IPv6 Multihoming

CITC Task force 13th Meeting

May 2013

Agenda

- IPv6 Multihoming Problem
- Proposed Solutions

IPv6 Multihoming Problem

Multi-Homing in IPv4

- **PI assignment**
  - ✓ Advertise your subnets to two providers
  - ✓ Granularity will depend on how many /24s and organization gets.
- **PA Assignment**
  - ✓ Same thing will be done as PI as long as organization gets /24 or more for the service provider
- **Saudi Special Requirements**
  - ✓ Filtering requirement is adding extra complexity.
  - ✓ Enterprise will have to maintain symmetrical traffic to avoid traffic drop.
  - ✓ Usually done by accepting default route from both subnet both providers and preferring one for upstream traffic.
  - ✓ For downstream traffic, the enterprise will advertise its subnets from one link and the same subnets from the other link with prepending (*There are other variations, keeping symmetry*)

Multi-Homing in IPv6 – The Problem

- The Methods discussed in IPv4 can't be applied to IPv6 as they are architecturally un-scalable.

- To Control routing table explosion
  - Current practice is filter on /48 prefix in the router facing customer.
  - Current practice is to filter on /32 on the internet transit provider.

- In **PA** assignment
  - A customer will be delegated two prefixes, one from ISP-A and one from ISP-B
  - Each Machine will have two IPs.
  - **ISP B** will not accept packets sourced from ISP-A and vise versa

Multi-Homing in IPv6 – The Problem (Cont.)

- In **PI** assignment
  - Current Policy from RIP is to assign /48
  - ISP will filter on /48
    - ✓ Upstream transit providers filter on /32
    - ✓ PI space will not be propagated to the internet
  - Can’t do traffic engineering and load sharing

RIPE Multi-homing Policy

RIPE Changed their policy on 23rd Jun, 2012

ORIGINAL TEXT	NEW TEXT
7. IPv6 Provider Independent (PI) Assignments To qualify for IPv6 PI address space, an organisation must: a) demonstrate that it will be multihomed b) meet the requirements of the policies described in the RIPE NCC document entitled " Contractual Requirements for Provider Independent Resources Holders in the RIPE NCC Service Region " The RIPE NCC will assign the prefix directly to the End User organisations upon a request properly submitted to the RIPE NCC, either directly or through a sponsoring LIR. The minimum size of the assignment is a /48. Organisations requesting a larger assignment (shorter prefix) must provide documentation justifying	*[Removed point a) in the first paragraph]* **7. IPv6 Provider Independent (PI) Assignments** To qualify for IPv6 PI address space, an organisation must meet the requirements of the policies described in the RIPE NCC document entitled " Contractual Requirements for Provider Independent Resources Holders in the RIPE NCC Service Region " The RIPE NCC will assign the prefix directly to the End User organisations upon a request properly submitted to the RIPE NCC, either directly or through a sponsoring LIR. The minimum size of the assignment is a /48. Organisations requesting a larger assignment (shorter prefix) must provide documentation justifying the need for additional subnets. Additional assignments may also be made when the

Reference No. (9 pt Arial)

Proposed Solutions

CIDR Boundary of /48

- This is proposing to use /48 for each site
  - ✓ Current RIPE Policy is to assign /48 for PI
  - ✓ One /32 has has **65,536** /48 subnets.
- Is an IPv6 multi-homing solution that does not support fine-grained traffic engineering that is currently available good enough for now?

Pro's	Con's
Every end-site can use their current Pv4 multi-homing knowledge to implement IPv6.	This solution does not allow for fine grained inter-AS traffic engineering as it does not allow a site to advertise more than one route
This solution provides end-sites the ability to multi-home and have fail-over capabilities.	This solution could cause some bloating in the routing tables, but will be limited to the number of AS's (currently 23,093 potentially 4,294,967,296).
Fast to implementation.	This solution could cost money in the future for a different solution to coexist for routing table bloat.
This is not an economically demanding solution in the near future.	

Reference No. (9 pt Arial)

CIDR Boundary of less than /48

- It will allow for more fine graining traffic engineering.
- Many have suggested that /49 would be good choice as it will allow for load sharing between two links.
- Assign a shorter prefix to static customer to avoid accidental redistribution of statically connected customers.

Pro's	Con's
This solution provides the ability for both multi-homing and route engineering.	This solution does not allow for fine grained inter-AS traffic engineering as it does not allow a site to advertise more than one route
	This solution could cause some bloating in the routing tables, but will be limited to the number of AS's (currently 23,093 potentially 4,294,967,296).
	This solution could cost money in the future for a different solution to coexist for routing table bloat.

Aggregation Slices

- On average in IPv4 there is an 8.4:1 route to AS ratio.
  - ✓ This means that we need 8 subnets to be advertised from one AS for proper traffic engineering
- Based on this it is required to increase the filter to /51.
- Keep the static customer in mind and give them smaller prefix.
- RIR policies will have to be rewritten to support this proposal.
- RIR need to be more strict in assigning ASN.

Pro's	Con's
This solution provides the ability for both multi-homing and route engineering.	This solution could cause bloating in the routing tables.
	This solution could cause some bloating in the routing tables, but will be limited to the number of AS's (currently 23,093 potentially 4,294,967,296).

Community Codes

- A BGP community attribute for tagging prefixes would be used
- The name of community would be MULTIHOME
- IANA should give a special code for this.
- The whole internet should accept any route with this community.
- This will help Multi-homing for PA space.

Pro's	Con's
This solution provides the ability for both multi-homing and route engineering for PA and PI space.	This solution could cause bloating in the routing tables.
	If someone fails to propagate the MULTIHOME community, then routes will be either not received or dropped.
	This solution is very prone to human error.
	This could be used as an easy way to hijack IP space.

Published List

- A list of approved prefixes for multi-homing would be published and filters would be opened to this approved list.
  - ✓ Similar to the AS Macro concept.
- This can be accomplished with the RIR's producing a list of approved multihoming blocks and then publishing it on their public Data Base as a loadable list.
- Opening of this list can only be a suggestion NOT mandated by RIR.
- Include PA and PI

Pro's	Con's
This solution provides the ability for both multi-homing and route engineering for PA and PI space.	This would require more work from the RIR
This solution would help control routing table bloat.	Frequency of updates may become an issue
This solution may solve several other issues such as hijacking, reduction of spam, reduction of viruses, and create some basic security.	This list of approved blocks might become excessively long
	The Internet could encounter a catastrophic failure if someone were to hack the DB.

Maximum Prefix

- Each origin AS can be limited to a certain number of prefixes.
- The result would be an Internet routing table that contains a maximum of X times larger the number of AS's.
- Drop the most specific routes of a given origin if the origin exceeds the configured limit.
- If the irresponsible ISP advertises more than 5 routes for a given origin AS, then the upstream Peer would discard all but five routes for the given origin AS.
- How Many prefixes should be allowed?

Pro's	Con's
This solution provides the ability for both multi-homing and route engineering.	This solution could cause bloating in the routing tables.

Selection of the Default route based on the Source Address (SDSA)

- ❑ Solve the issue of source address selection if the end use have two delegated PA prefixes.
- ❑ Selection will based on SDSA protocol.
- ✓ Each default route will is associated with one of the delegated prefixes

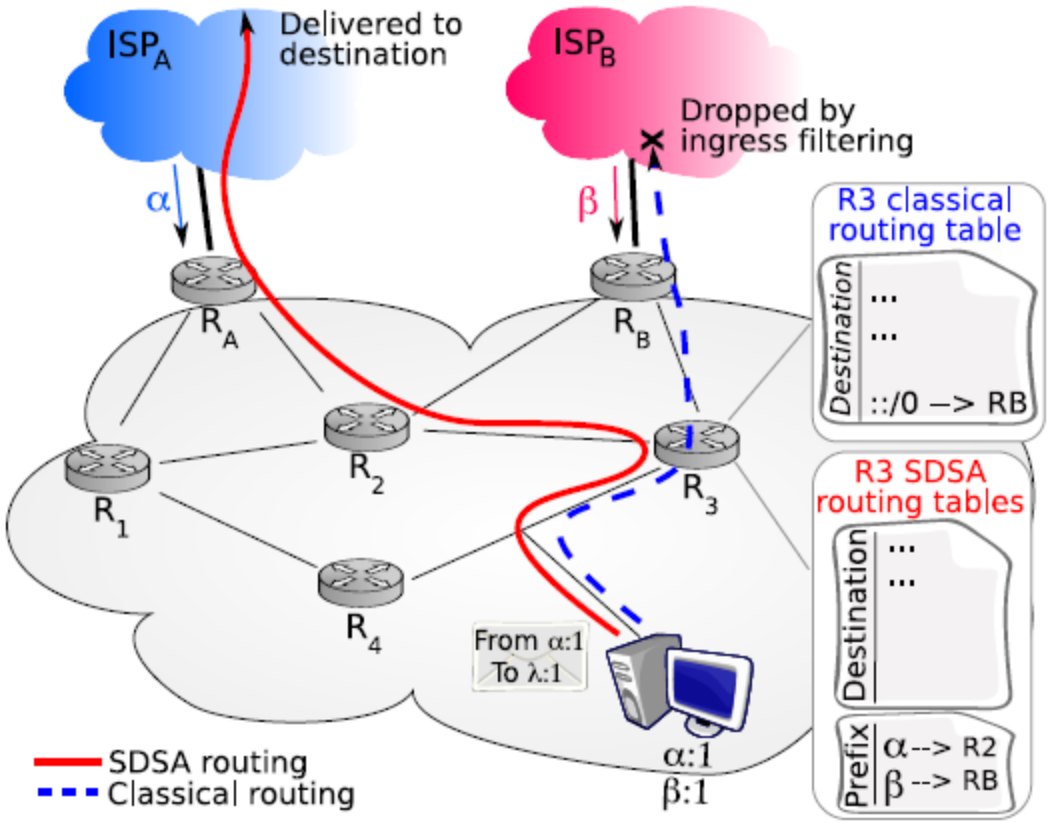

Source: Solving the Ingress Filtering Issue in an IPv6 multihued Home Network by Etienne Gallet de Santerre, Samih Jammoul and Laurent Toutain

Sources

- http://www.nro.net/wpcontent/uploads/2006/10/MultihomeIPv6procon.pdf
- http://www.ripe.net/ripe/docs/ripe-582
- http://tools.ietf.org/html/draft-ietf-v6ops-ipv6-multihoming-without-ipv6nat-03
- http://www.cs.auckland.ac.nz/~brian/multi6survey.pdf
- Solving the Ingress Filtering Issue in an IPv6 multihomed Network, by Etienne Gallet de Santerre, Samih Jammoul and Laurent Toutain

Q & A

Thank You