

# BGP Best Current Practices



ISP Training 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
```

← Replace with public ASN

```
bgp deterministic-med
```

```
distance bgp 200 200 200
```

← Make ebgp and ibgp distance the same

```
no synchronization
```

```
no auto-summary
```

- ❑ If supporting more than just IPv4 unicast neighbours

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

- is also very important and required

# 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
  neighbour 1.2.3.4 remote-as 64510
  neighbour 1.2.3.4 prefix-list as64510-in in
  neighbour 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 used internally (iBGP) and externally (eBGP)
- iBGP used to carry
  - some/all Internet prefixes across backbone
  - customer prefixes
- eBGP 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
- ❑ Unfortunately too many people are still thinking about class Cs, resulting in a proliferation of /24s in the Internet routing table

# 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

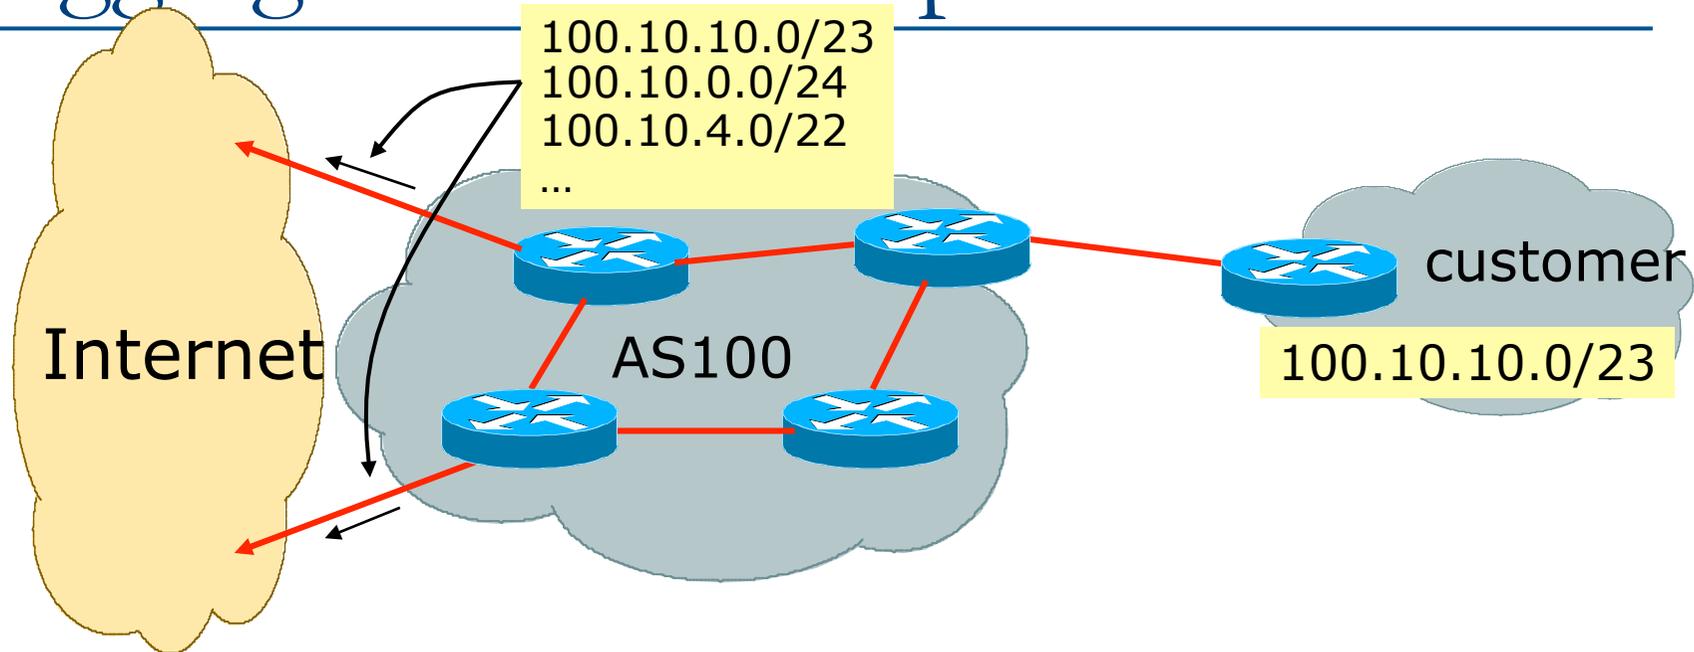
```
router bgp 64511
  network 101.10.0.0 mask 255.255.224.0
  neighbor 102.102.10.1 remote-as 101
  neighbor 102.102.10.1 prefix-list out-filter out
!
ip route 101.10.0.0 255.255.224.0 null0
!
ip prefix-list out-filter permit 101.10.0.0/19
ip prefix-list out-filter deny 0.0.0.0/0 le 32
```

# Announcing an Aggregate

---

- ❑ ISPs who don't and won't aggregate are held in poor regard by community
- ❑ Registries publish their minimum allocation size
  - Anything from a /20 to a /22 depending on RIR
  - Different sizes for different address blocks
- ❑ No real reason to see anything longer than a /22 prefix in the Internet
  - BUT there are currently (Nov 2011) >198000 /24s!
- ❑ But: APNIC changed (Oct 2010) its minimum allocation size on all blocks to /24
  - IPv4 run-out is starting to have an impact

# Aggregation – Example



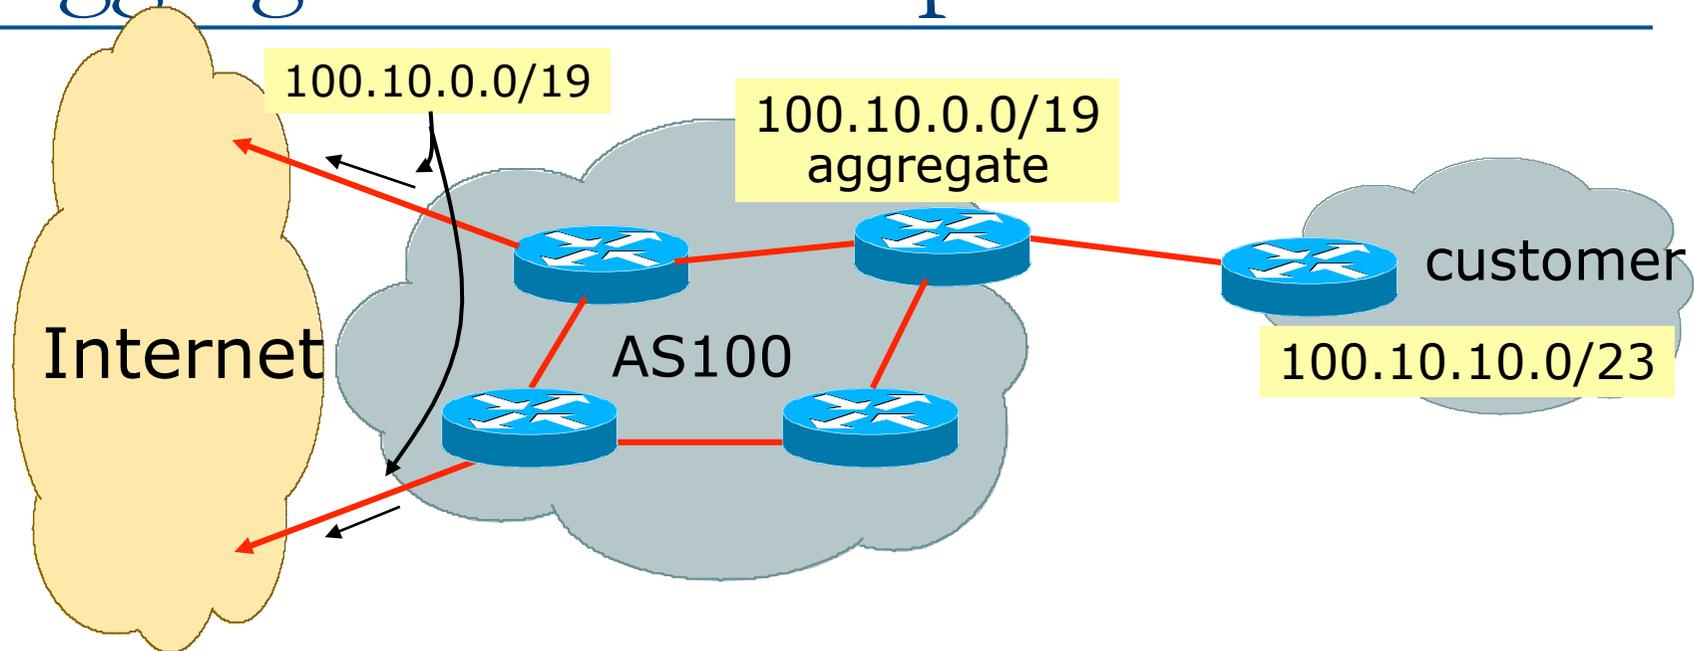
- ❑ Customer has /23 network assigned from AS100's /19 address block
- ❑ AS100 announces customers' individual networks to the Internet

# 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???

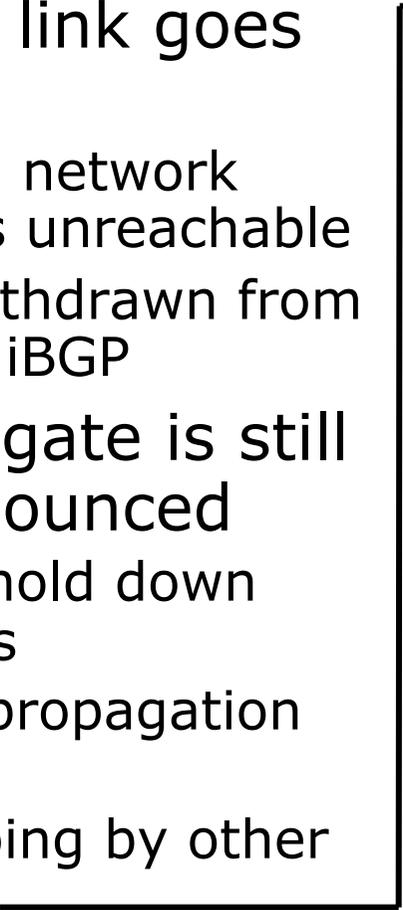
# Aggregation – Example



- ❑ 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 (Nov 2011)

---

## □ Current Internet Routing Table Statistics

■ BGP Routing Table Entries	382305
■ Prefixes after maximum aggregation	166896
■ Unique prefixes in Internet	187837
■ Prefixes smaller than registry alloc	161493
■ /24s announced	198683
■ ASes in use	39323

# Efforts to improve aggregation

---

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

# Efforts to Improve Aggregation

## The CIDR Report

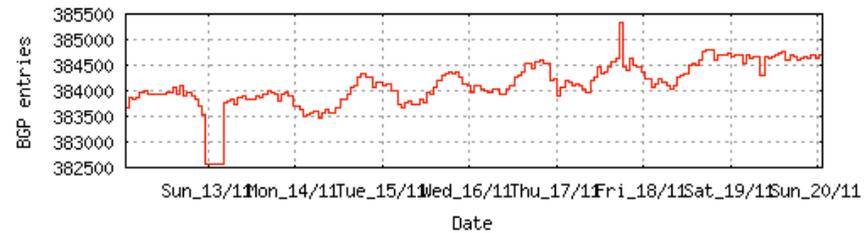
---

- ❑ 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

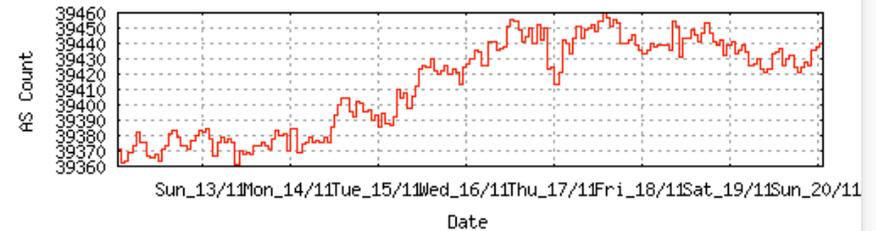
Date	Prefixes	CIDR Aggregated
13-11-11	382573	224878
14-11-11	383706	224973
15-11-11	384163	225105
16-11-11	384097	225263
17-11-11	384218	225441
18-11-11	384382	225289
19-11-11	384737	225361
20-11-11	384631	225359



Plot: [BGP Table Size](#)

## AS Summary

- 39439 Number of ASes in routing system
- 16631 Number of ASes announcing only one prefix
- 3486 Largest number of prefixes announced by an AS  
[AS6389](#): BELLSOUTH-NET-BLK - BellSouth.net Inc.
- 108833792 Largest address span announced by an AS (/32s)  
[AS4134](#): CHINANET-BACKBONE No.31,Jin-rong Street



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 \(from Registry WHOIS data\)](#)

## Aggregation Suggestions

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

Rank	AS	AS Name	Current	Wthdw	Aggte	Annce	Redctn	%
7	<a href="#">AS4755</a>	TATACOMM-AS TATA Communications formerly VSNL	1526	1318	35	243	1283	84.08%

Prefix	AS Path	Aggregation Suggestion
14.140.0.0/14	4777 2516 6453 4755	
14.140.0.0/22	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.4.0/23	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.6.0/23	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.16.0/22	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.20.0/22	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.24.0/22	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.32.0/23	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.40.0/21	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.48.0/21	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.56.0/21	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.64.0/21	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.72.0/22	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.80.0/23	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.82.0/23	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.84.0/22	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.88.0/21	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.96.0/22	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.104.0/21	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.112.0/22	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.116.0/23	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.118.0/23	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.120.0/21	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.128.0/23	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.130.0/23	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.144.0/21	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.152.0/22	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.160.0/22	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.164.0/22	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.168.0/21	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.176.0/22	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.180.0/22	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.184.0/23	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.186.0/23	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.188.0/23	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.190.0/23	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
14.140.196.0/23	4777 2516 6453 4755	- Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755

## Aggregation Suggestions

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

Rank	AS	AS Name	Current	Withdw	Aggte	Annce	Redctn	%
3	<a href="#">AS18566</a>	COVAD - Covad Communications Co.	2093	1882	195	406	1687	80.60%

Prefix	AS Path	Aggregation Suggestion
64.81.16.0/22	4777 2516 3356 18566	
64.81.20.0/22	4777 2516 4565 18566	
64.81.22.0/24	4777 2516 4565 18566	- Withdrawn - matching aggregate 64.81.20.0/22 4777 2516 4565 18566
64.81.24.0/21	4777 2516 3356 18566	+ Announce - aggregate of 64.81.24.0/22 (4777 2516 3356 18566) and 64.81.28.0/22 (4777 2516 3356 18566)
64.81.24.0/22	4777 2516 3356 18566	- Withdrawn - aggregated with 64.81.28.0/22 (4777 2516 3356 18566)
64.81.28.0/22	4777 2516 3356 18566	- Withdrawn - aggregated with 64.81.24.0/22 (4777 2516 3356 18566)
64.81.32.0/20	4777 2516 4565 18566	
64.81.32.0/24	4777 2516 4565 18566	- Withdrawn - matching aggregate 64.81.32.0/20 4777 2516 4565 18566
64.81.33.0/24	4777 2516 4565 18566	- Withdrawn - matching aggregate 64.81.32.0/20 4777 2516 4565 18566
64.81.34.0/24	4777 2516 4565 18566	- Withdrawn - matching aggregate 64.81.32.0/20 4777 2516 4565 18566
64.81.35.0/24	4777 2516 4565 18566	- Withdrawn - matching aggregate 64.81.32.0/20 4777 2516 4565 18566
64.81.36.0/24	4777 2516 4565 18566	- Withdrawn - matching aggregate 64.81.32.0/20 4777 2516 4565 18566
64.81.37.0/24	4777 2516 4565 18566	- Withdrawn - matching aggregate 64.81.32.0/20 4777 2516 4565 18566
64.81.38.0/24	4777 2516 4565 18566	- Withdrawn - matching aggregate 64.81.32.0/20 4777 2516 4565 18566
64.81.39.0/24	4777 2516 4565 18566	- Withdrawn - matching aggregate 64.81.32.0/20 4777 2516 4565 18566
64.81.40.0/24	4777 2516 4565 18566	- Withdrawn - matching aggregate 64.81.32.0/20 4777 2516 4565 18566
64.81.44.0/24	4777 2516 4565 18566	- Withdrawn - matching aggregate 64.81.32.0/20 4777 2516 4565 18566
64.81.48.0/20	4777 2516 3356 18566	
64.81.48.0/24	4777 2516 3356 18566	- Withdrawn - matching aggregate 64.81.48.0/20 4777 2516 3356 18566
64.81.49.0/24	4777 2516 3356 18566	- Withdrawn - matching aggregate 64.81.48.0/20 4777 2516 3356 18566
64.81.50.0/24	4777 2516 3356 18566	- Withdrawn - matching aggregate 64.81.48.0/20 4777 2516 3356 18566
64.81.51.0/24	4777 2516 3356 18566	- Withdrawn - matching aggregate 64.81.48.0/20 4777 2516 3356 18566
64.81.52.0/24	4777 2516 3356 18566	- Withdrawn - matching aggregate 64.81.48.0/20 4777 2516 3356 18566
64.81.53.0/24	4777 2516 3356 18566	- Withdrawn - matching aggregate 64.81.48.0/20 4777 2516 3356 18566
64.81.54.0/24	4777 2516 3356 18566	- Withdrawn - matching aggregate 64.81.48.0/20 4777 2516 3356 18566
64.81.55.0/24	4777 2516 3356 18566	- Withdrawn - matching aggregate 64.81.48.0/20 4777 2516 3356 18566
64.81.56.0/24	4777 2516 3356 18566	- Withdrawn - matching aggregate 64.81.48.0/20 4777 2516 3356 18566
64.81.57.0/24	4777 2516 3356 18566	- Withdrawn - matching aggregate 64.81.48.0/20 4777 2516 3356 18566
64.81.58.0/24	4777 2516 3356 18566	- Withdrawn - matching aggregate 64.81.48.0/20 4777 2516 3356 18566
64.81.59.0/24	4777 2516 3356 18566	- Withdrawn - matching aggregate 64.81.48.0/20 4777 2516 3356 18566
64.81.60.0/24	4777 2516 3356 18566	- Withdrawn - matching aggregate 64.81.48.0/20 4777 2516 3356 18566
64.81.61.0/24	4777 2516 3356 18566	- Withdrawn - matching aggregate 64.81.48.0/20 4777 2516 3356 18566
64.81.64.0/20	4777 2516 3356 18566	
64.81.64.0/24	4777 2516 3356 18566	- Withdrawn - matching aggregate 64.81.64.0/20 4777 2516 3356 18566
64.81.65.0/24	4777 2516 3356 18566	- Withdrawn - matching aggregate 64.81.64.0/20 4777 2516 3356 18566
64.81.66.0/24	4777 2516 3356 18566	- Withdrawn - matching aggregate 64.81.64.0/20 4777 2516 3356 18566
64.81.67.0/24	4777 2516 3356 18566	- Withdrawn - matching aggregate 64.81.64.0/20 4777 2516 3356 18566
64.81.68.0/24	4777 2516 3356 18566	- Withdrawn - matching aggregate 64.81.64.0/20 4777 2516 3356 18566

# Importance of Aggregation

---

- Size of routing table
  - Router Memory is not so much of a problem as it was in the 1990s
  - Routers can be specified to carry 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](http://bgpupdates.potaroo.net/instability/bgpupd.html)

# The BGP Instability Report

The BGP Instability Report is updated daily. This report was generated on 19 November 2011 06:13 (UTC+1000)

## 50 Most active ASes for the past 7 days

RANK	ASN	UPDs	%	Prefixes	UPDs/Prefix	AS NAME
1	8402	46821	2.79%	1717	27.27	CORBINA-AS OJSC "Vimpelcom"
2	9829	41675	2.48%	1166	35.74	BSNL-NIB National Internet Backbone
3	7552	28501	1.70%	1394	20.45	VIETEL-AS-AP Viettel Corporation
4	19743	28407	1.69%	7	4058.14	
5	15180	28138	1.68%	52	541.12	Diveo do Brasil Telecomunicacoes Ltda
6	32528	23379	1.39%	8	2922.38	ABBOTT Abbot Labs
7	27738	22482	1.34%	340	66.12	Ecuadortelecom S.A.
8	20632	19556	1.17%	39	501.44	PETERSTAR-AS PeterStar
9	6316	18930	1.13%	132	143.41	AS-PAETEC-NET - PaeTec Communications, Inc.
10	31148	16979	1.01%	410	41.41	FREENET-AS FreeNet ISP
11	9498	13610	0.81%	845	16.11	BBIL-AP BHARTI Airtel Ltd.
12	5800	12621	0.75%	246	51.30	DNIC-ASBLK-05800-06055 - DoD Network Information Center
13	11492	11764	0.70%	1154	10.19	CABLEONE - CABLE ONE, INC.
14	8866	10778	0.64%	465	23.18	BTC-AS Bulgarian Telecommunication Company Plc.
15	12880	10017	0.60%	96	104.34	DCI-AS Information Technology Company (ITC)
16	16322	8016	0.48%	123	65.17	PARSONLINE PARSONLINE Autonomous System
17	36923	7622	0.45%	192	39.70	SWIFTNG-ASN
18	9394	7597	0.45%	325	23.38	CRNET CHINA RAILWAY Internet(CRNET)
19	6072	7518	0.45%	17	442.24	UNISYS-6072 For routing issues, email hostmaster@unisys.com
20	16916	7196	0.43%	20	359.80	NETLOGIC-WEST - INFINIPLEX LLC DBA NETLOGIC
21	2708	7124	0.42%	130	54.80	Universidad de Guanajuato
22	25184	6951	0.41%	104	66.84	AFRANET AFRANET Co. Tehran, Iran
23	7029	6880	0.41%	3030	2.27	WINDSTREAM - Windstream Communications Inc

## 50 Most active Prefixes for the past 7 days

RANK	PREFIX	UPDs	%	Origin AS -- AS NAME
1	84.204.132.0/24	19385	1.08%	20632 -- PETERSTAR-AS PeterStar
2	130.36.35.0/24	11688	0.65%	32528 -- ABBOTT Abbot Labs
3	130.36.34.0/24	11687	0.65%	32528 -- ABBOTT Abbot Labs
4	66.248.104.0/21	10794	0.60%	6316 -- AS-PAETEC-NET - PaeTec Communications, Inc.
5	202.92.235.0/24	10583	0.59%	9498 -- BBIL-AP BHARTI Airtel Ltd.
6	213.16.48.0/24	10396	0.58%	8866 -- BTC-AS Bulgarian Telecommunication Company Plc.
7	206.80.93.0/24	7195	0.40%	16916 -- NETLOGIC-WEST - INFINIPLEX LLC DBA NETLOGIC
8	65.122.196.0/24	6750	0.38%	19743 --
9	66.248.120.0/21	6270	0.35%	6316 -- AS-PAETEC-NET - PaeTec Communications, Inc.
10	72.164.144.0/24	4333	0.24%	19743 --
11	66.238.91.0/24	4331	0.24%	19743 --
12	65.162.204.0/24	4331	0.24%	19743 --
13	66.89.98.0/24	4331	0.24%	19743 --
14	65.163.182.0/24	4331	0.24%	19743 --
15	208.98.239.0/24	3947	0.22%	25983 -- ENMAX-ENVISION - Enmax Envision Inc.
16	202.153.174.0/24	3207	0.18%	17408 -- ABOVE-AS-AP AboveNet Communications Taiwan
17	192.139.142.0/24	2268	0.13%	23498 -- CDSI - Cogeco Data Services Inc.
18	203.110.64.0/20	2161	0.12%	38528 -- LANIC-AS-AP Lao National Internet Committee
19	217.52.70.0/24	1926	0.11%	15475 -- NOL 36992 -- ETISALAT-MISR
20	217.52.130.0/24	1920	0.11%	15475 -- NOL 36992 -- ETISALAT-MISR
21	14.102.50.0/24	1917	0.11%	18002 -- WORLDPHONE-IN AS Number for Interdomain Routing
22	208.76.228.0/22	1747	0.10%	53558 -- TSG-3 - TCN Systems Group
23	208.54.82.0/24	1660	0.09%	701 -- UUNET - MCI Communications Services, Inc. d/b/a Verizon Business
24	66.248.96.0/21	1594	0.09%	6316 -- AS-PAETEC-NET - PaeTec Communications, Inc.
25	192.246.39.0/24	1310	0.07%	4565 -- MEGAPATH2-US - MegaPath Networks Inc.
26	216.110.67.0/24	1257	0.07%	4323 -- TWTC - tw telecom holdings, inc.

# 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 database)

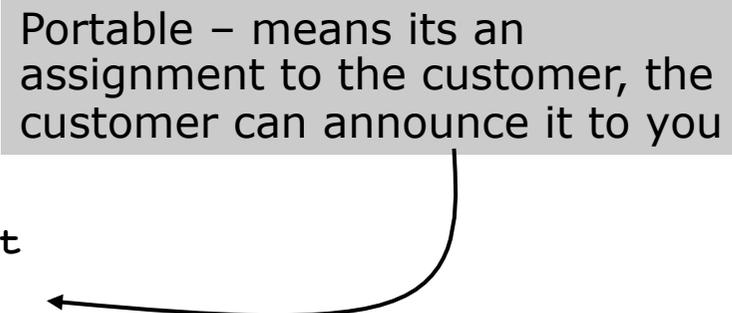
# Receiving Prefixes: From Customers

---

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

```
$ whois -h whois.apnic.net 202.12.29.0
inetnum:          202.12.28.0 - 202.12.29.255
netname:          APNIC-AP
descr:           Asia Pacific Network Information Centre
descr:           Regional Internet Registry for the Asia-Pacific
descr:           6 Cordelia Street
descr:           South Brisbane, QLD 4101
descr:           Australia
country:         AU
admin-c:         AIC1-AP
tech-c:          NO4-AP
mnt-by:          APNIC-HM
mnt-irt:         IRT-APNIC-AP
changed:         hm-changed@apnic.net
status:          ASSIGNED PORTABLE
changed:         hm-changed@apnic.net 20110309
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
inetnum:          193.128.0.0 - 193.133.255.255
netname:          UK-PIPEX-193-128-133
descr:           Verizon UK Limited
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
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
!
ip prefix-list customer permit 100.50.0.0/20
```

# 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  
[www.isc.org/sw/IRRToolSet/](http://www.isc.org/sw/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
- For IPv4:
  - Don't accept private (RFC1918) and certain special use prefixes:  
<http://www.rfc-editor.org/rfc/rfc5735.txt>
  - Don't accept prefixes longer than /24 (?)
- For IPv6:
  - Don't accept certain special use prefixes:  
<http://www.rfc-editor.org/rfc/rfc5156.txt>
  - Don't accept prefixes longer than /48 (?)

# Receiving Prefixes: From Upstream/Transit Provider

---

- ❑ Check Team Cymru's list of "bogons"  
[www.team-cymru.org/Services/Bogons/http.html](http://www.team-cymru.org/Services/Bogons/http.html)
- ❑ For IPv4 also consult:  
[datatracker.ietf.org/doc/draft-ietf-grow-no-more-unallocated-slash8s](http://datatracker.ietf.org/doc/draft-ietf-grow-no-more-unallocated-slash8s)
- ❑ For IPv6 also consult:  
[www.space.net/~gert/RIPE/ipv6-filters.html](http://www.space.net/~gert/RIPE/ipv6-filters.html)
- ❑ Bogon Route Server:  
[www.team-cymru.org/Services/Bogons/routeserver.html](http://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
ip prefix-list in-filter deny 10.0.0.0/8 le 32
ip prefix-list in-filter deny 101.10.0.0/19 le 32      ! Block local prefix
ip prefix-list in-filter deny 127.0.0.0/8 le 32
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
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
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       ! Block multicast
ip prefix-list in-filter deny 0.0.0.0/0 ge 25         ! Block 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 deny ::/0                ! Default
  ipv6 prefix-list v6in-filter deny ::/8 le 128
  ipv6 prefix-list v6in-filter permit 2001::/32        ! Teredo
  ipv6 prefix-list v6in-filter deny 2001::/32 le 128
  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
  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 deny fc00::/7 le 128    ! Unique Local
  ipv6 prefix-list v6in-filter deny fe80::/10 le 128   ! Link Local
  ipv6 prefix-list v6in-filter deny ff00::/8 le 128    ! Multicast
  ipv6 prefix-list v6in-filter permit 2000::/3 le 48   ! Global Unicast Block
  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
<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

# BGP Best Current Practices



ISP Training Workshops