

BGP Best Current Practices



ISP Workshops

Configuring BGP



Where do we start?

IOS Good Practices

- ❑ ISPs should start off with the following BGP commands as a basic template:

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

← Replace with public ASN

← Make ebgp and ibgp distance the same

- ❑ If supporting more than just IPv4 unicast neighbours

```
no bgp default ipv4-unicast
```

- Turns off IOS assumption that all neighbours will exchange IPv4 prefixes

Cisco IOS Good Practices

- ❑ BGP in Cisco IOS is **permissive** by default
- ❑ 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
- ❑ **Good practice is to ensure that each eBGP neighbour has inbound and outbound filter applied:**

```
router bgp 64511
  neighbor 1.2.3.4 remote-as 64510
  neighbor 1.2.3.4 prefix-list as64510-in in
  neighbor 1.2.3.4 prefix-list as64510-out out
```

What is BGP for??



What is an IGP not for?

BGP versus OSPF/ISIS

- ❑ Internal Routing Protocols (IGPs)
 - Examples are ISIS 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/ISIS

- ❑ 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/ISIS

- ❑ 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
 - April 2016: 324000 /24s in IPv4 table of 590000 prefixes
- ❑ **The same is happening for /48s with IPv6**
 - January 2016: 12900 /48s in IPv6 table of 27700 prefixes

Configuring Aggregation – Cisco IOS

- ❑ ISP has 101.10.0.0/19 address block
- ❑ To put into BGP as an aggregate:

```
router bgp 64511
  network 101.10.0.0 mask 255.255.224.0
  ip route 101.10.0.0 255.255.224.0 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
  network 101.10.0.0 mask 255.255.224.0
  neighbor 102.102.10.1 remote-as 101
  neighbor 102.102.10.1 prefix-list out-filter out
  !
ip route 101.10.0.0 255.255.224.0 null0
!
ip prefix-list out-filter permit 101.10.0.0/19
ip prefix-list out-filter deny 0.0.0.0/0 le 32
!
```

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:
 - ❑ Now ranging from a /20 to a /24 depending on RIR
 - ❑ Different sizes for different address blocks
 - ❑ (APNIC changed its minimum allocation to /24 in October 2010)
 - For IPv6:
 - ❑ /48 for assignment, /32 for allocation
- ❑ Until recently there was no real reason to see anything longer than a /22 IPv4 prefix in the Internet
 - Maybe IPv4 run-out is starting to have an impact?

Aggregation – Example

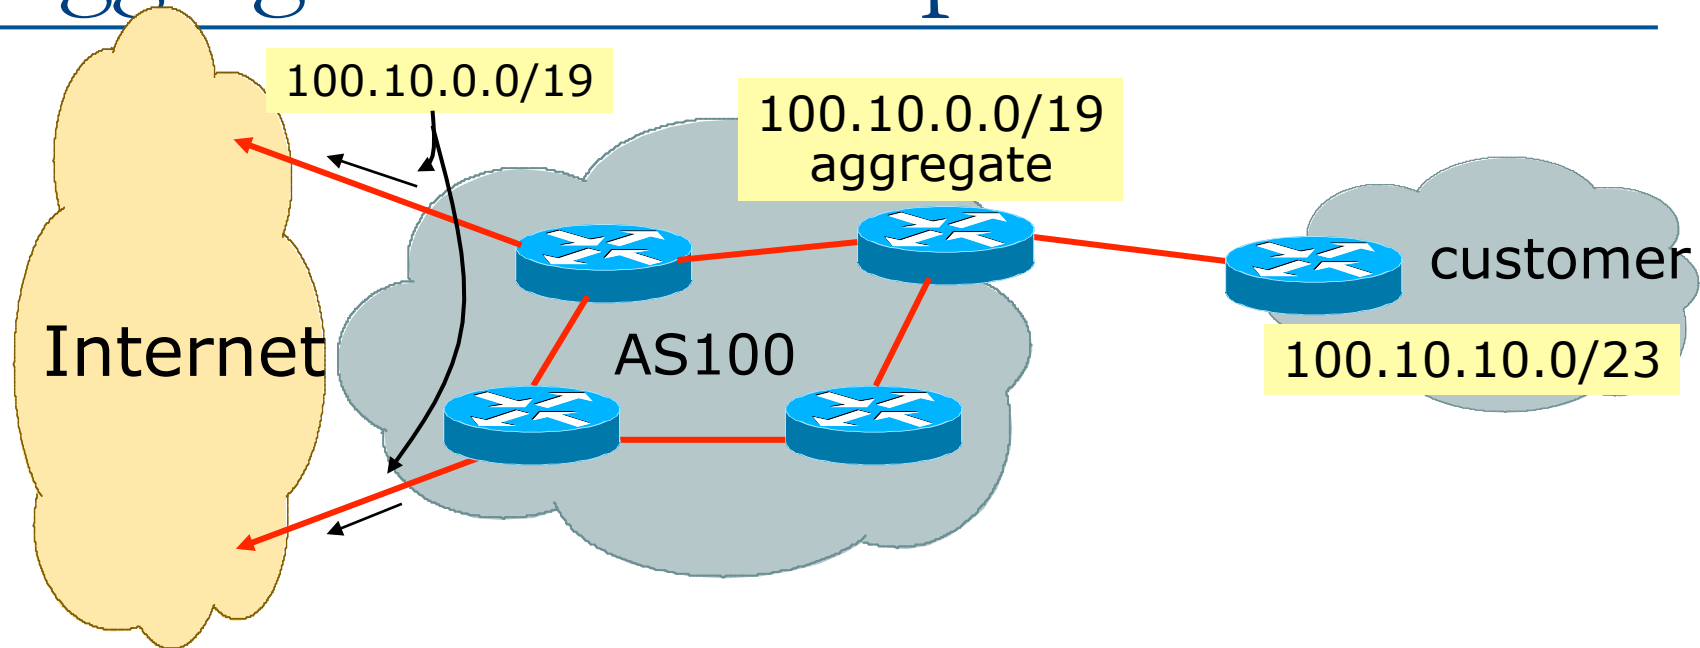


- ❑ Customer has /23 network assigned from AS100's /19 address block
- ❑ AS100 announces customers' individual networks to the Internet

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

(April 2016)

□ Current Internet Routing Table Statistics

- | | |
|--|--------|
| ■ BGP Routing Table Entries | 588570 |
| ■ Prefixes after maximum aggregation | 216761 |
| ■ Unique prefixes in Internet | 288585 |
| ■ Prefixes smaller than registry alloc | 192680 |
| ■ /24s announced | 324221 |
| ■ ASes in use | 53367 |
- (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
05-04-16	604203	337395
06-04-16	604667	337556
07-04-16	604749	337656
08-04-16	604799	337545
09-04-16	604798	336944
10-04-16	604687	337300
11-04-16	604810	336995
12-04-16	603298	337027

Plot: [BGP Table Size](#)



AS Summary

53612	Number of ASes in routing system
20955	Number of ASes announcing only one prefix
5608	Largest number of prefixes announced by an AS AS4538 : ERX-CERNET-BKB China Education and Research Network Center, CN
120999168	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)



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 1 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 11 April 2016 06:23 (UTC+1000)

50 Most active ASes for the past 7 days

RANK	ASN	UPDs	%	Prefixes	UPDs/Prefix	AS NAME
1	9829	267294	9.41%	2480	107.78	BSNL-NIB National Internet Backbone, IN
2	48340	123678	4.35%	10	12367.80	LIRA-S-IX-AS OOO Lira-S, RU
3	57731	92657	3.26%	109	850.06	DIVINETWORKS DiViNetworks LTD., IL
4	38197	64111	2.26%	1558	41.15	SUNHK-DATA-AS-AP Sun Network (Hong Kong) Limited, HK
5	13118	25123	0.88%	97	259.00	ASN-YARTELECOM PJSC Rostelecom, RU
6	17974	21977	0.77%	2923	7.52	TELKOMNET-AS2-AP PT Telekomunikasi Indonesia, ID
7	134438	20850	0.73%	1	20850.00	AIRAAIFUL-AS-AP Aira & Aiful Public Company Limited, TH
8	41691	19669	0.69%	41	479.73	SUMTEL-AS-RIPE Summa Telecom LLC, RU
9	27947	18473	0.65%	370	49.93	Telconet S.A, EC
10	56636	17727	0.62%	1	17727.00	ASVEDARU VEDA Ltd., RU
11	52931	17377	0.61%	17	1022.18	Netvip Telecomunicações Ltda, BR
12	37457	16288	0.57%	1391	11.71	Telkom-Internet, ZA
13	38264	14717	0.52%	317	46.43	WATEEN-IMS-PK-AS-AP National WiMAX/IMS environment, PK
14	131090	13318	0.47%	390	34.15	CAT-IDC-4BYTENET-AS-AP CAT TELECOM Public Company Ltd,CAT, TH
15	41480	12697	0.45%	16	793.56	SYSTEMEC-AS Systemec B.V., NL
16	15468	12453	0.44%	265	46.99	KLGELECS-AS PJSC Rostelecom, RU
17	33788	12399	0.44%	35	354.26	KANARTEL, SD
18	45899	12006	0.42%	984	12.20	VNPT-AS-VN VNPT Corp, VN
19	7545	11898	0.42%	3336	3.57	TPG-INTERNET-AP TPG Telecom Limited, AU
20	4800	11831	0.42%	300	39.44	LINTASARTA-AS-AP Network Access Provider and Internet Service Provider, ID
21	18403	11690	0.41%	1012	11.55	FPT-AS-AP The Corporation for Financing & Promoting Technology, VN
22	55832	11593	0.41%	112	103.51	HOMESYSTEM-AS-AP HOME SYSTEMS PVT.LTD, IN
23	43754	11431	0.40%	310	36.87	ASIATECH Asiatech Data Transfer Inc PLC, IR
24	45194	11397	0.40%	523	21.79	SIPI -AS Syscon Infoway Pvt Ltd IN

bgpupdates.potaroo.net				
The BGP Instability Report				
50 Most active Prefixes for the past 7 days				
RANK	PREFIX	UPDs	%	Origin AS -- AS NAME
1	93.181.192.0/19	21863	0.72%	13118 -- ASN-YARTELECOM PJSC Rostelecom, RU
2	110.170.17.0/24	20850	0.69%	134438 -- AIRAAIFUL-AS-AP Aira & Aiful Public Company Limited, TH
3	195.128.159.0/24	17727	0.58%	56636 -- ASVEDARU VEDA Ltd., RU
4	177.185.110.0/24	17137	0.56%	52931 -- Netvip Telecomunicações Ltda, BR
5	77.244.208.0/20	13766	0.45%	48340 -- LIRA-S-IX-AS OOO Lira-S, RU
6	92.53.64.0/20	13763	0.45%	48340 -- LIRA-S-IX-AS OOO Lira-S, RU
7	92.53.80.0/20	13754	0.45%	48340 -- LIRA-S-IX-AS OOO Lira-S, RU
8	78.155.192.0/20	13754	0.45%	48340 -- LIRA-S-IX-AS OOO Lira-S, RU
9	80.93.176.0/20	13748	0.45%	48340 -- LIRA-S-IX-AS OOO Lira-S, RU
10	46.229.208.0/22	13740	0.45%	48340 -- LIRA-S-IX-AS OOO Lira-S, RU
11	92.53.100.0/22	13722	0.45%	48340 -- LIRA-S-IX-AS OOO Lira-S, RU
12	92.53.108.0/22	13719	0.45%	48340 -- LIRA-S-IX-AS OOO Lira-S, RU
13	78.155.216.0/21	13711	0.45%	48340 -- LIRA-S-IX-AS OOO Lira-S, RU
14	61.7.155.0/24	12873	0.42%	131090 -- CAT-IDC-4BYTENET-AS-AP CAT TELECOM Public Company Ltd,CAT, TH
15	182.23.47.0/24	10263	0.34%	4800 -- LINTASARTA-AS-AP Network Access Provider and Internet Service Provider, ID
16	69.168.232.0/24	8409	0.28%	36344 -- ADVAN-CAST - Advan Technologies LLC, US
17	184.84.182.0/24	7849	0.26%	35994 -- AKAMAI-AS - Akamai Technologies, Inc., US
18	66.19.194.0/24	6185	0.20%	6316 -- AS-PAETEC-NET - PaeTec Communications, Inc., US
19	202.56.215.0/24	5963	0.20%	24560 -- AIRTELBROADBAND-AS-AP Bharti Airtel Ltd., Telemedia Services, IN
20	109.172.53.0/24	5871	0.19%	41691 -- SUMTEL-AS-RIPE Summa Telecom LLC, RU
21	109.172.51.0/24	5865	0.19%	41691 -- SUMTEL-AS-RIPE Summa Telecom LLC, RU
22	109.172.52.0/24	5863	0.19%	41691 -- SUMTEL-AS-RIPE Summa Telecom LLC, RU
23	203.252.142.0/24	5706	0.19%	9459 -- ASKONKUK Konkuk University, KR
24	196.43.234.0/24	5612	0.18%	30619 -- TDM-AS, MZ 37436 -- INTIC, MZ
25	185.53.196.0/22	5419	0.18%	8953 -- ASN-ORANGE-ROMANIA Orange Romania S.A., RO
26	203.190.142.0/24	4744	0.16%	9430 -- STPI-NOIDA Software Technology Parks of India,Block-IV, IN
27	223.239.16.0/20	4414	0.15%	45609 -- BHARTI-MOBILITY-AS-AP Bharti Airtel Ltd. AS for GPRS Service, IN
28	223.239.0.0/23	4407	0.15%	45609 -- BHARTI-MOBILITY-AS-AP Bharti Airtel Ltd. AS for GPRS Service, IN

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 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:      202.12.28.0 - 202.12.29.255
netname:      APNIC-SERVICES
descr:        Asia Pacific Network Information Centre
descr:        Regional Internet Registry for the Asia-Pacific Region
descr:        6 Cordelia Street
descr:        PO Box 3646
descr:        South Brisbane, QLD 4101
descr:        Australia
country:      AU
admin-c:      AIC1-AP
tech-c:       AIC1-AP
mnt-by:       APNIC-HM
mnt-irt:      IRT-APNIC-IS-AP
changed:      hm-changed@apnic.net
status:       ASSIGNED PORTABLE
changed:      hm-changed@apnic.net 20131128
source:      APNIC
```

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 whois.ripe.net 193.128.0.0/16
```

```
inetnum:      193.128.0.0 - 193.133.255.255
netname:      UK-PIPEX-193-128-133
country:      GB
org:          ORG-UA24-RIPE
admin-c:      WERT1-RIPE
tech-c:       UPHM1-RIPE
status:       ALLOCATED UNSPECIFIED
remarks:      Please send abuse notification to abuse@uk.uu.net
mnt-by:       RIPE-NCC-HM-MNT
mnt-lower:    AS1849-MNT
mnt-routes:   AS1849-MNT
mnt-routes:   WCOM-EMEA-RICE-MNT
mnt-irt:      IRT-MCI-GB
created:      2002-06-25T15:05:40Z
last-modified: 2010-04-28T09:08:15Z
source:       RIPE # Filtered
```

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 100.50.0.0/20 block
 - Should only announce this to upstreams
 - Upstreams should only accept this from them
- ❑ Configuration on upstream

```
router bgp 100
  neighbor 102.102.10.1 remote-as 101
  neighbor 102.102.10.1 prefix-list customer in
  neighbor 102.102.10.1 prefix-list default out
  !
ip prefix-list customer permit 100.50.0.0/20
  !
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 the IRRToolSet

<https://github.com/irrtoolset/irrtoolset>

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
  neighbor 102.102.10.1 remote-as 101
  neighbor 102.102.10.1 prefix-list my-peer in
!
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
```

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
  network 101.10.0.0 mask 255.255.224.0
  neighbor 101.5.7.1 remote-as 101
  neighbor 101.5.7.1 prefix-list infilter in
  neighbor 101.5.7.1 prefix-list outfilter out
!
ip prefix-list infilter permit 0.0.0.0/0
!
ip prefix-list outfilter permit 101.10.0.0/19
```


Receiving Prefixes: From Upstream/Transit Provider

▣ Upstream Router Configuration

```
router bgp 101
  neighbor 101.5.7.2 remote-as 100
  neighbor 101.5.7.2 default-originate
  neighbor 101.5.7.2 prefix-list cust-in in
  neighbor 101.5.7.2 prefix-list cust-out out
!
ip prefix-list cust-in permit 101.10.0.0/19
!
ip prefix-list cust-out permit 0.0.0.0/0
```

Receiving Prefixes:

From Upstream/Transit Provider

- ❑ If 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"
www.team-cymru.org/Services/Bogons/http.html
- ❑ For IPv4 also consult:
www.rfc-editor.org/rfc/rfc6441.txt (BCP171)
- ❑ For IPv6 also consult:
www.space.net/~gert/RIPE/ipv6-filters.html
- ❑ Bogon Route Server:
www.team-cymru.org/Services/Bogons/routeserver.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 101.5.7.1 remote-as 101
  neighbor 101.5.7.1 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 addr
  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 proto
  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 & Expmnt
  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 permit 2001::/32         ! Teredo
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 permit 2002::/16         ! 6to4
ipv6 prefix-list v6in-filter deny 2002::/16 le 128    ! 6to4 subnets
ipv6 prefix-list v6in-filter deny 2020:3030::/32 le 128 ! Local Prefix
ipv6 prefix-list v6in-filter deny 3ffe::/16 le 128    ! Old 6bone
ipv6 prefix-list v6in-filter permit 2000::/3 le 48    ! Global Unicast
ipv6 prefix-list v6in-filter deny ::/0 le 128
```

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 215.17.3.1 255.255.255.255
!
interface Serial 5/0
  ip unnumbered loopback 0
  ip verify unicast reverse-path
!
ip route 215.34.10.0 255.255.252.0 Serial 5/0
!
router bgp 100
  network 215.34.10.0 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 215.34.10.0 255.255.252.0 Serial 5/0
!
router bgp 100
 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 215.34.10.0/22 le 30
```

Injecting prefixes into iBGP

- Route-map ISP-block 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

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 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 15
```

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 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

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

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.org/documents.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 version 3 BGP 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!!

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
- ❑ It's all about scaling – if your network won't scale, then it won't be successful

BGP Best Current Practices



ISP Workshops