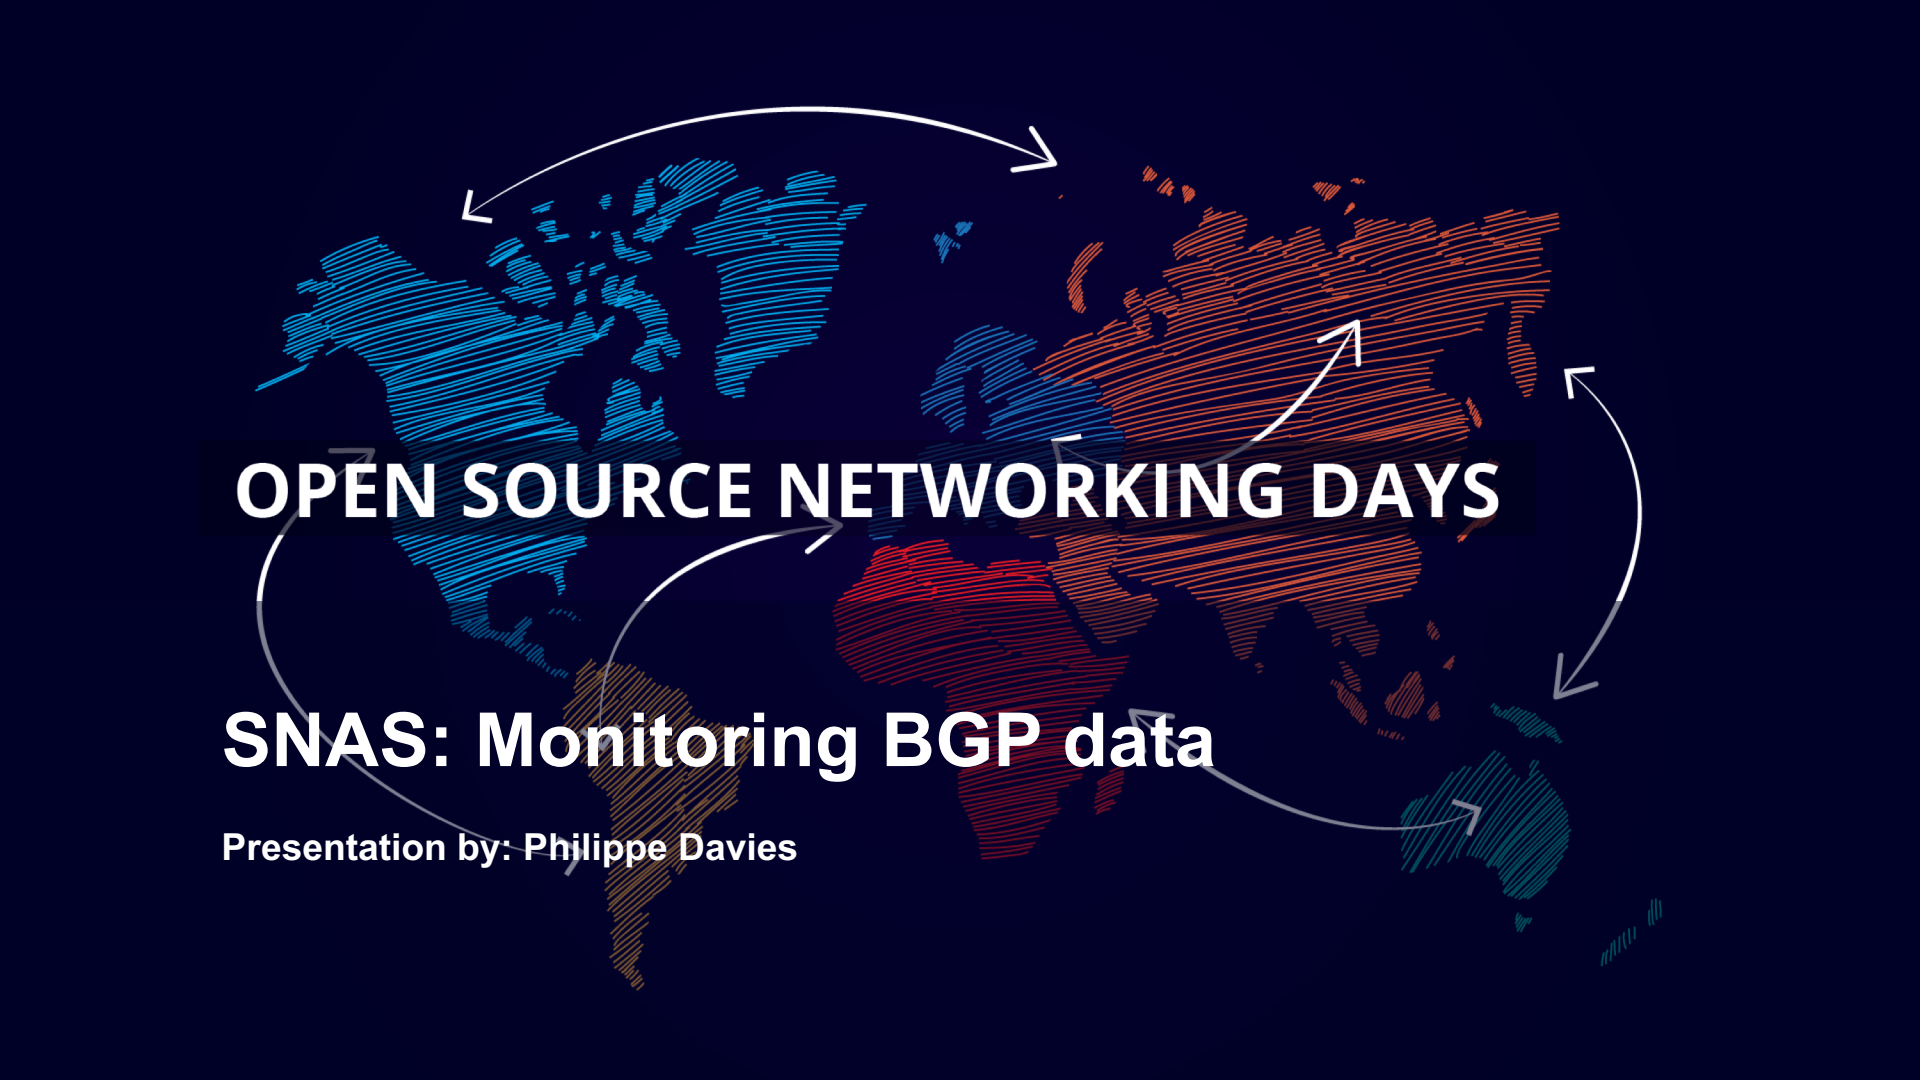

SNAS: Monitoring BGP data

Presentation by: Philippe Davies

Agenda

- A little bit about me
- What are BGP and BMP?
- SNAS introduction
- SNAS system architecture
- SNAS use cases
- History of BGP hijacks
- SNAS demo and alerts
- Next steps for the project

Who am I?

Philippe Davies

- Networking student at Carleton University
  - BIT: Networking program
- Intern at CENGN: Summer 2018
  - Implemented SNAS cluster for Network Analytics
- Contributor to the SNAS project

What is BGP?

- Border Gateway Protocol (RFC1654)
  - Invented in the late 1980s
  - External routing protocol
  - Connects Autonomous Systems (ASes) together
- Connects the entirety of the Internet
  - Primarily used by ISPs and multi homed corporations
- Route advertisements with prefix and AS path

What is BMP?

- BGP Monitoring Protocol (RFC 7854)
- BGP enabled devices forwarding routing data to a centralized location
- Inspection of active and inactive BGP routes

Internet Governance

AS and Prefix Validation

- IRR (Internet Routing Registry)
  - Regional database for public routing entries
  - RIRs such as ARIN, RIPE, AFNIC, APNIC, etc..

- RPKI (Resource Public Key Infrastructure)

RPKI

Resource Public Key Infrastructure

IP Address & AS Number | Digital Certificate

SNAS introduction

- SNAS (Streaming Network Analytics System)
- Linux Foundation Networking project
  - Created by Tim Evens - Principal Engineer at Cisco
- BMP collector
- BMP data analysis
- Grafana dashboards:
  - Peer analysis
  - Prefix history
  - BGP Security (RPKI/IRR and hijacks/leaks)
  - And many more graphs

SNAS System Architecture

- BMP collector (openbmp)
- Kafka bus (internal messaging)
- TimescaleDB (PostgreSQL)
- Grafana (graphs)

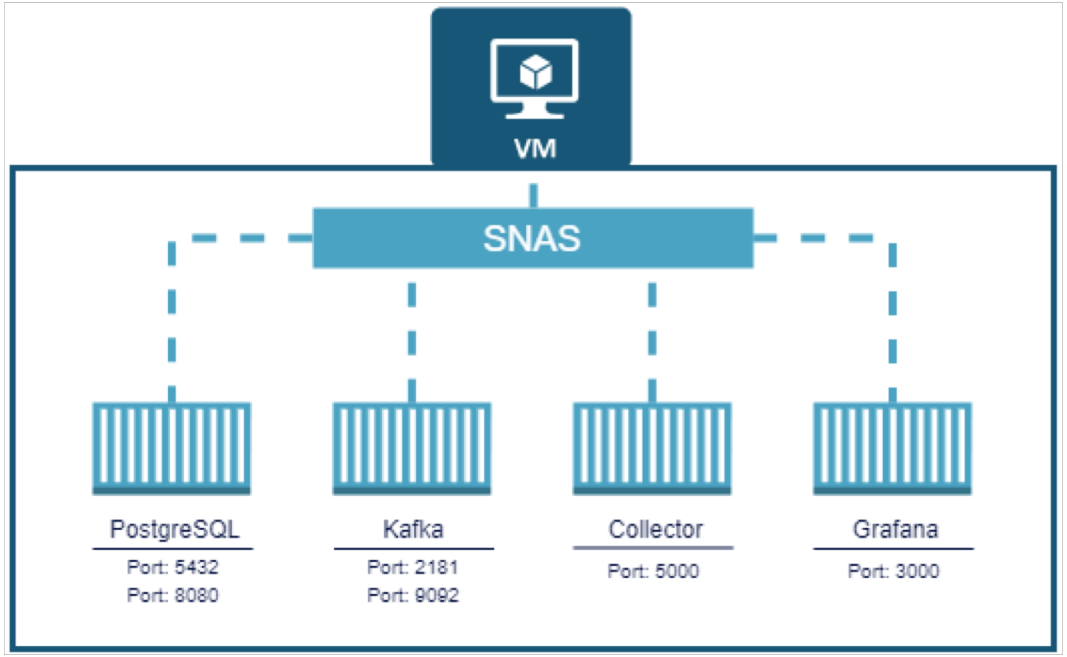

Use cases

- Identifying route leaks
- Identify noisy peers
- Track prefix history
- BGP security analysis

What is a BGP Route Hijack?

- Prefixes being redirected by malicious or unqualified network operators
- Prefix hijacked with more specific route
- Traffic rerouted for theft of information
- Inadvertent leak of country AS restrictions

Example Hijacks

- Youtube Hijack by Pakistan in 2008. Link
- Bitcoin mining hijack in 2014. Link
- DV-LINK (unknown Russian entity) rerouted traffic in 2017. Link
- Rostelecom (Russian telecom) rerouted financial traffic in 2017. Link
- China Telecom BGP hijacking from 2010 to 2017. Link

Youtube Hijack

- Youtube Hijack by Pakistan in 2008. Link

Since BGP relies on a transitive trust model, validation between customer and provider is important. In this case, PCCW (3491) did not validate Pakistan Telecom's (17557) advertisement for 208.65.153.0/24. By accepting this advertisement and readvertising to its peers and providers PCCW was propagating the wrong route. Those who saw this route from PCCW selected it since it was a more specific route. YouTube was advertising 208.65.152.0/22 before the event started and the /24 was a smaller (and more specific) advertisement. According to usual BGP route selection process, the /24 was then chosen, effectively completing the hijack.

China Telecom route manipulation

China Telecom BGP hijacking from 2010 to 2017. Link

Canada to Korea, 2016 – traffic to Government Site

Starting from February 2016 and for about 6 months, routes from Canada to Korean government sites were hijacked by China Telecom and routed through China. Figure 2a shows the shortest and normal route: Canada-US-Korea. As shown in figure 2b, however, the hijacked route started at the China Telecom PoP in Toronto, the traffic was then forwarded inside the Chinese network to their PoP on the US West Coast, from there to China, and finally to delivery in Korea. This is a perfect scenario for long term espionage, where the victim's local protections won't raise alaems about the long term traffic detours. Note that the shortest route between the originators and the destination is definitely not through two China Telcom PoPs in North America to China and only then to Korea. That this pattern continued for six months is good evidence that this was no short term misconfiguration or temporary internet conditions disruption. This attack repeated later for shorter time durations.

Figure 2a: The normal and shortest route from Canada to Korea before the attack.

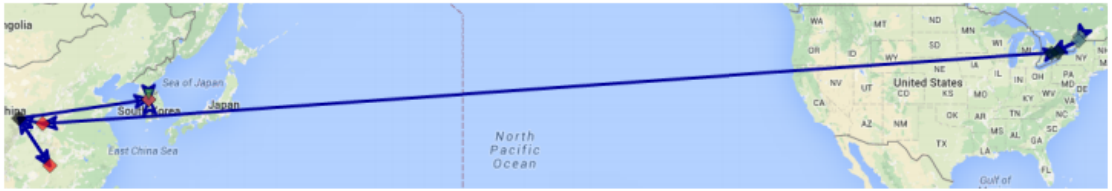

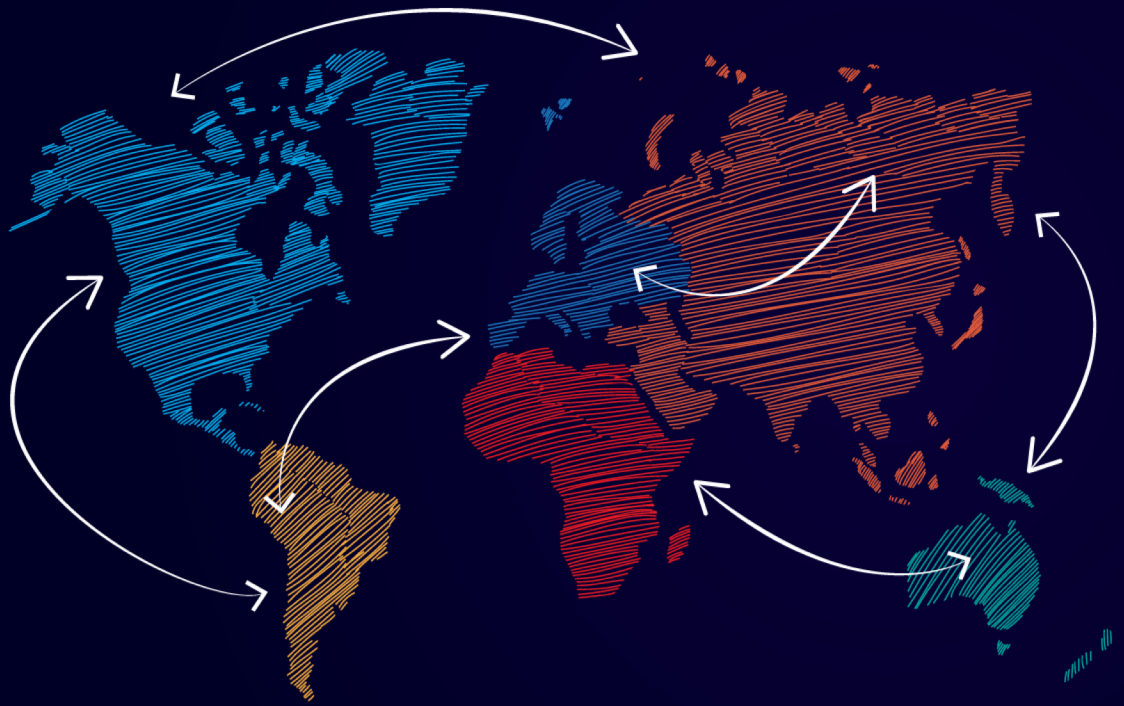

Peer Analysis Demo

Peer Analysis

Last 1 hour UTC

Acceptable Ranked Peers by Router

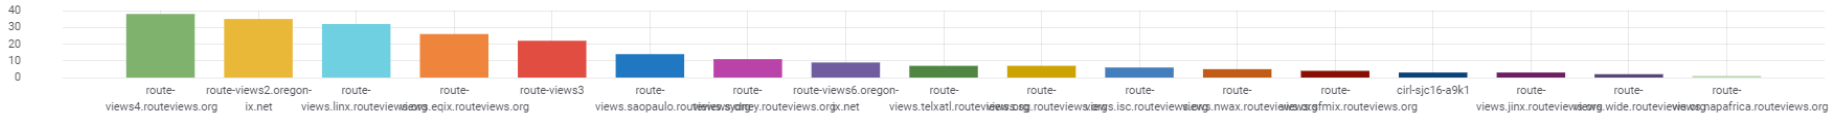

Peers by State

	current	percentage
up	431	52%
down	392	48%

IPv4 Peer RIB Sizes

	current	percentage
full (> 700k)	170	55%
customer_only (< 9K)	75	24%
poor (> 9k)	27	9%
partial (> 300K)	19	6%
weak (> 100K)	19	6%

Peers by AFI

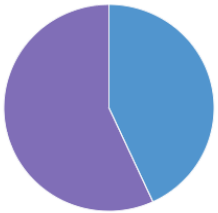

	current	percentage
IPv6	355	43%
IPv4	468	57%

IPv6 Peer RIB Sizes

	current	percentage
full (> 50K)	102	51%
partial (> 20K)	41	20%
customer_only (< 1K)	40	20%
poor (> 1k)	12	6%
weak (> 10K)	6	3%

Peer Rank

Rank	peer_ip_ver	router_name	peer_name	peer_as_name	peer_ip	peer_asn_match	state	rib_score	advertisements_score	withdraws_score	Connections	Transit Connections
1	IPv4	route-views.telxatl.routeviews.org	10gigabitethernet1-3.core1.atl1.he.net	HURRICANE	198.32.132.75	No Match	up	1000	1000	1000	8,759	7,377
1	IPv6	cirl-sjc16-a9k1	2001:470:1:4ea::1	HURRICANE	2001:470:1:4ea::1	No Match	up	1000	1000	1000	8,759	7,377
1	IPv4	route-views.linx.routeviews.org	40ge1-3.core1.lon2.he.net	HURRICANE	195.66.224.21	No Match	up	1000	1000	1000	8,759	7,377
1	IPv4	cirl-sjc16-a9k1	v416.core1.sjc1.he.net	HURRICANE	64.71.176.49	No Match	up	1000	1000	1000	8,759	7,377

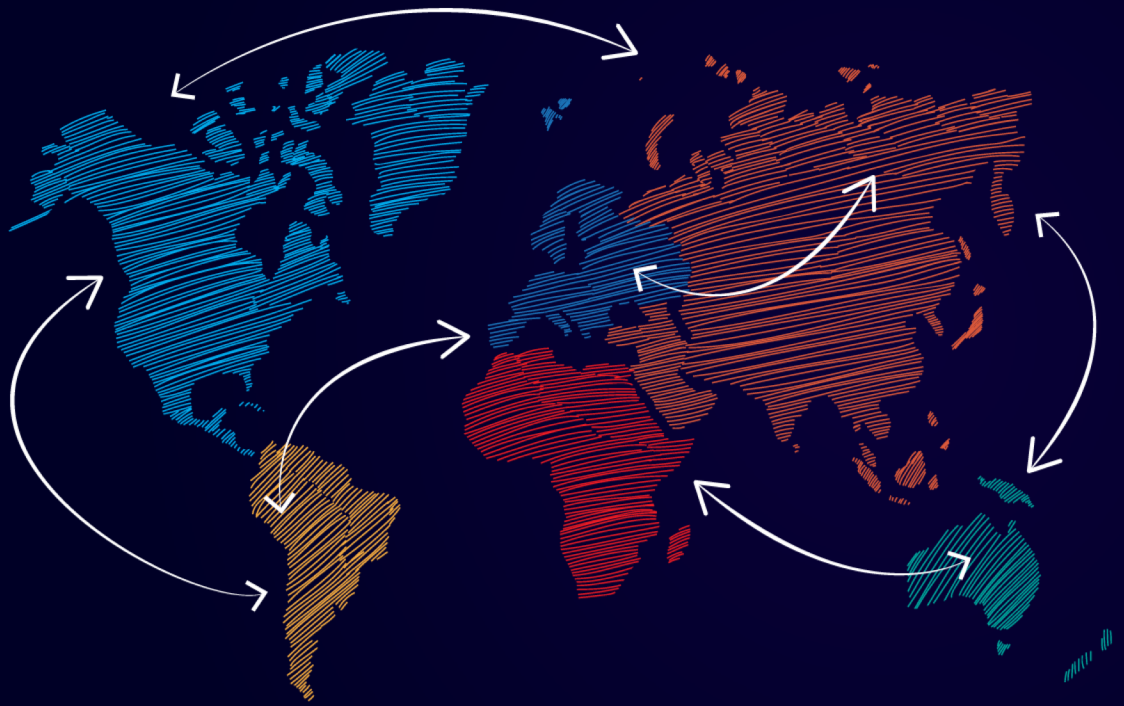

Prefix History Demo

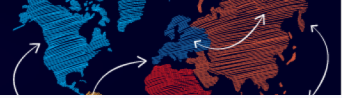

Prefix History (by Peer)

Last 1 hour UTC

Router: route-views2.oregon-ix.net

Peer: 0.lo0.GW66.DCA6.ALTER.NET

Prefix Advertisements & Withdrawals

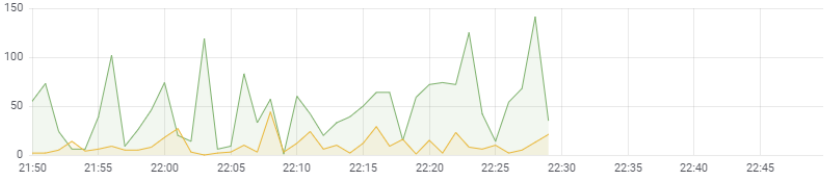

	max	avg	total
updates	141	49	1945
withdraws	44	10	399

Changes by Peer

	max	avg	total
0.lo0.GW66.DCA6.ALTER.NET	154	59	2344

Prefix History

Search:

lastmodified	event	routername	peername	prefix	origin_as	as_path	communities
2018-10-27T22:28:46Z	Advertised	route-views2.oregon-ix.net	0.lo0.GW66.DCA6.ALTER.NET	190.73.96.0/19	8048	701 2914 52320 8048	
2018-10-27T22:28:46Z	Advertised	route-views2.oregon-ix.net	0.lo0.GW66.DCA6.ALTER.NET	200.44.192.0/19	8048	701 6461 52320 8048 8048 8048 8048 8048 8048 8048 8048 8048	
2018-10-27T22:28:46Z	Advertised	route-views2.oregon-ix.net	0.lo0.GW66.DCA6.ALTER.NET	200.93.0.0/19	8048	701 6461 52320 8048 8048 8048 8048 8048 8048 8048 8048 8048	
2018-10-27T22:28:46Z	Advertised	route-views2.oregon-ix.net	0.lo0.GW66.DCA6.ALTER.NET	190.201.224.0/19	8048	701 2914 52320 8048 8048 8048 8048 8048 8048	
2018-10-27T22:28:46Z	Advertised	route-views2.oregon-ix.net	0.lo0.GW66.DCA6.ALTER.NET	200.90.64.0/19	8048	701 6461 52320 8048 8048 8048 8048 8048 8048 8048 8048 8048	
2018-10-27T22:28:46Z	Advertised	route-views2.oregon-ix.net	0.lo0.GW66.DCA6.ALTER.NET	190.205.224.0/19	8048	701 6461 52320 8048 8048 8048 8048 8048 8048 8048 8048 8048	
2018-10-27T22:28:47Z	Advertised	route-views2.oregon-ix.net	0.lo0.GW66.DCA6.ALTER.NET	186.90.224.0/19	8048	701 1299 23520 8048	
2018-10-27T22:28:47Z	Advertised	route-views2.oregon-ix.net	0.lo0.GW66.DCA6.ALTER.NET	177.137.219.0/24	263087	701 3356 3549 262733 263087	

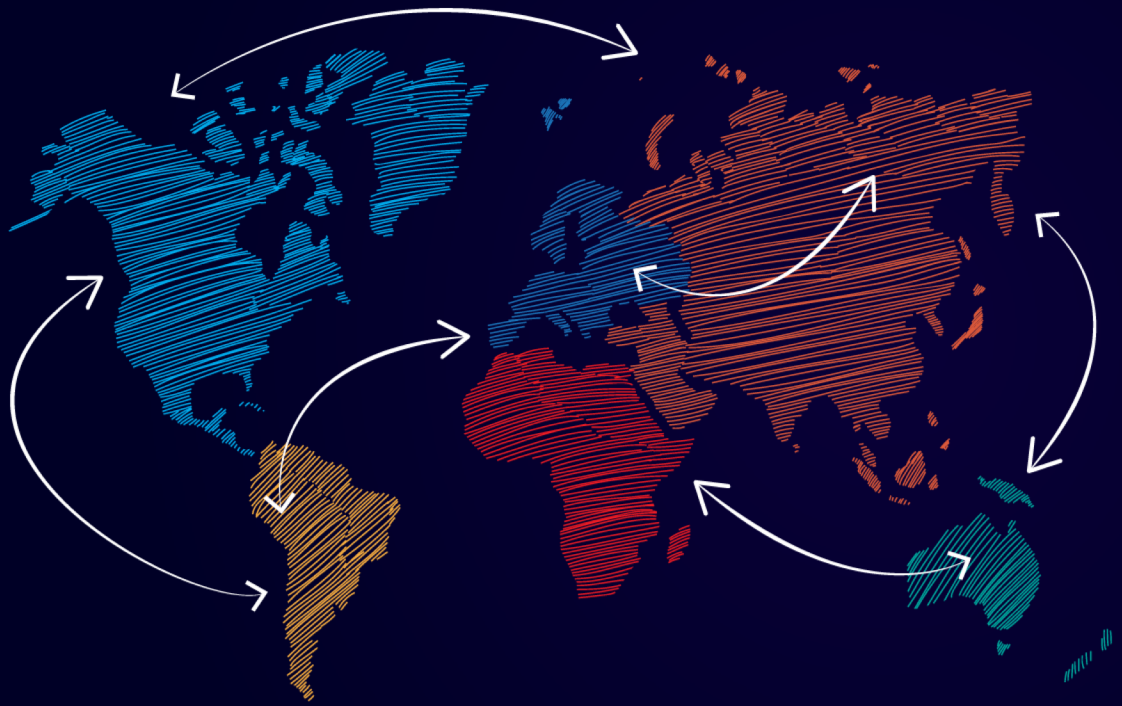

RPKI & IRR Demo

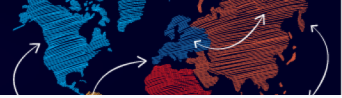

BGP Sec: RPKI & IRR

Last 3 hours UTC

RPKI Prefixes in Violations

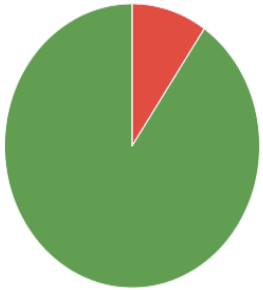

	current	percentage
mismatches	8024	10%
matches	75828	90%

IRR Prefixes in Violations

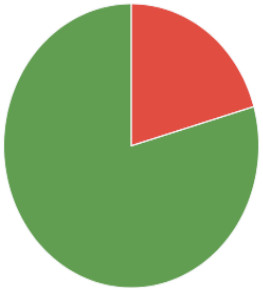

	current	percentage
mismatches	101272	21%
matches	391584	79%

RPKI ROAS

61,707

IRR Entries

1,882,758

Prefixes in Violation of IRR

Search:

timestamp	Prefix	Received Origin	Expected Origin	IRR Source
2018-10-27T20:24:06Z	42.225.128.0/17	4837	135091	nttcom
2018-10-27T20:24:06Z	42.225.0.0/17	4837	135091	nttcom
2018-10-27T20:32:49Z	103.235.222.0/23	23724	53587	arin
2018-10-27T20:32:49Z	103.235.220.0/23	23724	53587	arin
2018-10-27T20:48:11Z	59.151.64.0/19	23724	9308	radb
2018-10-27T20:48:11Z	59.151.96.0/19	23724	17428	apnic
2018-10-27T20:53:49Z	82.44.128.0/17	5089	5462	ripe
2018-10-27T20:54:51Z	75.65.128.0/17	7922	22258	level3
2018-10-27T20:54:51Z	75.65.0.0/17	7922	22258	level3
2018-10-27T21:03:28Z	177.66.240.0/22	65013	53005	radb
2018-10-27T21:05:45Z	148.164.0.0/24	5	23154	arin
2018-10-27T21:06:56Z	103.226.227.0/24	134026	59276	nttcom

Prefixes in Violation of RPKI

Search:

timestamp	Prefix	Received Origin	Expected Origin
2018-10-27T21:25:14Z	190.214.64.0/21	28011	14420
2018-10-27T22:07:02Z	181.60.106.0/23	10620	14080
2018-10-27T22:07:02Z	181.60.104.0/23	10620	14080

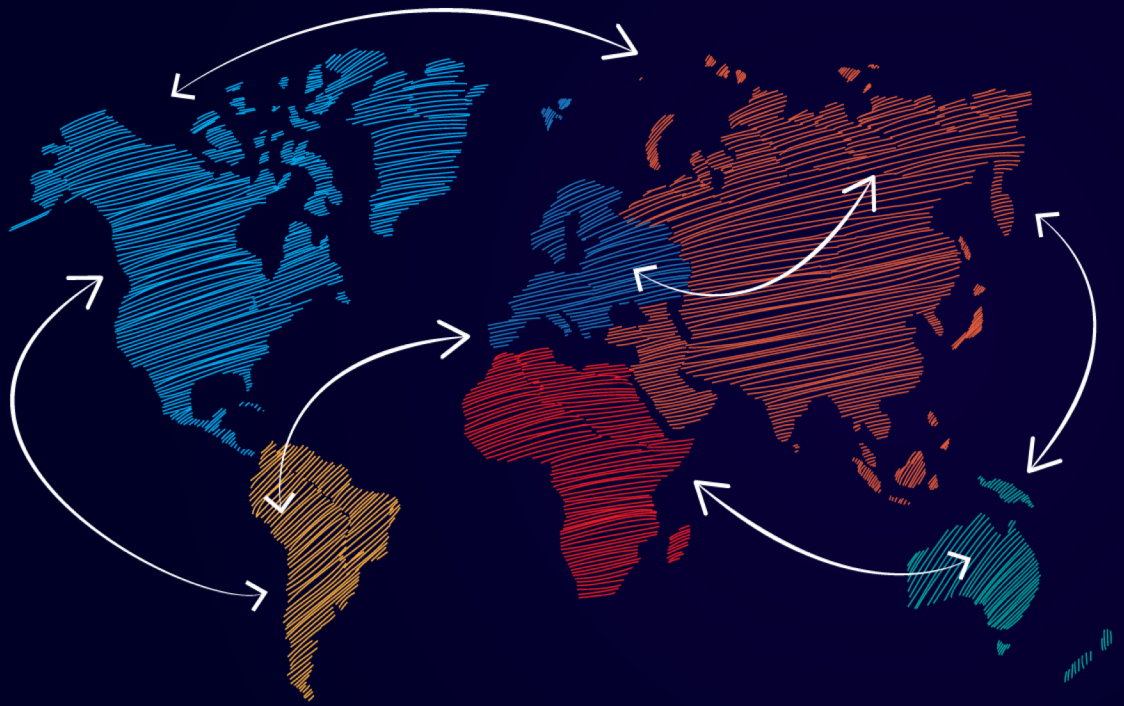

Hijack/Leak Demo

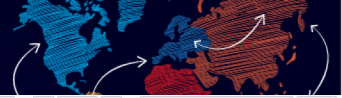

BGP Sec: Hijack/Leak

Last 3 hours UTC

ASN 0

Live Notifications

See the live notifications at SNAS/Alerts

Transit Leak/Hijacks	Upstream Leak/Hijacks	Origin HiJacks
1 Alerts	8 Alerts	3 Alerts

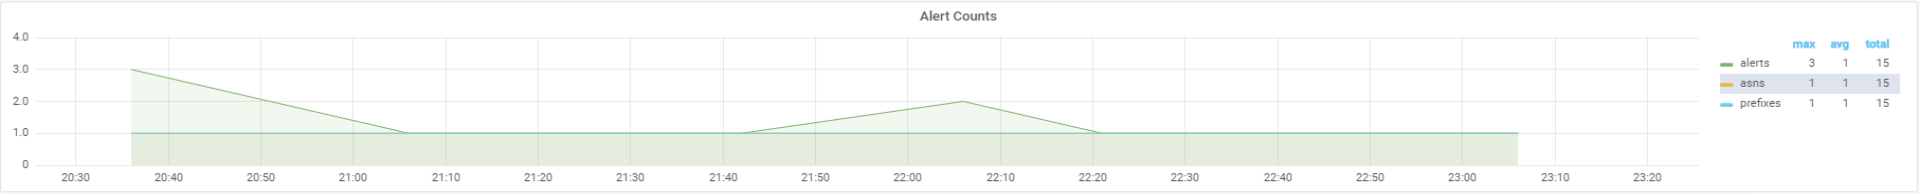

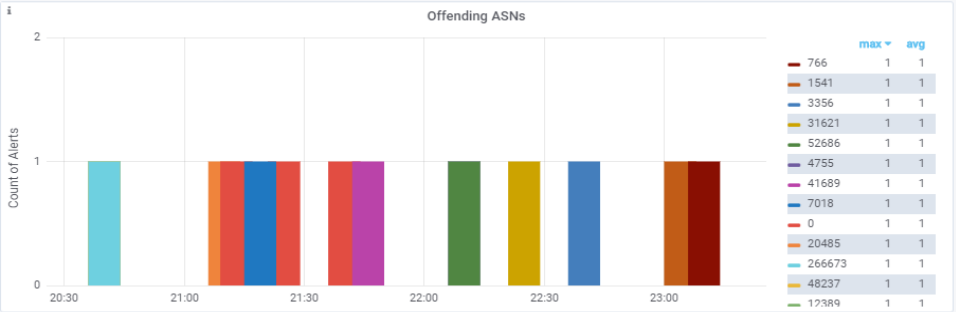

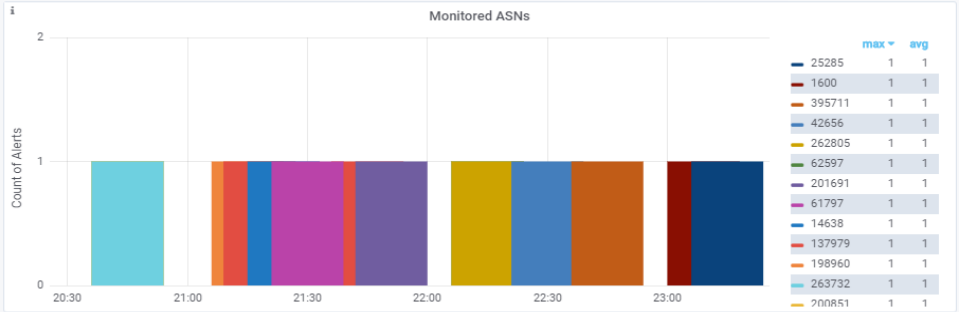

Alerts

Search:

Time	type	Monitored	Offending	message	History

Hijack alerts

- Proof of concept applications
  - Identification of hijacks/leaks in real time
  - Available on Gitter. link

Future of SNAS

- Add BGPSec support to SNAS
- Increased efficiency in collector scripts
- Private AS notifications
- Hijack/Leak notification (beyond gitter)

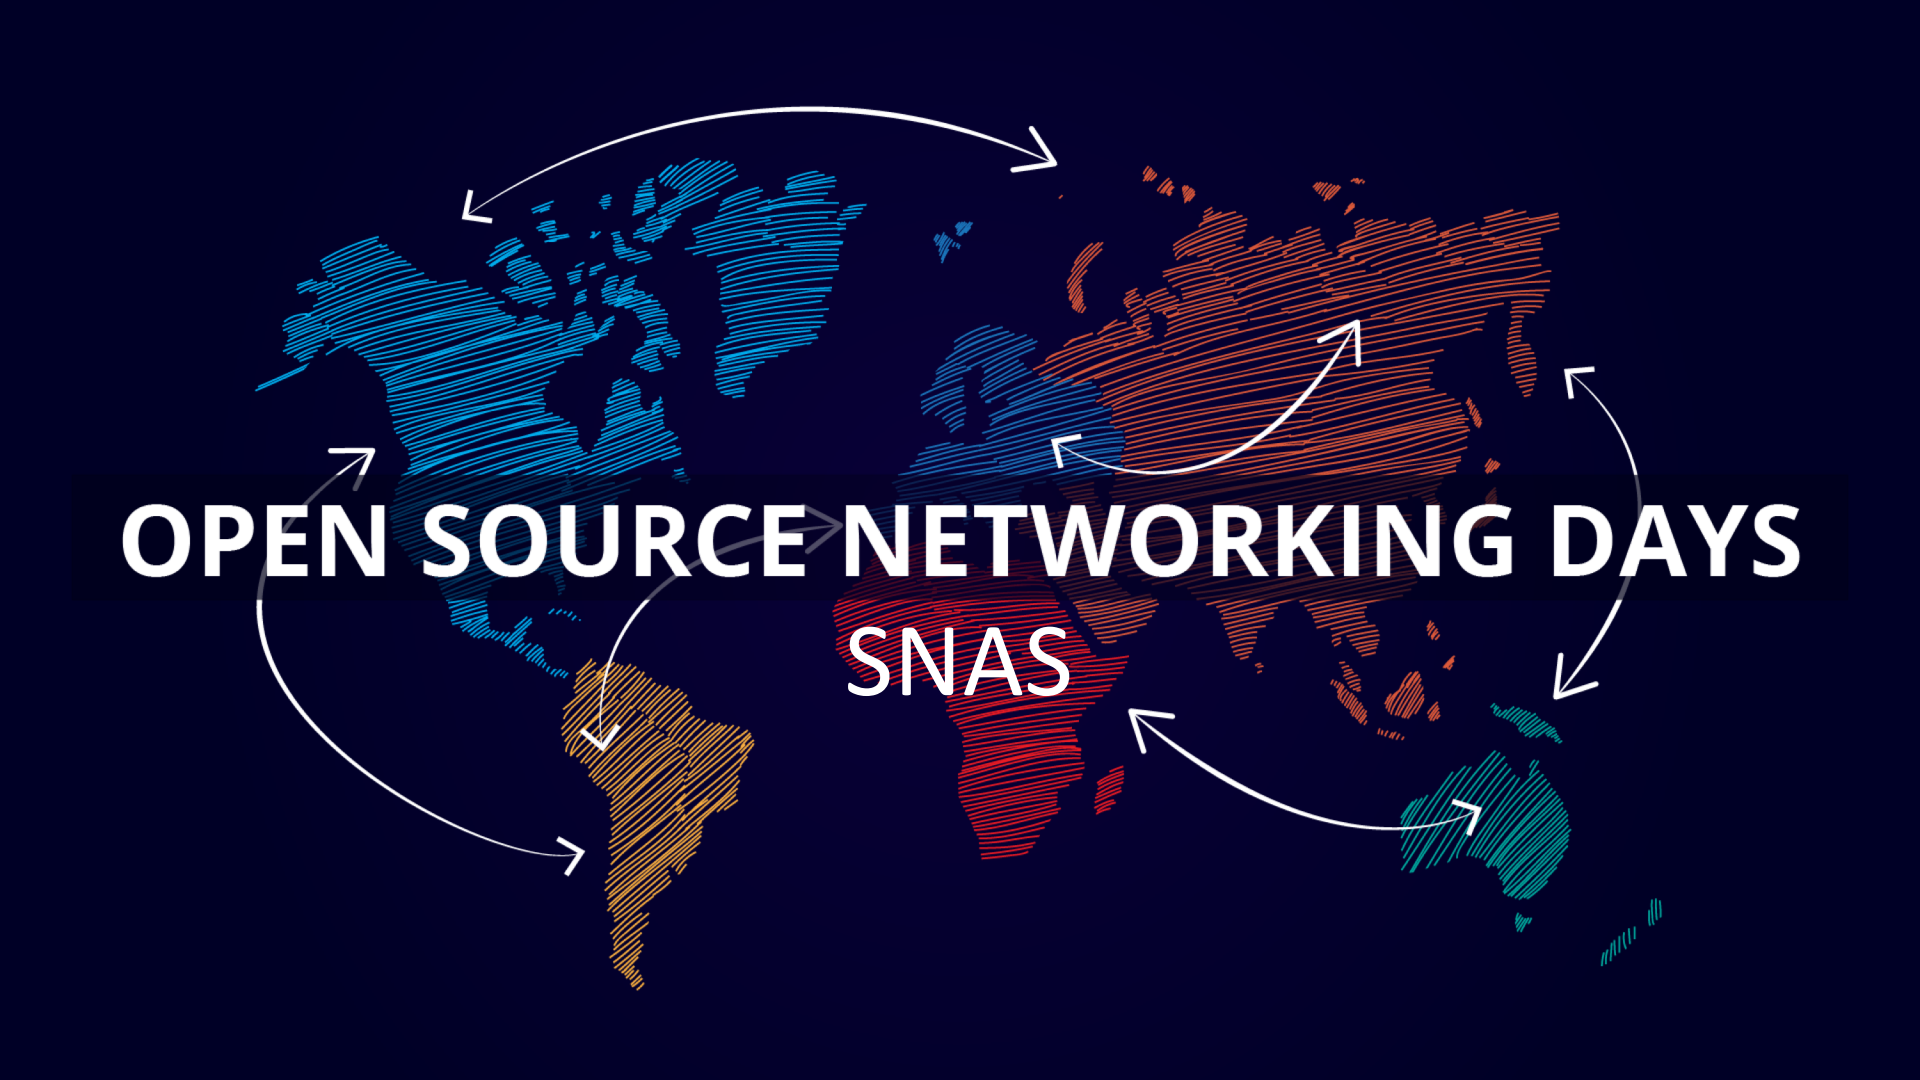
OPEN SOURCE NETWORKING DAYS
SNAS