


The logo for INIT7, featuring the letters 'INIT7' in a stylized, blue, blocky font. The '7' is slightly larger and more prominent. The logo is set against a light blue background with a subtle grid pattern.

IPv6. Now!



PLNOG #3, Sept. 10, 2009

(updated from "IPv6 deployment for the IPv4 clueful..". - LINX v6 Meeting, March 13, 2009 and "IPv6 now!!!" from SwiNOG #18, April 18, 2009)

Fredy Künzler, AS13030
kuenzler at init7 dot nospam dot net

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

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



**[this space is a
placeholder for one or
more marketing slides]**

Init Seven AG / Init7

- Internet Service Provider, based in Zurich / Switzerland
- privately owned company
- in Switzerland: mainly Business- and Wholesale Customers
- in Europe: MPLS and IP connectivity (full and partial table)
- own international dual-stacked v4 and v6 backbone (AS13030), mostly 10gig enabled (two rings: Zurich-Geneva-London-Amsterdam-Frankfurt-Zurich and Zurich-Basel-Frankfurt(2)-Zurich)
- connected to 20+ internet exchanges (i.E. PL-IX)



IPv4 & IPv6 Statistics

v4 Addresses
446,064,349 ↓

v4 /8s Left
10% (28/256)

v6 Networks
5% (1,643/32,497)

v6 Ready TLDs
79% (223/280)

v6 Glue
1,503

v6 Domains
1,517,089 ↑

705
Days remaining

HURRICANE ELECTRIC
INTERNET SERVICES

Face facts:

Less than ~~1000~~ **705** days until the free IPv4 pool gets exhausted.



IPv4 & IPv6 Statistics

v4 Addresses
446,064,349 ↓

v4 /8s Left
10% (28/256)

v6 Networks
5% (1,643/32,497)

v6 Ready TLDs
79% (223/280)

v6 Glue
1,503

v6 Domains
1,517,089 ↑

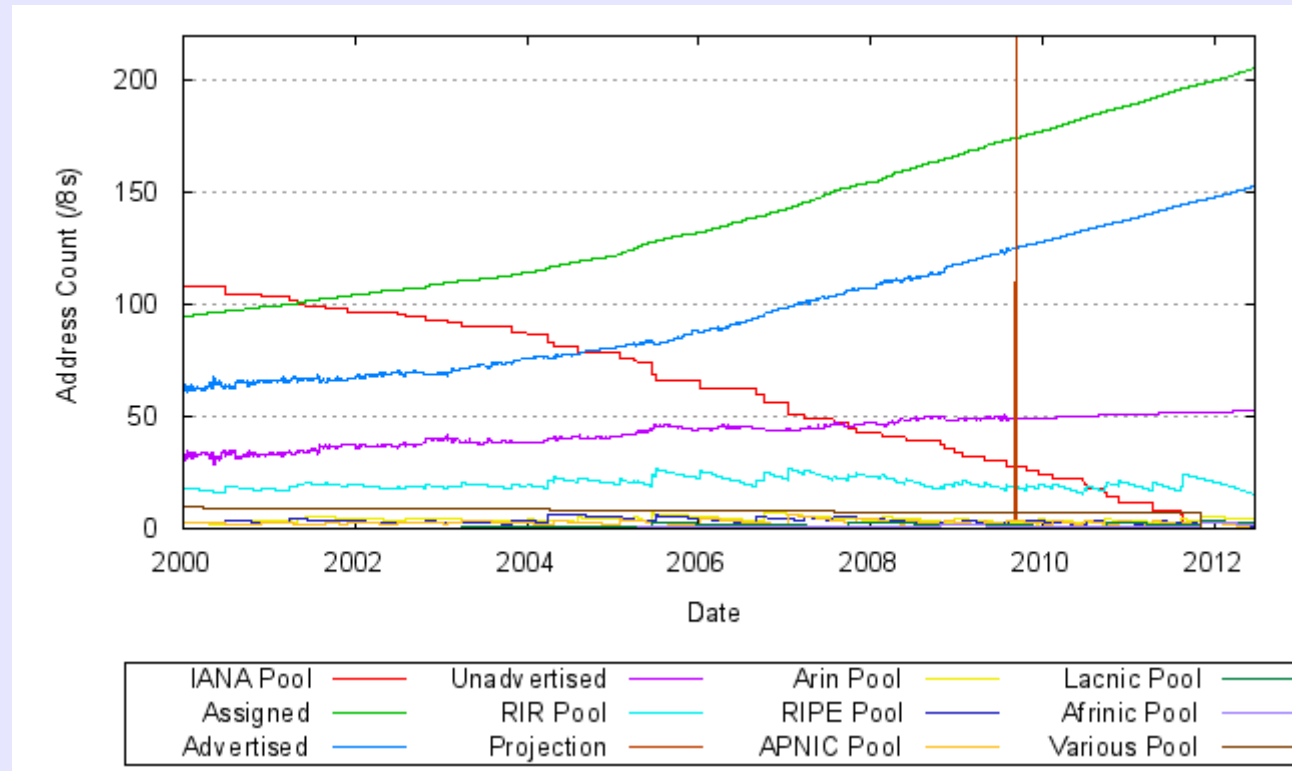
705

Days remaining

HURRICANE ELECTRIC
INTERNET SERVICES

- Prediction by Geoff Huston, APNIC
<http://www.potaroo.net/tools/ipv4/>

The available / unassigned pool of IPv4 /8 blocks will be empty in 2011.





IPv4 & IPv6 Statistics

v4 Addresses
446,064,349 ↓

v4 /8s Left
10% (28/256)

v6 Networks
5% (1,643/32,497)

v6 Ready TLDs
79% (223/280)

v6 Glue
1,503

v6 Domains
1,517,089 ↑

705
Days remaining

HURRICANE ELECTRIC
INTERNET SERVICES

Are you and your customers aware of the business continuity risk?

- 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 [insert KPMG / PriceWaterhouseCoopers etc.] start requiring disclosure of IPv4 exhaustion as a business continuity risk, as well as the presence or lack thereof of an IPv6 plan?



IPv4 & IPv6 Statistics

v4 Addresses

446,064,349 ↓

v4 /8s Left

10% (28/256)

v6 Networks

5% (1,643/32,497)

v6 Ready TLDs

79% (223/280)

v6 Glue

1,503

v6 Domains

1,517,089 ↑

705

Days remaining

HURRICANE ELECTRIC
INTERNET SERVICES

Excuses – why people think they don't have to act now ...

- IDW - Informatik Services of the City of Winterthur:

<http://fetch.stadt.winterthur.ch/upload/politik/weisungen/W09032.pdf>

[...] Die Stadt Winterthur verfügt somit über genügend öffentliche und private IP-Adressen, weshalb sich für sie die erwähnte Verknappung von IPv4-Adressen relativiert.

[...] Aktuell ist kein Handlungsbedarf vorhanden [...]

[...] The City of Winterthur has sufficient public and private IP addresses, which relativizes the aforementioned shortage of IPv4 addresses. [...] Currently there is no need for action [...]



IPv4 & IPv6 Statistics

v4 Addresses
446,064,349 ↓

v4 /8s Left
10% (28/256)

v6 Networks
5% (1,643/32,497)

v6 Ready TLDs
79% (223/280)

v6 Glue
1,503

v6 Domains
1,517,089 ↑

705
Days remaining

HURRICANE ELECTRIC
INTERNET SERVICES

Excuses – what people forget when they excuse their v6 ignorance:

- As soon as the first person (a citizen of your local city) gets only v6 connectivity due to v4 shortage, and is supposed to [fill in electronic tax form | vote electronically | access cultural activity calendar | etc.], then the local government is in trouble.
- think of thousands of other examples ...



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 an DSL/Cable enduser ... or a TIER-1 tunnel hustler ;-)

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 =`

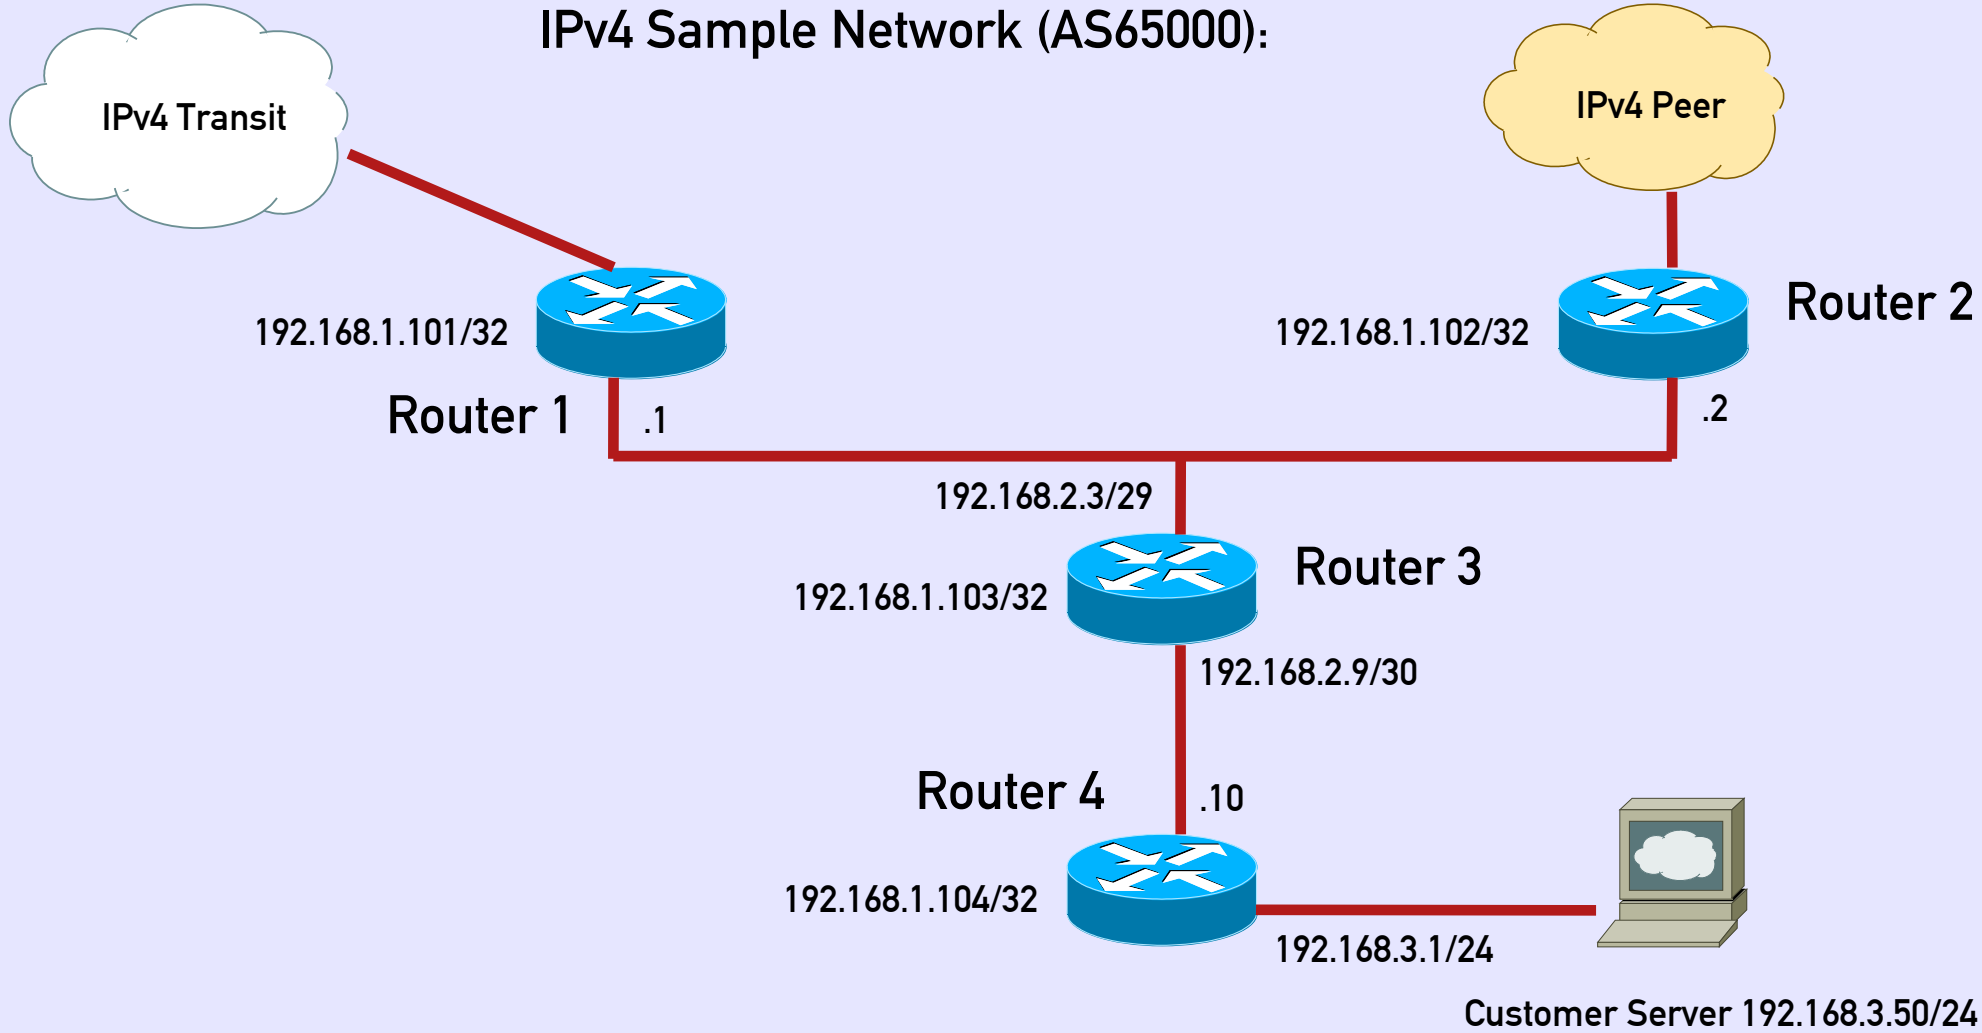
`2001:0db8:0000:0000:0000:0000:0000:0000/32`

- 0 are omitted
- two colons are replacing many 0

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



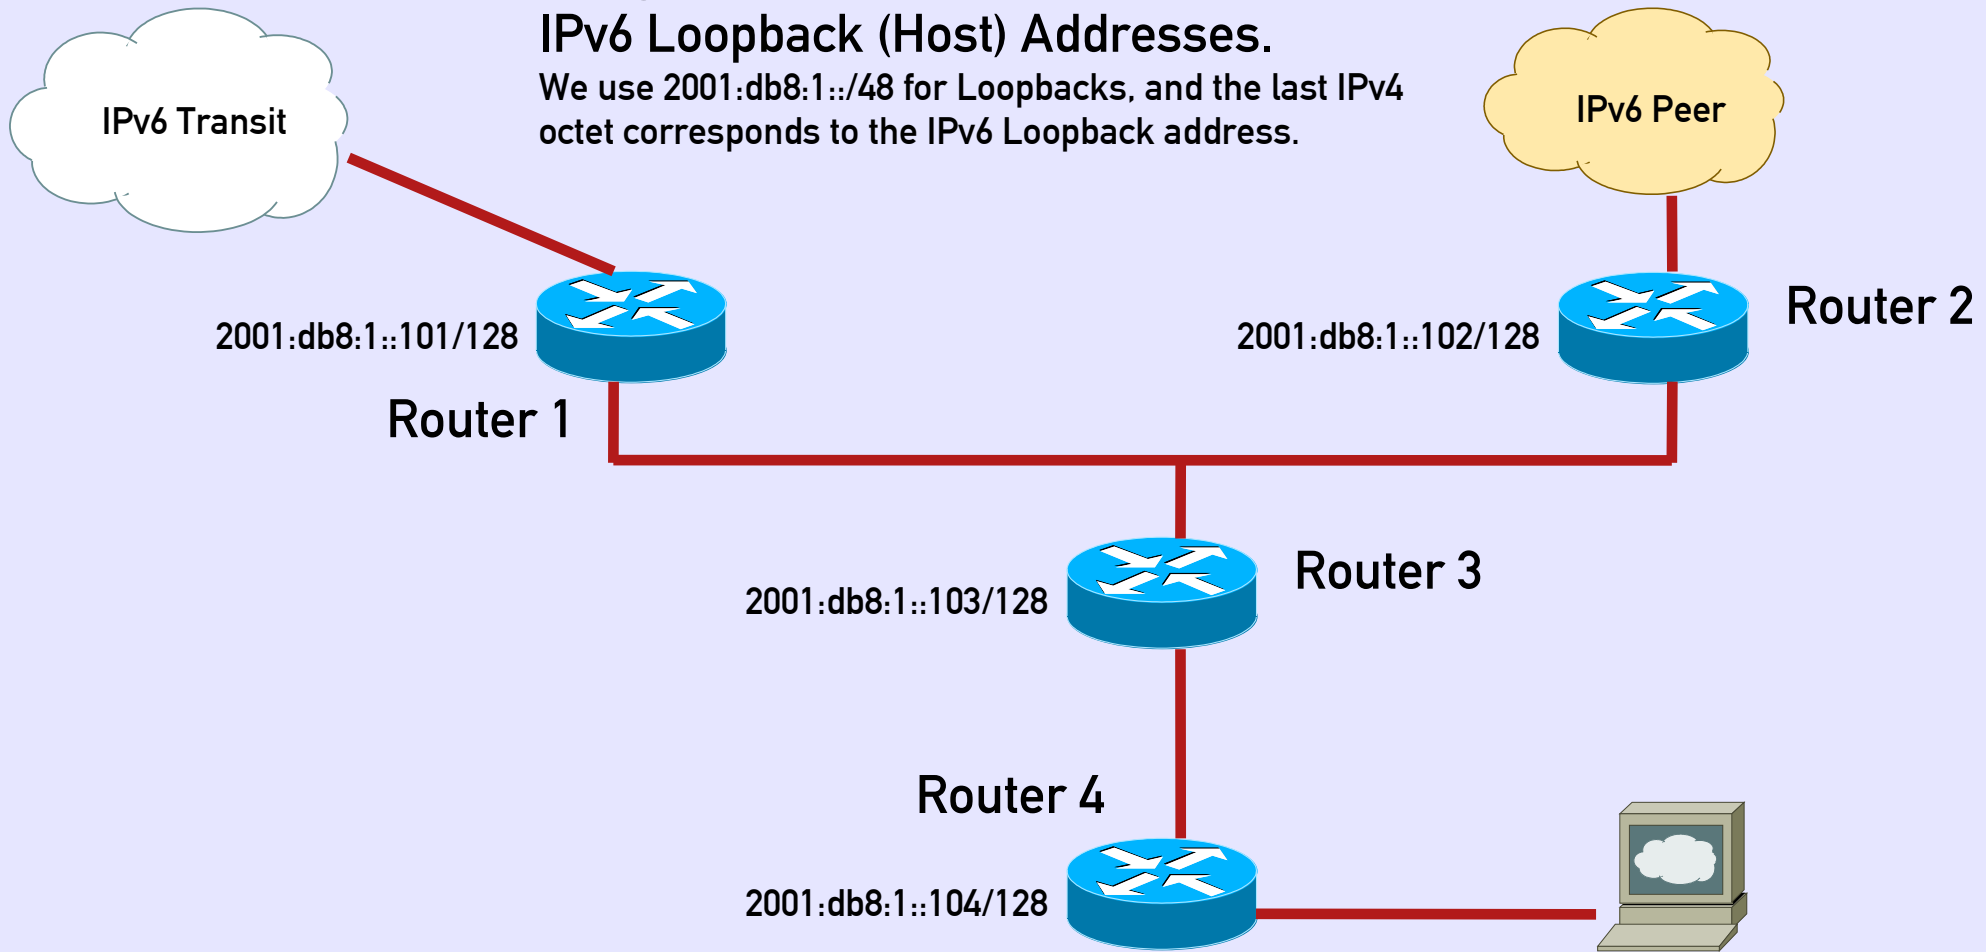
IPv4 Sample Network (AS65000):



Migration to IPv6 #1

IPv6 Loopback (Host) Addresses.

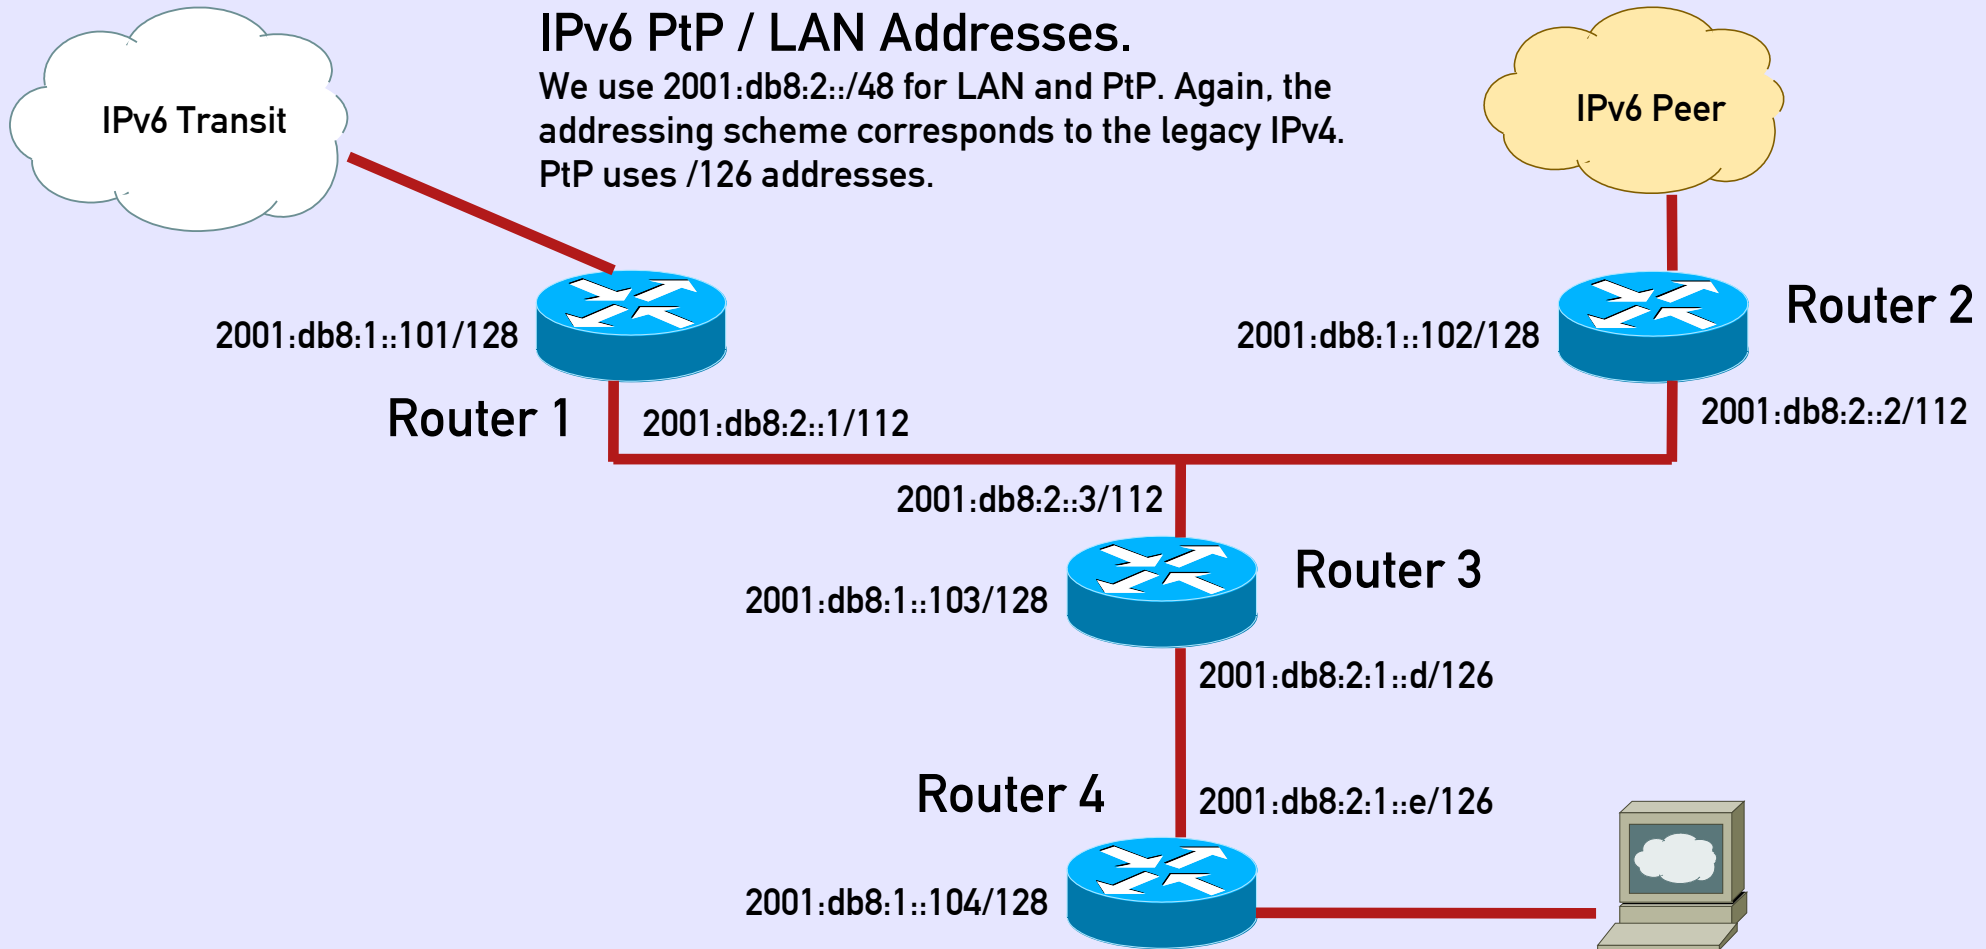
We use 2001:db8:1::/48 for Loopbacks, and the last IPv4 octet corresponds to the IPv6 Loopback address.



Migration to IPv6 #2

IPv6 PtP / LAN Addresses.

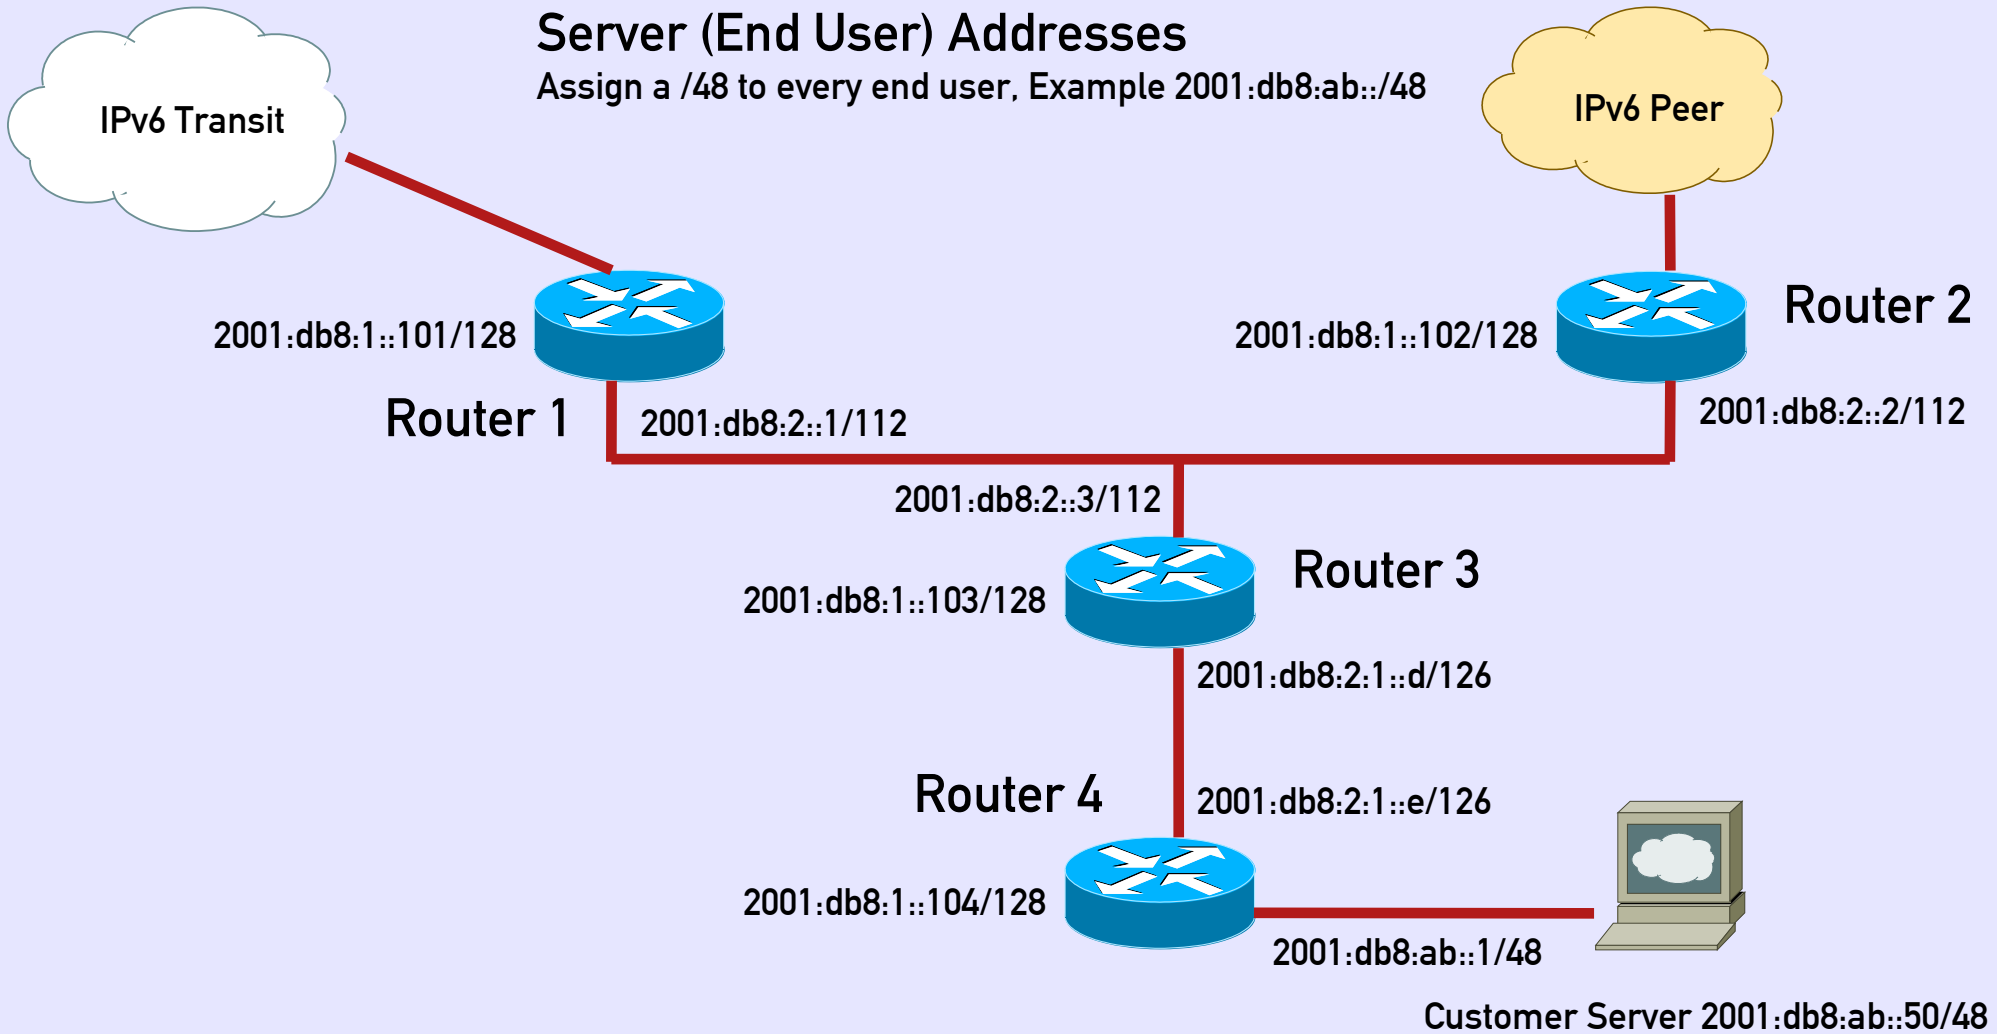
We use 2001:db8:2::/48 for LAN and PtP. Again, the addressing scheme corresponds to the legacy IPv4. PtP uses /126 addresses.



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  
!
```

Brocade (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  
!
```

Brocade (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:

Brocade (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]
```

Brocade (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 / Brocade (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 / Brocade (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 / Brocade (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/network/interfaces
```

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

Reboot or restart network stack ... 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  
!
```

Brocade (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

IPv6 Reverse DNS Zone Builder for BIND 8/9

<http://www.fpsn.net/index.cgi?pg=tools&tool=ipv6-inaddr>

```
;
; 2001:db8:abcd::/48
;
; Zone file built with the fpsn.net IPv6 Reverse DNS zone builder
; http://tools.fpsn.net/ipv6-inaddr
;
$TTL 3d; Default TTL (bind 8 needs this, bind 9 ignores it)
@      IN SOA d.c.b.a.8.b.d.0.1.0.0.2.ip6.int. noc.example.com. (
        200904010      ; Serial number (YYYYMMdd)
        24h            ; Refresh time
        30m           ; Retry time
        2d            ; Expire time
        3d            ; Default TTL (bind 8 ignores this, bind 9 needs it)
)

; Name server entries
IN     NS     ns1.example.com.
IN     NS     ns2.example.com.

; IPv6 PTR entries

; Subnet #1
$ORIGIN 0.0.0.0.d.c.b.a.8.b.d.0.1.0.0.2.ip6.int.

1.0.0.0.f.e.e.b.d.a.e.d.0.0.0.0      IN     PTR     gate.example.com.
2.0.0.0.f.e.e.b.d.a.e.d.0.0.0.0      IN     PTR     www.example.com.

;
; End of zone file.
; Thank you for using the fpsn.net IPv6 Reverse DNS zone builder
```



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/>

Some thoughts about v6 peering ...

We already see „Peering Wars“ in the IPv6 world, even though many responsible of larger networks think: „IPv6 is still a sandbox“.

Currently the IPv6 routing table contains ~2000 prefixes. Most larger v6 networks have some ~1800 to ~1950 prefixes (Source: route-views.linx.routeviews.org):

```
route-views.linx.routeviews.org> sh bgp su
```

Neighbor State/PfxRcd	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	
2001:7f8:4::b62:1	4	2914	84005	25654	0	0	0	01w1d21h	1673
2001:7f8:4::cb9:1	4	3257	127076	25654	0	0	0	01w1d21h	1863
2001:7f8:4::d1c:1	4	3356	35357	13039	0	0	0	01w1d21h	1899
2001:7f8:4::1a6a:1	4	6762	5884	1571	0	0	0	1d01h38m	1830
2001:7f8:4::1b1b:1	4	6939	2054890	13039	0	0	0	01w1d21h	1966
2001:7f8:4::32e6:1	4	13030	50969	13039	0	0	0	01w1d21h	1921



Largest v6 backbones:

(Source: <http://bgpmon.net/weathermap.php?inet=6>)

Top 25 Transit/Upstream AS numbers in the routing table based on OriginAS

This represents the number of distinct number of originAS's for which these AS's provide transit (i.e. are seen as upstream AS)

Transit/Upstream AS number	AS description	number of AS's for which this AS is providing Transit
AS6939	HURRICANE - Hurricane Electric, Inc.	531
AS3257	TINET-BACKBONE Tinet SpA	152
AS3549	GBLX Global Crossing Ltd.	126
AS2914	NTT-COMMUNICATIONS-2914 - NTT America, Inc.	115
AS6453	GLOBEINTERNET TATA Communications	74
AS13030	INIT7 Init Seven AG, Zurich, Switzerland	63
AS6175	SPRINTLINK9	62
AS3356	LEVEL3 Level 3 Communications	59
AS30071	ASN-TBONE - TowardEX Technologies International, Inc.	57
AS1299	TELIANET TeliaNet Global Network	48
AS12859	NL-BIT BIT BV	44
AS286	KPN KPN Internet Backbone	41
AS1273	CW Cable and Wireless plc	38
AS20965	GEANT The GEANT IP Service	27
AS13237	LAMBDA-NET-AS European Backbone of LambdaNet	24
AS11537	ABILENE - Internet2	24
AS2497	IJ Internet Initiative Japan Inc.	22
AS1257	TELE2	21
AS4826	VOCUS-BACKBONE-AS Vocus Connect International Backbone	16
AS12702	UUNET-EMEA-IPV6 -- Verizon Business EMEA - IPv6 only autonomous System	16
AS5539	SPACENET SpaceNET AG, Munich	15
AS9002	RETN-AS ReTN.net Autonomous System	14
AS2516	KDDI KDDI CORPORATION	13
AS38610	APNIC-JP-RD -- APNIC R&D Centre JP	13





IPv6 peering wars #1 ...

Recent quote from the ipv6-ops@lists.cluenet.de mailinglist:

`„We got the IPv6 details for our Cogent BGP session today, and yes, it's not a full table yet.`

`Cogent only has 781 routes, that they advertise to ourselves, so anyone not in there obviously would have no route to them.” (Aug 14, 2009)`

Why? Cogent (AS174) claims to be a so-called TIER-1 (better default-free network). They don't get a v6 fullfeed at all.

Other networks claim to have the same status (i.E. Hurricane Electric he.net AS6939). For presumably political reasons there is no v6 peering between AS174 and AS6939 (of course this is just one example).

IPv6 peering wars #2 ...

The result:

- Cogent (AS174) has a crippled v6 „fulltable“ with only ~800 IPv6 prefixes. So watch out before you sign for „v6 transit“ of [AS174 | any other vendor] ...
- he.net (AS6939) has no access to Cogent's v6 customers, unless they are otherwise connected. Cogent's own prefix (2001:550::/32) is not reachable from he.net:

```
route-server.he.net> sh bgp ipv6 2001:550::/32
% Network not in table
```

IPv6 peering wars #3 ...

v6 Peering Policies need to be revised:

- Some networks (mostly incumbents):

We peer v6 only with networks which have an existing v4 peering ... *(and my comment: regardless whether they can spell v6 or not ...)*

- some networks „leak“ full v6 tables to others, without transiting their prefixes properly ...

- many large networks don't see a business case (yet) in v6, especially in the United States – no surprise they don't care about a full table ...

IPv6 peering considerations ...

Get your v6 peerings now! 50%+ of the routing table via peering can be achieved rather easy. If you happen to be also on AMS-IX and DE-CIX, success will be bigger ...


- peer with v6 routeservers if available
- ask Hurricane Electric (AS6939) (~950+ prefixes)
- ask Init7 (AS13030) (~100+ prefixes)
- many other networks are still willing to peer openly IPv6
- transit some prefixes, convince your existing downstream customers to implement IPv6 – you may convince your potential peers if you're not carrying just one prefix ...
- and last but not least: get one or more decent v6 transits

- Lesson from the v4 peering history: „old“ ASNs have many legacy peers, despite of todays size of the network



„90% of the v6 traffic is PING!“

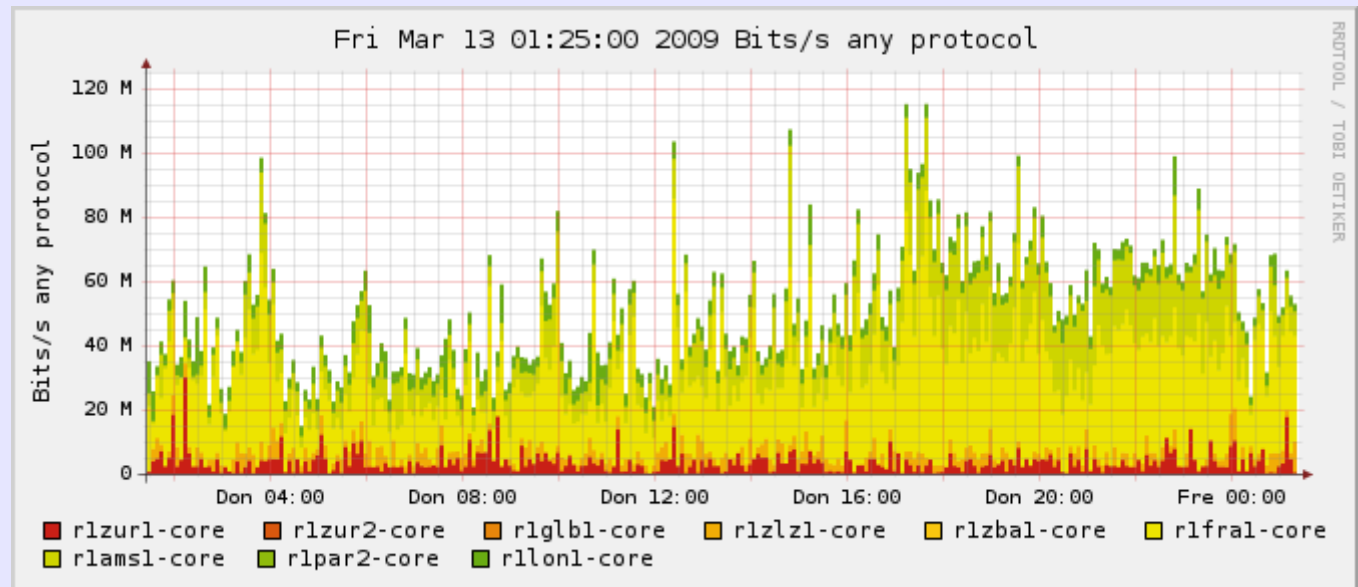
Not true!





Yes, there is real IPv6 Traffic!

Up to 100 mbps on the edges of the Init7 network:





Some 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 end 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 | ... |



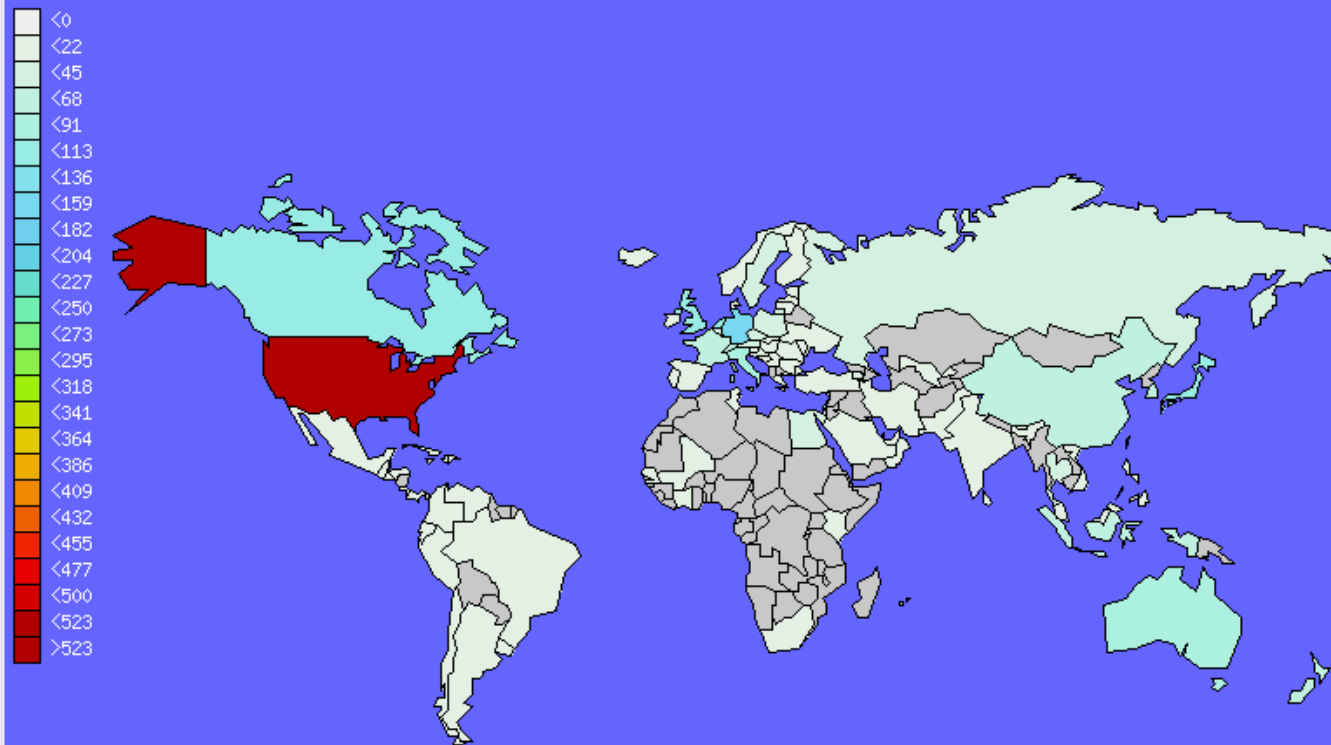
Largest v6 countries:

(Source: <http://bgpmon.net/weathermap.php?inet=6>)

Prefixes per country

the number represent the number of unique prefixes found in the routing table

1	United States	546
2	Germany	153
3	Japan	107
4	Canada	104
5	United Kingdom	93
6	Netherlands	92
7	Italy	86
8	Australia	71
9	China	62
10	France	57
11	Indonesia	55
12	Switzerland	50
13	Korea, Republic of	50
14	New Zealand	47
15	Russian Federation	40
16	Sweden	39
17	Poland	39
18	EU	35
19	Czech Republic	34
20	Hong Kong	30
21	Austria	29
22	Egypt	24
23	Thailand	24
24	Taiwan, Province of China	22





IPv6. Now!

Questions?

Fredy Künzler, AS13030
kuenzler at init7 dot nospam dot net

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