

IAB IPv6 Multi-Homing BOF

Jason Schiller
Senior Internet Network Engineer
IP Core Infrastructure Engineering
UUNET / MCI



Multi-homing Problems

- Inbound to the destination traffic engineering is needed
- Current multi-homing is site based not host based. Host based multi-homing does not lend itself to current operational processes
 - Large number of hosts
 - Complex routed network
 - End users do not own network or traffic engineering preferences
- Currently TE decisions are decided for and configured at the network level **not all end hosts!**
 - The Internet facing routers and end hosts may not be managed by the same group of operators
 - May want to manage the inter-AS traffic engineering policy in a few places (Internet facing routers) as opposed to every host
- Transit AS TE capabilities may be a requirement
- May make firewall filtering difficult

Site Multi-homing vs. Host Multi-homing

- Host multi-homing may be useful for consumer customers
 - Number of hosts at location is small
 - End user own host and network configuration
 - Routing equipment may have limited capabilities or is owned by service provider
- Site multi-homing is more useful for large commercial customers
 - Large number of hosts
 - Complex routed network
 - End users do not own network or traffic engineering preferences
- Do we need both? Is it worth the complexity?

Non-useful Transit AS Traffic Engineering

- Inter-AS traffic engineering in IPv4 is accomplished by sending more specific routes to the Internet, and allowing these more specific routes to be reachable across all connected ASes.
 - This allows each transit AS to make its own decision about what is the “best” path to take. Each transit AS can manipulate which is the best path by manipulating route announcements heard from its peers.
- In IPv6 transit ASes can only manipulate routing for an ISP aggregate affecting all customers using the ISP aggregate as the routing table lacks more specifics

Packet Filtering and Firewall Issues

- IP source and destination address will not change on non-shimmed packets
- IP source and destination may change on packets with a layer 3 shim
- Packet filters may need to match on IP source, IP destination, layer 3 source ULID, Layer 3 destination ULID, protocol, and port numbers
- Packet filters may require additional logic to map TCP established sessions when ULID is inserted and IP source or IP destination changes
- Stateful firewalls may need to match on IP source, IP destination, layer 3 source ULID, Layer 3 destination ULID, protocol, and port numbers
- Stateful firewalls will require additional logic to map sessions established with non-shimmed packets that migrate to shimmed packets with possible changing IP source and IP destination addresses.

Solution Space

1. Destination host sends extra information to the source host choosing the locators
2. Let the routers reach into the locator set exchange and add additional information or modify the locator set exchange in some way.
3. An 8+8 type solution, where the end hosts choose a locator, and allow the transit routers to rewrite network portion
4. Move the shim to be a router function

Destination Based TE

- Inbound TE to destination requires some communication from the destination to the source host who is choosing what locator to use
- Outbound TE from the source may be determined by host w/o IGP information by selecting source IP address due to upstream ISP filtering
- If destination host provides this information the all hosts will need to be configured with the network wide TE preferences
 - May be a good solution for consumer networks
 - Will be a bad solution for large commercial networks
- Configure TE policy at network level and push policy out to hosts
 - Requires TE server or routers to be configured with TE policy
 - Requires additional protocol and complexity
 - Additional protocols may create security issues
- Breaks transit AS TE capabilities

Routers Rewriting Shim6 Exchange

- Leave locator exchange on end hosts, but allow routers to insert TE preferences into the locator exchange
- Easily lends itself to network wide inbound TE policy
- Can leverage information about routing outages
- Will require routers to re-write shim6 locator exchanges to add TLVs (at least one per session)
 - Which routers to rewrite? Transit AS rewrite?
- Adds complexity to routers which may be difficult to support for consumer customers
- If transit ASes can reach into the locator set exchanges and further poison TLV metrics then locator ordering by the source can be influenced

Move Shim Insertion to Router

- Allow routers to insert shim on behalf of end hosts
 - May create additional security / authentication problems
- Allow routers to insert a additional “network” shim
 - May create additional security / authentication problems
 - Routers will need to recognize which packets should be shimmed
 - Not all hosts will be multi-homed (embedded devices with small IP stacks)
 - Will require routers to insert shim into all transit packets from multi-homed host at line rate
 - Will need to support line rate
 - Current routers support 600Mpps
 - Largest measured link 6Mpps throughput
- Which routers to do the shimming? Ingress? Transit? Egress?
Transit ASes?

Routers Rewriting IP Addresses (8+8)

- Leave the shim on the hosts and allow routers to rewrite network portion of the address
- Router will need to be able to map multiple networks to a given destination.
 - Requires IPv6NAT type solution -- routers keep locator set state
 - OR all hosts have to be multi-homed with same host address on all upstream ISPs
- Rewriting source may solve ISP filtering issues
- Transit ASes aware of a destination being reachable through alternate locators and can forward to alternates if they are better may allow transit AS TE

8+8 Problems

- Router will need to re-write packets at line rate
 - Current routers support 600Mpps
 - Largest measured link 6Mpps throughput
- Which router to do re-writing? Ingress? Egress? Transit?
 - Will all routers know about all source networks?
 - Transit routers may need to look past MPLS labels
- May break HBA/CGA
- Rewriting network address
 - Will break privacy addressing
 - Will break non-multi-homed hosts
- IPv6 NAT type solution
 - Requires all re-writing routers to keep source locator set state

And with that on with Ted From Sprint