


## How to prevent and detect BGP route leaks with RFC 9234 “BGP Roles”

My name is Alex Kozlov, but all the credit is due to the RFC’s authors:  
**A. Azimov** (Qrator Labs & Yandex), **E. Bogomazov** (Qrator Labs), **R. Bush** (IIJ & Arccus), **K. Patel** (Arccus) and **K. Sriram** (USA NIST)

**BGP Route Leak** is the redirection of traffic through an autonomous system that should not be on the route.

BGP Route Leaks could result in a redirection of traffic through an unintended path that may enable eavesdropping or traffic analysis, and may or may not result in an overload or complete drop (black hole) of the traffic. Route leaks can be accidental or malicious but most often arise from accidental misconfigurations.



## How frequent are Route Leaks?

Leakers

Year	Quarter	Uniq Leakers
2023	1	2605
2022	4	2775
2022	3	3030
2022	2	2914
2022	1	3235

Lots of BGP Route Leaks, smaller and larger, happens almost every day.

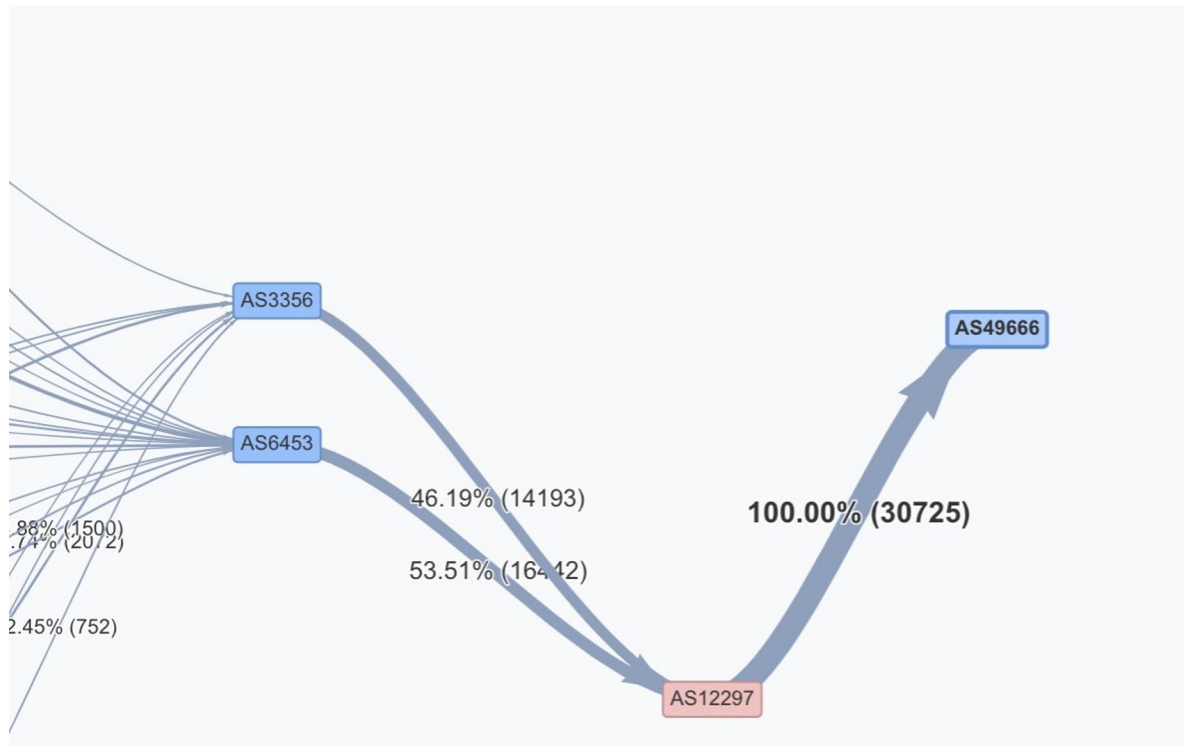
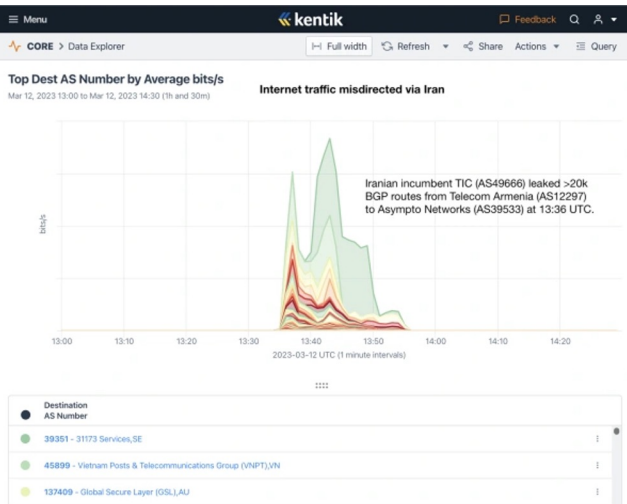
We estimate that nearly every announced ASN and prefix was affected by a small leak;

>10% of ASNs and prefixes were affected by a big one at least once.

Route Leaks

Year	Quarter	Total Leaks
2023	1	6565744
2022	4	3302804
2022	3	12103955
2022	2	3366094
2022	1	11695198

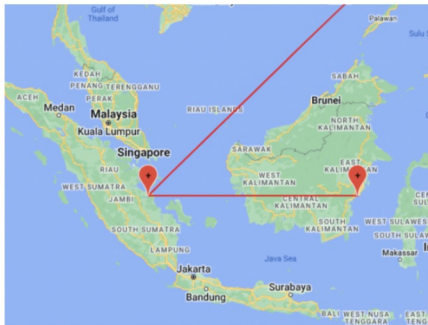

## The most recent example of a big BGP Route Leak (March 12)



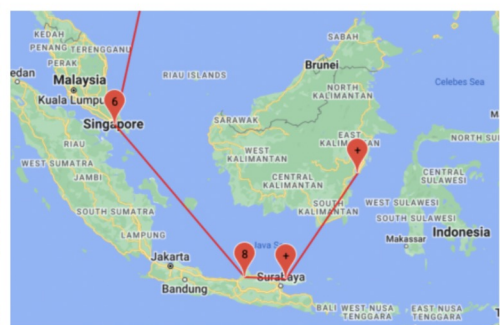
© Kentik

30725 prefixes affected,  
3183 ASNs in 142 countries.

Accepted Chart



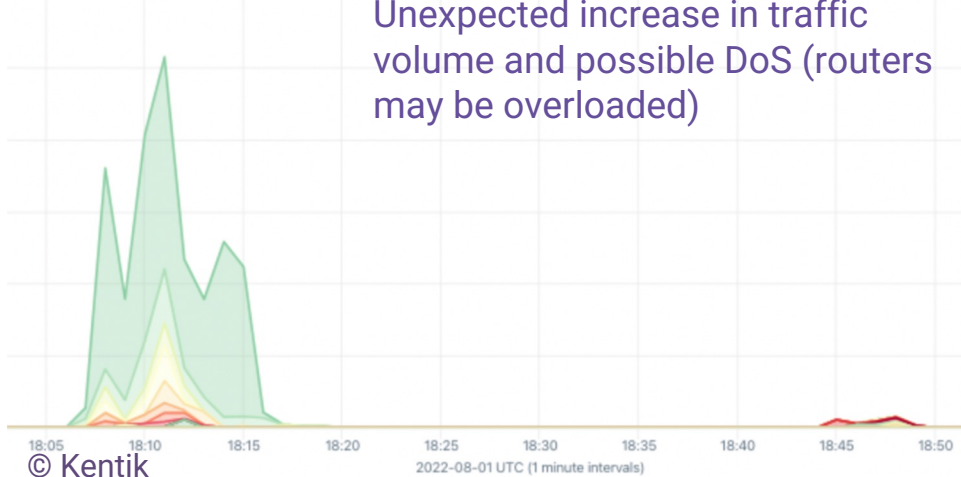
Normal trace



Leaked trace

### Change in latency (RTT)


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



```
Traceroute to 113.11.155.81 (113.11.155.81), 1 hops
 1 216.56.3.73 AS2381 9.123ms
 2 140.189.9.29 r-ummadison-isp-ae8.ip4.wiscnet 0.647ms
 3 140.189.9.77 r-222wwash-isp-ae2.ip4.wiscnet 4.034ms
 4 140.189.8.134 r-minneapolis-isp-ae7.ip4.wiscnet 6.087ms
 5 62.115.46.174 mini-b2-link.ip.twelve99.net 5.842ms
 6 62.115.143.225 omha-b1-link.ip.twelve99.net *
 7 *
 8 62.115.136.46 dls-b24-link.ip.twelve99.net *
 9 *
10 62.115.118.247 las-b22-link.ip.twelve99.net 327.119ms
11 213.248.76.163 telekomunikasi-svc074956-las 356.126ms
12 180.240.192.10 AS7713 214.981ms
13 *
14 36.89.254.161 AS7713 231.506ms
15 113.11.155.81 AS9326 230.03ms
```


```
Traceroute to 113.11.155.81 (113.11.155.81), 15 hops
 1 216.56.3.73 24.625ms
 2 140.189.9.29 0.647ms
 3 140.189.8.170 4.034ms
 4 140.189.8.125 6.087ms
 5 208.115.136.255 5.842ms
 6 103.14.246.174 214.147ms
 7 *
 8 103.146.188.130 327.119ms
 9 *
10 *
11 *
12 113.11.155.10 356.126ms
13 113.11.155.81 351.038ms
```

## Before BGP Roles



One mistake from failure

## After BGP Roles



Double set, double check.

Provider AS



Notification

Customer AS




Open(peer)

### 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)





The following ingress procedure applies to the processing of the OTC Attribute on route receipt:

1. If a route with the OTC Attribute is received from a Customer or an RS-Client, then it is a route leak and **MUST** be considered ineligible (see [Section 3](#)).
2. If a route with the OTC Attribute is received from a Peer (i.e., remote AS with a Peer Role) and the Attribute has a value that is not equal to the remote (i.e., Peer's) AS number, then it is a route leak and **MUST** be considered ineligible.
3. If a route is received from a Provider, a Peer, or an RS and the OTC Attribute is not present, then it **MUST** be added with a value equal to the AS number of the remote AS.

The following egress procedure applies to the processing of the OTC Attribute on route advertisement:

1. If a route is to be advertised to a Customer, a Peer, or an RS-Client (when the sender is an RS), and the OTC Attribute is not present, then when advertising the route, an OTC Attribute **MUST** be added with a value equal to the AS number of the local AS.
2. If a route already contains the OTC Attribute, it **MUST NOT** be propagated to Providers, Peers, or RSes.

## 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

```
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)
```

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.


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

## In the wild

<https://rfc.hashnode.dev/rfc9234-observed-in-the-wild>


© Mingwei Zhang

asn	as_name	org_name	org_country
-----	-----	-----	-----
6939	HURRICANE	Hurricane Electric LLC	US
15562	SNIJDERS	Job Snijders	NL
20555	WSISIZ-AS	Wyższa Szkoła Informatyki Stosowanej i Zarządzania	PL
212068	SHINRA-AS	Shinra Electric Power Company Limited	GB



**Daryll Swer** · Mar 16, 2023  
@DaryllSwer · [Follow](#)  
Replying to @heymingwei

RFC9234 is \*mandatory\* config on MikroTik. You likely saw that from MikroTik users on IXPs or even transits.



**Job Snijders**  
@JobSnijders · [Follow](#)

I'm pretty sure this comes from the YYCIX route servers: as part of the ASPA-filtering deployment we enabled RFC9234 roles too [mailman.nanog.org/pipermail/nano...](mailto:mailman.nanog.org/pipermail/nano...) (source: my ASN is showing up in the RIB dumps, and I manage the YYCIX route servers :-)

10:01 AM · Mar 16, 2023 ⓘ

♥ 4    💬 Reply    ↗ Share

[Read 1 reply](#)

## Thanks for listening!

I am not an engineer, but will do my best to answer your possible questions.

You can find more information about BGP Roles at Qrator Labs blog ([blog.qrator.net/en](http://blog.qrator.net/en)) and Qrator.Radar website ([radar.qrator.net](http://radar.qrator.net)).

If you have more complex questions about BGP Roles - feel free to drop us a line at [radar@qrator.net](mailto:radar@qrator.net)

Or contact me directly via [shapelez@qrator.net](mailto:shapelez@qrator.net)