

FORTHNETGROUP



forthnet



nova



IPv6

# deployment in Forthnet

**RIPE 65 (Amsterdam)**

Anastasios Chatzithomaoglou  
achatz@forthnetgroup.gr

**IP Engineering**

26-Sep-2012

the Greeks are



and are here to stay


Thank you very  
much RIPE for the  
donation in favor  
of Greece's debt

Congratulations! a

to: [Anastasios Chatzithomaoglou](#)  
from:

Hi Anastasios, Congratulations, your name was drawn as the winner for rating the plenary presentations. Enjoy your prize! Kind regards, The RIPE Programme Committee

gift amount:  
**£80.00**


claim code:  


[Start shopping](#)

Learn how to redeem your gift certificate

## Agenda

- ❑ Intro-addiction
- ❑ Why, When and Where
- ❑ Experience and Experiments
- ❑ Nanobots
- ❑ Game Consoles







## IPv6 Support Required for All IP-Capable Nodes

*Given the global lack of available IPv4 space, and limitations in IPv4 extension and transition technologies, this document advises that **IPv6 support is no longer considered optional**. It also cautions that there are places in existing IETF documents where the term "IP" is used in a way that could be misunderstood by implementers as **the term "IP" becomes a generic** which can mean IPv4 + IPv6, IPv6-only, or IPv4-only, depending on context and application.*

- ☐ *New IP implementations must support IPv6.*
- ☐ *Updates to current IP implementations should support IPv6.*
- ☐ *IPv6 support must be equivalent or better in quality and functionality when compared to IPv4 support in a new or updated IP implementation.*
- ☐ *New and updated IP networking implementations should support IPv4 and IPv6 coexistence (dual-stack), but must not require IPv4 for proper and complete function.*
- ☐ *Implementers are encouraged to update existing hardware and software to enable IPv6 wherever technically feasible.*

RFC 6540 (BCP, Apr 2012)

## Some interesting (or not) numbers

### General:

- ❑ Largest alternative fixed broadband provider in Greece
- ❑ ~500k Internet subscribers (Forthnet) & ~400k PayTV subscribers (Nova)

### IPv6 related (current):

- ❑ ~100 Dual-Stack active –pilot- subscribers (60% bridging, 40% CPE)
- ❑ ~2k Dual-Stack active subscribers (95% bridging, 5% CPE)
- ❑ ~5k Dual-Stack ready subscribers (100% CPE, to be activated in Nov 2012)
- ❑ ~140k Dual-Stack capable subscribers (100% CPE, to be activated after firmware upgrade)

### IPv6 related (future):

- ❑ ~30k Dual-Stack active subscribers until 21st Dec 2012
- ❑ ~150k Dual-Stack active subscribers until end of 2013 (Maya permitting)

### Rules:

- ❑ Since 2011 every new CPE provided to customers MUST support\* Dual-Stack
- ❑ Since 2012 every new CPE provided to customers MUST support\* Dual-Stack, DS-Lite
- ❑ From 2013 every new CPE provided to customers MUST support\* Dual-Stack, DS-Lite, PCP
- ❑ \* with or without firmware upgrade





## Why IPv6 in Forthnet

### ☐ **Business Continuity**

- ☐ How well would our business run if we could only talk to only a part of our customers?
- ☐ What would we do if we could no longer obtain a crucial part in our business sector?
- ☐ What if customers want IPv6 and we can only supply IPv4?
- ☐ **What if we can't support customer growth due to IPv4 exhaustion?**



### ☐ **Technological Imperative**

- ☐ Internet is evolving, do we really want to ignore our main business?
- ☐ How are we going to serve all these mobile/home devices in the future?

### ☐ **Global Competitiveness**

- ☐ All major carriers are already deploying it, why stay behind?

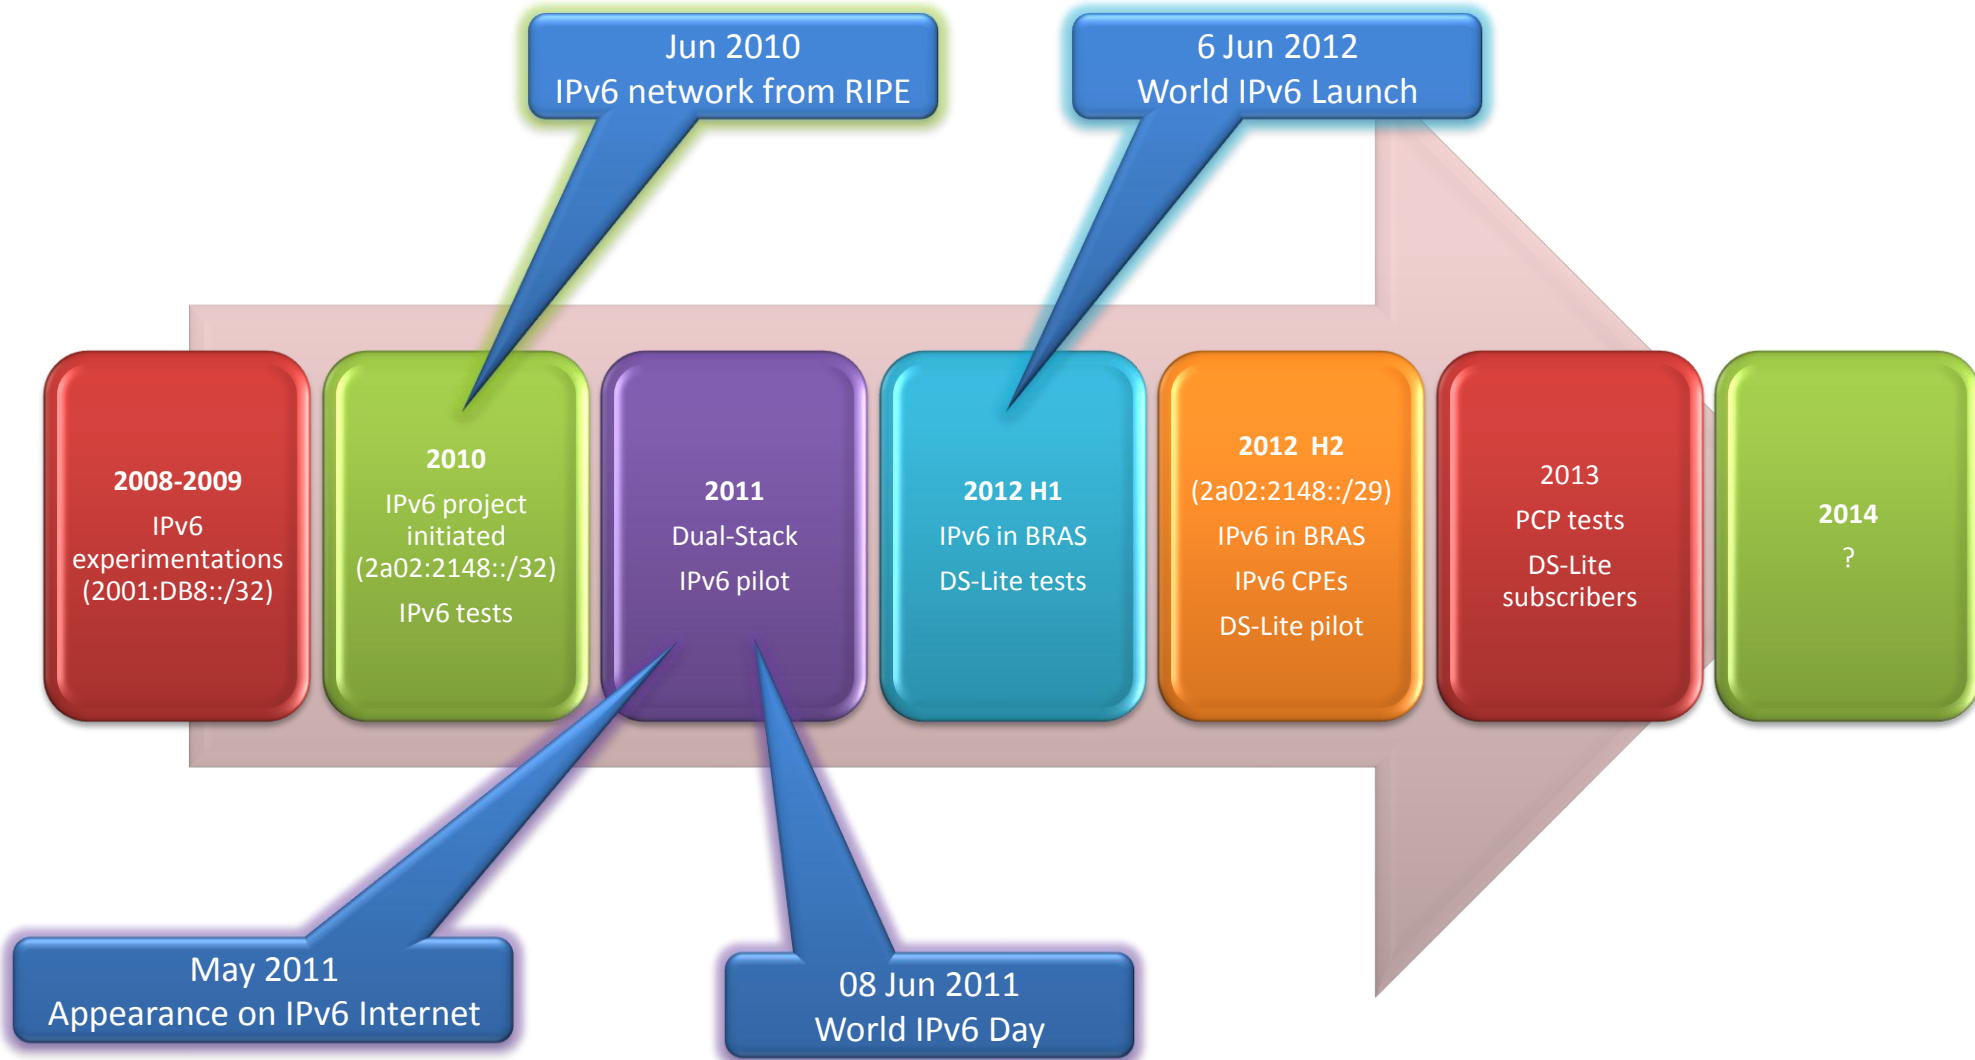
### ☐ **Government Requirement**

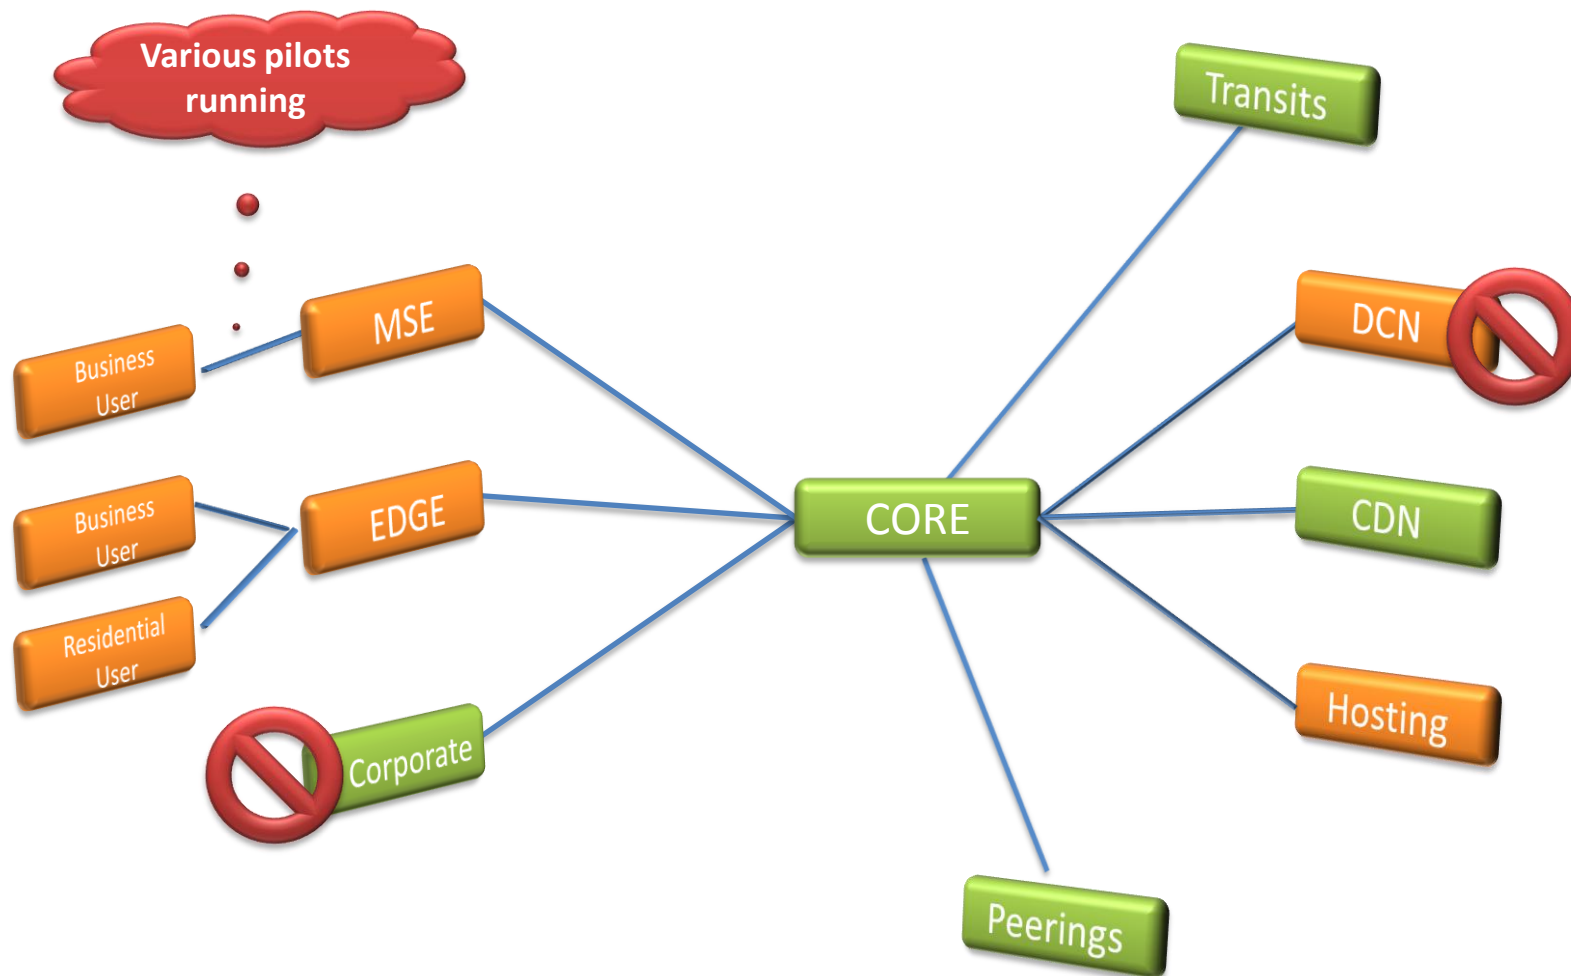
- ☐ What will happen if IPv6 becomes a national IT strategy?
- ☐ What about EU initiatives?

IPv6 Project was officially initiated by **Technical Division**  
and as such it met many difficulties until endorsed by all other Divisions



# When





Ready

Running

 **IT Alert**

Past is experience,  
present is  
experiment, future  
is expectation. Use

your experience in  
your experiments  
to achieve your  
expectations...




## Experience until now

RIPE IPv6 Training Courses  
Hellenic IPv6 Task Force Meetings/Workshops  
Internal IPv6 Trainings  
Vendor IPv6 Workshops  
World IPv6 Day/Launch  
IETF WGs

- Educate - Participate

IPv6 Tests on various network devices  
IPv6 Web & DNS tests  
IPv6 Pilot for Residential Customers  
IPv6 Pilot for Business Customers

- Explore - Activate



## World IPv6 Day (8-Jun-2011)

### Forthnet Actions


- ☐ Live IPv6 streaming through WebTV ([webtv.ipv6.forthnet.gr](http://webtv.ipv6.forthnet.gr))
- ☐ Web site with IPv6 information ([ipv6.forthnet.gr](http://ipv6.forthnet.gr))
- ☐ Web site with IPv6 test ([test-ipv6.forthnet.gr](http://test-ipv6.forthnet.gr))
- ☐ IPv6 pilot for Residential Customers (@ipv6forthnet.gr)






# Pilot Users

~1/5 of pilot users had actually managed to activate IPv6 after 1 week

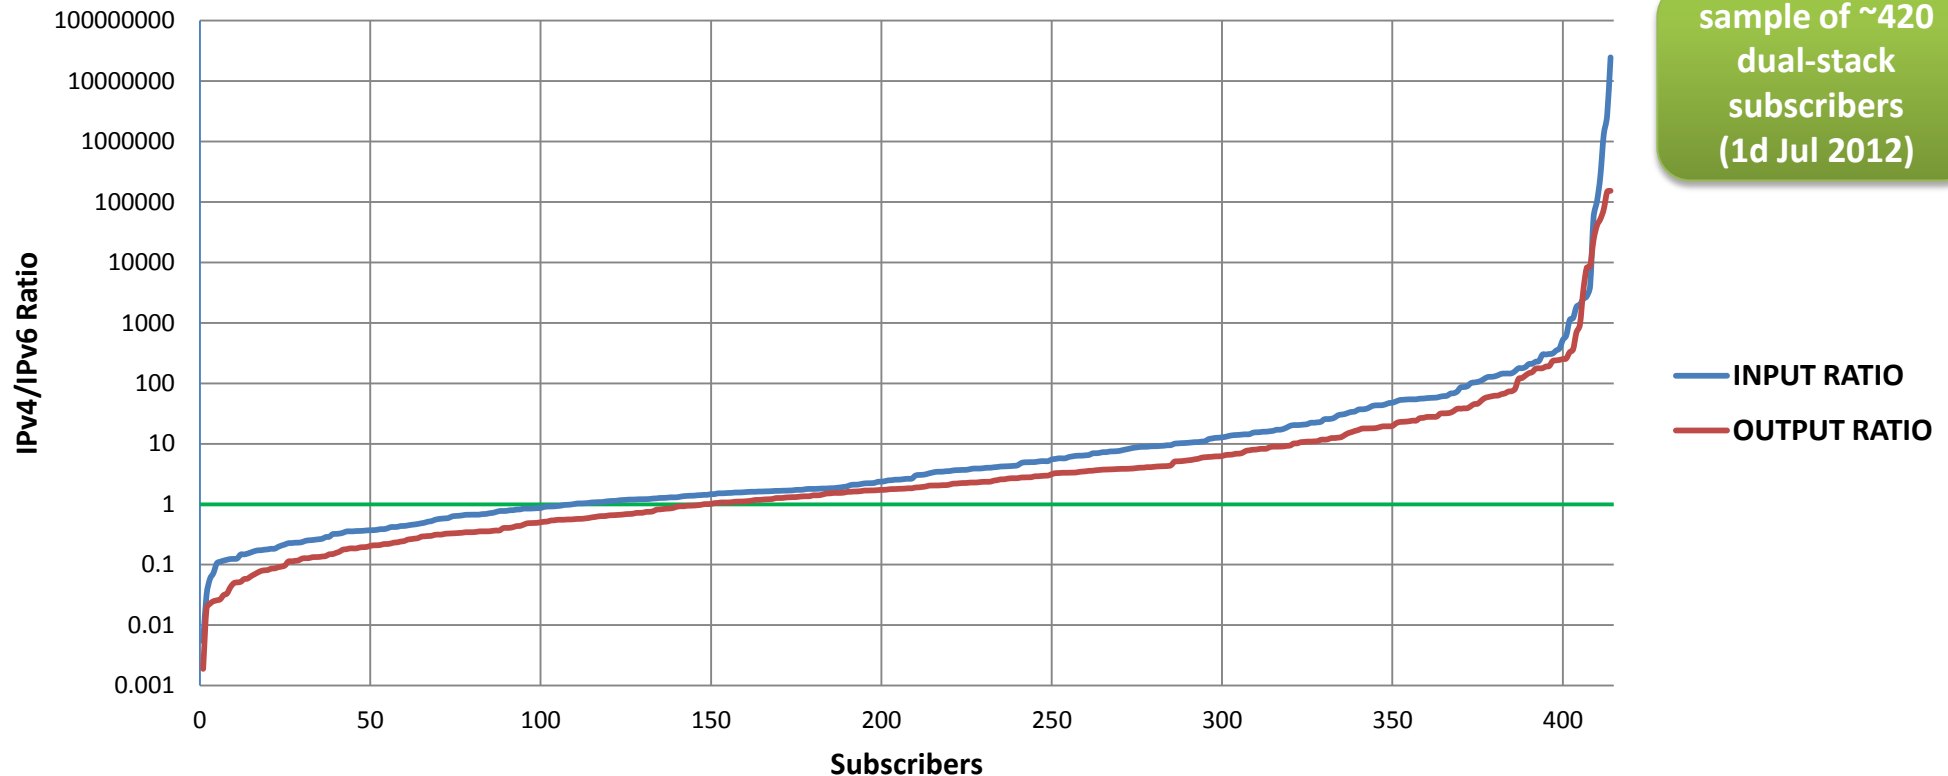


## 2 days after W6L

After enabling dual-stack in some of our production BRAS  
we started seeing major IPv6 traffic increase  
Many users had already IPv6 enabled in their desktops/CPEs







## IPv4/IPv6 Input & Output Bytes Ratio



**36%/28%** of dual-stack subscribers download/upload more IPv6 than IPv4  
 Largest download and/or upload ratio was **99,8% IPv6** and **0,2% IPv4**  
 Highest IPv6 download/upload of a single subscriber was **12/1,3 GB**  
 Highest IPv4 download/upload of a single subscriber was **49/27 GB**  
**...IPv6 traffic is still small compared to IPv4...IPv6 upload is even smaller...**

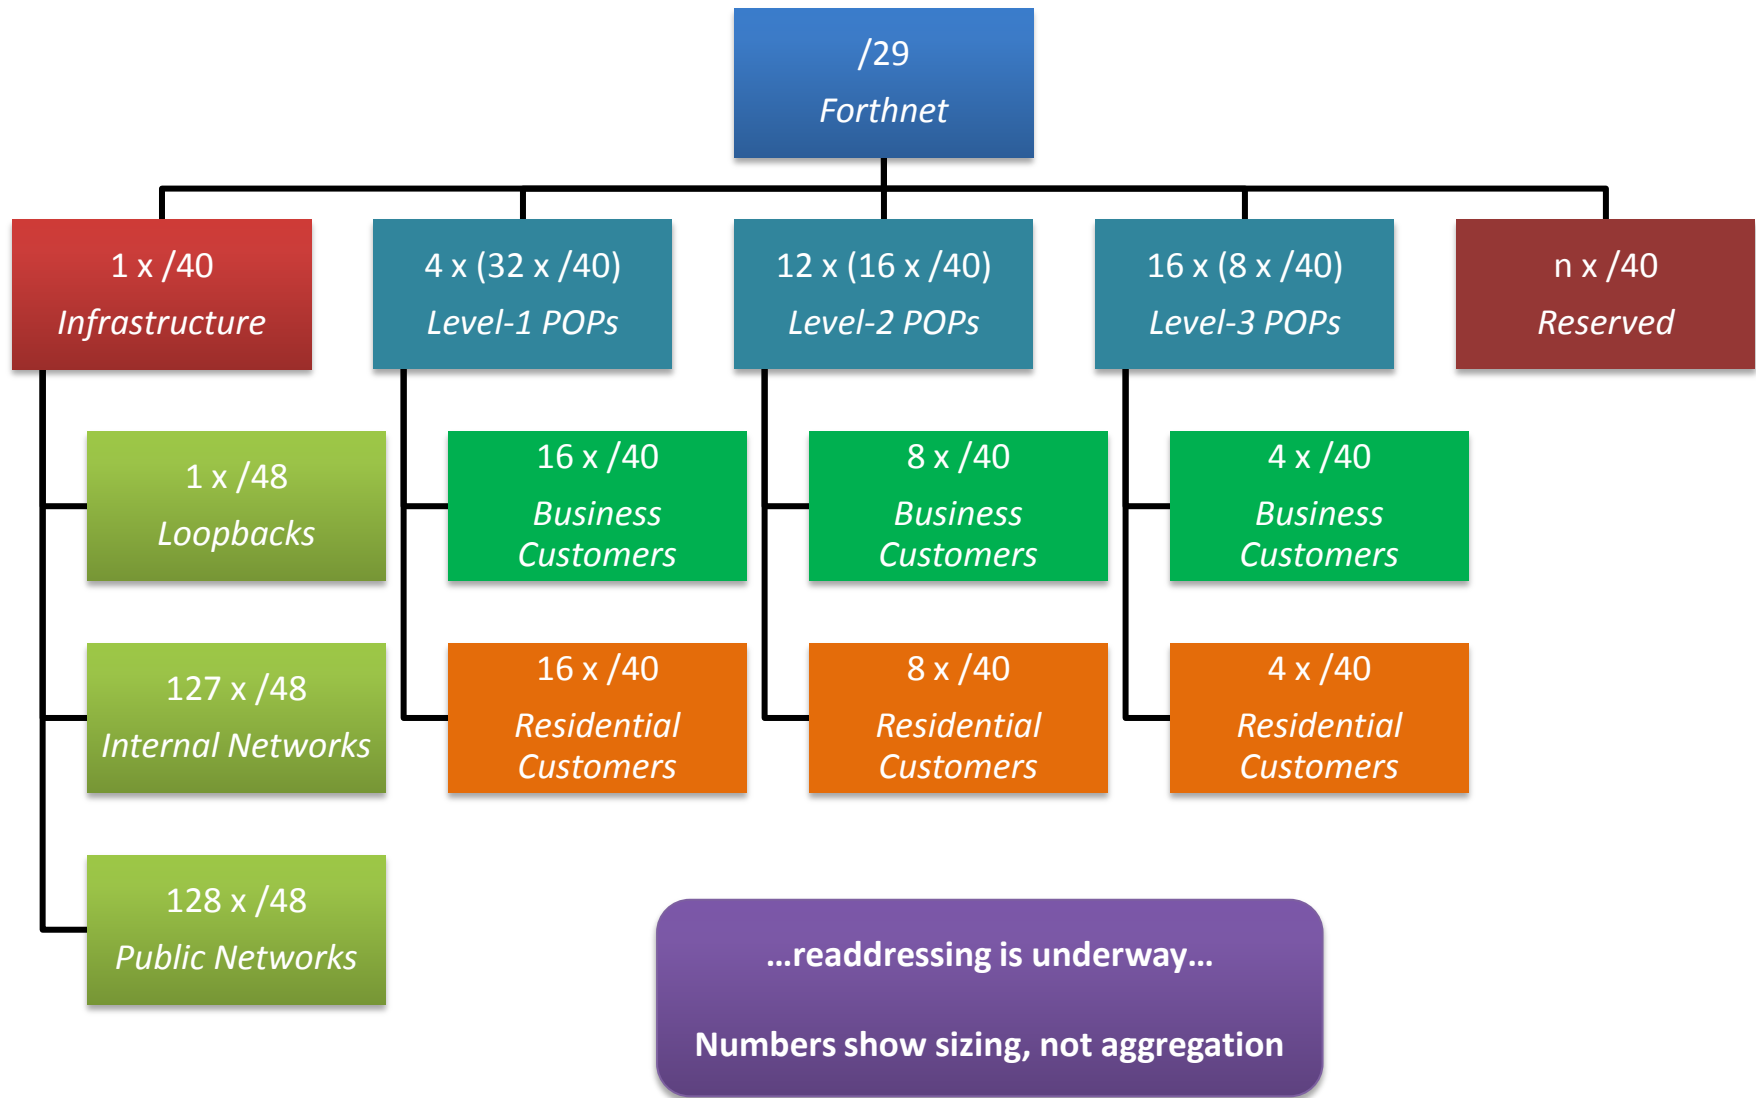


# Addressing

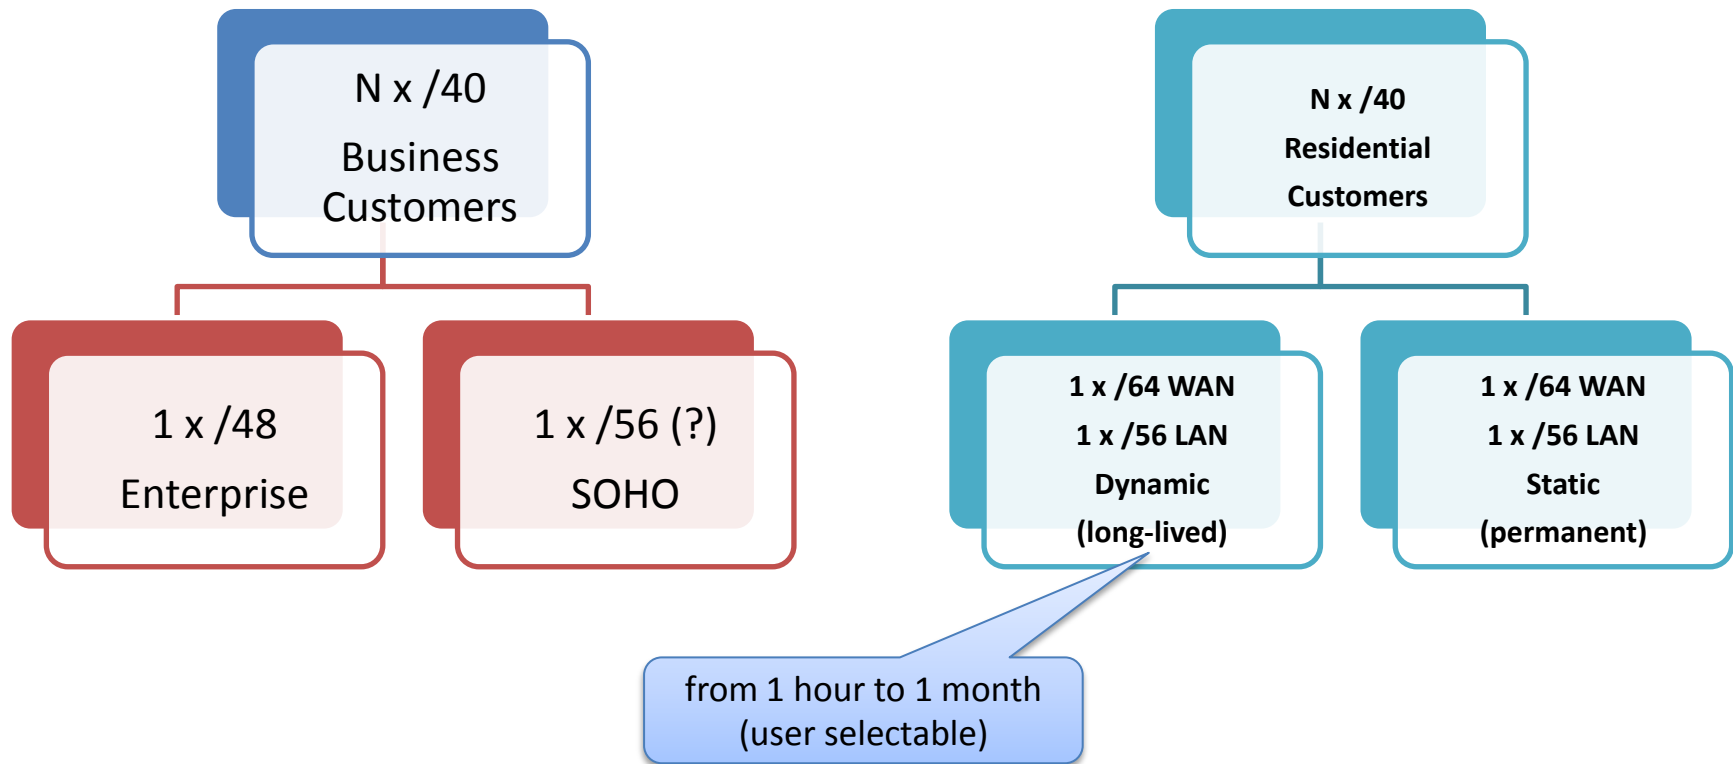


# Allocation - Assignment

# Addressing Plan



# Customer Addresses Allocation





## IPv6 Prefix Allocation/Assignment

### Issues:

- ☐ Keep the same IPv6 prefix per subscriber for as long as possible (=> *static is preferred*)
- ☐ De-aggregation of IPv6 routing table on BRAS (=> *dynamic is preferred*)

### Conflicts:

- ☐ Marketing doesn't like the static scenario
- ☐ Some subscribers do not like the static scenario (i.e. free downloading from file-hosting sites)
- ☐ Most technical people do not like the dynamic scenario

### Facts:

- ☐ According to our current BRAS clustering design, each subscriber might log in to a different bras on each try
- ☐ No DHCPv6 server available, everything must be done with current infrastructure
- ☐ At the time of authentication, radius server doesn't know if the subscriber is IPv6 enabled

**Solution:** *Give everyone what he wants...kind of.*

**Warning:** *Please do not try this at home!*

## IPv6 Prefix Allocation/Assignment

### Our solution:

- ❑ IPv6 prefixes are “stored” on radius/db using categories per BRAS id (i.e. nas-ip-address)
- ❑ A unique IPv6 prefix is reserved for X time for each subscriber, the 1<sup>st</sup> time the subscriber logs in to a BRAS
- ❑ As long as the X time hasn’t expired, the subscriber will get the same IPv6 prefix regardless of the BRAS he might log in to (during that X time)
- ❑ After the X time expires, the subscriber will get a new IPv6 prefix the next time he logs in to a BRAS (the same logic will apply to the new prefix and so on)
- ❑ Let the subscriber “define” the X time, current default is 1 week

### Technical details/issues:

- ❑ Create the appropriate procedures and tables in Sybase for assignment of IPv6 prefixes
- ❑ Create the appropriate code in Radius (Radiator) for handling the allocation/assignment
- ❑ Create a web interface to have the subscriber choose how often to change the IPv6 prefix (always, 1h, 1d, 1w, 1m, never)
- ❑ Radius response time per auth request increased from <10ms to >100ms (128 procs x 2 servers => 1 db cluster)
- ❑ After optimizing procedures and indexes in Sybase, above time fell to ~30ms
- ❑ Still not good enough, so we need to optimize/change db or try DHCPv6...or use static

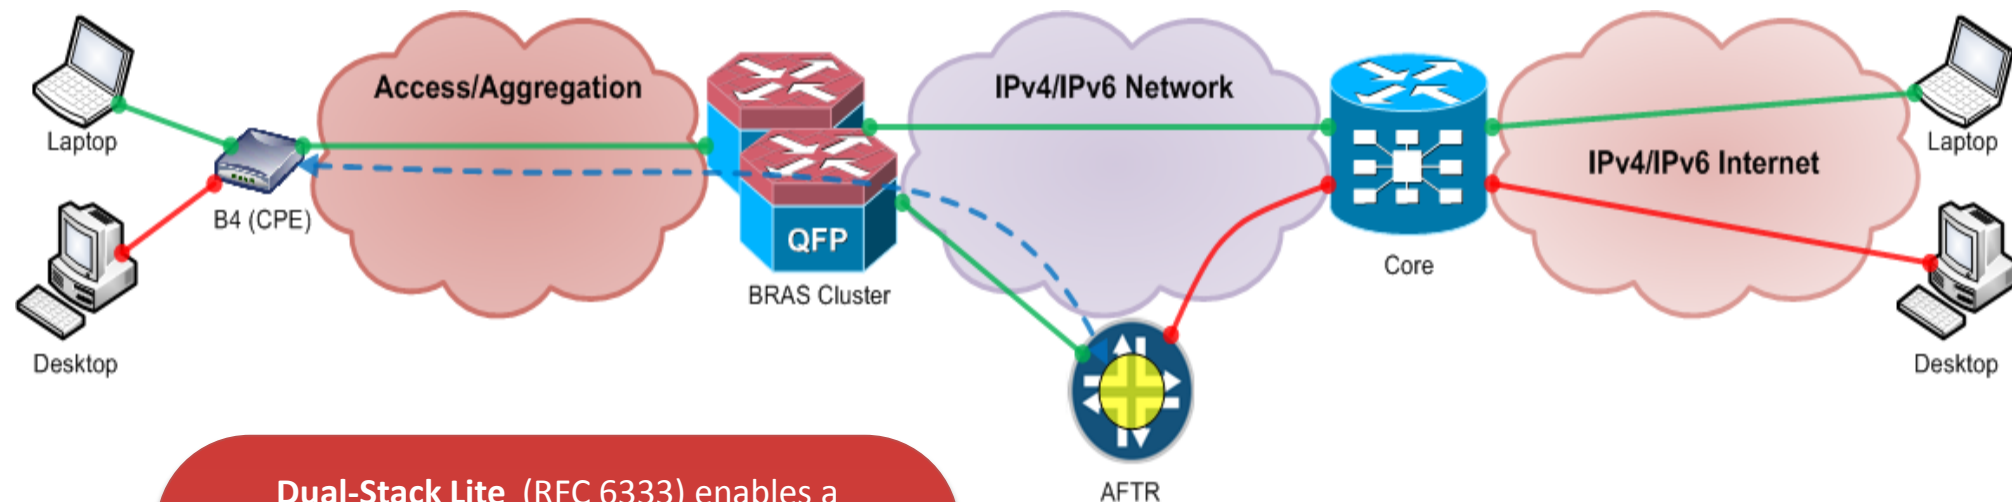
## Radius IPv6 Attributes

- ☐ WAN prefix: global /64 through SLAAC
- ☐ LAN prefix: global /56 through DHCPv6-PD
- ☐ AAA through Radius
- ☐ BRAS as DHCPv6-PD server

Authentication	Accounting
<del>Framed-IPv6-Prefix (WAN Prefix)</del>	Framed-IPv6-Prefix (WAN Prefix)
Delegated-IPv6-Prefix (LAN Prefix)	Delegated-IPv6-Prefix (LAN Prefix)
Framed-Interface-Id	Framed-Interface-Id
	IPv6-Acct-Input-Octets
	IPv6-Acct-Output-Octets

NINTENDO DS<sup>®</sup> lite





**Dual-Stack Lite** (RFC 6333) enables a broadband service provider to share IPv4 addresses among customers by combining two well-known technologies: IP in IP (IPv4-in-IPv6) and Network Address Translation (NAT)

DS-Lite was chosen because at the time we had to take a decision (Sep 2011) ,it was evaluated to be the best available solution for our IPv4 exhaustion problem.



## Tests

### **Platforms:**

- ☐ 1 AFTR (Juniper)
- ☐ 3 CPEs (Huawei, Technicolor, Gennet)

### **Issues on AFTR:**

- ☐ IPv4 MTU enforced to 1444 without any apparent reason
- ☐ Not clear whether to use physical or logical interface for tunnel termination
- ☐ No support for radius acct records with NAT mappings
- ☐ No support for deterministic Port Block Allocation (NAT) with DS-Lite
- ☐ Port usage sometimes shows unrealistic numbers
- ☐ Using port block allocation sometimes breaks NAT

### **Issues on CPEs:**

- ☐ Throughput decreased by 20%-40%
- ☐ MTU/MSS havoc (every vendor is following a unique approach)
- ☐ Default route expires due to RAs not sent
- ☐ Status/diagnostic pages miss a lot of information
- ☐ IPv6 firewall setup is not user friendly
- ☐ IPv6 access on the WAN interface cannot be blocked
- ☐ DS-Lite and IPv4 can be used simultaneously (btw, this works)

## Results

### **Works:**

- ☐ Anything http based
- ☐ Video/Audio streaming from various sites
- ☐ FTP
- ☐ Skype
- ☐ Torrent leeching
- ☐ SIP calls initiated from inside

### **Semi/No-Works:**

- ☐ Anything expecting incoming connections without first initiating them (servers)
- ☐ Torrent seeding (after a while the leecher learned the external ip/port of the seeder)
- ☐ SIP calls without an intermediate server
- ☐ Online gaming

## Softwires & Flows

```

achatz@mx-lab-kln-01> show services softwire flows ds-lite
Interface: sp-2/0/0, Service set: DSLITE-SERVICE-SET
Flow                               State      Dir      Frm count
DS-LITE2a02:2148:100:304::1234->2a02:2148:77:77:2::5 Forward  I      75819
TCP      192.168.1.69:53428 ->  92.252.240.8:51413 Forward  I      58
    NAT source  192.168.1.69:53428 ->  194.219.240.1:1078
    Softwire   2a02:2148:100:304::1234 -> 2a02:2148:77:77:2::5
UDP      122.160.26.199:24853 -> 194.219.240.1:1816 Forward  O        1
    NAT dest    194.219.240.1:1816 -> 192.168.1.69:42338
    Softwire   2a02:2148:77:77:2::5 -> 2a02:2148:100:304::1234
UDP      192.168.1.69:42338 ->109.195.156.143:35691 Forward  I        9
    NAT source  192.168.1.69:42338 -> 194.219.240.1:1290
    Softwire   2a02:2148:100:304::1234 -> 2a02:2148:77:77:2::5
...
UDP      178.122.93.189:38121 -> 194.219.240.1:1149 Forward  O       31
    NAT dest    194.219.240.1:1149 -> 192.168.1.69:42338
    Softwire   2a02:2148:77:77:2::5 -> 2a02:2148:100:304::1234
TCP      98.28.107.204:59420 -> 194.219.240.1:1347 Forward  O       83
    NAT dest    194.219.240.1:1347 -> 192.168.1.69:53460
    Softwire   2a02:2148:77:77:2::5 -> 2a02:2148:100:304::1234
TCP      94.19.140.192:58750 -> 194.219.240.1:1490 Forward  O      406
    NAT dest    194.219.240.1:1490 -> 192.168.1.69:53287
    Softwire   2a02:2148:77:77:2::5 -> 2a02:2148:100:304::1234
TCP      90.45.16.240:38317 -> 194.219.240.1:1671 Forward  O      105

```

## Translations & Ports

```
achatz@mx-lab-kln-01> show services nat pool detail DSLITE-POOL
Interface: sp-2/0/0, Service set: DSLITE-SERVICE-SET
NAT pool: DSLITE-POOL, Translation type: dynamic
Address range: 194.219.240.1-194.219.240.255
Port range: 1024-65535, Ports in use: 1089, Out of port errors: 0, Max ports used: 1333
```

```
achatz@mx-lab-kln-01> show services nat pool detail DSLITE-POOL
Interface: sp-2/0/0, Service set: DSLITE-SERVICE-SET
NAT pool: DSLITE-POOL, Translation type: dynamic
Address range: 194.219.240.1-194.219.240.255
Port range: 1024-65535, Ports in use: 54365, Out of port errors: 0, Max ports used: 54537
```

## IPv4 Address Sharing Issues

- ☐ Restricted allocations of outgoing ports will impact performance for end-users
- ☐ Incoming port negotiation mechanisms may fail
- ☐ Port discovery mechanisms will not work
- ☐ Inbound ICMP will fail in many cases
- ☐ Amplification of security issues will occur
- ☐ Service usage monitoring and abuse logging will be impacted
- ☐ Spam blacklisting will be affected
- ☐ Geo-location services will be impacted
- ☐ Authentication mechanisms may be impacted
- ☐ and many more...see **RFC 6269**
  
- ☐ Since our TR-069 platform can identify subscribers that do not use port forwarding (and UPnP) on their CPE, these subscribers can be candidates for address sharing. In any case, an option should be provided for disabling this after subscriber's request.
- ☐ Whether **PCP** is a viable solution will be proved in the following months
- ☐ Latest PCP draft (27) doesn't (and won't) support multiple ports in a single request/response, so we expect some issues with incoming connections and maybe suboptimal ip/port usage
- ☐ After running a "scan" on our network, we have chosen 1:6 as the initial ratio



Thank You!

Questions 

**We accept donations of any size of IPv4 addresses**