

tetaneutral.net - Evolution #125

Problemes de filtrage IPv6 sur annonce /40 PA

07/12/2011 09:56 - Laurent GUERBY

Statut:	Résolu	Début:	07/12/2011
Priorité:	Normal	Echéance:	
Assigné à:		% réalisé:	100%
Catégorie:		Temps estimé:	0.00 heure
Version cible:			
Description			
Certains sites comme afnic.fr ne sont pas accessibles en IPv6			

Historique

#1 - 07/12/2011 09:57 - Laurent GUERBY

```
bird> show route for 2001:0660:3003::/48
2001:660::/32 via 2001:7f8:43::6939:1 on eth0.127 [HE_EQUINIX_IX_6 2011-11-30 12:32:29] * (310) [AS2200i]
via 2001:7f8:54::19 on eth0.126 [FRANCE_IX_6_RS2 2011-12-05 13:07:57 from 2001:7f8:54::251] (150) [AS2200i]
via 2001:7f8:54::10 on eth0.126 [HE_FRANCE_IX_6 2011-11-30 12:32:29] (300) [AS2200i]
via 2001:7f8:54::19 on eth0.126 [FRANCE_IX_6_RS1 2011-11-28 14:51:38 from 2001:7f8:54::250] (100) [AS2200i]
```

J'ai essayé les différentes routes manuellement en ping6
2001:660:3003:2::4:20 mais ca coince a probablement l'entrée de Renater :

```
via 2001:7f8:43::6939:1
guerby@pc2:~$ mtr rwe 5 2001:660:3003:2::4:20
HOST: pe2
Loss% Snt Last Avg Best Wrst StDev
1|- myrys-h3-gw-ipv6.tetaneutral.net 0.0% 5 3.2 4.0 2.1 8.1 2.4
2.|- th2-gw-ipv6.tetaneutral.net 0.0% 5 21.7 20.7 19.1 23.1 1.6
3.|- 10gigabitethernet-2-2.par2.he.net 0.0% 5 20.5 20.7 19.9 21.2 0.5
4.|- 10gigabitethernet3-1.core1.par1.he.net 0.0% 5 20.6 23.1 20.6 25.3 1.9
5.|- ??? 100.0 5 0.0 0.0 0.0 0.0 0.0
```

```
via 2001:7f8:54::19
guerby@pc2:~$ mtr rwe 5 2001:660:3003:2::4:20
HOST: pe2
Loss% Snt Last Avg Best Wrst StDev
1|- myrys-h3-gw-ipv6.tetaneutral.net 0.0% 5 3.4 3.7 1.7 7.2 2.1
2.|- th2-gw-ipv6.tetaneutral.net 0.0% 5 21.6 21.4 20.3 22.2 0.7
3.|- ??? 100.0 5 0.0 0.0 0.0 0.0 0.0
```

```
via 2001:7f8:54::10
guerby@pc2:~$ mtr rwe 5 2001:660:3003:2::4:20
HOST: pe2
Loss% Snt Last Avg Best Wrst StDev
1|- myrys-h3-gw-ipv6.tetaneutral.net 0.0% 5 1.6 2.5 1.6 3.6 0.9
2.|- th2-gw-ipv6.tetaneutral.net 0.0% 5 19.6 20.2 19.6 21.2 0.6
3.|- hurricane-electric.franceix.net 0.0% 5 21.0 21.0 20.0 22.2 0.9
4.|- ??? 100.0 5 0.0 0.0 0.0 0.0 0.0
```

Notre routeur Toulousain myrys-h3-gw-ipv6.tetaneutral.net a une session avec Jaguar que nous utilisons en secours (AS prepending et faible local pref) :

```
bird> show route for 2001:0660:3003::/48
2001:660::/32 via 2a01:6600:8000::175 on eth0.3100 [TETANEUTRAL_6 Dec03] * (100) [AS2200i]
via 2a01:240:600::1 on eth0.4001 [JAGUAR_6 Dec06] (50) [AS2200i]
```

Si je force la route via Jaguar tout fonctionne :

```
root@h3:~# ip r a 2001:0660:3003::/48 via 2a01:240:600::1
guerby@pc2:~$ mtr rwe 20 2001:660:3003:2::4:20
HOST: pe2
Loss% Snt Last Avg Best Wrst StDev
1|- myrys-h3-gw-ipv6.tetaneutral.net 0.0% 20 2.3 2.7 1.4 3.9 0.8
2.|- 2a01:240:600::1 0.0% 20 2.2 5.7 2.1 53.0 11.2
3.|- ? 100.0 20 0.0 0.0 0.0 0.0 0.0
4.|- ? 100.0 20 0.0 0.0 0.0 0.0 0.0
5.|- ??? 100.0 20 0.0 0.0 0.0 0.0 0.0
6.|- afnic-primaire-vl387-te2-6-paris1-rtr-021.noc.renater.fr 0.0% 20 22.4 23.7 22.2 25.4 0.9
7.|- isg.sqy.interco.ipv6.nic.fr 0.0% 20 22.6 24.2 22.6 26.8 1.3
```

```
8.]-- rigolo.nic.fr 0.0% 20 24.2 23.6 22.6 26.3 0.9
guerby@pc2:~$ ping6 -c 3 2001:660:3003:2::4:20
PING 2001:660:3003:2::4:20(2001:660:3003:2::4:20) 56 data bytes
64 bytes from 2001:660:3003:2::4:20: icmp_seq=1 ttl=59 time=29.5 ms
64 bytes from 2001:660:3003:2::4:20: icmp_seq=2 ttl=59 time=24.0 ms
64 bytes from 2001:660:3003:2::4:20: icmp_seq=3 ttl=59 time=23.1 ms
```

```
--- 2001:660:3003:2::4:20 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2002ms
rtt min/avg/max/mdev = 23.131/25.577/29.598/2.871 ms
```

En cherchant les recommandations

<http://www.space.net/~gert/RIPE/ipv6-filters.html>

Renater semble appliquer la version "stricte" et non la "relaxed" qui permet jusqu'à /48 sur les PA avec le commentaire suivant "It accepts more specific routes up to a /48 from the RIR (2000::/3) ranges. I know that some people very harshly oppose announcing more specifics, but other people argue that it's better to do multihoming by announcing a /48 with a limited distribution (same region, for example) than to announce your own /32 globally."

#2 - 07/12/2011 10:00 - Laurent GUERBY

Proposition sur les PA /48 pour multihoming :

http://meetings.ripe.net/ripe-53/presentations/multihoming_paspape.pdf

#3 - 07/12/2011 10:05 - Laurent GUERBY

A noter que la recommandation RIPE actuellement en vigueur depuis le 7 novembre 2011 semble favoriser un routage des annonces PA jusqu'à /48 :

<http://www.ripe.net/ripe/docs/current-ripe-documents/ripe-532>

<<
RIPE Routing Working Group Recommendations on IPv6 Route Aggregation
(...)
Recommandation

It is suggested that prefix filters allow for prudent subdivision of an IPv6 allocation. The operator community will ultimately decide what degree of subdivision is supportable, but the majority of ISPs accept prefixes up to a length of /48 within PA space.
(...)

#4 - 07/12/2011 17:02 - Laurent GUERBY

<http://www.ripe.net/ripe/docs/ripe-538>

IPv6 Address Allocation and Assignment Policy
07 Dec 2011

#5 - 25/09/2015 21:29 - Matthieu Herrb

- Statut changé de Nouveau à Résolu

- % réalisé changé de 0 à 100

Résolu par annonce du /32 par FS non ? (et bientôt on aura notre /32).
Entre temps contacts avec Renater qui filtrait aussi trop strictement...