



Interconnecting Securely

Joanne Liew
Interconnection Manager
Cloudflare



Agenda

- The Internet Infrastructure & Routing Economics
- BGP & Its Security Risks
- Mitigation & Detection
- Working towards Secure Interconnections

The Internet Infrastructure & Routing Economics

- The Internet is a network of networks.
 - A network, or Autonomous System (AS), typically exchanges routing information with another AS through the Border Gateway Protocol (BGP)
- Economics of network for inter-domain routing.
 - Transit (provider-customer) & Peering - the usual types of relationship between two ASes
 - Policy based routing protocol → BGP
 - BGP decides how data is routed between ASes in alignment with the policy interests of two ASes and intermediary ASes
- BGP's role is critical in the global Internet infrastructure
 - But.. it's relying on trust-based configurations between ASes
 - This at times presents security risks in routing within the Internet

BGP and its Security Risks

- The common threats:

- Route leaks

- Propagation of routing announcements beyond their intended scope (RFC7908)

- BGP origin hijackings

- An attack when a malicious or misconfigured AS falsely claims ownership of IP prefixes that it does not own

- Misconfigurations / malicious peering

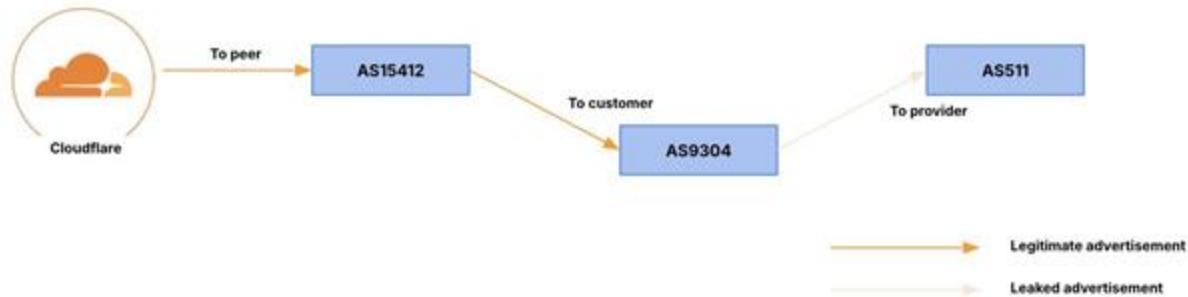
- Incorrect route filtering, missing BGP prefix limits, not honouring or incorrectly tagged BGP communities
 - MitM via peering, reserved prefix announcements

- Lack of origin / route validation

- The impact:
 - Traffic interception, blackholing, suboptimal routing, congestion, etc
 - Causing network outages & service interruptions
- It is critical for all networks to play their part, as routing security matters for both routes received and routes announced
 - Routing security is important in two directions..

- Example of a recent occurrence
 - Route leak due to BGP AS Path error

Example path "5746 5511 **9304** 15412 13335" for a Cloudflare originated prefix



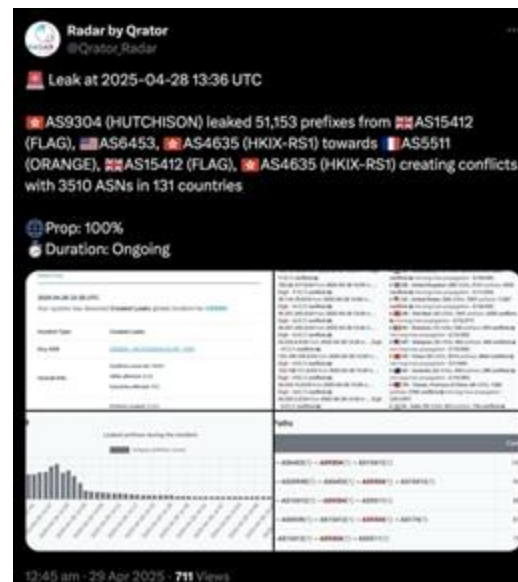
● Example of a recent occurrence

○ Route leak - detected by Cloudflare Radar and other tools

■ Detection on Cloudflare Radar



■ Detection on Qrator Radar



Sources:

<https://radar.cloudflare.com/routing/anomalies/leak-307159>
https://x.com/Qrator_Radar/status/1916896514292773128

- An event causing major impact to Cloudflare (1.1.1.1 incident June 2024)
 - Cloudflare's public DNS resolver 1.1.1.1 was unreachable or reachable with high latency to some users globally, due to a combination of BGP hijacking and a route leak

- BGP Hijack

- 1 AS267613 announced 1.1.1.1/32 to peers, providers, and customers
AS398465 accepted the hijacked route

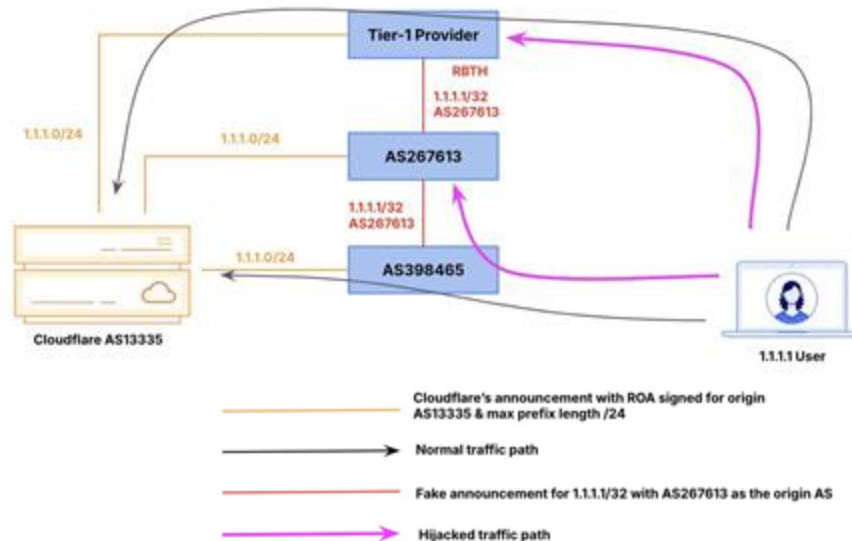
```
monocle search --start-ts 2024-06-27T18:51:00Z --end-ts 2024-06-27T18:55:00Z --prefix '1.1.1.1/32'

A11719514377.130203|206.126.236.209|398465|1.1.1.1/32|398465
267613|IGP|206.126.236.209|0|0|false||route-views.eqix
-
A11719514377.681932|206.82.104.185|398465|1.1.1.1/32|398465
267613|IGP|206.82.104.185|0|0|13538:1|false||route-views.ny
```

- 2 One Tier-1 transit provider accepted the 1.1.1.1/32 announcement as a RTBH (Remote-Triggered Blackhole) route from AS267613, discarding all traffic at their edge that would normally route to Cloudflare.

- An event causing major impact to Cloudflare (1.1.1.1 incident June 2024)
 - Cloudflare's public DNS resolver 1.1.1.1 was unreachable or reachable with high latency to some users globally, due to a combination of BGP hijacking and a route leak

- BGP Hijack



- An event causing major impact to Cloudflare (1.1.1.1 incident June 2024)
 - Cloudflare's public DNS resolver 1.1.1.1 was unreachable or reachable with high latency to some users globally, due to a combination of BGP hijacking and a route leak

- Route Leak

Example path "199524 1031 **262504** 267613 13335" for the Cloudflare originated prefix 1.1.1.0/24

AS262504 received the prefix from AS267613 and leaked to transit provider AS1031

AS1031 redistributed the prefix advertisement to their IX peers and route-servers, thus widening the impact of the leak



- An event causing major impact to Cloudflare (1.1.1.1 incident June 2024)
 - Cloudflare's public DNS resolver 1.1.1.1 was unreachable or reachable with high latency to some users globally, due to a combination of BGP hijacking and a route leak
 - Detection - when some users of 1.1.1.1 experienced disruption
 - Remediation - disabled peering location with AS267613 that is receiving traffic toward 1.1.1.0/24, engaged AS262504 regarding the route leak of 1.1.1.0/24 to their upstream providers
 - Possible mitigation - RPKI origin validation, prefix list filter for v4 prefixes longer than /24, IRR filtering by transit provider

- April 2021 - Vodafone Idea BGP Hijack
 - Vodafone Idea announced more than 30k routes (their own routes + routes that don't belong to them) by mistake.
- June 2019 - Verizon / DQE / Allegheny Technologies Route Leak
 - Allegheny learned thousands of IP prefixes from DQE and incorrectly announced to Verizon. Verizon accepted and propagated the leaked routes.
- June 2015 - Telekom Malaysia Route Leak
 - Telekom Malaysia advertised a large number of prefixes (of customers + other networks) to Global Crossing, who then accepted and announced those prefixes to their peers and customers.
- February 2008 - YouTube Hijack by Pakistan Telecom
 - PTCL hijacked YouTube's IP space using BGP to announce a route that it does not own and announced it to its upstream providers. These were then propagated by one of PTCL's upstream, PCCW.

Sources:

<https://manrs.org/2021/04/a-major-bgp-hijack-by-as55410-vodafone-idea-ltd/>

<https://blog.cloudflare.com/how-verizon-and-a-bgp-optimizer-knocked-large-parts-of-the-internet-offline-today/>

<https://www.bgpmon.net/massive-route-leak-cause-internet-slowdown/>

<https://www.ripe.net/about-us/news/youtube-hijacking-a-ripe-ncc-ris-case-study/>

Some ways to mitigate and detect these risks

- Mitigation

- Best practices

- Filtering - prefix limit, AS-Path filter, prefix filter
 - Up-to-date IRR and AS-SET records
 - E.g. Cloudflare performs leak test during peering turn up

- RPKI

- Sign ROA, enable RPKI validation

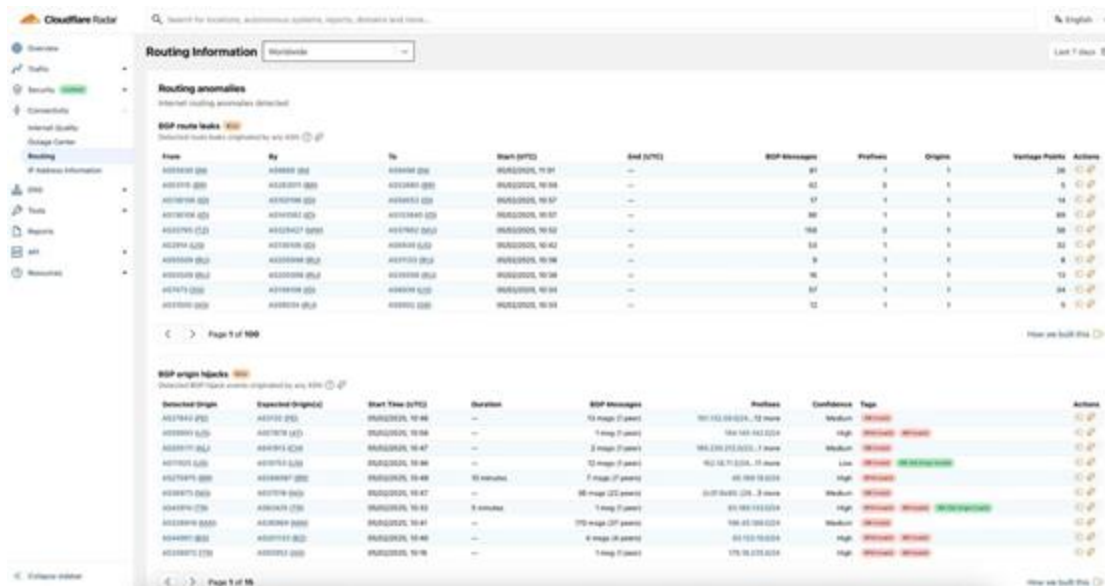
- Detection

- BGP monitoring tools detecting anomalies - RouteViews, Cloudflare Radar

Working towards secure interconnections

- **BGP ASPA** (tracked within <https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-aspa-verification>)
 - BGP AS_PATH Verification based on Autonomous System Provider Authorization (ASPA) objects in the Resource Public Key Infrastructure (RPKI).
- **RFC9234** (tracked within <https://datatracker.ietf.org/doc/rfc9234/>)
 - BGP roles are defined in this RFC - Provider, Customer, RS, RS-Client, Peer
 - Also define the Only to Customer (OTC) BGP attribute - when receiving the OTC attribute from a peer AS, the local AS should only propagate the route to customers.

- Best practices to protect routing on the Internet
- Universal RPKI adoption - <https://isbgpsafeyet.com/>
- Increase observability for data and insights on routing anomalies - <https://radar.cloudflare.com/routing>



The screenshot displays the Cloudflare Radar interface, which provides insights into routing anomalies and BGP origin hijacks. The interface includes a sidebar with navigation options like Overview, Traffic, Security, and DNS. The main content area is divided into two sections: Routing anomalies and BGP origin hijacks.

Routing anomalies

Internet routing anomalies detected

BGP route leaks

Detected route leaks originating by any ASN

From	By	To	Start (UTC)	End (UTC)	BGP Message	Prefixes	Origins	Verdict Points	Actions
AS15169 (US)	AS15169 (US)	AS15169 (US)	2023-03-15 10:01	---	91	1	1	26	
AS15169 (US)	AS15169 (US)	AS15169 (US)	2023-03-15 10:01	---	42	0	1	5	
AS15169 (US)	AS15169 (US)	AS15169 (US)	2023-03-15 10:01	---	17	1	1	14	
AS15169 (US)	AS15169 (US)	AS15169 (US)	2023-03-15 10:01	---	80	1	1	88	
AS15169 (US)	AS15169 (US)	AS15169 (US)	2023-03-15 10:01	---	160	0	1	58	
AS15169 (US)	AS15169 (US)	AS15169 (US)	2023-03-15 10:01	---	14	1	1	32	
AS15169 (US)	AS15169 (US)	AS15169 (US)	2023-03-15 10:01	---	9	1	1	8	
AS15169 (US)	AS15169 (US)	AS15169 (US)	2023-03-15 10:01	---	16	1	1	13	
AS15169 (US)	AS15169 (US)	AS15169 (US)	2023-03-15 10:01	---	17	1	1	24	
AS15169 (US)	AS15169 (US)	AS15169 (US)	2023-03-15 10:01	---	12	1	1	5	

Page 1 of 100

BGP origin hijacks

Detected BGP origin hijacks originating by any ASN

Observed Origin	Expected Origin(s)	Start Time (UTC)	Duration	BGP Message	Prefixes	Confidence	Tags	Actions
AS15169 (US)	AS15169 (US)	2023-03-15 10:01	---	10 message (1 peer)	101.153.101.0/24, 12 more	Medium		
AS15169 (US)	AS15169 (US)	2023-03-15 10:01	---	1 message (1 peer)	101.153.101.0/24, 1 more	High		
AS15169 (US)	AS15169 (US)	2023-03-15 10:01	---	1 message (1 peer)	101.153.101.0/24, 1 more	Medium		
AS15169 (US)	AS15169 (US)	2023-03-15 10:01	---	10 message (1 peer)	101.153.101.0/24, 10 more	Low		
AS15169 (US)	AS15169 (US)	2023-03-15 10:01	10 minutes	1 message (1 peer)	101.153.101.0/24	High		
AS15169 (US)	AS15169 (US)	2023-03-15 10:01	---	10 message (1 peer)	101.153.101.0/24, 10 more	Medium		
AS15169 (US)	AS15169 (US)	2023-03-15 10:01	5 minutes	1 message (1 peer)	101.153.101.0/24	High		
AS15169 (US)	AS15169 (US)	2023-03-15 10:01	---	170 message (17 peers)	101.153.101.0/24	Medium		
AS15169 (US)	AS15169 (US)	2023-03-15 10:01	---	4 message (4 peers)	101.153.101.0/24	High		
AS15169 (US)	AS15169 (US)	2023-03-15 10:01	---	1 message (1 peer)	101.153.101.0/24	High		

Page 1 of 16

Thank you

 peering@cloudflare.com | epp@cloudflare.com

 <https://www.cloudflare.com/partners/peering-portal/>