

# BGP Best Current Practices



These materials are licensed under the Creative Commons Attribution-NonCommercial 4.0 International license (<http://creativecommons.org/licenses/by-nc/4.0/>)



UNIVERSITY OF OREGON

Last updated 22<sup>nd</sup> May 2026



# Acknowledgements

- This material originated from the Cisco ISP/IXP Workshop Programme developed by Philip Smith & Barry Greene
- Use of these materials is encouraged as long as the source is fully acknowledged, and this notice remains in place
- Bug fixes and improvements are welcomed
  - Please email [learn \(at\) nsrc \(dot\) org](mailto:learn@nsrc.org)

# BGP Videos

- NSRC has made a video recording of this presentation, as part of a library of BGP videos for the whole community to use:
  - [https://learn.nsrc.org/bgp#bgp\\_best\\_practices](https://learn.nsrc.org/bgp#bgp_best_practices)

The screenshot shows the NSRC (Network Startup Resource Center) website. The header includes the NSRC logo, navigation links (Home, About, BGP for All, perfSONAR, ScienceDMZ, FedIdM, Contact Us), and a search bar. The main content area is divided into several sections:

- BGP for All**: A text-based introduction to BGP, explaining its role as the primary routing protocol for the Internet and autonomous systems. It describes how BGP maintains paths to different hosts, networks, and gateway routers, and how it determines routing decisions based on rules, filtering, weight, and community.
- Video Topics**: A list of video topics with color-coded buttons: BGP for All (orange), perfSONAR (yellow), ScienceDMZ (blue), and FedIdM (green).
- Introduction to Routing**: A list of video topics including Internet Routing, Routing Protocols, Introduction to IS-IS (UPDATED), IS-IS Levels, IS-IS Adjacencies, Best Configuration Practices for IS-IS on Cisco IOS, IS-IS Authentication, Default Routes and IPv6, Introduction to OSPF, OSPF Areas, OSPF Adjacencies, Best Configuration Practices for OSPF on Cisco IOS, OSPF Authentication, Default Routes and IPv6, Comparing OSPF and IS-IS, Choosing between OSPF and IS-IS, Migrating from OSPF to IS-IS, Migration Plan, and Finalizing Migration.
- Introduction to BGP**: A list of video topics including Introduction to Border Gateway Protocol, Transit and Peering, Autonomous Systems (UPDATED), How BGP works, Supporting Multiple Protocols, IBGP versus EBGP, Setting up EBGP, and Setting up IBGP.
- BGP Case Studies**: A list of video topics including Peering Priorities (NEW), Transit Provider Peering at an IXP (NEW), Customer Multihomed between two IXP members (NEW), Traffic Engineering for an ISP connected to two IXes (NEW), Traffic Engineering for an ISP with two interfaces on one IX LAN (NEW), and Traffic Engineering and CDNs (NEW).
- Communities**: A list of video topics including Communities: RFC 1998 Traffic Engineering, Communities: Simplifying Traffic Engineering, How to Apply Communities to Originated Routes, and How to Use Communities for Service Identification.

A video player is embedded in the center, showing a video titled "BGP for All" from the "Internet Routing" series. The video player includes a play button, a progress bar, and a "Watch on YouTube" button. The video player interface also shows "Watch later", "Share", and "1/158" (likely indicating the video's position in a playlist).

# Configuring BGP

Where do we start?



UNIVERSITY OF OREGON



# Cisco IOS Good Practices

- Network Operators should start off with the following BGP commands as a basic template:

```
router bgp 64511  
  bgp deterministic-med  
  no bgp default ipv4-unicast  
  distance bgp 200 200 200  
  no synchronization  
  no auto-summary
```

← Replace with public ASN

← Turn off IOS assumption that all neighbours will exchange IPv4 prefixes

← Make EBGP and IBGP distance the same & more than any IGP



# EBGP Default Behaviour

- Industry standard is described in RFC8212
  - <https://tools.ietf.org/html/rfc8212>
  - External BGP (EBGP) Route Propagation Behaviour without Policies
- NB: BGP in Cisco IOS is permissive by default
  - This is contrary to industry standard and RFC8212
- Configuring BGP peering without using filters means:
  - All best paths on the local router are passed to the neighbour
  - All routes announced by the neighbour are received by the local router
  - Can have disastrous consequences (see RFC8212)

# EBGP Default Behaviour

- Best practice is to ensure that each EBGP neighbour has inbound and outbound filter applied:

```
router bgp 64511
  address-family ipv4
    neighbor 100.64.0.1 remote-as 64510
    neighbor 100.64.0.1 prefix-list as64510-in in
    neighbor 100.64.0.1 prefix-list as64510-out out
    neighbor 100.64.0.1 activate
```



# EBGP Default Behaviour

- FRR supports RFC8212 by default:

– <https://frrouting.org/>

```
frr.pfs.lab(config)# router bgp 64512 view LAB
frr.pfs.lab(config-router)# bgp ?
<snip>
ebgp-requires-policy          Require in and out policy for eBGP peers (RFC8212)
<snip>
```

- No prefixes will be sent or received to external peers in the absence of inbound and outbound policy

# EBGP Default Behaviour

- Cisco IOS-XR supports RFC8212 by default:

```
TR1(config)# router bgp 64512
TR1(config-router)# bgp ?
<snip>
unsafe-ebgp-policy      Make eBGP neighbors with no policy pass all routes
<snip>
```

– BGP option to turn it off (please don't!)

- Cisco IOS-XE supports RFC8212 as from Release 17.2.1 only ☹️

```
csr8000v(config)# router bgp 64512
csr8000v(config-router)# bgp ?
<snip>
safe-ebgp-policy       Allow propagation of ebgp routes only with policy
                        configured. Command may trigger refresh on all sessions!)
<snip>
```



# What is BGP for??

## What is an IGP not for?



UNIVERSITY OF OREGON



# BGP versus OSPF/ISIS

- Internal Routing Protocols (IGPs)
  - Examples are IS-IS and OSPF
  - Used for carrying **infrastructure** addresses
  - **NOT** used for carrying Internet prefixes or customer prefixes
  - Design goal is to **minimise** number of **prefixes** in IGP to aid **scalability** and rapid **convergence**

# BGP versus OSPF/IS-IS

- BGP is used
  - Internally (IBGP)
  - Externally (EBGP)
- IBGP is used to carry:
  - Some/all Internet prefixes across backbone
  - Customer prefixes
- EBGP is used to:
  - Exchange prefixes with other ASes
  - Implement routing policy

# BGP versus OSPF/IS-IS

- DO NOT:
  - Distribute BGP prefixes into an IGP
  - Distribute IGP routes into BGP
  - Use an IGP to carry customer prefixes
- **YOUR NETWORK WILL NOT SCALE**

# Aggregation

# Aggregation

- Aggregation means announcing the address block received from the RIR to the other ASes connected to your network
- Subprefixes of this aggregate may be:
  - Used internally in the provider network
  - Announced to other ASes to aid with multihoming
- Too many operators are still thinking about class Cs, resulting in a proliferation of /24s in the Internet routing table
  - May 2026: 661010 /24s in IPv4 table of 1049364 prefixes
- **The same is happening for /48s with IPv6**
  - May 2026: 109815 /48s in IPv6 table of 235731 prefixes

# Configuring Aggregation – Cisco IOS

- Service Provider has 100.66.0.0/19 address block
- To put into BGP as an aggregate:

```
router bgp 64511
  address-family ipv4
    network 100.66.0.0 mask 255.255.224.0
  ip route 100.66.0.0 255.255.224.0 null0
```

- The static route is a “pull up” route
  - More specific prefixes within this address block ensure connectivity to Service Provider’s customers
  - “Longest match” lookup

# Aggregation

- Address block should be announced to the Internet as an aggregate
- Subprefixes of address block should **NOT** be announced to Internet unless for traffic engineering
  - See BGP Multihoming presentations
- Aggregate should be generated internally
  - Not on the network borders!

# Announcing Aggregate – Cisco IOS

- Configuration Example

```
router bgp 64511
  address-family ipv4
    network 100.66.0.0 mask 255.255.224.0
    neighbor 100.67.10.1 remote-as 64510
    neighbor 100.67.10.1 prefix-list out-filter out
    neighbor 100.67.10.1 prefix-list default in
    neighbor 100.67.10.1 activate
  !
ip route 100.66.0.0 255.255.224.0 null0
!
ip prefix-list out-filter permit 100.66.0.0/19
ip prefix-list out-filter deny 0.0.0.0/0 le 32
!
ip prefix-list default permit 0.0.0.0/0
```




# Announcing an Aggregate

- Network Operators who don't and won't aggregate are held in poor regard by community
- Registries publish their minimum allocation size
  - For IPv4:
    - /24
  - For IPv6:
    - /48 for assignment, /32 for allocation
- Until 2010, there was no real reason to see anything longer than a /22 IPv4 prefix on the Internet. But now?
  - IPv4 run-out is having an impact
  - It is expected that eventually the global IPv4 table will be mostly /24s

# Aggregation – Example

- Customer has /23 network assigned from AS64510's /19 address block
- AS64510 announces customers' individual networks to the Internet




# Aggregation – Bad Example

- Customer link goes down
  - Their /23 network becomes unreachable
  - /23 is withdrawn from AS64510's IBGP
- Their service provider doesn't aggregate its /19 network block
  - /23 network withdrawal announced to peers
  - Starts rippling through the Internet
  - Added load on all Internet backbone routers as network is removed from routing table
- Customer link returns
  - Their /23 network is now visible to their provider
  - Their /23 network is re-advertised to peers
  - Starts rippling through Internet
  - Load on Internet backbone routers as network is reinserted into routing table
  - Some service providers suppress the flaps
  - Internet may take 10-20 min or longer to be visible
  - Where is the Quality of Service???



# Aggregation – Example

- Customer has /23 network assigned from AS64510's /19 address block
- AS64510 announced /19 aggregate to the Internet



# Aggregation – Good Example

- Customer link goes down
  - Their /23 network becomes unreachable
  - /23 is withdrawn from AS64510's IBGP
- /19 aggregate is still being announced
  - No BGP hold down problems
  - No BGP propagation delays
  - No damping by other network operators
- Customer link returns
  - Their /23 network is visible again
    - The /23 is re-injected into AS64510's IBGP
  - The whole Internet becomes visible immediately
  - Customer has Quality of Service perception

# Aggregation – Summary

- Good example is what everyone should do!
  - Adds to Internet stability
  - Reduces size of routing table
  - Reduces routing churn
  - Improves Internet QoS for **everyone**
- Bad example is what too many still do!
  - Why? Lack of knowledge?
  - Laziness?

# Separation of IBGP and EBGP

- Many network operators do not understand the importance of separating IBGP and EBGP
  - IBGP is where all customer prefixes are carried
  - EBGP is used for announcing aggregate to Internet and for Traffic Engineering
- Do **NOT** do traffic engineering with customer originated IBGP prefixes
  - Leads to instability similar to that mentioned in the earlier bad example
  - Even though aggregate is announced, a flapping subprefix will lead to instability for the customer concerned
- **Generate traffic engineering prefixes on the Border Router**

# The Internet Today (May 2026)

- Current IPv4 Internet Routing Table Statistics

BGP Routing Table Entries	1049364
Prefixes after maximum aggregation	401359
Unique prefixes in Internet	515576
/24s announced	661010
ASNs in use	78429

- (maximum aggregation is calculated by Origin AS)
- (unique prefixes > max aggregation means that operators are announcing prefixes from their blocks without a covering aggregate)



# The Internet Today (May 2026)

- Current IPv6 Internet Routing Table Statistics

BGP Routing Table Entries	235731
/48s announced	109815
ASNs in use	36100



# Efforts to improve aggregation

- The CIDR Report
  - Initiated and operated for many years by Tony Bates
  - Now combined with Geoff Huston's routing analysis
    - [www.cidr-report.org](http://www.cidr-report.org)
    - (covers both IPv4 and IPv6 BGP tables)
  - Results e-mailed on a weekly basis to most operations lists around the world
  - Lists the top 30 service providers who could do better at aggregating
- RIPE Routing WG aggregation recommendations
  - IPv4: RIPE-399 — [www.ripe.net/ripe/docs/ripe-399.html](http://www.ripe.net/ripe/docs/ripe-399.html)
  - IPv6: RIPE-532 — [www.ripe.net/ripe/docs/ripe-532.html](http://www.ripe.net/ripe/docs/ripe-532.html)

# Efforts to Improve Aggregation

## The CIDR Report


- Also computes the size of the routing table assuming network operators performed optimal aggregation
- Website allows searches and computations of aggregation to be made on a per AS basis
  - Flexible and powerful tool to aid network operators
  - Intended to show how greater efficiency in terms of BGP table size can be obtained without loss of routing and policy information
  - Shows what forms of origin AS aggregation could be performed and the potential benefit of such actions to the total table size
  - Very effectively challenges the traffic engineering excuse

## Status Summary

### Table History

Date	Prefixes	CIDR Aggregated
28-04-26	1050826	601013
29-04-26	1051461	601869
30-04-26	1052088	602662
01-05-26	1052912	601878
02-05-26	1052965	601789
03-05-26	1052628	602022
04-05-26	1053060	602181
05-05-26	1052069	602158

Plot: [BGP Table Size](#)



### AS Summary

- 78577 Number of ASes in routing system
- 27631 Number of ASes announcing only one prefix
- 14911 Largest number of prefixes announced by an AS  
[AS16509](#): AMAZON-02 - Amazon.com, Inc., US
- 227730432 Largest address span announced by an AS (/32s)  
[AS749](#): DNIC-AS-00749 - United States Department of Defense DoD, US


Plot: [AS count](#)

Plot: [Average announcements per origin AS](#)

Report: [ASes ordered by originating address span](#)

Report: [ASes ordered by transit address span](#)

Report: [Autonomous System number-to-name mapping \(from Registry WHOIS data\)](#)



## Announced Prefixes

Rank	AS	Type	Originate	Addr Space (pfx)	Transit	Addr space (pfx)	Description
140	AS6389		ORG+TRN Originate:	4590336	/9.87	Transit:	256 /24.00 BELLSOUTH-NET-BLK - AT&T Enterprises, LLC, US

## Aggregation Suggestions

Filter: [Aggregates](#), [Specifics](#)

This report does not take into account conditions local to each origin AS in terms of policy or traffic engineering requirements, so this is an approximate guideline as to aggregation possibilities.

Rank	AS	AS Name	Current	Wthdw	Aggte	Annce	Redctn	%
4714	<a href="#">AS6389</a>	BELLSOUTH-NET-BLK - AT&T Enterprises, LLC, US	56	19	9	46	10	17.86%

Long term deaggregator, now mending their ways – BellSouth in the US

Prefix	AS Path	Aggregation Suggestion
12.130.209.0/24	4777 2497 7018 6389	
65.5.0.0/16	4777 2497 7018 6389	
65.13.0.0/16	4777 2497 7018 6389	
65.14.0.0/15	4777 2497 7018 6389	+ Announce - aggregate of 65.14.0.0/16 (4777 2497 7018 6389) and 65.15.0.0/16 (4777 2497 7018 6389)
65.14.0.0/16	4777 2497 7018 6389	- Withdrawn - aggregated with 65.15.0.0/16 (4777 2497 7018 6389)
65.15.0.0/16	4777 2497 7018 6389	- Withdrawn - aggregated with 65.14.0.0/16 (4777 2497 7018 6389)
65.81.0.0/16	4777 2497 7018 6389	
65.82.0.0/15	4777 2497 7018 6389	+ Announce - aggregate of 65.82.0.0/16 (4777 2497 7018 6389) and 65.83.0.0/16 (4777 2497 7018 6389)
65.82.0.0/16	4777 2497 7018 6389	- Withdrawn - aggregated with 65.83.0.0/16 (4777 2497 7018 6389)
65.83.0.0/16	4777 2497 7018 6389	- Withdrawn - aggregated with 65.82.0.0/16 (4777 2497 7018 6389)
65.83.236.0/22	4777 2497 7018 6389	- Withdrawn - matching aggregate 65.83.0.0/16 4777 2497 7018 6389
66.20.0.0/15	4777 2497 7018 6389	
67.32.0.0/15	4777 2497 7018 6389	+ Announce - aggregate of 67.32.0.0/16 (4777 2497 7018 6389) and 67.33.0.0/16 (4777 2497 7018 6389)
67.32.0.0/16	4777 2497 7018 6389	- Withdrawn - aggregated with 67.33.0.0/16 (4777 2497 7018 6389)
67.33.0.0/16	4777 2497 7018 6389	- Withdrawn - aggregated with 67.32.0.0/16 (4777 2497 7018 6389)
68.16.0.0/15	4777 2497 7018 6389	
68.152.0.0/15	4777 2497 7018 6389	
68.156.0.0/15	4777 2497 7018 6389	
68.208.0.0/16	4777 2497 7018 6389	
68.213.0.0/16	4777 2497 7018 6389	
68.214.0.0/16	4777 2497 7018 6389	
68.216.0.0/16	4777 2497 7018 6389	
68.222.0.0/16	4777 2497 7018 6389	
70.147.0.0/16	4777 2497 7018 6389	
70.148.0.0/16	4777 2497 7018 6389	
70.150.0.0/15	4777 2497 7018 6389	+ Announce - aggregate of 70.150.0.0/16 (4777 2497 7018 6389) and 70.151.0.0/16 (4777 2497 7018 6389)
70.150.0.0/16	4777 2497 7018 6389	- Withdrawn - aggregated with 70.151.0.0/16 (4777 2497 7018 6389)
70.150.196.0/22	4777 2497 174 29802 6389	
70.151.0.0/16	4777 2497 7018 6389	- Withdrawn - aggregated with 70.150.0.0/16 (4777 2497 7018 6389)
70.154.0.0/15	4777 2497 7018 6389	
70.158.0.0/15	4777 2497 7018 6389	
72.149.0.0/16	4777 2497 7018 6389	
72.151.0.0/16	4777 2497 7018 6389	
72.156.0.0/14	4777 2497 7018 6389	
74.164.0.0/15	4777 2497 7018 6389	+ Announce - aggregate of 74.164.0.0/16 (4777 2497 7018 6389) and 74.165.0.0/16 (4777 2497 7018 6389)

## Announced Prefixes

Rank	AS	Type	Originate	Addr Space (pfx)	Transit	Addr space (pfx)	Description
114	AS18403	ORG+TRN	Originate:	6163200 /9.44	Transit:	492032 /13.09	FPT-AS-AP FPT Telecom Company, VN

## Aggregation Suggestions

Filter: [Aggregates](#), [Specifics](#)

This report does not take into account conditions local to each origin AS in terms of policy or traffic engineering requirements, so this is an approximate guideline as to aggregation possibilities.

Rank	AS	AS Name	Current	Wthdw	Aggte	Annce	Redctn	%
8	<a href="#">AS18403</a>	FPT-AS-AP FPT Telecom Company, VN	4547	4222	24	349	4198	92.32%

Prefix	AS Path	Aggregation Suggestion
1.52.0.0/14	4777 2516 18403 18403	
1.52.0.0/18	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.0.0/20	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.0.0/23	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.2.0/23	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.4.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.6.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.7.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.8.0/23	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.10.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.13.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.14.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.15.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.16.0/20	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.16.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.17.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.18.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.19.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.20.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.21.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.22.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.23.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.24.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.25.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.26.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.27.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.28.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.29.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.32.0/20	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.32.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.33.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.34.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.35.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.36.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403
1.52.37.0/24	4777 2516 18403 18403 - Withdrawn - matching	aggregate 1.52.0.0/14 4777 2516 18403 18403

Long term deaggregator –  
FPT in Vietnam

## Announced Prefixes

Rank	AS	Type	Originate	Addr Space (pfx)	Transit	Addr space (pfx)	Description
119	AS7545	ORG+TRN	Originate:	5672192	/9.56	Transit:	3198464 /10.39 TPG-INTERNET-AP TPG Telecom Limited, AU

## Aggregation Suggestions

Filter: [Aggregates](#), [Specifics](#)

Long term deaggregator –  
TPG in Australia

This report does not take into account conditions local to each origin AS in terms of policy or traffic engineering requirements, so this is an approximate guideline as to aggregation possibilities.

Rank	AS	AS Name	Current	Wthdw	Aggte	Annce	Redctn	%
5	<a href="#">AS7545</a>	TPG-INTERNET-AP TPG Telecom Limited, AU	5884	5533	148	499	5385	91.52%

Prefix	AS Path	Aggregation Suggestion
14.2.0.0/17	4608 7575 7545	+ Announce - aggregate of 14.2.0.0/18 (4608 7575 7545) and 14.2.64.0/18 (4608 7575 7545)
14.2.0.0/19	4608 7575 7545	- Withdrawn - aggregated with 14.2.32.0/19 (4608 7575 7545)
14.2.32.0/19	4608 7575 7545	- Withdrawn - aggregated with 14.2.0.0/19 (4608 7575 7545)
14.2.32.0/21	4608 7575 7545	- Withdrawn - matching aggregate 14.2.32.0/19 4608 7575 7545
14.2.40.0/21	4608 7575 7545	- Withdrawn - matching aggregate 14.2.32.0/19 4608 7575 7545
14.2.48.0/21	4608 7575 7545	- Withdrawn - matching aggregate 14.2.32.0/19 4608 7575 7545
14.2.56.0/21	4608 7575 7545	- Withdrawn - matching aggregate 14.2.32.0/19 4608 7575 7545
14.2.64.0/19	4608 7575 7545	- Withdrawn - aggregated with 14.2.96.0/19 (4608 7575 7545)
14.2.96.0/19	4608 7575 7545	- Withdrawn - aggregated with 14.2.64.0/19 (4608 7575 7545)
14.2.128.0/18	4608 7575 7545	
14.2.192.0/20	4608 7575 7545	
14.200.0.0/14	4608 7575 7545	
14.200.0.0/24	4608 7575 7545	- Withdrawn - matching aggregate 14.200.0.0/14 4608 7575 7545
14.200.1.0/24	4608 7575 7545	- Withdrawn - matching aggregate 14.200.0.0/14 4608 7575 7545
14.200.2.0/24	4608 7575 7545	- Withdrawn - matching aggregate 14.200.0.0/14 4608 7575 7545
14.200.3.0/24	4608 7575 7545	- Withdrawn - matching aggregate 14.200.0.0/14 4608 7575 7545
14.200.4.0/24	4608 7575 7545	- Withdrawn - matching aggregate 14.200.0.0/14 4608 7575 7545
14.200.5.0/24	4608 7575 7545	- Withdrawn - matching aggregate 14.200.0.0/14 4608 7575 7545
14.200.7.0/24	4608 7575 7545	- Withdrawn - matching aggregate 14.200.0.0/14 4608 7575 7545
14.200.8.0/24	4608 7575 7545	- Withdrawn - matching aggregate 14.200.0.0/14 4608 7575 7545
14.200.9.0/24	4608 7575 7545	- Withdrawn - matching aggregate 14.200.0.0/14 4608 7575 7545
14.200.10.0/24	4608 7575 7545	- Withdrawn - matching aggregate 14.200.0.0/14 4608 7575 7545
14.200.11.0/24	4608 7575 7545	- Withdrawn - matching aggregate 14.200.0.0/14 4608 7575 7545
14.200.12.0/24	4608 7575 7545	- Withdrawn - matching aggregate 14.200.0.0/14 4608 7575 7545
14.200.13.0/24	4608 7575 7545	- Withdrawn - matching aggregate 14.200.0.0/14 4608 7575 7545
14.200.15.0/24	4608 7575 7545	- Withdrawn - matching aggregate 14.200.0.0/14 4608 7575 7545
14.200.16.0/24	4608 7575 7545	- Withdrawn - matching aggregate 14.200.0.0/14 4608 7575 7545
14.200.17.0/24	4608 7575 7545	- Withdrawn - matching aggregate 14.200.0.0/14 4608 7575 7545
14.200.18.0/24	4608 7575 7545	- Withdrawn - matching aggregate 14.200.0.0/14 4608 7575 7545
14.200.19.0/24	4608 7575 7545	- Withdrawn - matching aggregate 14.200.0.0/14 4608 7575 7545
14.200.20.0/24	4608 7575 7545	- Withdrawn - matching aggregate 14.200.0.0/14 4608 7575 7545
14.200.21.0/24	4608 7575 7545	- Withdrawn - matching aggregate 14.200.0.0/14 4608 7575 7545
14.200.22.0/24	4608 7575 7545	- Withdrawn - matching aggregate 14.200.0.0/14 4608 7575 7545
14.200.23.0/24	4608 7575 7545	- Withdrawn - matching aggregate 14.200.0.0/14 4608 7575 7545
14.200.24.0/24	4608 7575 7545	- Withdrawn - matching aggregate 14.200.0.0/14 4608 7575 7545

## Announced Prefixes

Rank AS Type Originate Addr Space (pfx) Transit Addr space (pfx) Description  
 47 AS12479 ORG+FRN Originate: 14579968 /8.20 Transit: 360448 /13.54 UNI2-AS, ES

Long term deaggregator – Orange in Spain

## Aggregation Suggestions

Filter: [Aggregates](#), [Specifics](#)

This report does not take into account conditions local to each origin AS in terms of policy or traffic engineering requirements, so this is an approximate guideline as to aggregation possibilities.

Rank	AS	AS Name	Current	Wthdw	Aggte	Annce	Redctn	%
4	<a href="#">AS12479</a>	UNI2-AS, ES	7187	6935	41	293	6894	95.92%

Prefix	AS Path	Aggregation Suggestion
1.178.224.0/19	4777 2497 5511 12479	
1.178.224.0/20	4777 2497 5511 12479	- Withdrawn - matching aggregate 1.178.224.0/19 4777 2497 5511 12479
1.178.240.0/20	4777 2497 5511 12479	- Withdrawn - matching aggregate 1.178.224.0/19 4777 2497 5511 12479
37.11.0.0/16	4777 2497 5511 12479	
37.11.0.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.8.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.12.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.16.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.20.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.24.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.28.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.32.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.36.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.40.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.44.0/24	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.45.0/24	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.46.0/23	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.48.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.52.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.56.0/23	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.58.0/23	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.60.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.64.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.68.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.72.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.76.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.80.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.84.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.88.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.92.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.96.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.100.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.104.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479
37.11.108.0/22	4777 2497 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4777 2497 5511 12479

# Importance of Aggregation

- Size of routing table
  - Router Memory is not so much of a problem as it was in the 1990s
  - Routers routinely carry over 2 million prefixes
- Convergence of the Routing System
  - This is a problem
  - Bigger table takes longer for CPU to process
  - BGP updates take longer to deal with
  - BGP Instability Report tracks routing system update activity
  - <https://bgp.potaroo.net/index-upd.html>



# The BGP Instability Report

The BGP Instability Report is updated daily. This report was generated on 04 May 2026 06:20 (UTC+1000)

## 50 Most active ASes for the past 14 days

RANK	ASN	UPDs	%	Prefixes	UPDs/Prefix	AS NAME
1	16509	262848	4.14%	15017	17.50	AMAZON-02 - Amazon.com, Inc., US
2	9294	229554	3.62%	390	588.60	GNETINC-AS-AP GNET INC., US
3	8151	197778	3.11%	12901	15.33	UNINET, MX
4	6648	138584	2.18%	265	522.96	BAYAN-TELECOMMUNICATIONS Bayan Telecommunications, Inc., PH
5	398781	86648	1.36%	1707	50.76	OSL-188 - OCULUS NETWORKS INC, US
6	7821	76290	1.20%	17	4487.65	ZAYO-MN - Inteliquent, inc., US
7	36492	67068	1.06%	61	1099.48	GOOGLEWIFI - Google, LLC, US
8	9829	66341	1.04%	1887	35.16	BSNL-NIB National Internet Backbone, IN
9	22773	63440	1.00%	4787	13.25	ASN-CXA-ALL-CCI-22773-RDC - Cox Communications Inc., US
10	45669	61514	0.97%	330	186.41	MOBILINK-AS-PK PMCL LDI IP TRANSIT, PK
11	39891	58855	0.93%	4943	11.91	ALJAWWALSTC-AS, SA
12	9155	57489	0.91%	438	131.25	QNET Kuwait, KW
13	24863	51090	0.80%	1293	39.51	LINKdotNET-AS, EG
14	335	49108	0.77%	232	211.67	DNIC-ASBLK-00306-00371 - United States Department of Defense DoD, US
15	36903	46993	0.74%	1232	38.14	MT-MPLS, MA
16	11404	39925	0.63%	1729	23.09	AS-WAVE-1 - Wave Broadband, US
17	197207	37507	0.59%	785	47.78	MCCI-AS, IR
18	36914	31956	0.50%	574	55.67	KENET-AS, KE
19	647	27849	0.44%	402	69.28	DNIC-ASBLK-00616-00665 - United States Department of Defense DoD, US
20	8796	26807	0.42%	404	66.35	FD-298-8796 - FASTNET DATA INC, US
21	4007	26656	0.42%	87	306.39	SUBISU-CABLENET-AS-AP Subisu Cablenet Pvt Ltd, Baluwatar, Kathmandu, Nepal, NP
22	10620	25127	0.40%	3430	7.33	Telmex Colombia S.A., CO
23	6147	23727	0.37%	1250	18.98	INTEGRATEL PERU S.A.A., PE
24	37294	23588	0.37%	145	162.68	TNM, MW
25	25019	23108	0.36%	1661	13.91	SAUDINETSTC-AS, SA

50 Most active Prefixes for the past 14 days

RANK	PREFIX	UPDs	%	Origin AS -- AS NAME
1	207.167.116.0/22	11905	0.18%	7954 -- IMMENSE-NETWORKS - Immense Networks, LLC, US
2	124.195.190.0/24	11171	0.17%	38684 -- CMBDAEJEON-AS-KR CMB Daejeon Broadcasting Co.,Ltd, KR
3	130.137.231.0/24	10700	0.16%	16509 -- AMAZON-02 - Amazon.com, Inc., US
4	55.74.56.0/24	9628	0.14%	335 -- DNIC-ASBLK-00306-00371 - United States Department of Defense DoD, US
5	55.74.57.0/24	9572	0.14%	335 -- DNIC-ASBLK-00306-00371 - United States Department of Defense DoD, US
6	55.67.56.0/24	7913	0.12%	335 -- DNIC-ASBLK-00306-00371 - United States Department of Defense DoD, US
7	55.27.217.0/24	7888	0.12%	335 -- DNIC-ASBLK-00306-00371 - United States Department of Defense DoD, US
8	55.27.56.0/24	7857	0.12%	335 -- DNIC-ASBLK-00306-00371 - United States Department of Defense DoD, US
9	130.137.108.0/24	7854	0.12%	16509 -- AMAZON-02 - Amazon.com, Inc., US
10	130.137.12.0/24	7768	0.12%	16509 -- AMAZON-02 - Amazon.com, Inc., US
11	173.82.125.0/24	7759	0.12%	16509 -- AMAZON-02 - Amazon.com, Inc., US
12	201.182.187.0/24	7719	0.11%	267548 -- W P DE LIMA - ME, BR
13	199.47.232.0/22	7358	0.11%	13341 -- TRANQUILITY - Tranquility Internet Services Inc., US
14	84.205.66.0/24	6762	0.10%	12654 -- RIPE-NCC-RIS-AS Reseaux IP Europeens Network Coordination Centre RIPE NCC, NL
15	130.137.219.0/24	6612	0.10%	16509 -- AMAZON-02 - Amazon.com, Inc., US
16	130.137.118.0/24	6459	0.10%	8987 -- AWS-GOVCLOUD AWS-GOVCLOUD, IE
17	173.82.66.0/24	6440	0.10%	16509 -- AMAZON-02 - Amazon.com, Inc., US
18	173.82.63.0/24	6407	0.10%	16509 -- AMAZON-02 - Amazon.com, Inc., US
19	130.137.80.0/24	6399	0.10%	16509 -- AMAZON-02 - Amazon.com, Inc., US
20	173.82.108.0/24	6375	0.09%	16509 -- AMAZON-02 - Amazon.com, Inc., US
21	173.82.67.0/24	6355	0.09%	16509 -- AMAZON-02 - Amazon.com, Inc., US
22	130.137.28.0/24	6257	0.09%	16509 -- AMAZON-02 - Amazon.com, Inc., US
23	130.137.79.0/24	6246	0.09%	16509 -- AMAZON-02 - Amazon.com, Inc., US
24	195.114.122.0/24	6241	0.09%	8369 -- INTERSVYAZ-AS 38-B, Komsomolsky prospekt, RU
25	195.114.123.0/24	6236	0.09%	8369 -- INTERSVYAZ-AS 38-B, Komsomolsky prospekt, RU
26	194.0.225.0/24	6203	0.09%	201983 -- ANSLUTEN-AS, SE
27	193.105.156.0/24	6199	0.09%	8369 -- INTERSVYAZ-AS 38-B, Komsomolsky prospekt, RU
28	130.137.47.0/24	6120	0.09%	16509 -- AMAZON-02 - Amazon.com, Inc., US
29	194.0.231.0/24	6110	0.09%	201983 -- ANSLUTEN-AS, SE
30	130.137.150.0/24	6034	0.09%	16509 -- AMAZON-02 - Amazon.com, Inc., US

# The BGP IPv6 Instability Report

This report is updated daily. The current report was generated on 5 May 2026 01:17 (UTC+1000)

## 50 Most active ASes for the past 14 days

RANK	ASN	UPDs	%	Prefixes	UPDs/Prefix	AS NAME
1	<a href="#">32098</a>	281006	7.34%	1898	148.05	<a href="#">TRANSTELCO-INC - Transtelco Inc, US</a>
2	<a href="#">20473</a>	258689	6.75%	4414	58.61	<a href="#">AS-VULTR - The Constant Company, LLC, US</a>
3	<a href="#">13999</a>	173822	4.54%	388	447.99	<a href="#">Mega Cable, S.A. de C.V., MX</a>
4	<a href="#">7552</a>	114571	2.99%	4014	28.54	<a href="#">VIETEL-AS-AP Viettel Group, VN</a>
5	<a href="#">16509</a>	82298	2.15%	5605	14.68	<a href="#">AMAZON-02 - Amazon.com, Inc., US</a>
6	<a href="#">199340</a>	79353	2.07%	888	89.36	<a href="#">HONG-GUANWAN, HK</a>
7	<a href="#">7296</a>	72444	1.89%	3	24148.00	<a href="#">AS-DYNASCALE-LAX - Dynascale Technologies, Inc., US</a>
8	<a href="#">40138</a>	69088	1.80%	39	1771.49	<a href="#">MDNET - Mach Dilemma, LLC, US</a>
9	<a href="#">270746</a>	60292	1.57%	3	20097.33	<a href="#">HBC TELECOM MULTIMIDIA LTDA, BR</a>
10	<a href="#">211894</a>	53905	1.41%	8	6738.12	<a href="#">UNGLEICH-PLACE18, CH</a>
11	<a href="#">28669</a>	53889	1.41%	4	13472.25	<a href="#">VERO S.A, BR</a>
12	<a href="#">262197</a>	53061	1.39%	448	118.44	<a href="#">MILLICOM CABLE COSTA RICA S.A., CR</a>
13	<a href="#">11172</a>	53055	1.39%	2519	21.06	<a href="#">Alestra, S. de R.L. de C.V., MX</a>
14	<a href="#">38266</a>	51992	1.36%	1315	39.54	<a href="#">VIL-AS-AP Vodafone Idea Ltd, IN</a>
15	<a href="#">202256</a>	49142	1.28%	326	150.74	<a href="#">LAWLIETNET, CN</a>
16	<a href="#">22085</a>	45755	1.19%	193	237.07	<a href="#">Claro SA, BR</a>
17	<a href="#">58955</a>	41539	1.08%	7	5934.14	<a href="#">BANGMODENTERPRISE-TH Bangmod Enterprise Co., Ltd., TH</a>
18	<a href="#">214451</a>	33471	0.87%	5	6694.20	<a href="#">R3IDO101, GB</a>
19	<a href="#">9287</a>	32194	0.84%	8	4024.25	<a href="#">TIDCC-AS-AP True Internet Data Center Company Limited, TH</a>
20	<a href="#">23736</a>	30526	0.80%	51	598.55	<a href="#">TTSSB-MY TM TECHNOLOGY SERVICES SDN. BHD., MY</a>
21	<a href="#">18403</a>	30378	0.79%	3512	8.65	<a href="#">FPT-AS-AP FPT Telecom Company, VN</a>
22	<a href="#">29484</a>	30095	0.79%	6	5015.83	<a href="#">RUB-AS, DE</a>
23	<a href="#">137409</a>	29444	0.77%	1125	26.17	<a href="#">GSLNETWORKS-AS-AP GSL Networks Pty LTD, AU</a>
24	<a href="#">28548</a>	26154	0.68%	208	125.74	<a href="#">Cablevision, S.A. de C.V., MX</a>
25	<a href="#">28573</a>	26138	0.68%	1907	13.71	<a href="#">Claro NXT Telecomunicacoes Ltda, BR</a>

## 50 Most active Prefixes for the past 14 days

RANK	PREFIX	UPDs	%	Origin AS -- AS NAME
1	<a href="#">2606:6e00:8000::/35</a>	72442	1.73%	<a href="#">7296 -- AS-DYNASCALE-LAX - Dynascale Technologies, Inc., US</a>
2	<a href="#">2804:4bec::/32</a>	53865	1.29%	<a href="#">28669 -- VERO S.A, BR</a>
3	<a href="#">2407:7900:ffff:/48</a>	28527	0.68%	<a href="#">58955 -- BANGMODENTERPRISE-TH Bangmod Enterprise Co., Ltd., TH</a>
4	<a href="#">2a05:3e04::/32</a>	23583	0.56%	<a href="#">29484 -- RUB-AS, DE</a>
5	<a href="#">2804:70c0::/48</a>	23090	0.55%	<a href="#">270746 -- HBC TELECOM MULTIMIDIA LTDA, BR</a>
6	<a href="#">2804:248:100::/40</a>	21879	0.52%	<a href="#">28598 -- DB3 SERVICOS DE TELECOMUNICACOES S.A, BR</a>
7	<a href="#">2a14:3f87:9800::/38</a>	21666	0.52%	<a href="#">64289 -- MACARNE - Macarne.com, US</a>
8	<a href="#">2001:67c:1914::/48</a>	21442	0.51%	<a href="#">20708 -- SKODA-AUTO Tr. Vaclava Klementa 869, CZ</a>
9	<a href="#">2804:70c0:32::/48</a>	20820	0.50%	<a href="#">270746 -- HBC TELECOM MULTIMIDIA LTDA, BR</a>
10	<a href="#">2602:fa7c:100::/40</a>	20492	0.49%	<a href="#">26930 -- PARALLAX - PureVoltage Hosting Inc., US</a>
11	<a href="#">2a0c:b641:302::/47</a>	19775	0.47%	<a href="#">204210 -- ZEUSPACKAGINGLTD, IE</a>
12	<a href="#">2a05:79c0:1::/48</a>	16723	0.40%	<a href="#">27381 -- CASALE-MEDIA - Index Exchange Inc., CA</a>
13	<a href="#">2804:70c0::/32</a>	16382	0.39%	<a href="#">270746 -- HBC TELECOM MULTIMIDIA LTDA, BR</a>
14	<a href="#">2804:87bc:dc01::/48</a>	14933	0.36%	<a href="#">61609 -- NextHop Solutions, BR</a>
15	<a href="#">2804:6a0c:ff00::/40</a>	14568	0.35%	<a href="#">270309 -- TURBONET - Provedor de Acesso a Internet, BR</a>
16	<a href="#">2a0b:a8c0:2::/48</a>	14454	0.35%	<a href="#">202181 -- TAMCOM-AS, RU</a>
17	<a href="#">2407:5440::/48</a>	13017	0.31%	<a href="#">141145 -- GIGANET-AS-ID PT Giga Digital Nusantara, ID</a>
18	<a href="#">2a0b:f8c0::/29</a>	12858	0.31%	<a href="#">44297 -- ICODIA ICODIA NETWORK INTEGRITY, FR</a>
19	<a href="#">2602:fb0d:aa::/48</a>	12831	0.31%	<a href="#">61228 -- 64b22a8bb6ab181373c78aa1d1d42cd07844d95e, US</a>
20	<a href="#">2a0f:85c1:e14::/48</a>	12146	0.29%	<a href="#">206098 -- CC1-ROSE-NET, VN</a>
21	<a href="#">2400:cb00:1185::/48</a>	11714	0.28%	<a href="#">14789 -- CLOUDFLARENET - Cloudflare, Inc., US</a>
22	<a href="#">2407:7900:2000::/40</a>	11698	0.28%	<a href="#">58955 -- BANGMODENTERPRISE-TH Bangmod Enterprise Co., Ltd., TH</a>
23	<a href="#">2806:202::/32</a>	11432	0.27%	<a href="#">28458 -- IENTC S DE RL DE CV, MX</a>
24	<a href="#">2804:9524::/32</a>	11355	0.27%	<a href="#">61625 -- Grupo We Dev LTDA, BR</a>
25	<a href="#">240d:c010:165::/48</a>	11248	0.27%	<a href="#">139341 -- ACE-AS-AP ACE, SG</a>
26	<a href="#">2400:7400:11::/48</a>	11176	0.27%	<a href="#">23736 -- TTSSB-MY TM TECHNOLOGY SERVICES SDN. BHD., MY</a>
27	<a href="#">2a0f:85c1:846::/48</a>	11155	0.27%	<a href="#">214451 -- R3IDO101, GB</a>
28	<a href="#">2a0b:9e40::/31</a>	11155	0.27%	<a href="#">214451 -- R3IDO101, GB</a>
29	<a href="#">2a0f:85c1:870::/44</a>	11153	0.27%	<a href="#">214451 -- R3IDO101, GB</a>
30	<a href="#">2401:2260::/32</a>	11053	0.26%	<a href="#">151650 -- SWITCHFIBER-AS-AP Switch Fiber, PH</a>

# Aggregation: Summary

- Aggregation on the Internet could be **MUCH** better
  - 50% saving on Internet routing table size is quite feasible
  - Tools are available
  - Commands on the routers are not hard
  - CIDR-Report webpage

# Receiving Prefixes

# Receiving Prefixes

- There are three scenarios for receiving prefixes from other ASes
  - Customer talking BGP
  - Peer talking BGP
  - Upstream/Transit talking BGP
- Each has different filtering requirements and need to be considered separately

# Receiving Prefixes: From Customers

- Network Operators must only accept prefixes which have been assigned or allocated to their downstream customer
- If the Network Operator has assigned address space to its customer, then the customer IS entitled to announce it back to their provider
- If the Network Operator has NOT assigned address space to its customer, then:
  - Check in the five RIR databases to see if this address space really has been assigned to the customer
  - The tool: `whois -h jwhois.apnic.net x.x.x.0/24`



UNIVERSITY OF OREGON

- (jwhois is “joint whois” and queries the 5 RIR databases)



# Receiving Prefixes: From Customers

- Example use of whois to check if customer is entitled to announce address space:

```
$ whois -h jwhois.apnic.net 202.12.29.0
```

```
inetnum: 202.12.29.0 - 202.12.29.255
netname:  APNIC-SERVICES-AU
descr:   Asia Pacific Network Information Centre
descr:   Regional Internet Registry for the Asia-Pacific Region
descr:   6 Cordelia Street
descr:   South Brisbane
country: AU
geofeed: https://geofeed.apnic.net/geofeed.csv
org:     ORG-APNI1-AP
admin-c: AIC1-AP
tech-c:  AIC1-AP
abuse-c: AA1589-AP
status:  ASSIGNED PORTABLE
mnt-by:  APNIC-HM
mnt-routes: MAINT-APNIC-IS-AP
mnt-irt:  IRT-APNIC-IS-AP
last-modified: 2024-03-01T05:17:46Z
source:  APNIC
```

inetnum – means it is an address delegation to an entity

Portable – means it is an assignment to the customer, the customer can announce it to you

# Receiving Prefixes: From Customers

- Example use of whois to check if customer is entitled to announce address space:

```
$ whois -h jwhois.apnic.net 194.15.141.0

inetnum:      194.15.141.0 - 194.15.141.255
netname:      INETTECH
country:      SE
org:          ORG-ITAS2-RIPE
admin-c:      KEL5-RIPE
tech-c:       KEL5-RIPE
status:       ASSIGNED PI
mnt-by:       RIPE-NCC-END-MNT
mnt-by:       KURTIS-PP-MNT
mnt-routes:   KURTIS-PP-MNT
mnt-domains:  KURTIS-PP-MNT
created:      2003-12-04T09:33:09Z
last-modified: 2016-04-14T08:21:55Z
source:       RIPE
sponsoring-org: ORG-NIE1-RIPE
```

inetnum – means it is an address delegation to an entity

Assigned PI – means its an assignment to the customer, the customer can announce it to you



# Receiving Prefixes: From Customers

- Example use of whois to check if customer is entitled to announce address space:

```
$ whois -h jwhois.apnic.net 193.128.0.0/22
```

```
inetnum:      193.128.0.0 - 193.128.6.255
netname:      UK-PIPEX-19931014
country:      GB
org:          ORG-UA24-RIPE
admin-c:      WERT1-RIPE
tech-c:       UPHM1-RIPE
status:       ALLOCATED PA
remarks:      Please send abuse notification to abuse@uk.uu.net
mnt-by:       RIPE-NCC-HM-MNT
mnt-by:       AS1849-MNT
mnt-routes:   AS1849-MNT
mnt-routes:   WCOM-EMEA-RICE-MNT
mnt-irt:      IRT-MCI-GB
created:      2018-07-30T09:42:04Z
last-modified: 2018-07-30T09:42:04Z
source:       RIPE # Filtered
```

inetnum – means it is an address delegation to an entity

ALLOCATED – means that this is Provider Aggregatable address space and can only be announced by the service provider holding the allocation (in this case Verizon UK)



# Receiving Prefixes from customer: Cisco IOS

- For Example:
  - Downstream has 100.69.0.0/20 block
  - Should only announce this to upstreams
  - Upstreams should only accept this from them
- Configuration on upstream

```
router bgp 100
  address-family ipv4
    neighbor 100.67.10.1 remote-as 101
    neighbor 100.67.10.1 prefix-list customer in
    neighbor 100.67.10.1 prefix-list default out
    neighbor 100.67.10.1 activate
  !
ip prefix-list customer permit 100.69.0.0/20
!
ip prefix-list default permit 0.0.0.0/0
```

# Receiving Prefixes: From Peers

- A peer is a Network Operator with whom you agree to exchange prefixes you originate into the Internet routing table
  - Prefixes you accept from a peer are only those they have indicated they will announce
  - Prefixes you announce to your peer are only those you have indicated you will announce

# Receiving Prefixes: From Peers

- Agreeing what each will announce to the other:
  - Exchange of e-mail documentation as part of the peering agreement, and then ongoing updates
- OR
- Use of the Internet Routing Registry and configuration tools such as:
  - IRRToolSet: <https://github.com/irrtolset/irrtolset>
  - bgpq4: <https://github.com/bgp/bgpq4>  
(uses NTT's IRR database by default)



# Receiving Prefixes from peer: Cisco IOS

- For Example:
  - Peer has 220.50.0.0/16, 61.237.64.0/18 and 81.250.128.0/17 address blocks
- Configuration on local router

```
router bgp 100
  address-family ipv4
    neighbor 100.67.10.1 remote-as 101
    neighbor 100.67.10.1 prefix-list my-peer in
    neighbor 100.67.10.1 prefix-list my-prefix out
    neighbor 100.67.10.1 activate
  !
ip prefix-list my-peer permit 220.50.0.0/16
ip prefix-list my-peer permit 61.237.64.0/18
ip prefix-list my-peer permit 81.250.128.0/17
ip prefix-list my-peer deny 0.0.0.0/0 le 32
!
ip prefix-list my-prefix permit 100.67.16.0/20
```

# Receiving Prefixes: From Upstream/Transit Provider

- Upstream/Transit Provider is a Network Operator who you pay to give you transit to the **WHOLE** Internet
  - Receiving prefixes from them is not desirable unless really necessary
    - Traffic Engineering – see BGP Multihoming presentations
  - Ask upstream/transit provider to either:
    - originate a default-route
- OR
- announce one prefix you can use as default

# Receiving Prefixes: From Upstream/Transit Provider

- Downstream Router Configuration

```
router bgp 100
  address-family ipv4
    network 100.66.0.0 mask 255.255.224.0
    neighbor 100.65.7.1 remote-as 101
    neighbor 100.65.7.1 prefix-list infilter in
    neighbor 100.65.7.1 prefix-list outfilter out
    neighbor 100.65.7.1 activate
!
ip prefix-list infilter permit 0.0.0.0/0
!
ip prefix-list outfilter permit 100.66.0.0/19
```



# Receiving Prefixes: From Upstream/Transit Provider

- Upstream Router Configuration

```
router bgp 101
  address-family ipv4
    neighbor 100.65.7.2 remote-as 100
    neighbor 100.65.7.2 default-originate
    neighbor 100.65.7.2 prefix-list cust-in in
    neighbor 100.65.7.2 prefix-list cust-out out
    neighbor 100.65.7.2 activate
!
ip prefix-list cust-in permit 100.66.0.0/19
!
ip prefix-list cust-out permit 0.0.0.0/0
```



# Receiving Prefixes: From Upstream/Transit Provider

- Note that the previous configuration originated default route from the upstream's access router
  - This is fine as long as the upstream always has a route to the Global Internet
  - But if the upstream loses their entire Internet access, originating the default route from the access route will blackhole customer traffic
- To avoid problems with blackholing customer traffic, the upstream instead propagates the default route they learn from their upstreams
  - If their upstream links go down, the default will be withdrawn, and the customers will not receive the default any more

# Receiving Prefixes: From Upstream/Transit Provider

- If it is necessary to receive prefixes from any provider, care is required.
  - Don't accept default (unless you need it)
  - Don't accept your own prefixes
- Special use prefixes for IPv4 and IPv6:
  - [www.rfc-editor.org/rfc/rfc6890.txt](http://www.rfc-editor.org/rfc/rfc6890.txt)
- For IPv4:
  - Don't accept prefixes longer than /24 (?)
    - /24 was the historical class C
- For IPv6:
  - Don't accept prefixes longer than /48 (?)
    - /48 is the design minimum delegated to a site

# Receiving Prefixes: From Upstream/Transit Provider

- Check Team Cymru's list of "bogons"
  - [www.team-cymru.com/bogon-reference-http](http://www.team-cymru.com/bogon-reference-http)
- For IPv4 also consult:
  - [www.rfc-editor.org/rfc/rfc6441.txt](http://www.rfc-editor.org/rfc/rfc6441.txt) (BCP171)
- Bogon Route Server:
  - [www.team-cymru.com/bogon-reference-bgp](http://www.team-cymru.com/bogon-reference-bgp)
  - Supplies a BGP feed (IPv4 and/or IPv6) of address blocks which should not appear in the BGP table

# Receiving IPv4 Prefixes

```
router bgp 100
  network 101.10.0.0 mask 255.255.224.0
  neighbor 100.65.7.1 remote-as 101
  neighbor 100.65.7.1 prefix-list in-filter in
!
ip prefix-list in-filter deny 0.0.0.0/0           ! Default
ip prefix-list in-filter deny 0.0.0.0/8 le 32     ! RFC1122 local host
ip prefix-list in-filter deny 10.0.0.0/8 le 32    ! RFC1918
ip prefix-list in-filter deny 100.64.0.0/10 le 32  ! RFC6598 shared address
ip prefix-list in-filter deny 101.10.0.0/19 le 32  ! Local prefix
ip prefix-list in-filter deny 127.0.0.0/8 le 32   ! Loopback
ip prefix-list in-filter deny 169.254.0.0/16 le 32 ! Auto-config
ip prefix-list in-filter deny 172.16.0.0/12 le 32 ! RFC1918
ip prefix-list in-filter deny 192.0.0.0/24 le 32  ! RFC6598 IETF protocol
ip prefix-list in-filter deny 192.0.2.0/24 le 32  ! TEST1
ip prefix-list in-filter deny 192.168.0.0/16 le 32 ! RFC1918
ip prefix-list in-filter deny 198.18.0.0/15 le 32 ! Benchmarking
ip prefix-list in-filter deny 198.51.100.0/24 le 32 ! TEST2
ip prefix-list in-filter deny 203.0.113.0/24 le 32 ! TEST3
ip prefix-list in-filter deny 224.0.0.0/3 le 32   ! Multicast & Experimental
ip prefix-list in-filter deny 0.0.0.0/0 ge 25     ! Prefixes >/24
ip prefix-list in-filter permit 0.0.0.0/0 le 32
```

# Receiving IPv6 Prefixes

```
router bgp 100
  network 2020:3030::/32
  neighbor 2020:3030::1 remote-as 101
  neighbor 2020:3030::1 prefix-list v6in-filter in
!
ipv6 prefix-list v6in-filter permit 64:ff9b::/96          ! RFC6052 v4v6trans
ipv6 prefix-list v6in-filter deny 2001::/23 le 128       ! RFC2928 IETF prot
ipv6 prefix-list v6in-filter deny 2001:2::/48 le 128    ! Benchmarking (RFC5180)
ipv6 prefix-list v6in-filter deny 2001:10::/28 le 128   ! ORCHID
ipv6 prefix-list v6in-filter deny 2001:db8::/32 le 128  ! Documentation (RFC3849)
ipv6 prefix-list v6in-filter deny 2002::/16 le 128     ! Deny all 6to4
ipv6 prefix-list v6in-filter deny 2020:3030::/32 le 128 ! Local Prefix
ipv6 prefix-list v6in-filter deny 3ffe::/16 le 128     ! Formerly 6bone
ipv6 prefix-list v6in-filter deny 3fff::/20 le 128     ! Documentation (new)
ipv6 prefix-list v6in-filter permit 2000::/3 le 48     ! Global Unicast
ipv6 prefix-list v6in-filter deny ::/0 le 128
```

Note: These filters block Teredo (serious security risk) and 6to4 (deprecated by RFC7526)

# Receiving Prefixes

- Paying attention to prefixes received from customers, peers and transit providers assists with:
  - The integrity of the local network
  - The integrity of the Internet
- Responsibility of all Network Operators to be good Internet citizens

# Receiving BGP attributes



UNIVERSITY OF OREGON



# Receiving BGP attributes

- BGP attributes are sent as part of the BGP updates for each prefix
- Common attributes operators need to be aware of, for routing best practice, are:
  - MED
  - AS numbers (only public ASNs are routable)
  - BGP Communities
  - AS\_SET

# Receiving Prefixes: MEDs?

- MEDs are used by EBGP neighbours to indicate preferred entry point into their network over two or more links with their neighbour
  - Allows the operator to determine entry path into their network
    - Might have unintended consequences within their peer's network
  - Many operators will override MEDs attached to BGP announcements by setting their own local-preference values

# Receiving Prefixes: Bogon ASNs?

- What about prefixes originated by bogon AS numbers?
  - Public ranges are 1-64495 (excluding 23456) and 131072-458751
    - IANA is distributing AS blocks to the RIRs from the latter range
  - All other ASNs are either for documentation, or for private use, or are unassigned
    - And any prefixes originating from those need to be dropped
    - Configuration error? Malicious intent?
- What would the AS\_PATH filter look like?
  - Challenging with regular expression (as per IOS)
  - Easier with AS ranges (as per Bird or JunOS)

# Filtering bogon ASNs – BIRD

- Here is a function showing how to filter bogon ASNs, as described previously:

```
function as_path_contains_bogons()
int set invalid_asns;
{
    invalid_asns = [
        0,                # Reserved
        23456,            # Transition AS
        64496..64511,    # Documentation ASNs
        64512..65534,    # Private ASNs
        65535,           # Reserved
        65536..65551,    # Documentation ASNs
        65552..131071,   # Reserved
        458752..419999999, # IANA Reserved
        4200000000..4294967294, # Private ASNs
        4294967295       # Reserved
    ];
    return bgp_path ~ invalid_asns;
}
```

# Filtering bogon ASNs – FRR

- Here is an AS-PATH regexp showing how to filter bogon ASNs:

```
bgp as-path access-list Bogon_ASNs deny _0_  
bgp as-path access-list Bogon_ASNs deny _23456_  
bgp as-path access-list Bogon_ASNs deny _6449[6-9]_  
bgp as-path access-list Bogon_ASNs deny _64[5-9][0-9][0-9]_  
bgp as-path access-list Bogon_ASNs deny _6[5-9][0-9][0-9][0-9]_  
bgp as-path access-list Bogon_ASNs deny _[7-9][0-9][0-9][0-9][0-9]_  
bgp as-path access-list Bogon_ASNs deny _1[0-2][0-9][0-9][0-9][0-9]_  
bgp as-path access-list Bogon_ASNs deny _130[0-9][0-9][0-9]_  
bgp as-path access-list Bogon_ASNs deny _1310[0-6][0-9]_  
bgp as-path access-list Bogon_ASNs deny _13107[0-1]_  
bgp as-path access-list Bogon_ASNs deny _45875[2-9]_  
bgp as-path access-list Bogon_ASNs deny _4587[6-9][0-9]_  
bgp as-path access-list Bogon_ASNs deny _458[8-9][0-9][0-9]_  
bgp as-path access-list Bogon_ASNs deny _459[0-9][0-9][0-9]_  
bgp as-path access-list Bogon_ASNs deny _4[6-9][0-9][0-9][0-9][0-9]_  
bgp as-path access-list Bogon_ASNs deny _[5-9][0-9][0-9][0-9][0-9][0-9]_  
bgp as-path access-list Bogon_ASNs deny _[0-9][0-9][0-9][0-9][0-9][0-9][0-9]_  
bgp as-path access-list Bogon_ASNs deny _[0-9][0-9][0-9][0-9][0-9][0-9][0-9][0-9]_  
bgp as-path access-list Bogon_ASNs deny _[0-9][0-9][0-9][0-9][0-9][0-9][0-9][0-9][0-9]_  
bgp as-path access-list Bogon_ASNs deny _[0-9][0-9][0-9][0-9][0-9][0-9][0-9][0-9][0-9][0-9]_  
bgp as-path access-list Bogon_ASNs permit .*
```



# Receiving Prefixes: BGP Communities?

- BGP communities are attached to BGP announcements to indicate:
  - Internal policy within an AS
  - External policy supported by a peer, for:
    - Onward routing policy/traffic engineering
    - Filtering (eg Remotely Triggered Blackhole Filtering)
    - Traffic engineering between the two networks
- Different BGP implementations have different default BGP community behaviours – consult:
  - Vendor documentation
  - [www.rfc-editor.org/rfc/rfc8642.txt](http://www.rfc-editor.org/rfc/rfc8642.txt) for discussion of some of the issues operators need to be aware of

# Receiving Prefixes: BGP Communities

- Do **NOT** accept community values that are not expected
  - Match expected values
  - Overwrite received community values with your own default value

```
ip community-list standard lp-250 permit 65534:250
!  
route-map ebgp-import permit 5  
  description Set high preference  
  match community lp-250  
  set local-preference 250  
  set community 65534:100  
!  
route-map ebgp-import permit 10  
  description Set our default community  
  set community 65534:100  
!
```

We only expect this community from our EBGP neighbour

Cisco IOS: this overwrites all incoming community values



# Receiving Prefixes: BGP Communities

- Do **NOT** send community values that are not needed by the peer
  - This avoids propagating your internal communities to other networks
    - Propagating your internal communities leaves you open to DoS or worse!

```
route-map ebgp-export permit 5
  description Tell upstream to set local-pref 250
  set community 65534:250 ←
!
```

Cisco IOS: this overwrites all other community values

- Propagate all communities within the AS (by IBGP)
  - This may need changes to your equipment's default!

# Receiving Prefixes: AS\_SETs?

- The AS\_SET in BGP is introduced by the router that is performing aggregation (by proxy) on behalf of its own and customer (downstream) ASes

- Described in RFC4271 (the BGP standard)

- It would appear in the BGP table like this:

```
76.165.0.0/16      11537 32440 {2048,2055,10349,22985,23207,23294,23366,26002,26303,26333,30564,54529,396992,401290}
```

- This states that AS32440 router is generating the 76.165.0.0/16 aggregate to cover subnets of this aggregate introduced by the ASNs within the curly brackets

- Normal Internet operations this century has simply introduced the aggregate within the network the address space belongs

- Subnets used by other (customer/downstream) ASes would simply appear as originated by those ASes



# Receiving Prefixes: AS\_SETs?

- There are several well-known and well-documented problems with the use of AS\_SETs
  - <https://www.rfc-editor.org/rfc/rfc6472.txt>
  - Operators mistakenly include Private ASNs in the list
  - Most networks are signing ROAs to support Route Origin Validation; and ROAs have to include the AS originating the prefix, otherwise ROV would fail
- As from 2011 their use has not been recommended, and as of May 2025, their use has been fully deprecated:
  - <https://www.rfc-editor.org/rfc/rfc9774>
  - This means that the AS\_SET must not be used in any BGP configuration and implementations no longer have the command set to support them

# Receiving Prefixes: AS\_SETs?

- But there are still plenty of AS\_SETs visible on the Internet today
- A complete IPv4 list:
  - <https://bgp.nsrc.org/REN/OIX/ipv4-as-set-table>
- A complete IPv6 list (even though IPv6 deployment has mostly happened **after** AS\_SETs were deprecated):
  - <https://bgp.nsrc.org/REN/OIX/ipv6-as-set-table>
- What should a network operator be doing?
  - Dropping routes with AS\_SETs in the path?
  - OR
  - Propagate them as normal?

# Receiving BGP attributes

- Care is needed when receiving prefixes, to be aware of some of the optional BGP attributes that may be attached
  - BGP communities are only intended for policy decisions within an AS or between two peering ASes
  - MEDs may have unexpected consequences for traffic flows on the peer's network
  - Bogon ASNs, like bogon address space, must never be used or announced to the global Internet
  - AS\_SETs are fully deprecated, yet are still visible



# Prefixes into IBGP

# Injecting prefixes into IBGP

- Use IBGP to carry customer prefixes
  - Don't use IGP
- Point static route to customer router address (next-hop)
- Use BGP network statement
- As long as static route exists (interface active), prefix will be in BGP

# Router Configuration: network statement

- Example:

```
interface loopback 0
  ip address 100.64.3.1 255.255.255.255
!
interface GigabitEthernet 5/0/0.20
  description Customer p-t-p link
  ip address 100.65.0.1 255.255.255.252
  ip verify unicast reverse-path
!
ip route 100.71.10.0 255.255.252.0 100.65.0.2
!
router bgp 100
  address-family ipv4
    network 100.71.10.0 mask 255.255.252.0
!
```



# Injecting prefixes into IBGP

- Interface flap will result in prefix withdraw and reannounce
  - use “`ip route . . . permanent`”
- Many network operators redistribute static routes into BGP rather than using the network statement
  - Only do this if you understand why

# Router Configuration: redistribute static

- Example:

```
ip route 100.71.10.0 255.255.252.0 100.65.0.2
!
router bgp 100
  address-family ipv4
    redistribute static route-map static-to-bgp
<snip>
!
route-map static-to-bgp permit 10
  match ip address prefix-list ISP-block
  set origin igp
  set community 100:1000
<snip>
!
ip prefix-list ISP-block permit 100.71.10.0/22 le 30
```



# Injecting prefixes into IBGP

- Route-map `static-to-bgp` can be used for many things:
  - Setting communities and other attributes
  - Setting origin code to IGP, etc
- Be careful with prefix-lists and route-maps
  - Absence of either/both means all statically routed prefixes go into IBGP

# Summary

- Best Practices Covered:
  - When to use BGP
  - When to use ISIS/OSPF
  - Aggregation
  - Receiving Prefixes
  - Prefixes into BGP

# Interconnection Best Practices

## PeeringDB and the Internet Routing Registry



UNIVERSITY OF OREGON



# Interconnection Best Practices

- Types of Peering
- Using the PeeringDB and IXPDB
- Using the Internet Routing Registry

# Types of Peering (1)

- Private Peering
  - Where two network operators agree to interconnect their networks, and exchange their respective routes, for the purpose of ensuring their customers can reach each other directly over the peering link
- Settlement Free Peering
  - No traffic charges
  - The most common form of peering
- Paid Peering
  - Where two operators agree to exchange traffic charges for a peering relationship

# Types of Peering (2)

- Bi-lateral Peering
  - Very similar to Private Peering, but usually takes place at a public peering point (IXP)
- Multilateral Peering
  - Takes place at Internet Exchange Points, where operators all peer with each other via a Route Server
- Mandatory Multilateral Peering
  - Where operators are forced to peer with each other as condition of IXP membership
  - **Strongly discouraged: Has no record of success**

# Types of Peering (3)


- Open Peering
  - Where a network operator publicly states that they will peer with all parties who approach them for peering
  - Commonly found at IXPs where the network operator participates via the Route Server (RS)
- Selective Peering
  - Where a network operator's peering policy depends on the nature of the operator who requests peering with them
  - At IXPs, the operator will not peer with RS but will only peer bilaterally
- Restrictive Peering
  - Where a network operator decides who its peering partners are, and is generally not approachable to considering peering opportunities

# Types of Peering (4)

- The Peering Database documents network operator peering policies
  - <https://www.peeringdb.com>
- All operators with an AS are recommended to register in the PeeringDB
  - All operators who are considering peering or are peering must be in the PeeringDB to enhance their peering opportunities
- Participation in peering fora is encouraged too
  - Global Peering Forum (GPF) – (for North American peering)
  - Regional Peering Fora (Europe, Middle East, Africa, Asia, Caribbean, Latin America)
  - Many countries now have their own Peering Fora

# Types of Peering (5)

- The IXPDB documents IXPs and their participants around the world
  - <https://ixpdb.euro-ix.net/en/>
- All Internet Exchange Point operators should register their IXP in the database
  - IXPs using IXP Manager will have this happen as part of the IXP Manager set up
  - Provides the LAN IP addresses of each member to facilitate automation



The IXP Database (IXPDB) is a live public source of data related to Internet Exchange Points (IXPs). It automatically collects data directly from IXP management systems in order to provide an accurate up-to-date view of the global interconnection landscape.

The IXPDB contains a wide range of data related to IXP architecture, membership, hardware, capacity, and utilisation. As a result, it can uniquely satisfy a wide range of use cases including commercial analysis, network automation, and historical benchmarking. The data can be explored via this website and exported in full detail via our free API.



UNIVERSITY OF OREGON



## HKIX Gold Sponsor

Peers **261**   Connections **337**   Open Peers **157**   Total Speed **17.9T**   % with IPv6 **86**

EXPORT

Organization	<a href="#">Hong Kong Internet eXchange Limited</a>
Also Known As	HKIX
Long Name	Hong Kong Internet Exchange
City	Hong Kong
Country	HK
Continental Region	Asia Pacific
Service Level	Not Disclosed
Terms	Not Disclosed
Last Updated	2025-09-10T03:35:27Z
Notes	
Logo	

### Contact Information

Company Website	<a href="https://www.hkix.net/">https://www.hkix.net/</a>
Traffic Stats Website	<a href="https://www.hkix.net/hkix/stat/aggt/hkix-aggregate.html">https://www.hkix.net/hkix/stat/aggt/hkix-aggregate.html</a>
Technical Email	<a href="mailto:noc@hkix.net">noc@hkix.net</a>
Technical Phone	+85239439900
Policy Email	<a href="mailto:info@hkix.net">info@hkix.net</a>
Policy Phone	+85239438800
Sales Email	<a href="mailto:order@hkix.net">order@hkix.net</a>
Sales Phone	+85239438800
Health Check	

### Peers at this Exchange Point

Add your network

Peer Name  IPv4	ASN IPv6	Speed Port Location	Policy
<a href="#">2012 Limited</a> 123.255.90.135	4658 2001:7fa:0:1::ca28:a087	10G	Selective
<a href="#">2012 Limited</a> 123.255.90.122	4658 2001:7fa:0:1::ca28:a07a	1G	Selective
<a href="#">ACE CDN</a> 123.255.91.67	139341 2001:7fa:0:1::ca28:a143	200G	Open
<a href="#">ACE CDN</a> 123.255.91.194	139341 2001:7fa:0:1::ca28:a1c2	200G	Open
<a href="#">ACE CDN</a> 123.255.92.38	139341 2001:7fa:0:1::ca28:a226	200G	Open
<a href="#">ACE CDN</a> 123.255.91.79	139341 2001:7fa:0:1::ca28:a14f	200G	Open
<a href="#">ACME Universal</a> 123.255.91.24	56190	1G	Open
<a href="#">ADVANCED HOSTING</a> 123.255.91.178	39572 2001:7fa:0:1::ca28:a1b2	100G	Selective
<a href="#">Advanced Wireless Network Co. Ltd.(IIG)</a> 123.255.92.80	45430 2001:7fa:0:1::ca28:a250	100G	Selective
<a href="#">AgotoZ HK</a> 123.255.90.175	141167 2001:7fa:0:1::ca28:a0af	10G	Open
<a href="#">Akamai Prolexic DDoS Mitigation</a> 123.255.91.26	32787 2001:7fa:0:1::ca28:a11a		Selective

## Amazon.com Diamond Sponsor

EXPORT

Organization	<a href="#">Amazon.com, Inc.</a>
Also Known As	Amazon Web Services
Long Name	
Company Website	<a href="https://www.amazon.com">https://www.amazon.com</a>
ASN	16509
IRR as-set/route-set	AS16509:AS-AMAZON
Route Server URL	
Looking Glass URL	
Network Types	Enterprise
IPv4 Prefixes	16000
IPv6 Prefixes	8000
Traffic Levels	Not Disclosed
Traffic Ratios	Balanced
Geographic Scope	Global
Protocols Supported	<input checked="" type="checkbox"/> Unicast IPv4 <input type="checkbox"/> Multicast <input checked="" type="checkbox"/> IPv6 <input checked="" type="checkbox"/> Never via route servers
Last Updated	2026-04-21T03:05:40Z
Public Peering Info Updated	2026-04-21T03:05:33Z
Peering Facility Info Updated	2026-04-14T06:07:41Z
Contact Info Updated	2026-03-03T10:22:08Z
Notes	<p><b>AWS Peering - <a href="https://peering.aws/">https://peering.aws/</a></b></p> <p><b>Peering requests</b></p> <p><b>IX</b></p> <ul style="list-style-type: none"> <li>Please use our peering portal to request IX peering sessions: <a href="https://interconnect.amazon/">https://interconnect.amazon/</a></li> <li>Note: If you're about to request IX peering sessions with</li> </ul>

### Public Peering Exchange Points

Exchange AZ IPv4	ASN IPv6	Speed Port Location	RS Peer	BFD Support
<a href="#">1-IX EU</a> 185.0.42.91	16509 2001:7f8:115:1::91	200G	<input type="radio"/>	<input type="radio"/>
<a href="#">AIX</a> 43.254.83.9	16509 2401:91e0:31::16:50 9:1	100G	<input type="radio"/>	<input type="radio"/>
<a href="#">AIX</a> 43.254.83.10	16509 2401:91e0:31::16:50 9:2	100G	<input type="radio"/>	<input type="radio"/>
<a href="#">AKL-IX (Auckland NZ)</a> 43.243.21.112	16509 2001:7fa:11:6:0:407 d:0:1	100G	<input type="radio"/>	<input type="radio"/>
<a href="#">AKL-IX (Auckland NZ)</a> 43.243.21.113	16509 2001:7fa:11:6:0:407 d:0:2	100G	<input type="radio"/>	<input type="radio"/>
<a href="#">AMS-IX</a> 80.249.210.100	16509 2001:7f8:1::a501:650 9:1	600G	<input type="radio"/>	<input type="radio"/>
<a href="#">AMS-IX</a> 80.249.210.217	16509 2001:7f8:1::a501:650 9:2	600G	<input type="radio"/>	<input type="radio"/>
<a href="#">AMS-IX Bangkok</a> 103.100.140.204	16509 2402:b740:15:388:0: a501:6509:2	100G	<input type="radio"/>	<input type="radio"/>
<a href="#">AMS-IX Bangkok</a> 103.100.140.204	16509 2402:b740:15:388:0: a501:6509:2	100G	<input type="radio"/>	<input type="radio"/>

### Interconnection Facilities

## Arelion (Twelve99)

EXPORT

Organization	<a href="#">Arelion</a>
Also Known As	f/k/a Telia Carrier
Long Name	
Company Website	<a href="https://www.arelion.com/">https://www.arelion.com/</a>
ASN	1299
IRR as-set/route-set ?	RIPE::AS1299:AS-TWELVE99
Route Server URL	
Looking Glass URL	<a href="https://lg.twelve99.net/">https://lg.twelve99.net/</a>
Network Types	NSP
IPv4 Prefixes ?	
IPv6 Prefixes ?	
Traffic Levels ?	100+Tbps
Traffic Ratios	Balanced
Geographic Scope	Global
Protocols Supported	<input checked="" type="checkbox"/> Unicast IPv4 <input type="checkbox"/> Multicast <input checked="" type="checkbox"/> IPv6 <input checked="" type="checkbox"/> Never via route servers ?
Last Updated	2025-12-11T14:05:02Z
Public Peering Info Updated	
Peering Facility Info Updated	2026-03-03T05:58:05Z
Contact Info Updated	2023-06-20T13:36:16Z
Notes ?	<p><b>All trouble ticket requests or support related emails should be sent to <a href="mailto:support@arelion.com">support@arelion.com</a>.</b></p> <p>AS1299 is matching RPKI validation state and reject invalid prefixes from peers and customers. Our looking-glass marks validation state for all prefixes. Please review your registered ROAs to reduce number of invalid prefixes.</p>

### Public Peering Exchange Points

Exchange AZ v IPv4	ASN IPv6	Speed Port Location	RS Peer	BFD Support
-----------------------	-------------	------------------------	---------	-------------

No filter matches.  
You may filter by **Exchange**, **ASN** or **Speed**.

### Interconnection Facilities

Facility AZ v ASN	Country City
<a href="#">123.NET - DC1 - 24700 Northwestern Hwy.</a> 1299	United States of America Southfield
<a href="#">1530 SWIFT - NOCIX</a> 1299	United States of America North Kansas City
<a href="#">1623 Farnam - Omaha</a> 1299	United States of America Omaha
<a href="#">165 Halsey Meet-Me Room</a> 1299	United States of America Newark
<a href="#">2Connect DC10</a> 1299	Hungary Budapest
<a href="#">365 Data Centers Detroit (DT1)</a> 1299	United States of America Southfield
<a href="#">3U Rechenzentrum Berlin</a> 1299	Germany Berlin
<a href="#">Aitelecom Mérida</a> 1299	Mexico Mérida
<a href="#">aixit Frankfurt</a> 1299	Germany Frankfurt
<a href="#">Arelion Dusseldorf DDF/B</a> 1299	Germany Düsseldorf
<a href="#">Arelion London HEX</a>	United Kingdom

# Internet Routing Registry

- Many major transit providers and several content providers pay attention to what is contained in the Internet Routing Registry
  - There are many IRRs operating, the most commonly used being those hosted by the Regional Internet Registries, RADB, and some transit providers
- Best practice for any AS holder is to document their routing policy in the IRR
  - A route-object is the absolute minimum requirement

# Internet Routing Registry

- IRR objects can be created via the database web-interfaces or submitted via email
- Policy language used to be known as RPSL
- Problems:
  - IRR contains a lot of outdated information
  - Network operators not following best practices
- Some network operators now using RPKI and ROAs to securely indicate the origin AS of their routes
  - Takes priority over IRR entries
  - RPKI and ROAs covered in other presentations

# Which Internet Routing Registry database to use?

- Members of a Regional Internet Registry are strongly encouraged to use their RIR's Internet Routing Registry instance
  - Usually managed via the RIR's member portal giving easy access for creation and update of objects
  - Provided as part of the RIR's services to its members
- Operators who do not belong to any RIR generally use:
  - Their upstream transit provider's Routing Registry (if provided)
  - The RADB (<https://www.radb.net>)
    - Placing objects in the RADB requires an annual subscription fee
    - RADB now uses IRRDv4 – objects with RPKI Invalid cannot be created; existing RPKI Invalid objects will NOT be visible in a query, nor can they be modified



# Route Object: Purpose

- Documents which Autonomous System number is originating the route listed
- Required by many major transit providers
  - They build their customer and peer filter based on the route-objects listed in the IRR
  - Referring to at least the 5 RIR routing registries and the RADB
  - Some operators run their own Routing Registry
    - May require their customers to place a Route Object there (if not using the 5 RIR or RADB versions of the IRR)

# Route Object: Examples

```
route:      202.144.128.0/20
descr:     DRUKNET-BLOCK-A1
country:   BT
notify:    ioc@bt.bt
mnt-by:    MAINT-BT-DRUKNET
origin:    AS18024
last-modified: 2018-09-18T09:37:40Z
source:    APNIC
```

This declares that AS18024 is the origin of 202.144.128.0/20

```
route6:    2405:D000::/32
descr:     DRUKNET-IPV6-BLOCK
origin:    AS17660
notify:    netops@bt.bt
mnt-by:    MAINT-BT-DRUKNET
last-modified: 2010-07-21T03:46:02Z
source:    APNIC
```

This declares that AS17660 is the origin of 2405:D000::/32



# AS Object: Purpose

- Documents peering policy with other Autonomous Systems
  - Lists network information
  - Lists contact information
  - Lists routes announced to neighbouring autonomous systems
  - Lists routes accepted from neighbouring autonomous systems
- Some operators pay close attention to what is contained in the AS Object
  - Some configure their border router BGP policy based on what is listed in the AS Object

# AS Object: Example

```
aut-num:          AS17660
as-name:          DRUKNET-AS
descr:           DrukNet ISP, Bhutan Telecom, Thimphu
country:         BT
org:             ORG-BTL2-AP
import:          from AS6461      action pref=100;      accept ANY
export:          to AS6461        announce AS-DRUKNET-TRANSIT
import:          from AS2914      action pref=150;      accept ANY
export:          to AS2914        announce AS-DRUKNET-TRANSIT
<snip>
import:          from AS135666    action pref=250;      accept AS135666
export:          to AS135666      announce {0.0.0.0/0} AS-DRUKNET-TRANSIT
admin-c:         DNO1-AP
tech-c:          DNO1-AP
notify:          netops@bt.bt
mnt-irt:         IRT-BTTELECOM-BT
mnt-by:          APNIC-HM
mnt-lower:       MAINT-BT-DRUKNET
mnt-routes:      MAINT-BT-DRUKNET
last-modified:   2019-06-09T22:40:10Z
source:          APNIC
```

Examples of inbound and  
outbound policies – RPSL

# AS-Set: Purpose

- The AS-Set is used by network operators to group AS numbers they provide transit for in an easier to manage form
  - Convenient for more complicated policy declarations
  - Used mostly by network operators who build their EBGP filters from their IRR entries
  - Commonly used at Internet Exchange Points to handle large numbers of peers

# AS-Set: Example

```
as-set:      AS-DRUKNET-TRANSIT
descr:      DrukNet transit networks
members:    AS17660
members:    AS132232
members:    AS134715
members:    AS135666
members:    AS137925
members:    AS59219
members:    AS18024
members:    AS18025
members:    AS137994
members:    AS140695
members:    AS151498
members:    AS151955
members:    AS152317
members:    AS138558
admin-c:    DNO1-AP
tech-c:     DNO1-AP
notify:     netops@bt.bt
mnt-by:     MAINT-BT-DRUKNET
last-modified: 2024-09-16T04:35:58Z
source:     APNIC
```

← Lists all the autonomous systems within the AS-DRUKNET-TRANSIT group

# Hierarchical AS-Set

- The usage of hierarchical AS-Set (RFC2622) is strongly recommended now (and required for APNIC IRR) – this helps resolve name collisions

```
as-set:      AS-GEMNET
descr:      GEMNET LLC
country:    MN
members:    AS9934, AS9484, AS10219, AS9789,
            AS38038, AS24496, AS24559, AS4850,
<snip>
tech-c:     GA263-AP
admin-c:    GA263-AP
mnt-by:     MAINT-GEMNET-MN
mnt-lower:  MAINT-GEMNET-MN
last-modified: 2023-09-26T01:25:15Z
source:     APNIC
```

VS

```
as-set:      AS-GEMNET
descr:      GEMNET s.r.o. ASes
members:    AS59479
members:    AS202733
tech-c:     DUMY-RIPE
admin-c:    DUMY-RIPE
mnt-by:     GEMNETCZ-MNT
created:    2013-08-19T09:49:13Z
last-modified: 2024-08-27T14:09:27Z
source:     RIPE
```

- Solution: AS-Set name changes to AS45204:AS-GEMNET
- Consult [https://sanog.org/resources/sanog41/SANOG41\\_Conference-Recent-IRR-changes\\_Maz.pdf](https://sanog.org/resources/sanog41/SANOG41_Conference-Recent-IRR-changes_Maz.pdf) for more information and migration steps

# Summary

- PeeringDB
  - An industry Best Practice so that:
    - Network operators can promote the interconnects they participate in and attract more peering partners
- IXPDB
  - An industry Best Practice so that:
    - Internet Exchange Points can show their participants and help make the interconnect more attractive for potential participants
- IRR
  - An industry Best Practice:
    - So that network operators can document which autonomous system is originating their prefixes
    - Used by network operators to filter prefixes received from their customers and peers

# Route Origin Authorisation

Steps to securing the Routing System



UNIVERSITY OF OREGON



# Route Origin Authorisation

- Essential first step to secure the global routing system
- Covered in detail in separate presentation slide deck:
  - <https://nsrc.org/activities/agendas/en/riso-5-days/networking/routing-security/en/presentations/BGP-Origin-Validation.pdf>
- But there are some important best practices
  - Signing ROAs
  - Implementing ROV to drop "invalids"

# Route Origin Authorisation (ROA)

- A digital object that contains a list of address prefixes and one AS number
- It is an authority created by a prefix holder to authorise an AS Number to originate one or more specific route advertisements
- Publish a ROA using your RIR member portal
  - Consult your RIR for how to use their member portal to publish your ROAs

# Route Origin Authorisation

- A typical ROA would look like this:

<b>Prefix</b>	10.10.0.0/16
<b>Max-Length</b>	/18
<b>Origin-AS</b>	AS65534

- There can be more than one ROA per address block
  - Allows the operator to originate prefixes from more than one AS
  - Caters for changes in routing policy or prefix origin
- **NB: Only create ROAs for the aggregate and the exact subnets expected in the routing table!! (See RFC9319)**

# Route Origin Validation

- Route Origin Validation means checking if the prefix received has a valid ROA
  - Valid ROA means that the prefix (and subnet) is being originated from the correct origin AS
  - See the “BGP Origin Validation” presentation for more in-depth content
- Implementing ROV means checking the validation database with what is learned from BGP peers:
  - Valid – allow; Invalid – drop; NotFound – allow (at lower preference?)
- **Problem:** how is this implemented in routers today?

# Route Origin Validation

- The ideal would be to write directly to the active BGP table
- Some implementations use existing EBGP policy handling routines
  - ADJ-RIB-IN: table of all prefixes received prior to policy being applied
  - Route Refresh (RFC2918)
- Routers which maintain the ADJ-RIB-IN:
  - Apply the ROV policy to the stored received BGP table
  - Updates are applied “automatically” to the BGP table and therefore the FIB
  - No impact on any BGP peers (Route Refresh not needed)

# Route Origin Validation

- Routers which do NOT maintain the ADJ-RIB-IN:
  - Apply the ROV policy by sending a Route Refresh to peers
  - When there are a large number of ROAs (April 2025 sees over 670k), and frequent changes or updates of ROAs:
    - Routers are sending frequent Route Refresh requests to peers (typically every few minutes)
    - Peers are being “bombarded” by Route Refresh requests: significant resource burden when they send the full or a large portion of the BGP table
    - Severe control plane CPU impact on the peer router (effectively a Denial of Service on the peer router)
  - As more and more ROAs are created and altered globally, this problem becomes significantly more serious!



# Route Refresh: Route Origin Validation

- JunOS implements ADJ-RIB-IN by default
  - ROA updates do not cause a problem when operating ROV
- Cisco does not implement ADJ-RIB-IN by default:
  - Applies to all versions of Cisco IOS/IOS-XE and older versions of IOS-XR
  - **MUST turn on soft-reconfiguration if running ROV on the router**
  - Soft-reconfiguration is similar concept to ADJ-RIB-IN
    - Note that Route Refresh CLI seems to be no longer accessible

# Enabling Cisco's Soft Reconfiguration

```
router bgp 64510
  address-family ipv4
    neighbor 100.64.1.1 remote-as 64511
    neighbor 100.64.1.1 route-map infiltrer in
    neighbor 100.64.1.1 soft-reconfiguration inbound
```

- When the policy needs to be changed:

```
clear ip bgp 100.64.1.1 soft in
```

- Note:
  - When “soft-reconfiguration” is enabled, there is no access to the route-refresh capability CLI
  - `clear ip bgp 100.64.1.1 in` also does a soft refresh



# Using Cisco's Soft-Reconfiguration

- Strongly recommended when deploying Route Origin Validation
- Operators will also use soft-reconfiguration when troubleshooting EBGP peer problems
  - Soft reconfiguration enabled on an EBGP session means that the operator can see which prefixes were sent by a neighbour **before** any policy is applied
  - This helps save arguments between operators about whose BGP filters may have configuration errors!



# Configuration Tips

Of passwords, tricks and templates



UNIVERSITY OF OREGON



# IBGP and IGP Reminder!

- Make sure loopback is configured on router
  - IBGP between loopbacks, NOT physical interfaces
- Make sure IGP carries loopback IPv4 /32 and IPv6 /128 address
- Consider the DMZ nets:
  - Use unnumbered interfaces?
  - Use next-hop-self on IBGP neighbours
  - Or carry the DMZ IPv4 /30s and IPv6 /127s in the IBGP
  - Basically, keep the DMZ nets out of the IGP!

# IBGP: Next-hop-self

- BGP speaker announces external network to IBGP peers using router's local address (loopback) as next-hop
- Used by many service providers on edge routers
  - Preferable to carrying DMZ point-to-point link addresses in the IGP
  - Reduces size of IGP to just core infrastructure
  - Alternative to using unnumbered interfaces
  - Helps scale network
  - Many service providers consider this “best practice”

# Limiting AS\_Path Length

- Some BGP implementations have problems with long AS\_PATHS
  - Memory corruption
  - Memory fragmentation
- Even using AS\_PATH prepends, it is not normal to see more than 20 ASNs in a typical AS\_PATH in the Internet Routing Table today
  - The Internet is around 5 ASes deep on average
  - Largest AS\_PATH is usually 16-20 ASNs

```
neighbor x.x.x.x maxas-limit 20
```



# Limiting AS\_Path Length

- Some announcements have ridiculous lengths of AS-paths

- This example is an error in one IPv6 implementation

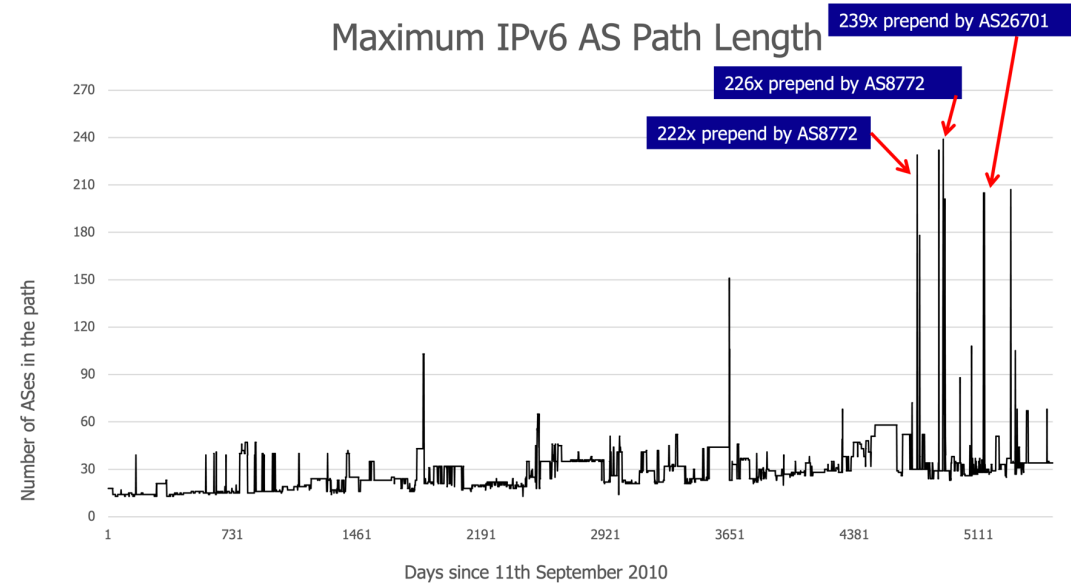
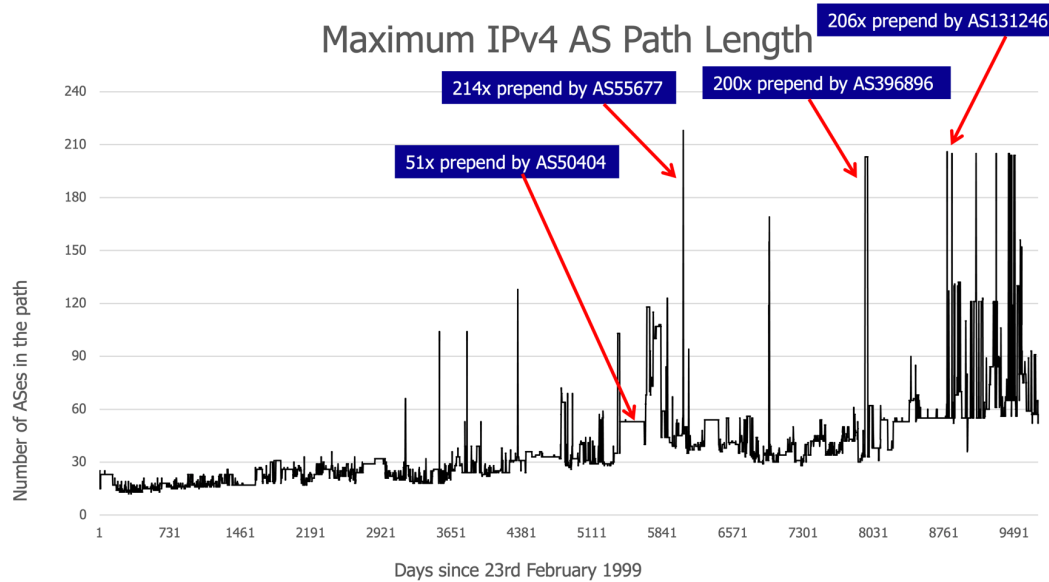
```
*> 3FFE:1600::/24      22 11537 145 12199 10318 10566 13193 1930 2200 3425 293
5609 5430 13285 6939 14277 1849 33 15589 25336 6830 8002 2042 7610 i
```

- This example shows 100 prepends (for no obvious reason)

```
*>i193.105.15.0      2516 3257 50404 50404 50404 50404 50404 50404 50404 50404 50404
50404 50404 50404 50404 50404 50404 50404 50404 50404 50404 50404 50404 50404
50404 50404 50404 50404 50404 50404 50404 50404 50404 50404 50404 50404 50404
50404 50404 50404 50404 50404 50404 50404 50404 50404 50404 50404 50404 50404
50404 50404 50404 50404 50404 50404 50404 50404 50404 50404 50404 50404 50404
50404 50404 50404 50404 50404 50404 50404 50404 50404 50404 50404 50404 50404
50404 50404 50404 50404 50404 50404 50404 50404 50404 50404 50404 50404 50404
50404 50404 i
```

- If your implementation supports it, consider limiting the maximum AS-path length you will accept

# AS\_PATH Length over the years



# BGP Maximum Prefix Tracking

- Allow configuration of the maximum number of prefixes a BGP router will receive from a peer
  - Two level control:
    - Warning threshold: log warning message
    - Maximum: tear down the BGP peering, manual intervention required to restart
- ```
neighbor <x.x.x.x> maximum-prefix <max> [restart N] [<threshold>] [warning-only]
```
- **restart** is an optional keyword which will restart the BGP session N minutes after being torn down
  - **threshold** is an optional parameter between 1 to 100
    - Specify the percentage of <max> that will cause a warning message to be generated. Default is 75%.
  - **warning-only** is an optional keyword which allows log messages to be generated but peering session will not be torn down

# Private-AS – Application


- A network operator with end-sites multihomed on their backbone (RFC2270)

or

- A corporate network with several regions but connections to the Internet only in the core

or

- Within a BGP Confederation



# Private-AS – Removal


- Private ASNs MUST be removed from all prefixes announced to the public Internet
  - Include configuration to remove private ASNs in the EBGP template
- As with RFC1918 address space, private ASNs are intended for internal use
  - They must not be leaked to or used on the public Internet
- Cisco IOS

```
neighbor x.x.x.x remove-private-AS
```



# BGP TTL “hack”

- Implement RFC5082 on BGP peerings
  - (Generalised TTL Security Mechanism)
  - Neighbour sets TTL to 255
  - Local router expects TTL of incoming BGP packets to be 254
  - No one apart from directly attached devices can send BGP packets which arrive with TTL of 254, so any possible attack by a remote miscreant is dropped due to TTL mismatch



# BGP TTL “hack”

- TTL Hack:
  - Both neighbours must agree to use the feature
  - TTL check is much easier to perform than MD5
  - (Called BTSH – BGP TTL Security Hack)
- Provides “security” for BGP sessions
  - In addition to packet filters of course
  - MD5 should still be used for messages which slip through the TTL hack
  - See <https://www.nanog.org/meetings/nanog27/presentations/meyer.pdf> for more details

# BGP TTL “hack”

- Configuration example:

```
neighbor 100.121.0.2 ttl-security hops 1
```

- BGP neighbour status:

```
Router# sh ip bgp neigh 100.121.0.2
...
Minimum incoming TTL 254, Outgoing TTL 255
Local host: 100.121.0.1, Local port: 41103
Foreign host: 100.121.0.2, Foreign port: 179
```

- The neighbour must set the same configuration
  - If they don't, the BGP session will not come up

# Templates

- Good practice to configure templates for everything
  - Vendor defaults tend not to be optimal or even very useful for service providers
  - Service providers create their own defaults by using configuration templates
- EBGP and IBGP examples follow
  - Also see Team Cymru's BGP templates
    - [www.team-cymru.com/community-services](http://www.team-cymru.com/community-services)

# IBGP Template Example

- IBGP between loopbacks!
- Next-hop-self
  - Keep DMZ and external point-to-point out of IGP
- Always send communities in IBGP
  - Otherwise BGP policy accidents will happen
  - (Default on some vendor implementations, optional on others)
- Hardwire BGP to version 4
  - Yes, this is being paranoid!
  - Prevents accidental configuration of BGP version 3 which is still supported in some implementations

# IBGP Template Example (continued)

- Use passwords on IBGP session
  - Not being paranoid, **VERY** necessary
  - It's a secret shared between you and your peer
  - If arriving packets don't have the correct MD5 hash, they are ignored
  - Helps defeat miscreants who wish to attack BGP sessions
- Powerful preventative tool, especially when combined with filters and the TTL "hack"

# EBGP Template Example

- BGP damping
  - Do **NOT** use it unless you understand the impact
  - Do **NOT** use the vendor defaults without thinking
- Cisco's Soft Reconfiguration
  - Do **NOT** use unless troubleshooting or doing Route Origin Validation – it will consume considerable amounts of extra memory for BGP
- Remove private ASNs from announcements
  - Common omission today
- Use extensive filters, with “backup”
  - Use AS-path filters to backup prefix filters
  - Keep policy language for implementing policy, rather than basic filtering

# EBGP Template Example (continued)

- Use password agreed between you and peer on EBGP session
- Use maximum-prefix tracking
  - Router will warn you if there are sudden increases in BGP table size, bringing down EBGP if desired
- Limit maximum as-path length inbound
- Log changes of neighbour state
  - ...and monitor those logs!
- Make BGP admin distance higher than that of any IGP
  - Otherwise, prefixes heard from outside your network could override your IGP!!

# Mutually Agreed Norms for Routing Security

Industry Best Practices  
to ensure Security of  
the Routing System



MANRS



UNIVERSITY OF OREGON



# Routing Security

- Implement the recommendations in [www.manrs.org](http://www.manrs.org)
  - Prevent propagation of incorrect routing information
    - Filter BGP peers, in & out!
  - Prevent traffic with spoofed source addresses
    - BCP38 – Unicast Reverse Path Forwarding
  - Facilitate communication between network operators
    - NOC to NOC Communication
    - Up-to-date details in Route and AS Objects, and PeeringDB
  - Facilitate validation of routing information
    - Route Origin Authorisation using RPKI



MANRS



UNIVERSITY OF OREGON



# MANRS 1)

- Filtering prefixes inbound and outbound
  - RFC8212 requires all EBGP implementations to reject prefixes received and announced in the absence of any policy
- Advice: **Never** set up an EBGP session without inbound and outbound prefix filters
  - If full table required, block at least the bogons (see earlier)

# MANRS 2)

- Implementing BCP 38
  - Unicast Reverse Path Forwarding
  - (Deny outbound traffic from customers which has spoofed source addresses)
- Advice: implement uRPF on ***all*** single-homed customer facing interfaces
  - Cheaper (CPU & RAM) than implementing packet filters

# MANRS 3)

- Facilitate NOC to NOC communication
  - Know the **direct** NOC contacts for your customer Network Operators, your peer Network Operators, and your upstream Network Operators
  - This is not calling their “customer support line”
  - Make sure NOC contact info is part of any service contract
  - Up to date info in Route and AS Objects
  - Up to date AS info in PeeringDB
- Advice: NOC contact info for all connected Autonomous Networks is known to your NOC

# MANRS 4)

- Facilitate validation of Routing Information
  - RPKI and Route Origin Authorisation (ROA)
  - All routes originated need to be signed to indicate that your AS is authorised to originate these routes
    - Helps secure the global routing system
- Advice: Sign ROAs for all originated routes using RPKI
  - And make sure all customer originated routes are also signed
  - Validate received routes from all peers
    - High priority for validated routes
    - Discard invalid routes
    - Low priority for unsigned routes

# MANRS summary

- If your organisation supports and implements all 4 techniques in your network
  - Then join MANRS
  - [www.manrs.org/join/](http://www.manrs.org/join/)
- MANRS for Operators
- MANRS for IXPs
- MANRS for CDN & Cloud Providers



MANRS



UNIVERSITY OF OREGON



# Summary

- Use configuration templates
- Standardise the configuration
- Be aware of standard “tricks” to avoid compromise of the BGP session
- Anything to make your life easier, network less prone to errors, network more likely to scale
- Implement the four fundamentals of MANRS
- It’s all about scaling – if your network won’t scale, then it won’t be successful

# Questions?



UNIVERSITY OF OREGON

