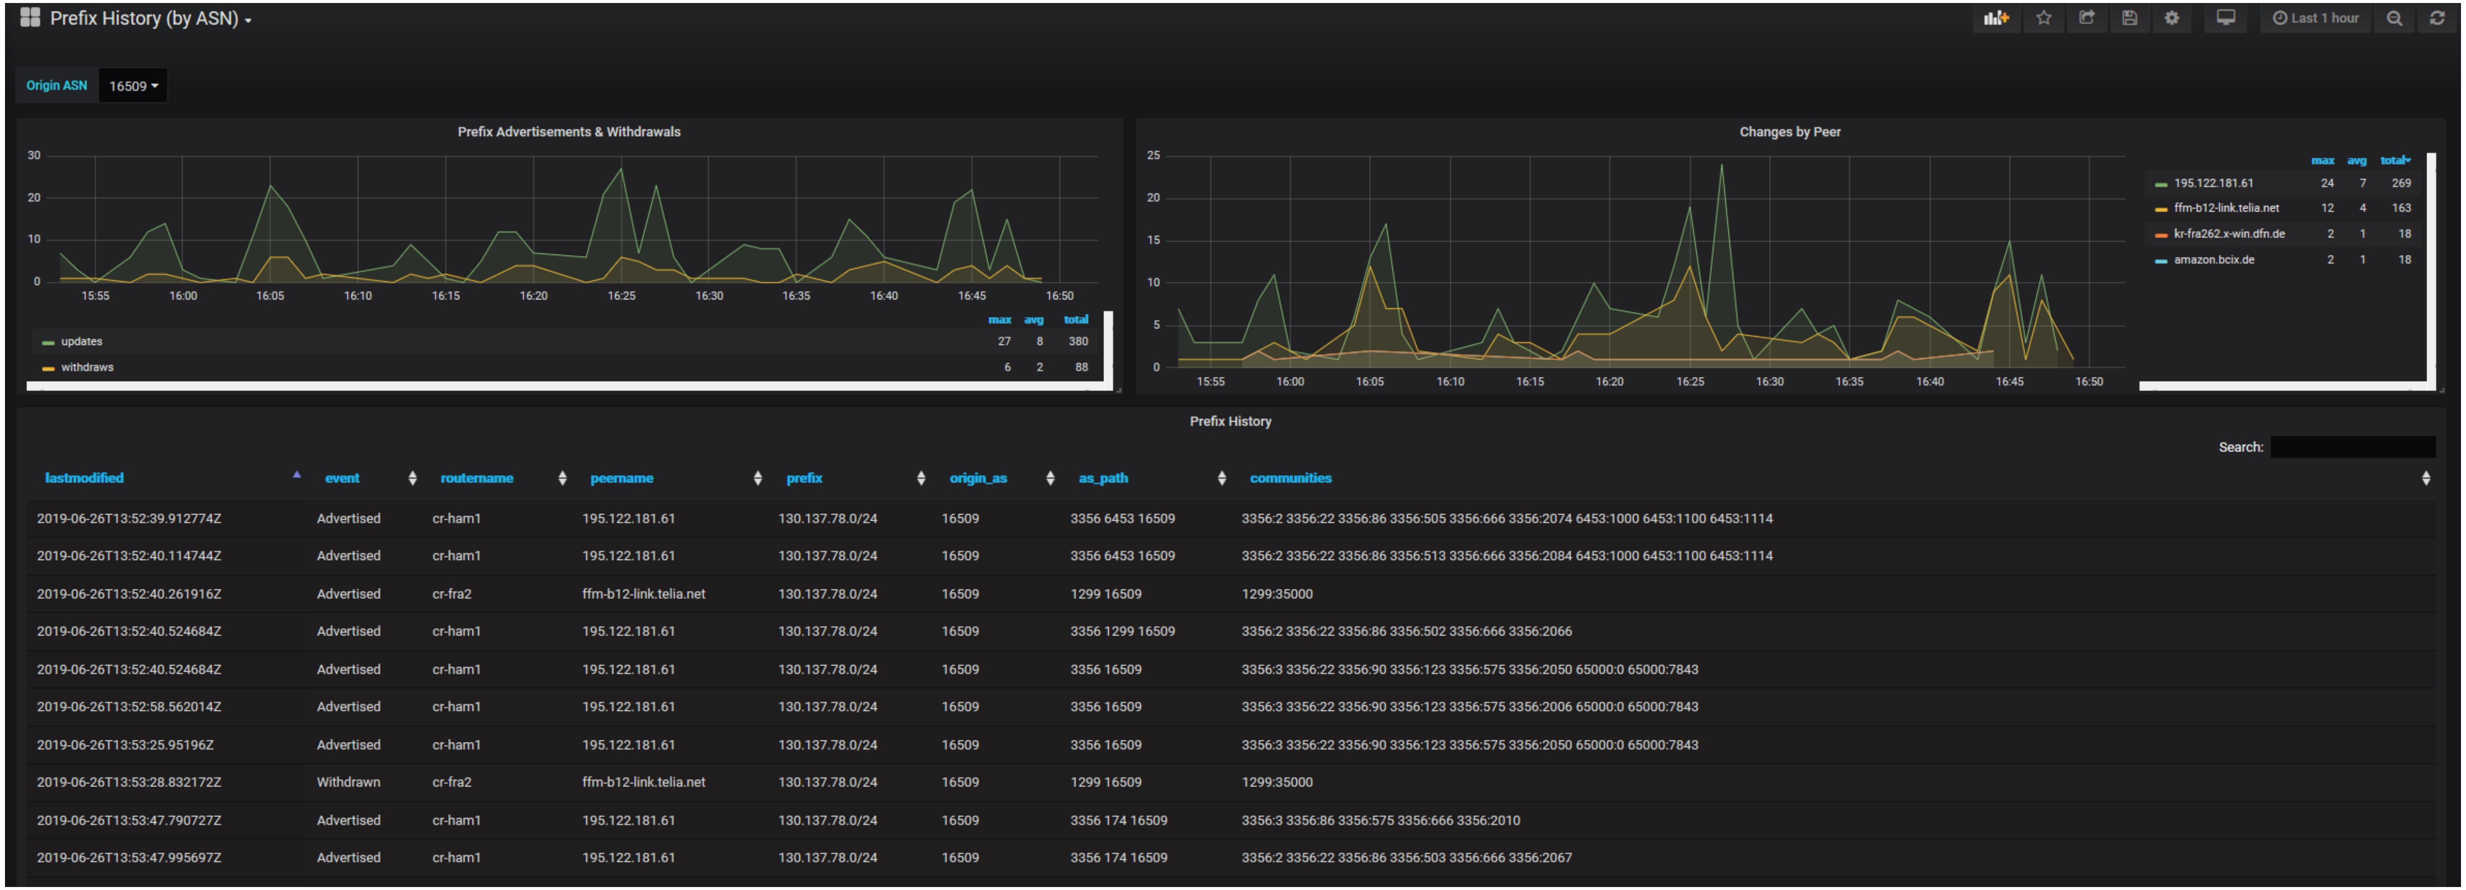

Prefix History (by ASN)

Origin ASN: 16509

Prefix Advertisements & Withdrawals

	max	avg	total
🟩 updates	27	8	380
🟨 withdraws	6	2	88

Changes by Peer

	max	avg	total
🟩 195.122.181.61	24	7	269
🟨 ffm-b12-link.telia.net	12	4	163
🟧 kr-fra262.x-win.dfn.de	2	1	18
🟦 amazon.bcix.de	2	1	18

Prefix History

Search:

lastmodified	event	routername	peername	prefix	origin_as	as_path	communities
2019-06-26T13:52:39.912774Z	Advertised	cr-ham1	195.122.181.61	130.137.78.0/24	16509	3356 6453 16509	3356:2 3356:22 3356:86 3356:505 3356:666 3356:2074 6453:1000 6453:1100 6453:1114
2019-06-26T13:52:40.114744Z	Advertised	cr-ham1	195.122.181.61	130.137.78.0/24	16509	3356 6453 16509	3356:2 3356:22 3356:86 3356:513 3356:666 3356:2084 6453:1000 6453:1100 6453:1114
2019-06-26T13:52:40.261916Z	Advertised	cr-fra2	ffm-b12-link.telia.net	130.137.78.0/24	16509	1299 16509	1299:35000
2019-06-26T13:52:40.524684Z	Advertised	cr-ham1	195.122.181.61	130.137.78.0/24	16509	3356 1299 16509	3356:2 3356:22 3356:86 3356:502 3356:666 3356:2066
2019-06-26T13:52:40.524684Z	Advertised	cr-ham1	195.122.181.61	130.137.78.0/24	16509	3356 16509	3356:3 3356:22 3356:90 3356:123 3356:575 3356:2050 65000:0 65000:7843
2019-06-26T13:52:58.562014Z	Advertised	cr-ham1	195.122.181.61	130.137.78.0/24	16509	3356 16509	3356:3 3356:22 3356:90 3356:123 3356:575 3356:2006 65000:0 65000:7843
2019-06-26T13:53:25.95196Z	Advertised	cr-ham1	195.122.181.61	130.137.78.0/24	16509	3356 16509	3356:3 3356:22 3356:90 3356:123 3356:575 3356:2050 65000:0 65000:7843
2019-06-26T13:53:28.832172Z	Withdrawn	cr-fra2	ffm-b12-link.telia.net	130.137.78.0/24	16509	1299 16509	1299:35000
2019-06-26T13:53:47.790727Z	Advertised	cr-ham1	195.122.181.61	130.137.78.0/24	16509	3356 174 16509	3356:3 3356:86 3356:575 3356:666 3356:2010
2019-06-26T13:53:47.995697Z	Advertised	cr-ham1	195.122.181.61	130.137.78.0/24	16509	3356 174 16509	3356:2 3356:22 3356:86 3356:503 3356:666 3356:2067

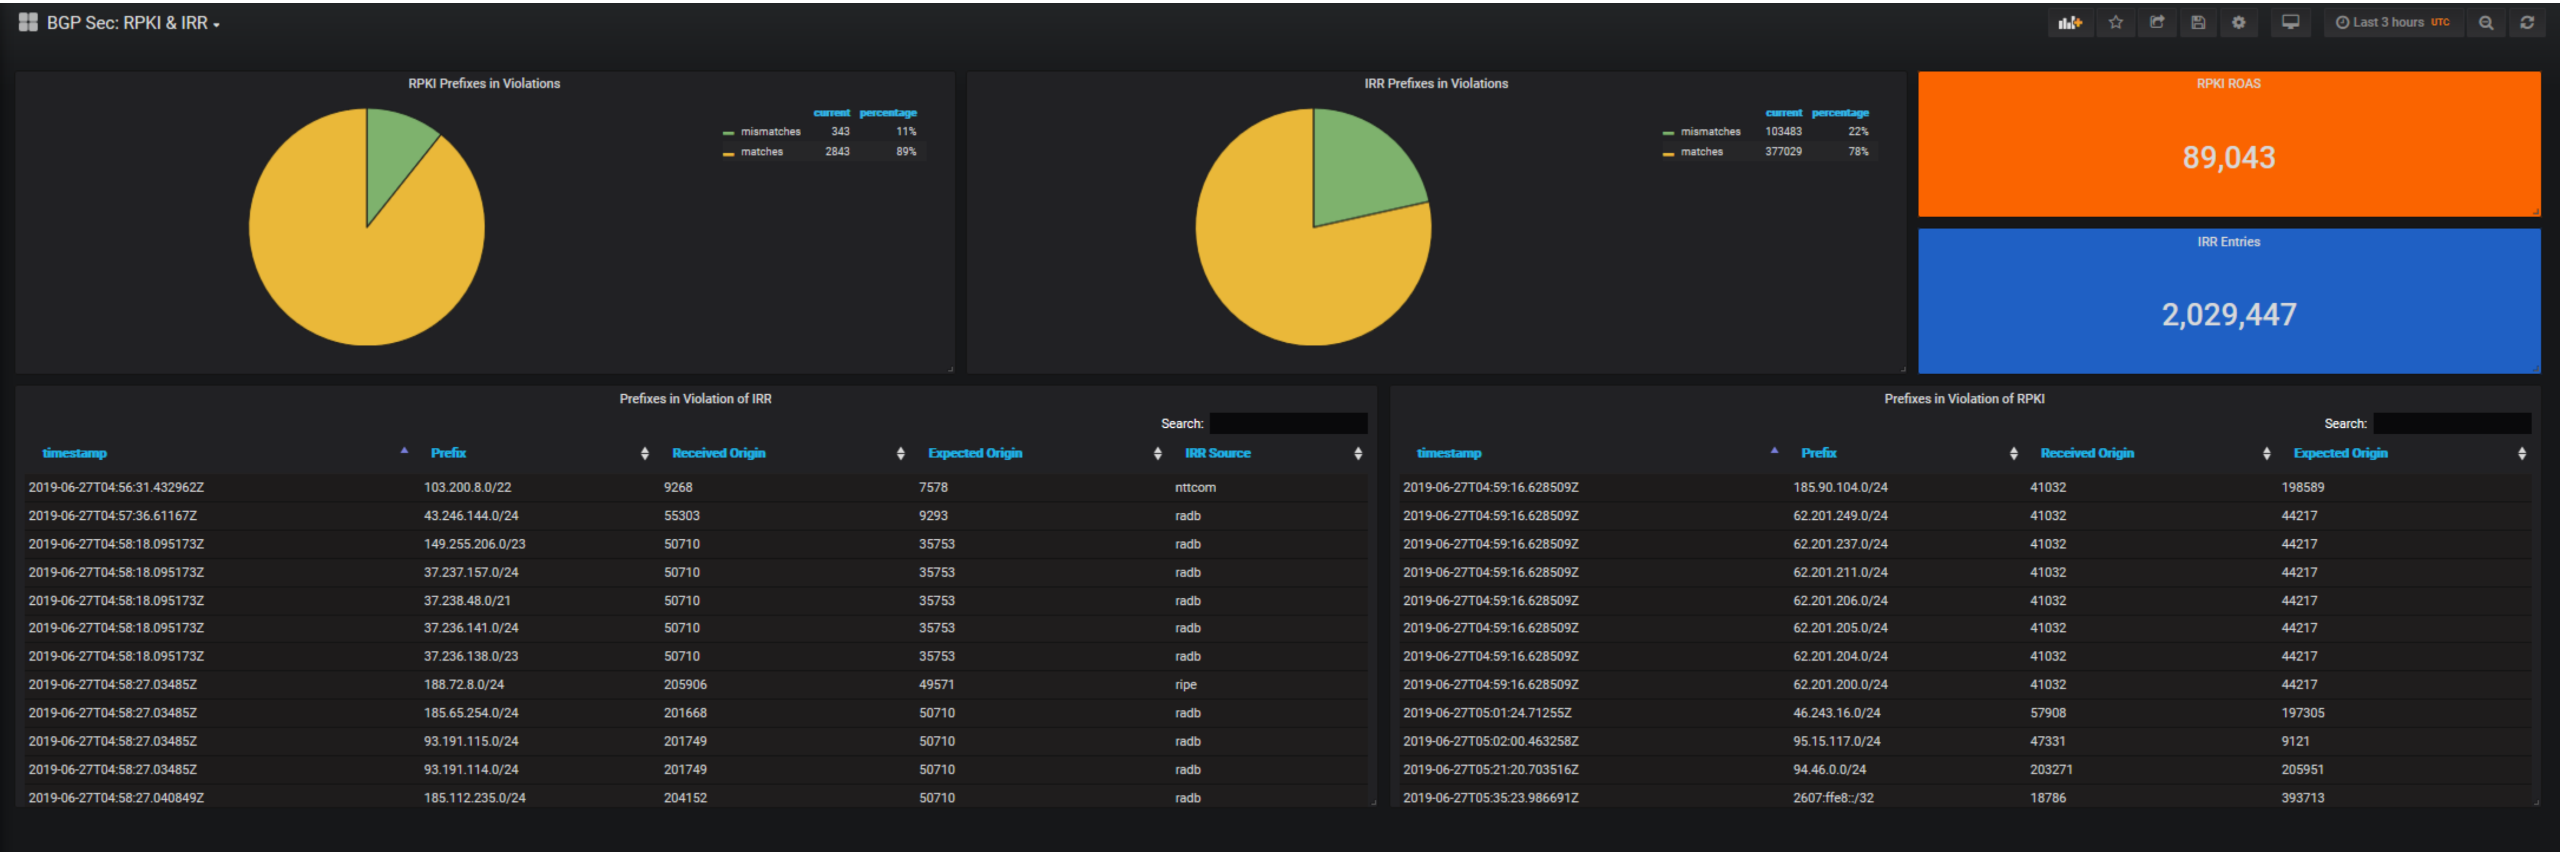

RPKI Prefixes in Violations

	current	percentage
🟩 mismatches	343	11%
🟨 matches	2843	89%

IRR Prefixes in Violations

	current	percentage
🟩 mismatches	103483	22%
🟨 matches	377029	78%

RPKI ROAS

89,043

IRR Entries

2,029,447

Prefixes in Violation of IRR

Search:

timestamp	Prefix	Received Origin	Expected Origin	IRR Source
2019-06-27T04:56:31.432962Z	103.200.8.0/22	9268	7578	nttcom
2019-06-27T04:57:36.61167Z	43.246.144.0/24	55303	9293	radb
2019-06-27T04:58:18.095173Z	149.255.206.0/23	50710	35753	radb
2019-06-27T04:58:18.095173Z	37.237.157.0/24	50710	35753	radb
2019-06-27T04:58:18.095173Z	37.238.48.0/21	50710	35753	radb
2019-06-27T04:58:18.095173Z	37.236.141.0/24	50710	35753	radb
2019-06-27T04:58:18.095173Z	37.236.138.0/23	50710	35753	radb
2019-06-27T04:58:27.03485Z	188.72.8.0/24	205906	49571	ripe
2019-06-27T04:58:27.03485Z	185.65.254.0/24	201668	50710	radb
2019-06-27T04:58:27.03485Z	93.191.115.0/24	201749	50710	radb
2019-06-27T04:58:27.03485Z	93.191.114.0/24	201749	50710	radb
2019-06-27T04:58:27.040849Z	185.112.235.0/24	204152	50710	radb

Prefixes in Violation of RPKI

Search:

timestamp	Prefix	Received Origin	Expected Origin
2019-06-27T04:59:16.628509Z	185.90.104.0/24	41032	198589
2019-06-27T04:59:16.628509Z	62.201.249.0/24	41032	44217
2019-06-27T04:59:16.628509Z	62.201.237.0/24	41032	44217
2019-06-27T04:59:16.628509Z	62.201.211.0/24	41032	44217
2019-06-27T04:59:16.628509Z	62.201.206.0/24	41032	44217
2019-06-27T04:59:16.628509Z	62.201.205.0/24	41032	44217
2019-06-27T04:59:16.628509Z	62.201.204.0/24	41032	44217
2019-06-27T04:59:16.628509Z	62.201.200.0/24	41032	44217
2019-06-27T05:01:24.71255Z	46.243.16.0/24	57908	197305
2019-06-27T05:02:00.463258Z	95.15.117.0/24	47331	9121
2019-06-27T05:21:20.703516Z	94.46.0.0/24	203271	205951
2019-06-27T05:35:23.986691Z	2607:ffe8::/32	18786	393713

psql database

- + timescale DB
- accessible with standard management tools e.g. pgAdmin
- well structured and useful to build your own dashboards
    - customer views, iBGP, etc.

pros and cons

- + scales well
  - currently monitoring almost 2000 BGP sessions
    - Dell PowerEdge M640, 2 x Intel Xeon Silver 4110, 64GB, 500GB SSD
  - lightweight on the routers
    - no visible additional load when turning on BMP
- + easy to adapt to your specific needs
  - database has all the information you want
- - no easy „full routing table view per router“
  - focus on changes **per neighbor**
  - BGP-table, not RIB
  - better use quagga and dump BGP table with caida tools BGPstream BGPreader for this purpose
  - stream mrt data to openbmp collector: mrt2bmp
- - development stopped?
  - no updates since 6months, gitter chat quiet
    - but already a mature and good product

new monitoring DMON

- early stage. The dream:
  - full integrated monitoring and subsitution of existing monitoring tools: cacti, mrtg, log, ...
  - full view over all network layers
  - full root cause analysis of failures
  - full alarming features
  - great GUI
- source of truth: GIS
  - central database, OSS/BSS for DFN
  - all dependencies and processes already modeled

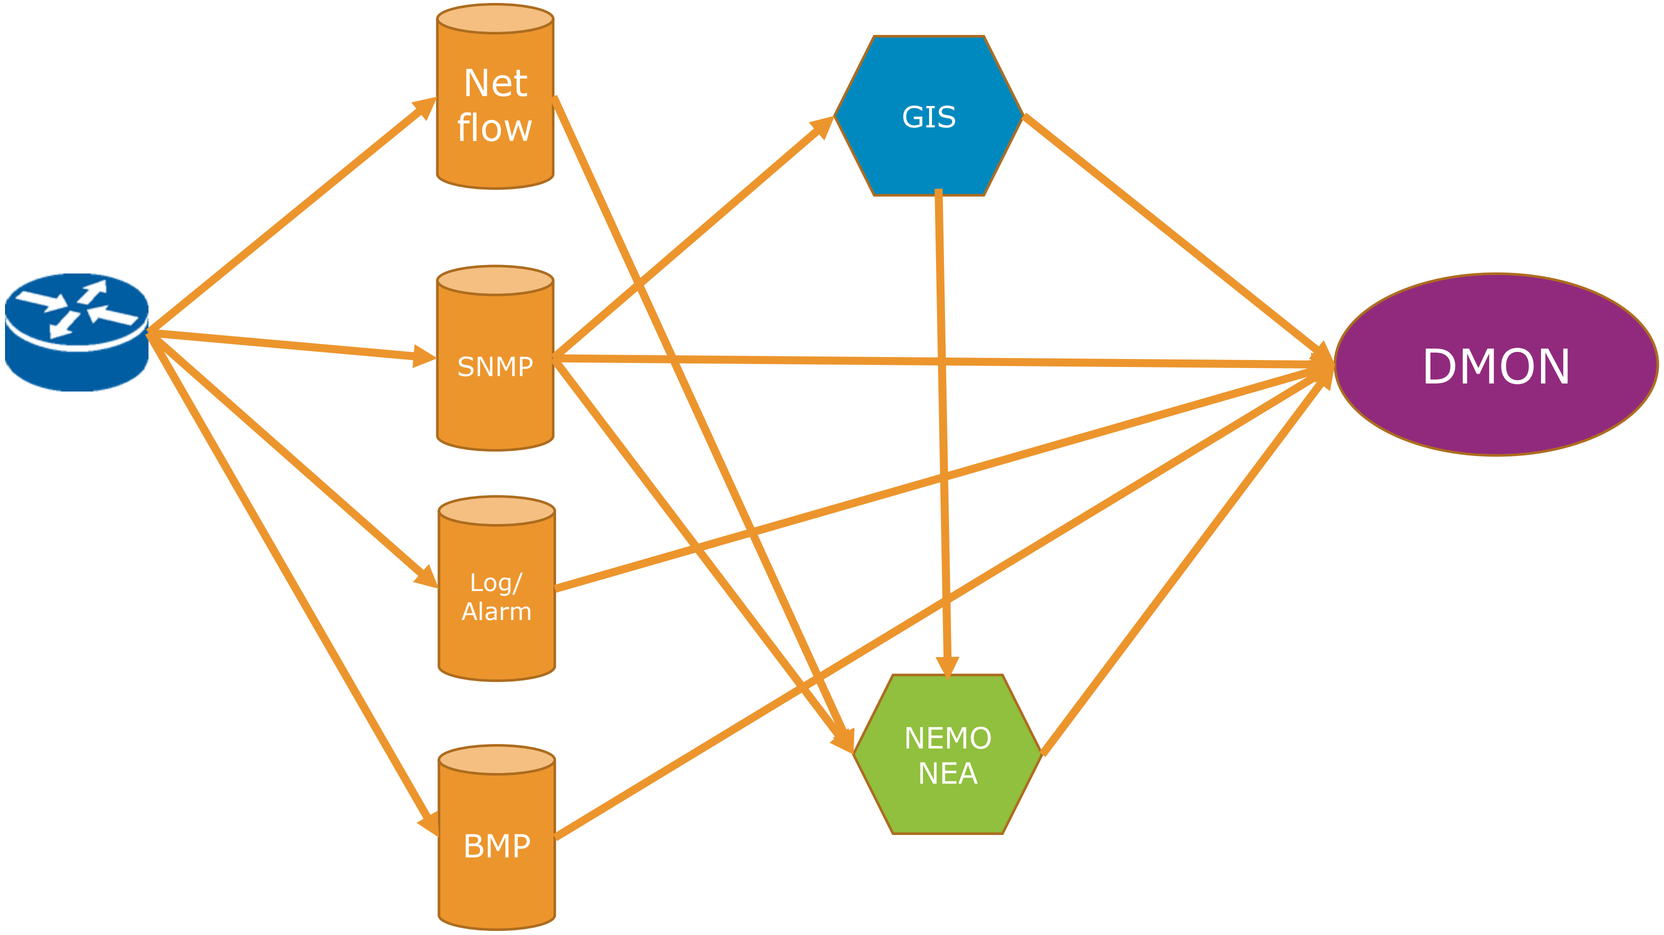
Net flow
SNMP
Log/ Alarm
BMP
GIS
NEMO NEA
DMON

challenges

- GIS architecture outdated
  - 20 years
  - Adabas DB
  - single threaded
  - middleware CORBA
  - missing real-time APIs
    - mainly XML exports
- Telemetry later

Questions?