Yet Another Step for Origin Validation

Eugene Bogomazov
Qrator Labs

Who are we?

- Qrator Labs
  - A DDoS attack mitigation company
- Radar
  - BGP monitoring
  - IETF BGP security drafts: Roles/OTC, ASPA

State of Affairs: POV pushing

- Filtering on:
  - Some Tier-1s
  - Huge IXPs
- Updates for a RPKI-RTR cache
- RIRs support RPKI hosted systems
- ROA is needed for the ASPA drafts

RPKI Deployment

- AT&T rejects invalids on peering sessions
- Nordunet rejects invalids on all EBGP sessions
- KPN / AS 286 rejects invalids on customer sessions
- Seacomm & Workonline drop invalids per April 2019
- INEX, AMS-IX, DE-CIX, France-IX, Netnod
- MSK-IX
- XS4ALL
- **THE RIPE MEETING NETWORK!!!**
- IX.br (.... soon? :-)
- **You.... ?**

**From Job Sniders' last RIPE presentation*

What else do I want?!

This is not about increasing ISP motivation

Why we need POV?

- It's all about traffic redirection
- Who is the real owner?
- What can we trust?

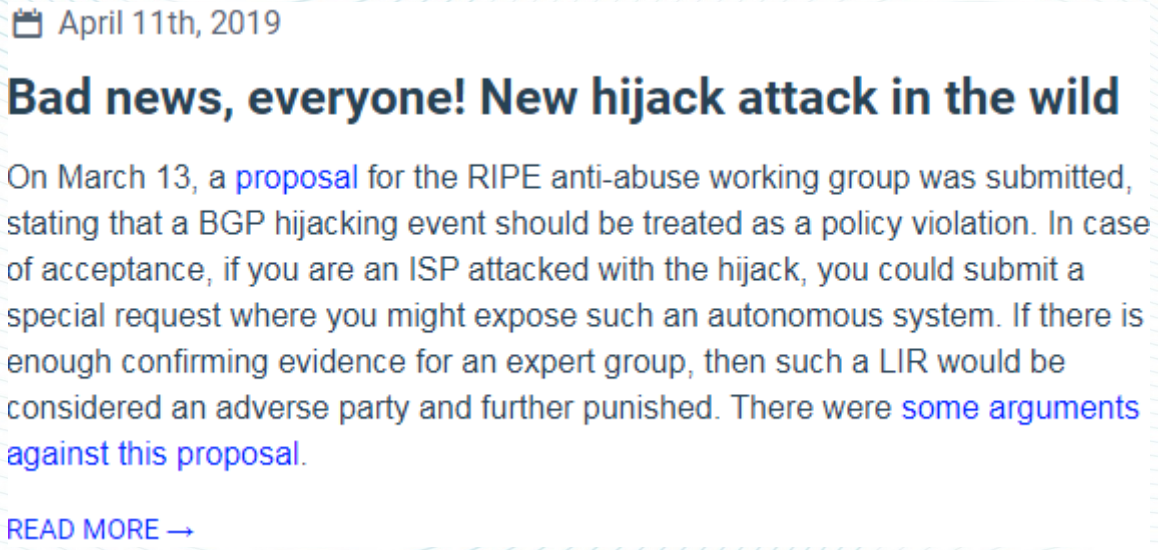

April 11th, 2019

Bad news, everyone! New hijack attack in the wild

On March 13, a proposal for the RIPE anti-abuse working group was submitted, stating that a BGP hijacking event should be treated as a policy violation. In case of acceptance, if you are an ISP attacked with the hijack, you could submit a special request where you might expose such an autonomous system. If there is enough confirming evidence for an expert group, then such a LIR would be considered an adverse party and further punished. There were some arguments against this proposal.

READ MORE →

Trust in registries

- There are no fake registration objects!
  - With a few known exceptions
  - There are no fake **RPKI objects!**

**RADB objects for some ISPs*

Showrooms (for AS197068)

AS Info | Graph v4 | Prefixes v4 | Peers v4 | Whois | IRR

Prefix	Description
45.8.210.0/24	Avento Mt Limited
45.116.91.0/24	Lazada South East Asia Pte. Ltd.
78.155.198.0/24	HLL LLC
178.248.232.0/21	HLL LLC
178.248.232.0/23	HLL LLC
178.248.234.0/23	HLL LLC
178.248.236.0/23	HLL LLC

ACTIVE AT 08/10/2019 (18) | NEW (3) | LEFT (2) | RETURNING (1)

Prefix	ROA	Route Object	Graph	Active now	History
45.8.210.0/24	Unknown	Valid	Graph →	✓	
45.116.91.0/24	Unknown	Valid	Graph →	✓	
78.155.198.0/24	Valid	Valid	Graph →	✓	
178.248.232.0/21	Valid	Valid	Graph →	✓	
178.248.232.0/23	Valid	Valid	Graph →	✓	
178.248.234.0/23	Valid	Valid	Graph →	✓	

prefix	In RIS	RIPE IRR	Other IRRs	RPKI
178.248.232.0/21	yes	yes	no	☺
178.248.232.0/23	yes	yes	no	☺
178.248.232.0/24	no	yes	no	☺
178.248.232.0/32	no	yes	no	☺
178.248.232.1/32	no	yes	no	☺
178.248.232.10/32	no	yes	no	☺
178.248.232.100/32	no	yes	no	☺
178.248.232.101/32	no	yes	no	☺
178.248.232.102/32	no	yes	no	☺
178.248.232.103/32	no	yes	no	☺

Showing 1 to 10 of 2,080 entries

ASN	Prefix	Max Length	Trust Anchors
197068	78.155.198.0/24	24	RIPE NCC RPKI Root
197068	178.248.232.0/21	32	RIPE NCC RPKI Root
197068	185.65.148.0/22	32	RIPE NCC RPKI Root
197068	185.94.108.0/22	32	RIPE NCC RPKI Root
197068	192.166.48.0/24	24	RIPE NCC RPKI Root
197068	2a03:70c0::/32	64	RIPE NCC RPKI Root

Key features (for the previous slide)

- We sign our own address space
  - Though our clients don't sign ROAs
- We use a max maxLength
  - Using 32 for ROAs
  - Having route objects for blackhole
- Adopt the ROA validation status to Route objects

Our position

- Using prefixes as a regular ISP
- Updating functionality on our web site
- Investigating bad cases

- Sometimes we don't agree with the current approaches

Too many questions

- Are current filtering policies good?
- What to do with the maxLength?
- Do we need a maxLength analog for IRR objects?
- Is it time for the RIR policy to check with IRR?

**We will cover them step by step*

Actually...

There are even more questions
(Maybe not for today)

QRATOR LABS

And so it begins

Digging a ground for a ground truth

The basics

- Were covered yesterday on MANRS Tutorial
- Let's dive into details

*Great introduction to BGP area

Route object vs ROA

- Both are <prefix, origin ASN> pairs
  - In addition, ROA has a maxLength
  - But there is an exact/covered match type for IRR
- With very different usage semantics

**Are more-specifics allowed? Exact – not, covered – yes*

How different?

	Route objects	ROA
Format	<Prefix, Origin ASN>	
MaxLength	Exact/cover match type	Min/Max/Intermediate
Filtering type	Prefix Whitelist	Pair Blacklist
Leading role	Transit ISP	Prefix owner
Trust roots	Set of databases	One hierarchical tree
Trust level	Depends on database	Full trust
Sub-prefix attacks	Possible if only target is in CC	Possible without AS_PATH check

IRR for a transit ISP

- Transit ISP wants to create a filter
- Finds ASNs in its Customer Cone (with an AS_SET)
- Chooses registries
- Extracts prefixes
- Chooses an exact/covered match type

**Many places to choose*

About exact match

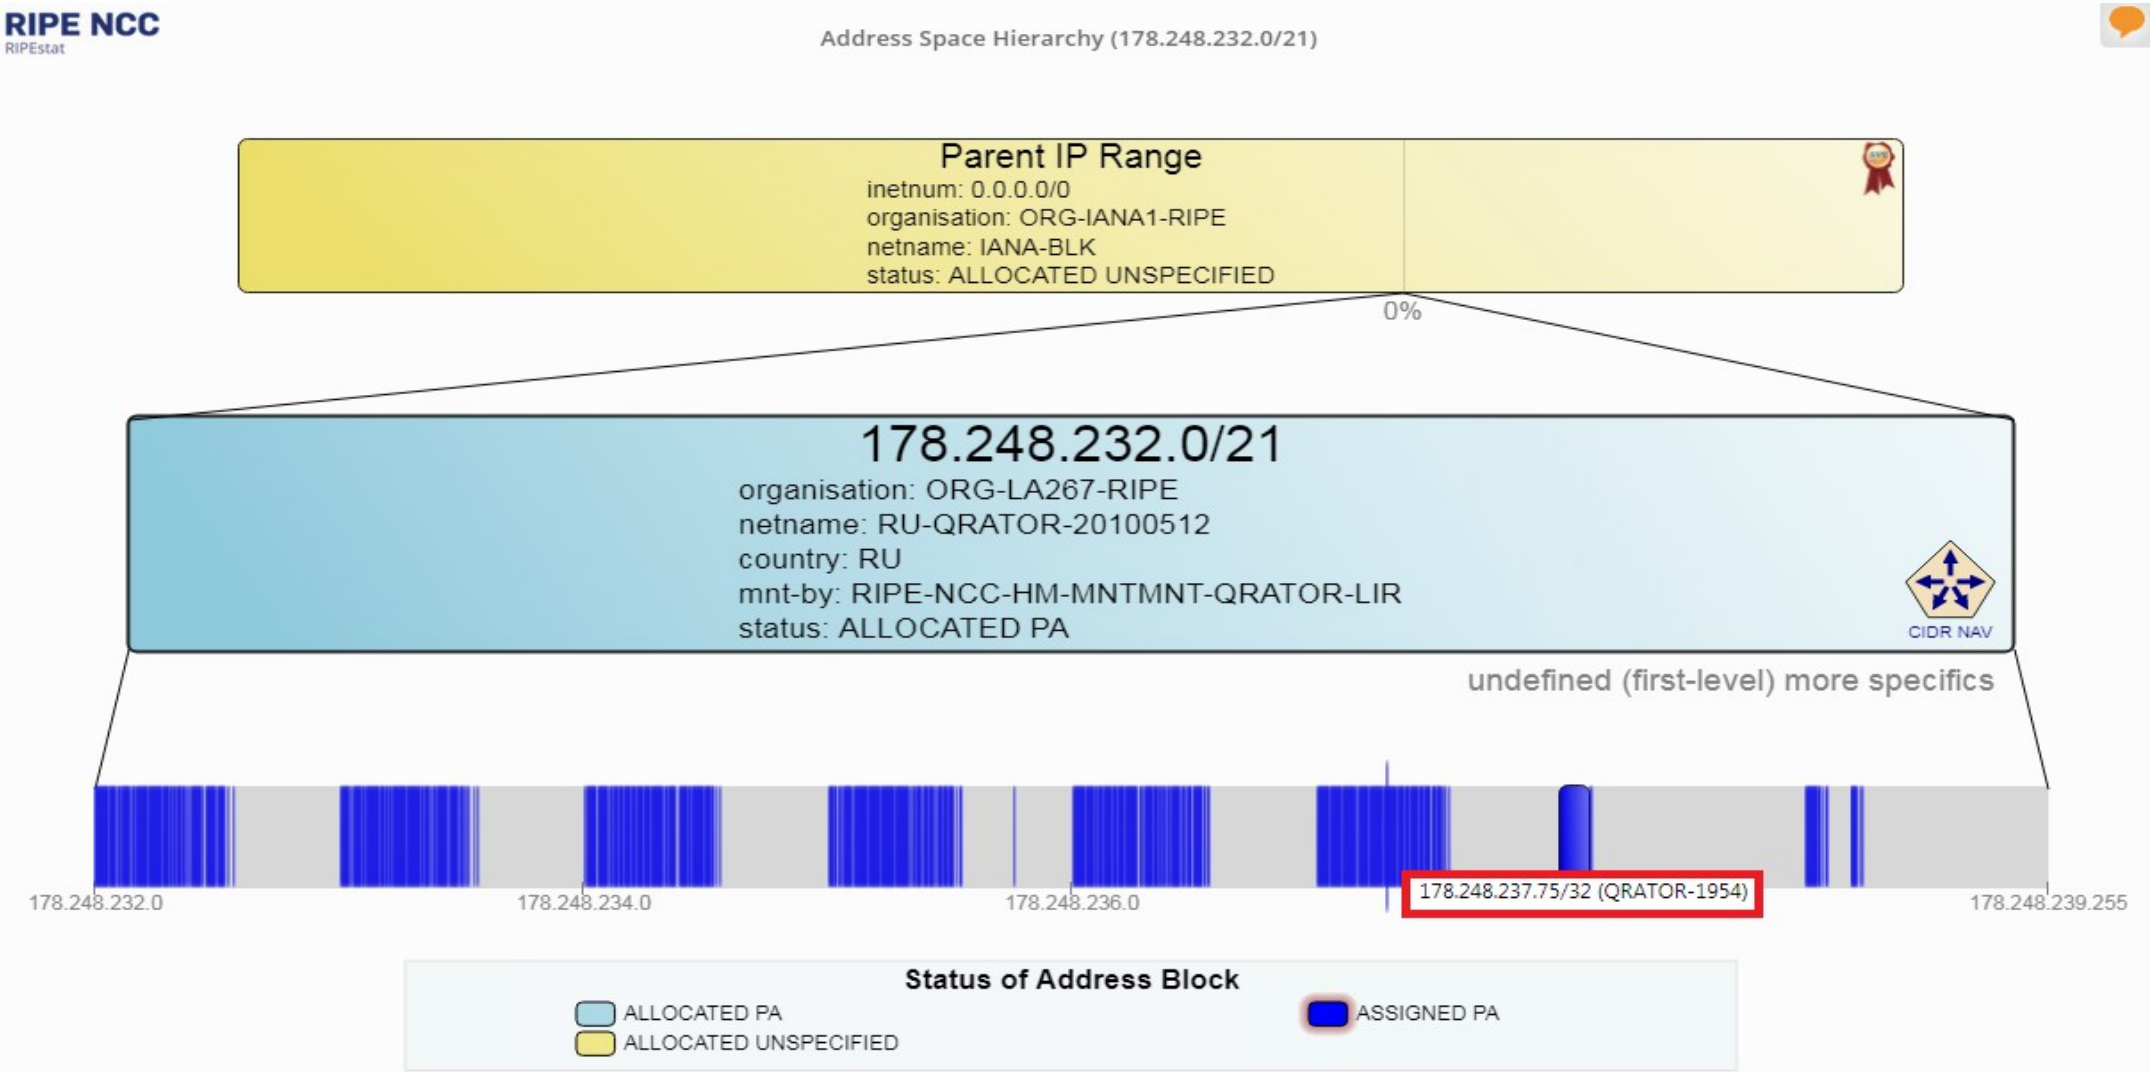

IRR for a prefix owner

- Creates objects for its own traffic
- Wants to bypass existing filters
  - Not to stop others
  - “Do and forget” -> outdated information

**Or tries to prepare for future filters*

ROA motivation

- For a prefix owner
  - Create objects to claim a possession
  - Stop others from your address space misuse
- For a transit ISP
  - Drop all the bad routes that all the prefix owners want

Which pill to choose: IRR

- A Prefix Whitelist
- Filters can be placed on customer and peer links
- Owners create objects to pass upstream filters
- AS-SET - place of errors
- Bad registrars

- One cannot attack a target outside of transit CC

Which pill to choose: ROA

- A Pair Blacklist
- Filters can be placed on any links
- Owners create objects to stop the others
  - To stop them on filtering points
- Cryptographic load (?)

maxLength

> One advantage of minimal ROA length is that the forged origin attack does not work for sub-prefixes that are not covered by overly long max length. For example, if, instead of 10.0.0.0/16-24, one issues 10.0.0.0/16 and 10.0.42.0/24, a forged origin attack cannot succeed against 10.0.666.0/24. They must attack the whole /16, which is more likely to be noticed because of its size.

*From RFC7115

IRL — CloudFlare case

- Main features:
  - It was a sub-prefix attack
  - ROAs had a min maxLength
  - It was a “Drop Invalid” policy on the upstreams

- But we will look at even simpler case

The deep-dive into how Verizon and a BGP Optimizer Knocked Large Parts of the Internet Offline Monday

Martin J Levy

27 июня 2019 г., 01:22 AM

A recap on what happened Monday

QRATOR LABS

Before

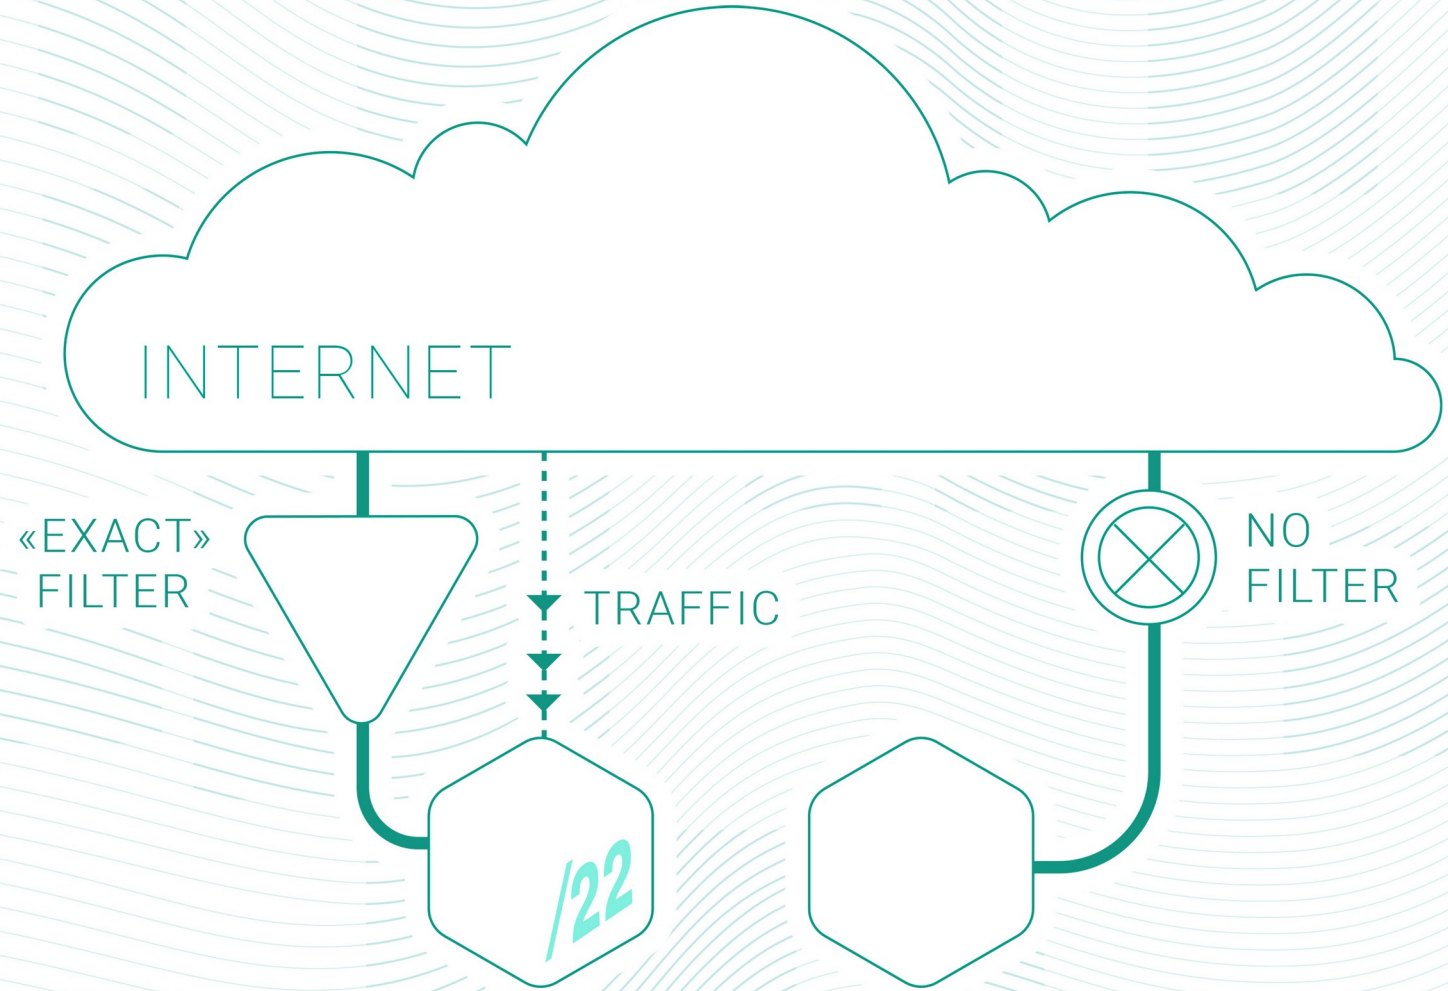

QRATOR LABS

The attack

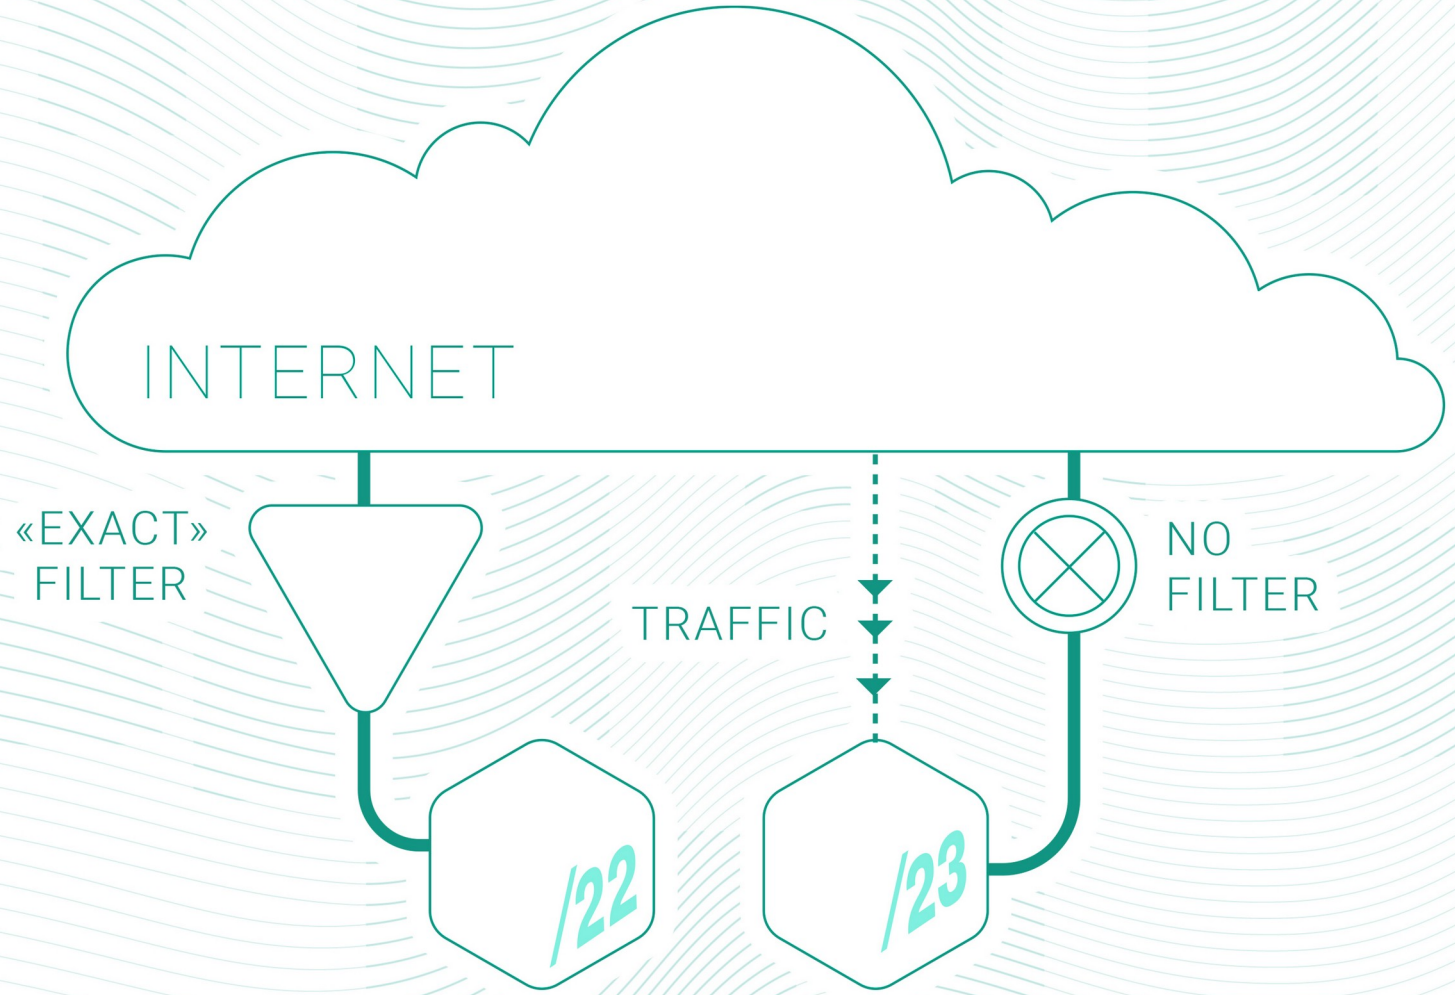

More-specifics do not work

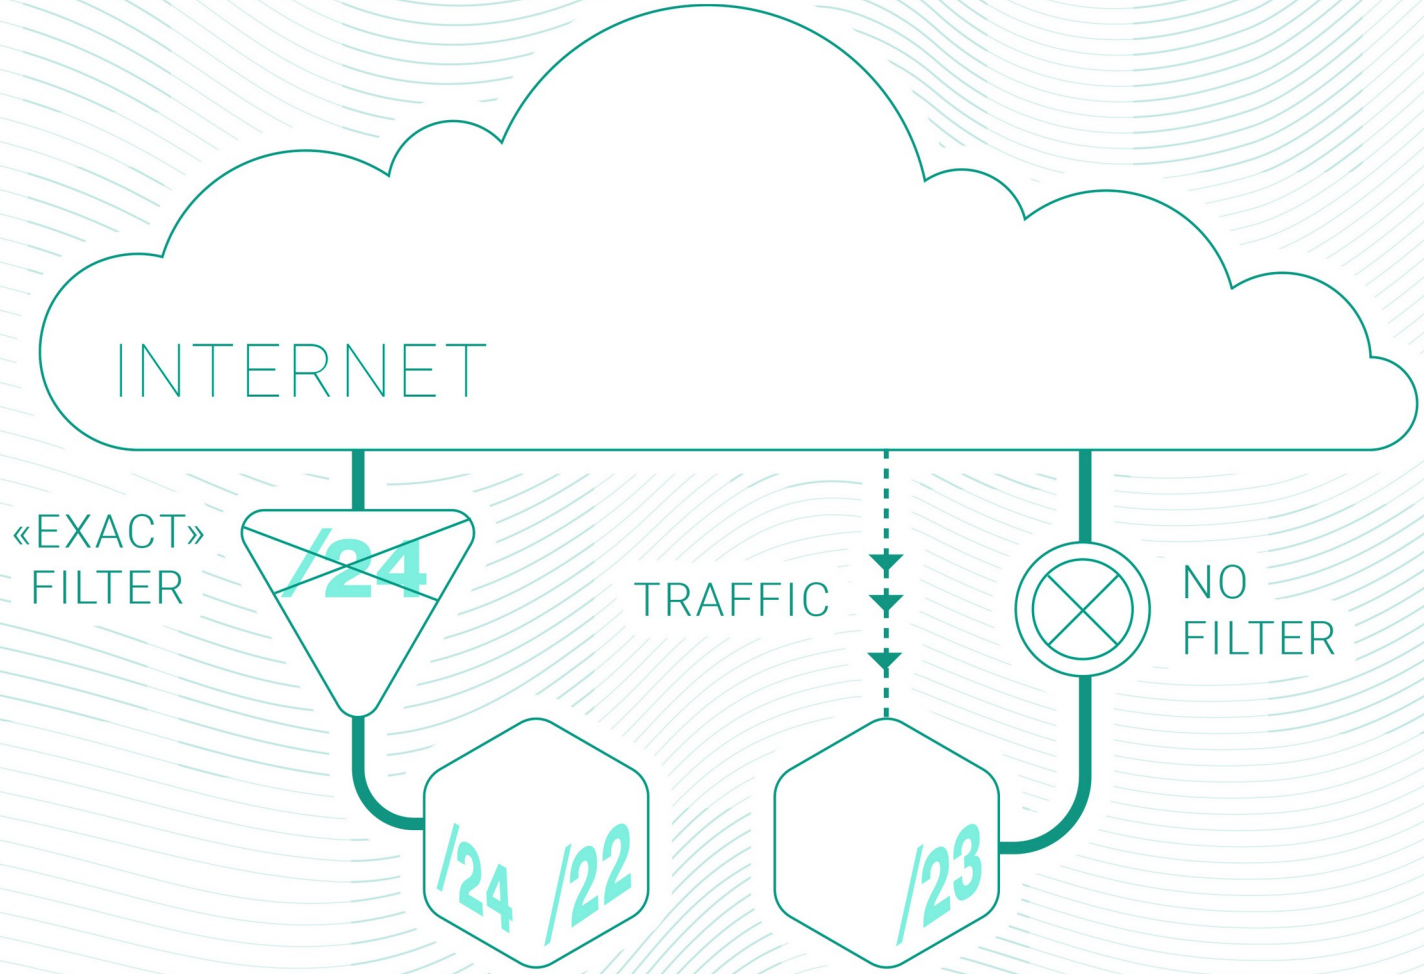

Moral

- Community conclusion: policies are great
  - Discussion about ROA deployment
  - Not about policies problems during partial deployment
- mailman.nanog.org for June/Jule
  - Keyword: CloudFlare

☐ ☆	Dmitry .. Matt, Mark 48	Inbox	Nanog	CloudFlare issues? - On 6/Jul/19 23:44, Matt Corallo wrote: > On my test net I take ROA_INVAL...	Jul 7
☐ ☆	Francois .. Job 13	Inbox	Nanog	Re: CloudFlare issues? - > Anyway, you can now enjoy https://rpki.net/s/rpki-test even more! :-)...	Jul 4
☐ ☆	Stephen .. Bryan 13	Inbox	Nanog	Intermittent "bad gateway" - > > Cloudflare did fall over for a bit this morning. > > mdr > -- > So...	Jul 2
☐ ☆	Patrick .. Randy 8	Inbox	Nanog	Are network operators morons? [was: CloudFlare issues?] - >> perhaps the good side of this sa...	Jun 26
☐ ☆	Job, Tom, TIM, Alex 4	Inbox	Nanog	BGP filtering study resources (Was: CloudFlare issues?) - Was: CloudFlare issues?) > > Job als...	Jun 25
☐ ☆	Martin J. Levy	Inbox	Nanog	How Verizon and a BGP Optimizer Knocked Large Parts of the Internet Offline Today - Cloudfla...	Jun 25
☐ ☆	Jared, ML, Max 5	Inbox	Nanog	Verizon Routing issue - . Why Cloudflare did not immediately announced all their address spac...	Jun 24

Another POV

How to Sign a ROA?

It's simple:

1. https://my.ripe.net/#/rpki
2. Sign only aggregates (inetnum);
3. Set max_length to 32 in IPv4 (128 in IPv6);

**From the last ROA signing party*

Which MaxLength to choose?

		Min	Average	Max
Sub-prefix attacks	**Filter bypass?**	Not Possible	Possible	
	Traffic return (over filter)	Not Possible or Not Immediate	Depends	Always possible
	Monitoring	Not needed	Is needed	
Room for TE		Create new ROA	Yes	Yes
Blackhole		Possible		From the box

**In the case of a "Drop Invalid" policy*

3 ways of POV

- «Valid» - everything is going along with a policy
- «Invalid» - some policy is violated
- «Not Found» - there is no suitable policy

“Drop Invalid” as a default

- In ideal world sub-prefix attacks would be gone
- Make explicit exceptions if necessary
  - Thus supporting good trends (RFC8212)
- Maybe look at less specifics?
  - Just not a good idea

- Real question — are there any other bad cases?
  - The similar discussion is going for an ASPA

#404

- ROA still covers only 17% of prefixes
  - Thus «Not Found» routes cannot be dropped!
- «Less-specifics are less harmful» is a mistake
  - Especially for non-routable address space
  - Or for uRPF

**Moral: “Not Found” status is not equal to valid*

Return of the IRR

- A Prefix Whitelist
- A filter policy is defined by a filtering point
- Independent decisions
  - Prefix owners have no influence on their transit ISPs

IRR Filtering BCP

We don't have one

Do we need it?

IRR filtering problems

- Different level of trust to a RIR/RADB/whatever
- Conflicts with a RPKI POV (first step)
- Prefix delegation to a **multihome** customer
  - A sub-allocation problem
- Who, where and why should deploy it

Multihome with a PA block

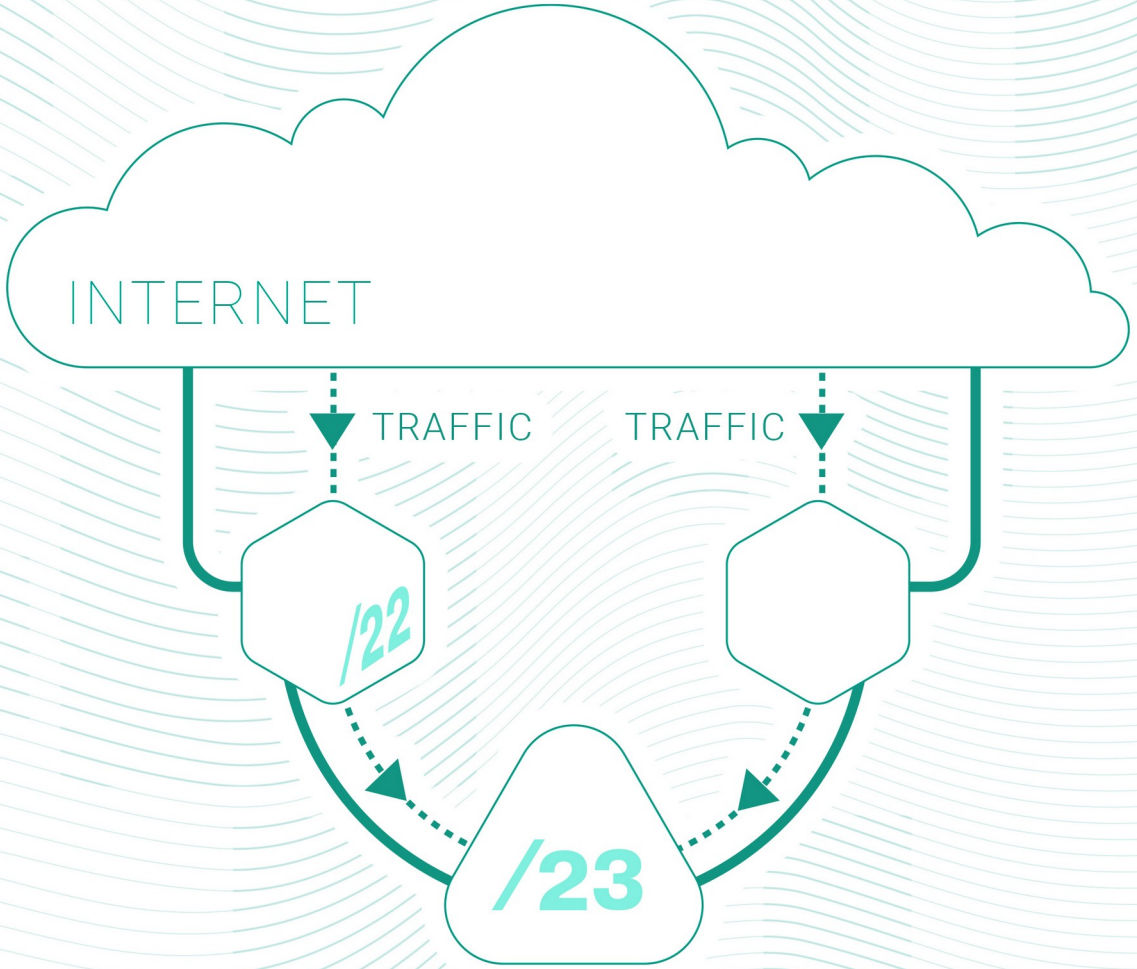

QRATOR LABS

Your space – your rules?

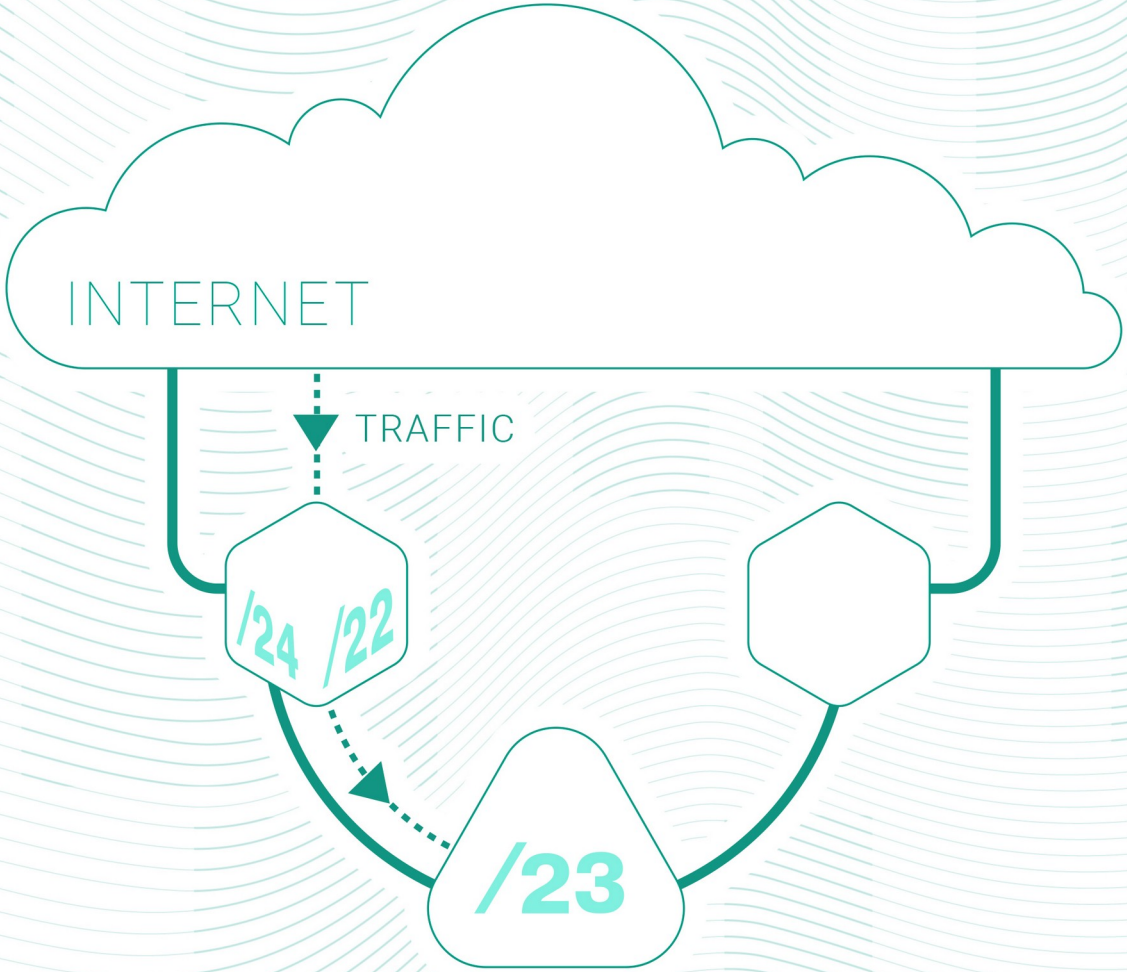

A sub-allocation state

- ISP had a /22
- He gave the /23 from it to the multihomed client
- And announced the /24
  - To became the only transit for this client

**MaxLength also doesn't resolve the problem in general*

My questions

- Are there any bad cases of «Drop Invalid» policy?
- Which maxLength is worth using now (min or max)?
  - What to do in Route objects case?
- Are we ready for RIR policy for validation with IRR?
  - Maybe it's time to return leading role to prefix owners?
  - Should we adopt ROA “Pair Blacklist” approach for Route objects?

- **To which RIPE/IETF WGs this questions should be brought to continue a discussion?**

QRATOR LABS

Your questions?

Contacts: eb@qrator.net


SWAT
Government Agent: We got a guy out here asking a lot of questions.