

IPv6 Multihoming For Enterprises

What's the Problem?

*Jen Linkova
May 2016*

What's so Special About Enterprise Networks?

- ✓ Very conservative in network designs
- ✓ A lot of legacy equipment/hosts (need to be supported)
- ✓ Often need to enforce network policies

If multihoming in IPv6 is hard - they would not deploy IPv6...

Ideal World

End2End (Smart Host, Stupid Network)

- ✓ Host chooses src IP address(es)
- ✓ Network gets the packets to the destination

Problems:

- ✓ Network policies (egress link choice based on...\$\$\$\$ etc)
- ✓ Stupid hosts

Dumb Host, Smart Net: Option #1

Enterprise LIR (...or get a PI...)

Pros: problem solved

Cons:

- ✓ BGP between ISPs and the customer (3G? DLS?)
- ✓ Inflates global routing table (*one or more /48 for 5 users?*)
- ✓ More specifics from PA can be filtered
- ✓ \$\$\$\$\$\$\$\$ of becoming an LIR
- ✓ Operational overhead

Dumb Host, Smart Net: Option #2.1

PA sub assignments + BGP

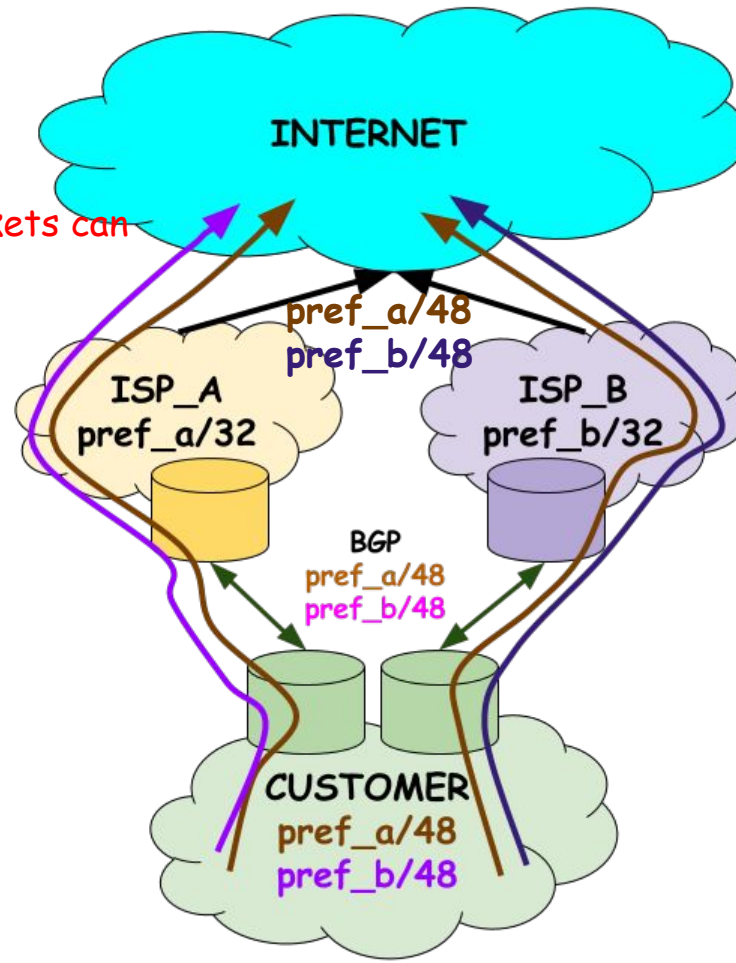
Pros: problem solved (+TE)

Cons:

- ✓ BGP again
- ✓ Inflates global routing table*
- ✓ More specific can be blocked
- ✓ BCP38 filtering on upper levels
- ✓ Operational overhead

(*) if more specifics are filtered - no redundancy

Routes/packets can
be filtered
upstream



Dumb Host, Smart Net: Option #2.2

No BGP, Just Accept Packets

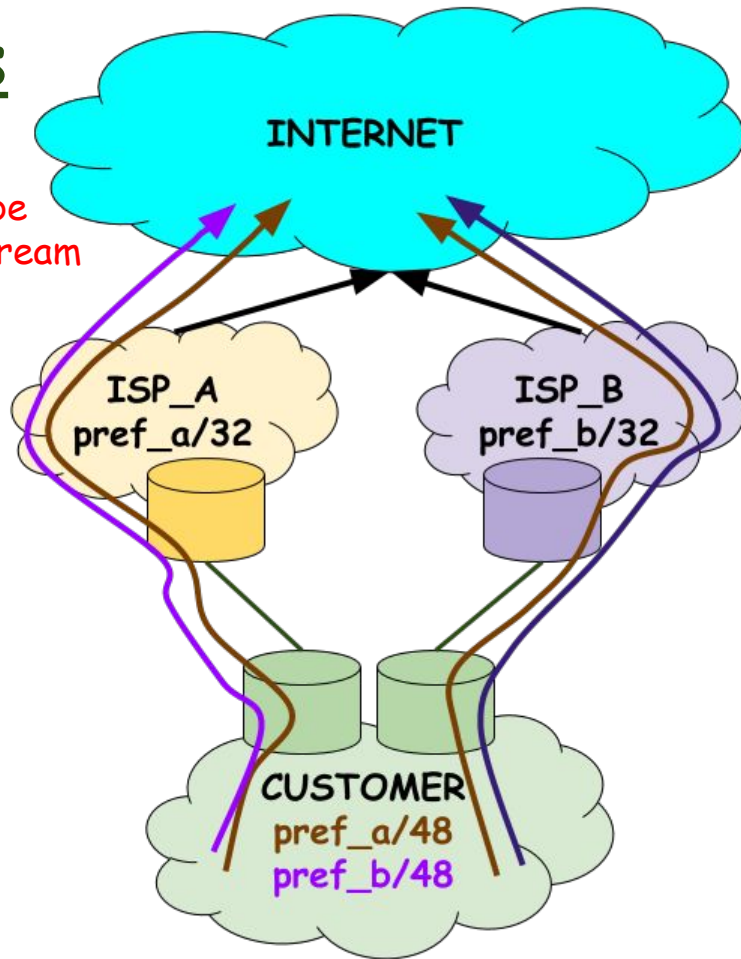
Pros:

- ✓ problem solved (if it only worked..)
- ✓ No routing table inflation

Cons:

- ✓ BCP38 filtering on upper levels
 - IRRs data will never be up-to-date
- ✓ ISPs need to cooperate
 - no going to happen

Packets can be
filtered upstream



To Make Multihoming Really Work

Choose the uplink based on src IP address

~~OR~~ AND

Choose the src IP based on the desired/available path

network



Hosts (with some input from the network)

Choosing the Uplink Based on Src IP Address

Currently can be done by dirty hacks (aka PBR)

Better solution required:

- ✓ Zero-touch configuration
- ✓ Ability to apply policies
 - Load-balancing vs primary/backup
 - traffic-engineering
- ✓ Fault tolerant

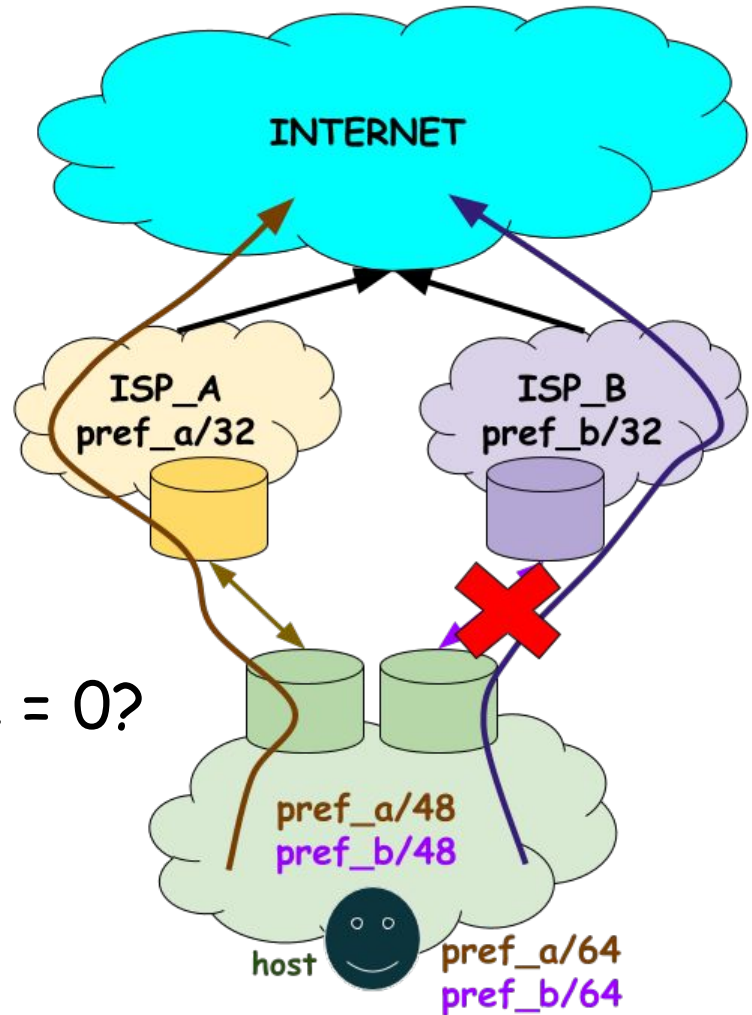
Network needs
to influence
src ip selection
on hosts

Src IP Selection on Hosts: Failure Scenario

How to signal to hosts

that **pref_b** should not be used?

- ✓ ICMP Error Message?
 - No way to signal when link is back up
- ✓ RA PIO with preferred lifetime = 0?
 - What if no more GUA left?
- ✓ Anything else?



QUESTIONS?