


We Love Route Leaks

?

Alexander Azimov, Yandex
Eugene Bogomazov, Qrator

Classic Route Leak



The real route goes wrong

Disrupts cash flow

Types 1-4 RFC 7908

Classic Route Leak Effects

- Traffic moving in the wrong way results in:
 - Delays
 - Packet loss
 - Eavesdropping/Sniffing
- Leaker overload
 - Drop in traffic quality

Are They Often?

- ~4 500 unique leakers during this year
- ~10 global leaks

Uniq Leakers/Quarter			Uniq Leakers/Month		
Year	Quarter	Uniq Leakers	Year	Month	Uniq Leakers
2022	3	2197	2022	8	1265
2022	2	2914	2022	7	1924
2022	1	3235	2022	6	1949
2021	4	3180	2022	5	1832
2021	3	3031	2022	4	1885
2021	2	2998	2022	3	2249
2021	1	2982	2022	2	2024
			2022	1	1938

Are they often?

Nearly every ASN and prefix was affected by a small leak

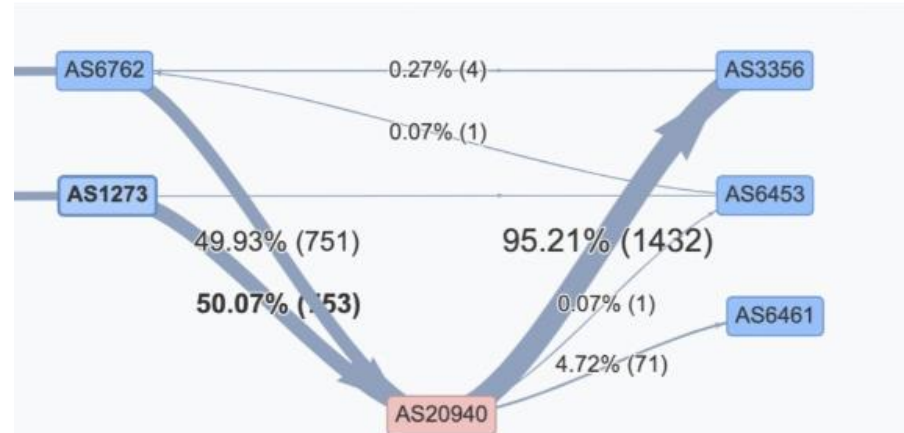
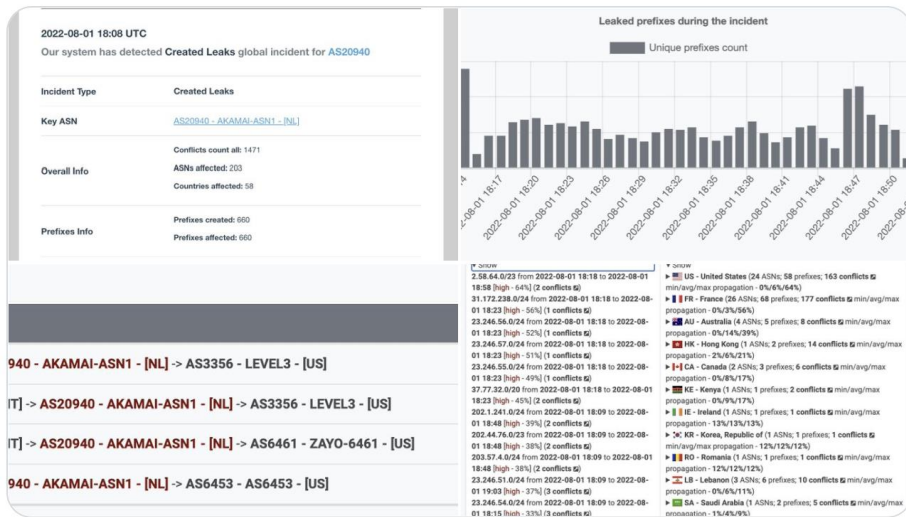
>10% of them were affected by a big one at least once

This Month Example



@Qrator_Radar

August 1, 2022 — AS20940 — AKAMAI-ASN1 [NL] — leaked 660 prefixes creating 1471 conflicts with 203 ASNs in 58 countries. Maximum propagation: 64%. Duration: 3 hours 9 minute.

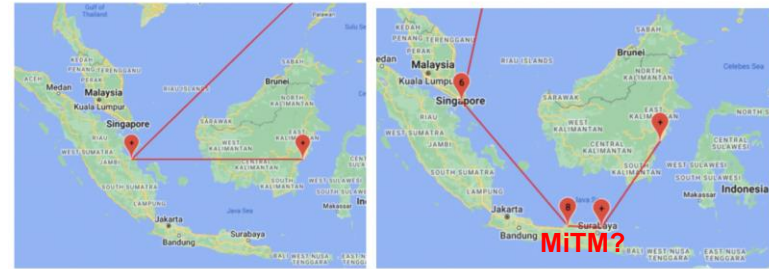


The Route Leak Consequence Example

Unexpected propagation of leaked prefixes



Changing of a traffic trace




Normal trace


Leaked trace

Unexpected increase in traffic volume and possible DoS (routers may be overloaded)

Traffic spike from 0 to Gigabits!



Changing of latency (RTT)



It's a Route Leak!

:

1. Find a problem
2. Find a responsible party
3. Find their abuse email contact
4. Write a complaint
5. Wait
6. ...
7. Wait
8. Profit! (or not)

Bully the Leaker

Before:	After:
ASZ - LeakerAS - ASY - ASX - YourAS	ASY - ASX - YourAS - LeakerAS - Your AS

How does it work?

BGP Loop prevention mechanism

Your AS at the end for ROA check

Your AS in the middle for neighbor check

Prefix Deaggregation

If you are a big guy:

- Directly connect to the most significant region ISPs
- Create ROAs with sub-prefix ability
- If your prefix in the leak:
 - Directly announce sub-prefix to affected parties
- You are amazing, you return a big amount of traffic back

Why Do We Love Route Leaks?

- We can configure devices to prevent them;
- We can write new monitoring tools;
- We can create action plans to fight them (of course, with drills);
- And there is always data to present at NOG meetings 😊

Route Leak Prevention and Detection Using Roles in UPDATE and OPEN Messages RFC 9234

Status

[IESG evaluation record](#)

[IESG writeups](#)


[Email expansions](#)

[History](#)

Versions:

[00](#) [01](#) [02](#) [03](#) [04](#) [05](#) [06](#) [07](#) [08](#) [09](#) [10](#) [11](#) [12](#) [13](#) [14](#) [15](#) [16](#) [17](#) [18](#) [19](#) [20](#) [21](#) [22](#) [23](#) [24](#)

draft-ymbk-idr-bgp-open-policy




draft-ietf-idr-bgp-open-policy




rfc9234


rfc9234



Route Leak Prevention: Communities



Route Leak Prevention: Communities



One mistake from failure

One Role To Rule Them All

Role – a new configuration option that

- Automates leak prevention;
- Provides leak detection;
- Controls your neighbor's configuration.


BGP Roles Negotiation




Allowed roles:

- Provider - sender is a transit provider to neighbor;
- Customer - sender is transit customer of neighbor;
- RS - sender is a Route Server, usually at internet exchange point (IX);
- RS-Client - sender is client of RS;
- Peer - sender and neighbor are peers.


Only-To-Customer Attribute (OTC)



Route Leak Prevention & Detection: OTC



Route Leak Prevention & Detection: OTC



Double set, double check.

OTC Setting

Egress policy:

- If route is sent to customer, peer or RS-client and the OTC attribute is not set it **MUST** be added with value equal to AS number of the sender;

Ingress policy:

- If a route is received from a Provider, Peer or RS and the OTC attribute has not been set it **MUST** be added with value equal to AS number of the neighbor (sender).

OTC Checking

Egress policy (before egress marking):

- A route with the OTC attribute set **MUST NOT** be sent to providers, peers, or RS(s).

Ingress policy (before ingress marking):

- If a route with OTC attribute is received from Customer or RS-client - it's a route leak;
- If a route with OTC attribute is received from Peer and its value isn't equal to the neighbor's ASN - it's a route leak.

What Should We Do with Route Leaks?

The only acceptable mitigation policy – route leaks **MUST** be rejected.

This mitigation policy **SHOULD** be used.

Configuring Roles

BIRD

```
protocol bgp {
  local as 65001;
  neighbor 127.20.0.1 as 65000;
  multihop;
  source address 127.20.0.2;
  strict bind on;
  ipv4 {
    import all;
    export all;
  };
  local role customer;
}
```

FRR

```
router bgp 64502
  neighbor 172.16.200.101 remote-as 64501
  neighbor 172.16.200.101 ebgp-multihop
  neighbor 172.16.200.101 passive
  neighbor 172.16.200.101 local-role customer
```

In case of
error/misconfiguration

```
bird> show protocol
Name      Proto    Table    State    Since          Info
device1   Device   ---      up       13:40:00.329
bgp1      BGP      ---      start    13:40:04.884  Idle          BGP Error: Role mismatch
bgp2      BGP      ---      up       13:40:04.335  Established
bird>
```

Routes are automatically tagged with the OTC attribute

Only to Customer

```
BGP routing table entry for 192.0.2.0/24, version 1
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  64501
    172.16.200.101 from 172.16.200.101 (172.16.200.101)
      Origin IGP, metric 0, valid, external, otc 64501, best (First path received)
```


BGP Roles & OTC

You configure only BGP Roles, OTC configuration is done in code;

- BGP Roles are negotiated;
- OTC is set on both ingress and egress;
- OTC is checked on both ingress and egress;
- OTC is an attribute – it is unlikely to be stripped;
- Detecting route leaks even several hops away from the source.

Vendor Support

Solution	Status	Version
BIRD	+	Will appear in 2.0.11
FRR	+	Will appear in 8.4
OpenBGPD	+	7.5
Mikrotik	Reduced Functionality	Appeared even before RFC

If We Don't Really Love Route Leaks

- If you are using opensource tools for routing – set up roles!
- Send feature request to your favorite vendor;
- Contribute to opensource tools (BMP parsers, bgpdump, etc.);
- And make nice slides about your user experience at NOG meetings! 😊