

IPv6. Now!

DE-CIX Tekkie Meeting #9 Nov 13, 2008

Fredy Künzler, AS13030
kuenzler at init7 dot nospam dot net

Init Seven AG
Elias-Canetti-Strasse 7
CH-8050 Zürich

www.init7.net
www.blogg.ch
www.bgp-and-beyond.com

Disclaimer: the configuration examples in this presentation are probably inaccurate. Use it on your own risk. If you find 5 errors in the presentation, let me know and you qualify for a free Init7 geek T-Shirt (please note your size) ...



Face facts:

Less than ~~1000~~ **800** days until the free IPv4 pool gets exhausted.

- Prediction by Geoff Huston, APNIC

<http://www.potaroo.net/tools/ipv4/>

- Quoting Mike Leber, Hurricane Electric (he.net):

Do you have an IPv6 plan?

How long do you think it will be until Sarbanes Oxley and SAS 70 auditors start requiring disclosure of IPv4 exhaustion as a business continuity risk, as well as the presence or lack thereof of an IPv6 plan?



The good news #1:

- If you are a LIR, request your IPv6 space now. It's just a formality:

<https://lirportal.ripe.net/>

-> Request Forms

-> IPv6 First Allocation Request Form



After receiving the IPv6 allocation, put the appropriate route object into the RIPE database.

The good news #2:

You don't have to reinvent the wheel. Everything you know about

- BGP4
- Subnet Masks
- OSPF
- Route-Maps, Prefix-lists, AS-Path, Filter-Lists is still valid. More or less.
- DualStack (IPv4/IPv6 on the same equipment) works well.
- Do not think of deploying IPv6-Tunnels (unless you are a DSL/Cable user ... or a TIER-1 ;-)

Initial IPv6 Deployment in a typical small provider network is less than 1 day of work.



IPv6 Address Plan #1:

- RIPE's allocation to the LIR is a /32
- Customer Assignments are usually a /48
(don't bother anymore how many addresses the customer should get)
- IPv4 /30 (PtP) becomes IPv6 /126
- Host Address (IPv4 /32) becomes IPv6 /128

Deploying an address plan:

- Stick to your existing IPv4 address scheme, for convenience.
- Use a /48 for loopbacks
- Use another /48 for Point-to-Point links

IPv6 Address Plan #2:

Remember the notation of IPv6:

`2001:db8::/32 =`

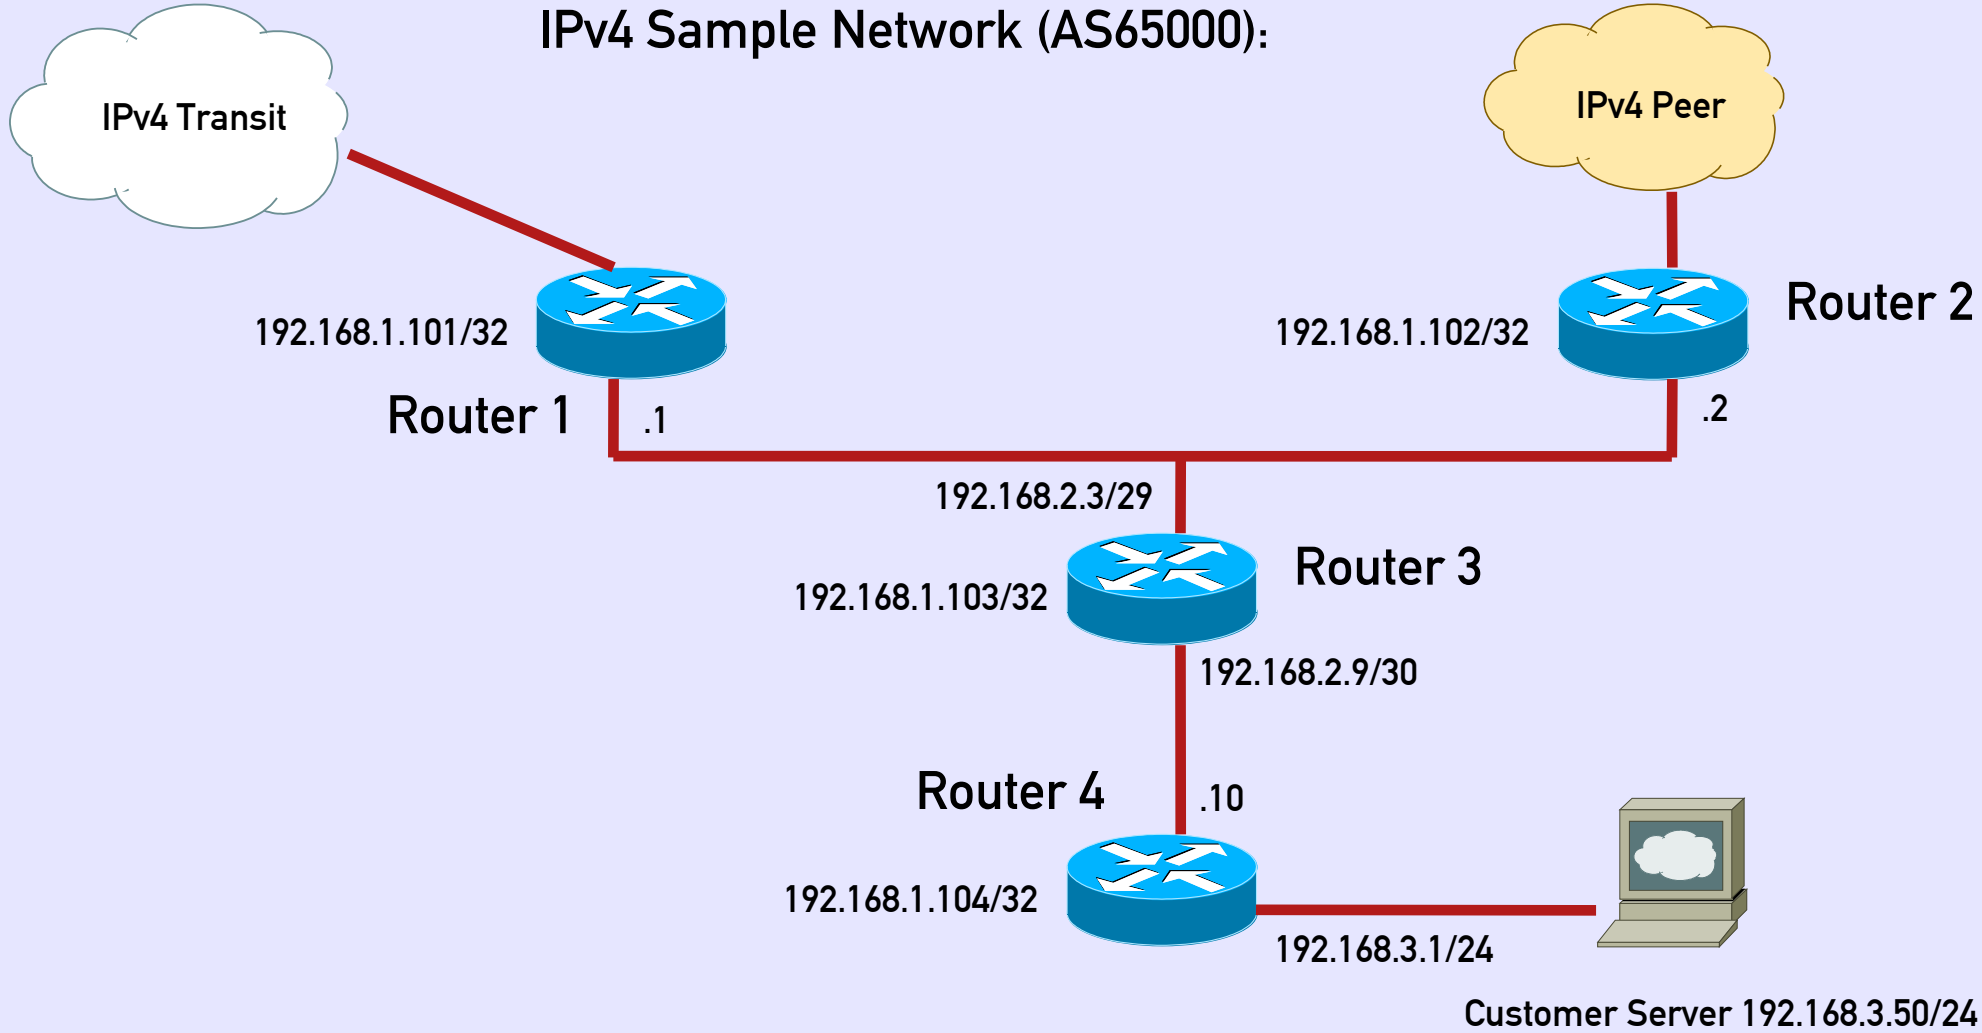
`2001:0db8:0000:0000:0000:0000:0000:0000/32`

- 0 are omitted
- two colons are replacing many 0

Note: `2001:db8::/32` should be used for any sample documentation.



IPv4 Sample Network (AS65000):

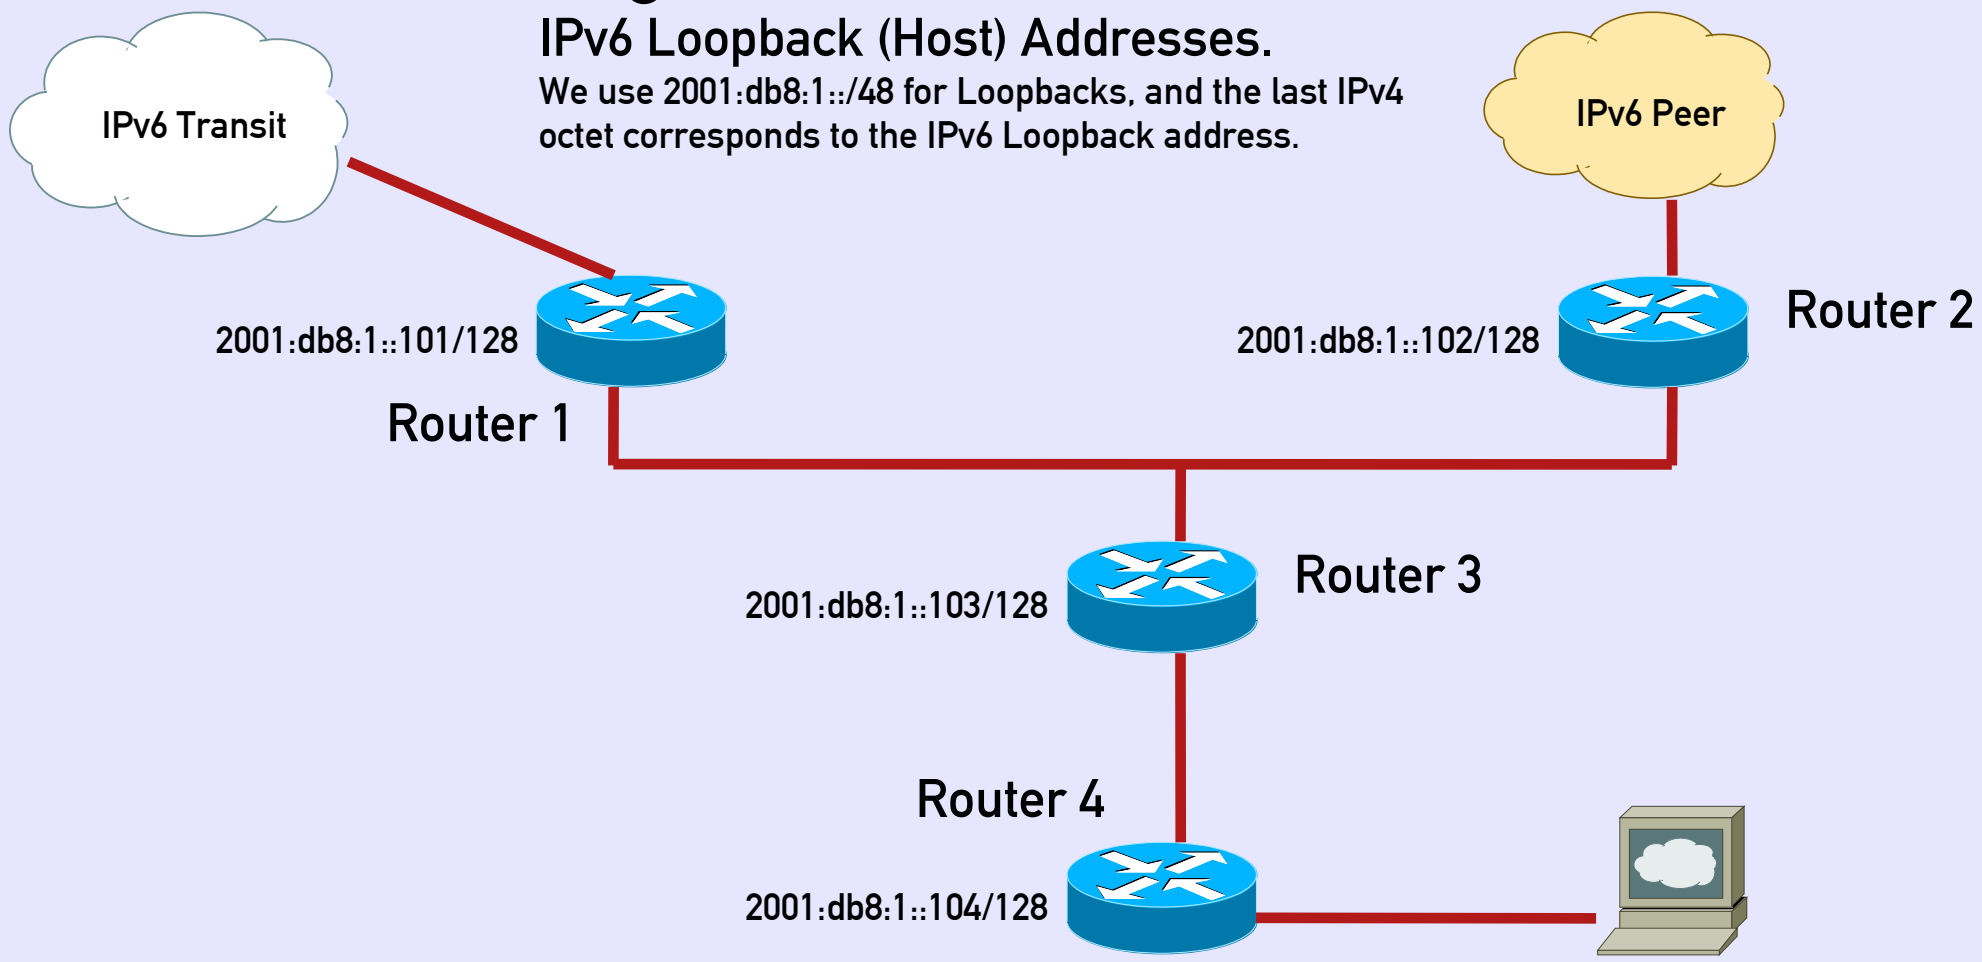




Migration to IPv6 #1

IPv6 Loopback (Host) Addresses.

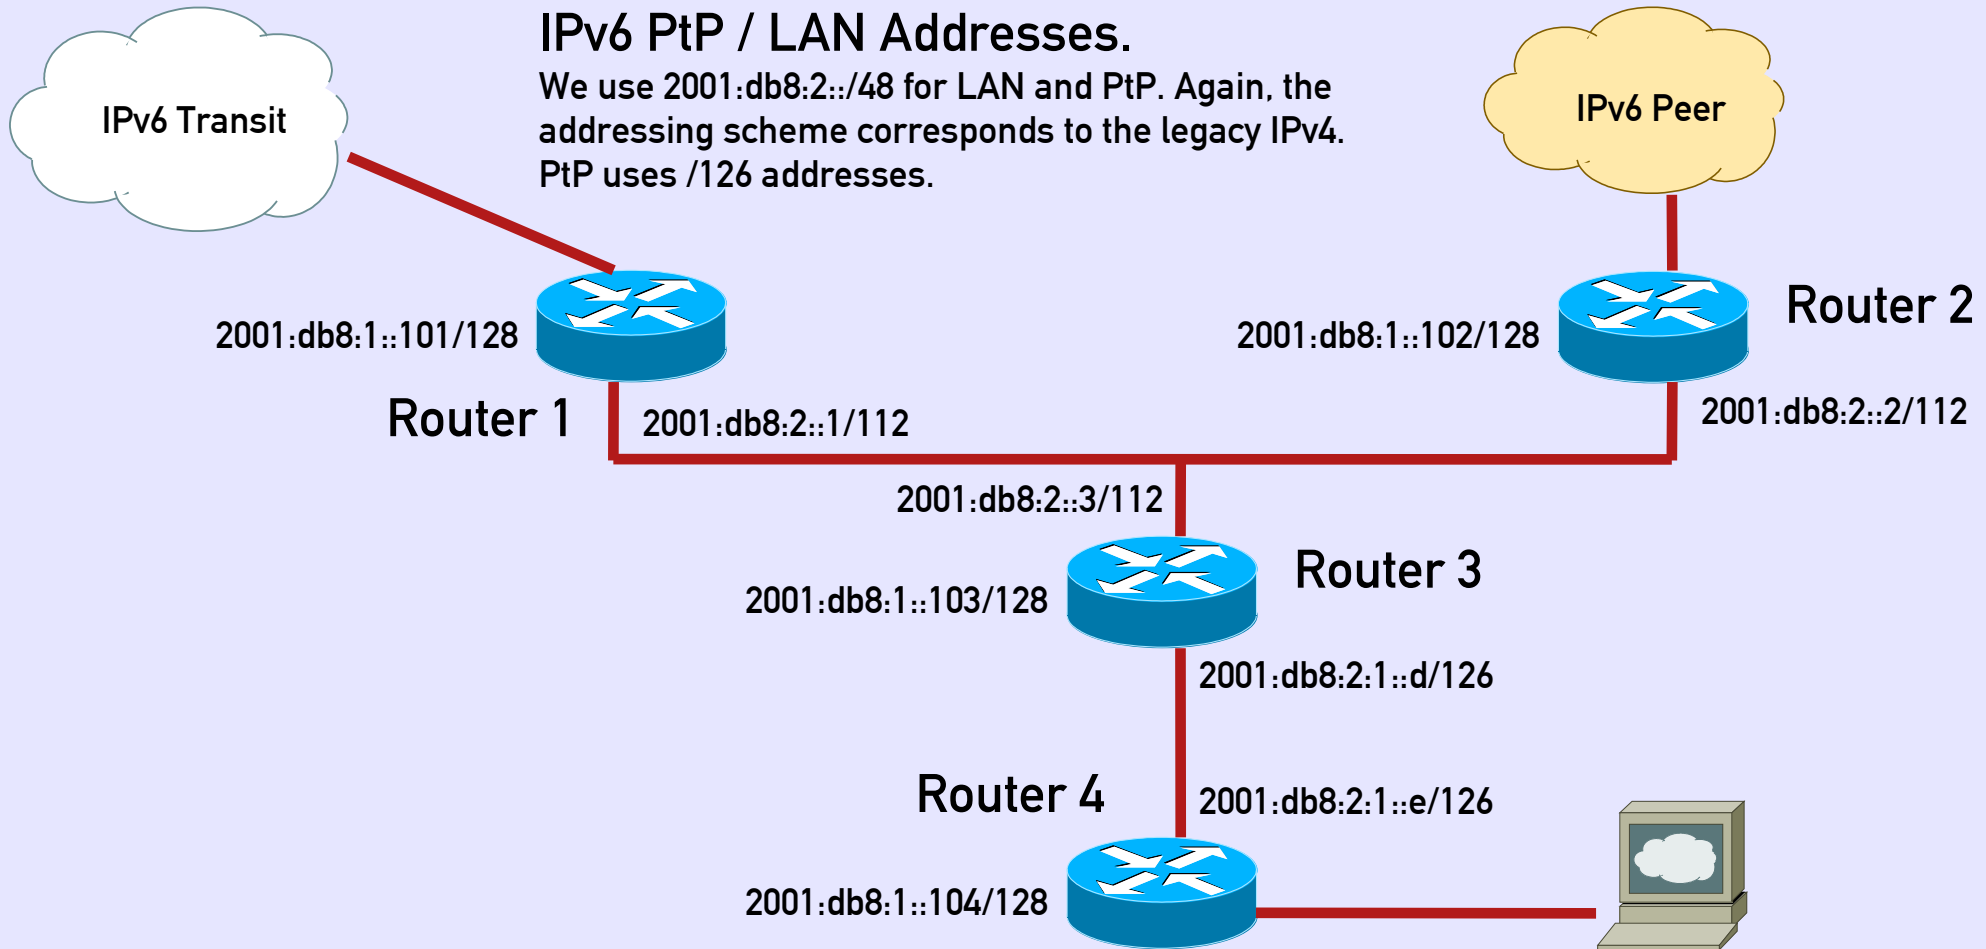
We use 2001:db8:1::/48 for Loopbacks, and the last IPv4 octet corresponds to the IPv6 Loopback address.



Migration to IPv6 #2

IPv6 PtP / LAN Addresses.

We use 2001:db8:2::/48 for LAN and PtP. Again, the addressing scheme corresponds to the legacy IPv4. PtP uses /126 addresses.





Migration to IPv6 #3

Server (End User) Addresses

Assign a /48 to every end user, Example 2001:db8:ab::/48





IPv6 Deployment #1:

Enable IPv6 on routers:

Cisco:

```
!  
ip cef    !required by ipv6 cef  
!  
ipv6 unicast-routing  
ipv6 cef  
!
```

Foundry Networks:

```
!  
ipv6 enable-acl-cam-sharing  
ipv6 unicast-routing  
!
```



IPv6 Deployment #2:

Configure Interfaces:

Cisco:

```
!  
interface lo0  
ipv6 enable  
ipv6 address 2001:db8:1::101/128  
!
```

Foundry Networks:

```
!  
interface lo 1  
ipv6 address 2001:db8:1::102/128  
ipv6 enable  
!
```



IPv6 Deployment #3:

Configure IPv6 OSPF:

Cisco:

```
!  
ipv6 router ospf 1  
  passive-interface default  
  no passive-interface x/y  
!  
interface x/y  
  ipv6 ospf 1 area 0  
!
```

Make sure that you are not talking OSPF to external interfaces (passive-interface default).



IPv6 Deployment #4:

Configure IPv6 OSPF:

Foundry Networks:

```
!  
ipv6 router ospf 1  
  area 0  
!  
interface lo 1  
  ipv6 ospf enable  
  ipv6 ospf area 0  
!  
interface ethernet x/y  
  ipv6 ospf enable  
  ipv6 ospf area 0  
  ipv6 ospf passive !do not send OSPF multicast packets here  
!
```

IPv6 Deployment #5:

Check OSPF adjacencies and the IPv6 routing table:

Cisco:

```
# sh ipv6 ospf neighbor
```

```
# sh ipv6 route [ospf|connected|static]
```

Foundry Networks:

```
# sh ipv6 ospf neighbor
```

```
# sh ipv6 route [ospf|connect|static]
```

IPv6 Deployment #6:

Configure IPv6 iBGP. Between Router 1/2/3 configure a full mesh, while Router 4 is configured as a route-reflector-client. All iBGP sessions should be configured between the loopback addresses (what else?).

Cisco / Foundry Networks:

```
!  
router bgp 65000  
  neighbor 2001:db8:1::102 remote-as 65000  
  neighbor 2001:db8:1::102 description iBGP to Router 2  
  neighbor 2001:db8:1::102 update-source loopback 0  
  !  
  address-family ipv4 unicast  
  no neighbor 2001:db8:1::102 activate  
  !  
  address-family ipv6 unicast  
  neighbor 2001:db8:1::102 activate  
  neighbor 2001:db8:1::102 send-community  
  neighbor 2001:db8:1::102 next-hop-self      !depends on config.  
  Neighbor 2001:db8:1::102 soft-reconfiguration inbound  
  !
```

IPv6 Deployment #7:

Configure IPv6 eBGP Transit / Peers. Neighbor address is provided by Transit Provider or Internet Exchange.

Cisco / Foundry Networks:

```
!  
ipv6 prefix-list MYSELF seq 5 permit 2001:db8::/32  
!  
router bgp 65000  
  neighbor x:x:x:x::x remote-as 123  
  neighbor x:x:x:x::x description IPv6 Transit Provider  
  !  
  address-family ipv4 unicast  
  no neighbor x:x:x:x::x activate  
  !  
  address-family ipv6 unicast  
  neighbor x:x:x:x::x activate  
  neighbor x:x:x:x::x send-community !depends on config  
  neighbor x:x:x:x::x next-hop-self !implicit  
  neighbor x:x:x:x::x soft-reconfiguration inbound  
  neighbor x:x:x:x::x prefix-list MYSELF out  
!
```



IPv6 Deployment #8:

Start propagating your IPv6 netblock. Advise your Transit Provider in order to adjust his inbound filter list. Don't forget the static route with high distance value...

Cisco / Foundry Networks:

```
!  
ipv6 route 2001:db8::/32 null0 [distance] 250  
!  
router bgp 65000  
  network 2001:db8::/32  
!
```

Check your propagation in various IPv6 Looking Glass.

Linux Server / Apache 2

Configure your Linux Server with IPv6:

```
add into /etc/rc.local or /etc/network/interfaces
```

```
#  
ifconfig eth0 inet6 add 2001:db8:ab::50/48  
#
```

Reboot or execute `/etc/rc.local` ... check connectivity with `ping6`.

Apache 1.x is not IPv6 compatible. If not done yet, migrate to Apache 2.

Good reading is Peter Bieringer's IPv6 Howto:

<http://www.bieringer.de/linux/IPv6/IPv6-HOWTO/>

For Apache 2 see section #4.

No default gateway anymore ...

Neighbor Discovery Protocol (ND) – no definition of the default gateway is required anymore. A router propagates himself as the Default Gateway to v6-Hosts [On by default]. Please turn it off at interfaces facing an Internet Exchange:

Cisco:

```
!  
interface lo0  
ipv6 enable  
ipv6 address 2001:db8:1::101/128  
ipv6 nd suppress-ra  
!
```

Foundry Networks:

```
!  
interface lo 1  
ipv6 address 2001:db8:1::102/128  
ipv6 enable  
ipv6 nd suppress-ra  
!
```



DNS Server

Configure your DNS record with IPv6:

```
;  
servername.com.      A           192.168.3.50  
;  
servername.com.      AAAA        2001:db8:ab::50  
www                  CNAME      servername.com.
```

Note: DNS servers don't need to be accessible via IPv6 to resolve AAAA records, however IPv6-only clients may have problems to resolve.

IPv6 Reverse DNS Server

This is a nasty thing. I suppose we are going to postpone it... don't we always have a task list with priorities?!



Check your IPv6 Connectivity

If you don't have an IPv6-enabled client, try the tools of SixXS:

<http://ipv4gate.sixxs.net/>

<http://www.sixxs.net/>



More IPv6 Links:

Tunnelbrokers (IPv6 for IPv4-only broadband users):

- <http://tunnelbroker.net/> (by Hurricane Electric)
- <http://www.sixxs.net/faq/account/?faq=10steps>

From RIPE #56

- Structural Problems in the IPv6 Routing:

http://rosie.ripe.net/ripe/meetings/ripe-56/presentations/uploads/Tuesday/Plenary%2016:00/upl/Schmidt-Structural_problems

- IPv6 PI Space – refer to Gert Dörings Presentation:

http://rosie.ripe.net/ripe/meetings/ripe-56/presentations/uploads/Thursday/Address%20Policy%202/upl/Doering-IPv6_Routing

- Measurement of the IPv6 deployment by Geoff Huston:

http://rosie.ripe.net/ripe/meetings/ripe-56/presentations/uploads/Wednesday/Plenary%2009:00/upl/Huston-Measuring_IPv6_I

IPv6 BOGON information (use the “Relaxed” template):

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



History of IPv6 @ SwiNOG

- SwiNOG #2 (March 21, 2001) by Simon Leinen, SWITCH:
IPv6 Operational Experiences

<http://www.swinog.ch/meetings/swinog2/leinen-ipv6.ps>

- SwiNOG #8 (March 24, 2004) by Simon Leinen, SWITCH:
IPv6 Peering Group

http://www.swinog.ch/meetings/swinog8/ipv6_swinog.pdf

- SwiNOG #11 (October 20, 2005) by Jeroen Massar
SixXS: Looking at deploying IPv6

<http://www.swinog.ch/meetings/swinog11/SwiNOG11-DeployingIPv6.pdf>



Help to deploy IPv6 in Switzerland Europe!

- Obtain your IPv6 allocation now
- Init7 offers IPv6 connectivity to all existing Colo/Transit customers free of charge according to their current CDR (BGP4 feed or static routing).
- Init7 offers IPv6 transit connectivity to non-customers (BGP4 feed only) free of charge until mid of 2009 (Cabling on prospects expense) [Disclaimer: up to a reasonable volume of traffic].
- DSL/Cable customers: no decent IPv6 CPE devices available yet, use Tunnelbrokers until further notice. Blame ... | Fritz | Zyxel | Netopia | ... |



IPv6. Now!

Questions?

Fredy Künzler, AS13030
kuenzler at init7 dot nospam dot net

www.init7.net
www.blogg.ch
www.bgp-and-beyond.com