# BGP Best Current Practices

AfNOG 2012 AR-E Workshop

# 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
```

  - 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 (May 2012) >213000 /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

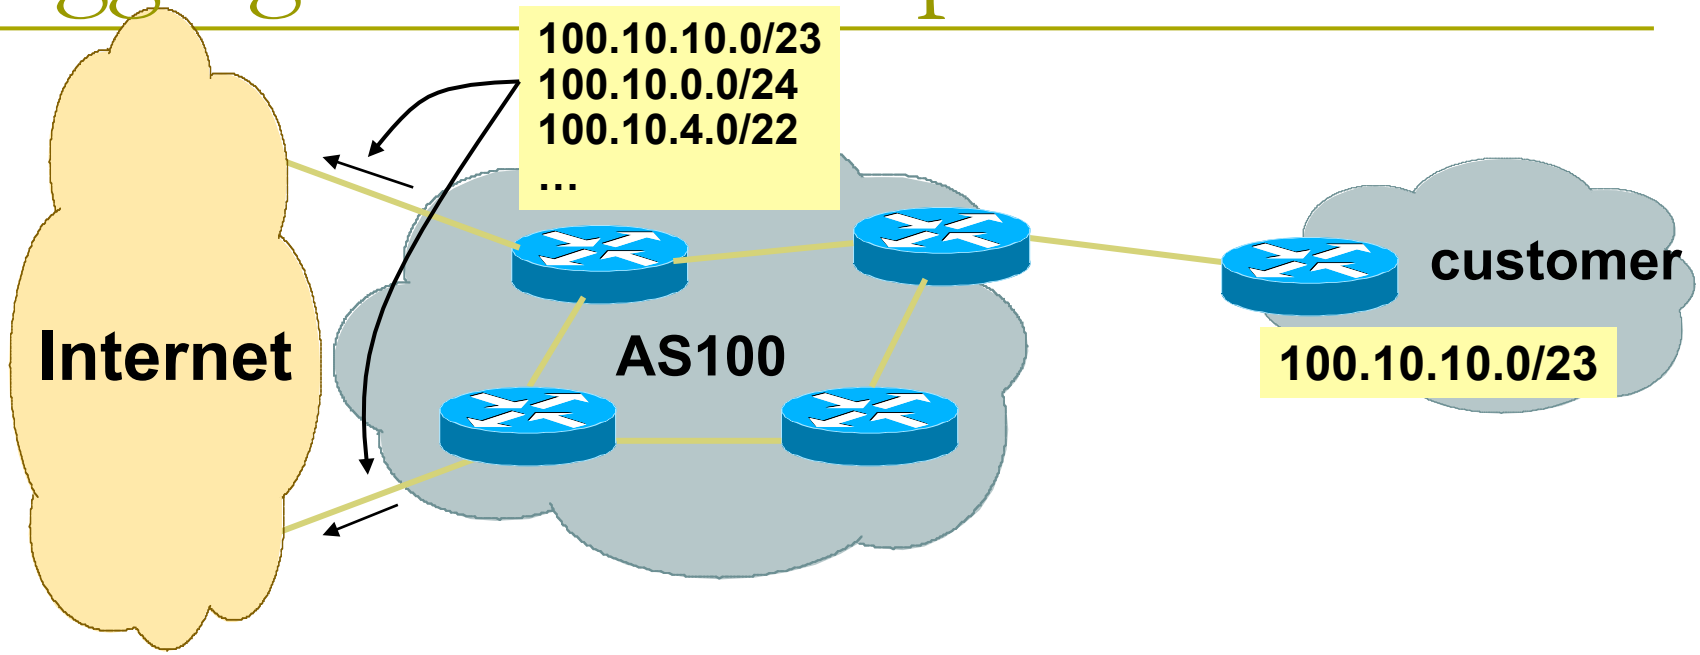


- Customer has /23 network assigned from AS100's /19 address block
- AS100 announces customers' individual networks to the Internet

# 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

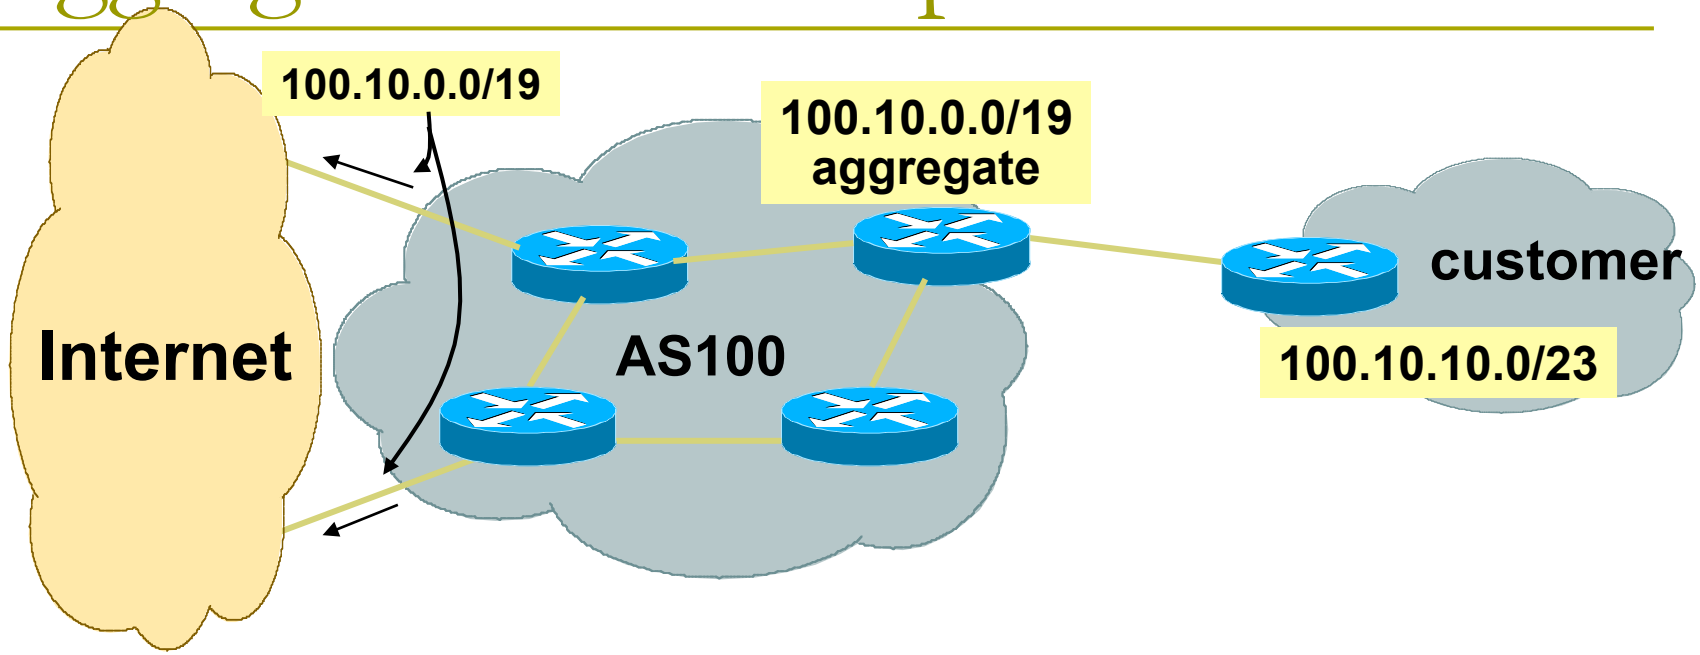


- Customer has /23 network assigned from AS100's /19 address block
- AS100 announced /19 aggregate to the Internet

# 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 (May 2012)

- Current Internet Routing Table Statistics
  - BGP Routing Table Entries 407854
  - Prefixes after maximum aggregation 173143
  - Unique prefixes in Internet 198376
  - Prefixes smaller than registry alloc 173910
  - /24s announced 213402
  - ASes in use 40876

# 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 recommendation
  - RIPE-399 — 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

CIDR Report

http://www.cidr-report.org/as2.0/

Google

Apple Smart Bookmarks Popular (537) Philip TinyURL! Networking Miscellaneous

AR-E Timetable - AfNOG 2012 CIDR Report

# Status Summary

## Table History

| Date | Prefixes | CIDR Aggregated |
|---|---|---|
| 02-05-12 | 415226 | 239890 |
| 03-05-12 | 410389 | 239955 |
| 04-05-12 | 410788 | 239947 |
| 05-05-12 | 410006 | 239925 |
| 06-05-12 | 409830 | 239992 |
| 07-05-12 | 409942 | 240125 |
| 08-05-12 | 409732 | 239935 |
| 09-05-12 | 409471 | 240133 |

Plot: BGP Table Size

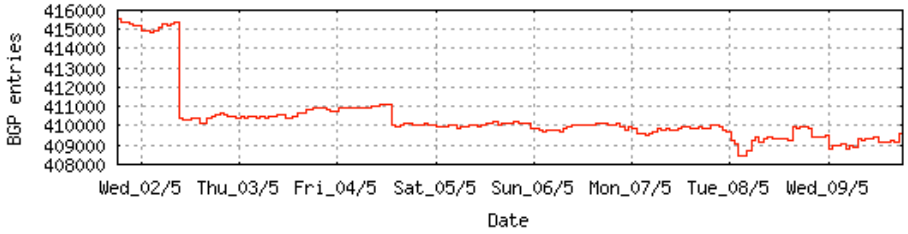


## AS Summary

| | |
|---|---|
| 41050 | Number of ASes in routing system |
| 17152 | Number of ASes announcing only one prefix |
| 3420 | Largest number of prefixes announced by an AS<br>AS6389: BELLSOUTH-NET-BLK - BellSouth.net Inc. |
| 112444384 | Largest address span announced by an AS (/32s)<br>AS4134: CHINANET-BACKBONE No.31,Jin-rong Street |

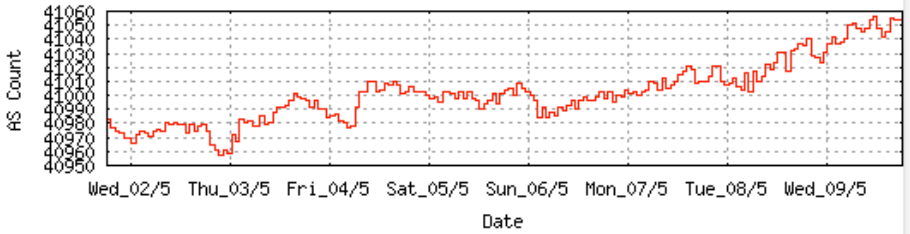


Plot: AS count
Plot: Average announcements per origin AS
Report: ASes ordered by originating address span
Report: ASes ordered by transit address span
Report: Autonomous System number-to-name mapping (from Registry WHOIS data)

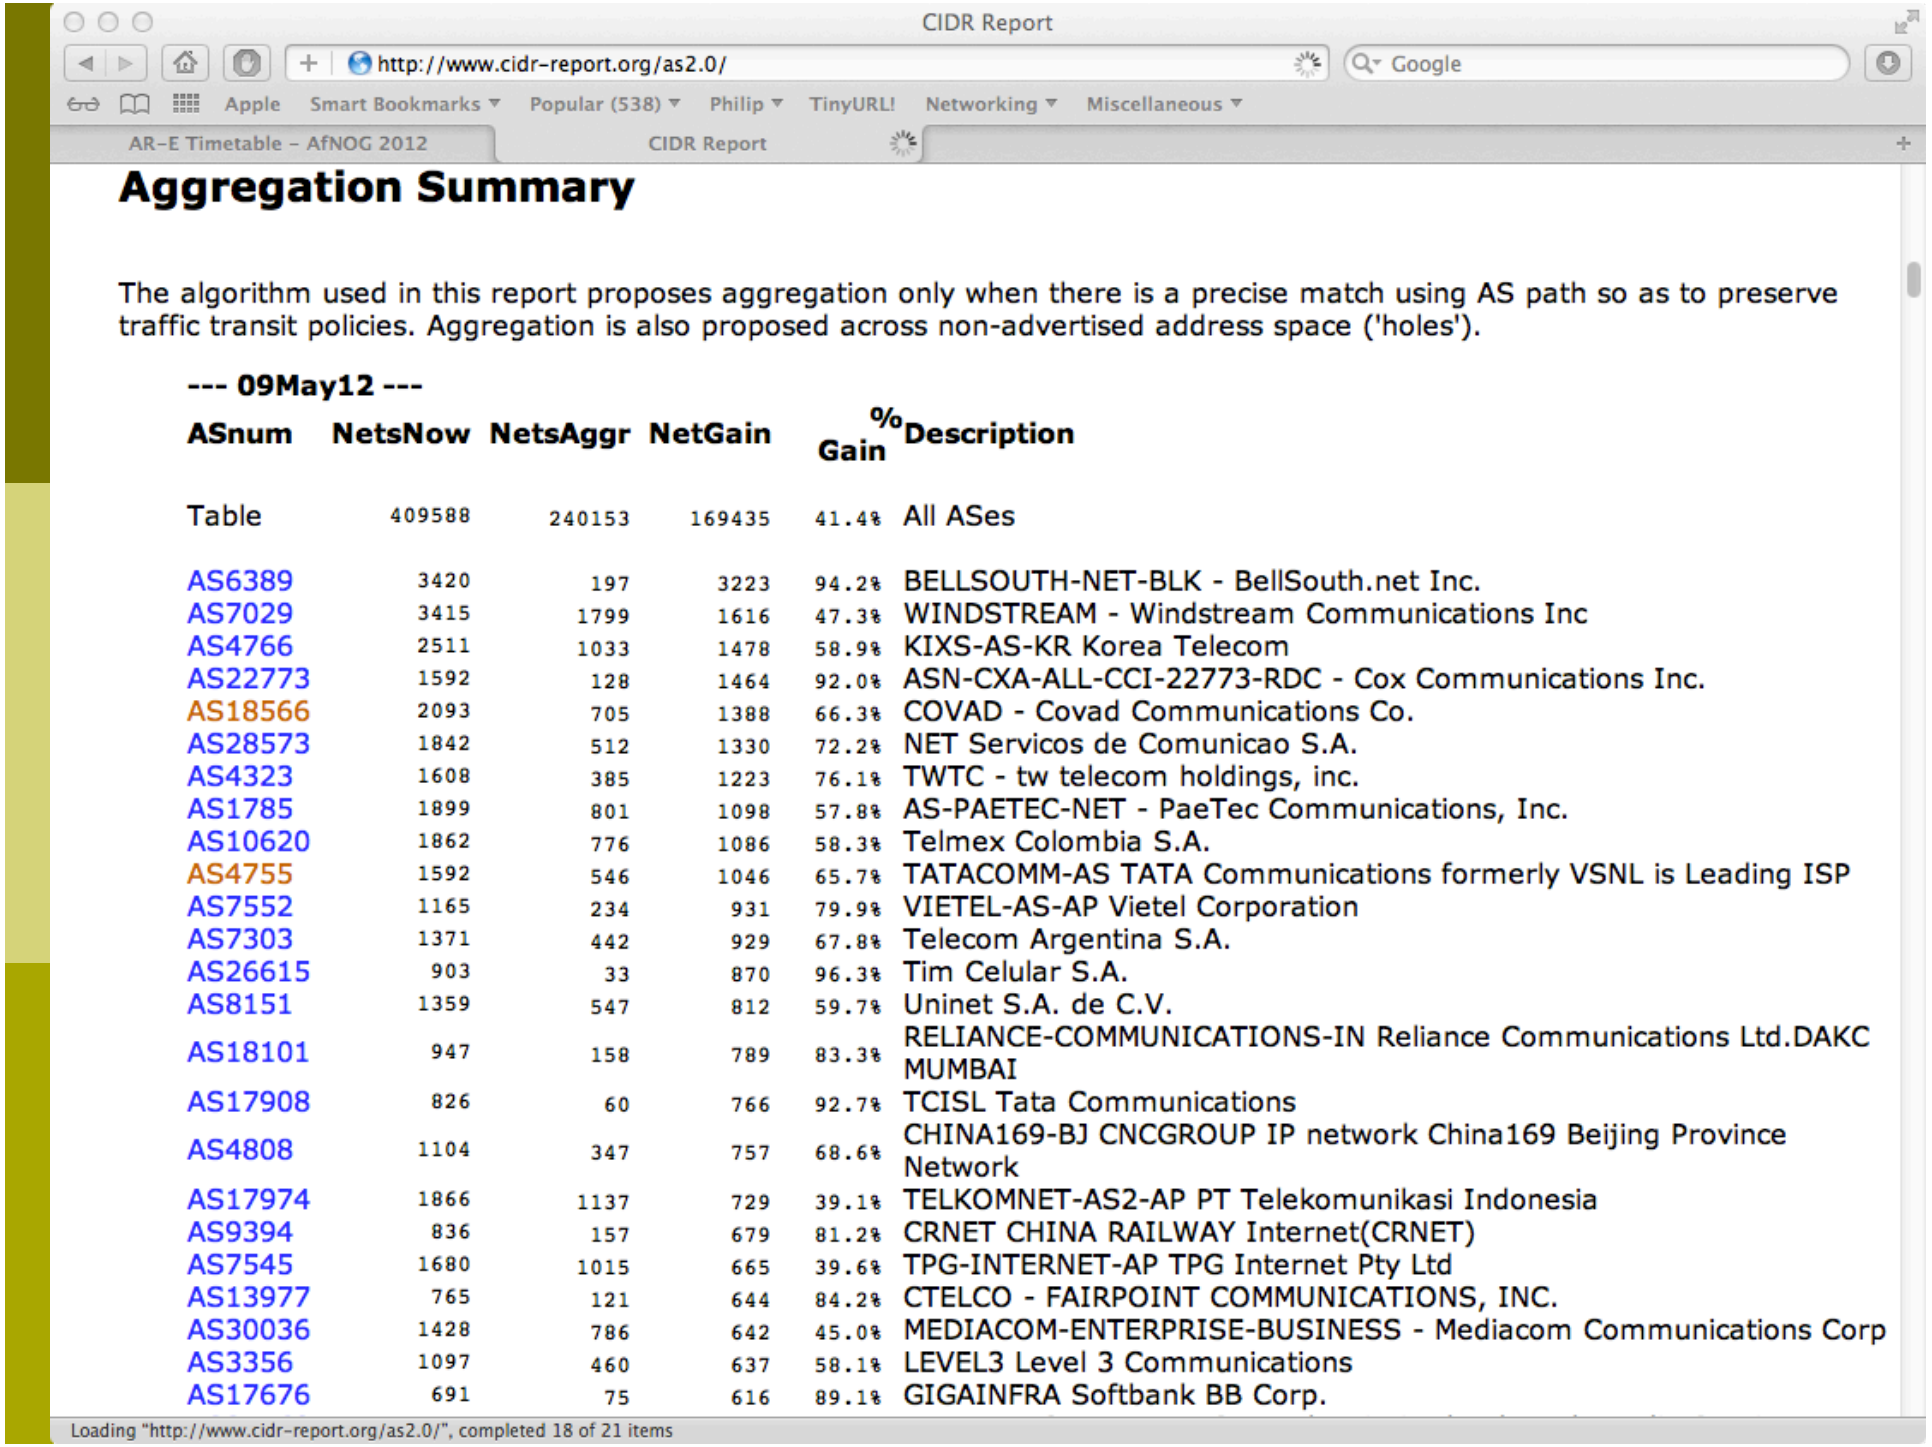


# Aggregation Summary

The algorithm used in this report proposes aggregation only when there is a precise match using AS path so as to preserve traffic transit policies. Aggregation is also proposed across non-advertised address space ('holes').

**--- 09May12 ---**

| ASnum | NetsNow | NetsAggr | NetGain | % Gain | Description |
|---|---|---|---|---|---|
| Table | 409588 | 240153 | 169435 | 41.4% | All ASes |
| AS6389 | 3420 | 197 | 3223 | 94.2% | BELLSOUTH-NET-BLK - BellSouth.net Inc. |
| AS7029 | 3415 | 1799 | 1616 | 47.3% | WINDSTREAM - Windstream Communications Inc |
| AS4766 | 2511 | 1033 | 1478 | 58.9% | KIXS-AS-KR Korea Telecom |
| AS22773 | 1592 | 128 | 1464 | 92.0% | ASN-CXA-ALL-CCI-22773-RDC - Cox Communications Inc. |
| AS18566 | 2093 | 705 | 1388 | 66.3% | COVAD - Covad Communications Co. |
| AS28573 | 1842 | 512 | 1330 | 72.2% | NET Servicos de Comunicao S.A. |
| AS4323 | 1608 | 385 | 1223 | 76.1% | TWTC - tw telecom holdings, inc. |
| AS1785 | 1899 | 801 | 1098 | 57.8% | AS-PAETEC-NET - PaeTec Communications, Inc. |
| AS10620 | 1862 | 776 | 1086 | 58.3% | Telmex Colombia S.A. |
| AS4755 | 1592 | 546 | 1046 | 65.7% | TATACOMM-AS TATA Communications formerly VSNL is Leading ISP |
| AS7552 | 1165 | 234 | 931 | 79.9% | VIETEL-AS-AP Vietel Corporation |
| AS7303 | 1371 | 442 | 929 | 67.8% | Telecom Argentina S.A. |
| AS26615 | 903 | 33 | 870 | 96.3% | Tim Celular S.A. |
| AS8151 | 1359 | 547 | 812 | 59.7% | Uninet S.A. de C.V. |
| AS18101 | 947 | 158 | 789 | 83.3% | RELIANCE-COMMUNICATIONS-IN Reliance Communications Ltd.DAKC MUMBAI |
| AS17908 | 826 | 60 | 766 | 92.7% | TCISL Tata Communications |
| AS4808 | 1104 | 347 | 757 | 68.6% | CHINA169-BJ CNCGROUP IP network China169 Beijing Province Network |
| AS17974 | 1866 | 1137 | 729 | 39.1% | TELKOMNET-AS2-AP PT Telekomunikasi Indonesia |
| AS9394 | 836 | 157 | 679 | 81.2% | CRNET CHINA RAILWAY Internet(CRNET) |
| AS7545 | 1680 | 1015 | 665 | 39.6% | TPG-INTERNET-AP TPG Internet Pty Ltd |
| AS13977 | 765 | 121 | 644 | 84.2% | CTELCO - FAIRPOINT COMMUNICATIONS, INC. |
| AS30036 | 1428 | 786 | 642 | 45.0% | MEDIACOM-ENTERPRISE-BUSINESS - Mediacom Communications Corp |
| AS3356 | 1097 | 460 | 637 | 58.1% | LEVEL3 Level 3 Communications |
| AS17676 | 691 | 75 | 616 | 89.1% | GIGAINFRA Softbank BB Corp. |

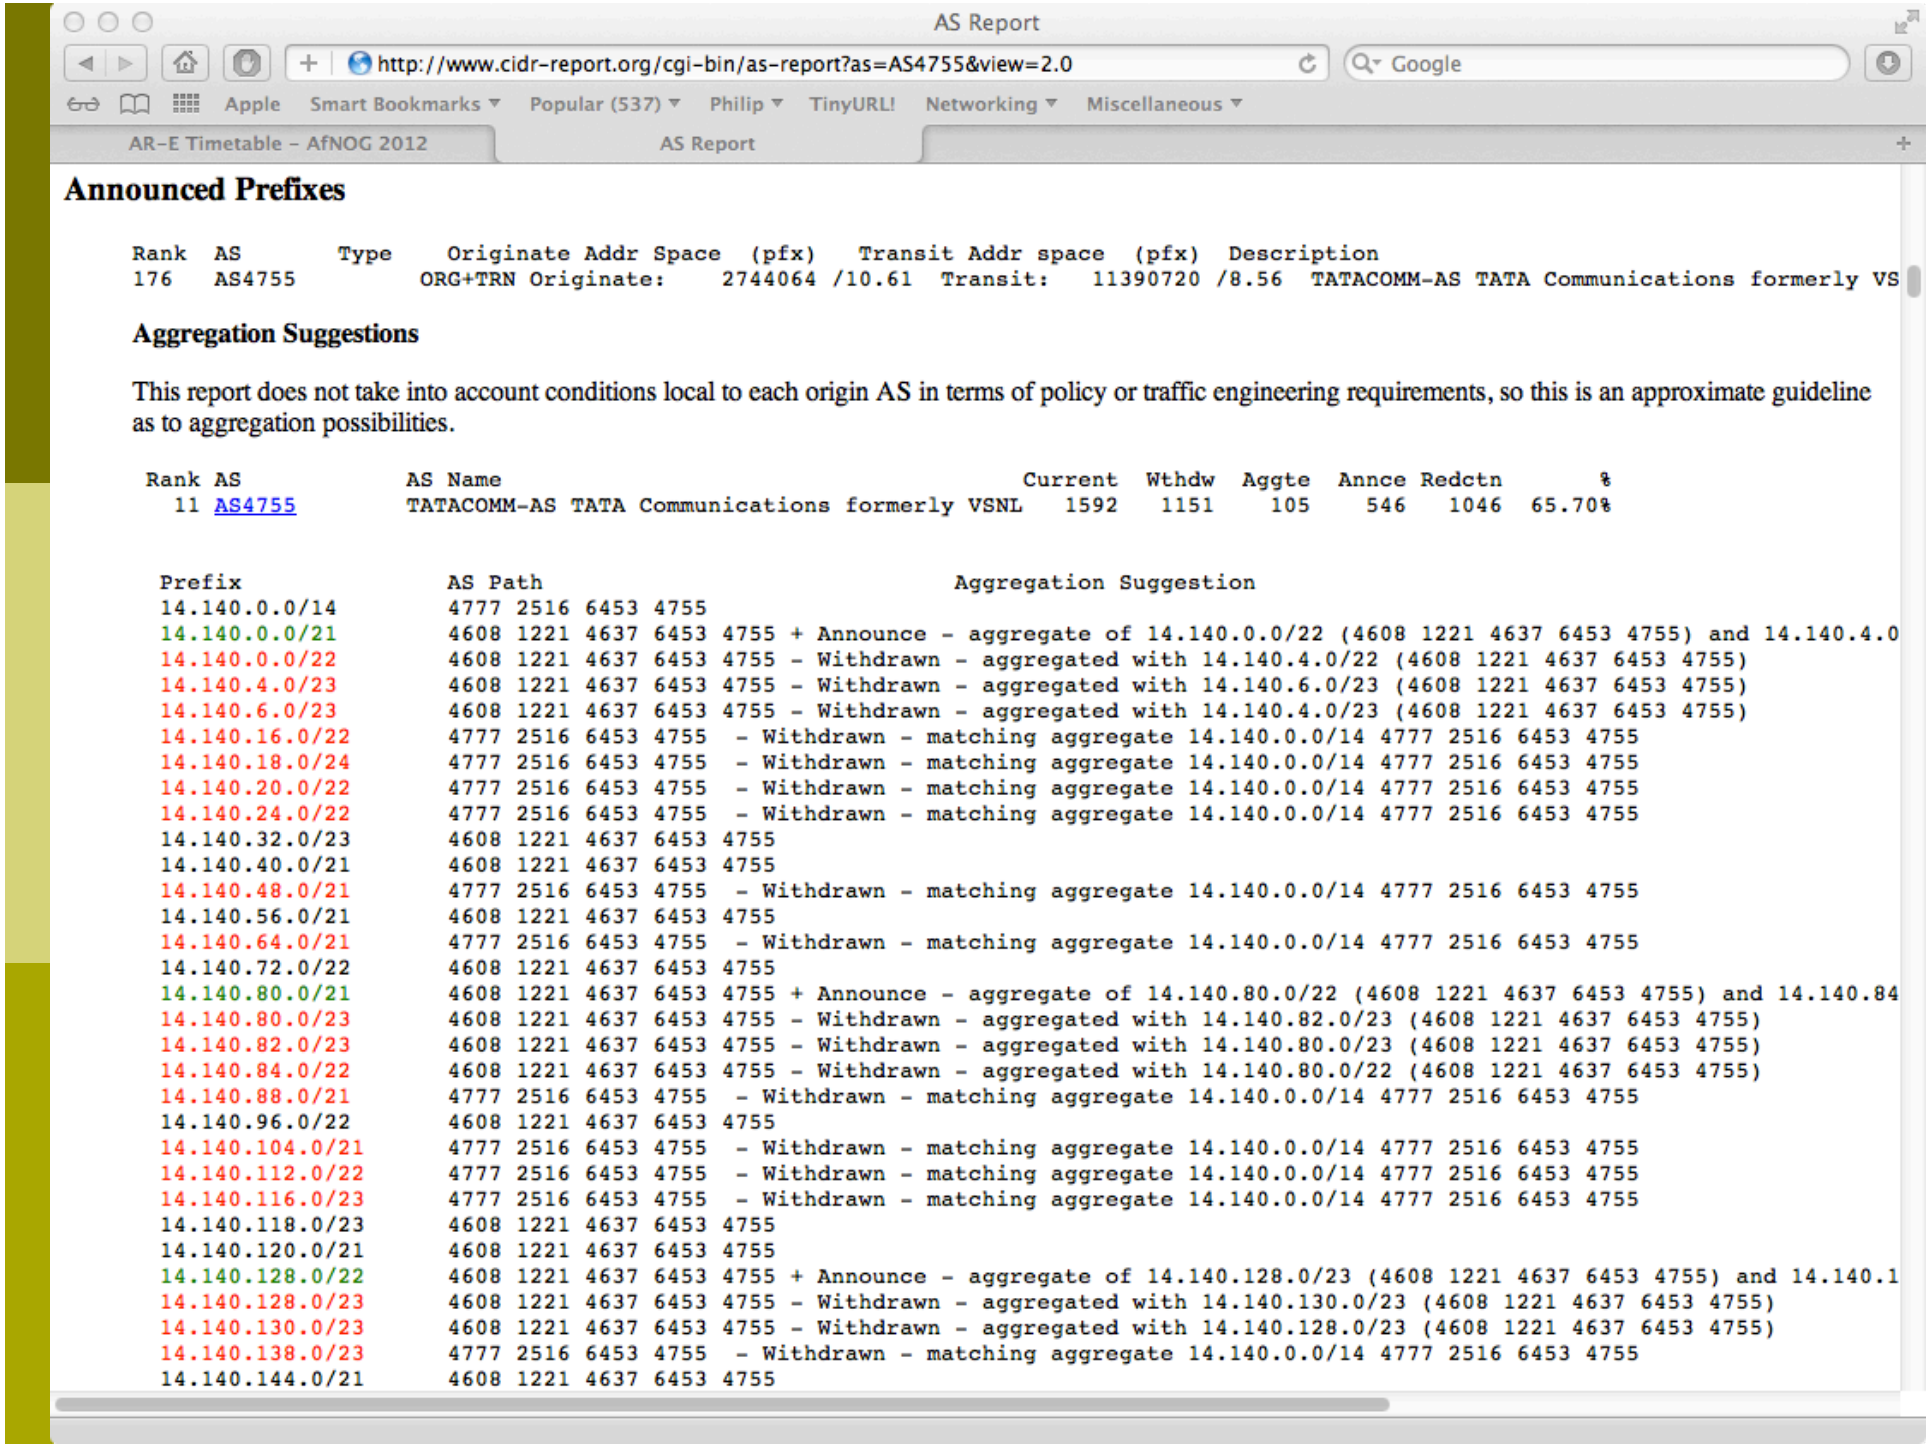


## Announced Prefixes

```
Rank  AS       Type    Originate Addr Space  (pfx)   Transit Addr space  (pfx)  Description
176   AS4755         ORG+TRN Originate:    2744064 /10.61  Transit:   11390720 /8.56  TATACOMM-AS TATA Communications formerly VS
```

### 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       %
  11 AS4755        TATACOMM-AS TATA Communications formerly VSNL   1592   1151    105    546   1046  65.70%


 Prefix              AS Path                              Aggregation Suggestion
 14.140.0.0/14       4777 2516 6453 4755
 14.140.0.0/21       4608 1221 4637 6453 4755 + Announce - aggregate of 14.140.0.0/22 (4608 1221 4637 6453 4755) and 14.140.4.0
 14.140.0.0/22       4608 1221 4637 6453 4755 - Withdrawn - aggregated with 14.140.4.0/22 (4608 1221 4637 6453 4755)
 14.140.4.0/23       4608 1221 4637 6453 4755 - Withdrawn - aggregated with 14.140.6.0/23 (4608 1221 4637 6453 4755)
 14.140.6.0/23       4608 1221 4637 6453 4755 - Withdrawn - aggregated with 14.140.4.0/23 (4608 1221 4637 6453 4755)
 14.140.16.0/22      4777 2516 6453 4755  - Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
 14.140.18.0/24      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      4608 1221 4637 6453 4755
 14.140.40.0/21      4608 1221 4637 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      4608 1221 4637 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      4608 1221 4637 6453 4755
 14.140.80.0/21      4608 1221 4637 6453 4755 + Announce - aggregate of 14.140.80.0/22 (4608 1221 4637 6453 4755) and 14.140.84
 14.140.80.0/23      4608 1221 4637 6453 4755 - Withdrawn - aggregated with 14.140.82.0/23 (4608 1221 4637 6453 4755)
 14.140.82.0/23      4608 1221 4637 6453 4755 - Withdrawn - aggregated with 14.140.80.0/23 (4608 1221 4637 6453 4755)
 14.140.84.0/22      4608 1221 4637 6453 4755 - Withdrawn - aggregated with 14.140.80.0/22 (4608 1221 4637 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      4608 1221 4637 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     4608 1221 4637 6453 4755
 14.140.120.0/21     4608 1221 4637 6453 4755
 14.140.128.0/22     4608 1221 4637 6453 4755 + Announce - aggregate of 14.140.128.0/23 (4608 1221 4637 6453 4755) and 14.140.1
 14.140.128.0/23     4608 1221 4637 6453 4755 - Withdrawn - aggregated with 14.140.130.0/23 (4608 1221 4637 6453 4755)
 14.140.130.0/23     4608 1221 4637 6453 4755 - Withdrawn - aggregated with 14.140.128.0/23 (4608 1221 4637 6453 4755)
 14.140.138.0/23     4777 2516 6453 4755  - Withdrawn - matching aggregate 14.140.0.0/14 4777 2516 6453 4755
 14.140.144.0/21     4608 1221 4637 6453 4755
```

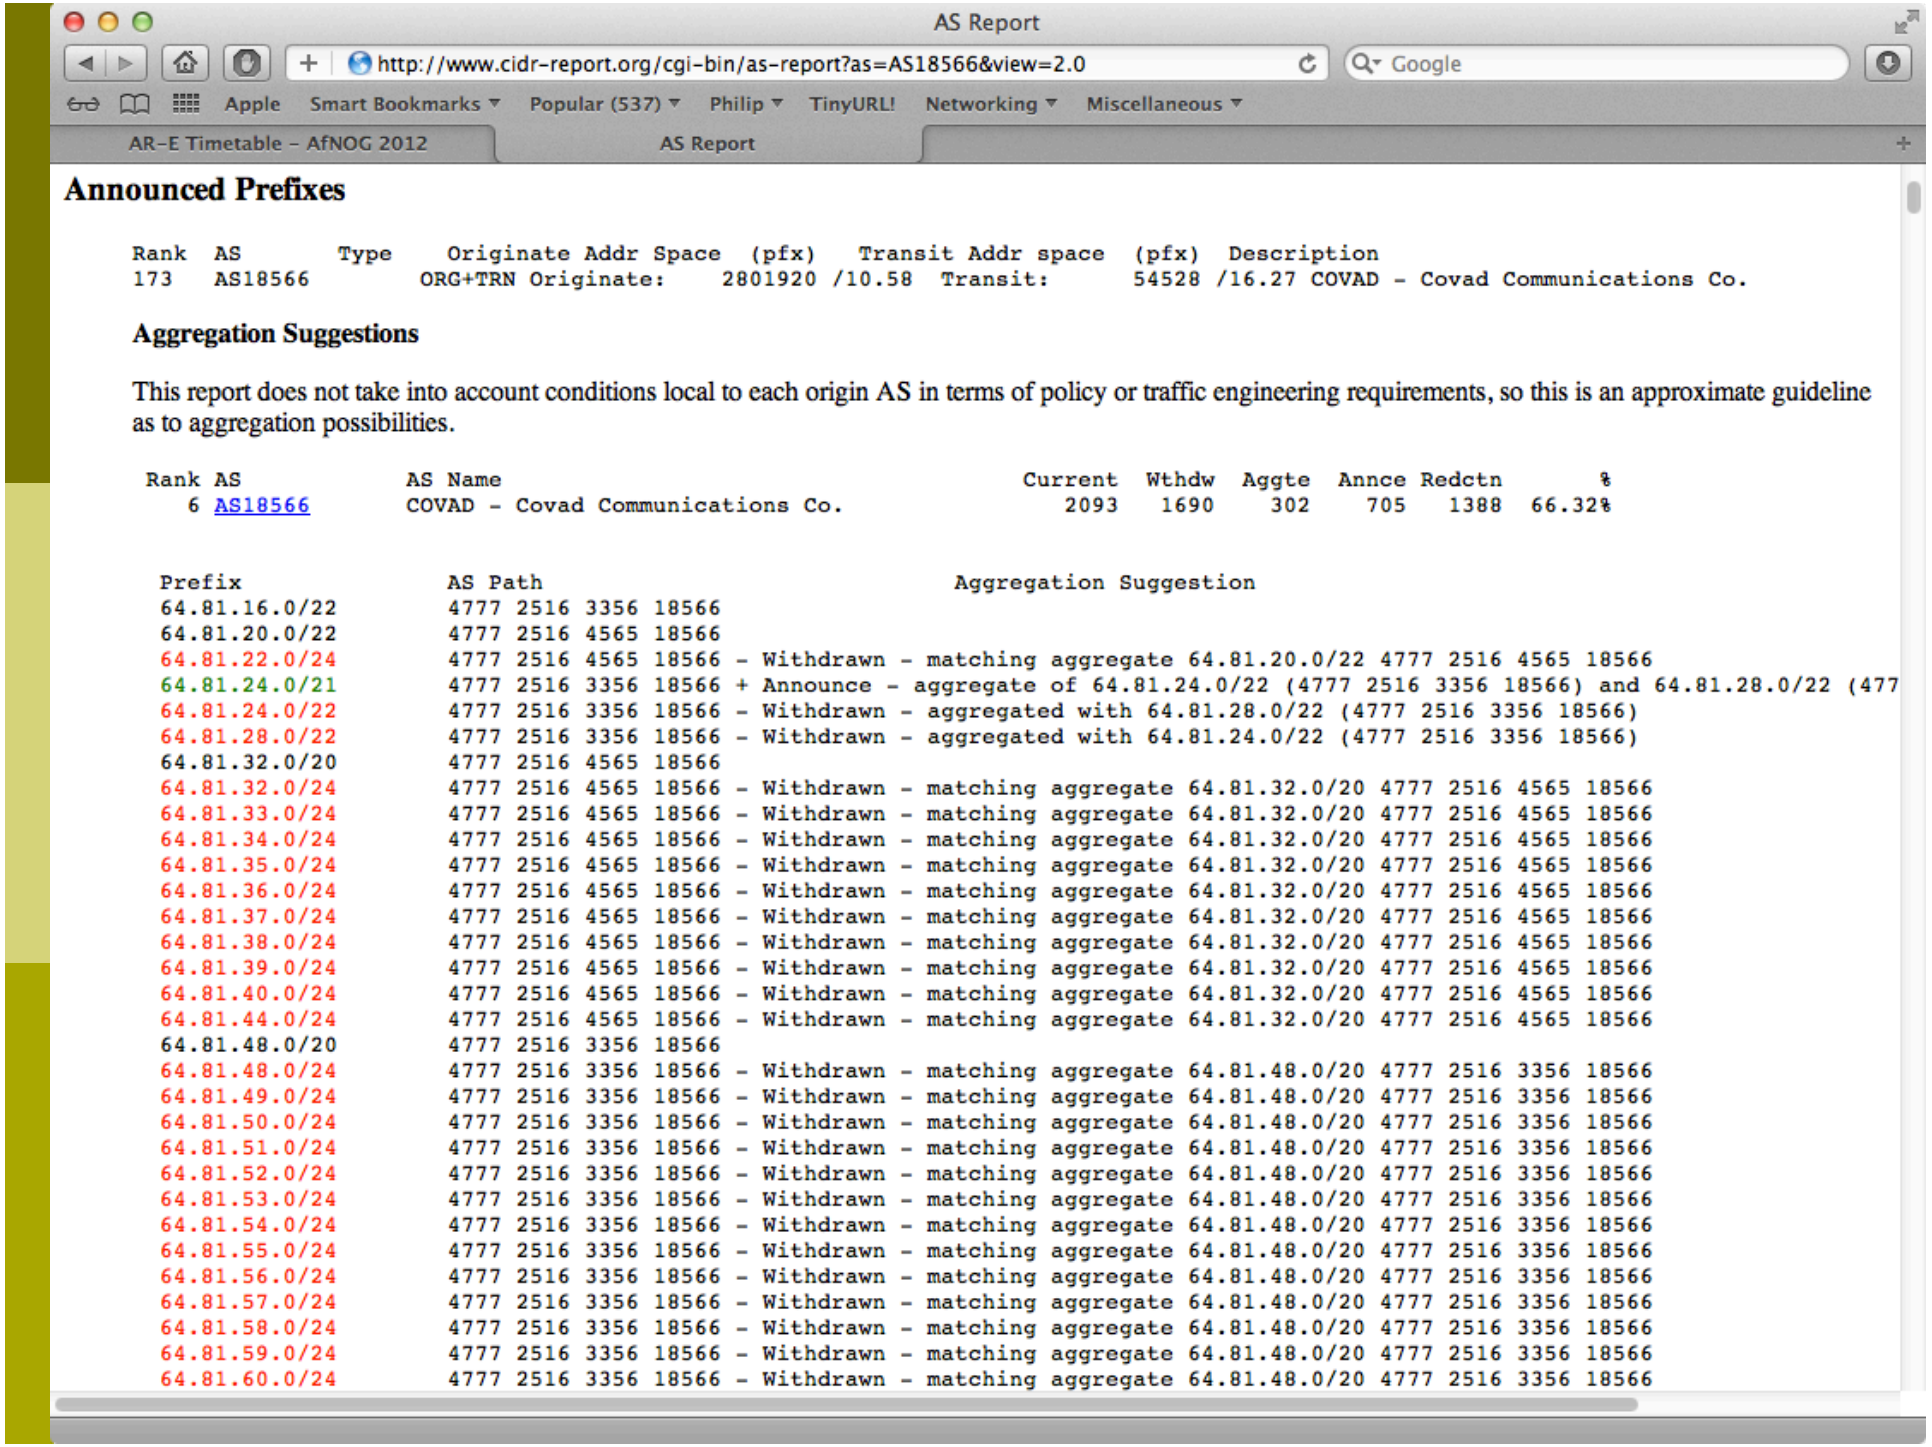


## Announced Prefixes

```
Rank  AS       Type    Originate Addr Space  (pfx)   Transit Addr space  (pfx)  Description
173   AS18566        ORG+TRN Originate:    2801920 /10.58  Transit:      54528 /16.27 COVAD - Covad Communications Co.
```

### 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       %
   6 AS18566      COVAD - Covad Communications Co.            2093   1690    302    705   1388  66.32%
```

```
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 (477
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
```

# 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

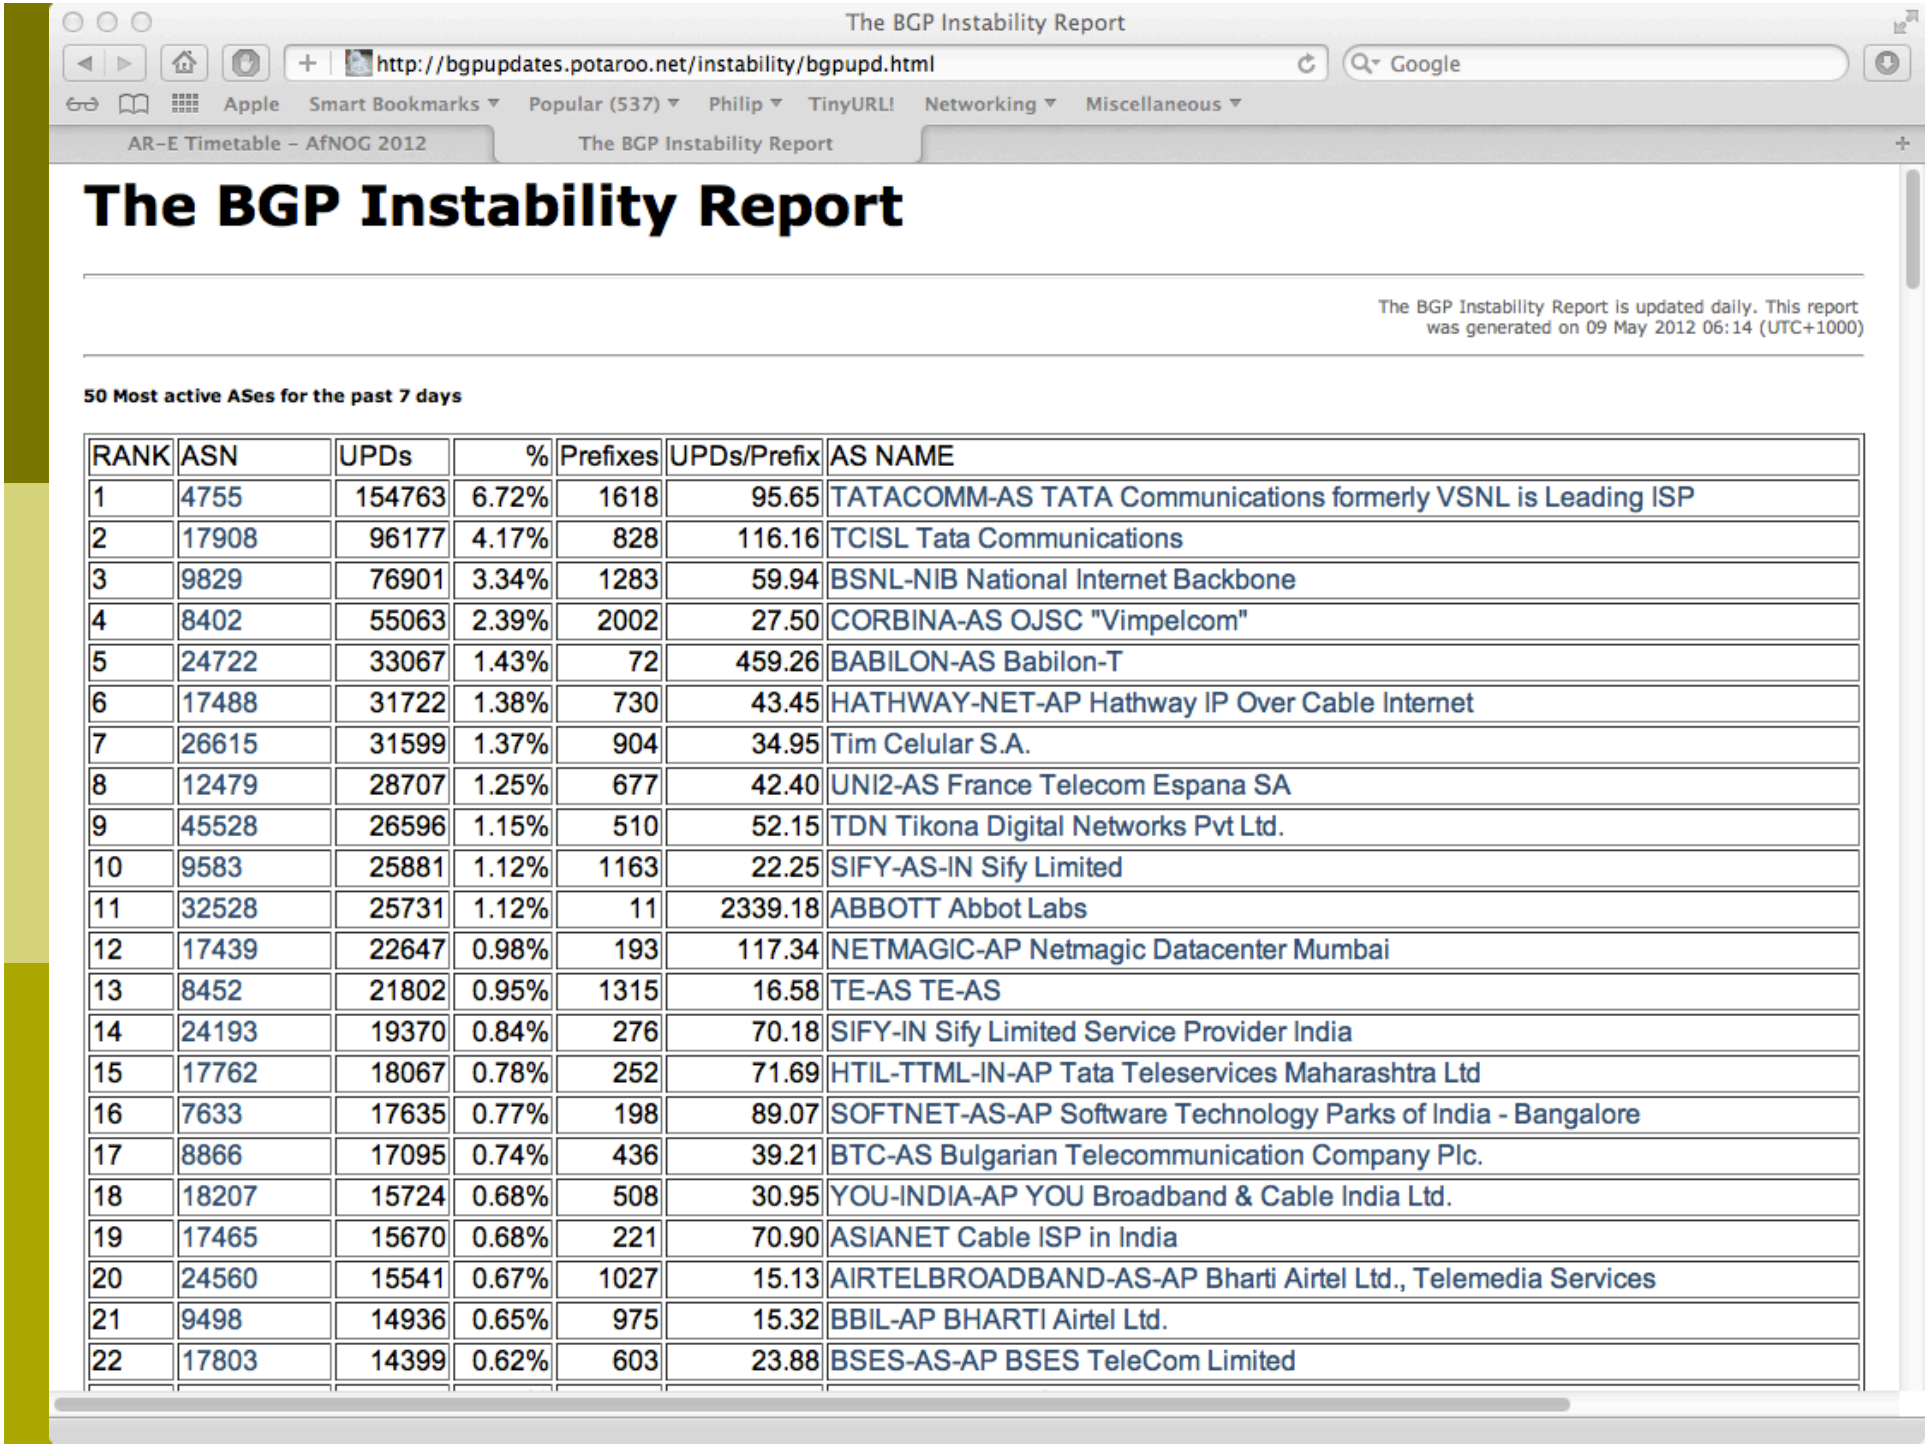


# The BGP Instability Report

The BGP Instability Report is updated daily. This report was generated on 09 May 2012 06:14 (UTC+1000)

**50 Most active ASes for the past 7 days**

| RANK | ASN | UPDs | % | Prefixes | UPDs/Prefix | AS NAME |
|---|---|---|---|---|---|---|
| 1 | 4755 | 154763 | 6.72% | 1618 | 95.65 | TATACOMM-AS TATA Communications formerly VSNL is Leading ISP |
| 2 | 17908 | 96177 | 4.17% | 828 | 116.16 | TCISL Tata Communications |
| 3 | 9829 | 76901 | 3.34% | 1283 | 59.94 | BSNL-NIB National Internet Backbone |
| 4 | 8402 | 55063 | 2.39% | 2002 | 27.50 | CORBINA-AS OJSC "Vimpelcom" |
| 5 | 24722 | 33067 | 1.43% | 72 | 459.26 | BABILON-AS Babilon-T |
| 6 | 17488 | 31722 | 1.38% | 730 | 43.45 | HATHWAY-NET-AP Hathway IP Over Cable Internet |
| 7 | 26615 | 31599 | 1.37% | 904 | 34.95 | Tim Celular S.A. |
| 8 | 12479 | 28707 | 1.25% | 677 | 42.40 | UNI2-AS France Telecom Espana SA |
| 9 | 45528 | 26596 | 1.15% | 510 | 52.15 | TDN Tikona Digital Networks Pvt Ltd. |
| 10 | 9583 | 25881 | 1.12% | 1163 | 22.25 | SIFY-AS-IN Sify Limited |
| 11 | 32528 | 25731 | 1.12% | 11 | 2339.18 | ABBOTT Abbot Labs |
| 12 | 17439 | 22647 | 0.98% | 193 | 117.34 | NETMAGIC-AP Netmagic Datacenter Mumbai |
| 13 | 8452 | 21802 | 0.95% | 1315 | 16.58 | TE-AS TE-AS |
| 14 | 24193 | 19370 | 0.84% | 276 | 70.18 | SIFY-IN Sify Limited Service Provider India |
| 15 | 17762 | 18067 | 0.78% | 252 | 71.69 | HTIL-TTML-IN-AP Tata Teleservices Maharashtra Ltd |
| 16 | 7633 | 17635 | 0.77% | 198 | 89.07 | SOFTNET-AS-AP Software Technology Parks of India - Bangalore |
| 17 | 8866 | 17095 | 0.74% | 436 | 39.21 | BTC-AS Bulgarian Telecommunication Company Plc. |
| 18 | 18207 | 15724 | 0.68% | 508 | 30.95 | YOU-INDIA-AP YOU Broadband & Cable India Ltd. |
| 19 | 17465 | 15670 | 0.68% | 221 | 70.90 | ASIANET Cable ISP in India |
| 20 | 24560 | 15541 | 0.67% | 1027 | 15.13 | AIRTELBROADBAND-AS-AP Bharti Airtel Ltd., Telemedia Services |
| 21 | 9498 | 14936 | 0.65% | 975 | 15.32 | BBIL-AP BHARTI Airtel Ltd. |
| 22 | 17803 | 14399 | 0.62% | 603 | 23.88 | BSES-AS-AP BSES TeleCom Limited |

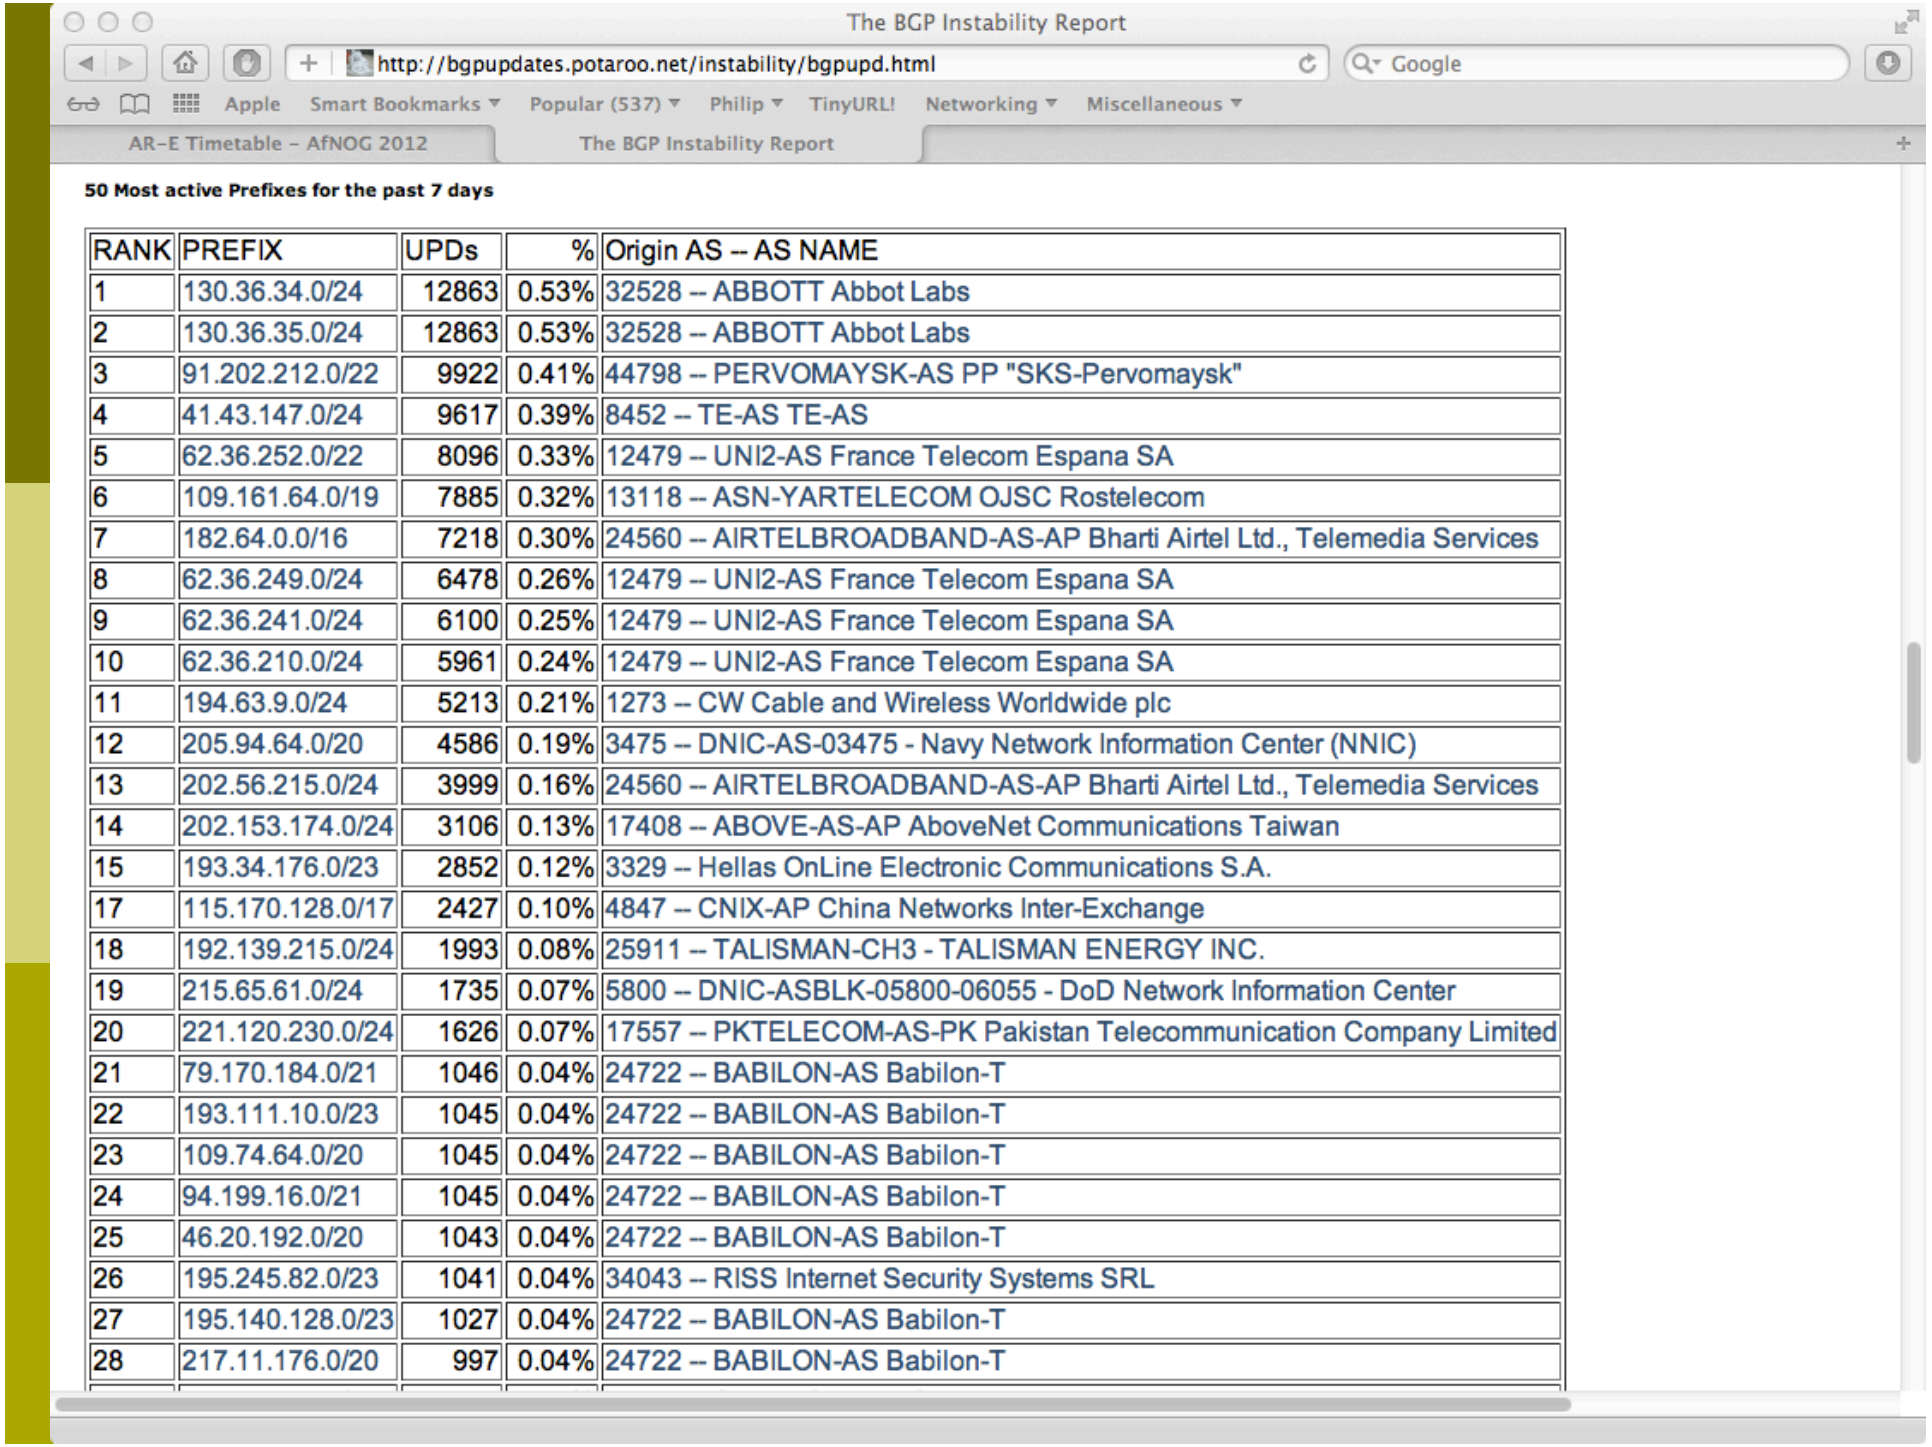


**50 Most active Prefixes for the past 7 days**

| RANK | PREFIX | UPDs | % | Origin AS -- AS NAME |
|---|---|---|---|---|
| 1 | 130.36.34.0/24 | 12863 | 0.53% | 32528 -- ABBOTT Abbot Labs |
| 2 | 130.36.35.0/24 | 12863 | 0.53% | 32528 -- ABBOTT Abbot Labs |
| 3 | 91.202.212.0/22 | 9922 | 0.41% | 44798 -- PERVOMAYSK-AS PP "SKS-Pervomaysk" |
| 4 | 41.43.147.0/24 | 9617 | 0.39% | 8452 -- TE-AS TE-AS |
| 5 | 62.36.252.0/22 | 8096 | 0.33% | 12479 -- UNI2-AS France Telecom Espana SA |
| 6 | 109.161.64.0/19 | 7885 | 0.32% | 13118 -- ASN-YARTELECOM OJSC Rostelecom |
| 7 | 182.64.0.0/16 | 7218 | 0.30% | 24560 -- AIRTELBROADBAND-AS-AP Bharti Airtel Ltd., Telemedia Services |
| 8 | 62.36.249.0/24 | 6478 | 0.26% | 12479 -- UNI2-AS France Telecom Espana SA |
| 9 | 62.36.241.0/24 | 6100 | 0.25% | 12479 -- UNI2-AS France Telecom Espana SA |
| 10 | 62.36.210.0/24 | 5961 | 0.24% | 12479 -- UNI2-AS France Telecom Espana SA |
| 11 | 194.63.9.0/24 | 5213 | 0.21% | 1273 -- CW Cable and Wireless Worldwide plc |
| 12 | 205.94.64.0/20 | 4586 | 0.19% | 3475 -- DNIC-AS-03475 - Navy Network Information Center (NNIC) |
| 13 | 202.56.215.0/24 | 3999 | 0.16% | 24560 -- AIRTELBROADBAND-AS-AP Bharti Airtel Ltd., Telemedia Services |
| 14 | 202.153.174.0/24 | 3106 | 0.13% | 17408 -- ABOVE-AS-AP AboveNet Communications Taiwan |
| 15 | 193.34.176.0/23 | 2852 | 0.12% | 3329 -- Hellas OnLine Electronic Communications S.A. |
| 17 | 115.170.128.0/17 | 2427 | 0.10% | 4847 -- CNIX-AP China Networks Inter-Exchange |
| 18 | 192.139.215.0/24 | 1993 | 0.08% | 25911 -- TALISMAN-CH3 - TALISMAN ENERGY INC. |
| 19 | 215.65.61.0/24 | 1735 | 0.07% | 5800 -- DNIC-ASBLK-05800-06055 - DoD Network Information Center |
| 20 | 221.120.230.0/24 | 1626 | 0.07% | 17557 -- PKTELECOM-AS-PK Pakistan Telecommunication Company Limited |
| 21 | 79.170.184.0/21 | 1046 | 0.04% | 24722 -- BABILON-AS Babilon-T |
| 22 | 193.111.10.0/23 | 1045 | 0.04% | 24722 -- BABILON-AS Babilon-T |
| 23 | 109.74.64.0/20 | 1045 | 0.04% | 24722 -- BABILON-AS Babilon-T |
| 24 | 94.199.16.0/21 | 1045 | 0.04% | 24722 -- BABILON-AS Babilon-T |
| 25 | 46.20.192.0/20 | 1043 | 0.04% | 24722 -- BABILON-AS Babilon-T |
| 26 | 195.245.82.0/23 | 1041 | 0.04% | 34043 -- RISS Internet Security Systems SRL |
| 27 | 195.140.128.0/23 | 1027 | 0.04% | 24722 -- BABILON-AS Babilon-T |
| 28 | 217.11.176.0/20 | 997 | 0.04% | 24722 -- BABILON-AS Babilon-T |

# 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**
    - eg: **whois –h jwhois.apnic.net x.x.x.0/24**

# 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/**

# 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

- For IPv4 also consult:

  datatracker.ietf.org/doc/draft-vegoda-no-more-unallocated-slash8s

- 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
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

# Scaling the network

How to get out of carrying all prefixes in IGP

# Why use BGP rather than IGP?

- IGP has Limitations:
  - The more routing information in the network
    - Periodic updates/flooding “overload”
    - Long convergence times
    - Affects the core first
  - Policy definition
    - Not easy to do

# Preparing the Network

- We want to deploy BGP now…
- BGP will be used therefore an ASN is required
- If multihoming to different ISPs is intended in the near future, a public ASN should be obtained:
  - Either go to upstream ISP who is a registry member, or
  - Apply to the RIR yourself for a one off assignment, or
  - Ask an ISP who is a registry member, or
  - **Join the RIR and get your own IP address allocation too (this option strongly recommended)!**

# Preparing the Network Initial Assumptions

- The network is not running any BGP at the moment
  - single statically routed connection to upstream ISP
- The network is not running any IGP at all
  - Static default and routes through the network to do “routing”

# Preparing the Network
# First Step: IGP

- Decide on an IGP: OSPF or ISIS ☺
- Assign loopback interfaces and /32 address to each router which will run the IGP
  - Loopback is used for OSPF and BGP router id anchor
  - Used for iBGP and route origination
- Deploy IGP (e.g. OSPF)
  - IGP can be deployed with NO IMPACT on the existing static routing
  - e.g. OSPF distance might be 110; static distance is 1
  - Smallest distance wins

# Preparing the Network IGP (cont)

- Be prudent deploying IGP – keep the Link State Database Lean!
  - Router loopbacks go in IGP
  - WAN point to point links go in IGP
  - (In fact, any link where IGP dynamic routing will be run should go into IGP)
  - Summarise on area/level boundaries (if possible) – i.e. think about your IGP address plan

# Preparing the Network IGP (cont)

- Routes which don't go into the IGP include:
  - Dynamic assignment pools (DSL/Cable/Dial)
  - Customer point to point link addressing
    - (using next-hop-self in iBGP ensures that these do NOT need to be in IGP)
  - Static/Hosting LANs
  - Customer assigned address space
  - Anything else not listed in the previous slide

# Preparing the Network Second Step: iBGP

- Second step is to configure the local network to use iBGP
- iBGP can run on
  - all routers, or
  - a subset of routers, or
  - just on the upstream edge
- ***iBGP must run on all routers which are in the transit path between external connections***

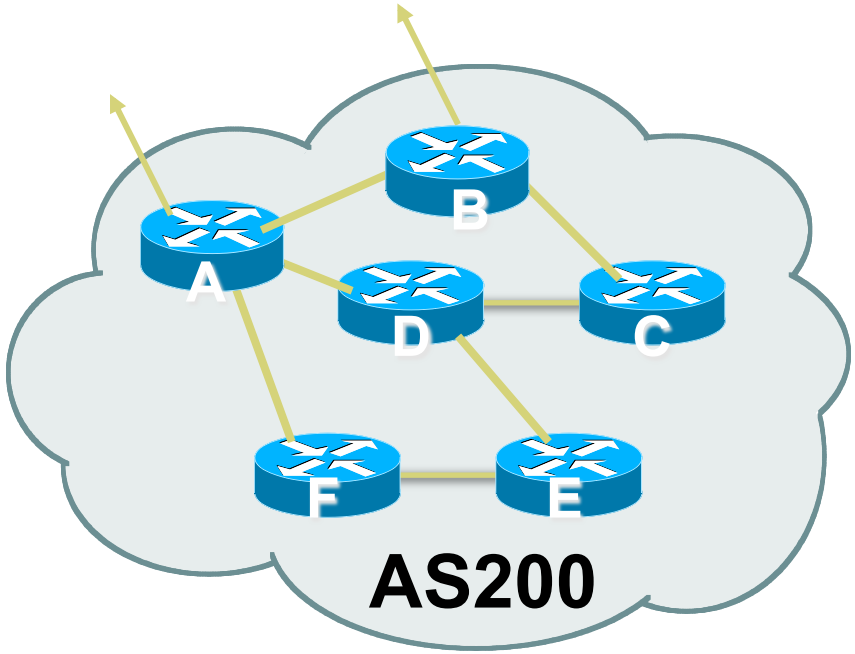

# Preparing the Network
# Second Step: iBGP (Transit Path)

- ***iBGP must run on all routers which are in the transit path between external connections***
- Routers C, E and F are not in the transit path
  - Static routes or IGP will suffice
- Router D is in the transit path
  - Will need to be in iBGP mesh, otherwise routing loops will result

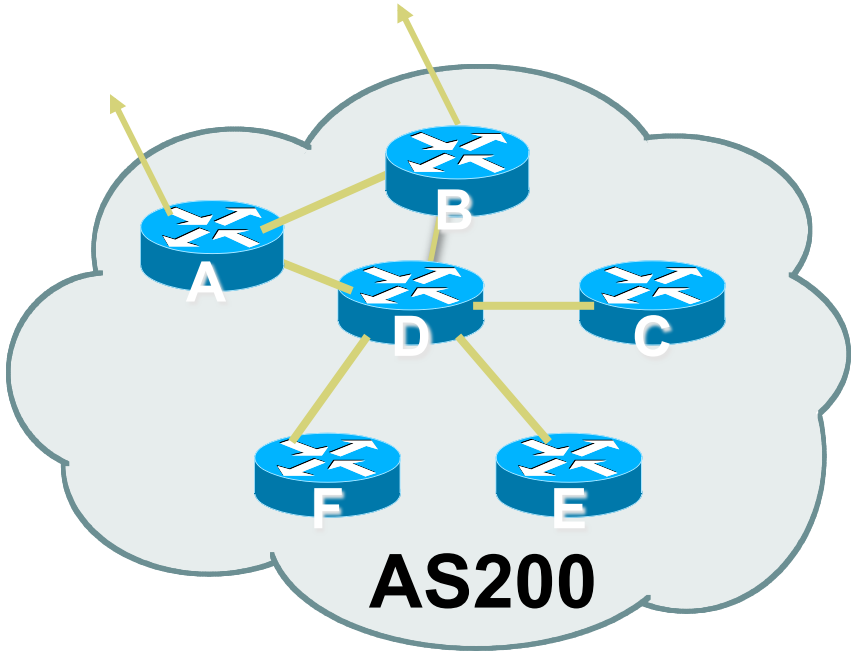

# Preparing the Network Layers

- Typical SP networks have three layers:
  - Core – the backbone, usually the transit path
  - Distribution – the middle, PoP aggregation layer
  - Aggregation – the edge, the devices connecting customers

# Preparing the Network Aggregation Layer

- iBGP is optional
  - Many ISPs run iBGP here, either partial routing (more common) or full routing (less common)
  - Full routing is not needed unless customers want full table
  - Partial routing is cheaper/easier, might usually consist of internal prefixes and, optionally, external prefixes to aid external load balancing
    - Communities and peer-groups make this administratively easy
- Many aggregation devices can't run iBGP
  - Static routes from distribution devices for address pools
  - IGP for best exit

# Preparing the Network Distribution Layer

- Usually runs iBGP
  - Partial or full routing (as with aggregation layer)
- But does not have to run iBGP
  - IGP is then used to carry customer prefixes (does not scale)
  - IGP is used to determine nearest exit
- Networks which plan to grow large should deploy iBGP from day one
  - Migration at a later date is extra work
  - No extra overhead in deploying iBGP, indeed IGP benefits

# Preparing the Network
# Core Layer

- Core of network is usually the transit path
- iBGP necessary between core devices
  - Full routes or partial routes:
    - Transit ISPs carry full routes in core
    - Edge ISPs carry partial routes only
- Core layer includes AS border routers

# Preparing the Network iBGP Implementation

Decide on:

- Best iBGP policy
  - Will it be full routes everywhere, or partial, or some mix?
- iBGP scaling technique
  - Community policy?
  - Route-reflectors?
  - Techniques such as peer groups and peer templates?

# Preparing the Network iBGP Implementation

- Then deploy iBGP:
  - Step 1: Introduce iBGP mesh on chosen routers
    - make sure that iBGP distance is greater than IGP distance (it usually is)
  - Step 2: Install “customer” prefixes into iBGP

    **Check!** Does the network still work?
  - Step 3: Carefully remove the static routing for the prefixes now in IGP and iBGP

    **Check!** Does the network still work?
  - Step 4: Deployment of eBGP follows

# Preparing the Network iBGP Implementation

***Install "customer" prefixes into iBGP?***

- Customer assigned address space
  - Network statement/static route combination
  - Use unique community to identify customer assignments
- Customer facing point-to-point links
  - Redistribute connected through filters which only permit point-to-point link addresses to enter iBGP
  - Use a unique community to identify point-to-point link addresses (these are only required for your monitoring system)
- Dynamic assignment pools & local LANs
  - Simple network statement will do this
  - Use unique community to identify these networks

# Preparing the Network iBGP Implementation

***Carefully remove static routes?***

- Work on one router at a time:
  - Check that static route for a particular destination is also learned by the iBGP
  - If so, remove it
  - If not, establish why and fix the problem
  - (Remember to look in the RIB, not the FIB!)
- Then the next router, until the whole PoP is done
- Then the next PoP, and so on until the network is now dependent on the IGP and iBGP you have deployed

# Preparing the Network Completion

- Previous steps are NOT flag day steps
  - Each can be carried out during different maintenance periods, for example:
  - Step One on Week One
  - Step Two on Week Two
  - Step Three on Week Three
  - And so on
  - And with proper planning will have NO customer visible impact at all

# Preparing the Network Configuration Summary

- IGP essential networks are in IGP
- Customer networks are now in iBGP
  - iBGP deployed over the backbone
  - Full or Partial or Upstream Edge only
- BGP distance is greater than any IGP
- Now ready to deploy eBGP

# BGP Best Current Practices

AfNOG 2012 AR-E Workshop