Multihoming Case Study

ISP Workshops

Last updated 10 October 2007

Multihoming Case Study

- Set of slides based on work assisting an ISP with their multihoming needs between 2000 and 2002
  - Should be taken as an indicative example only

Case Study

First Visit (2000)

Case Study – Requirements (1)

- ISP needs to multihome:
  - To AS5400 in Europe
  - To AS2516 in Japan
  - /19 allocated by APNIC
  - AS 17660 assigned by APNIC
  - 1Mbps circuits to both upstreams

Case Study – Requirements (2)

- ISP wants:
  - Symmetric routing and equal link utilisation in and out (as close as possible)
    - international circuits are expensive
  - Has two Cisco 2600 border routers with 64Mbytes memory
    - Cannot afford to upgrade memory or hardware on border routers or internal routers
- “Philip, make it work, please”

Case Study

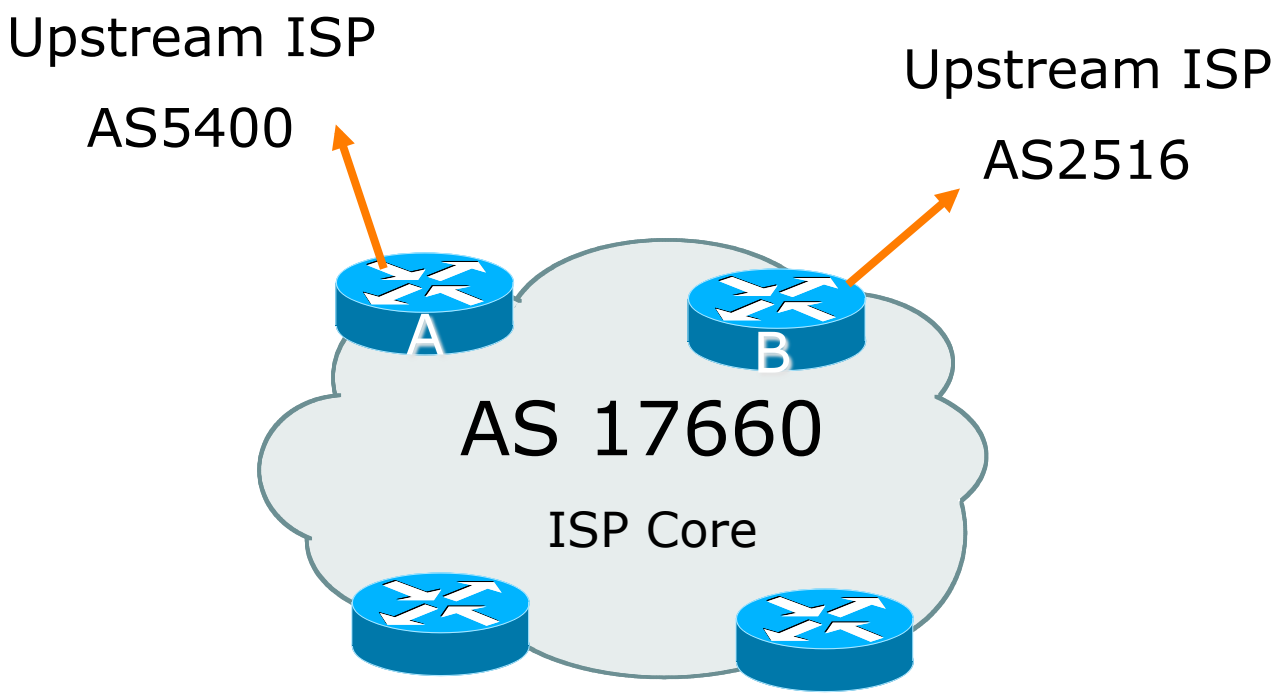

Allocated /19 from APNIC

Circuit to AS5400 is 1Mbps, circuit to AS2516 is 1Mbps

Case Study

- Both providers stated that routers with 128Mbytes memory required for AS17660 to multihome
  - Those myths again ☹
  - Full routing table is rarely required or desired
- Solution:
  - Accept default from one upstream
  - Accept partial prefixes from the other

Case Study – Inbound Loadsharing

- First cut: Went to a few US Looking Glasses
  - Checked the AS path to AS5400
  - Checked the AS path to AS2516
  - AS2516 was one hop “closer”
  - Sent AS-PATH prepend of one AS on AS2516 peering

Case Study – Inbound Loadsharing

- Refinement
  - Did not need any
  - First cut worked, seeing on average 600kbps inbound on each circuit
  - Does vary according to time of day, but this is as balanced as it can get, given customer profile
  - ☺

Case Study – Outbound Loadsharing

- First cut:
  - Requested default from AS2516
  - Requested full routes from AS5400
- Then looked at my Routing Report
  - Picked the top 5 ASNs and created a filter-list
    - If 701, 1, 7018, 1239 or 7046 are in AS-PATH, prefixes are discarded
    - Allowed prefixes originated by AS5400 and up to two AS hops away
  - Resulted in 32000 prefixes being accepted in AS17660

Case Study – Outbound Loadsharing

- Refinement
  - 32000 prefixes quite a lot, seeing more outbound traffic on the AS5400 path
  - Traffic was very asymmetric
    - out through AS5400, in through AS2516
  - Added the next 3 ASNs from the Top 20 list
    - 209, 2914 and 3549
  - Now seeing 14000 prefixes
  - Traffic is now evenly loadshared outbound
    - Around 200kbps on average
    - Mostly symmetric

Case Study
MRTG Graphs

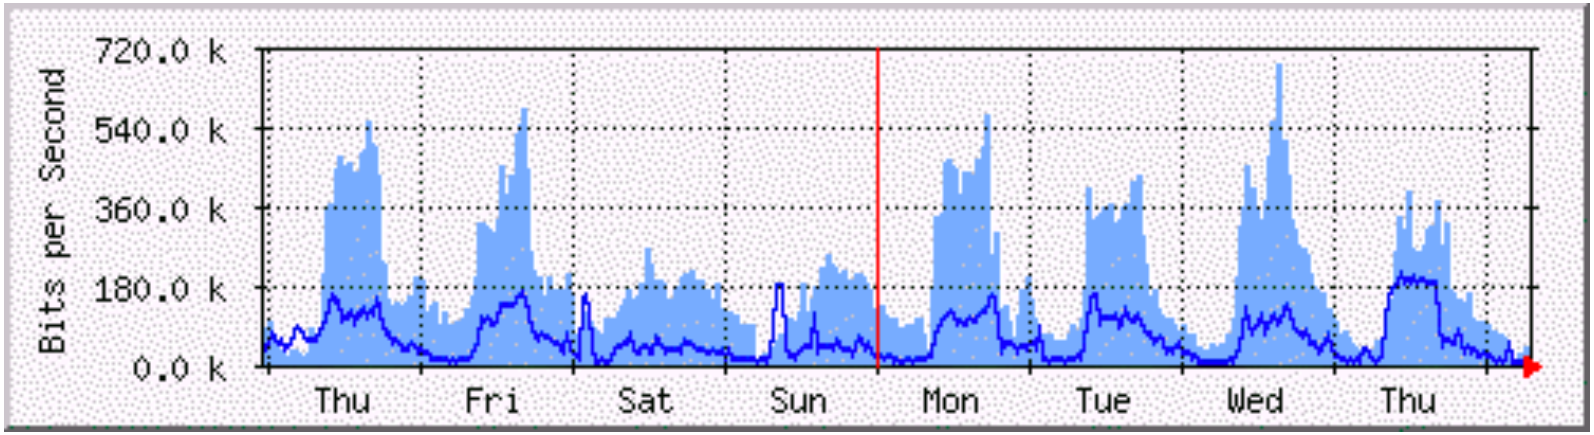

Router A to AS5400

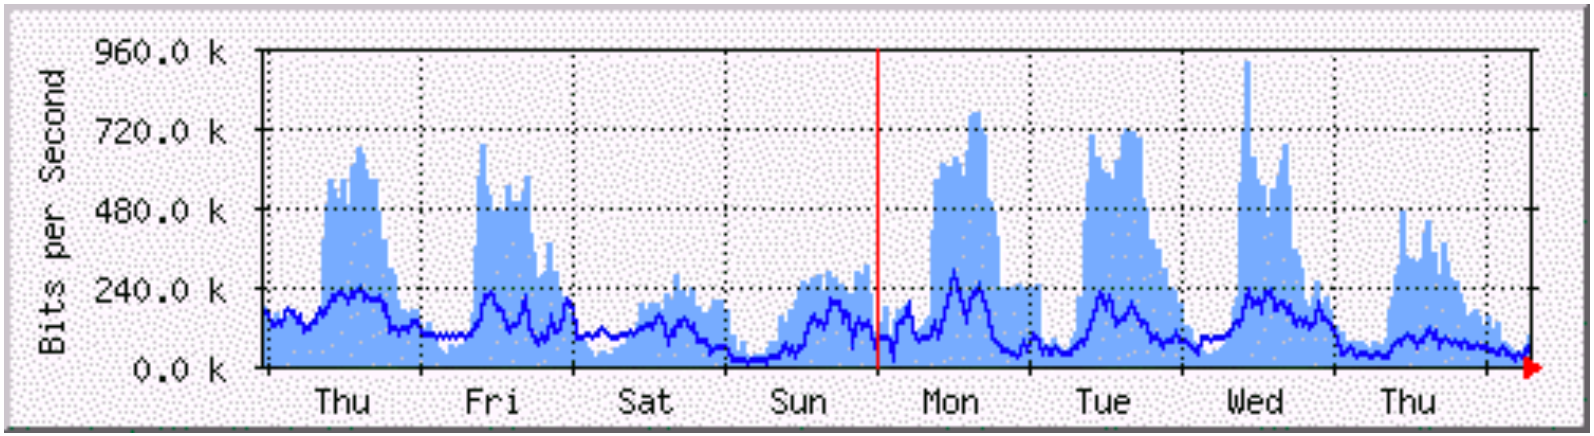

Router B to AS2516

Case Study
Configuration Router A

```
router ospf 100
 log-adjacency-changes
 passive-interface default
 no passive-interface Ethernet0/0
 default-information originate metric 20
!
router bgp 17660
 no synchronization
 no bgp fast-external-fallover
 bgp log-neighbor-changes
 bgp deterministic-med
```

...next slide

Case Study
Configuration Router A

```
neighbor 166.49.165.13 remote-as 5400
neighbor 166.49.165.13 descr eBGP multihop to AS5400
neighbor 166.49.165.13 ebgp-multihop 5
neighbor 166.49.165.13 update-source Loopback0
neighbor 166.49.165.13 prefix-list in-filter in
neighbor 166.49.165.13 prefix-list out-filter out
neighbor 166.49.165.13 filter-list 1 in
neighbor 166.49.165.13 filter-list 3 out
!
prefix-list in-filter  deny rfc1918etc in
prefix-list out-filter permit 202.144.128.0/19
!
ip route 0.0.0.0 0.0.0.0 serial 0/0 254
```

...next slide

Case Study
Configuration Router A

```
ip as-path access-list 1 deny _701_
ip as-path access-list 1 deny _1_
ip as-path access-list 1 deny _7018_
ip as-path access-list 1 deny _1239_
ip as-path access-list 1 deny _7046_
ip as-path access-list 1 deny _209_
ip as-path access-list 1 deny _2914_
ip as-path access-list 1 deny _3549_
ip as-path access-list 1 permit _5400$
ip as-path access-list 1 permit _5400_[0-9]+$
ip as-path access-list 1 permit _5400_[0-9]+_[0-9]+$
ip as-path access-list 1 deny .*
ip as-path access-list 3 permit ^$
!
```

Case Study
Configuration Router B

```
router ospf 100
 log-adjacency-changes
 passive-interface default
 no passive-interface Ethernet0/0
 default-information originate
!
router bgp 17660
 no synchronization
 no auto-summary
 no bgp fast-external-fallover
```

...next slide

Case Study
Configuration Router B

```
 bgp log-neighbor-changes
 bgp deterministic-med
 neighbor 210.132.92.165 remote-as 2516
 neighbor 210.132.92.165 description eBGP peering
 neighbor 210.132.92.165 soft-reconfiguration inbound
 neighbor 210.132.92.165 prefix-list default-route in
 neighbor 210.132.92.165 prefix-list out-filter out
 neighbor 210.132.92.165 route-map as2516-out out
 neighbor 210.132.92.165 maximum-prefix 100
 neighbor 210.132.92.165 filter-list 2 in
 neighbor 210.132.92.165 filter-list 3 out
!
```

...next slide

Case Study
Configuration Router B

```
!
prefix-list default-route permit 0.0.0.0/0
prefix-list out-filter permit 202.144.128.0/19
!
ip as-path access-list 2 permit _2516$
ip as-path access-list 2 deny .*
ip as-path access-list 3 permit ^$
!
route-map as2516-out permit 10
 set as-path prepend 17660
!
```

Configuration Summary

- Router A
  - Hears full routing table – throws away most of it
  - AS5400 BGP options are all or nothing
  - Static default pointing to serial interface – if link goes down, OSPF default removed
- Router B
  - Hears default from AS2516
  - If default disappears (BGP goes down or link goes down), OSPF default is removed

Case Study Summary

- Multihoming is not hard, really!
  - Needs a bit of thought, a bit of planning
  - Use this case study as an example strategy
  - Does not require sophisticated equipment, big memory, fast CPUs…

Case Study

Second Visit (2002)

Case Study – Current Status

- ISP currently multihomes:
  - To AS5400 in the UK
  - To AS2516 in Japan
  - /19 allocated by APNIC
  - AS 17660 assigned by APNIC
  - 1Mbps circuits to both upstreams

Case Study – Requirements

- ISP wants:
  - To add a new satellite connection, a 640K link to AS22351 in Germany to support the AS5400 link to UK
  - Still want symmetric routing and equal link utilisation in and out (as close as possible)
    - international circuits are expensive
  - Has upgraded to two Cisco 3725 border routers with plenty of memory
- Despite the working previous configuration with “sparse routing table”, wanted full prefixes
- Talked them out of that, and here is how…

Case Study

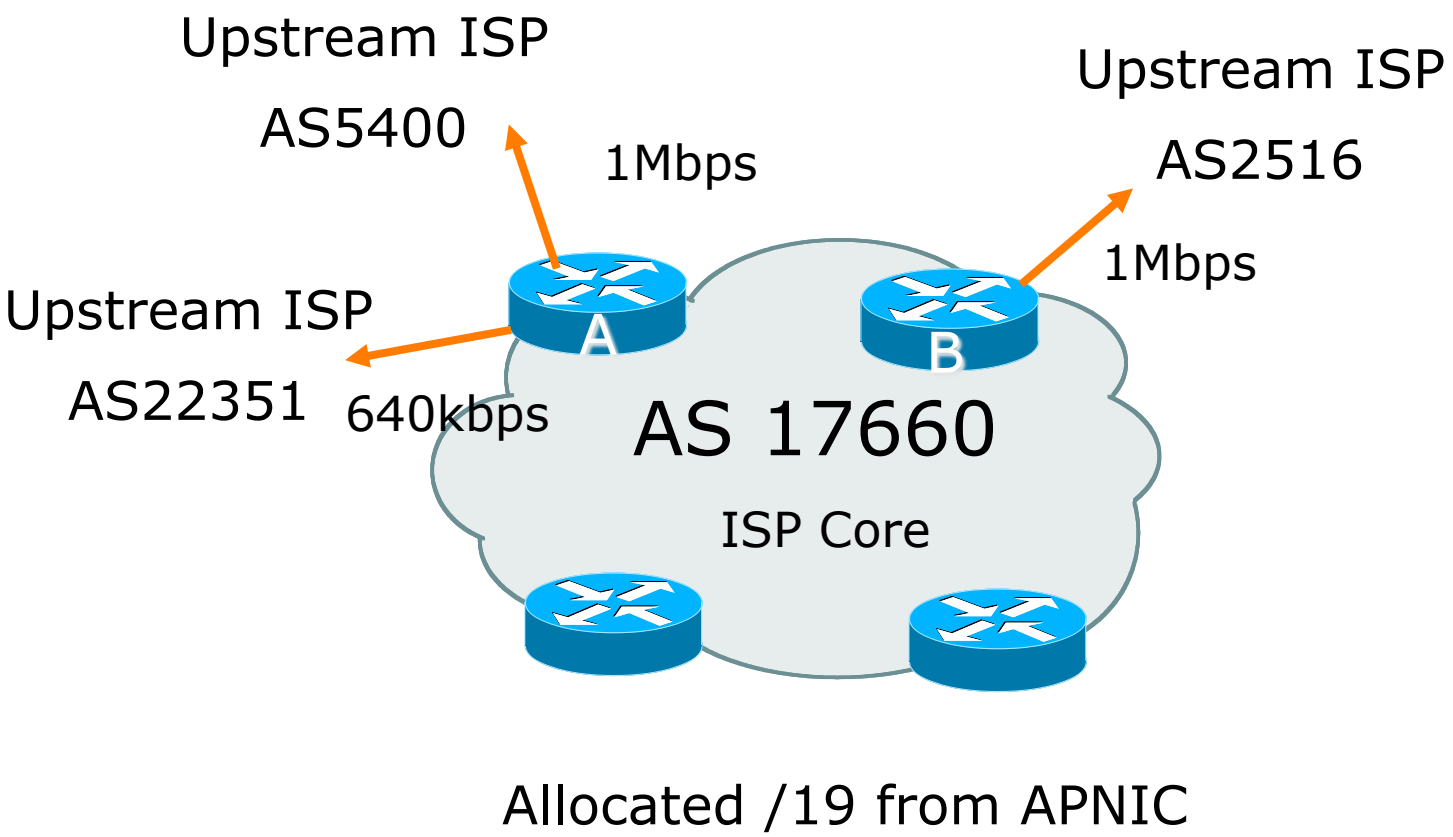

Allocated /19 from APNIC

Case Study – Inbound Loadsharing

- First cut: Went to a few US Looking Glasses
  - Checked the AS path to AS5400
  - Checked the AS path to AS2516
  - Checked the AS path to AS22351
  - AS2516 was one hop “closer” than the other two
  - Sent AS-PATH prepend of one AS on AS2516 peering
  - this is unchanged from two years ago

Case Study – Inbound Loadsharing

- Refinement
  - Needed some – AS5400 seemed to be always preferred over AS22351
  - AS5400 now supports RFC1998 style communities for customer use
    - see whois –h whois.ripe.net AS5400
  - Sent AS5400 some communities to insert prepends towards specific peers
    - Now saw some traffic on AS22351 link but not much
  - Sent a /23 announcement out AS22351 link
    - Now saw more traffic on AS22351 link

Case Study – Inbound Loadsharing

- Results:
  - Around 600kbps on the AS5400 link
  - Around 750kbps on the AS2516 link
  - Around 300kbps on the AS22351 link
  - Inbound traffic fluctuates quite substantially based on time of day
- Status:
  - Situation left pending monitoring by the ISP's NOC

Case Study – Outbound Loadsharing

- First cut:
  - Already receiving default from AS2516
  - Receiving full routes from AS5400
  - Requested full routes from AS22351 – the only option
- Retained the AS5400 configuration
  - Discard prefixes which had top 5 ASNs in the path
- AS22351 configuration uses similar ideas to AS5400 configuration
  - But only accepted prefixes originated from AS22351 or their immediate peers

Case Study – Outbound Loadsharing

- Results:
  - Around 35000 prefixes from AS5400
  - Around 2000 prefixes from AS22351
  - Around 200kbps on both the AS5400 and AS2516 links
  - Around 50kbps on the AS22351 link
  - Outbound traffic fluctuates quite substantially based on time of day
- Status:
  - Situation left pending monitoring by the ISP's NOC

Case Study
MRTG Graphs

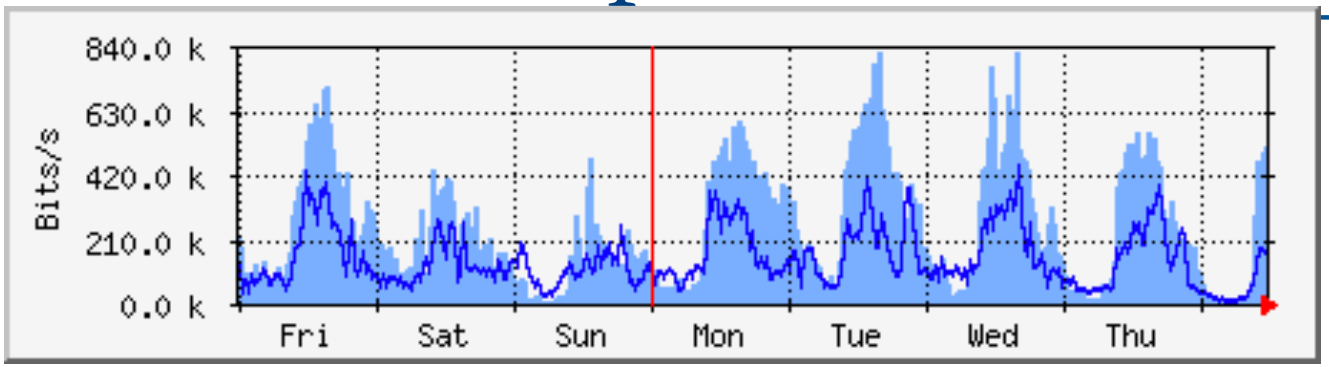

Router A to AS5400

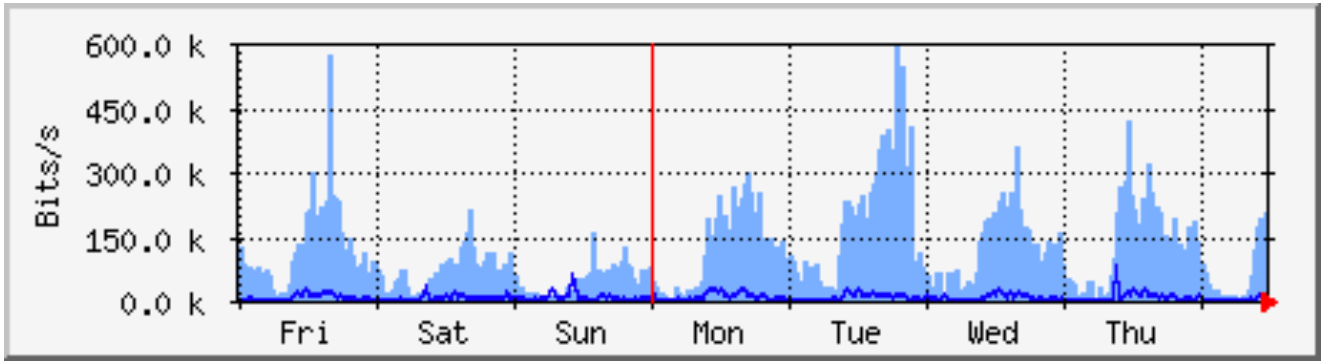

Router A to AS22351

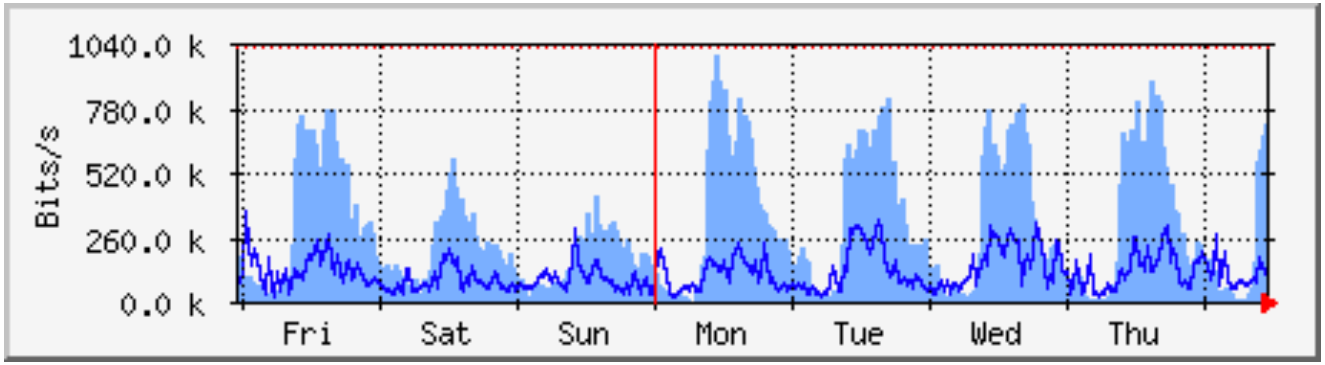

Router B to AS2516

Case Study
Configuration Router A

```
router bgp 17660
 no synchronization
 no bgp fast-external-fallover
 bgp log-neighbor-changes
 bgp deterministic-med
 neighbor 80.255.39.241 remote-as 22351
 neighbor 80.255.39.241 description ebgp peer to AS22351
 neighbor 80.255.39.241 send-community
 neighbor 80.255.39.241 prefix-list in-filter in
 neighbor 80.255.39.241 prefix-list out-filter-as22351 out
 neighbor 80.255.39.241 route-map as22351-out out
 neighbor 80.255.39.241 maximum-prefix 120000 95 warning-only
 neighbor 80.255.39.241 filter-list 3 in
 neighbor 80.255.39.241 filter-list 5 out
```

...next slide

Case Study
Configuration Router A

```
 neighbor 166.49.165.13 remote-as 5400
 neighbor 166.49.165.13 description eBGP multihop to AS5400
 neighbor 166.49.165.13 ebgp-multihop 5
 neighbor 166.49.165.13 update-source Loopback0
 neighbor 166.49.165.13 send-community
 neighbor 166.49.165.13 prefix-list in-filter in
 neighbor 166.49.165.13 prefix-list out-filter out
 neighbor 166.49.165.13 route-map as5400-out out
 neighbor 166.49.165.13 filter-list 1 in
 neighbor 166.49.165.13 filter-list 5 out
!
ip prefix-list in-filter  deny rfc1918 prefixes etc
ip prefix-list out-filter permit 202.144.128.0/19
ip prefix-list out-filter-as22351 permit 202.144.128.0/19
ip prefix-list out-filter-as22351 permit 202.144.158.0/23
```

...next slide

Case Study
Configuration Router A

```
ip as-path access-list 1 deny _701_
ip as-path access-list 1 deny _1_
ip as-path access-list 1 deny _7018_
ip as-path access-list 1 deny _1239_
ip as-path access-list 1 deny _7046_
ip as-path access-list 1 permit _5400$
ip as-path access-list 1 permit _5400_[0-9]+$
ip as-path access-list 1 permit _5400_[0-9]+_[0-9]+$
ip as-path access-list 1 deny .*
ip as-path access-list 3 permit _22351$
ip as-path access-list 3 permit _22351_[0-9]+$
ip as-path access-list 3 deny .*
ip as-path access-list 5 permit ^$
!
route-map as5400-out permit 10
 set community 5400:2001 5400:2101 5400:2119 5400:2124 5400:2128
route-map as22351-out permit 10
```

Case Study
Configuration Router B

```
router bgp 17660
 no synchronization
 no auto-summary
 no bgp fast-external-fallover
 bgp log-neighbor-changes
 bgp deterministic-med
 neighbor 210.132.92.165 remote-as 2516
 neighbor 210.132.92.165 descr eBGP Peering with AS2516
 neighbor 210.132.92.165 send-community
 neighbor 210.132.92.165 prefix-list default-route in
 neighbor 210.132.92.165 prefix-list out-filter out
 neighbor 210.132.92.165 route-map as2516-out out
 neighbor 210.132.92.165 maximum-prefix 100
 neighbor 210.132.92.165 filter-list 2 in
 neighbor 210.132.92.165 filter-list 5 out
```

...next slide

Case Study
Configuration Router B

```
!
prefix-list default-route permit 0.0.0.0/0
prefix-list out-filter permit 202.144.128.0/19
!
ip as-path access-list 2 permit _2516$
ip as-path access-list 2 deny .*
ip as-path access-list 5 permit ^$
!
route-map as2516-out permit 10
 set as-path prepend 17660
!
```

Interesting Aside

- Prior to installation of the new 640kbps link, ISP was complaining that both 1Mbps links were saturated inbound
  - Hence the requirement for the new 640kbps circuit
- Research using NetFlow, cflowd and FlowScan showed that Kazaa was to blame!
  - Kazaa is a peer to peer file sharing utility
  - Consumes all available bandwidth
  - Found that many customers were using Kazaa for file sharing, saturating the links inbound

Interesting Aside

- Solution
  - A time of day filter which blocked Kazaa during working hours, 8am to 8pm
  - Inbound and outbound ACLs on border routers had tcp/1214 filters added

```
access-list 100 deny tcp any any eq 1214 time-range OfficeHrs
access-list 101 deny tcp any any eq 1214 time-range OfficeHrs
!
time-range OfficeHrs
 periodic weekdays 8:00 to 20:00
```

  - The result: inbound traffic on external links dropped by 50%
  - And complaints about “the ‘net” being slow have reduced

Interesting Aside

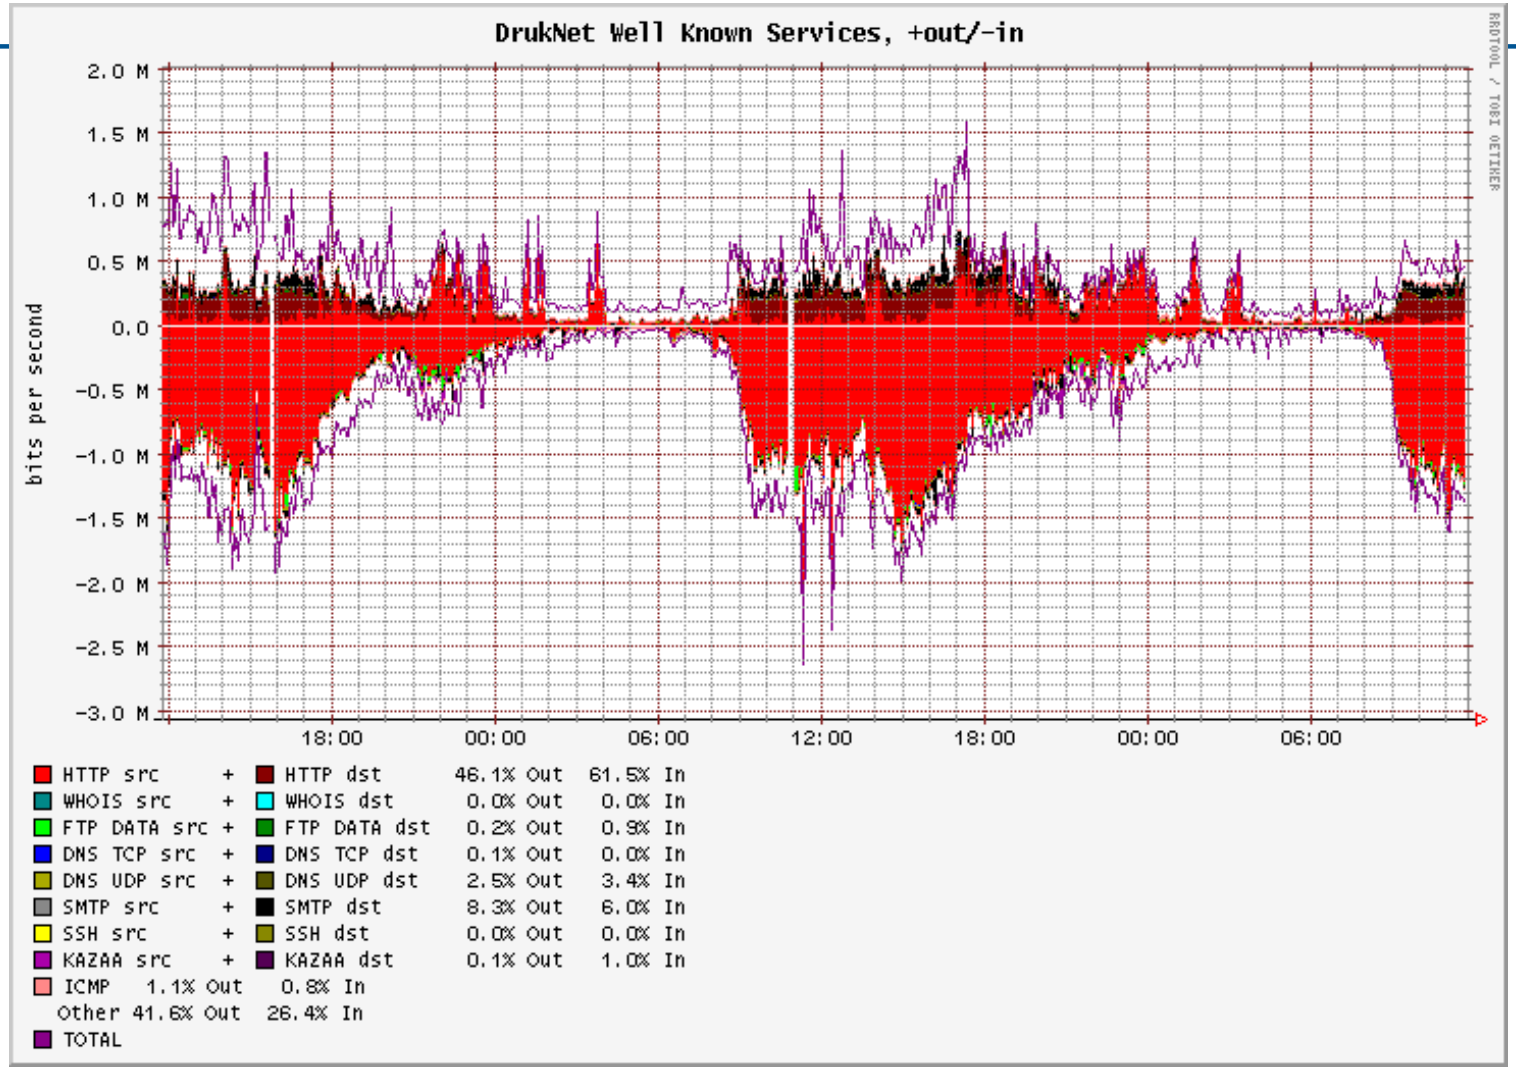

Typical FlowScan graph – no longer showing the effects of Kazaa

Summary

- Multihoming solution with three links of different bandwidths works well
  - Fluctuates significantly during the day time, maybe reflecting users browsing habits?
  - NOC is monitoring the situation
  - NOTE: Full routing table is not required ☺

Multihoming Case Study

ISP Workshops