

BGP Best Current Practices

ISP Workshops



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

Last updated 14th November 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

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

The screenshot displays the NSRC (Network Startup Resource Center) website. The navigation bar includes links for Home, About, BGP for All (highlighted), perfSONAR, ScienceDMZ, FedIdM, and Contact Us, along with a search bar. The main content area is divided into three columns:

- BGP for All:** A text-based introduction to BGP, explaining its role as the primary routing protocol for the Internet and autonomous systems. It also mentions that understanding routing options can lead to efficiencies and collaboration opportunities.
- 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.

On the right side, there is a video player for "BGP for All" with a play button and a "Watch on YouTube" button. Below the video player, there are sections for "BGP Case Studies" and "Communities".

BGP Case Studies:

- 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}
- Traffic Engineering and CDNs ^{NEW}

Communities:

- Communities: RFC 1998 Traffic Engineering
- Communities: Simplifying Traffic Engineering
- How to Apply Communities to Originated Routes
- How to Use Communities for Service Identification

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
```

EBGP Default Behaviour

- FRR turns on RFC8212 support 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

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
 - October 2021: 506533 /24s in IPv4 table of 868499 prefixes
- ❑ **The same is happening for /48s with IPv6**
 - October 2021: 62628 /48s in IPv6 table of 133799 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 (October 2021)

□ Current IPv4 Internet Routing Table Statistics

| | |
|------------------------------------|--------|
| BGP Routing Table Entries | 868499 |
| Prefixes after maximum aggregation | 329923 |
| Unique prefixes in Internet | 419756 |
| /24s announced | 506533 |
| ASNs in use | 72121 |

- (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 |
|----------|----------|-----------------|
| 18-10-21 | 898602 | 496196 |
| 19-10-21 | 898535 | 496201 |
| 20-10-21 | 897984 | 483474 |
| 21-10-21 | 899215 | 484272 |
| 22-10-21 | 899051 | 484837 |
| 23-10-21 | 899713 | 485141 |
| 24-10-21 | 899753 | 485867 |
| 25-10-21 | 900089 | 485932 |

Plot: [BGP Table Size](#)



AS Summary

| | |
|-----------|---|
| 72368 | Number of ASes in routing system |
| 25298 | Number of ASes announcing only one prefix |
| 8896 | Largest number of prefixes announced by an AS AS8151 : Uninet S.A. de C.V., MX |
| 211326208 | Largest address span announced by an AS (/32s) AS749 : DNIC-AS-00749, 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 |
|------|--------|------|--------------------|------------------|---------|------------------|------------------------------|
| 74 | AS6389 | | ORG+TRN Originate: | 9726464 | /8.79 | 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 aggreg

| Rank | AS | AS Name | Current | Withdw | Aggte | Annce | Redctn | % |
|------|------------------------|-----------------------|---------|--------|-------|-------|--------|--------|
| 84 | AS6389 | BELLSOUTH-NET-BLK, US | 703 | 586 | 24 | 141 | 562 | 79.94% |

| Prefix | AS Path | Aggregation Suggestion |
|-----------------|--------------------------|---|
| 12.81.120.0/24 | 4608 7575 2914 7018 6389 | |
| 12.130.209.0/24 | 4608 7575 2914 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 |
| 65.5.64.0/22 | 4608 7575 2914 7018 6389 | - Withdrawn - matching aggregate 65.4.0.0/14 4608 7575 2914 7018 6389 |
| 65.5.88.0/21 | 4608 7575 2914 7018 6389 | - Withdrawn - matching aggregate 65.4.0.0/14 4608 7575 2914 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 |
| 65.5.160.0/22 | 4608 7575 2914 7018 6389 | - Withdrawn - matching aggregate 65.4.0.0/14 4608 7575 2914 7018 6389 |
| 65.5.164.0/22 | 4608 7575 2914 7018 6389 | - Withdrawn - matching aggregate 65.4.0.0/14 4608 7575 2914 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 2914 7018 6389 | - Withdrawn - matching aggregate 65.4.0.0/14 4608 7575 2914 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 2914 7018 6389 | - Withdrawn - matching aggregate 65.4.0.0/14 4608 7575 2914 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 ²⁹ (4608 4826 3257 7018 6389) |
| 65.7.116.0/24 | 4608 4826 3257 7018 6389 | - Withdrawn - aggregated with 65.7.117.0/24 (4608 4826 3257 7018 6389) |
| 65.7.117.0/24 | 4608 4826 3257 7018 6389 | - Withdrawn - aggregated with 65.7.116.0/24 (4608 4826 3257 7018 6389) |

Long term deaggregator - BellSouth in the US

Announced Prefixes

| Rank | AS | Type | Originate | Addr Space (pfx) | Transit | Addr space (pfx) | Description |
|------|---------|--------|------------|------------------|---------|------------------|------------------------|
| 211 | AS18566 | ORIGIN | Originate: | 2825216 | /10.57 | Transit: | 0 /0.00 MEGAPATH5-, US |

Aggregation Suggestions

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

Long term deaggregator –
Megapath in the US

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 aggreg

| Rank | AS | AS Name | Current | Wthdw | Aggte | Annce | Redctn | % |
|------|-------------------------|----------------|---------|-------|-------|-------|--------|--------|
| 25 | AS18566 | MEGAPATH5-, US | 1963 | 1619 | 120 | 464 | 1499 | 76.36% |

| Prefix | AS Path | Aggregation Suggestion |
|----------------|----------------------|---|
| 64.6.160.0/23 | 4608 4826 3257 18566 | |
| 64.6.164.0/22 | 4608 4826 3257 18566 | + Announce - aggregate of 64.6.164.0/23 (4608 4826 3257 18566) and 64.6.166.0/23 (4608 4826 3257 18 |
| 64.6.164.0/23 | 4608 4826 3257 18566 | - Withdrawn - aggregated with 64.6.166.0/23 (4608 4826 3257 18566) |
| 64.6.166.0/24 | 4608 4826 3257 18566 | - Withdrawn - aggregated with 64.6.167.0/24 (4608 4826 3257 18566) |
| 64.6.167.0/24 | 4608 4826 3257 18566 | - Withdrawn - aggregated with 64.6.166.0/24 (4608 4826 3257 18566) |
| 64.50.206.0/23 | 4608 4826 3257 18566 | |
| 64.51.126.0/23 | 4608 4826 3257 18566 | |
| 64.81.0.0/16 | 4608 4826 3356 18566 | |
| 64.81.4.0/24 | 4608 4826 3257 18566 | |
| 64.81.16.0/20 | 4608 4826 3257 18566 | + Announce - aggregate of 64.81.16.0/21 (4608 4826 3257 18566) and 64.81.24.0/21 (4608 4826 3257 18 |
| 64.81.16.0/22 | 4608 4826 3257 18566 | - Withdrawn - aggregated with 64.81.20.0/22 (4608 4826 3257 18566) |
| 64.81.20.0/22 | 4608 4826 3257 18566 | - Withdrawn - aggregated with 64.81.16.0/22 (4608 4826 3257 18566) |
| 64.81.22.0/24 | 4608 4826 3257 18566 | - Withdrawn - matching aggregate 64.81.20.0/22 4608 4826 3257 18566 |
| 64.81.24.0/22 | 4608 4826 3257 18566 | - Withdrawn - aggregated with 64.81.28.0/22 (4608 4826 3257 18566) |
| 64.81.28.0/22 | 4608 4826 3257 18566 | - Withdrawn - aggregated with 64.81.24.0/22 (4608 4826 3257 18566) |
| 64.81.32.0/19 | 4608 4826 3257 18566 | + Announce - aggregate of 64.81.32.0/20 (4608 4826 3257 18566) and 64.81.48.0/20 (4608 4826 3257 18 |
| 64.81.32.0/20 | 4608 4826 3257 18566 | - Withdrawn - aggregated with 64.81.48.0/20 (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 |
| 64.81.39.0/24 | 4608 4826 3257 18566 | - Withdrawn - matchinq aqqreqate 64.81.32.0/20 4608 4826 3257 18566 |

Announced Prefixes

| Rank | AS | Type | Originate | Addr Space (pfx) | Transit | Addr space (pfx) | Description |
|------|--------|------|--------------------|------------------|---------|------------------|--|
| 141 | AS7545 | | ORG+TRN Originate: | 5187072 | /9.69 | Transit: 3157248 | /10.41 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 aggreg:

| Rank | AS | AS Name | Current | Wthdw | Aggte | Annce | Redctn | % |
|------|------------------------|---|---------|-------|-------|-------|--------|--------|
| 6 | AS7545 | TPG-INTERNET-AP TPG Telecom Limited, AU | 5832 | 4985 | 179 | 1026 | 4806 | 82.41% |

| 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 |
| 14.200.7.0/24 | 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 |

Long term deaggregator –
TPG in Australia

Announced Prefixes

| Rank | AS | Type | Originate | Addr Space (pfx) | Transit | Addr space (pfx) | Description |
|------|---------|------|--------------------|------------------|----------|------------------|-------------|
| 54 | AS12479 | | ORG+TRN Originate: | 14209792 /8.24 | Transit: | 282112 /13.89 | UNI2-AS, ES |

Aggregation Suggestions

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

Long term deaggregator – Orange in Spain

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 aggreg:

| Rank | AS | AS Name | Current | Wthdw | Aggte | Annce | Redctn | % |
|------|-------------------------|-------------|---------|-------|-------|-------|--------|--------|
| 4 | AS12479 | UNI2-AS, ES | 6864 | 6410 | 75 | 529 | 6335 | 92.29% |

| Prefix | AS Path | Aggregation Suggestion |
|-----------------------|----------------------|--|
| 1.178.224.0/19 | 4608 4826 5511 12479 | |
| 1.178.248.0/21 | 4608 4826 5511 12479 | - Withdrawn - matching aggregate 1.178.224.0/19 4608 4826 5511 12479 |
| 37.11.0.0/16 | 4608 4826 5511 12479 | |
| 37.11.0.0/22 | 4608 7575 3356 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 |
| 37.11.52.0/22 | 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 |
| 37.11.60.0/22 | 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 |

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

50 Most active ASes for the past 14 days

| RANK | ASN | UPDs | % | Prefixes | UPDs/Prefix | AS NAME |
|------|--------|--------|-------|----------|-------------|--|
| 1 | 64050 | 333079 | 2.54% | 874 | 381.10 | BCPL-SG BGPNET Global ASN, SG |
| 2 | 8151 | 242971 | 1.86% | 8932 | 27.20 | Uninet S.A. de C.V., MX |
| 3 | 16509 | 180063 | 1.38% | 5880 | 30.62 | AMAZON-02, US |
| 4 | 7713 | 161053 | 1.23% | 3526 | 45.68 | TELKOMNET-AS-AP PT Telekomunikasi Indonesia, ID |
| 5 | 38197 | 128878 | 0.98% | 257 | 501.47 | SUNHK-DATA-AS-AP Sun Network Hong Kong Limited - HongKong Backbone, HK |
| 6 | 132839 | 118902 | 0.91% | 404 | 294.31 | POWERLINE-AS-AP POWER LINE DATACENTER, HK |
| 7 | 5972 | 117090 | 0.89% | 2170 | 53.96 | DNIC-ASBLK-05800-06055, US |
| 8 | 37133 | 102803 | 0.79% | 72 | 1427.82 | airtel-tz-as, TZ |
| 9 | 174 | 90268 | 0.69% | 3086 | 29.25 | COGENT-174, US |
| 10 | 58224 | 85515 | 0.65% | 1701 | 50.27 | TCI, IR |
| 11 | 47331 | 82797 | 0.63% | 8393 | 9.87 | TTNET, TR |
| 12 | 7155 | 80461 | 0.61% | 3951 | 20.36 | VIASAT-SP-BACKBONE, US |
| 13 | 9829 | 76048 | 0.58% | 1863 | 40.82 | BSNL-NIB National Internet Backbone, IN |
| 14 | 7552 | 70613 | 0.54% | 3627 | 19.47 | VIETEL-AS-AP Viettel Group, VN |
| 15 | 17794 | 66389 | 0.51% | 14 | 4742.07 | HTCL-ORANGE-HK-AP Hutchison Telephone Company Limited, HK |
| 16 | 9583 | 66216 | 0.51% | 1746 | 37.92 | SIFY-AS-IN Sify Limited, IN |
| 17 | 36903 | 62936 | 0.48% | 1183 | 53.20 | MT-MPLS, MA |
| 18 | 27651 | 56173 | 0.43% | 837 | 67.11 | ENTEL CHILE S.A., CL |
| 19 | 42337 | 52155 | 0.40% | 412 | 126.59 | RESPINA-AS, IR |
| 20 | 8551 | 49632 | 0.38% | 3446 | 14.40 | BEZEQ-INTERNATIONAL-AS Bezeqint Internet Backbone, IL |
| 21 | 13188 | 49254 | 0.38% | 1611 | 30.57 | TRIOLAN, UA |
| 22 | 9919 | 46780 | 0.36% | 738 | 63.39 | NCIC-TW New Century InfoComm Tech Co., Ltd., TW |
| 23 | 4780 | 45130 | 0.34% | 716 | 63.03 | SEEDNET Digital United Inc., TW |
| 24 | 47730 | 44669 | 0.34% | 2 | 22334.50 | JEANDELESTRE, FR |
| 25 | 9785 | 44398 | 0.34% | 77 | 576.60 | JASATELNET-AS-ID PT Berca Hardayaperkasa, ID |
| 26 | 27871 | 44177 | 0.34% | 10 | 4417.70 | Telecom Argentina S.A., AR |
| 27 | 24835 | 42970 | 0.33% | 1571 | 27.35 | RAYA-AS, EG |

50 Most active Prefixes for the past 14 days

| RANK | PREFIX | UPDs | % | Origin AS -- AS NAME |
|------|------------------|-------|-------|--|
| 1 | 103.79.118.0/24 | 37670 | 0.28% | 135490 -- BAL-AS-AP Business Automation Ltd., BD |
| 2 | 197.186.0.0/15 | 36184 | 0.27% | 37133 -- airtel-tz-as, TZ |
| 3 | 145.236.90.0/24 | 35021 | 0.26% | 5483 -- MAGYAR-TELEKOM-MAIN-AS Magyar Telekom Nyrt., HU |
| 4 | 196.46.120.0/23 | 33120 | 0.25% | 37133 -- airtel-tz-as, TZ |
| 5 | 41.75.208.0/20 | 33012 | 0.25% | 37133 -- airtel-tz-as, TZ |
| 6 | 67.211.53.0/24 | 29091 | 0.22% | 26405 -- HDCS, US |
| 7 | 64.68.236.0/22 | 24129 | 0.18% | 13904 -- COSLINK, US |
| 8 | 45.148.160.0/24 | 22758 | 0.17% | 47730 -- JEANDELESTRE, FR |
| 9 | 203.145.78.0/24 | 22140 | 0.16% | 17794 -- HTCL-ORANGE-HK-AP Hutchison Telephone Company Limited, HK |
| 10 | 203.145.74.0/24 | 22128 | 0.16% | 17794 -- HTCL-ORANGE-HK-AP Hutchison Telephone Company Limited, HK |
| 11 | 202.45.88.0/24 | 22121 | 0.16% | 17794 -- HTCL-ORANGE-HK-AP Hutchison Telephone Company Limited, HK |
| 12 | 45.148.160.0/22 | 21911 | 0.16% | 47730 -- JEANDELESTRE, FR |
| 13 | 138.207.66.0/24 | 21084 | 0.16% | 12025 -- IMDC-AS12025, US |
| 14 | 138.207.67.0/24 | 20966 | 0.16% | 12025 -- IMDC-AS12025, US |
| 15 | 170.79.222.0/23 | 18766 | 0.14% | 263864 -- CLIC RAPIDO TELECOMUNICACAO LTDA, BR |
| 16 | 41.242.72.0/24 | 15666 | 0.12% | 37605 -- GTNL1-AS, NG |
| 17 | 200.110.229.0/24 | 15160 | 0.11% | 27871 -- Telecom Argentina S.A., AR |
| 18 | 200.110.228.0/24 | 15155 | 0.11% | 27871 -- Telecom Argentina S.A., AR |
| 19 | 209.22.66.0/24 | 14694 | 0.11% | 2046 -- DNIC-AS-02046, US |
| 20 | 209.22.67.0/24 | 14692 | 0.11% | 2046 -- DNIC-AS-02046, US |
| 21 | 141.98.139.0/24 | 14078 | 0.10% | 211975 -- WOHLERT, DE |
| 22 | 131.0.4.0/22 | 13587 | 0.10% | 61855 -- NOVA NET TECNOLOGIA LTDA - ME, BR |
| 23 | 216.194.128.0/21 | 13510 | 0.10% | 27871 -- Telecom Argentina S.A., AR |
| 24 | 182.23.171.0/24 | 11550 | 0.09% | 137346 -- CNI-AS-ID PT Cyber Network Indonesia, ID |
| 25 | 45.169.136.0/22 | 10593 | 0.08% | 268117 -- RAFAEL FERNANDES DE MEDEIROS, BR |
| 26 | 200.36.204.0/24 | 9530 | 0.07% | 271284 -- MP TELECOMUNICACOES EIRELI, BR |
| 27 | 130.137.78.0/24 | 9313 | 0.07% | 16509 -- AMAZON-02, US |
| 28 | 130.137.79.0/24 | 9271 | 0.07% | 16509 -- AMAZON-02, US |
| 29 | 130.137.108.0/24 | 9214 | 0.07% | 16509 -- AMAZON-02, US |
| 30 | 130.137.140.0/24 | 9186 | 0.07% | 16509 -- AMAZON-02, US |
| 31 | 130.137.90.0/24 | 9129 | 0.07% | 16509 -- AMAZON-02, US |
| 32 | 130.137.89.0/24 | 9125 | 0.07% | 14618 -- AMAZON-AES, US |

The BGP IPv6 Instability Report

50 Most active ASes for the past 14 days

| RANK | ASN | UPDs | % Prefixes | UPDs/Prefix | AS NAME |
|------|------------------------|---------|------------|-------------|--|
| 1 | 58336 | 3387439 | 58.49% | 304 | 11142.89 IXTS-AS, CN |
| 2 | 20473 | 827896 | 14.30% | 7040 | 117.60 AS-CHOOPA, US |
| 3 | 27951 | 166087 | 2.87% | 58 | 2863.57 Media Commerce Partners S.A., CO |
| 4 | 11172 | 87338 | 1.51% | 3385 | 25.80 Alestra, S. de R.L. de C.V., MX |
| 5 | 12208 | 58796 | 1.02% | 18 | 3266.44 TRUVISTA, US |
| 6 | 262983 | 44153 | 0.76% | 21 | 2102.52 Net Barretos Tecnologia LTDA - ME, BR |
| 7 | 15133 | 42692 | 0.74% | 404 | 105.67 EDGECAST, US |
| 8 | 31514 | 39061 | 0.67% | 2 | 19530.50 INF-NET-AS, RU |
| 9 | 262742 | 38760 | 0.67% | 14 | 2768.57 Fundacao Universidade Federal do ABC - UFABC, BR |
| 10 | 136440 | 36871 | 0.64% | 1 | 36871.00 SASPL-AS-AP Sungard Availability Services India Private Limited, IN |
| 11 | 36223 | 31094 | 0.54% | 9 | 3454.89 SPANISHFORK-COMMUNITY-NETWORK, US |
| 12 | 22356 | 29089 | 0.50% | 5 | 5817.80 Durand do Brasil Ltda, BR |
| 13 | 60982 | 28978 | 0.50% | 3 | 9659.33 WARPCON-AS, RO |
| 14 | 208753 | 28765 | 0.50% | 36 | 799.03 NETSIX, IT |
| 15 | 32629 | 26283 | 0.45% | 2 | 13141.50 CITY-OF-CHARLOTTE-ASN, US |
| 16 | 41495 | 25697 | 0.44% | 6 | 4282.83 FAELIX, GB |
| 17 | 267638 | 25041 | 0.43% | 1 | 25041.00 Wind Telecomunicacao do Brasil Ltda - ME, BR |
| 18 | 268976 | 22973 | 0.40% | 8 | 2871.62 Weclix Telecom SA, BR |
| 19 | 262865 | 22026 | 0.38% | 8 | 2753.25 Ired Internet, BR |
| 20 | 268275 | 21742 | 0.38% | 1 | 21742.00 GOOD TELECOM PROVEDOR DE INTERNET LTDA, BR |
| 21 | 61665 | 20176 | 0.35% | 8 | 2522.00 ROMICROS BRASIL LTDA, BR |
| 22 | 12654 | 19265 | 0.33% | 39 | 493.97 RIPE-NCC-RIS-AS Reseaux IP Europeens Network Coordination Centre RIPE NCC, NL |
| 23 | 52861 | 19097 | 0.33% | 7 | 2728.14 SN Internet Navegantes Ltda ME, BR |
| 24 | 52817 | 18994 | 0.33% | 9 | 2110.44 NEWCENTER TELECOM, BR |
| 25 | 28201 | 18400 | 0.32% | 7 | 2628.57 Companhia Itabirana Telecomunicacoes Ltda, BR |
| 26 | 14840 | 17826 | 0.31% | 24 | 742.75 BR Digital, BR |
| 27 | 269822 | 17755 | 0.31% | 6 | 2959.17 COLOMBIA MAS TV S.A.S, CO |

50 Most active Prefixes for the past 14 days

| RANK | PREFIX | UPDs | % | Origin AS -- AS NAME |
|------|--------------------------------------|-------|-------|---|
| 1 | 2400:dc40::/32 | 36871 | 0.55% | 136440 -- SASPL-AS-AP Sungard Availability Services India Private Limited, IN |
| 2 | 2801:17:4800::/48 | 27692 | 0.41% | 27951 -- Media Commerce Partners S.A, CO |
| 3 | 2620:0:2f0::/48 | 26280 | 0.39% | 32629 -- CITY-OF-CHARLOTTE-ASN, US |
| 4 | 2804:448c::/32 | 25041 | 0.37% | 267638 -- Wind Telecomunicacao do Brasil Ltda - ME, BR |
| 5 | 2804:4ea8::/32 | 21742 | 0.32% | 268275 -- GOOD TELECOM PROVEDOR DE INTERNET LTDA, BR |
| 6 | 2801:1b6::/44 | 21308 | 0.32% | 27951 -- Media Commerce Partners S.A, CO |
| 7 | 2801:174:3::/48 | 21287 | 0.32% | 27951 -- Media Commerce Partners S.A, CO |
| 8 | 2801:172:2::/48 | 21215 | 0.31% | 27951 -- Media Commerce Partners S.A, CO |
| 9 | 2a01:9e02::/32 | 19958 | 0.30% | 41495 -- FAELIX, GB |
| 10 | 2a05:3181::/32 | 19532 | 0.29% | 31514 -- INF-NET-AS, RU |
| 11 | 2a05:3181:ffff::/48 | 19529 | 0.29% | 31514 -- INF-NET-AS, RU |
| 12 | 2804:b:8000::/34 | 17495 | 0.26% | 14840 -- BR Digital, BR |
| 13 | 2801:186::/44 | 16865 | 0.25% | 27951 -- Media Commerce Partners S.A, CO |
| 14 | 2405:4000:800:8::/64 | 14929 | 0.22% | 38082 -- IIT-TIG-AS-AP True International Gateway Co., Ltd., TH |
| 15 | 2a04:1bc3::/32 | 13858 | 0.21% | 60982 -- WARPON-AS, RO |
| 16 | 2a04:1bc7::/32 | 13857 | 0.21% | 60982 -- WARPON-AS, RO |
| 17 | 2a05:dfc7:44::/48 | 13691 | 0.20% | 208753 -- NETSIX, IT |
| 18 | 2a05:dfc7:47::/48 | 13650 | 0.20% | 208753 -- NETSIX, IT |
| 19 | 2a04:5b05::/32 | 13646 | 0.20% | 203974 -- ADS, RO |
| 20 | 2a0e:b107:b85::/48 | 13474 | 0.20% | 212995 -- TAN-NET, CH |
| 21 | 2a0f:9400:7409::/48 | 13272 | 0.20% | 58336 -- IXTS-AS, CN |
| 22 | 2a0f:9400:743e::/48 | 13271 | 0.20% | 58336 -- IXTS-AS, CN |
| 23 | 2a0f:9400:744d::/48 | 13269 | 0.20% | 58336 -- IXTS-AS, CN |
| 24 | 2a0f:9400:7408::/48 | 13268 | 0.20% | 58336 -- IXTS-AS, CN |
| 25 | 2a0f:9400:744c::/48 | 13264 | 0.20% | 58336 -- IXTS-AS, CN |
| 26 | 2a0f:9400:7460::/48 | 13264 | 0.20% | 58336 -- IXTS-AS, CN |
| 27 | 2a0f:9400:7442::/48 | 13263 | 0.20% | 58336 -- IXTS-AS, CN |
| 28 | 2a0f:9400:7438::/48 | 13263 | 0.20% | 58336 -- IXTS-AS, CN |
| 29 | 2a0f:9400:740d::/48 | 13262 | 0.20% | 58336 -- IXTS-AS, CN |
| 30 | 2a0f:9400:7456::/48 | 13262 | 0.20% | 58336 -- IXTS-AS, CN |
| 31 | 2a0f:9400:7465::/48 | 13260 | 0.20% | 58336 -- IXTS-AS, CN |
| 32 | 2a0f:9400:7403::/48 | 13257 | 0.20% | 58336 -- IXTS-AS, CN |

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

- 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)
- 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)

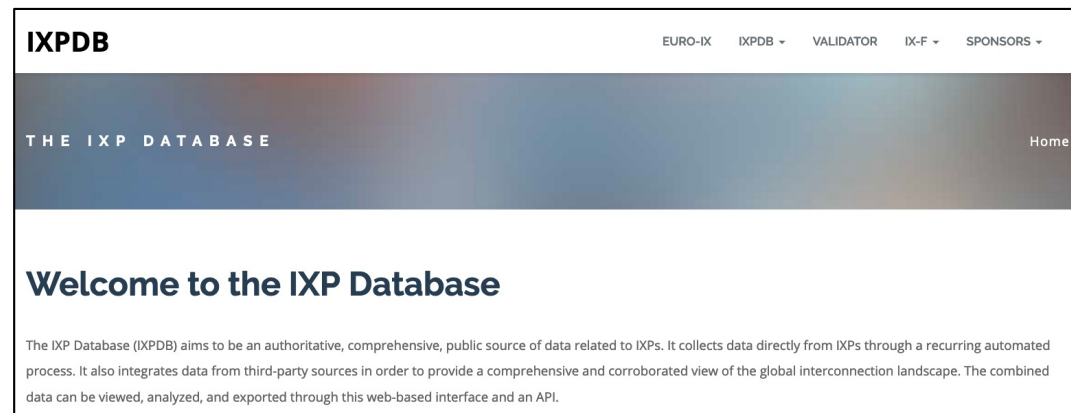
- Open Peering
 - Where an ISP publicly states that they will peer with all parties who approach them for peering
 - Commonly found at IXPs where ISP participates via the Route Server
- Selective Peering
 - Where an ISP's peering policy depends on the nature of the operator who requests peering with 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 AS operators 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



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 | <input checked="" type="radio"/> Unicast IPv4 <input type="radio"/> Multicast <input checked="" type="radio"/> 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 | <input type="radio"/> |
| IPv6 | 2001:7fa:0:1::/64 |
| IPv4 | 123.255.88.0/21 |

Local Facilities

| 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

| 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 Peerina LAN | 123.255.91.98 | 10G |

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 | <input checked="" type="radio"/> Unicast IPv4 <input type="radio"/> Multicast <input checked="" type="radio"/> IPv6 <input type="radio"/> Never via route servers |
| Last Updated | 2019-12-29T14:56:38Z |
| Notes ? | <p>If you have a connectivity issue to Amazon then please visit:</p> <ul style="list-style-type: none"> • IPv4: http://ec2-reachability.amazonaws.com/ • IPv6: http://ipv6.ec2-reachability.amazonaws.com/ <p>And include detail on prefixes you think you have a problem with if you contact our Ops alias. This will reduce time with troubleshooting.</p> <p>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.)</p> <ul style="list-style-type: none"> • Seattle • Palo Alto • San Jose • Los Angeles • Dallas |

Public Peering Exchange Points

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

Private Peering Facilities

| Facility ▼ ASN | Country City |
|--|-------------------|
| 151 Front Street West Toronto 16509 | Canada Toronto |
| 25 Lake Street / 250 Front Street West | Canada |

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 | <input checked="" type="radio"/> Unicast IPv4 <input type="radio"/> Multicast <input checked="" type="radio"/> IPv6 <input type="radio"/> Never via route servers |
| Last Updated | 2020-02-05T11:43:25Z |
| Notes ? | <p>IPv4 + IPv6 Prefixes above would be actuals, not proposed max- prefix values.</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> <p>All trouble ticket requests or support related emails should be sent to carrier-csc@teliacompany.com.</p> |

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

| Exchange ▼ ASN | IPv4 IPv6 | Speed RS Peer |
|-------------------|--------------|------------------|
|-------------------|--------------|------------------|

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

Private Peering Facilities

| Facility ▼ ASN | Country City |
|--|--|
| 365 Data Centers Buffalo (BU1) 1299 | United States of America Buffalo |
| 365 Data Centers Detroit (DT1) 1299 | United States of America Southfield |
| 365 Data Centers Nashville (NA1) 1299 | United States of America Nashville |
| 365 Data Centers Tampa (TA1) 1299 | United States of America Tampa |
| 3U Rechenzentrum Berlin 1299 | Germany Berlin |
| Altus IT 1299 | Croatia Zagreb |
| Borovaya 57 1299 | Russia St. Petersburg |
| CE Colo Prague 1299 | Czechia Prague |
| CINECA - DC NaMeX 1299 | Italy Roma |
| COD BM-18 1299 | Russia St.Petersburg |
| Caldera21 1299 | Italy Milan |
| CarrierColo Berlin Luetzow (I/P/B/ site B) 1299 | Germany Berlin |
| Cologix MTL3 1299 | Canada Montreal |
| Cologix TOR1 1299 | Canada Toronto |



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

Internet Routing Registry

- Which IRR database to use?
 - Members of a Regional Internet Registry are recommended 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>
 - Note: Placing objects in the RADB requires an annual subscription fee

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



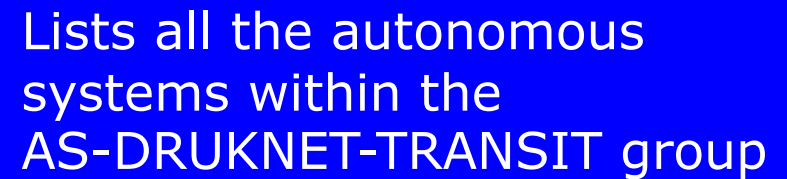
AS-Set: Purpose

- 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:        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
- But there are some important best practices
 1. Signing ROAs
 2. 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:

| | |
|-------------------|--------------|
| Prefix | 10.10.0.0/16 |
| Max-Length | /18 |
| Origin-AS | 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!!**

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 day?

Route Origin Validation

- The ideal would be to write directly to the active BGP table
- But most 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 (November 2021 sees over 290k), and frequent changes or updates of ROAs:
 - Routers are sending frequent Route Refresh requests to peers (typically every few minutes)
 - Peers are being “bombed” 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 Origin Validation

- Most modern BGP implementations have ADJ-RIB-IN by default
 - ROA updates do not cause a problem when operating ROV

- For those that do not implement ADJ-RIB-IN by default:
 - For example, all Cisco IOS/IOS-XE/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 saves arguments between operators about whose BGP filters may have configuration errors!

Configuration Tips



Of passwords, tricks and templates

IBGP and IGP

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 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 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
...
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 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 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

Routing Security

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

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
 - MANRS for CDN & Cloud Providers



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