

Risks & Dangers in BGP

What can we do about it?

Job Snijders

NTT Ltd

Global IP Network (AS 2914)

job@ntt.net

Agenda

- Traffic Engineering challenges in BGP
- “*bGp OtImlsEr*” Misery Deconstructed
 - Post-leak data-plane wasteland
- Mitigating risks
 - NO_EXPORT
 - Direct Peering/interconnection & Origin Validation
 - RPKI ROA MaxLength Considerations
- Outlook on industry trends & projects

Traffic engineering options – “in” vs “out”

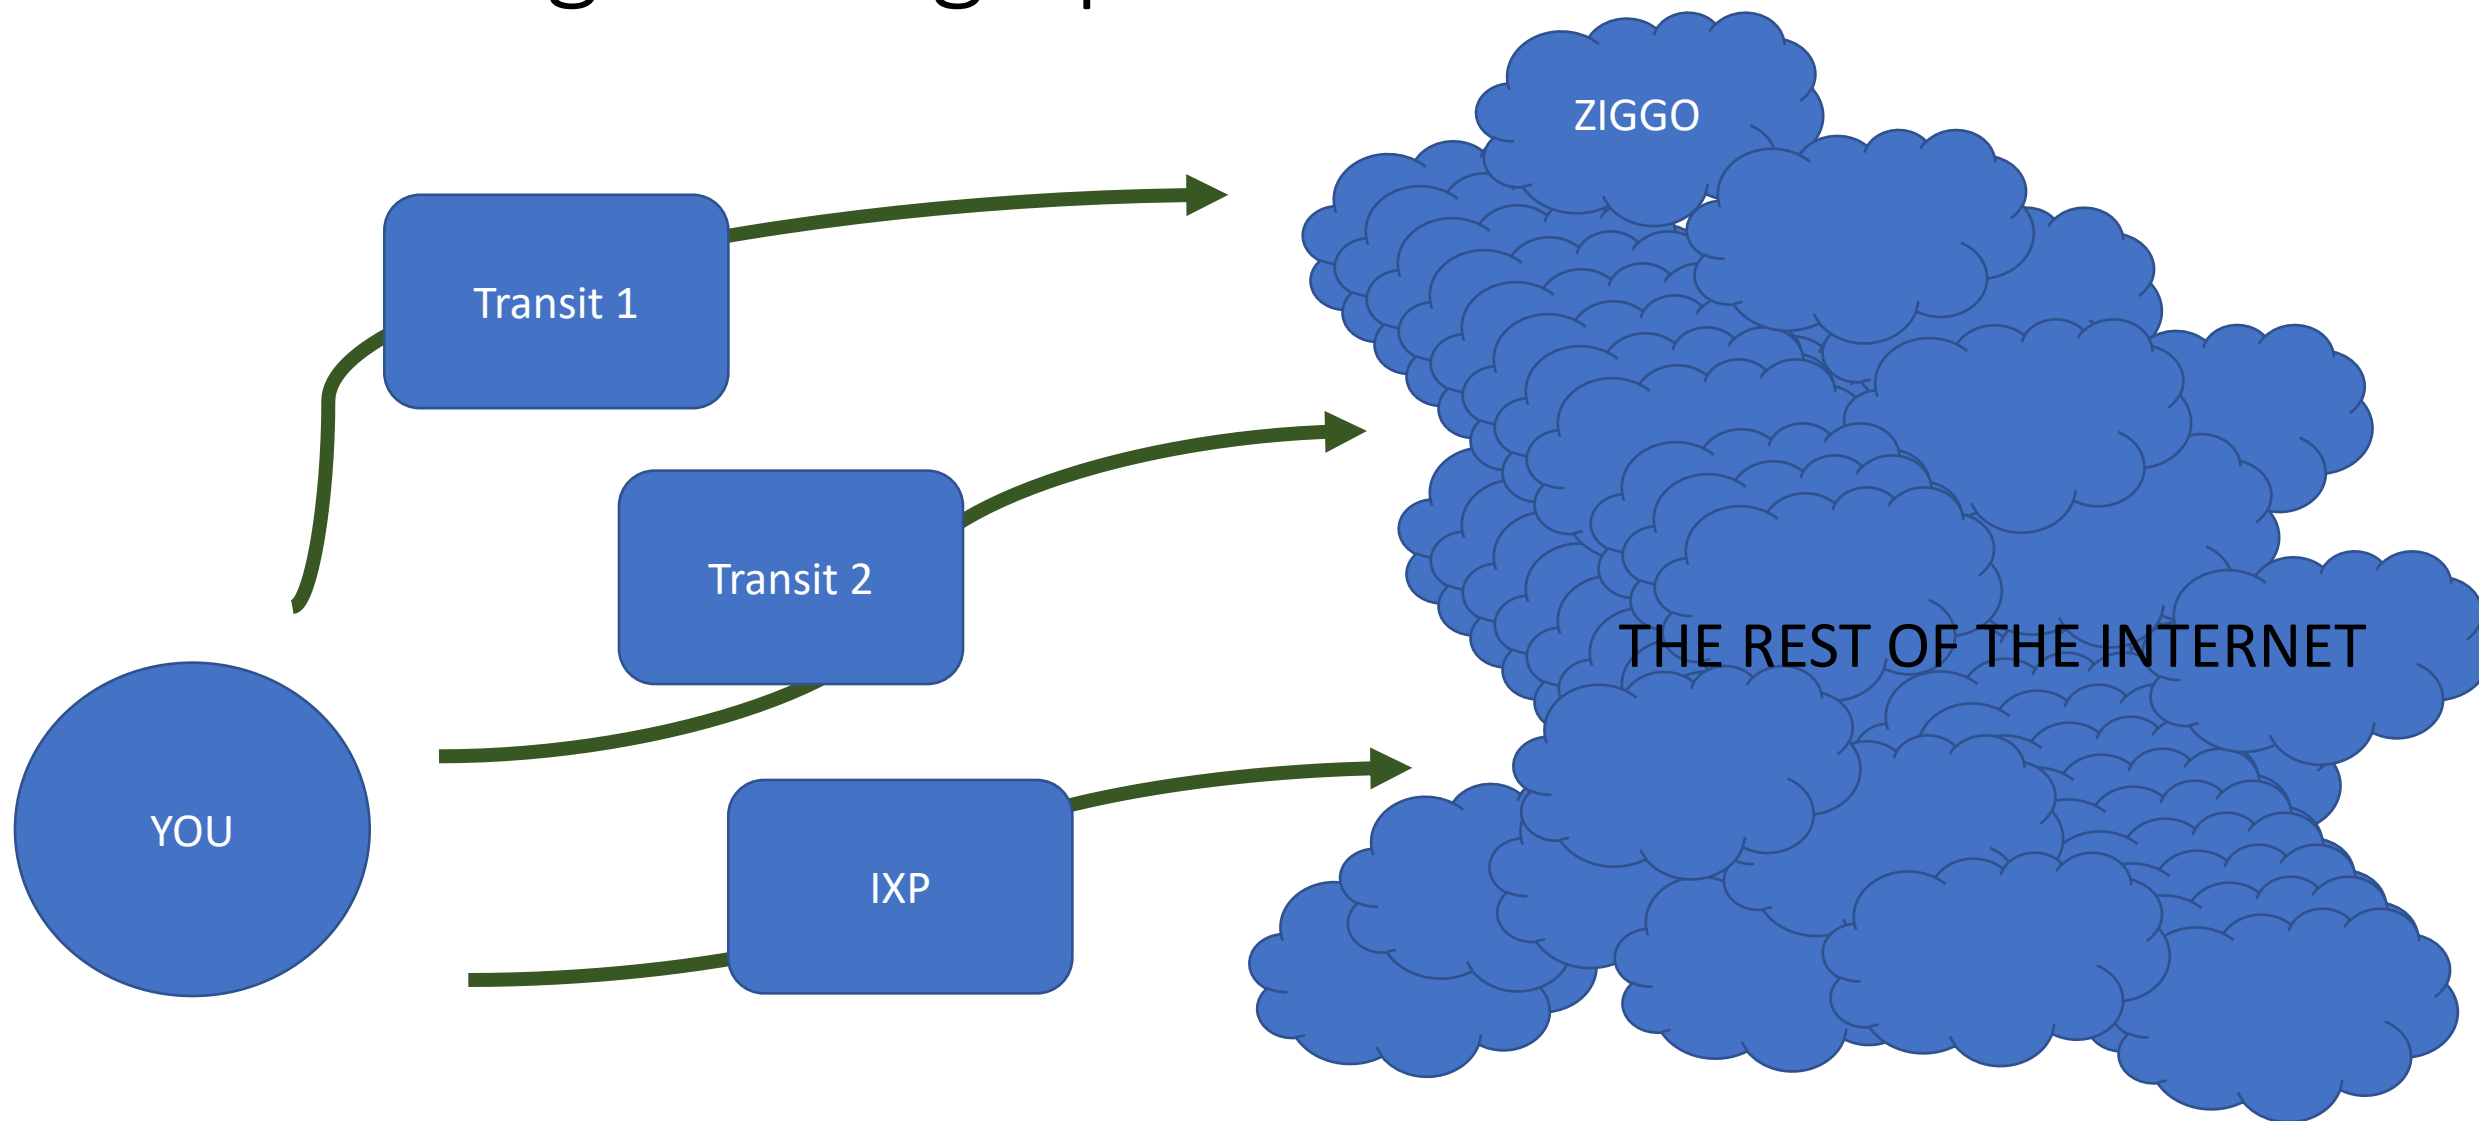
“**in**” means “**from the Internet, to you**”

“**out**” means “**from you, to the Internet**”

In general you can't really mandate where traffic comes in, but you do have a lot of control over where you send traffic out.


Some carriers use the words “push” and “pull” instead of “in” and “out”, always confirm the direction!

Traffic engineering options – “out”



Challenges (... in some places)


- Your company may not have a lot of choice of upstream providers
- Your company may have to navigate a fine line between multi-year traffic level commitments on circuits
- The organization doesn't have full time network engineers



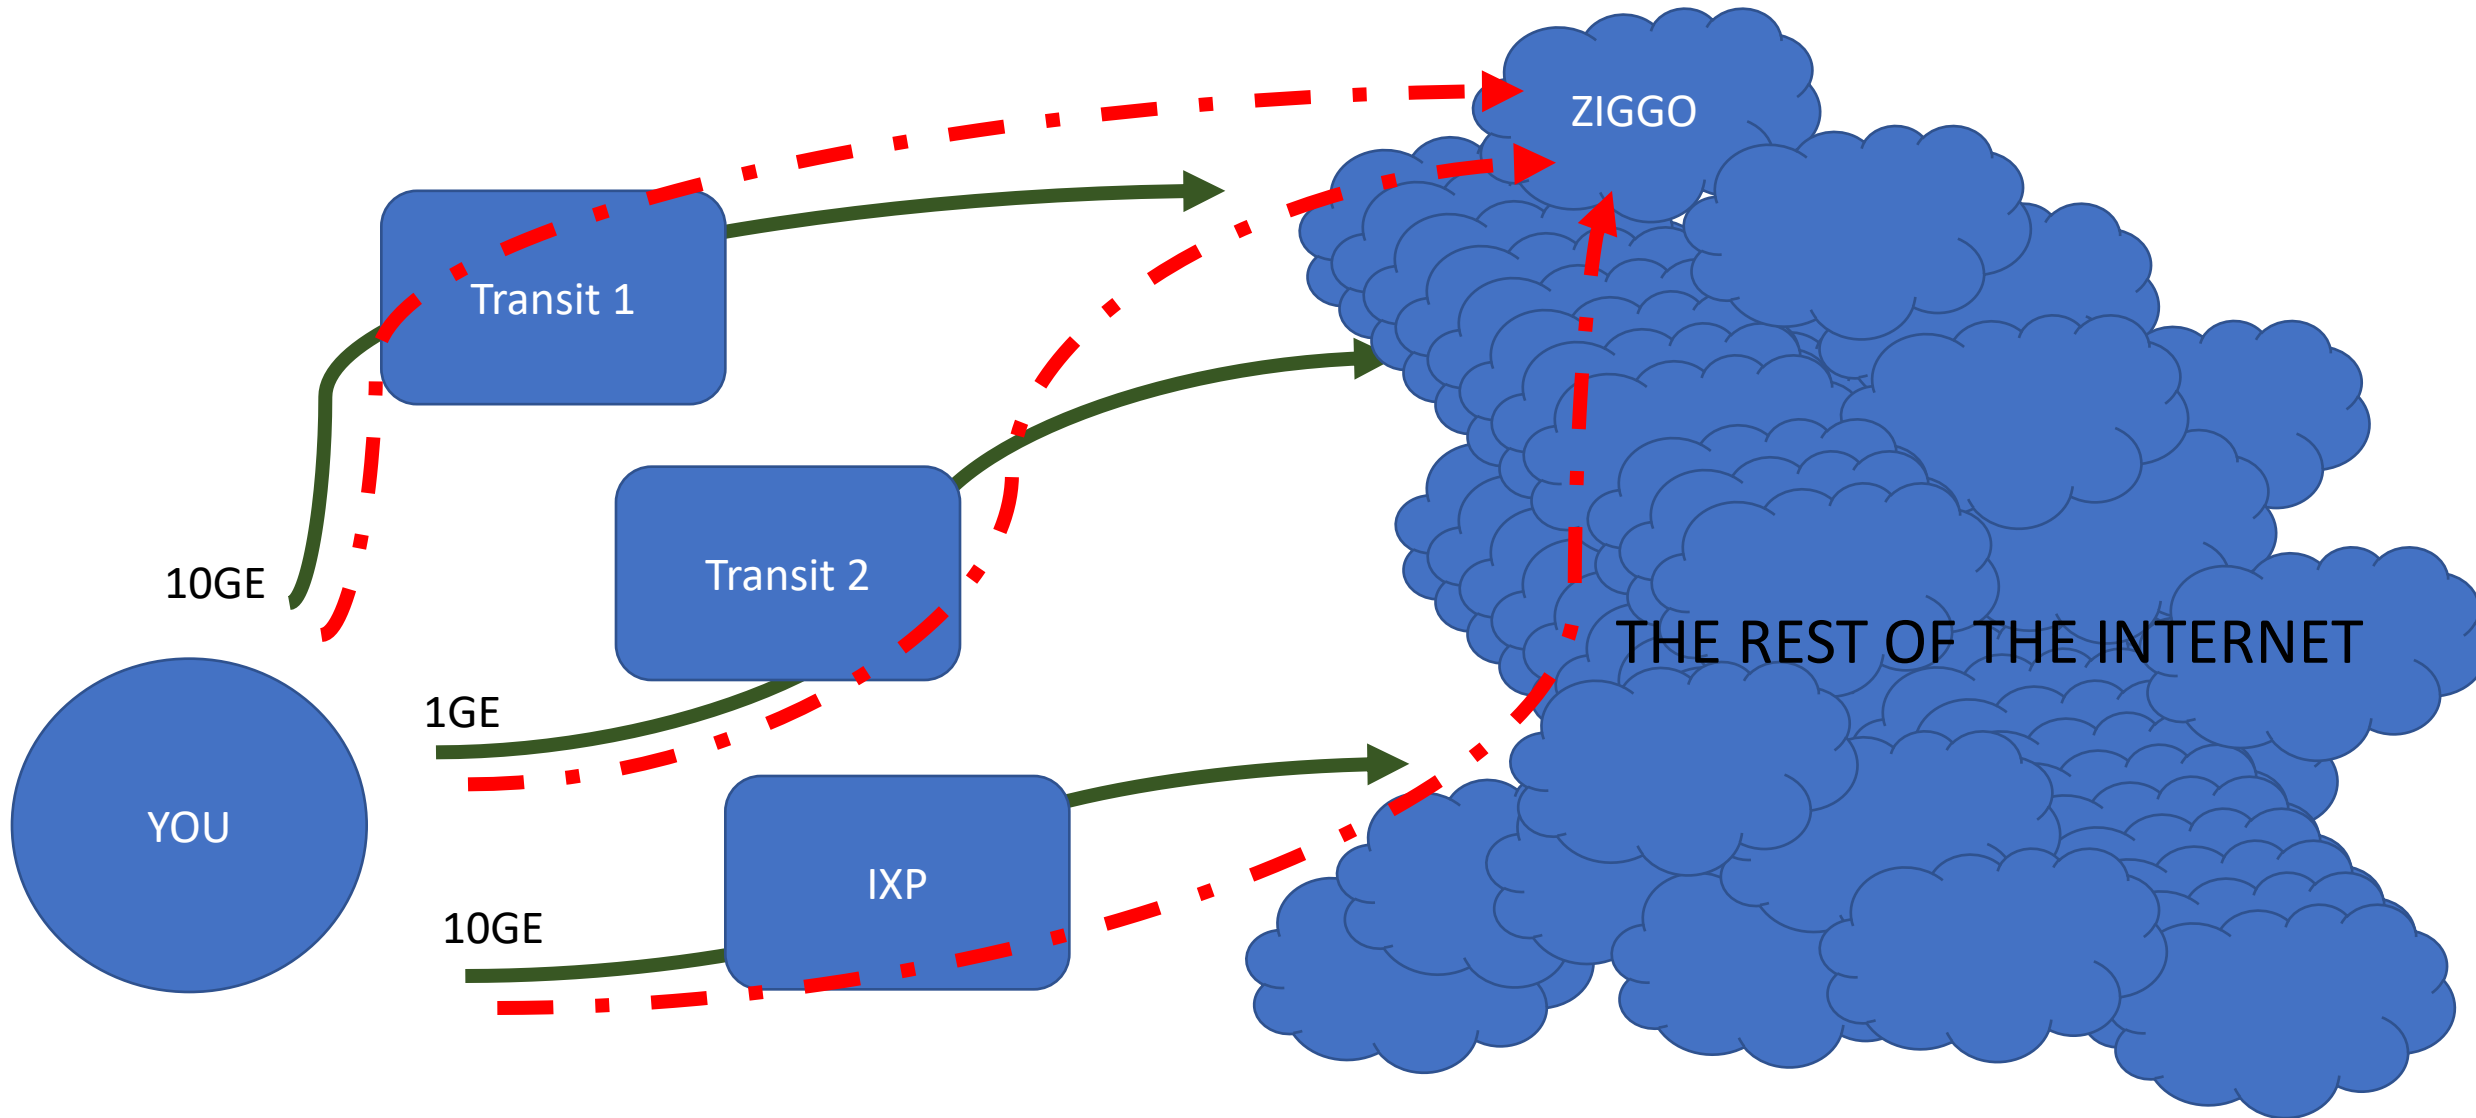
Challenges (... sometimes)

Results in software was sold to the company that:


- *Improves performance !111eleven11!!*
- *Reduces operational cost*
- *Adds intelligence*
- *Leverages existing connectivity*
- *Monitor 24/7*




Steering traffic “Out” -> “performance probing”



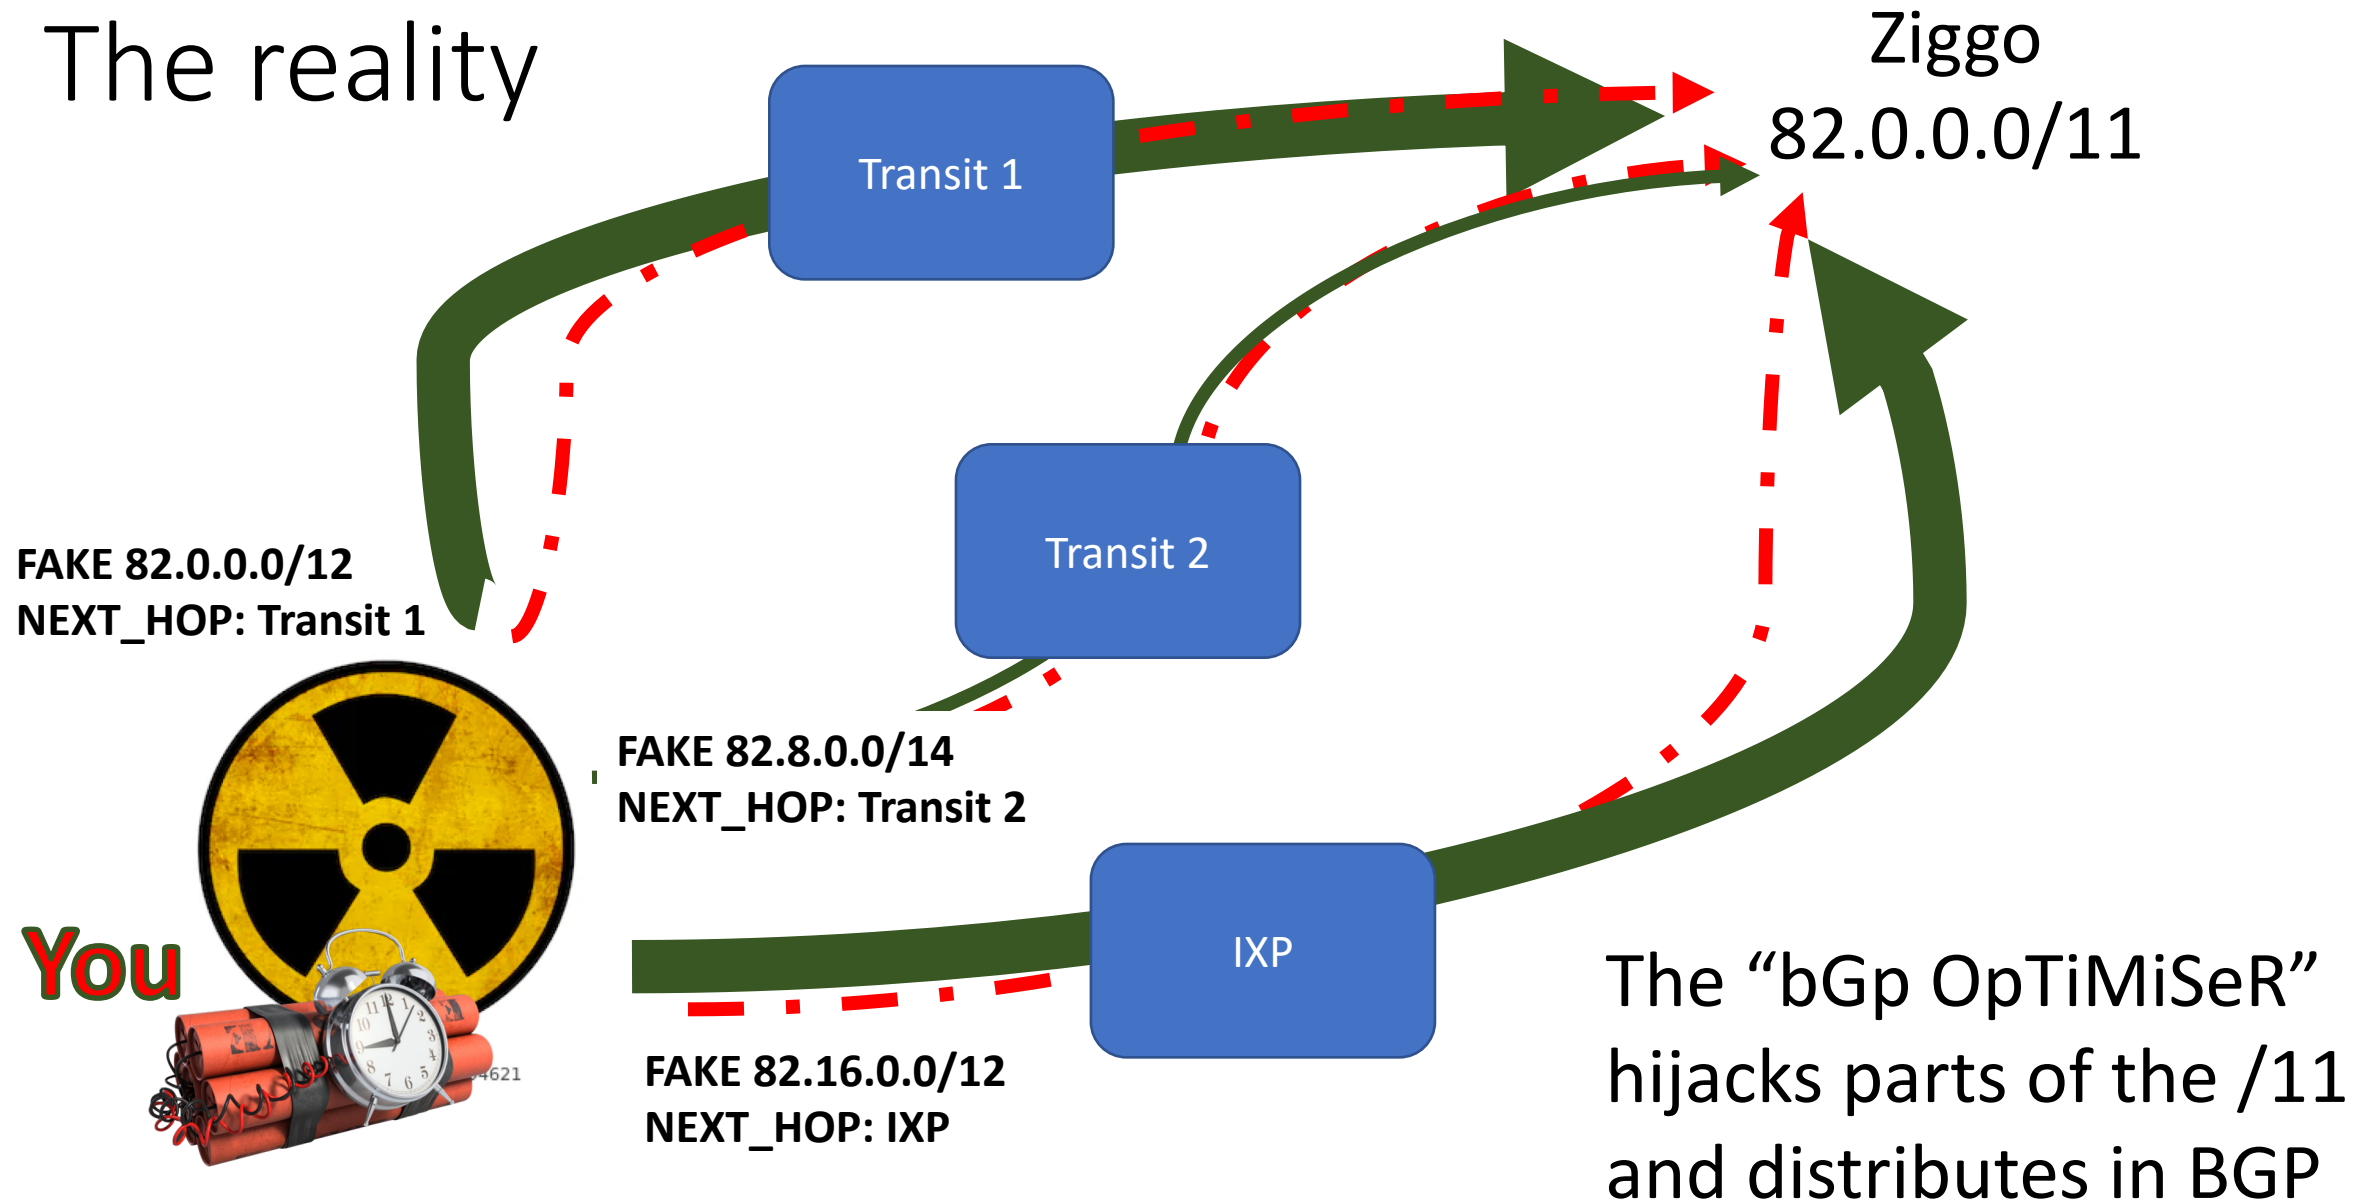
The promise: perfect balancing



The reality



The reality



When those lies leak out....



10GE

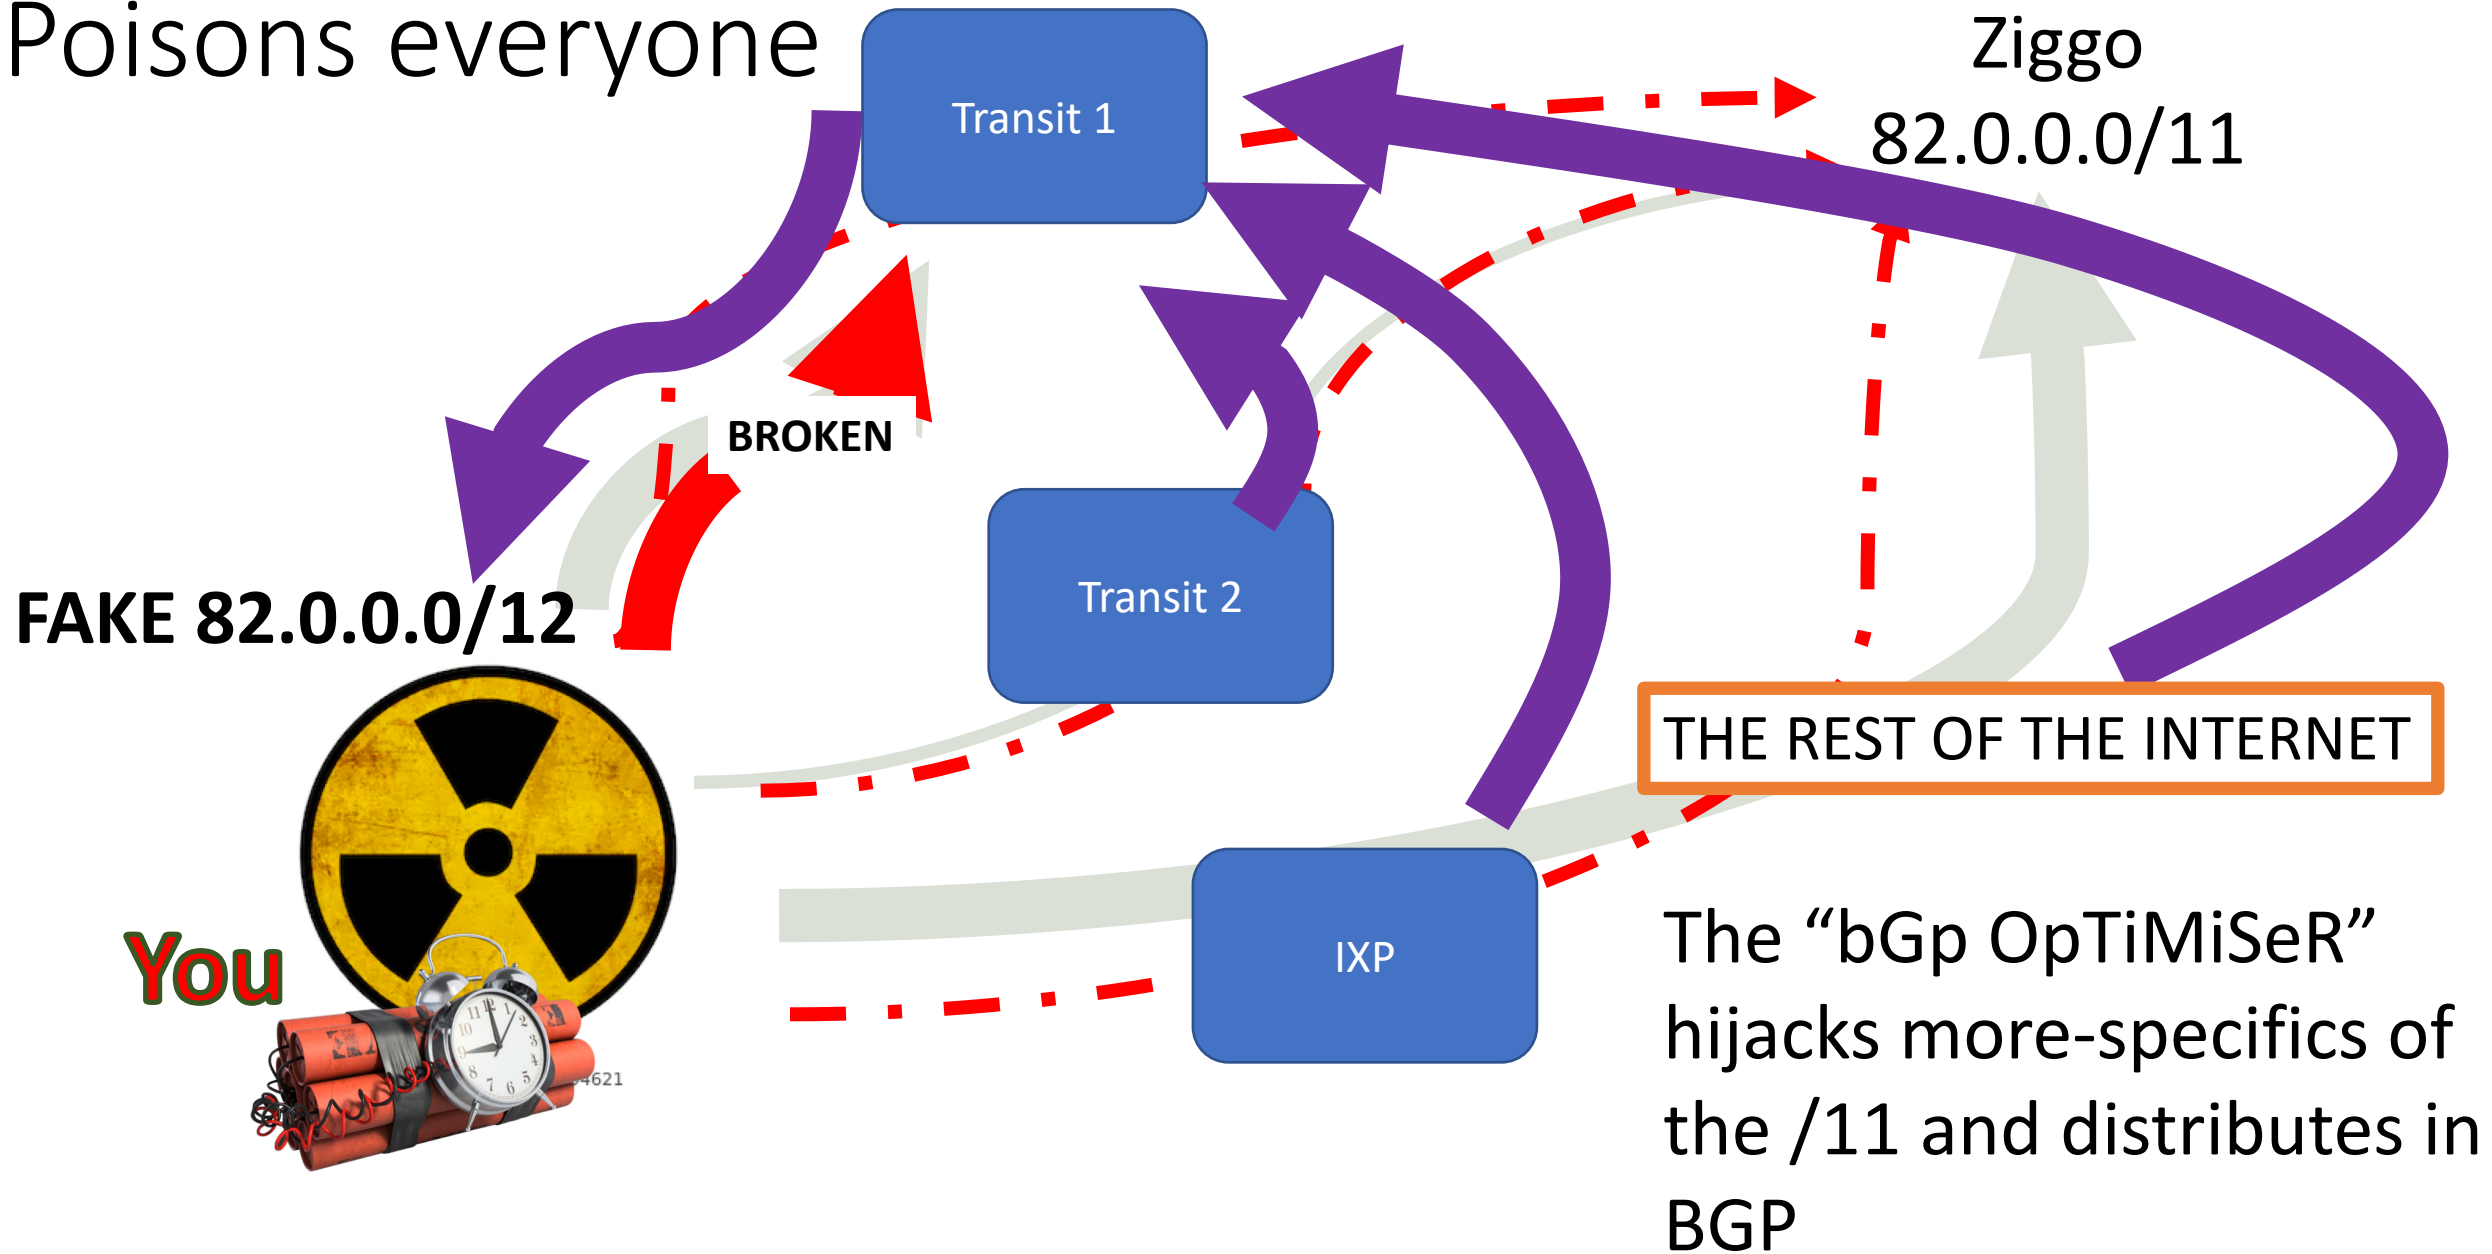
You



4621



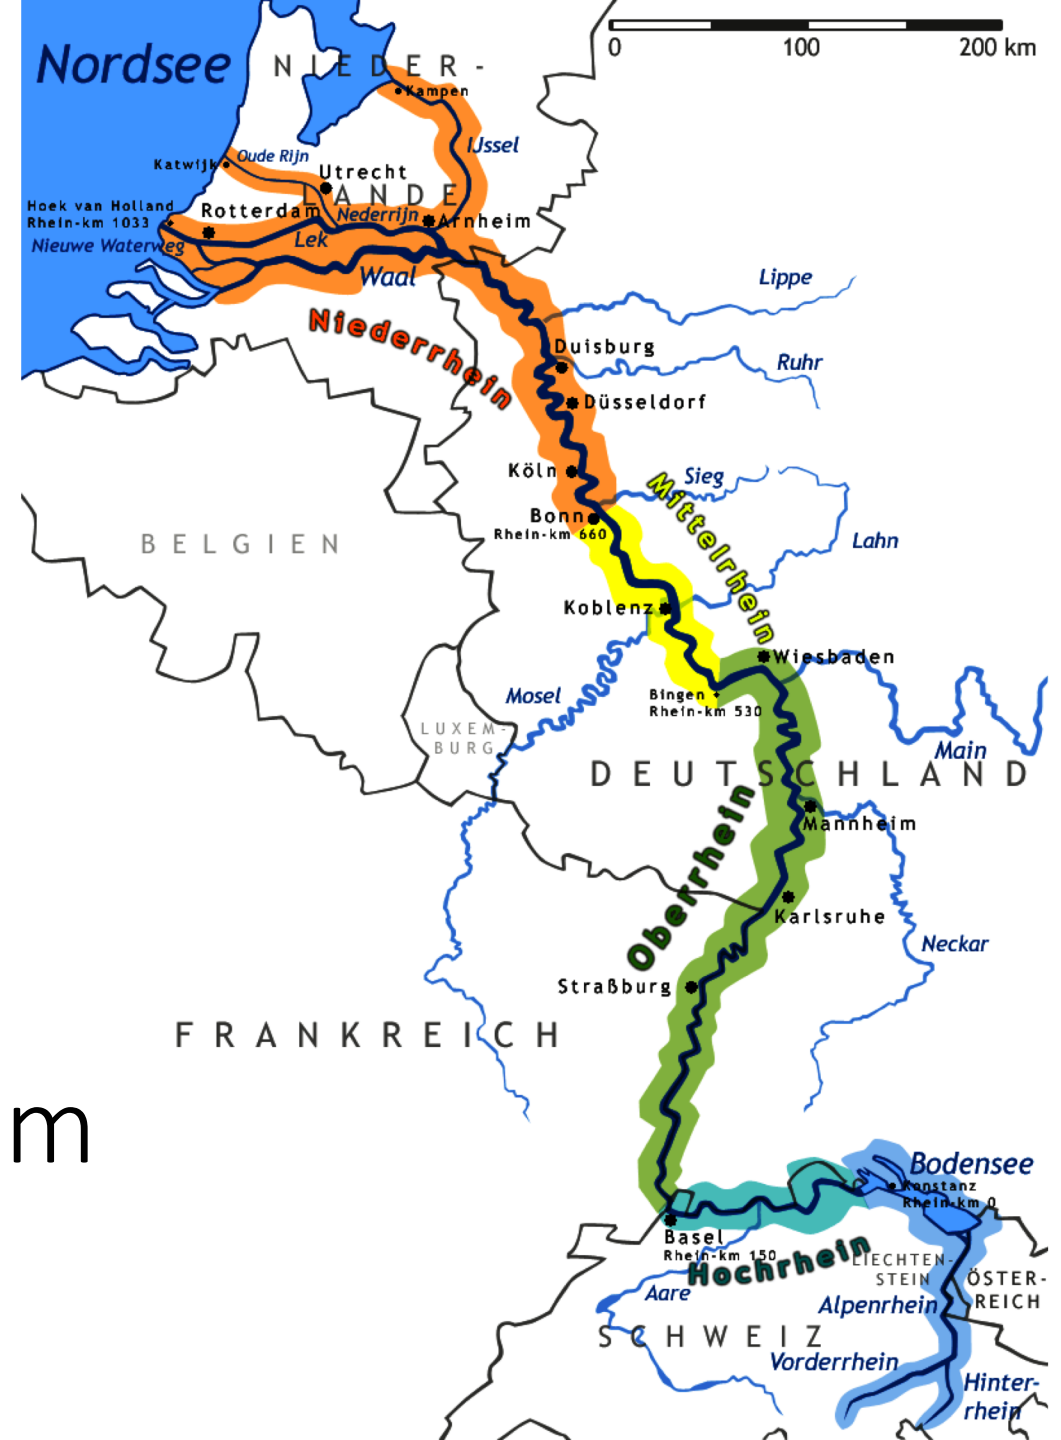
Poisons everyone



If the Default-Free Zone
is a shared resource...

bGp OtiMiSeRs are cause
for “water conflicts”

Problems upstream
cause problems downstream






Why do these things happen?

- Not all BGP Operators use NO_EXPORT
- Not all BGP implementations are bug free
- There is no “traffic-engineering” AFI in BGP. All routes are conflated in same IPv4/IPv6 unicast family
- In anticompetitive markets, the cost of Internet access may force people to be creative
- False sense of security *“If we don’t monitor and optimize, who will?!”*


Why does this happen?

twitter.com/PowerDNS_Bert/status/1143252745257979905

 **Tweet**





 **Bert Hubert** 
@PowerDNS_Bert

"Noction have been facilitating Internet outages for years and years and the best thing they can say in response is that it is technically possible to use their product responsibly, they just don't ship it that way." - anonymous specialist -> blog.cloudflare.com/how-verizon-an...

 **Noction** @noction · Jun 24
Replying to @tompaseka @beufanet and 3 others
We do have no export community support and have done for many years. The use of more specifics is also optional. Neither replaces the need for filters.

10:21 PM · Jun 24, 2019 · [Twitter Web Client](#)

26 Retweets 40 Likes

USE NO_EXPORT
DARN IT! F !!!!!
SEE ME AFTER CLASS

What can you do against this? R * P * K * I

1. Help others protect you:

-> **create RPKI ROAs for your space**

2. Protect yourself and others:

-> **Deploy RPKI based BGP Origin Validation**

(make sure you set the MaxLength attribute correctly!)

Pro-tip: Peer directly with each other, because...

When an attacker spoofs the origin & same prefix length ...
.... their AS_PATH is longer

When an attacker originates more-specifics...
..... you reject the route based on ROA

(Direct peering is sort of "AS_PATH Validation ", for just 1 AS Hop)

Outlook on industry trends

- ARIN TAL
- Pmacct / Kentik business intelligence
- IRRd 4 project
- RIPE 2018-06 Policy Proposal
- BGPalerter
- IXP Manager / Arouteserver

ARIN TAL Politics

- No real news
- Relaying party agreement is still the same




Search

Amsterda...

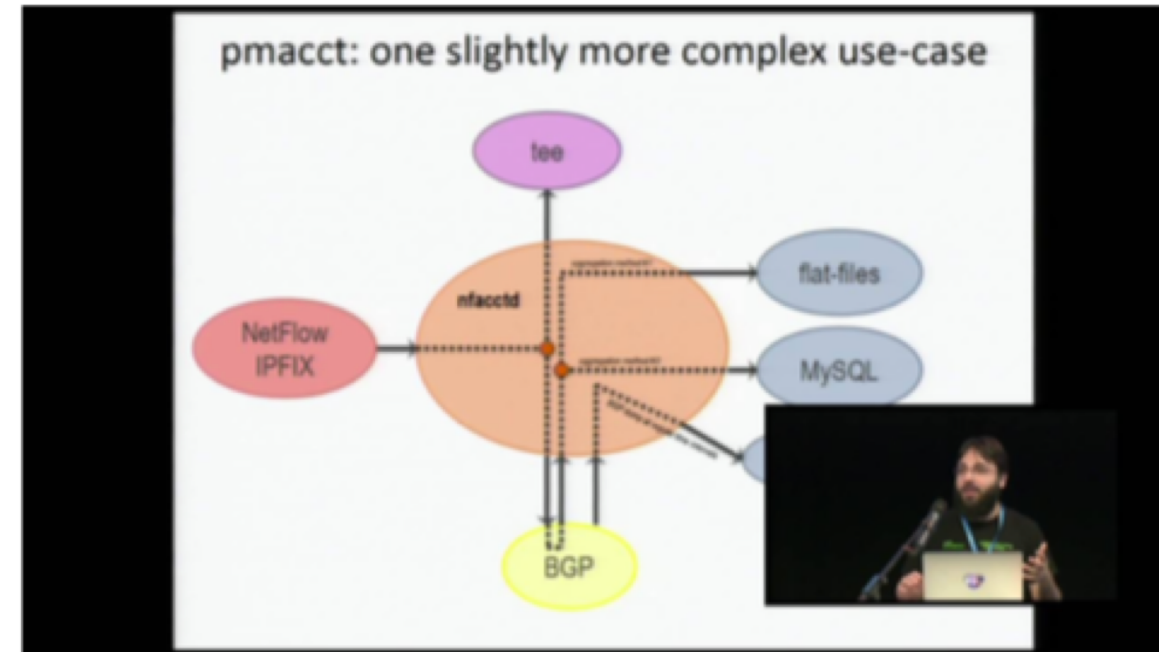
RPKI based traffic analysis with pmacct



14



15



14

A view from AS 2914 / NTT's global backbone

Free!

<http://pmacct.net/>

Free!

<http://pmacct.net/>

Free!

<http://pmacct.net/>

Free!

<http://pmacct.net/>

Free!



Cloudflare were about to announce their involvement in the RPKI community initiative last year.

Thanks also to **Job Snijders** ([twitter.com/jobSnijders](#)) and **Paulo Lucente** from **NTT** and the **PMACCT project**, who kindly pointed us to the PMACCT code to guide us through the implementation.

Lastly, to **Aaron Weintraub** from **Cogent** for his time in reviewing and testing our early, iterative versions for accuracy and usefulness. Thanks!

IRRD version 4 released!

NTT funded initiative
developed by Dashcare

<https://github.com/irrdnet/irrd4>

<https://irrd4.readthedocs.io>

Now in production at rr.ntt.net !!



dashcare



Watch last year's update: <https://www.youtube.com/watch?v=3BAwBClazzWc>

Slides

12 OK - no current status

13 Using RPKI to clean up the RIR

14 The "RR" system access

15 Improving security at the "aggregators"

16 Proposal: Let RPKI "own out" conflicting RIR

17 RPKI filter at the IRR instances

18 How are RPKI and RIR different?

19 RPKI supporting conflicting RIR advantages

20 OK - no current status

RPKI filter at the IRRd instances

```
graph LR; subgraph Data_sources [Data sources]; RIPE_IRR[RIPE IRR]; NTTCOM[NTTCOM]; RADB[RADB]; APNIC[APNIC]; Dots[...]; end; subgraph Aggregators; whois_radb_net([whois.radb.net]); rr_ntt_net([rr.ntt.net]); end; subgraph Clients; bgpq3[bgpq3]; end; Data_sources --> Filter(( )); Filter --> Aggregators; Aggregators --> Clients;
```

Slide 17 of 38

Find

Find All Match Case

Title and Content

1.20 / 7.90 0.00 x 0.00

English (USA)

80%

The screenshot shows a web browser window with the URL <https://www.ripe.net/participate/policies/proposals/2018-06>. The page is titled "RIPE NCC IRR Database Non-Authoritative Route Object Clean-up". The navigation bar includes links for "Manage IPs and ASNs", "Analyse", "Participate" (which is highlighted), "Get Support", "Publications", and "About Us". The breadcrumb trail reads: "You are here: Home > Participate > Policy Development > Policy Proposals > RIPE NCC IRR Database Non-Authoritative Route Object Clean-up". On the left sidebar, under "Participate", the "Policy Development" section is expanded, showing links to "Current Policy Proposals", "Archived Policy Proposals", "PDP Glossary", "Policy Proposal Template", "Policy Implementation Status", and "RIPE Document: Policy". The main content area has two tabs: "Policy Proposal" (selected) and "Draft Document". The title "RIPE NCC IRR Database Non-Authoritative Route Object Clean-up" is prominently displayed. Below the title, it states "First version: 11 Oct 2018 - This version: 22 Jul 2019". The "Summary of Proposal" section explains the goal: to change the business rules for the RIPE Internet Routing Registry (IRR) service to allow for a better hierarchy and to document this, giving resource holders the ability to modify their intention regarding which AS will announce their IP resources. It also mentions that using this update, various tools will produce a much cleaner filter set with a high level of trust in the actual resource holder's intent on the routing of their resources. On the right side, a sidebar for "2018-06" provides details: "Authors: Erik Bais, A2B Internet; Job Snijders, NTT; Martin Levy, Cloudflare"; "Proposal Version: 3.0 — 22 July 2019"; "All Versions: 1.0 — 11 October 2018, 2.0 — 17 May 2019, 3.0 — 22 July 2019"; "Current Phase Started: 14 August 2019"; and "Current Phase Ends: 12 September 2019". At the bottom, the "Working Group" is identified as the "Routing Working Group".

RIPE NCC IRR Database Non-A

ripe.net/participate/policies/proposals/2018-06

RIPE NETWORK COORDINATION CENTRE

Search IP Address or ASN

Manage IPs and ASNs > Analyse > **Participate** > Get Support > Publications > About Us >

You are here: Home > Participate > Policy Development > Policy Proposals > RIPE NCC IRR Database Non-Authoritative Route Object Clean-up

Participate <<

RIPE NCC Membership >

RIPE Community >

Meetings and Events >

Policy Development v

Current Policy Proposals

Archived Policy Proposals

PDP Glossary

Policy Proposal Template

Policy Implementation Status

RIPE Document: Policy

Policy Proposal Draft Document

RIPE NCC IRR Database Non-Authoritative Route Object Clean-up

First version: 11 Oct 2018 - This version: 22 Jul 2019

Summary of Proposal

The goal of this proposal is to change the business rules for the RIPE Internet Routing Registry (IRR) service to allow for a better hierarchy and to document this. This will give the resource holder the ability to modify their intention regarding which AS will announce their IP resources.

Using this update, the various tools will produce a much cleaner filter set with a high level of trust in the actual resource holder's intent on the routing of their resources.

2018-06

Authors:
Erik Bais, A2B Internet
Job Snijders, NTT
Martin Levy, Cloudflare

Proposal Version:
3.0 — 22 July 2019

All Versions:
1.0 — 11 October 2018
2.0 — 17 May 2019
3.0 — 22 July 2019

Current Phase Started:
14 August 2019

Current Phase Ends:
12 September 2019

Participation in the RIPE Policy Development

Working Group:
Routing Working Group

<https://www.ripe.net/participate/policies/proposals/2018-06>



When your favorite startup is a



honestnetworker.net/2015/06/28/when-your-favorite-s...



Incognito



Jun
28

When your favorite startup is acquired by Cisco

Posted by [ohseuch4aeji4xar](#)



00:03



00:15



nttgin/BGPalerter: Software to monitor streams of BGP data. Pre-configured for real-time visibility loss and hijacks detection.

Manage topics

116 commits 1 branch 0 releases 6 contributors

Branch: master New pull request Create new file Upload files Find File

massimocandela raise error in case of duplicated keys in monitoring rules Latest commit

bin new dist 19.9.5.2

connectors introduced support for AS_SET (which it will be introduced in next RI...

inputs raise error in case of duplicated keys in monitoring rules

mor

<https://github.com/nttgin/BGPalerter>

BGPalerter in a nutshell

- Uses data from RIPE RIS collectors (data inputs are “connectors”)
- Then push data through “monitors” for evaluation and alert event detection
- Finally Spit it out via “reports” (slack, email, IRC, etc..)

Use BGPalerter to monitor for hijacks & outages

<https://github.com/nttgin/BGPalerter>

Free
As fast as possible
Free
As fast as possible
Free
As fast as possible
Free

Start bgpalerter

RIPE RIS LIVE MAGIC



```
BGPalerter, version: 19.8.30.2 environment: production
Loaded config: /home/rinse/bgpalerter/config.yml
Subscribing to: 2001:9e0::/32
Subscribing to: 91.212.184.0/24
Subscribing to: 185.117.108.0/22
Subscribing to: 185.83.32.0/22
Subscribing to: 217.26.96.0/20
Subscribing to: 5.199.144.0/20
Subscribing to: 81.24.96.0/20
Subscribing to: 212.84.128.0/19
Subscribing to: 213.134.224.0/19
Subscribing to: 212.45.32.0/19
Subscribing to: 213.233.192.0/18
Subscribing to: 83.247.0.0/17
Subscribing to: 87.195.0.0/16
```

Observe DDOS mitigation in progress

```
rinse@mml:~/bgpalerter$ tail -f logs/reports-2019-09-02.log
2019-09-02T14:15:16.534Z [production] verbose: Possible change of configuration. A new prefix 213.134.227.0/24 is announced by
AS12414. It is a more specific of 213.134.224.0/19 (Solcon Internetdiensten B.V.)
2019-09-02T14:46:12.292Z [production] verbose: Possible change of configuration. A new prefix 213.134.227.0/24 is announced by
AS12414. It is a more specific of 213.134.224.0/19 (Solcon Internetdiensten B.V.).
```

Example contributed by Rinse Kloek

BGPalerter can email (of course)

```
—job@ntt.net—[ Folder: =INBOX ]-bgpalerter—[*148268 messages (617 new) ]—
Date: Wed, 14 Aug 2019 20:59:18 +0200
From: bgpalerter@bgpalerter.meerval.net
To: job@instituut.net
Subject: BGP alert: newprefix

Possible change of configuration. A new prefix 209.24.8.0/21 is announced by
AS15562. It is a more specific of 209.24.0.0/16 (Job).

DETAILS:
-----
Monitored prefix:      209.24.0.0/16
Prefix Description:    Job
Usually announced by: AS15562
Event type:            prefix-detection
Detected new prefix:   209.24.8.0/21
Announced by:         AS15562
When event started:    2019-08-14 06:58:54 UTC
Last event:            2019-08-14 06:58:54 UTC
Detected by peers:     1
See in BGPlay:
https://stat.ripe.net/widget/bgplay#w.resource=209.24.0.0/16&w.ignoreReannouncements=true&w.starttime=1565808834&w.endtime=1565809134&w.rrcs=0,1,2,5,6,7,10,11,13,14,15,16,18,20&w.type=bgp
~
```

BGPalerter can send to Slack too... (AOL or MSN?)

o.slack.com/client/T03CKAZCU/CFW2RF9MJ/details/members



#bgpalerts

☆ | 👤 6 | 📌 0 | ✎ Add a topic



Massimo Candela 2:51 PM

(I added coverage in the tests also)



arjen 2:55 PM

I disconnected my AS yesterday

I'll hook it up again tonight and see what happens, thanks! 😊 (edited)



incoming-webhook APP 4:19 PM

Possible change of configuration. A new prefix 203.105.76.0/23 is announced by AS3949. It is a more specific of 203.105.76.0/23 (subas 64600).


A new prefix 120.29.148.0/24 is announced by AS30103. It should be instead 120.29.144.0/21 (subas 64600)

A new prefix 117.104.145.0/24 is announced by AS394265. It should be instead 117.104.128.0/19 (subas 64600) announced by AS2914

Possible change of configuration. A new prefix 180.12.186.0/24 is announced by AS4713. It is a more specific of 180.0.0.0/10 (OCN).

Possible change of configuration. A new prefix 180.12.181.0/24 is announced by AS4713. It is a more specific of 180.0.0.0/10 (OCN).

Today





⚠ Not Secure | printablechecklist.org

Printable Checklist

For Internet Exchanges



RPKI support in Arouteserver




RPKI support in IXP Manager




RPKI support in BIRD



RPKI support in OpenBGPD



IXP management software



Route Server softwares

+ Add item

<<< time for questions! >>>

Also feel free to ask questions later today,
or via cyber

job@ntt.net

<https://twitter.com/jobsnijders>

**RELYING
ON AS_PATH
LOOP DETECTION**

**OUTBOUND
EBGP
PREFIX FILTERS**

**INBOUND EBGP
PREFIX FILTERS**

**RPKI ORIGIN
VALIDATION**

