


Mapping Address + Port

Alain Fiocco, Sr Director IPv6 HIP

Andrew Yourtchenko, Technical Leader, NOSTG

June 4, 2013

Cisco Knowledge Network: TechAdvantage Webinar



IPv4 & IPv6 Statistics

RIR v4 IPs Left

AfrNIC	62,583,986
APNIC	14,538,451
ARIN	38,859,068
LACNIC	42,312,892
RIPE	15,133,694

v6 ASNs
15% (7,053/44,375)

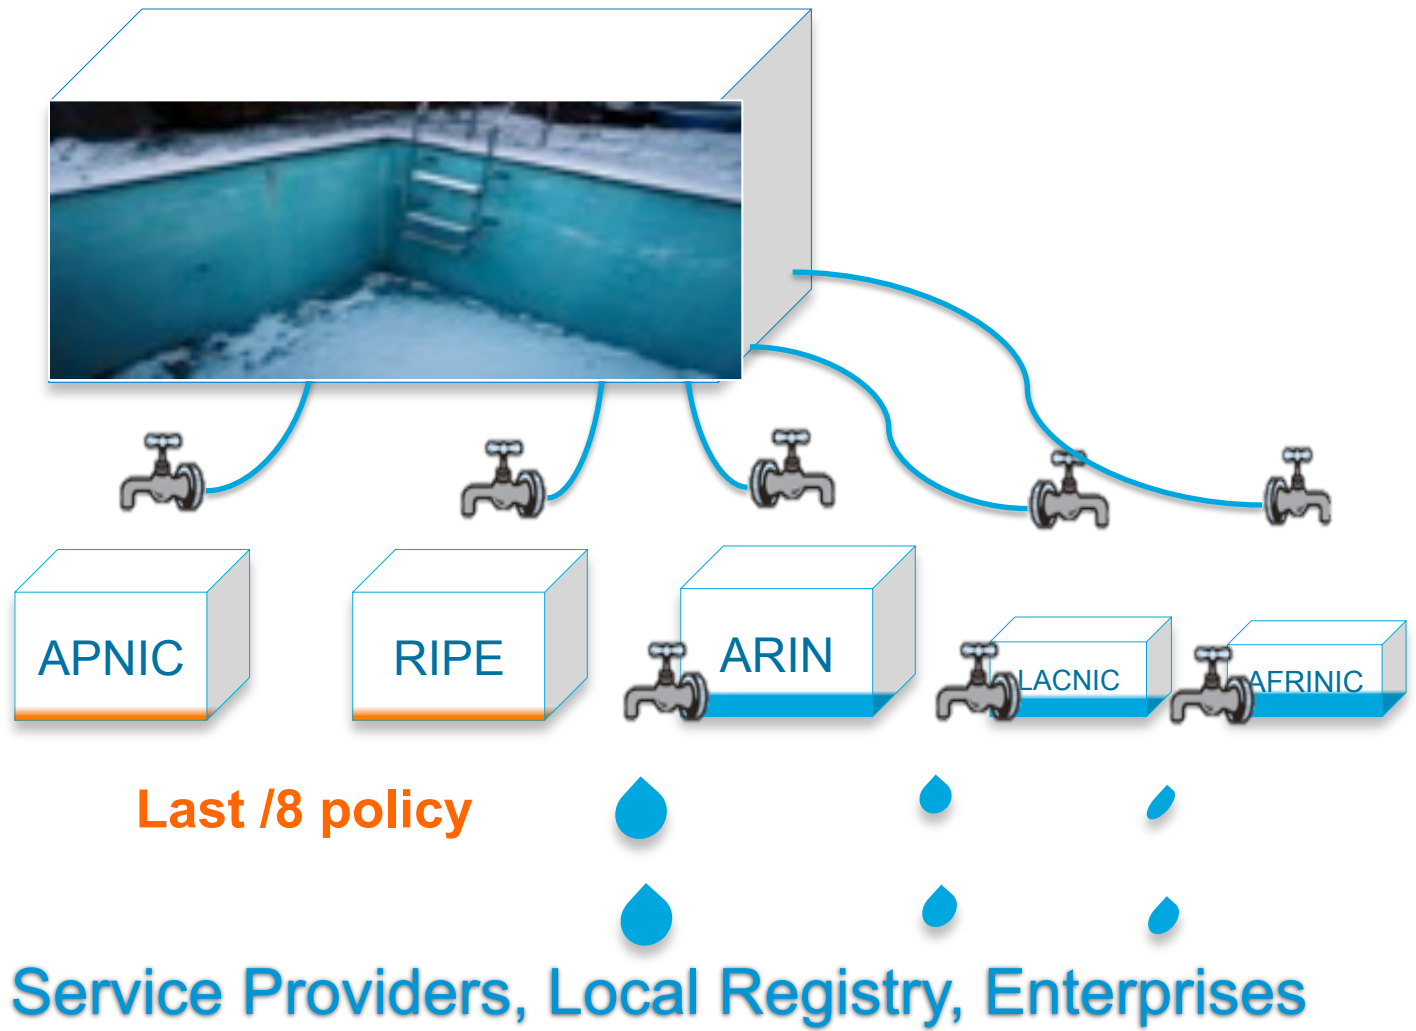
v6 Ready TLDs
87% (276/317)

v6 Glues
14,563

v6 Domains
4,754,437 ↑

0
days remaining
IANA exhausted

HURRICANE ELECTRIC
INTERNET SERVICES



<http://ipv6.he.net/statistics/>

IPv6 Network Transition options

1-Enable Core

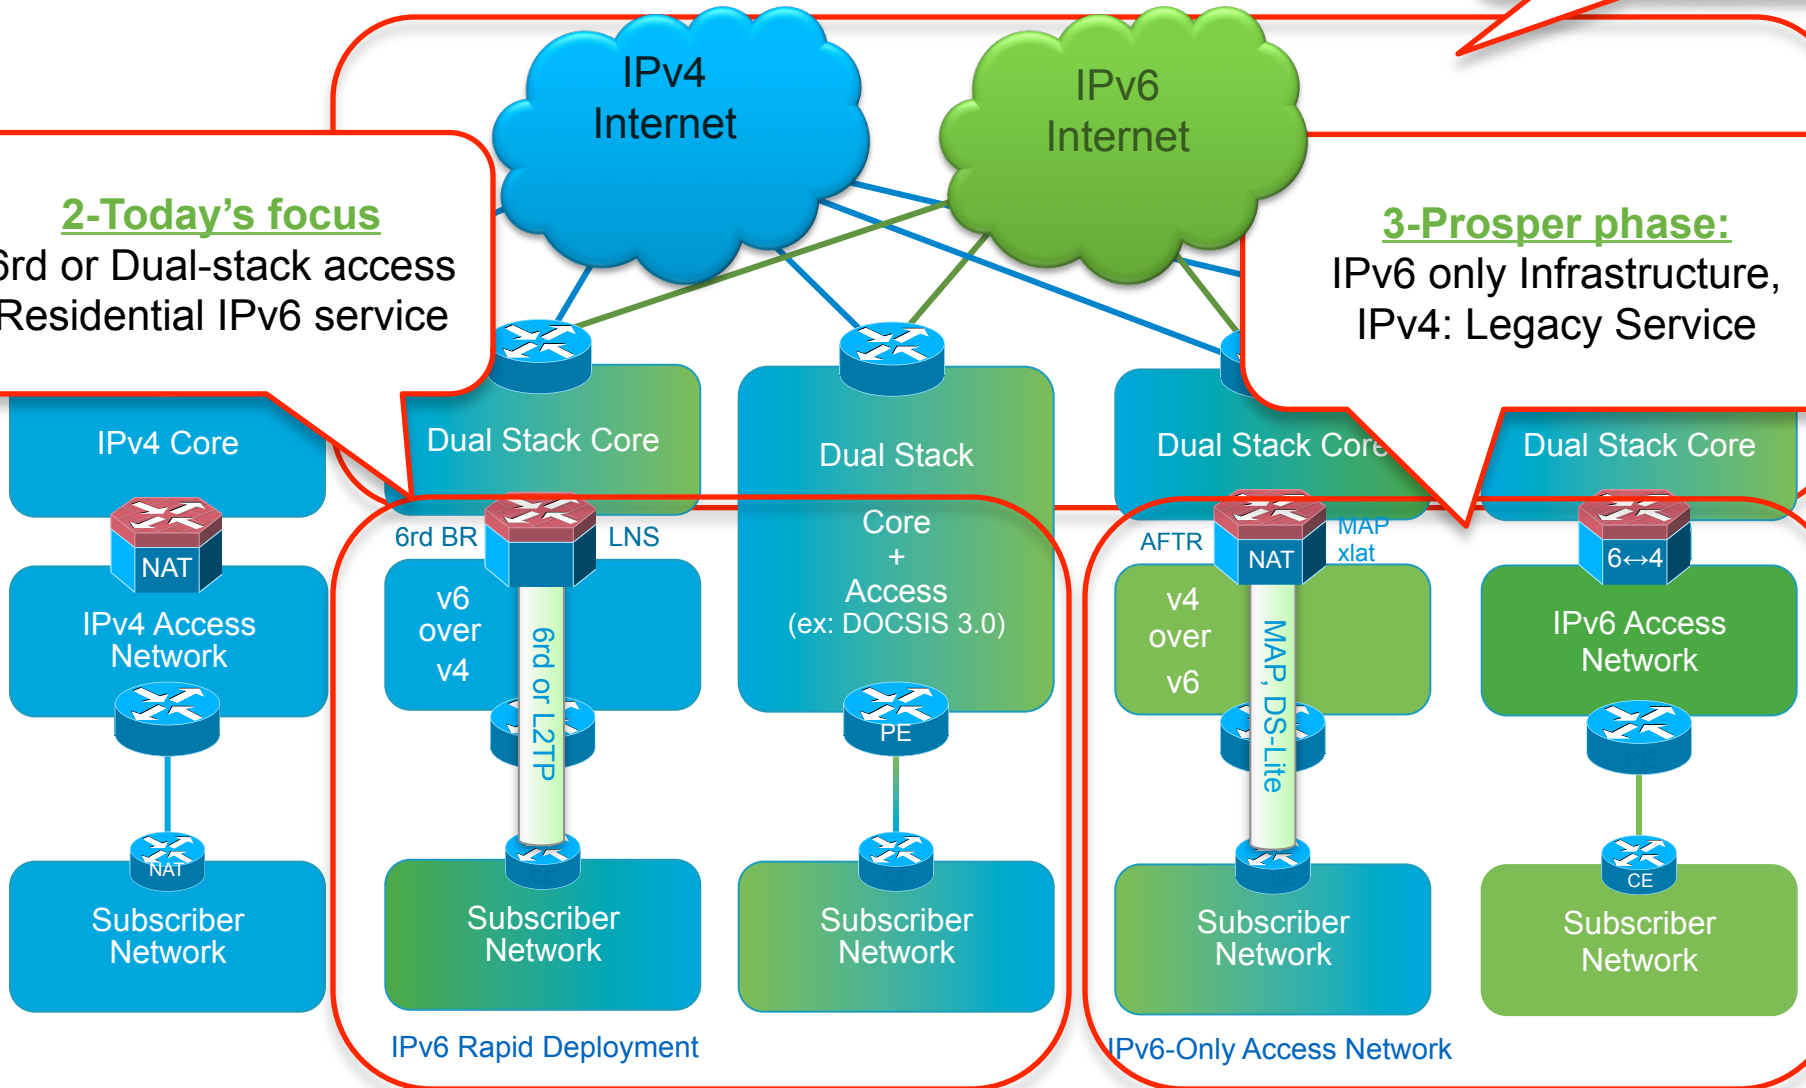
- Dual-Stack core
- MPLS/6(v)PE

2-Today's focus

6rd or Dual-stack access
Residential IPv6 service

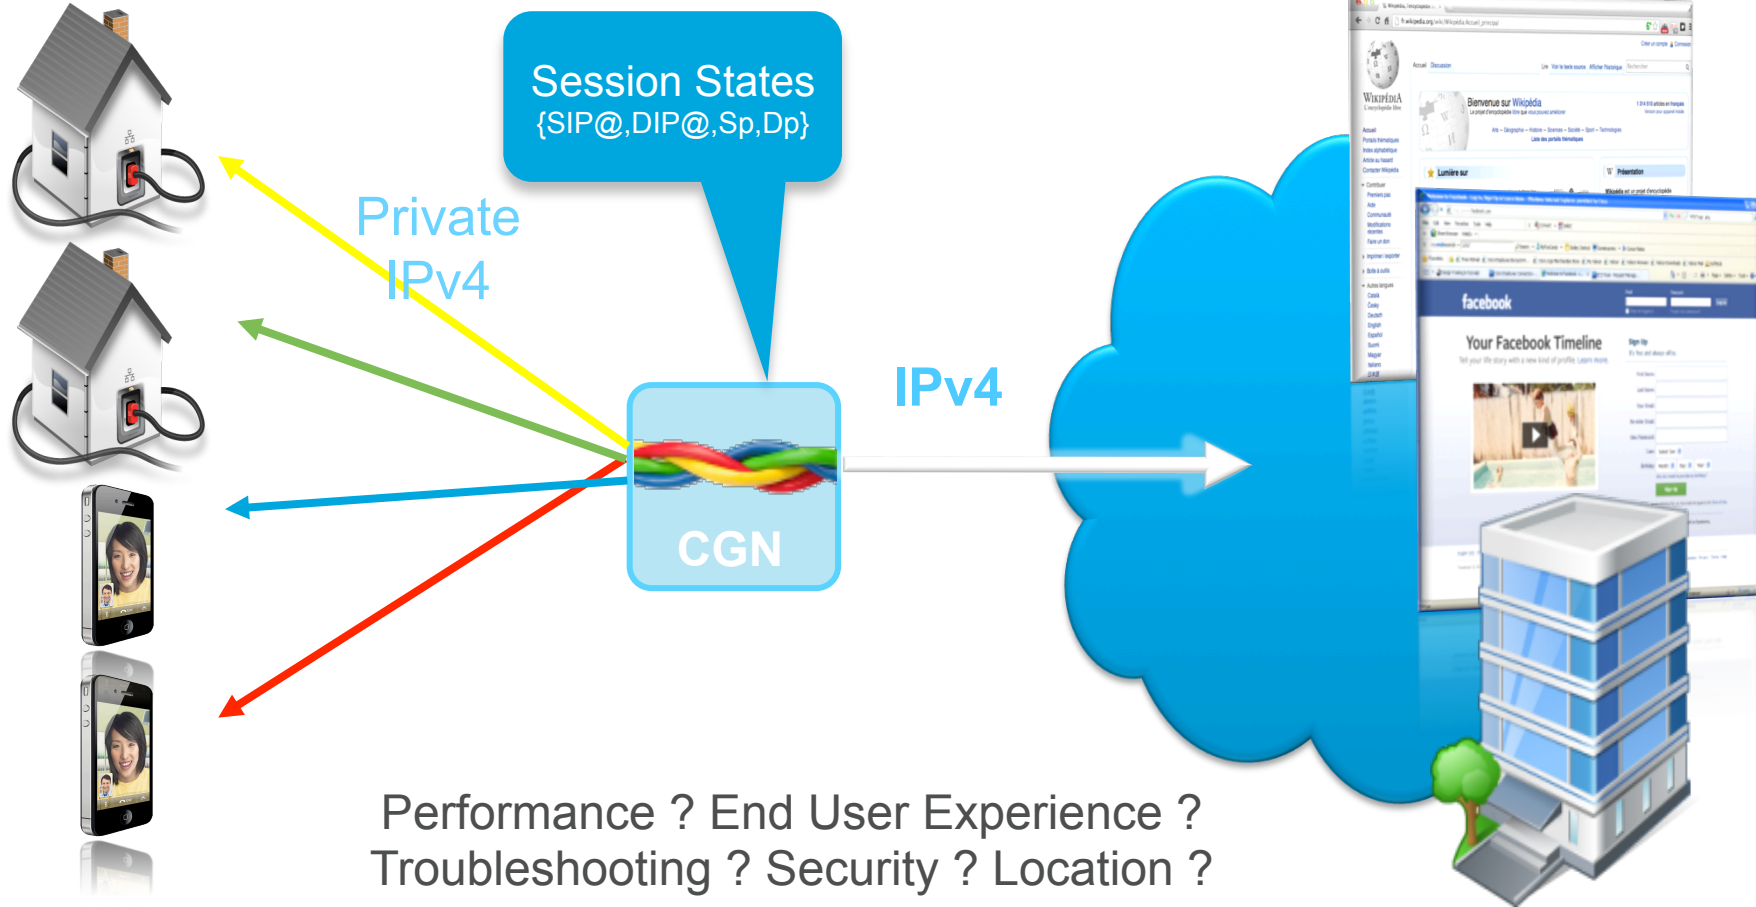
3-Prosper phase:

IPv6 only Infrastructure,
IPv4: Legacy Service



Sharing public IPv4 addresses

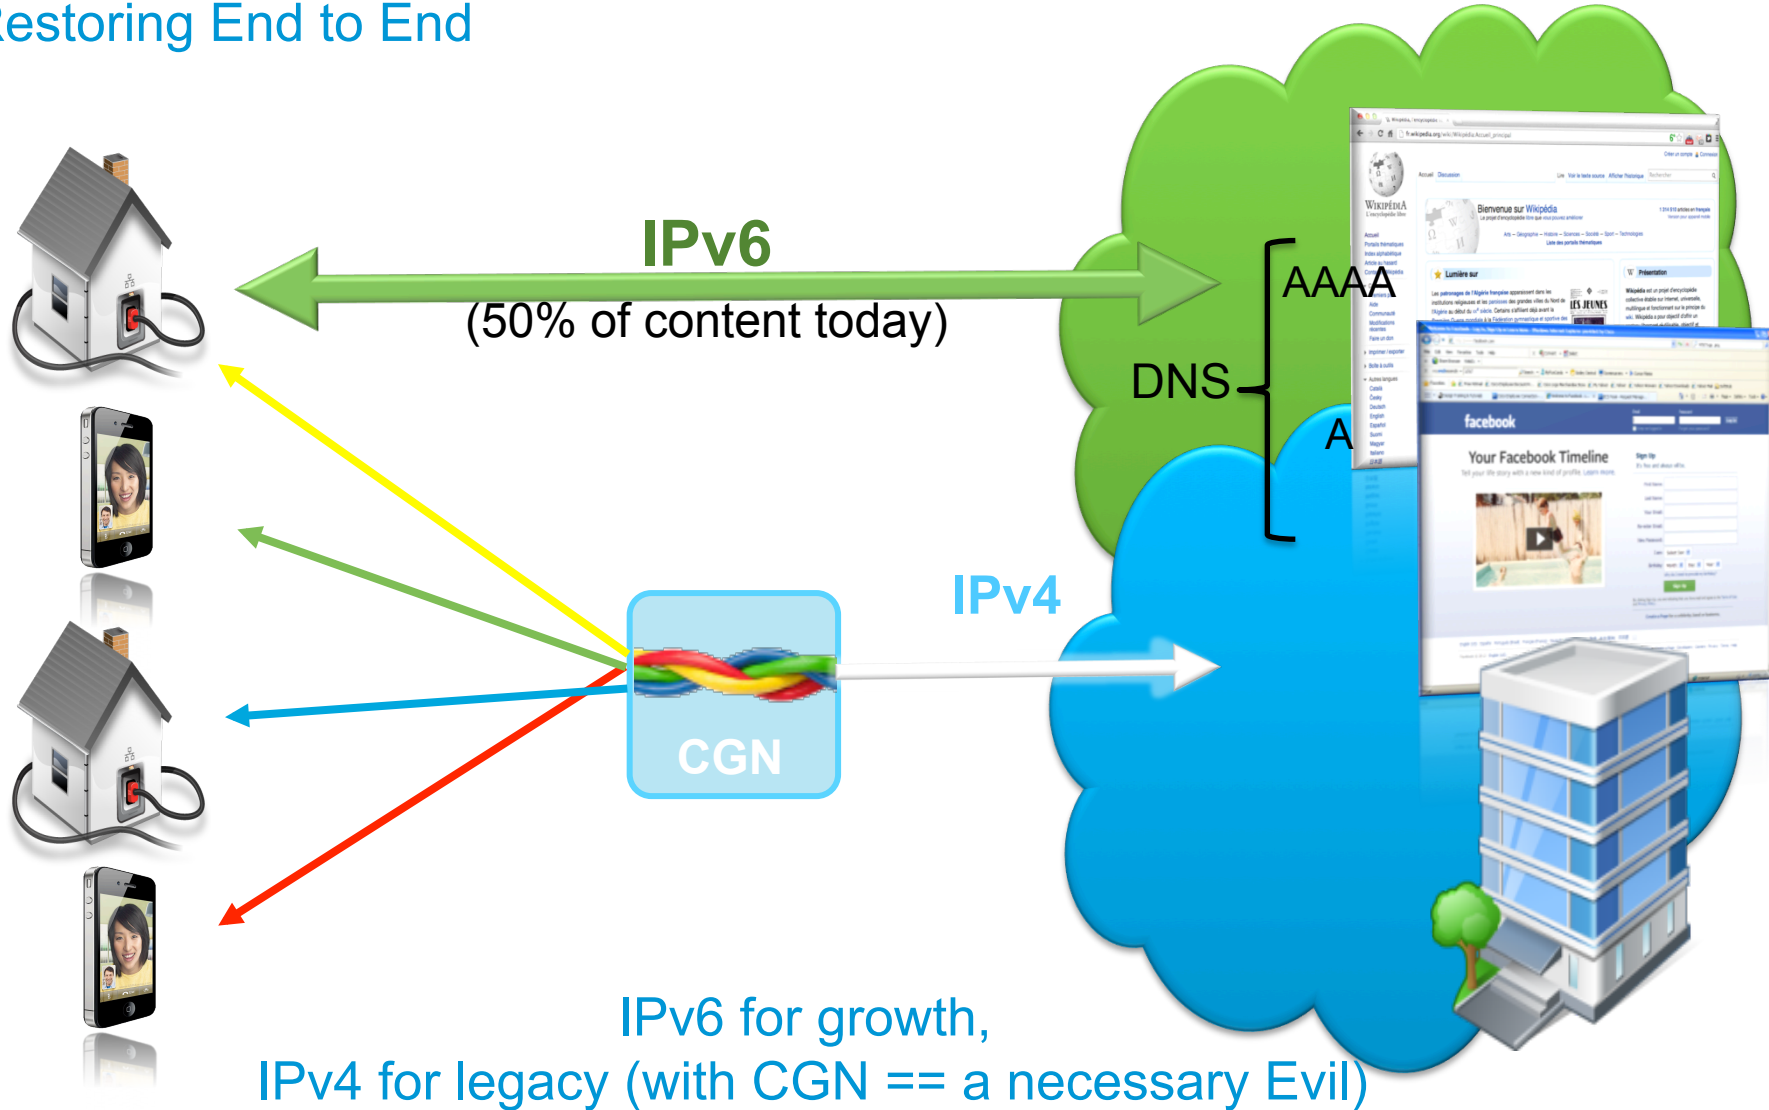
Makes the Internet Statefull ! ... Really ?



Performance ? End User Experience ?
Troubleshooting ? Security ? Location ?
User Logging/data retention vs too much information ?

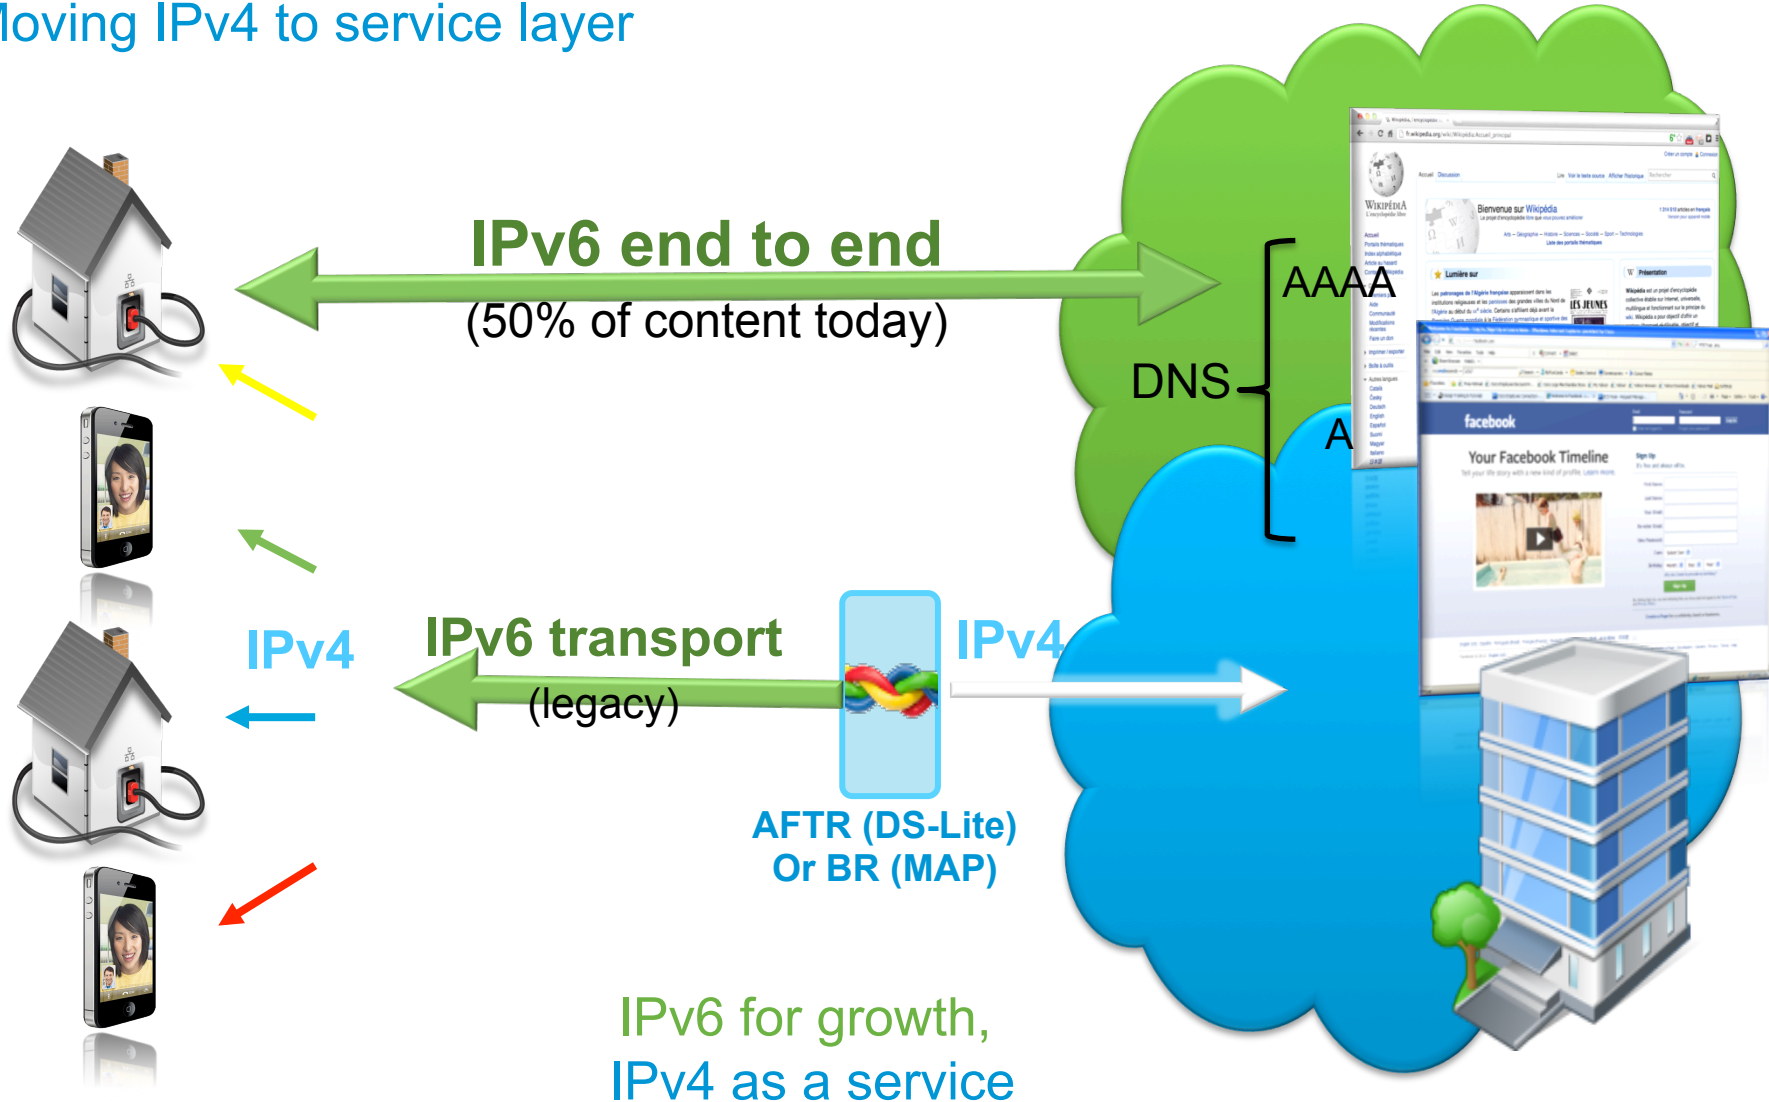
IPv6 – “Full Spectrum” Internet

Restoring End to End




IPv6 only network

Moving IPv4 to service layer

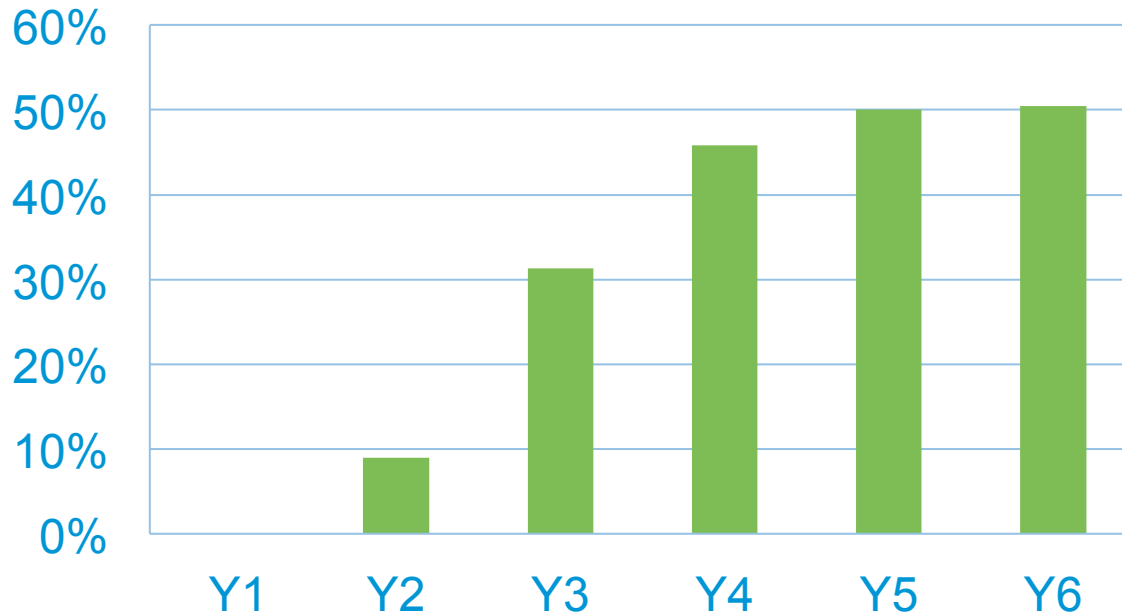


IPv6 only network, is the user ready ?



- IPv6 only access network
- Growing share of IPv6 reachable natively
Continue to promote “end to end” IPv6
- IPv4 legacy content ?
Option 0 => continue to promote IPv6 to WEB site
Option 1: NAT64/DNS64 (exclude some apps)
Option 2: IPv4 to end-user + share IPv4@
Fixed: MAP (Stateless) or DS-Lite (Statefull)
Mobile: XLAT464
- Professional network (ex: Emergency Response)
Specialized Users devices and mobile networks
IPv6 only environment


MAP Benefits: TCO Savings vs DS-Lite




Based on ASR9K+ISM
 - MAP is switched in LC
 - DS-Lite require state/ISM

	Y1	Y2	Y3	Y4	Y5	Y6
Connections (Thousands)	0	500000	2,500,000	5,000,000	8,000,000	9,500,000
Peak BW/Subs (Mbits/sec)	0.15	0.22	0.3	0.4	0.5	0.6
IPv6 Content	30%	50%	70%	80%	90%	95%
IPv6 Home Devices	80%	85%	88%	90%	91%	92%
Avg Sessions/household	200	220	242	266	293	322

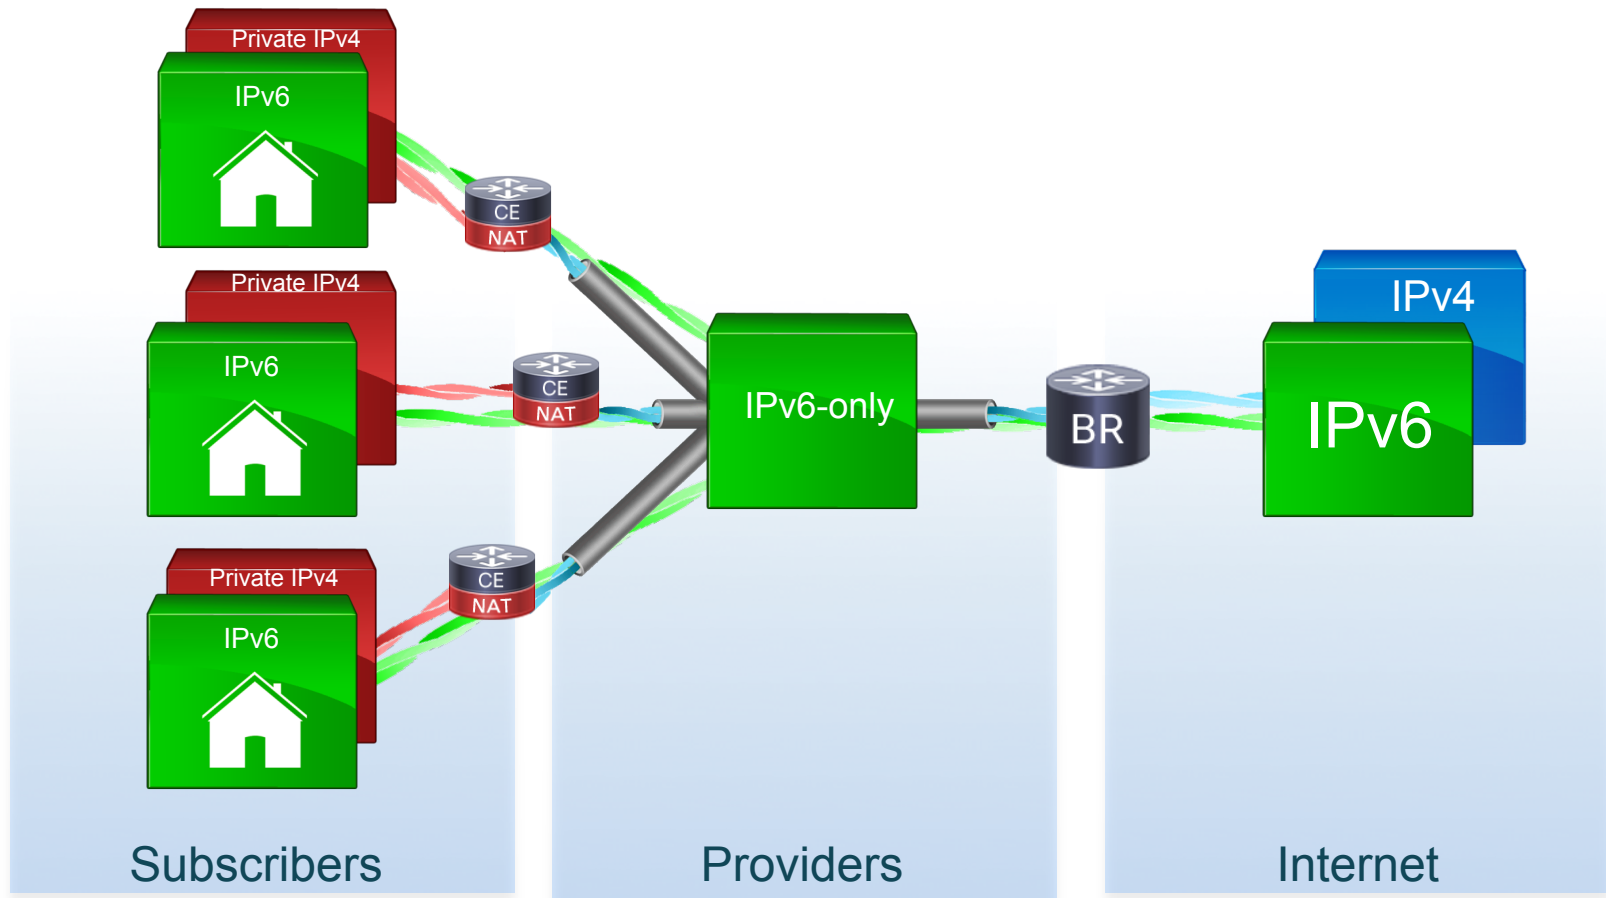
Dual Stack Lite (DS-Lite)




“Lightweight 4 over 6” (also “Public 4over6”)



Mapping Address + Port (MAP)



Imagine the Internet without any IP address aggregation



IPv4 CIDR Chart **RIPE NCC**

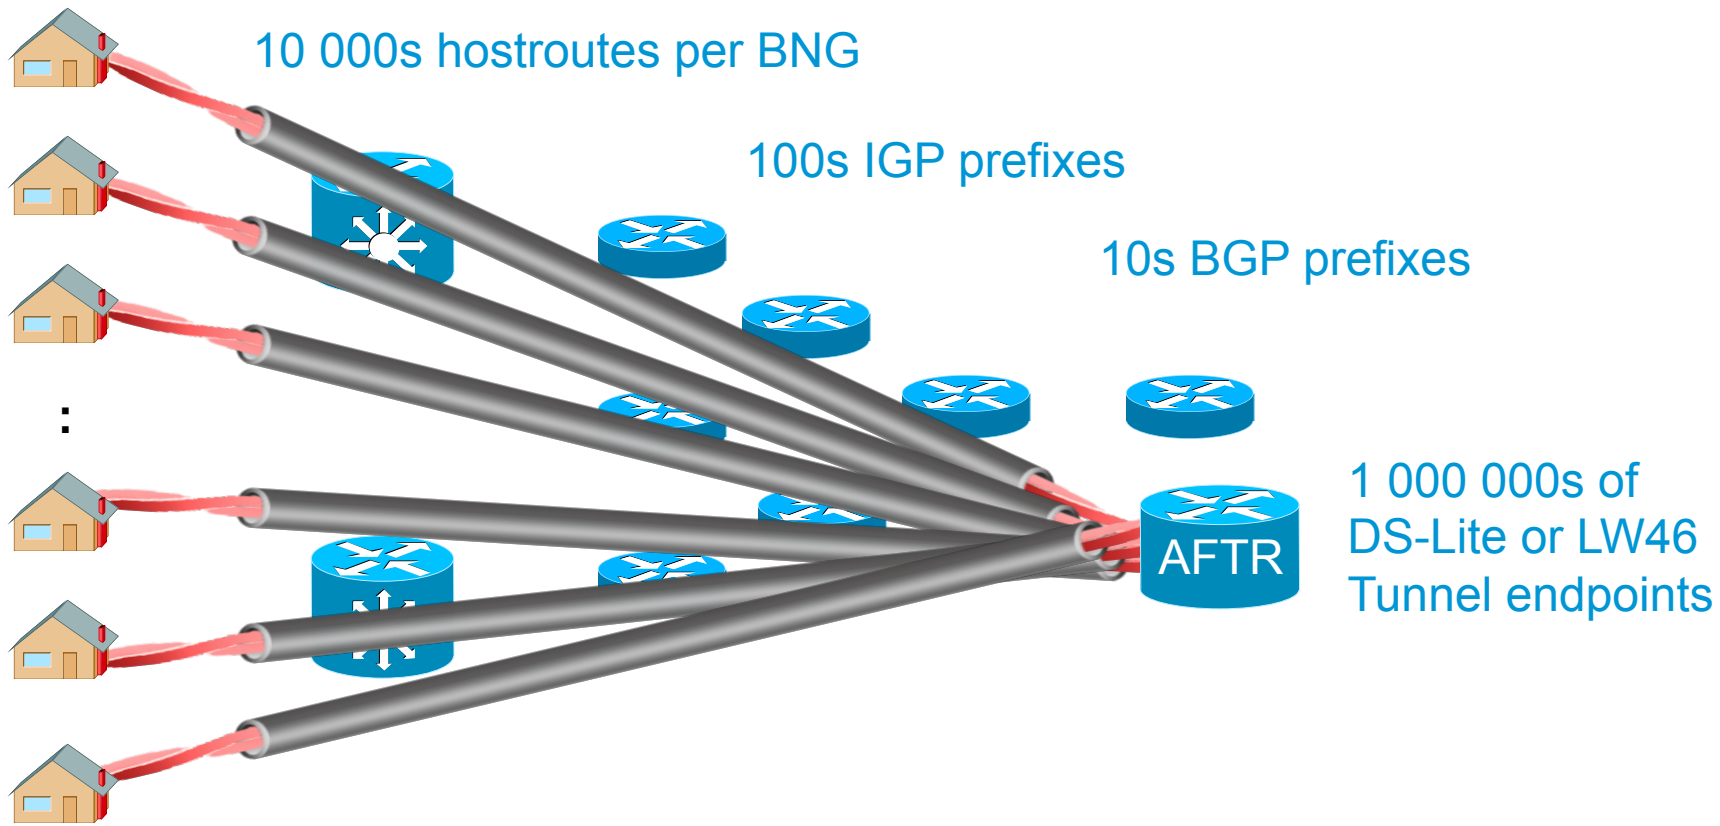
IP Addresses	Bits	Prefix	Subnet Mask
1	0	/32	255.255.255.255
2	1	/31	255.255.255.254
4	2	/30	255.255.255.252
8	3	/29	255.255.255.248
16	4	/28	255.255.255.240
32	5	/27	255.255.255.224
64	6	/26	255.255.255.192
128	7	/25	255.255.255.128
256	8	/24	255.255.255.0
512	9	/23	255.255.252.0
1 K	10	/22	255.255.248.0
2 K	11	/21	255.255.240.0
4 K	12	/20	255.255.240.0
8 K	13	/19	255.255.224.0
16 K	14	/18	255.255.192.0
32 K	15	/17	255.255.128.0
64 K	16	/16	255.255.0.0
128 K	17	/15	255.252.0.0
256 K	18	/14	255.252.0.0
512 K	19	/13	255.248.0.0
1 M	20	/12	255.240.0.0
2 M	21	/11	255.224.0.0
4 M	22	/10	255.192.0.0
8 M	23	/9	255.128.0.0
16 M	24	/8	255.0.0.0
32 M	25	/7	254.0.0.0
64 M	26	/6	252.0.0.0
128 M	27	/5	248.0.0.0
256 M	28	/4	240.0.0.0
512 M	29	/3	224.0.0.0
1024 M	30	/2	192.0.0.0
2048 M	31	/1	128.0.0.0
4096 M	32	/0	0.0.0.0

K = 1,024 • M = 1,048,576

Contact Registration Services:
hostmaster@ripe.net • lir-help@ripe.net www.ripe.net

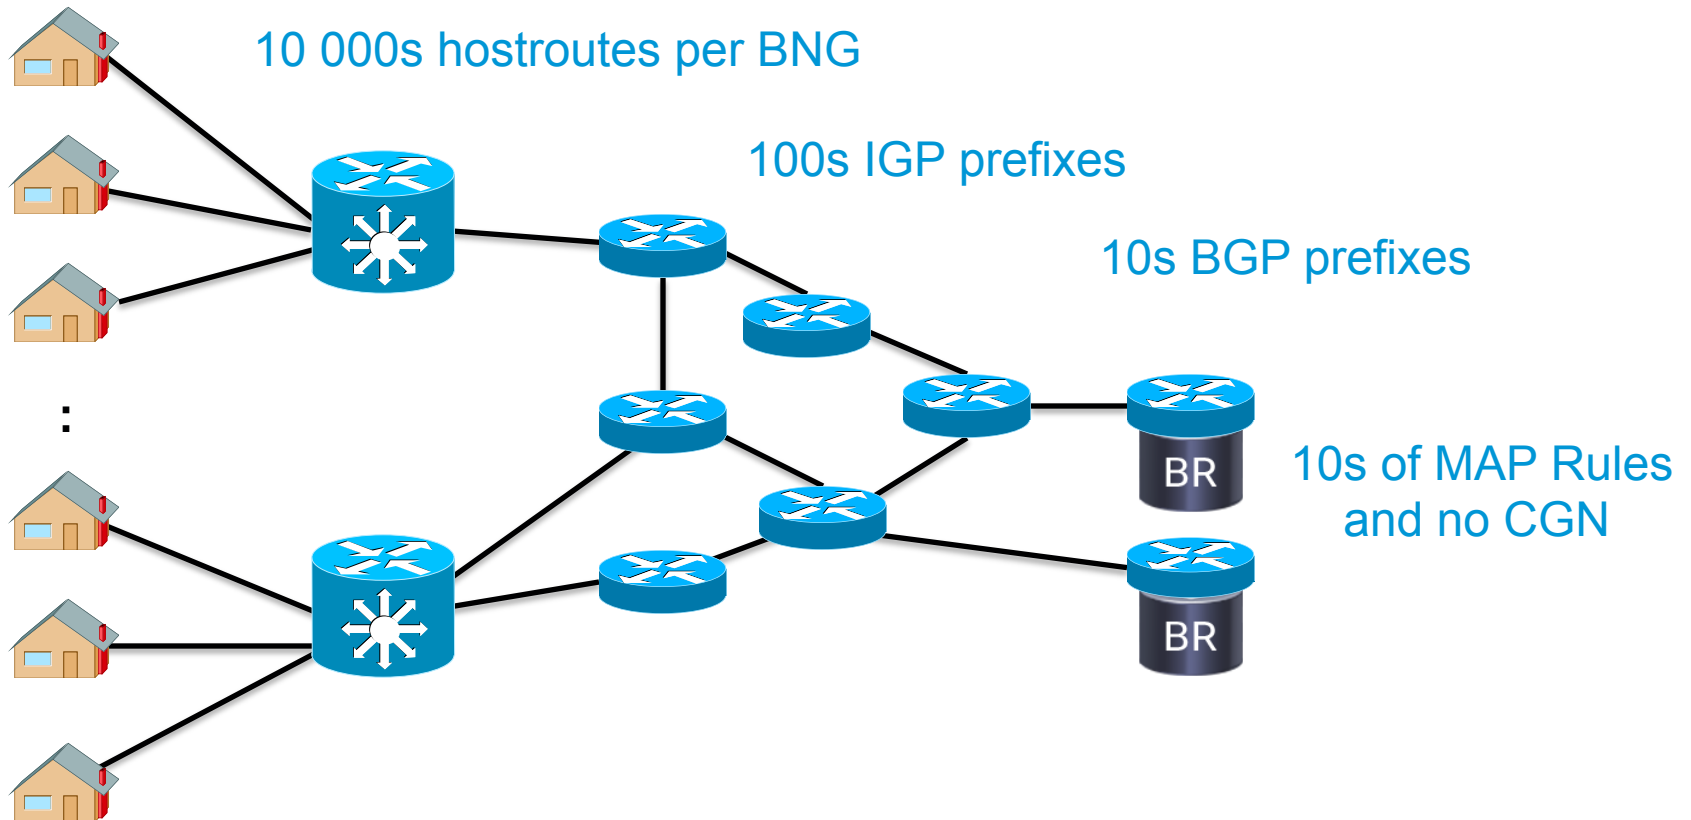
DS-Lite/LW46/Public 4over6 – Per-subscriber tunnels

1 000 000s of subscribers



MAP Exploits Aggregation in IPv6 Routing

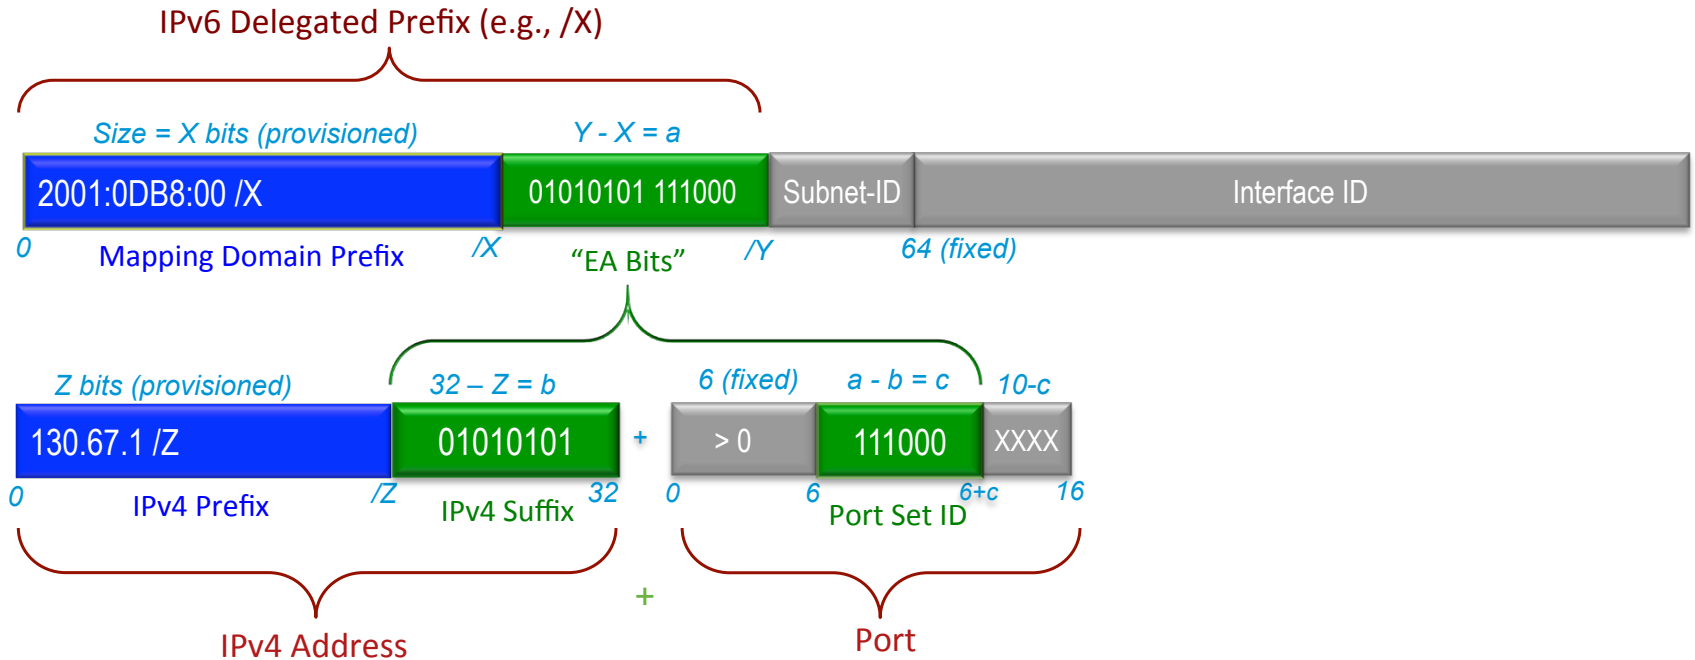
1 000 000s of subscribers



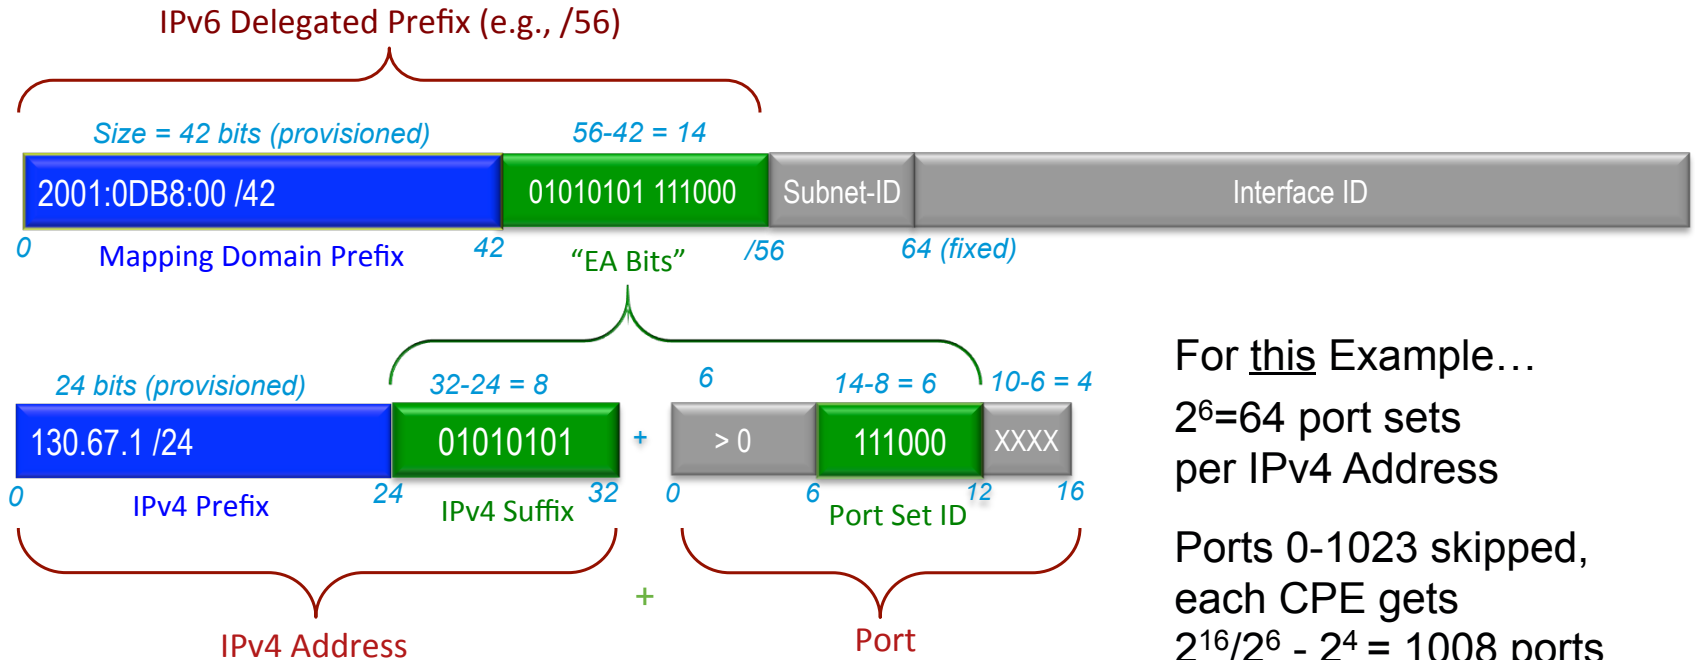
MAP: Easy as 1-2-3

- ① IPv6 to IPv4+Port Mapping
- ② Stateless Border Relay
- ③ Packet Flow and Forwarding

① IPv6 → IPv4 + Port Mapping



① IPv6 → IPv4 + Port Mapping



For this Example...

$2^6=64$ port sets
per IPv4 Address

Ports 0-1023 skipped,
each CPE gets
 $2^{16}/2^6 - 2^4 = 1008$ ports

One IPv4 /24 serves
 $2^{(6+8)} \approx 16,384$ (vs. ≈ 256)
subscribers



MAP Simulation Tool (beta)

[Video tutorial](#)
[Highly editable elements](#)
[Add a new MAP rule](#)
[Remove all MAP rules](#)
[Load rules from text](#)
[Save rules to text](#)
[Create a link to these rules](#)

Paste previously saved set of rules here.

Rule 0

[Delete](#)
[Advanced](#)
[Example](#)
/56

IPv6

2001:db8:9500:0

/40

EA Bits
(16 = 8 + 8)
Subnet
(8)

Interface ID (64)

IPv4 : Port

198.51.100.0 /24

Suffix
(8)

:

(4)

PSID
(8)

(4)

256 IPv4 addresses, 65536 users, 240 ports each (1:256)

In order to help us understand how this tool is being used and to improve it in the future, it will periodically save anonymous usage information for analysis. This does NOT include your IP address or any other information not needed by the tool itself. If you wish, you may override this by unchecking the box below.

 Data collection is currently on.

MAP Simulation tool created by [Arthur Lacoste](#) of Cisco Systems based on [this IETF draft](#).

A [quick video tutorial](#) for this tool is available on youtube.

Please send comments, bug reports, and other feedback to : [map46-tool-feedback\[at\]external.cisco.com](mailto:map46-tool-feedback[at]external.cisco.com)

Last updated: 6/19/2012

http://6lab.cisco.com/map



Search

- SHOP
- MY MUSIC
- MY BOOKS
- MY MAGAZINES
- MY MOVIES & TV
- MY ANDROID APPS





MAP Calculator
Cisco IPv6 Project



★★★★★ (2)

INSTALL

Users who viewed this also viewed

- 
IP calculator (IPv4 and IPv6)
 NOIR.DE
 ★★★★★ (97)
 Free
- 
IP calculator (Samsung)
 NOIR.DE
 ★★★★★ (20)
 Free
- 
IP Network Calculator
 ORBITINGPLUTO
 ★★★★★ (455)
 Free
- 
Shooter
 SEANKNDY
 ★★★★★ (1,046)
 \$9.99

Users who installed this also installed

- 
IPv4 Subnet Calculator
 NULLSETZ
 ★★★★★ (33)
 Free

- OVERVIEW
- USER REVIEWS
- WHAT'S NEW
- PERMISSIONS

Description

MAP Calculator is a tool designed to help you deal with MAP (Mapping Address and Port), IETF's last IPv4 - IPv6 stateless tunneling solution. The spec for MAP is being edited by the Softwire Working Group and is available here : <http://datatracker.ietf.org/doc/draft-ietf-softwire-map/>

MAP Calculator features :

- Immediately understand how IPv6 and IPv4+Port are mapped thanks to an intuitive graphical representation
- Get a feeling of the trade-offs through a clear display of core information (addresses used, customers supported, service level)

[Visit Developer's Website](#) [Email Developer](#)

MORE

 +1 1
 Tweet

ABOUT THIS APP

RATING:
★★★★★ (2)

UPDATED:
June 20, 2012

CURRENT VERSION:
1.00

REQUIRES ANDROID:
4.0 and up

CATEGORY:
Tools


INSTALLS:
50 - 100

SIZE:
91k

PRICE:
Free


CONTENT RATING:
Everyone

App Screenshots



User Reviews [Write a Review >](#)

No fans or critics yet? Be the first!



Average rating:

5.0

★★★★★
2

iTunes Preview

What's New What is iTunes What's on iTunes iTunes Charts How To

Cisco MAP Calculator

[View More By This Developer](#)

By Cisco

Open iTunes to buy and download apps.



[View In iTunes](#)

+ This app is designed for both iPhone and iPad

Free

Category: [Business](#)
Released: Sep 17, 2012
Version: 1.0
Size: 0.6 MB
Language: English
Seller: Cisco
© 2012 Cisco Systems, Inc.
All Rights Reserved.
[Rated 4+](#)

Requirements: Compatible with iPhone 3GS, iPhone 4, iPhone 4S, iPhone 5, iPod touch (3rd generation), iPod touch (4th generation), iPod touch (5th generation) and iPad. Requires iOS 5.1 or later.

Customer Ratings

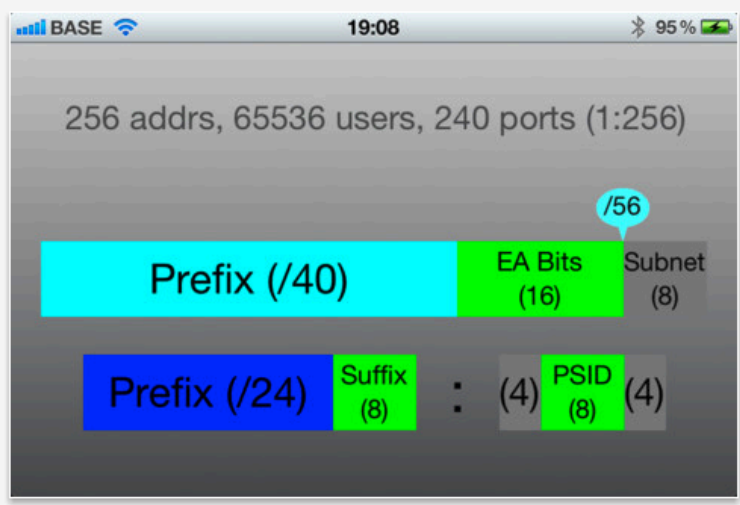
Description

MAP Calculator allows you to quickly visualize different layouts of the address space partitioning for the MAP port mapping algorithm

[Cisco MAP Calculator Support](#)

Screenshots

iPhone | iPad




Customer Reviews

haji ★★★★★
by Haji Danger
good work

② Stateless Border Relays

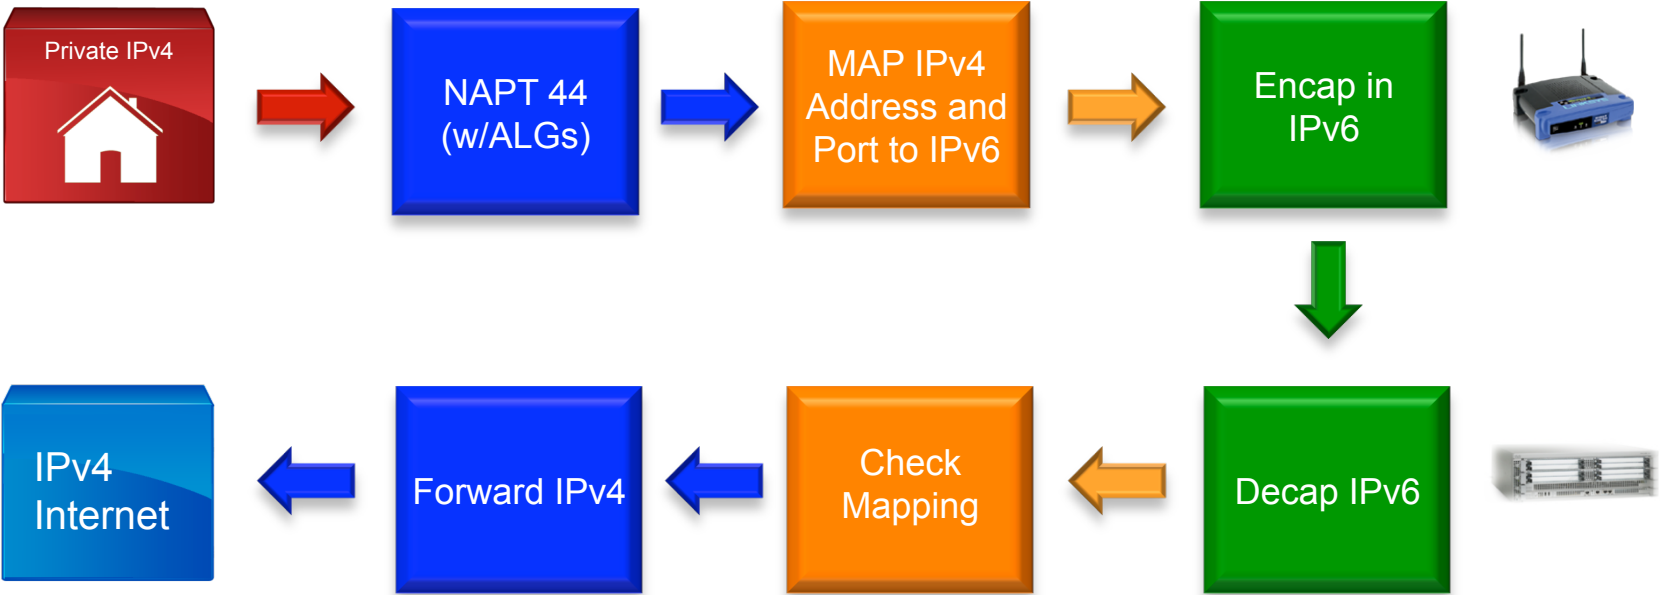
- Handle traffic to/from a given MAP domain
- Reachable via anycast, “built-in” load-balancing
- Each MAP rule is similar to a single LW46 entry
 - but MAP rules allow for aggregation and LW46 entries do not
- Can be processed inline with normal IP traffic

③ Packet Flow and Forwarding

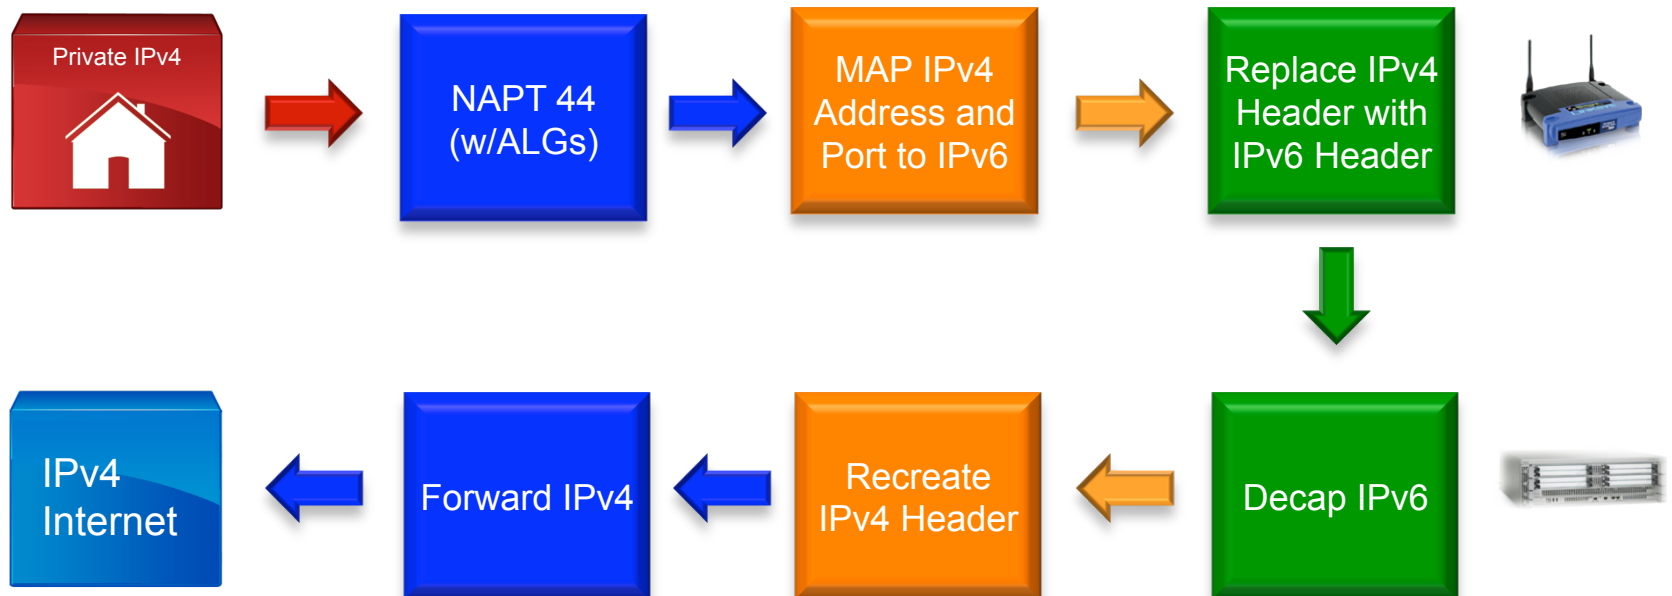


- IPv4 follows IPv6 routing within a domain (traffic destined to another subscriber does not traverse the BR)
- All other traffic sent via anycast to any MAP BR
- Forwarding is handled either by double translation (MAP-T) or encapsulation (MAP-E)


Forwarding (Encapsulation, MAP-E)



Forwarding (Translation, MAP-T)




Encapsulation or Translation – Boils down to 20 bytes



MAP-E

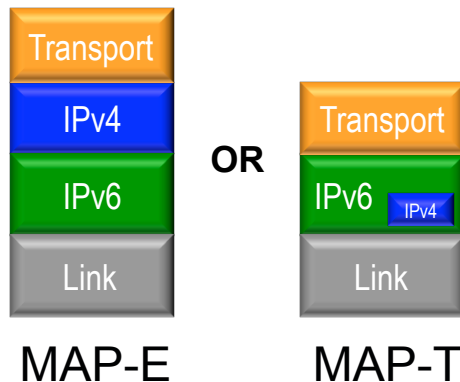

OR



MAP-T



Encapsulation or Translation – Boils down to 20 bytes



- The softwires WG was for a very long-time wedged with entrenched parties on all sides of MAP-E vs. T
- Encapsulation:
Well-understood, simple, transparent, same as stateful dual-stack lite
- Translation:
Native IPv6 ACLs and DPI functionality not masked by IPv4 header. NAT64 code reuse. Feels like “Real IPv6.”
- Arguments gravitate towards speculation about what future IPv6 deployments will require and what feature availability will be
- Architecturally, both are still TUNNELING

Standardizing MAP in the IETF


- MAP-E will be a Standards Track RFC
<http://tools.ietf.org/html/draft-ietf-softwire-map-07>
- MAP-T, 4rd, etc. will be Experimental or Informational
<http://tools.ietf.org/html/draft-ietf-softwire-map-t-01>
- LW46/Pubilc4over6 can be viewed as “special cases” of MAP
- Goal: One unified standard for CPE vendors
- Stretch Goal: One unified standard for BR/AFTR vendors

Running code

MAP testing by NIC.br

- “The working applications had no need of a special configuration to work.”
- Most of the applications work OK
- FTP active mode does not work.
- More info:
<http://tools.ietf.org/html/draft-cordeiro-experience-mapt-testing-00>

European Advanced Networking Test Center



**IPv6 MAP Testing
at Multi-Vendor Interoperability Test Event 2013**

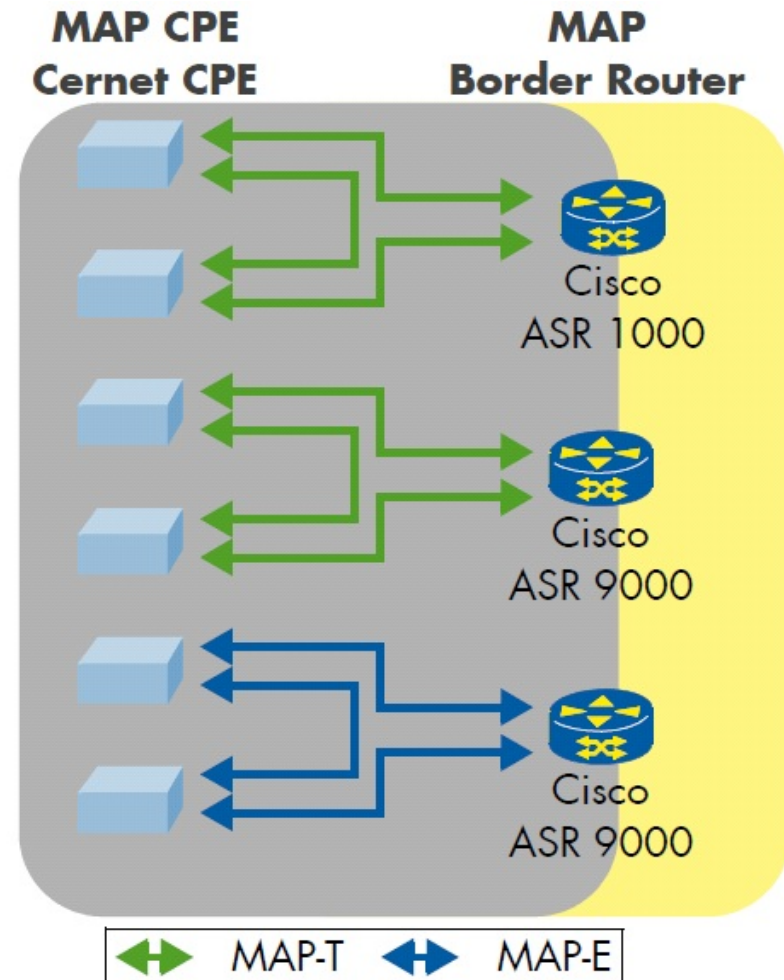
Mapping of Address and Port (MAP)

IPv6 Tests

- Stateless counterpart to DS-Lite
- Designed to be used without Carrier-Grade NAT
- Cisco ASR1000, ASR9000 and Cernet (CPE) participated


Successfully tested:

- Mapping of Address and Port with Encapsulation (MAP-E)
- Mapping of Address and Port using Translation (MAP-T)



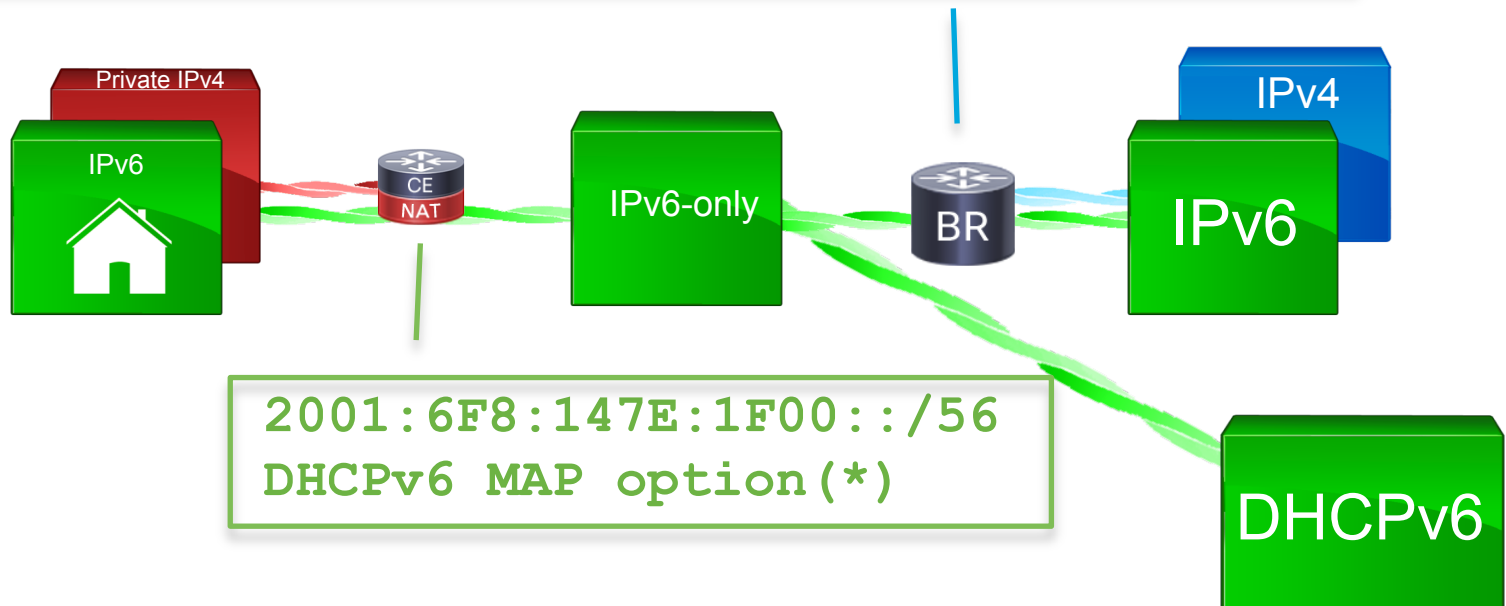
MAP on ASR 9K

- MAP does not route traffic through the ISM Blade, yielding line rate performance.
 - Using A9K-24x10G line cards = 240 Gbps per slot!
 - $7 \times 240 = 1.68$ Tbps on a 9010 chassis.
- DS-Lite routes traffic through the ISM Blade
 - 14Gbps per slot



MAP-T demo configuration

```
nat64 map-t domain 1
  default-mapping-rule 2610:D0:1208:CAFE::/64
  basic-mapping-rule
    ipv6-prefix 2001:6F8:147E:1000::/52
    ipv4-prefix 153.16.17.83/32
    port-parameters share-ratio 16
```




MAP CPE: open source

← → ↻ mapt.ivi2.org:8039/mapt.html

 **MAP (MAP-T/MAP-E) Configuration Examples**

Testing Topology



- [Configuration parameters](#)

- Linux (FC 11)
- OpenWRT Linksys 54GL
- TP-Link WR1043ND
- Source code

<http://mapt.ivi2.org:8039/mapt.html>


http://github.com/cernet/MAP

The screenshot shows the GitHub interface for the repository `cernet/`**MAP**. The browser address bar displays `https://github.com/cernet/`**MAP**. The repository is public and has 0 pull requests, 0 issues, and 1 branch (master). The description states: "An open source CPE implementation of MAP-E/MAP-T which can be run on Linux and Openwrt. — [Read more](#)". Cloning options include "Clone in Mac", "ZIP", "HTTP", "SSH", and "Git Read-Only" with the URL `https://github.com/cernet/`**MAP**`.git`. The commit history shows two recent commits:

Commit Message	Author	Time	Details
fix a small comment mistake	cernet	an hour ago	fix a small comment mistake [cernet]
MAP-1.0 version	cernet	4 hours ago	MAP-1.0 version [cernet]

Summary:

- You must have deployed IPv6 to use any of this!
- LW46 lighter than DS-Lite, both are heavier than MAP
- MAP works with no state and high performance



Thank you.

