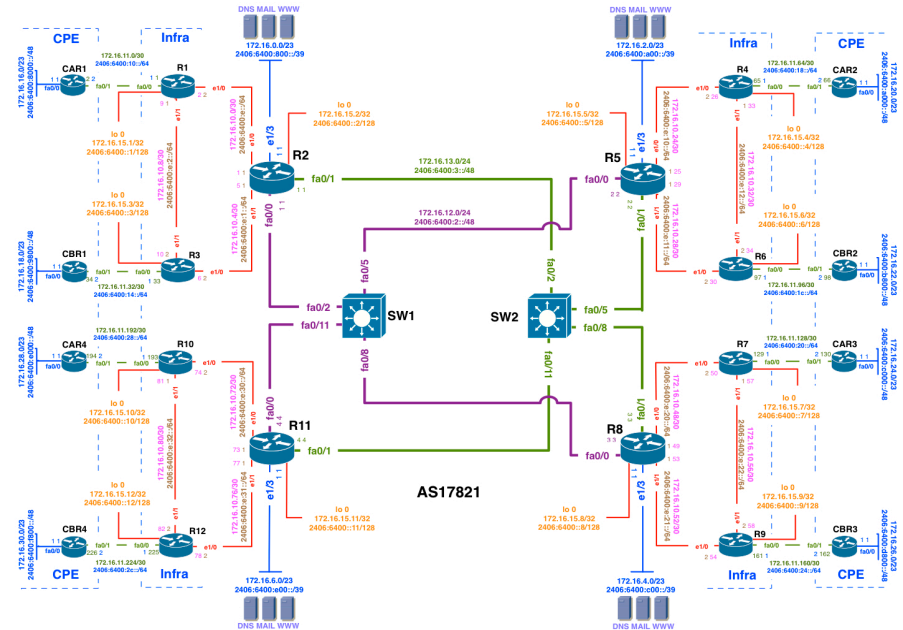


Route Filtering

Network Security Workshop


Route Filtering

- Types of prefixes in IP core network:
 - Internal Prefixes
 - External prefixes
 - Downstream customers
 - Internet prefixes



Route Filtering

- Internal prefixes originated in IP core network
 - Loopback
 - Transport
 - Connect inter-regional networks
 - Point-to-point
 - Infrastructure point-to-point
 - Customer side point-to-point
 - Data centre
 - Some ISP originate from separate AS if it is a large public hosting operation and multihomed DC



Route Filtering

- Loopback Prefix
 - Prefix size /128
 - Advertised in IGP i.e. OSPF
 - Scope within IP core network
 - Can be summarize in IGP i.e. OSPF if the number of loopback prefixes are large within the region

Route Filtering

Loopback prefixes in Training ISP network:

```
Router1#sh ipv6 route
IPv6 Routing Table - default - 51 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
        B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
        I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
        D - EIGRP, EX - EIGRP external, ND - Neighbor Discovery, I - LISP
        O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
        ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
```

```
LC 2406:6400::1/128 [0/0]
   via Loopback0, receive
OI 2406:6400::2/128 [110/10]
   via FE80::C801:1FF:FEAE:1C, Ethernet1/0
O  2406:6400::3/128 [110/10]
   via FE80::C802:1FF:FEAE:1D, Ethernet1/1
OI 2406:6400::4/128 [110/21]
   via FE80::C801:1FF:FEAE:1C, Ethernet1/0
OI 2406:6400::5/128 [110/11]
   via FE80::C801:1FF:FEAE:1C, Ethernet1/0
OI 2406:6400::6/128 [110/21]
   via FE80::C801:1FF:FEAE:1C, Ethernet1/0
OI 2406:6400::7/128 [110/21]
   via FE80::C801:1FF:FEAE:1C, Ethernet1/0
OI 2406:6400::8/128 [110/11]
   via FE80::C801:1FF:FEAE:1C, Ethernet1/0
OI 2406:6400::9/128 [110/21]
   via FE80::C801:1FF:FEAE:1C, Ethernet1/0
OI 2406:6400::10/128 [110/21]
   via FE80::C801:1FF:FEAE:1C, Ethernet1/0
OI 2406:6400::11/128 [110/11]
   via FE80::C801:1FF:FEAE:1C, Ethernet1/0
OI 2406:6400::12/128 [110/21]
   via FE80::C801:1FF:FEAE:1C, Ethernet1/0
```

Route Filtering

- Transport Prefix
 - Prefix size can be /64~/48
 - /48 is preferred if BGP traffic engineering required in future
 - Advertised in IGP i.e. OSPF
 - Scope within IP core network

Route Filtering

Transport prefixes in Training
ISP network:

```
Router1#sh ipv6 route
IPv6 Routing Table - default - 51 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
        B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
        I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
        D - EIGRP, EX - EIGRP external, ND - Neighbor Discovery, I - LISP
        O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
        ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
```

```
OI 2406:6400:2::/48 [110/11]
   via FE80::C801:1FF:FEAE:1C, Ethernet1/0
OI 2406:6400:3::/48 [110/11]
   via FE80::C801:1FF:FEAE:1C, Ethernet1/0
```

Route Filtering

- Prefixes advertised/originated in IP core network
 - Point-to-point
 - Infrastructure point-to-point
 - Prefix size /64 [/127 on interface configuration according to rfc-6164]
 - Advertised in IGP i.e. OSPF
 - Scope within IP core network
 - Can be summarize in IGP i.e. OSPF if the number of p-to-p prefixes are large within the region
 - Customer side point-to-point
 - Prefix size /64 [/127 on interface configuration according to rfc-6164]
 - Advertise from EGP i.e. iBGP (Not OSPF)
 - Scope within IP core network
 - Summarization in iBGP using network statement and pull up route [Atomic summarization]

Route Filtering

Infrastructure p-to-p prefixes
in Training ISP network:

```
Router1#sh ipv6 route
```

```
IPv6 Routing Table - default - 51 entries
```

```
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
```

```
B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
```

```
I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
```

```
D - EIGRP, EX - EIGRP external, ND - Neighbor Discovery, I - LISP
```

```
O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
```

```
ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
```

```
C 2406:6400:E::/64 [0/0]
  via Ethernet1/0, directly connected
O 2406:6400:E:1::/64 [110/20]
  via FE80::C802:1FF:FEAE:1D, Ethernet1/1
  via FE80::C801:1FF:FEAE:1C, Ethernet1/0
C 2406:6400:E:2::/64 [0/0]
  via Ethernet1/1, directly connected
OI 2406:6400:E:10::/64 [110/21]
  via FE80::C801:1FF:FEAE:1C, Ethernet1/0
OI 2406:6400:E:11::/64 [110/21]
  via FE80::C801:1FF:FEAE:1C, Ethernet1/0
OI 2406:6400:E:12::/64 [110/31]
  via FE80::C801:1FF:FEAE:1C, Ethernet1/0
OI 2406:6400:E:20::/64 [110/21]
  via FE80::C801:1FF:FEAE:1C, Ethernet1/0
OI 2406:6400:E:21::/64 [110/21]
  via FE80::C801:1FF:FEAE:1C, Ethernet1/0
OI 2406:6400:E:22::/64 [110/31]
  via FE80::C801:1FF:FEAE:1C, Ethernet1/0
OI 2406:6400:E:30::/64 [110/21]
  via FE80::C801:1FF:FEAE:1C, Ethernet1/0
OI 2406:6400:E:31::/64 [110/21]
  via FE80::C801:1FF:FEAE:1C, Ethernet1/0
OI 2406:6400:E:32::/64 [110/31]
  via FE80::C801:1FF:FEAE:1C, Ethernet1/0
```

Route Filtering

Customer p-to-p prefixes in
Training ISP network:

```
Router1#sh ipv6 route
```

```
IPv6 Routing Table - default - 51 entries
```

```
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
```

```
      B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
```

```
      I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
```

```
      D - EIGRP, EX - EIGRP external, ND - Neighbor Discovery, I - LISP
```

```
      O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
```

```
      ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
```

```
S 2406:6400:10::/48 [1/0]
  via Null0, directly connected
B 2406:6400:14::/48 [200/0]
  via 2406:6400::3
B 2406:6400:18::/48 [200/0]
  via 2406:6400::4
B 2406:6400:1C::/48 [200/0]
  via 2406:6400::6
B 2406:6400:20::/48 [200/0]
  via 2406:6400::7
B 2406:6400:24::/48 [200/0]
  via 2406:6400::9
B 2406:6400:28::/48 [200/0]
  via 2406:6400::10
B 2406:6400:2C::/48 [200/0]
  via 2406:6400::12
```

Route Filtering

- Data Centre Prefix
 - Prefix assignment can be /48 or a number of /48 if need more
 - /48 is preferred as it will support specific BGP network advertisement for traffic engineering purpose
 - Usually advertised in iBGP but ISP can prefer to advertise from separate AS using eBGP if DC is multihome, has separate routing policy the IP core and providing public hosting service
 - Scope within IP core network if single home
 - For multihoming case origin AS is different and ISP will allow transit only

Route Filtering

Data center prefixes in
Training ISP network:

IPv6 Routing Table - default - 51 entries

Codes: C - Connected, L - Local, S - Static, U - Per-user Static route

B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP

I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary

D - EIGRP, EX - EIGRP external, ND - Neighbor Discovery, I - LISP

O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2

ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2

```
B 2406:6400:800::/48 [200/0]
  via 2406:6400::2
B 2406:6400:A00::/48 [200/0]
  via 2406:6400::5
B 2406:6400:C00::/48 [200/0]
  via 2406:6400::8
B 2406:6400:E00::/48 [200/0]
  via 2406:6400::11
```

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. **Legitimacy of Address (LoA)** check
 - 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 2406:6400::/32
```

```
Inet6num:      2406:6400::/32
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 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 Upstream

- 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

Route Filtering Case study

- External Prefixes

- Customer prefix coming from downstream
- Internet prefix coming from the upstream



Route Filtering

Downstream customer
prefixes in Training ISP
network:

```
Router1#sh ipv6 route
```

```
IPv6 Routing Table - default - 51 entries
```

```
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
```

```
      B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
```

```
      I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
```

```
      D - EIGRP, EX - EIGRP external, ND - Neighbor Discovery, I - LISP
```

```
      O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
```

```
      ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
```

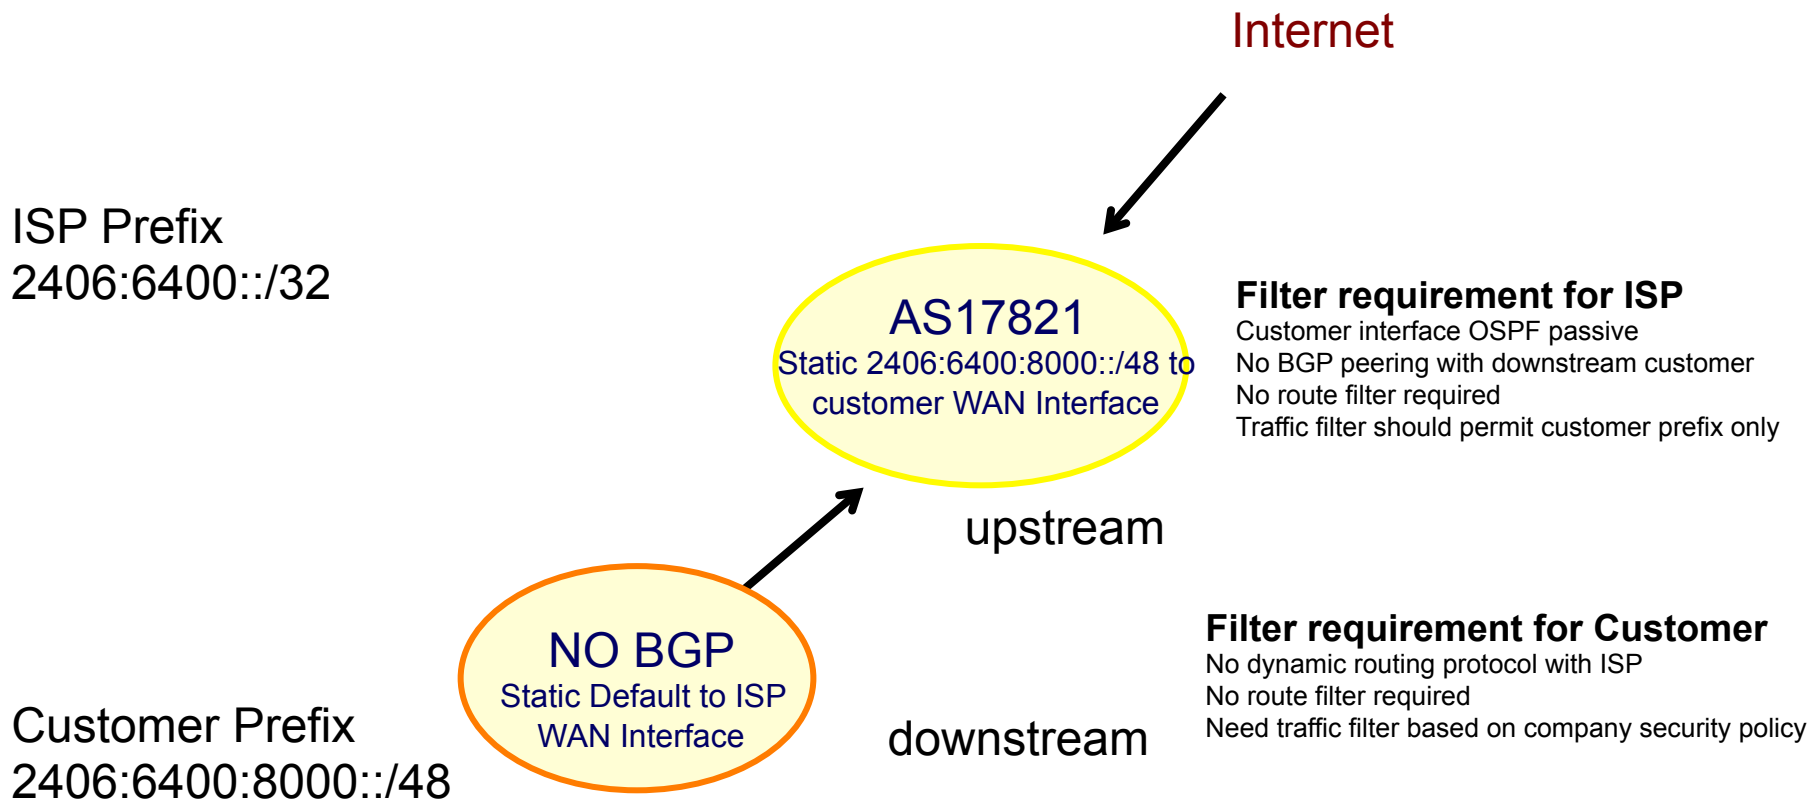
```
B 2406:6400:8000::/48 [20/0]
   via FE80::C80C:1FF:FEAF:6, FastEthernet0/0
B 2406:6400:9800::/48 [200/0]
   via 2406:6400::3
B 2406:6400:A000::/48 [200/0]
   via 2406:6400::4
B 2406:6400:B800::/48 [200/0]
   via 2406:6400::6
B 2406:6400:C000::/48 [200/0]
   via 2406:6400::7
B 2406:6400:D800::/48 [200/0]
   via 2406:6400::9
B 2406:6400:E000::/48 [200/0]
   via 2406:6400::10
B 2406:6400:F800::/48 [200/0]
   via 2406:6400::12
```

Route Filtering

- Customer prefix coming from downstream:
 - Option 1: Customer **single home** and **non portable prefix**
 - Customer is not APNIC member prefix received from upstream ISP
 - Option 2: Customer **single home** and **portable prefix**
 - Customer is APNIC member receive allocation as service provider but no AS number yet
 - Option 3: Customer **multihome** and **non portable prefix**
 - Customer is not APNIC member both prefix and ASN received from upstream ISP
 - Option 4: Customer **multihome** and **portable prefix**
 - Customer is APNIC member both prefix and ASN received from APNIC

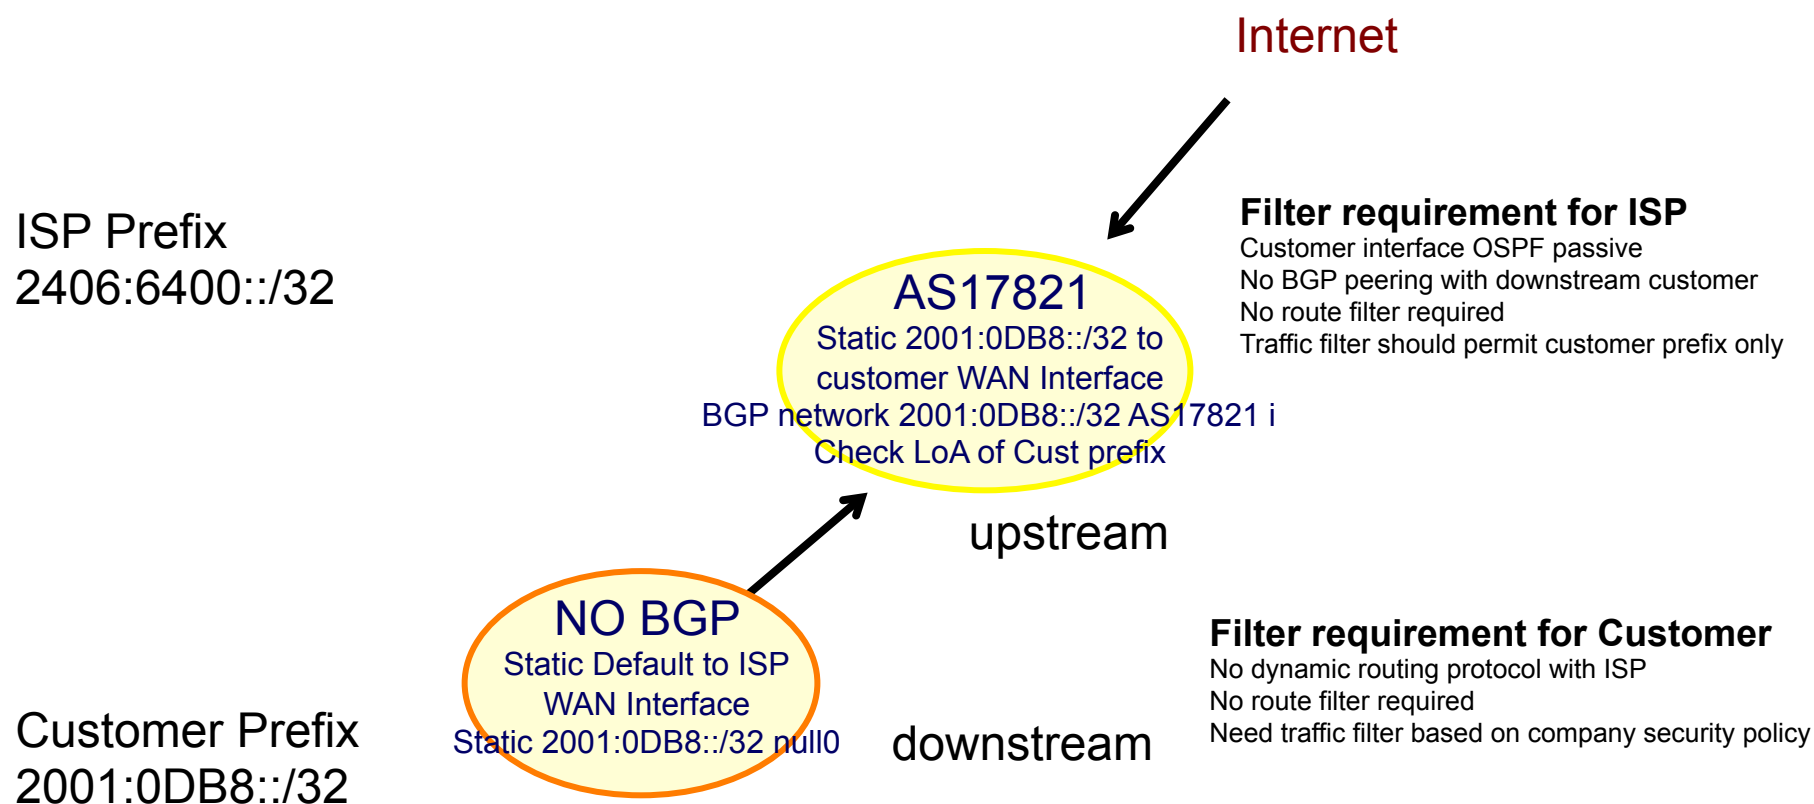
Route Filtering

- Option 1: Customer **single home** and **non portable prefix**



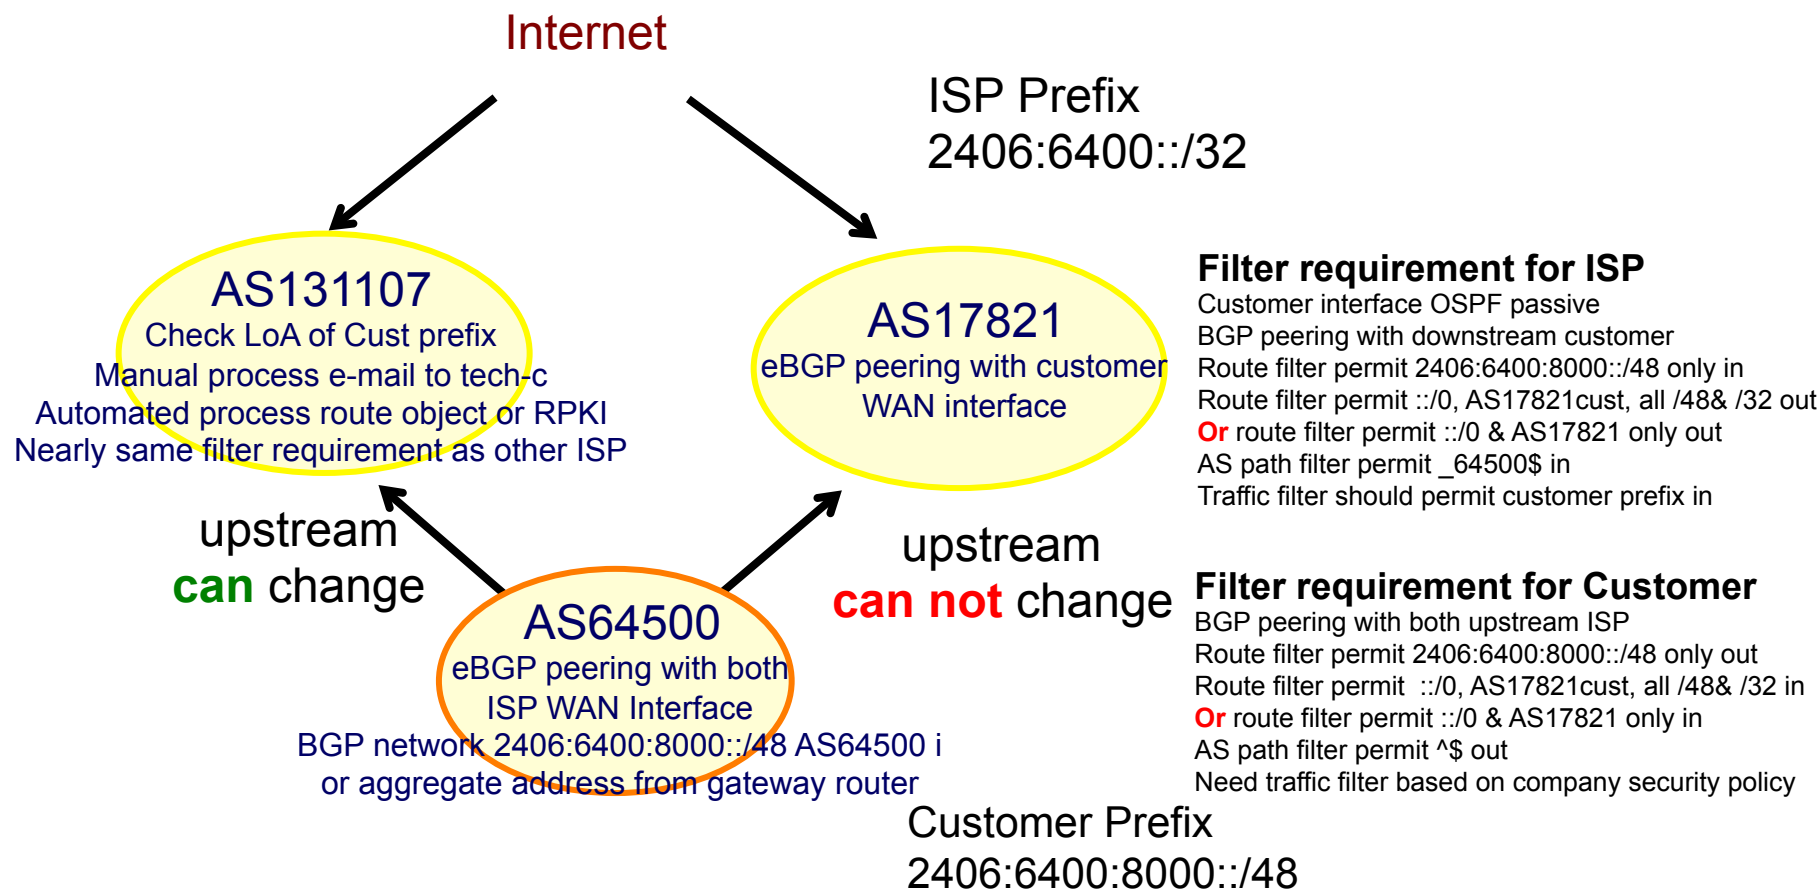
Route Filtering

- Option 2: : Customer **single home** and **portable prefix**



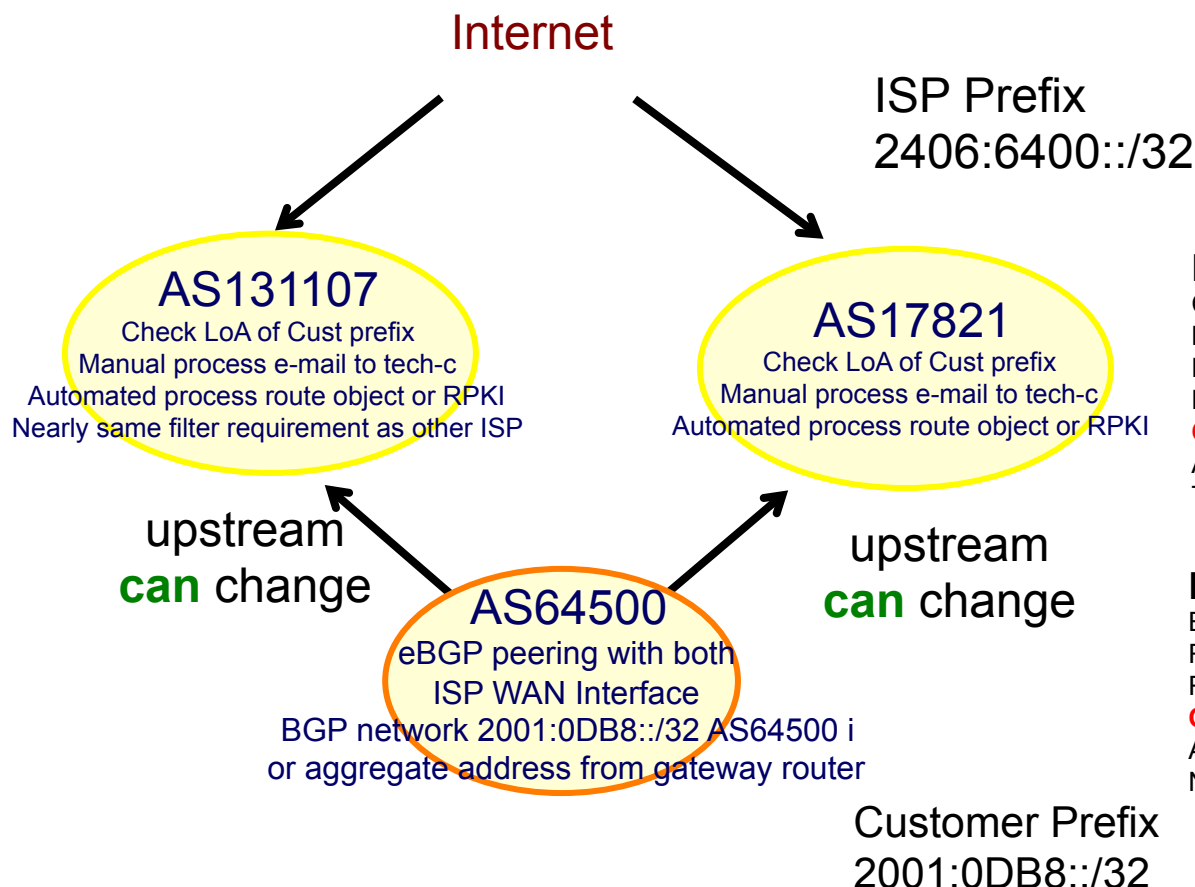
Route Filtering

- Option 3: Customer **multihome** and **non portable prefix**



Route Filtering

- Option 4: Customer **multihome** and **portable prefix**



Filter requirement for ISP

Customer interface OSPF passive
BGP peering with downstream customer
Route filter permit 2001:0DB8::/32 only in
Route filter permit ::/0, AS17821cust, all /48& /32 out
Or route filter permit ::/0 & AS17821 only out
AS path filter permit _64500\$ in
Traffic filter should permit customer prefix in

Filter requirement for Customer

BGP peering with both upstream ISP
Route filter permit 2001:0DB8::/32 only out
Route filter permit ::/0, AS17821cust, all /48& /32 in
Or route filter permit ::/0 & AS17821 only in
AS path filter permit ^\$ out
Need traffic filter based on company security policy

Route Filtering BCP

- **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 (?)

Route Filtering Plan in Training Lab

- We will use **option 3**: Config on ISP Edge router **(In)**
 - Receive individual customer prefix
 - i.e. On R1 From R13 2406:6400:8000::/48
 - On R3 From R14 2406:6400:9800::/48
 - On R4 From R15 2406:6400:a000::/48
 - On R6 From R16 2406:6400:b800::/48
 - On R7 From R17 2406:6400:c000::/48
 - On R9 From R18 2406:6400:d800::/48
 - On R10 From R19 2406:6400:e000::/48
 - On R11 From R20 2406:6400:f800::/48
 - And prefix originated by customer AS

Route Filtering Plan in Training Lab

- We will use **option 3**: Config on ISP Edge router (**Out**)
 - Send default prefix to customer i.e. `::/0`
 - Send aggregated ISP prefix i.e. `2406:6400::/32`
 - Send all individual customer prefix i.e.
 - `2406:6400:8000::/48`
 - `2406:6400:9800::/48`
 - `2406:6400:a000::/48`
 - `2406:6400:b800::/48`
 - `2406:6400:c000::/48`
 - `2406:6400:d800::/48`
 - `2406:6400:e000::/48`
 - `2406:6400:f800::/48`
 - Send all Internet prefix with prefix length $>/32$, $/32$ and $/48$ only

Route Filtering Plan in Training Lab

- We will use **option 3**: Config on CPE router (**IN**)
 - Receive default prefix to customer i.e. `::/0`
 - Receive aggregated ISP prefix i.e. `2406:6400::/32`
 - Receive all individual cust prefix i.e.
 - `2406:6400:8000::/48`
 - `2406:6400:9800::/48`
 - `2406:6400:a000::/48`
 - `2406:6400:b800::/48`
 - `2406:6400:c000::/48`
 - `2406:6400:d800::/48`
 - `2406:6400:e000::/48`
 - `2406:6400:f800::/48`
 - Receive all Internet prefix with prefix length $>/32$, $/32$ and $/48$ only

Route Filtering Plan in Training Lab

- We will use **option 3: Config on CPE router (Out)**
 - Send individual customer prefix only
 - i.e. From R13 To R1 2406:6400:8000::/48
 - From R14 To R3 2406:6400:9800::/48
 - From R15 To R4 2406:6400:a000::/48
 - From R16 To R6 2406:6400:b800::/48
 - From R17 To R7 2406:6400:c000::/48
 - From R18 To R9 2406:6400:d800::/48
 - From R19 To R10 2406:6400:e000::/48
 - From R20 To R12 2406:6400:f800::/48
 - Send that prefix originated customer AS number

Questions