BGP Best Current Practices

ISP Workshops

Last updated 4th February 2021

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 *workshop (at) bgp4all.com*

Philip Smith

Configuring BGP

Where do we start?

Cisco IOS Good Practices

- ISPs 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
```

What is BGP for??

What is an IGP not for?

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 ISP 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
  - January 2021: 488425 /24s in IPv4 table of 841044 prefixes
- The same is happening for /48s with IPv6
  - January 2021: 51923 /48s in IPv6 table of 105863 prefixes

Configuring Aggregation – Cisco IOS

- ISP 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 ISP'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 101
  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

- ISPs 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

Aggregation – Example

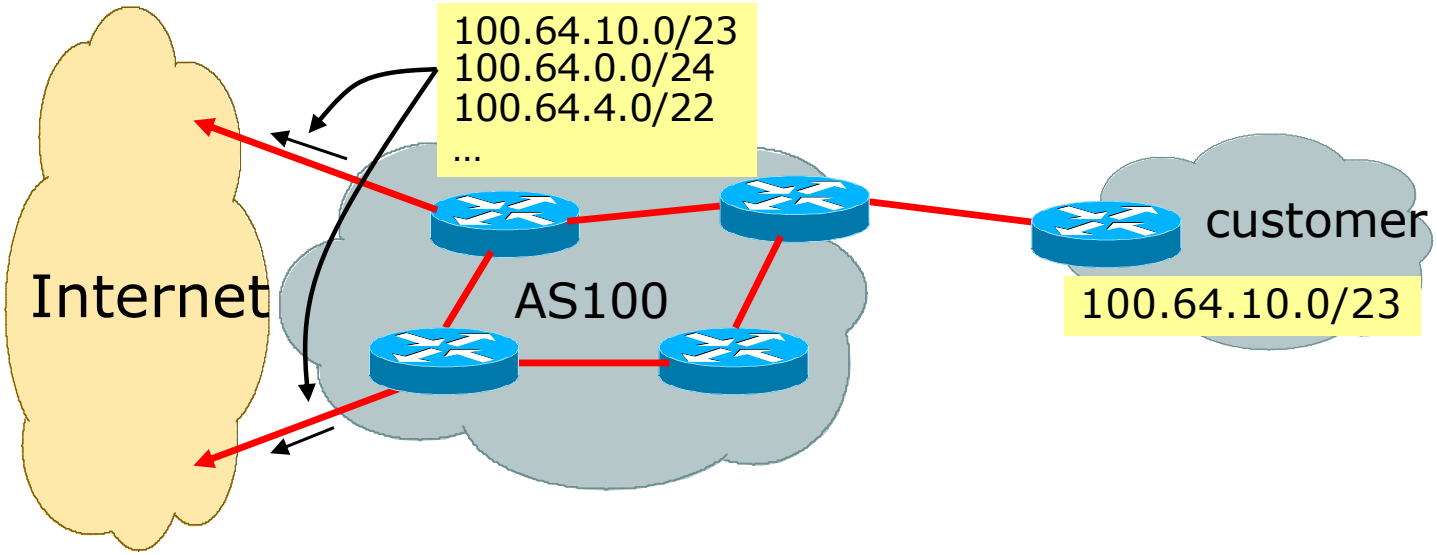

- Customer has /23 network assigned from AS100's /19 address block
- AS100 announces customers' individual networks to the Internet

Aggregation – Bad Example

- Customer link goes down
  - Their /23 network becomes unreachable
  - /23 is withdrawn from AS100's IBGP
- Their ISP 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 ISP
  - 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 ISP's suppress the flaps
  - Internet may take 10-20 min or longer to be visible
  - Where is the Quality of Service???

Aggregation – Example

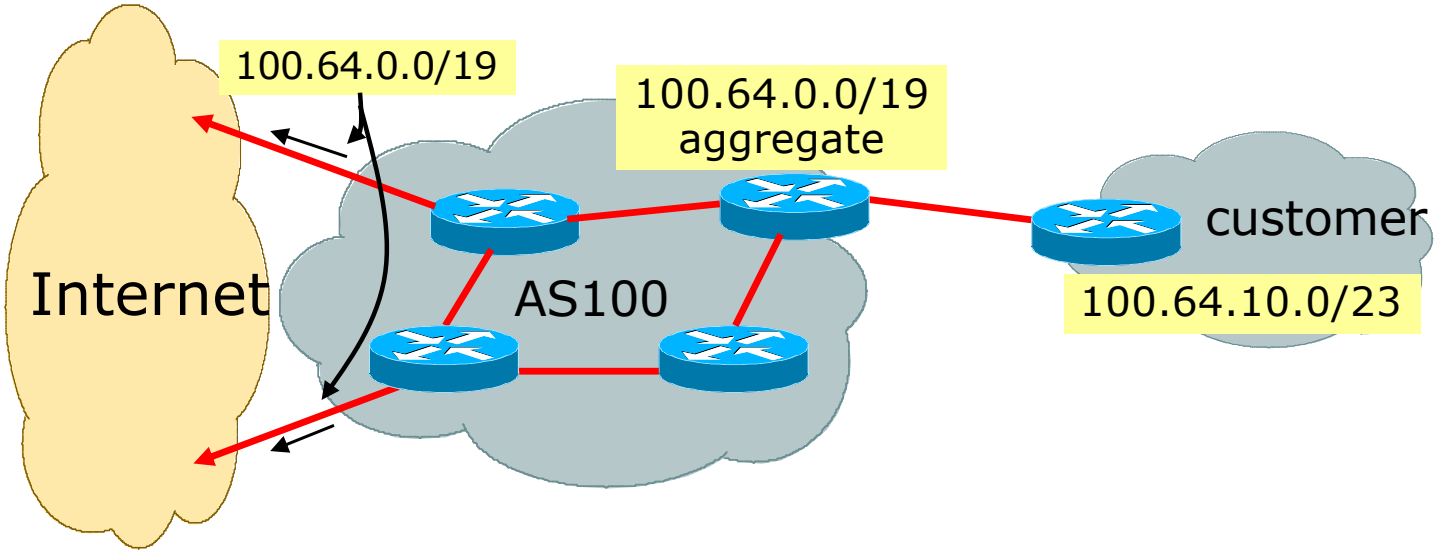

- Customer has /23 network assigned from AS100's /19 address block
- AS100 announced /19 aggregate to the Internet

Aggregation – Good Example

- Customer link goes down
  - Their /23 network becomes unreachable
  - /23 is withdrawn from AS100's IBGP
- /19 aggregate is still being announced
  - No BGP hold down problems
  - No BGP propagation delays
  - No damping by other ISPs

- Customer link returns
- Their /23 network is visible again
  - The /23 is re-injected into AS100'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 ISPs 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 (January 2021)

- Current IPv6 Internet Routing Table Statistics

BGP Routing Table Entries	841044
Prefixes after maximum aggregation	321740
Unique prefixes in Internet	401597
/24s announced	488425
ASNs in use	70404

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

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
    - (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
  - IPv6: RIPE-532 — www.ripe.net/ripe/docs/ripe-532.html

Efforts to Improve Aggregation
The CIDR Report

- Also computes the size of the routing table assuming ISPs performed optimal aggregation
- Website allows searches and computations of aggregation to be made on a per AS basis
  - Flexible and powerful tool to aid ISPs
  - 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
23-01-21	861036	471276
24-01-21	860570	472295
25-01-21	860495	471659
26-01-21	860455	471018
27-01-21	859974	471324
28-01-21	860772	471640
29-01-21	860836	471620
30-01-21	861315	471906

Plot: BGP Table Size

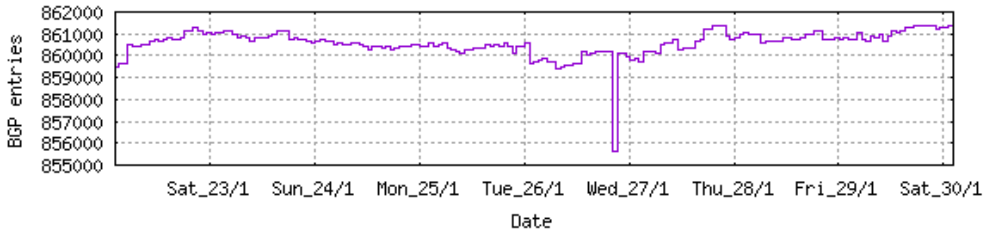

AS Summary

70632	Number of ASes in routing system
24860	Number of ASes announcing only one prefix
8180	Largest number of prefixes announced by an AS AS8151: Uninet S.A. de C.V., MX
116004352	Largest address span announced by an AS (/32s) AS4134: CHINANET-BACKBONE No.31,Jin-rong Street, CN

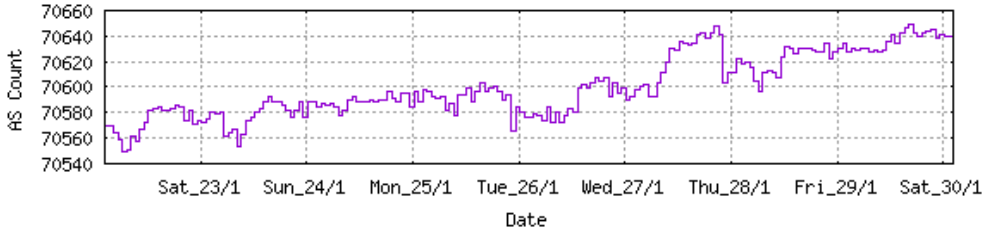

Plot: AS count
Plot: Average announcements per origin AS
Report: ASes ordered by originating address span
Report: ASes ordered by transit address span
Report: Autonomous System number-to-name mapping (from Registry WHOIS data)

Announced Prefixes

```
Rank  AS       Type    Originate Addr Space  (pfx)   Transit Addr space  (pfx)  Description
66    AS6389          ORG+TRN Originate:   10628608 /8.66   Transit:      60160 /16.12 BELLSOUTH-NET-BLK, 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 possibili

```
 Rank AS            AS Name                                  Current  Wthdw  Aggte  Annce Redctn       %
   74 AS6389        BELLSOUTH-NET-BLK, US                       1104    605     46    545    559  50.63%
```

```
  Prefix              AS Path                              Aggregation Suggestion
  12.81.120.0/24      4608 7575 2914 7018 6389
  12.130.209.0/24     4608 7575 6461 7018 6389 6389 6389 6389
  65.4.0.0/14         4608 7575 2914 7018 6389
  65.4.0.0/19         4608 7575 2914 7018 6389 - Withdrawn - matching aggregate 65.4.0.0/14 4608 7575 2914 7018 6389
  65.5.1.0/24         4608 7575 6461 7018 6389
  65.5.12.0/22        4608 7575 6461 7018 6389
  65.5.20.0/23        4608 7575 6461 7018 6389
  65.5.21.0/24        4608 7575 6461 7018 6389 - Withdrawn - matching aggregate 65.5.20.0/23 4608 7575 6461 7018 6389
  65.5.32.0/20        4608 7575 6461 7018 6389
  65.5.34.0/24        4608 4826 3257 7018 6389
  65.5.46.0/24        4608 7575 6461 7018 6389 - Withdrawn - matching aggregate 65.5.32.0/20 4608 7575 6461 7018 6389
  65.5.64.0/22        4608 7575 2914 7018 6389 - Withdrawn - matching aggregate 65.4.0.0/14 4608 7575 2914 7018 6389
  65.5.68.0/22        4608 7575 6461 7018 6389
  65.5.88.0/21        4608 7575 6461 7018 6389
  65.5.118.0/23       4608 7575 2914 7018 6389 - Withdrawn - matching aggregate 65.4.0.0/14 4608 7575 2914 7018 6389
  65.5.136.0/22       4608 7575 6461 7018 6389
  65.5.140.0/23       4608 7575 6461 7018 6389
  65.5.141.0/24       4608 7575 6461 7018 6389 - Withdrawn - matching aggregate 65.5.140.0/23 4608 7575 6461 7018 6389
  65.5.150.0/23       4608 7575 6461 7018 6389
  65.5.152.0/21       4608 7575 6461 7018 6389
  65.5.156.0/22       4608 7575 6461 7018 6389 - Withdrawn - matching aggregate 65.5.152.0/21 4608 7575 6461 7018 6389
  65.5.160.0/21       4608 7575 6461 7018 6389 + Announce - aggregate of 65.5.160.0/22 (4608 7575 6461 7018 6389) and 65.5.164.0/22 (4608 7575 6461 7018 6389)
  65.5.160.0/22       4608 7575 6461 7018 6389 - Withdrawn - aggregated with 65.5.164.0/22 (4608 7575 6461 7018 6389)
  65.5.164.0/22       4608 7575 6461 7018 6389 - Withdrawn - aggregated with 65.5.160.0/22 (4608 7575 6461 7018 6389)
  65.5.172.0/22       4608 7575 6461 7018 6389
  65.5.200.0/21       4608 7575 2914 7018 6389 - Withdrawn - matching aggregate 65.4.0.0/14 4608 7575 2914 7018 6389
  65.5.228.0/22       4608 7575 6461 7018 6389
  65.5.232.0/22       4608 7575 6461 7018 6389
  65.5.236.0/22       4608 7575 2914 7018 6389 - Withdrawn - matching aggregate 65.4.0.0/14 4608 7575 2914 7018 6389
  65.5.240.0/22       4608 7575 2914 7018 6389 - Withdrawn - matching aggregate 65.4.0.0/14 4608 7575 2914 7018 6389
  65.5.244.0/22       4608 7575 2914 7018 6389 - Withdrawn - matching aggregate 65.4.0.0/14 4608 7575 2914 7018 6389
  65.5.248.0/22       4608 7575 6461 7018 6389
  65.5.252.0/22       4608 7575 2914 7018 6389 - Withdrawn - matching aggregate 65.4.0.0/14 4608 7575 2914 7018 6389
  65.6.192.0/22       4608 7575 2914 7018 6389 - Withdrawn - matching aggregate 65.4.0.0/14 4608 7575 2914 7018 6389
  65.6.196.0/22       4608 7575 2914 7018 6389 - Withdrawn - matching aggregate 65.4.0.0/14 4608 7575 2914 7018 6389
  65.7.64.0/18        4608 7575 2914 7018 6389
  65.7.116.0/22       4608 4826 3257 7018 6389 + Announce - aggregate of 65.7.116.0/23 (4608 4826 3257 7018 6389) and 65.7.118.0/23 (4608 4826 3257 7018 6389)
```

Announced Prefixes

Rank	AS	Type	Originate Addr Space	(pfx)	Transit Addr space	(pfx)	Description
211	AS18566	ORG+TRN	Originate:	2838528 /10.56	Transit:	7936 /19.05	MEGAPATH5-, 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 possib

Rank	AS	AS Name	Current	Wthdw	Aggte	Annce	Redctn	%
38	AS18566	MEGAPATH5-, US	1998	1160	219	1057	941	47.10%

Prefix	AS Path	Aggregation Suggestion
64.6.160.0/23	4608 4826 3257 18566	
64.6.164.0/23	4608 4826 3257 18566	
64.6.166.0/23	4777 2516 3257 18566	+ Announce - aggregate of 64.6.166.0/24 (4777 2516 3257 18566) and 64.6.167.0/24 (4777 2516 3257 18566)
64.6.166.0/24	4777 2516 3257 18566	- Withdrawn - aggregated with 64.6.167.0/24 (4777 2516 3257 18566)
64.6.167.0/24	4777 2516 3257 18566	- Withdrawn - aggregated with 64.6.166.0/24 (4777 2516 3257 18566)
64.50.206.0/23	4777 2516 3257 18566	
64.51.126.0/23	4608 4826 3257 18566	
64.81.0.0/16	4777 2516 3356 18566	
64.81.4.0/24	4608 4826 3257 18566	
64.81.16.0/20	4777 2516 3257 18566	+ Announce - aggregate of 64.81.16.0/21 (4777 2516 3257 18566) and 64.81.24.0/21 (4777 2516 3257 18566)
64.81.16.0/22	4777 2516 3257 18566	- Withdrawn - aggregated with 64.81.20.0/22 (4777 2516 3257 18566)
64.81.20.0/22	4777 2516 3257 18566	- Withdrawn - aggregated with 64.81.16.0/22 (4777 2516 3257 18566)
64.81.22.0/24	4777 2516 3257 18566	- Withdrawn - matching aggregate 64.81.20.0/22 4777 2516 3257 18566
64.81.24.0/22	4777 2516 3257 18566	- Withdrawn - aggregated with 64.81.28.0/22 (4777 2516 3257 18566)
64.81.28.0/22	4777 2516 3257 18566	- Withdrawn - aggregated with 64.81.24.0/22 (4777 2516 3257 18566)
64.81.32.0/20	4608 4826 3257 18566	
64.81.32.0/24	4608 4826 3257 18566	- Withdrawn - matching aggregate 64.81.32.0/20 4608 4826 3257 18566
64.81.33.0/24	4608 4826 3257 18566	- Withdrawn - matching aggregate 64.81.32.0/20 4608 4826 3257 18566
64.81.34.0/24	4608 4826 3257 18566	- Withdrawn - matching aggregate 64.81.32.0/20 4608 4826 3257 18566
64.81.35.0/24	4608 4826 3257 18566	- Withdrawn - matching aggregate 64.81.32.0/20 4608 4826 3257 18566
64.81.36.0/24	4608 4826 3257 18566	- Withdrawn - matching aggregate 64.81.32.0/20 4608 4826 3257 18566
64.81.37.0/24	4608 4826 3257 18566	- Withdrawn - matching aggregate 64.81.32.0/20 4608 4826 3257 18566
64.81.38.0/24	4608 4826 3257 18566	- Withdrawn - matching aggregate 64.81.32.0/20 4608 4826 3257 18566
64.81.39.0/24	4608 4826 3257 18566	- Withdrawn - matching aggregate 64.81.32.0/20 4608 4826 3257 18566
64.81.40.0/24	4608 4826 3257 18566	- Withdrawn - matching aggregate 64.81.32.0/20 4608 4826 3257 18566
64.81.44.0/24	4608 4826 3257 18566	- Withdrawn - matching aggregate 64.81.32.0/20 4608 4826 3257 18566
64.81.48.0/20	4777 2516 3257 18566	
64.81.48.0/24	4777 2516 3257 18566	- Withdrawn - matching aggregate 64.81.48.0/20 4777 2516 3257 18566
64.81.50.0/24	4777 2516 3257 18566	- Withdrawn - matching aggregate 64.81.48.0/20 4777 2516 3257 18566
64.81.53.0/24	4777 2516 3257 18566	- Withdrawn - matching aggregate 64.81.48.0/20 4777 2516 3257 18566
64.81.54.0/24	4777 2516 3257 18566	- Withdrawn - matching aggregate 64.81.48.0/20 4777 2516 3257 18566
64.81.57.0/24	4777 2516 3257 18566	- Withdrawn - matching aggregate 64.81.48.0/20 4777 2516 3257 18566
64.81.60.0/24	4777 2516 3257 18566	- Withdrawn - matching aggregate 64.81.48.0/20 4777 2516 3257 18566
64.81.64.0/20	4777 2516 3257 18566	
64.81.66.0/24	4777 2516 3257 18566	- Withdrawn - matching aggregate 64.81.64.0/20 4777 2516 3257 18566
64.81.67.0/24	4777 2516 3257 18566	- Withdrawn - matching aggregate 64.81.64.0/20 4777 2516 3257 18566

Announced Prefixes

Rank	AS	Type	Originate Addr Space	(pfx)	Transit Addr space	(pfx)	Description
139	AS7545	ORG+TRN	Originate:	5040640 /9.73	Transit:	392448 /13.42	TPG-INTERNET-AP TPG Telecom Limited, AU

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 possi

Rank	AS	AS Name	Current	Wthdw	Aggte	Annce	Redctn	%
4	AS7545	TPG-INTERNET-AP TPG Telecom Limited, AU	5708	5138	141	711	4997	87.54%

Prefix	AS Path	Aggregation Suggestion
14.2.0.0/19	4608 4739 7545	
14.2.32.0/19	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
14.2.64.0/18	4608 4739 7545	+ Announce - aggregate of 14.2.64.0/19 (4608 4739 7545) and 14.2.96.0/19 (4608 4739 7545)
14.2.64.0/19	4608 4739 7545	- Withdrawn - aggregated with 14.2.96.0/19 (4608 4739 7545)
14.2.96.0/19	4608 4739 7545	- Withdrawn - aggregated with 14.2.64.0/19 (4608 4739 7545)
14.2.128.0/18	4608 7575 7545	
14.2.192.0/20	4608 4739 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
14.200.6.0/24	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
14.200.14.0/24	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
55	AS12479	ORG+TRN	Originate: 13956352 /8.27		Transit: 279808 /13.91		UNI2-AS, ES

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 possibi

Rank	AS	AS Name	Current	Wthdw	Aggte	Annce	Redctn	%
3	AS12479	UNI2-AS, ES	6260	5874	62	448	5812	92.84%

Prefix	AS Path	Aggregation Suggestion
1.178.224.0/19	4608 4826 5511 12479	
37.11.0.0/16	4608 4826 5511 12479	
37.11.0.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.4.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.8.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.12.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.16.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.20.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.24.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.28.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.32.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.36.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.40.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.44.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.48.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.56.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.64.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.68.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.72.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.76.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.80.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.84.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.88.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.92.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.96.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.100.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.104.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.108.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.112.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.116.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.120.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.124.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.128.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.132.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.136.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479
37.11.140.0/22	4608 4826 5511 12479	- Withdrawn - matching aggregate 37.11.0.0/16 4608 4826 5511 12479

Announced Prefixes

```
Rank  AS       Type    Originate Addr Space  (pfx)   Transit Addr space  (pfx)  Description
65    AS9829          ORG+TRN Originate:   10739456 /8.64   Transit:      34048 /16.94 BSNL-NIB National Internet Backbone, IN
```

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 possibi

```
 Rank AS            AS Name                                    Current  Wthdw  Aggte  Annce Redctn       %
   58 AS9829        BSNL-NIB National Internet Backbone, IN        1877    896    244   1225    652  34.74%


  Prefix              AS Path                              Aggregation Suggestion
  23.6.208.0/20       4608 9498 9829 9829 9829
  23.34.252.0/22      4608 7575 6461 6453 9829
  23.47.126.0/24      4777 2516 1273 9829
  23.48.225.0/24      4777 2516 174 9829
  23.60.169.0/24      4608 4826 1299 9829
  23.65.96.0/20       4608 9498 9829 9829 9829
  23.65.112.0/22      4777 2516 6453 4755 9829
  23.67.148.0/23      4777 2516 6453 4755 9829
  23.212.5.0/24       4777 2516 6453 4755 9829
  43.242.104.0/22     4608 24115 45489 9829
  59.88.0.0/13        4608 9498 9829 9829  + Announce - aggregate of 59.88.0.0/14 (4608 9498 9829 9829) and 59.92.0.0/14 (4608 9498 9829 9829)
  59.88.0.0/16        4608 9498 9829 9829  - Withdrawn - aggregated with 59.89.0.0/16 (4608 9498 9829 9829)
  59.88.0.0/20        4608 9498 9829 9829 9829
  59.88.16.0/20       4777 2516 6453 4755 9829
  59.88.32.0/20       4608 4826 3257 6453 9829
  59.88.48.0/20       4777 2516 6453 4755 9829
  59.88.64.0/20       4777 2516 174 9829
  59.88.96.0/20       4777 2516 6453 4755 9829
  59.88.112.0/20      4608 4826 3257 6453 9829
  59.88.128.0/20      4777 2516 174 9829
  59.88.144.0/20      4777 2516 6453 4755 9829
  59.88.160.0/20      4777 2516 1273 9829
  59.88.176.0/20      4777 2516 174 9829
  59.88.192.0/20      4777 2516 3491 9829
  59.88.208.0/20      4777 2516 6453 4755 9829
  59.88.224.0/19      4608 24115 45489 9829 + Announce - aggregate of 59.88.224.0/20 (4608 24115 45489 9829) and 59.88.240.0/20 (4608 24115 45489 9829)
  59.88.224.0/20      4608 24115 45489 9829 - Withdrawn - aggregated with 59.88.240.0/20 (4608 24115 45489 9829)
  59.88.240.0/20      4608 24115 45489 9829 - Withdrawn - aggregated with 59.88.224.0/20 (4608 24115 45489 9829)
  59.89.0.0/16        4608 9498 9829 9829  - Withdrawn - aggregated with 59.88.0.0/16 (4608 9498 9829 9829)
  59.89.0.0/20        4777 2516 6453 4755 9829
  59.89.16.0/20       4608 1221 4637 6453 9829
  59.89.32.0/20       4777 2516 174 9829
  59.89.48.0/20       4777 2516 6453 4755 9829
  59.89.64.0/20       4608 1221 4637 6453 9829
  59.89.80.0/20       4777 2516 6453 4755 9829
  59.89.96.0/19       4777 2516 6453 4755 9829 + Announce - aggregate of 59.89.96.0/20 (4777 2516 6453 4755 9829) and 59.89.112.0/20 (4777 2516 6453 4755 9829)
```

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
  - bgpupdates.potaroo.net/instability/bgpupd.html

The BGP Instability Report

The BGP Instability Report is updated daily. This report was generated on 29 January 2021 06:59 (UTC+1000)

50 Most active ASes for the past 14 days

RANK	ASN	UPDs	%	Prefixes	UPDs/Prefix	AS NAME
1	40065	638036	2.87%	549	1162.18	CNSERVERS, US
2	8551	272046	1.23%	3276	83.04	BEZEQ-INTERNATIONAL-AS Bezeqint Internet Backbone, IL
3	28573	212635	0.96%	1391	152.86	CLARO S.A., BR
4	11830	210091	0.95%	2565	81.91	Instituto Costarricense de Electricidad y Telecom., CR
5	34984	206378	0.93%	2179	94.71	TELLCOM-AS, TR
6	58224	187662	0.85%	1668	112.51	TCI, IR
7	23969	168716	0.76%	1950	86.52	TOT-NET TOT Public Company Limited, TH
8	7155	154008	0.69%	4038	38.14	VIASAT-SP-BACKBONE, US
9	33548	150008	0.68%	18	8333.78	UNWIRED-NOC, US
10	134548	146963	0.66%	1632	90.05	DXTL-HK DXTL Tseung Kwan O Service, HK
11	28202	145927	0.66%	804	181.50	Rede Brasileira de Comunicacao S/A, BR
12	9583	145626	0.66%	1709	85.21	SIFY-AS-IN Sify Limited, IN
13	16509	144426	0.65%	4786	30.18	AMAZON-02, US
14	9829	113600	0.51%	2043	55.60	BSNL-NIB National Internet Backbone, IN
15	11492	109931	0.50%	4793	22.94	CABLEONE, US
16	24000	106406	0.48%	98	1085.78	LIHGL-AS-AP 24.hk global BGP, HK
17	18004	84140	0.38%	167	503.83	WIRELESSNET-ID PT WIRELESS INDONESIA (WIN), ID
18	7011	83332	0.38%	529	157.53	FRONTIER-AND-CITIZENS, US
19	8764	80576	0.36%	622	129.54	TELIA-LIETUVA, LT
20	9009	74842	0.34%	2350	31.85	M247, GB
21	8881	69898	0.31%	446	156.72	VERSATEL, DE
22	12479	69841	0.31%	6338	11.02	UNI2-AS, ES
23	13904	69416	0.31%	173	401.25	COSLINK, US
24	36903	67580	0.30%	979	69.03	MT-MPLS, MA
25	8151	67026	0.30%	8181	8.19	Uninet S.A. de C.V., MX
26	7459	64950	0.29%	410	158.41	GRANDECOM-AS1, US
27	17794	64770	0.29%	14	4626.43	HTCL-ORANGE-HK-AP Hutchison Telephone Company Limited, HK

50 Most active Prefixes for the past 14 days

RANK	PREFIX	UPDs	%	Origin AS -- AS NAME
1	67.211.53.0/24	41496	0.18%	26405 -- HDCS, US
2	101.51.56.0/24	41328	0.18%	23969 -- TOT-NET TOT Public Company Limited, TH
3	66.152.121.0/24	39884	0.18%	14477 -- FLTG, US
4	216.238.254.0/23	37416	0.17%	13904 -- COSLINK, US
5	64.68.236.0/22	31818	0.14%	13904 -- COSLINK, US
6	162.217.44.0/22	30548	0.14%	13614 -- ALLWEST, US
7	193.104.187.0/24	30131	0.13%	212415 -- BRECHT-AS, DE
8	185.137.56.0/22	28530	0.13%	29119 -- SERVIHOSTING-AS AireNetworks, ES
9	140.17.208.0/24	26817	0.12%	7224 -- AMAZON-AS, US
10	64.192.0.0/24	26643	0.12%	33548 -- UNWIRED-NOC, US
11	64.192.1.0/24	26566	0.12%	33548 -- UNWIRED-NOC, US
12	64.192.2.0/24	26499	0.12%	33548 -- UNWIRED-NOC, US
13	64.192.4.0/24	26159	0.12%	33548 -- UNWIRED-NOC, US
14	64.192.3.0/24	25907	0.12%	33548 -- UNWIRED-NOC, US
15	185.174.212.0/23	22430	0.10%	29119 -- SERVIHOSTING-AS AireNetworks, ES
16	118.127.80.0/22	22423	0.10%	18117 -- HARBOURMSP-AU-AP NTT Communications ICT Solutions, AU
17	140.25.9.0/24	19951	0.09%	668 -- DNIC-AS-00668, US
18	63.92.224.0/20	19788	0.09%	714 -- APPLE-ENGINEERING, US
19	63.92.224.0/19	19719	0.09%	714 -- APPLE-ENGINEERING, US
20	103.214.100.0/24	18421	0.08%	64027 -- SMARTPLUS-AS-ID PT SURYA TEKNIKA PRATAMA, ID
21	103.214.101.0/24	18421	0.08%	64027 -- SMARTPLUS-AS-ID PT SURYA TEKNIKA PRATAMA, ID
22	64.192.5.0/24	18200	0.08%	33548 -- UNWIRED-NOC, US
23	99.194.200.0/22	17518	0.08%	22561 -- CENTURYLINK-LEGACY-LIGHTCORE, US
24	119.235.64.0/19	16422	0.07%	4638 -- IS-FJ-AS Telecom Fiji Limited, FJ
25	202.45.88.0/24	16187	0.07%	17794 -- HTCL-ORANGE-HK-AP Hutchison Telephone Company Limited, HK
26	203.145.78.0/24	16186	0.07%	17794 -- HTCL-ORANGE-HK-AP Hutchison Telephone Company Limited, HK
27	203.145.73.0/24	16181	0.07%	17794 -- HTCL-ORANGE-HK-AP Hutchison Telephone Company Limited, HK
28	203.145.74.0/24	16178	0.07%	17794 -- HTCL-ORANGE-HK-AP Hutchison Telephone Company Limited, HK
29	44.135.180.0/24	14930	0.07%	63479 -- HAMWAN, US
30	216.18.240.0/20	14372	0.06%	20101 -- MICHWAVE, US
31	192.206.204.0/24	13944	0.06%	53853 -- JAMA-ASN, US
32	199.192.180.0/22	13172	0.06%	20101 -- MICHWAVE, US

The BGP IPv6 Instability Report

This report is updated daily. The current report
was generated on 30 January 2021 01:32 (UTC+1000)

50 Most active ASes for the past 14 days

RANK	ASN	UPDs	%	Prefixes	UPDs/Prefix	AS NAME
1	11172	543493	5.40%	2807	193.62	Alestra, S. de R.L. de C.V., MX
2	36992	417345	4.15%	380	1098.28	ETISALAT-MISR, EG
3	18881	411438	4.09%	54	7619.22	TELEFONICA BRASIL S.A, BR
4	11664	354384	3.52%	229	1547.53	Techtel LMDS Comunicaciones Interactivas S.A., AR
5	27651	352515	3.50%	124	2842.86	ENTEL CHILE S.A., CL
6	26615	288941	2.87%	153	1888.50	TIM S/A, BR
7	270684	270457	2.69%	16	16903.56	ROGERIO SOUZA CORTES EIRELI, BR
8	28573	238459	2.37%	1400	170.33	CLARO S.A., BR
9	16509	214626	2.13%	1537	139.64	AMAZON-02, US
10	265594	189177	1.88%	175	1081.01	Television Internacional, S.A. de C.V., MX
11	28335	179875	1.79%	9	19986.11	MAXCOMM LTDA EPP, BR
12	61678	155730	1.55%	16	9733.12	NETWAY INFORMATICA LTDA, BR
13	270590	153387	1.52%	33	4648.09	M Andrade dos Santos, BR
14	7303	147889	1.47%	81	1825.79	Telecom Argentina S.A., AR
15	27951	129086	1.28%	46	2806.22	Media Commerce Partners S.A, CO
16	13999	120487	1.20%	106	1136.67	Mega Cable, S.A. de C.V., MX
17	205908	105046	1.04%	6	17507.67	PETERCXY-NETWORKS PeterCxy Networks, DE
18	28073	101410	1.01%	80	1267.62	Cooperativa Electrica Trenque Lauquen, AR
19	20940	92782	0.92%	673	137.86	AKAMAI-ASN1, NL
20	61932	86885	0.86%	4	21721.25	JCN COMUNICACOES E INFORMATICA LTDA ME, BR
21	52850	83571	0.83%	12	6964.25	Oxente.net Solucoes Tecnologicas Eireli, BR
22	210107	80015	0.80%	80	1000.19	PLUSWEB, TR
23	263252	78753	0.78%	6	13125.50	GIGASAT SERVICOS DE PROCESSAMENTOS DE DADOS LTDA, BR
24	50463	70426	0.70%	41	1717.71	TRIPLEC-ASN, IL
25	269575	69554	0.69%	17	4091.41	XIS 1 INTERNET FIBRA EIRELI, BR
26	19551	67557	0.67%	501	134.84	INCAPSULA, US
27	270252	67285	0.67%	17	3957.94	TVF INTERNET RAPIDA LTDA, BR

50 Most active Prefixes for the past 14 days

RANK	PREFIX	UPDs	%	Origin AS -- AS NAME
1	2405:4000:800:8::/64	47085	0.45%	38082 -- IIT-TIG-AS-AP True International Gateway Co., Ltd., TH
2	2400:dc40::/32	42773	0.41%	136440 -- SASPL-AS-AP Sungard Availability Services (India) Private Limited, IN
3	2a0e:b107:2f0::/48	35040	0.33%	205908 -- PETERCXY-NETWORKS PeterCxy Networks, DE
4	2a0e:b107:2f0::/44	35003	0.33%	205908 -- PETERCXY-NETWORKS PeterCxy Networks, DE
5	2a06:e881:2502::/48	34973	0.33%	205908 -- PETERCXY-NETWORKS PeterCxy Networks, DE
6	2a01:bfe0::/32	34571	0.33%	199997 -- IPPON_HOSTING-AS, FR
7	2804:371c:8000::/48	33775	0.32%	266390 -- Tajo Tecnologia Ltda, BR
8	2a00:9d20:5353::/48	29006	0.28%	60626 -- LEASEWEBCDN, NL
9	2a00:9d20:53::/48	28940	0.28%	60626 -- LEASEWEBCDN, NL
10	2402:53c0:2000::/36	27968	0.27%	137144 -- PASSLIR-AS Passit Media And Communication Pvt. Ltd., IN
11	2801:17:4800::/48	23798	0.23%	27951 -- Media Commerce Partners S.A, CO
12	2804:459c::/33	23188	0.22%	266938 -- LT SOLUCOES, BR
13	2804:459c:8000::/33	23128	0.22%	266938 -- LT SOLUCOES, BR
14	2804:db8:e000::/35	22699	0.22%	28335 -- MAXCOMM LTDA EPP, BR
15	2804:db8:8000::/35	22651	0.22%	28335 -- MAXCOMM LTDA EPP, BR
16	2804:db8::/35	22617	0.22%	28335 -- MAXCOMM LTDA EPP, BR
17	2804:db8:6000::/35	22577	0.21%	28335 -- MAXCOMM LTDA EPP, BR
18	2804:db8:2000::/35	22556	0.21%	28335 -- MAXCOMM LTDA EPP, BR
19	2801:1b6::/44	22504	0.21%	27951 -- Media Commerce Partners S.A, CO
20	2804:5f38::/32	22406	0.21%	269084 -- Global Tech Telecom Ltda - ME, BR
21	2801:174:3::/48	22403	0.21%	27951 -- Media Commerce Partners S.A, CO
22	2804:db8:a000::/35	22348	0.21%	28335 -- MAXCOMM LTDA EPP, BR
23	2804:db8:c000::/35	22216	0.21%	28335 -- MAXCOMM LTDA EPP, BR
24	2804:db8:4000::/35	22160	0.21%	28335 -- MAXCOMM LTDA EPP, BR
25	2804:1870:8000::/34	22160	0.21%	61932 -- JCN COMUNICACOES E INFORMATICA LTDA ME, BR
26	2801:172:2::/48	22127	0.21%	27951 -- Media Commerce Partners S.A, CO
27	2804:1b0:e000::/35	22078	0.21%	18881 -- TELEFONICA BRASIL S.A, BR
28	2804:5350::/32	22050	0.21%	268576 -- CONECT TELECOM LTDA - ME, BR
29	2804:1b2:d100::/46	22025	0.21%	18881 -- TELEFONICA BRASIL S.A, BR
30	2804:1b2:d000::/40	22019	0.21%	18881 -- TELEFONICA BRASIL S.A, BR
31	2801:80:350::/48	22004	0.21%	262524 -- COAMO AGROINDUSTRIAL COOPERATIVA, BR
32	2804:1b2:a000::/36	21746	0.21%	18881 -- TELEFONICA BRASIL S.A, BR

Receiving Prefixes

Receiving Prefixes

- There are three scenarios for receiving prefixes from other ASNs
  - 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

- ISPs should only accept prefixes which have been assigned or allocated to their downstream customer
- If ISP has assigned address space to its customer, then the customer IS entitled to announce it back to his ISP
- If the ISP 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
    - (jwhois is "joint whois" and queries all 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
geoloc:         27.4731138 153.0141194
country:        AU
admin-c:        AIC1-AP
tech-c:         AIC1-AP
mnt-by:         APNIC-HM
mnt-irt:        IRT-APNIC-IS-AP
status:         ASSIGNED PORTABLE
changed:        hm-changed@apnic.net 20170327
changed:        hm-changed@apnic.net 20170331
source:         APNIC
```

inetnum – means it is an address delegation to an entity

Portable – 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 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 ISP 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 an ISP 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/irrtoolset/irrtoolset
    - bgpq3:
      https://github.com/snar/bgpq3

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 an ISP 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

- 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:
  - 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"
  - http://www.team-cymru.com/bogon-reference.html
- For IPv4 also consult:
  - https://www.rfc-editor.org/rfc/rfc6441.txt (BCP171)
- For IPv6 also consult:
  - http://www.space.net/~gert/RIPE/ipv6-filters.html
- Bogon Route Server:
  - https://www.team-cymru.com/bogon-reference-bgp.html
  - 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
ipv6 prefix-list v6in-filter deny 2001:10::/28 le 128     ! ORCHID
ipv6 prefix-list v6in-filter deny 2001:db8::/32 le 128    ! Documentation
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 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 ISPs to be good Internet citizens

Prefixes into IBGP

Injecting prefixes into IBGP

- Use IBGP to carry customer prefixes
  - Don’t use IGP
- Point static route to customer interface
- 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 Serial 5/0
 ip unnumbered loopback 0
 ip verify unicast reverse-path
!
ip route 100.71.10.0 255.255.252.0 Serial 5/0
!
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 ISPs 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 Serial 5/0
!
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

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

Types of Peering (4)

- The Peering Database documents ISPs peering policies
  - https://www.peeringdb.com
- All operators of ASNs should 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 (European, Middle Eastern, Asian, Caribbean, Latin American)
  - 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

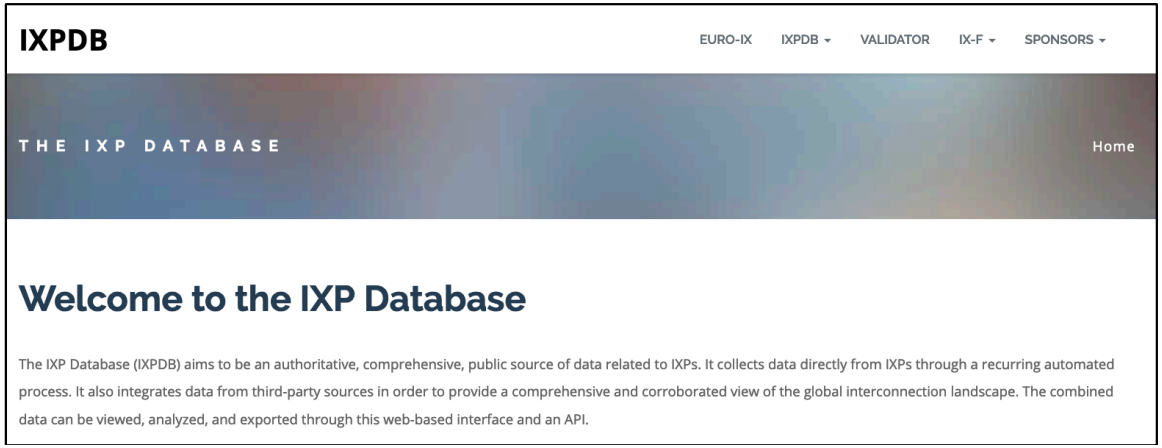

PeeringDB

Search here for a network, IX, or facility.

Advanced Search

pfsinoz

HKIX

Organization	Hong Kong Internet eXchange Limited
Long Name	Hong Kong Internet Exchange
City	Hong Kong
Country	HK
Continental Region	Asia Pacific
Media Type	Ethernet
Protocols Supported	✓ Unicast IPv4 ○ Multicast ✓ IPv6
Notes	

Contact Information

Company Website	https://www.hkix.net/
Traffic Stats Website	https://www.hkix.net/hkix/stat/aggt/hkix-aggregate.html
Technical Email	noc@hkix.net
Technical Phone	+85239439900
Policy Email	info@hkix.net
Policy Phone	+85239438800

LAN

MTU	1500
DOT1Q	○
IPv6	2001:7fa:0:1::/64
IPv4	123.255.88.0/21

Local Facilities

Filter

Facility	Country	City
CUHK	Hong Kong	Hong Kong
MEGA Two (iAdvantage Hong Kong)	Hong Kong	Hong Kong
MEGA-i (iAdvantage Hong Kong)	Hong Kong	Hong Kong

Peers at this Exchange Point

Filter

Peer Name ASN	IPv4 IPv6	Speed Policy
ASGCNET HKIX Peering LAN 24167	123.255.91.53 2001:7fa:0:1::ca28:a135	10G Open
Asia Pacific Telecom HKIX Peering LAN 17709	123.255.91.86 2001:7fa:0:1::ca28:a156	10G Open
ASLINE HKIX Peering LAN 18013	123.255.92.13 2001:7fa:0:1::ca28:a20d	10G Open
AT&T AP - AS2687 HKIX Peering LAN 2687	123.255.91.46 2001:7fa:0:1::ca28:a12e	10G Selective
Automattic HKIX Peering LAN 2635	123.255.90.71 2001:7fa:0:1::ca28:a047	10G Open
Badoo Ltd HKIX Peering LAN 12678	123.255.90.220 None	2G Open
Baidu HKIX Peering LAN 55967	123.255.90.131 2001:7fa:0:1::ca28:a083	10G Open
Baidu HKIX Peering LAN 55967	123.255.91.61 2001:7fa:0:1::ca28:a13d	10G Open
Bayan Telecommunications Inc. HKIX Peering LAN 6648	123.255.91.45 2001:7fa:0:1::ca28:a12d	3G Open
BGP Network Limited HKIX Peering LAN 64050	123.255.91.177 2001:7fa:0:1::ca28:a1b1	100G Open
BIGHUB-ISP HKIX Peering LAN 137989	123.255.90.207 2001:7fa:0:1::ca28:a0cf	1G Open
BIGHUB-ISP HKIX Peering LAN	123.255.91.98	10G

PeeringDB

Search here for a network, IX, or facility.

Advanced Search

pfsinoz

Amazon.com Diamond Sponsor

Organization	Amazon.com
Also Known As	Amazon Web Services
Company Website	http://www.amazon.com
Primary ASN	16509
IRR as-set/route-set	AS-AMAZON
Route Server URL	
Looking Glass URL	
Network Type	Enterprise
IPv4 Prefixes	5000
IPv6 Prefixes	2000
Traffic Levels	Not Disclosed
Traffic Ratios	Balanced
Geographic Scope	Global
Protocols Supported	✓ Unicast IPv4 ○ Multicast ✓ IPv6 ○ Never via route servers
Last Updated	2019-12-29T14:56:38Z
Notes	If you have a connectivity issue to Amazon then please visit: • IPv4: http://ec2-reachability.amazonaws.com/ • IPv6: http://ipv6.ec2-reachability.amazonaws.com/ And include detail on prefixes you think you have a problem with if you contact our Ops alias. This will reduce time with troubleshooting. The following Amazon US locations and associated IX's carry routes/traffic specific only to the services with infrastructure in that metro. For example, Jacksonville is CloudFront only, whereas Ashburn is CloudFront, EC2, S3, etc.) • Seattle • Palo Alto • San Jose • Los Angeles • Dallas

Public Peering Exchange Points

Filter

Exchange ASN	IPv4 IPv6	Speed RS Peer
AMS-IX 16509	80.249.210.100 2001:7f8:1::a501:6509:1	400G ○
AMS-IX 16509	80.249.210.217 2001:7f8:1::a501:6509:2	400G ○
AMS-IX Chicago 16509	206.108.115.36 2001:504:38:1:0:a501:6509:1	100G ○
AMS-IX Hong Kong 16509	103.247.139.10 2001:df0:296::a501:6509:1	100G ○
AMS-IX India 16509	223.31.200.29 2001:e48:44:100b:0:a501:6509:2	10G ○
AMS-IX India 16509	223.31.200.30 2001:e48:44:100b:0:a501:6509:1	10G ○
BBIX Osaka 16509	218.100.9.24 2001:de8:c:2:0:1:6509:1	40G ○
BBIX Tokyo 16509	218.100.6.52 2001:de8:c::1:6509:1	200G ○
BBIX Tokyo 16509	218.100.6.207 2001:de8:c::1:6509:2	200G ○
BCIX BCIX Peering LAN 16509	193.178.185.95 2001:7f8:19:1::407d:1	200G ○
BIX.BG Main 16509	193.169.198.87 2001:7f8:58::407d:0:1	100G ○
RNIX	194.53.172.122	100G

Private Peering Facilities

Filter

Facility ASN	Country City
151 Front Street West Toronto 16509	Canada Toronto

PeeringDB

Search here for a network, IX, or facility.

Advanced Search

pfsinoz

Telia Carrier

Organization	Telia Group
Also Known As	TeliaSonera, Telia, TSIC
Company Website	http://www.teliacarrier.com/
Primary ASN	1299
IRR as-set/route-set	RIPE::AS-TELIANET RIPE::AS-TELIANET-V6
Route Server URL	
Looking Glass URL	https://lg.telia.net/
Network Type	NSP
IPv4 Prefixes	426000
IPv6 Prefixes	40000
Traffic Levels	1 Tbps+
Traffic Ratios	Balanced
Geographic Scope	Global
Protocols Supported	⊘ Unicast IPv4 ◯ Multicast ⊘ IPv6 ◯ Never via route servers
Last Updated	2020-02-05T11:43:25Z
Notes	IPv4 + IPv6 Prefixes above would be actuals, not proposed max- prefix values. 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. All trouble ticket requests or support related emails should be sent to carrier-csc@teliacompany.com.

Peering Policy Information

Peering Policy	https://www.teliacarrier.com/dam/jcr:d1e83942-3db1-4334-a5f8-431578633d26/Telia_Carrier_Global_Peering_Policy.pdf
General Policy	Restrictive

Public Peering Exchange Points

Filter

Exchange ▾ / ASN	IPv4 / IPv6	Speed / RS Peer

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

Private Peering Facilities

Filter

Facility ▾ / ASN	Country / City
365 Data Centers Buffalo (BU1)	United States of America
1299	Buffalo
365 Data Centers Detroit (DT1)	United States of America
1299	Southfield
365 Data Centers Nashville (NA1)	United States of America
1299	Nashville
365 Data Centers Tampa (TA1)	United States of America
1299	Tampa
3U Rechenzentrum Berlin	Germany
1299	Berlin
Altus IT	Croatia
1299	Zagreb
Borovaya 57	Russia
1299	St. Petersburg
CE Colo Prague	Czechia
1299	Prague
CINECA - DC NaMeX	Italy
1299	Roma
COD BM-18	Russia
1299	St.Petersburg
Caldera21	Italy
1299	Milan
CarrierColo Berlin Luetzow (I/P/B/ site B)	Germany
1299	Berlin
Cologix MTL3	Canada
1299	Montreal

Screenshot

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 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 later in the presentation

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 instance of the IRR as well
    - 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:         AS38004
members:         AS132232
members:         AS134715
members:         AS135666
members:         AS137925
members:         AS59219
members:         AS18024
members:         AS18025
members:         AS137994
admin-c:         DNO1-AP
tech-c:          DNO1-AP
notify:          netops@bt.bt
mnt-by:          MAINT-BT-DRUKNET
last-modified:   2019-01-15T08:51:21Z
source:          APNIC
```

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

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

Route Origin Authorisation

- Essential first step to secure the global routing system
- Covered in detail in separate presentation slide deck:
  - http://www.bgp4all.com.au/pfs/_media/workshops/02-rpki.pdf

Configuration Tips

Of passwords, tricks and templates

IBGP and IGPs Reminder!

- Make sure loopback is configured on router
  - IBGP between loopbacks, NOT real 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 ISPs 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 ISPs 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 ASes in a typical AS_PATH in the Internet 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 50404
50404 i
```

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

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

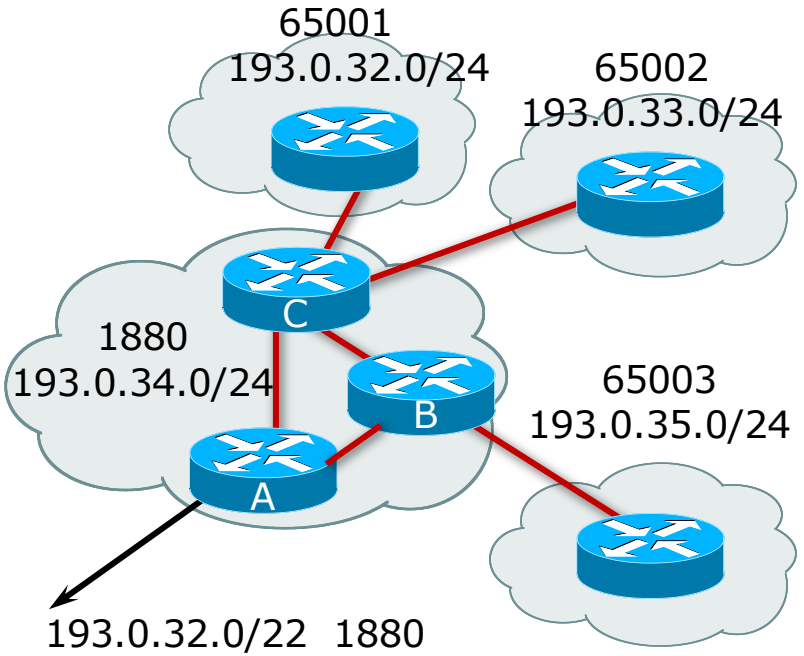

Private-AS – Removal

- 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

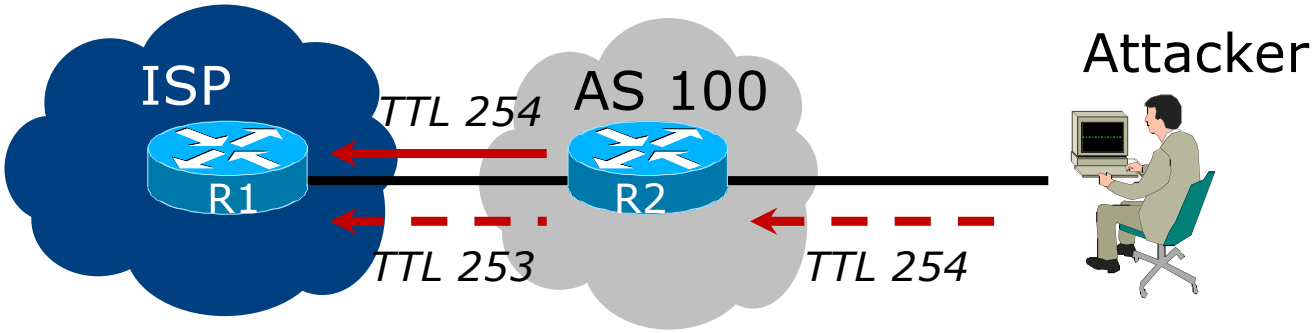

BGP TTL "hack"

- 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
...
Mininum 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 ISPs
  - ISPs create their own defaults by using configuration templates
- EBGP and IBGP examples follow
  - Also see Team Cymru's BGP templates
    - http://www.team-cymru.com/community-services.html

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 – it will consume considerable amounts of extra memory for BGP
- Remove private ASes 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

Routing Security

- Implement the recommendations in https://www.manrs.org/manrs

1. Prevent propagation of incorrect routing information
   - Filter BGP peers, in & out!
2. Prevent traffic with spoofed source addresses
   - BCP38 – Unicast Reverse Path Forwarding
3. Facilitate communication between network operators
   - NOC to NOC Communication
   - Up-to-date details in Route and AS Objects, and PeeringDB
4. Facilitate validation of routing information
   - Route Origin Authorisation using RPKI

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
  - https://www.manrs.org/join/
  - MANRS for Operators
  - MANRS for IXPs

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

BGP Best Current Practices

ISP Workshops