


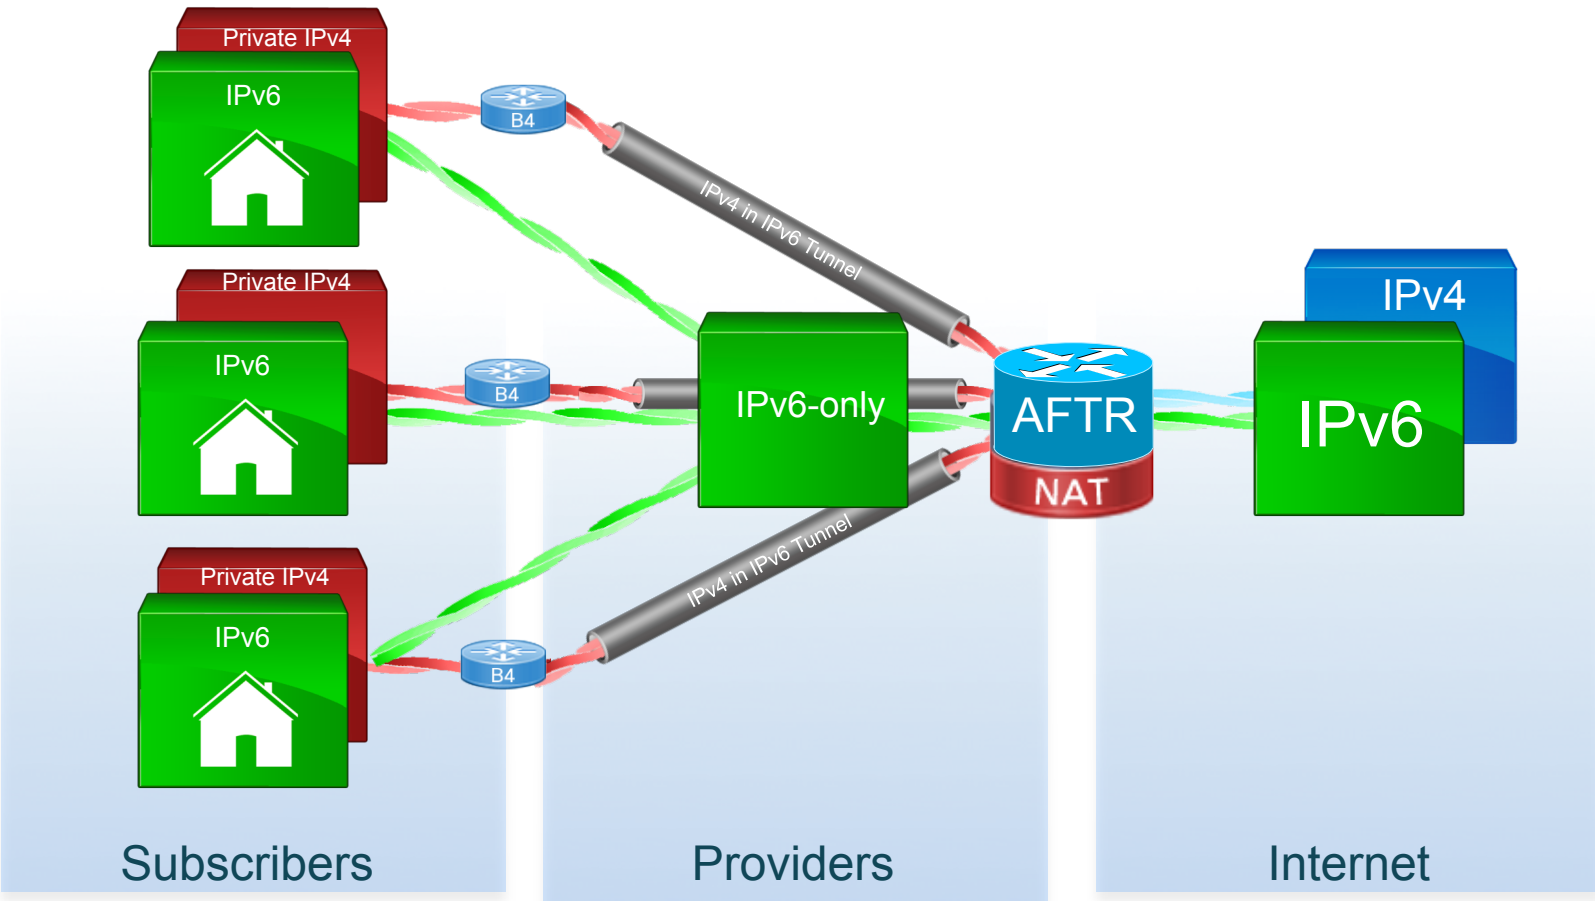
Mapping Address + Port

Mark Townsley
Cisco Fellow


APNIC 36, August 2013



Dual Stack Lite (DS-Lite)




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




*draft-cui-software-b4-translated-ds-lite-08

Mapping Address + Port (MAP)



Imagine the Internet without any IP address aggregation



The image shows a table titled "IPv4 CIDR Chart" from RIPE NCC. The table lists IP address ranges, their bit lengths, prefixes, and subnet masks. A large red 'X' is drawn over the entire table, indicating that the scenario of no IP aggregation is being rejected or is not the intended state.

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	255.255.254.0
1 K	10	/22	255.255.252.0
2 K	11	/21	255.255.248.0
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	255.254.0.0
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

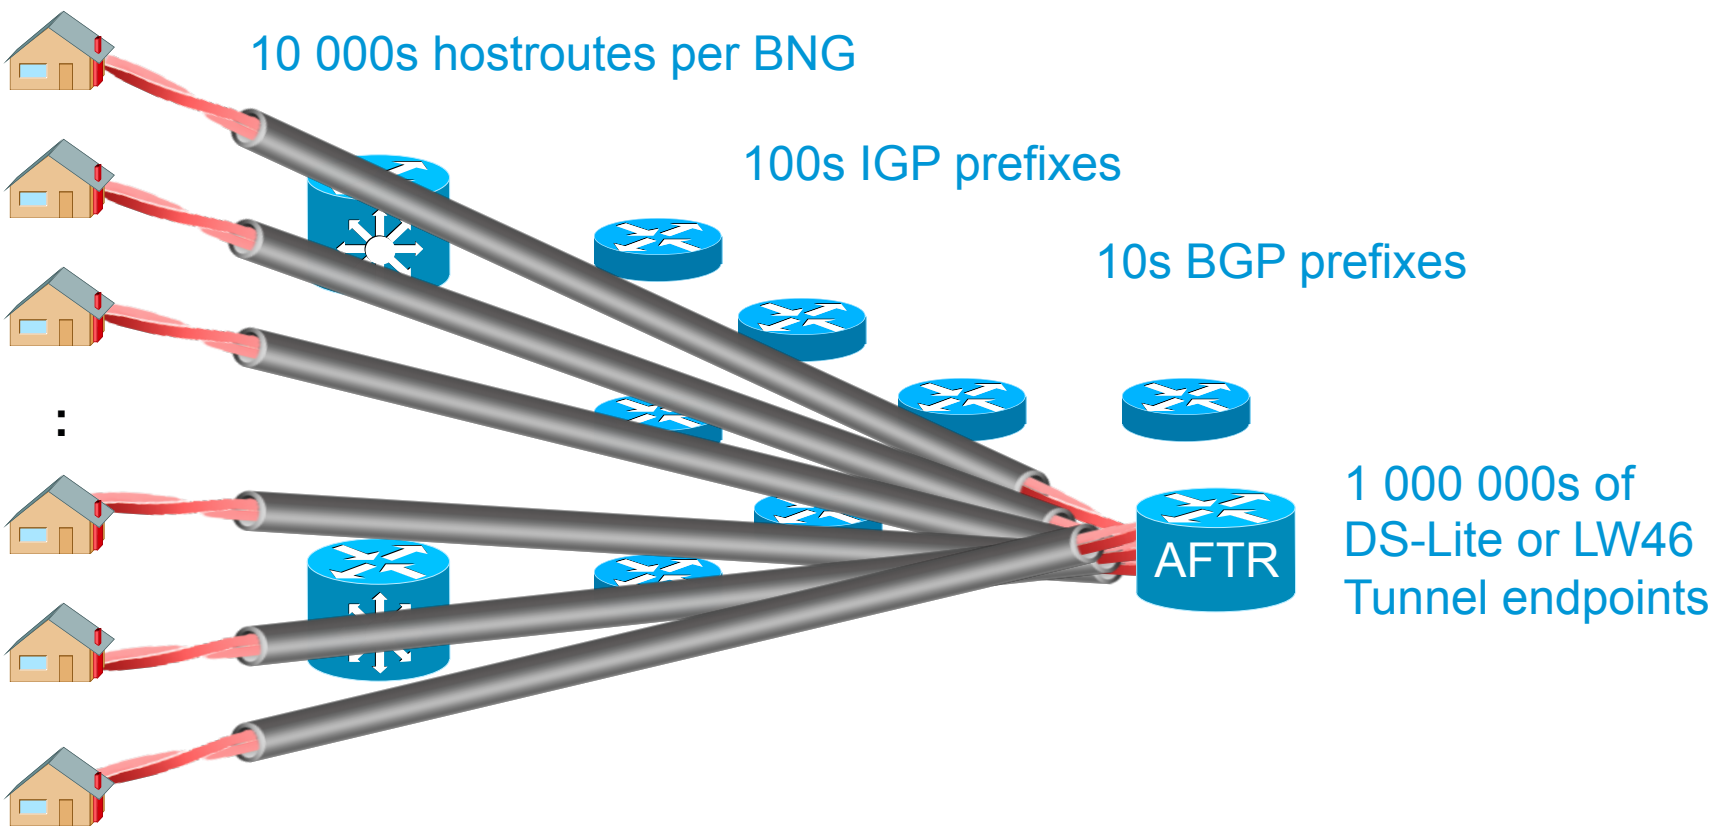
K = 1,024 • M = 1,048,576

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

www.ripe.net

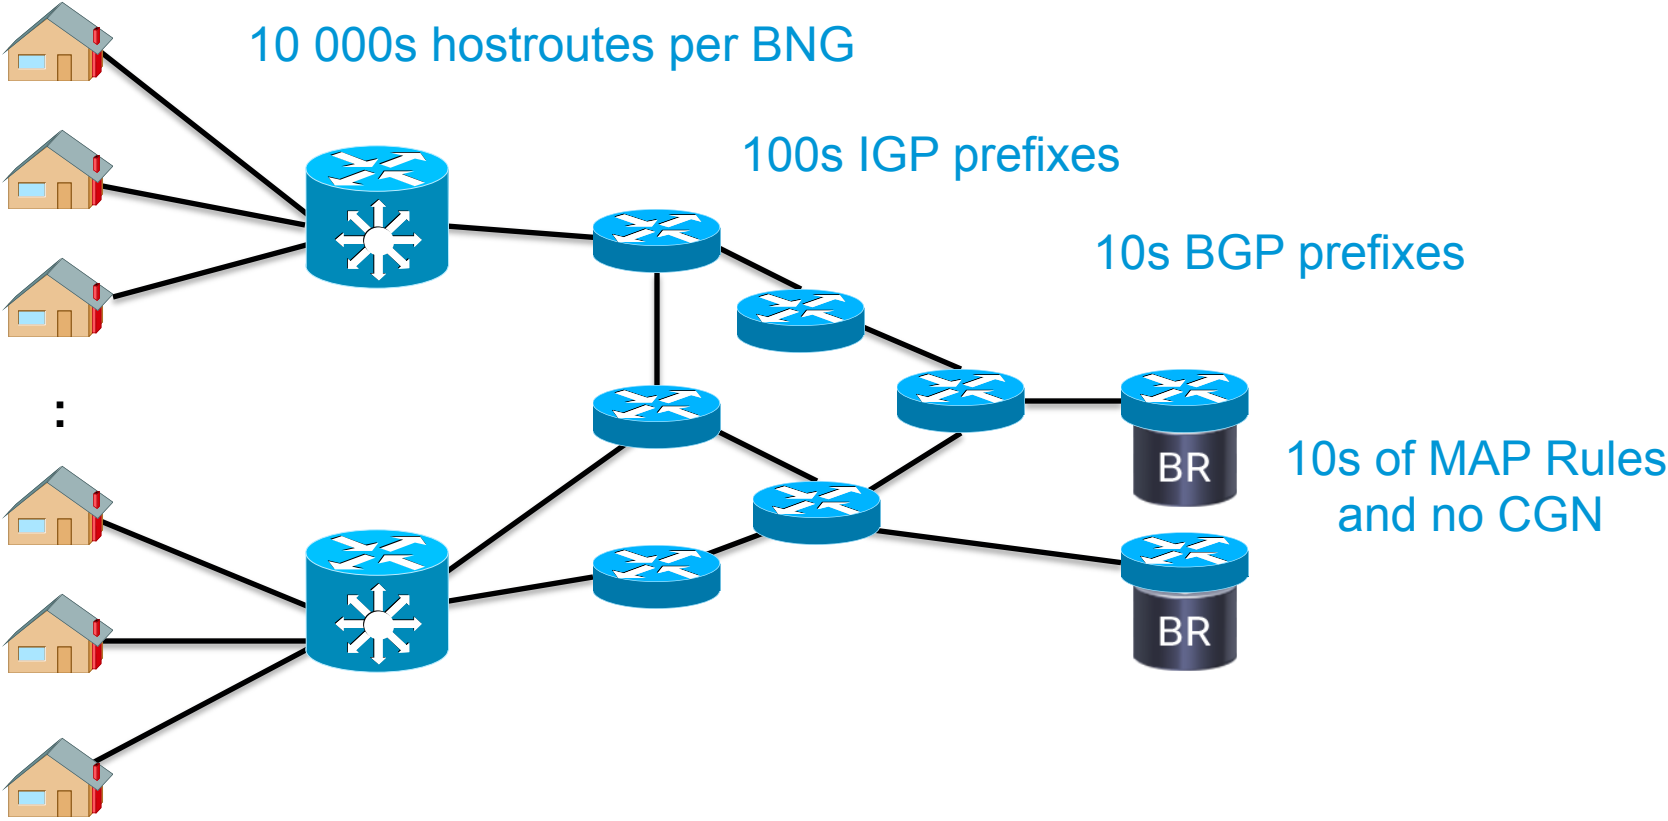
DS-Lite/LW46/Public 4over6 – No Aggregation, just lots of tunnels!

1 000 000s of subscribers



MAP Exploits IPv6 Aggregation

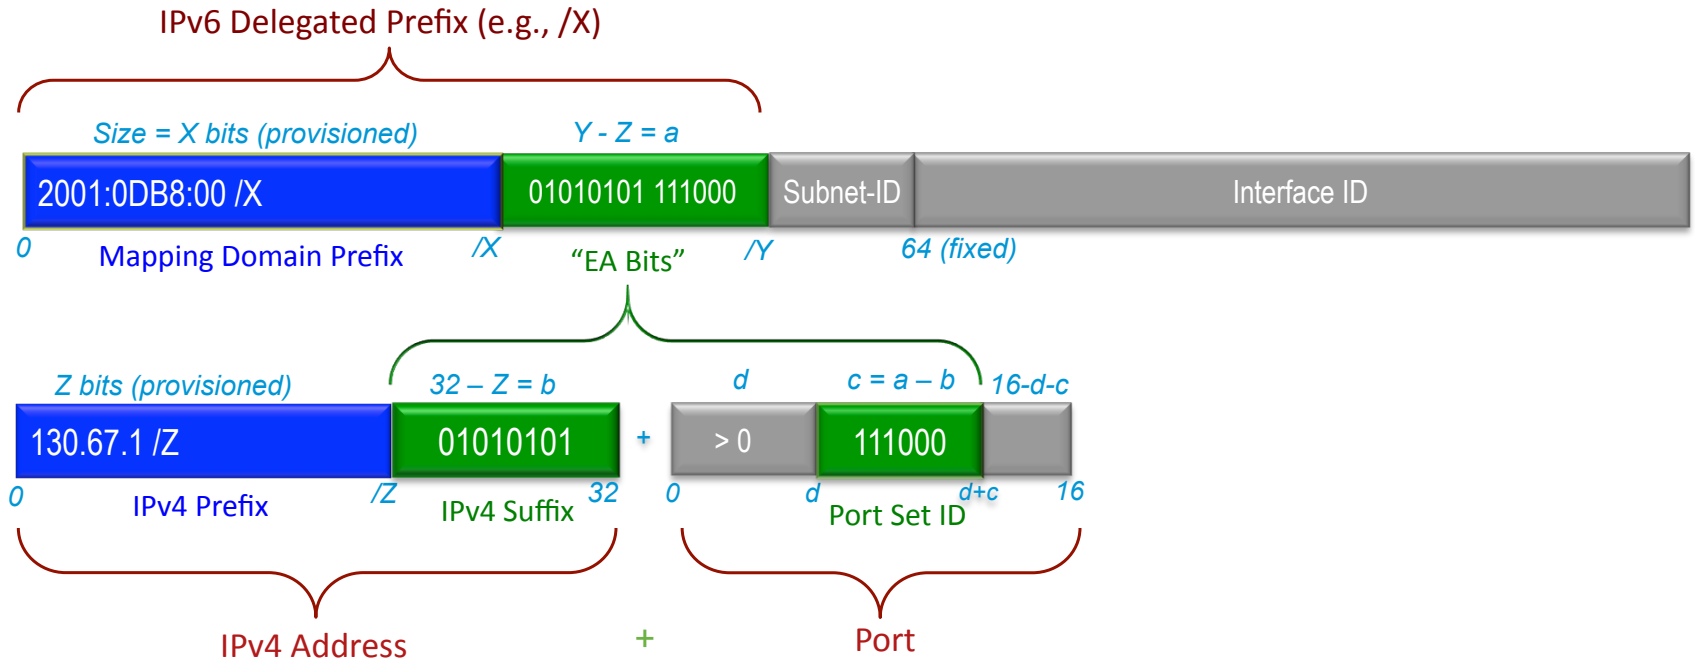
1 000 000s of subscribers



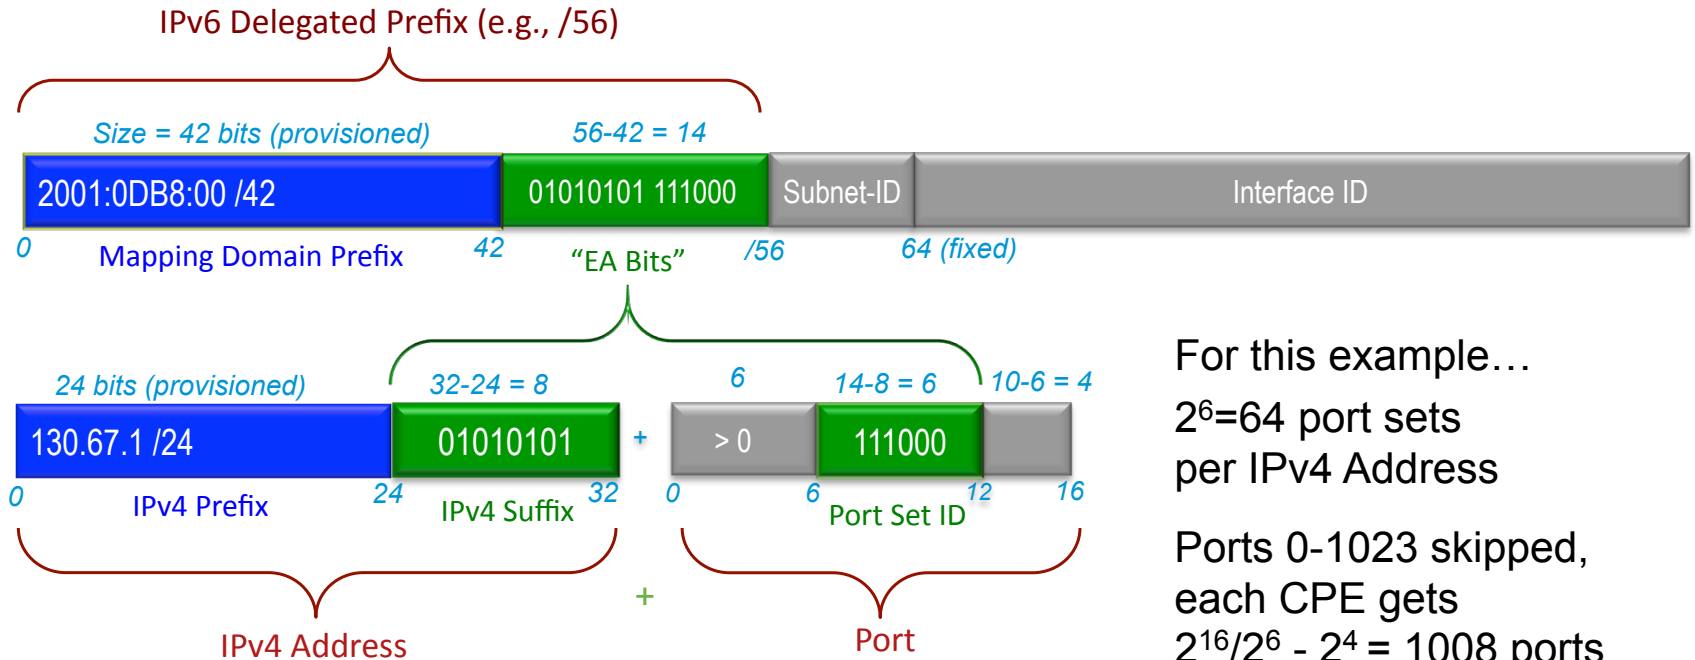
MAP: Easy as 1-2-3

- 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

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



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@external.cisco.com

Last updated: 6/19/2012

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

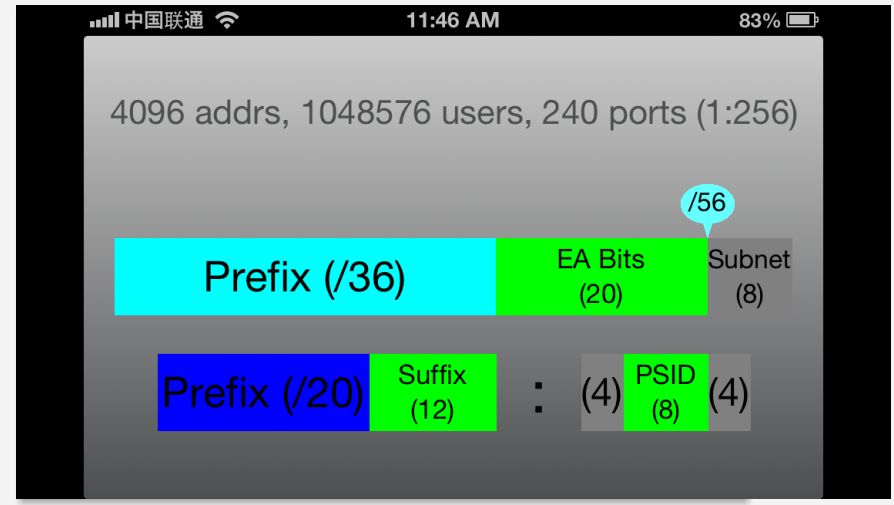
Description

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
- Scales according to traffic and number of rules only, not number of users or number of users per rule


③ Packet Flow and Forwarding



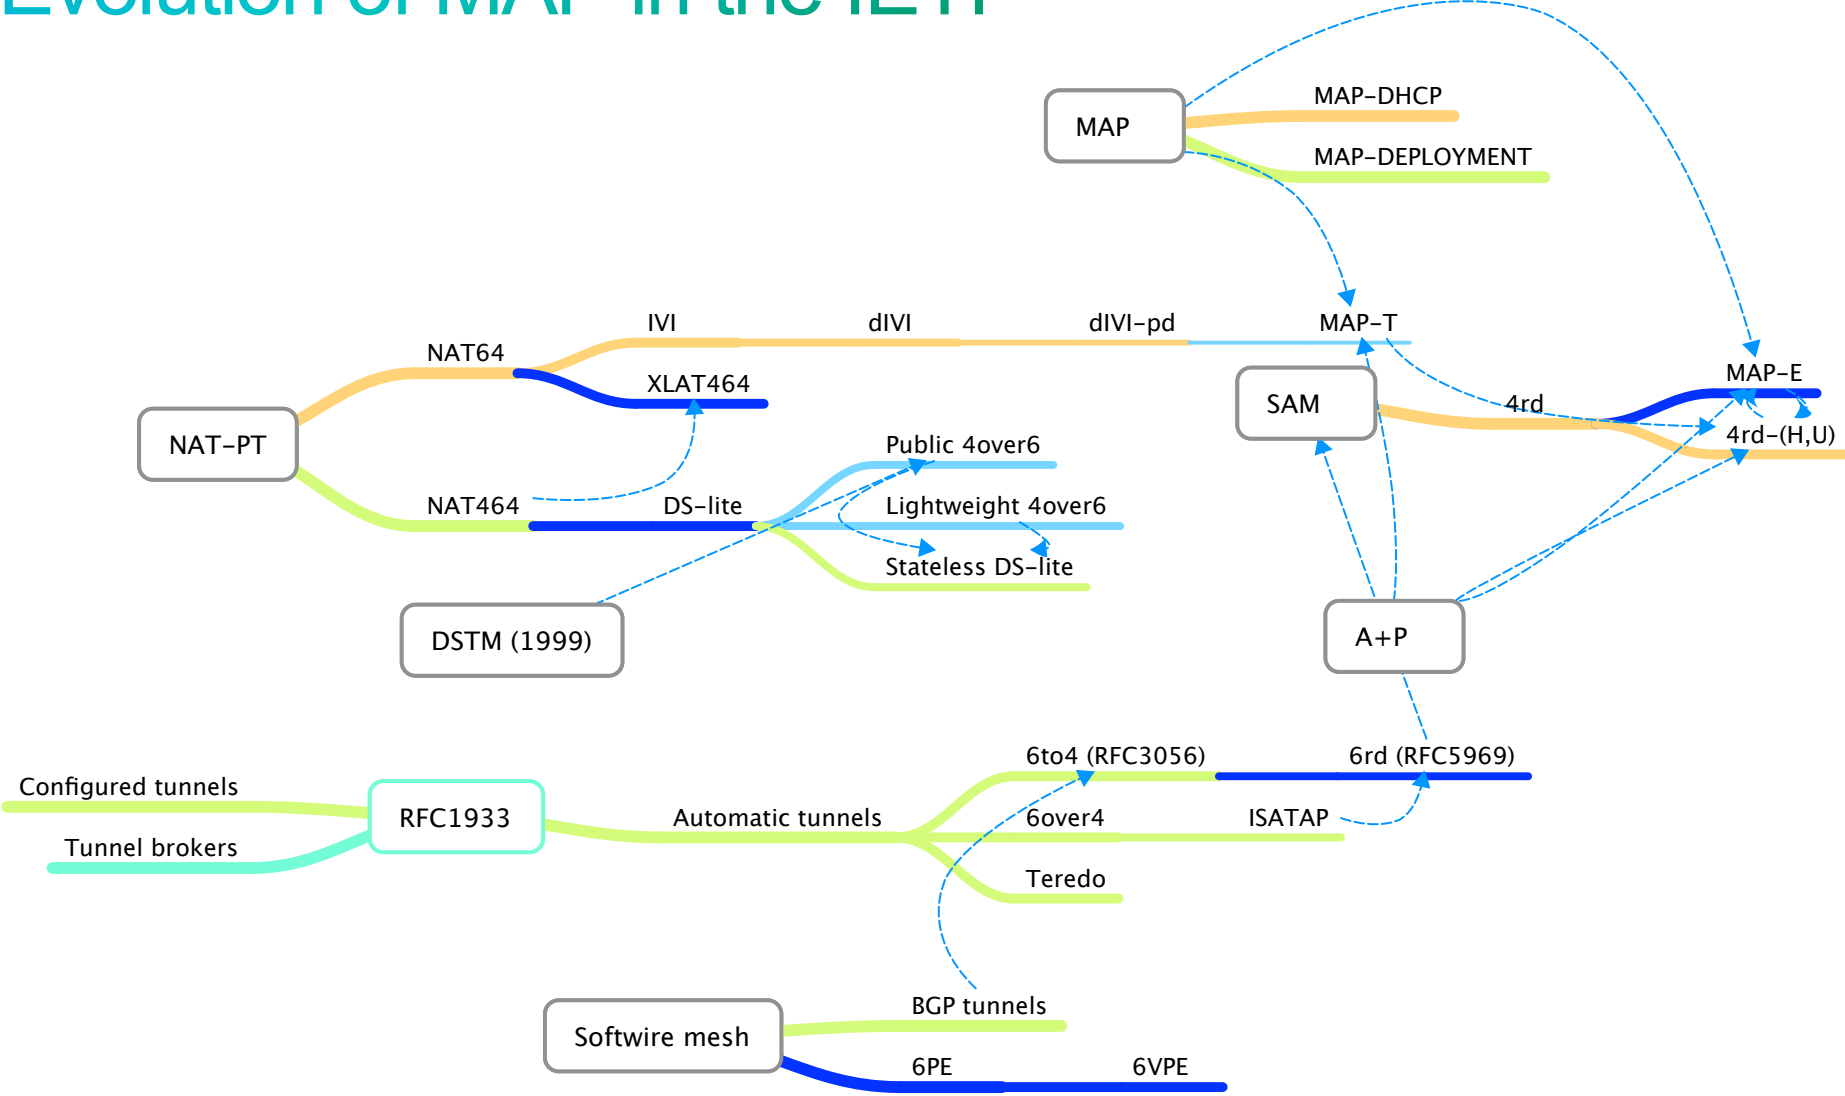
- IPv4 follows IPv6 routing within a domain (traffic destined to another subscriber does not traverse the BR)
- All other traffic sent via anycast to any MAP BR
- Forwarding is handled either by double translation (MAP-T) or encapsulation (MAP-E)

MAP support @ Cisco – ASR 9K

- MAP (T or E) does not route traffic through the ISM Blade
 - 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



Evolution of MAP in the IETF




IETF Standardization Path (softwires WG)

- MAP-E will be published as a Standards Track RFC
- MAP-T, 4rd, Public4over6, etc. will be published as Experimental or Informational, *after* MAP-E is finished.
- Working hard now on defining a “Unified CPE” – Aims to allow any “MAP” CPE to work with MAP-T, MAP-E, LW46, etc.

In Sum:

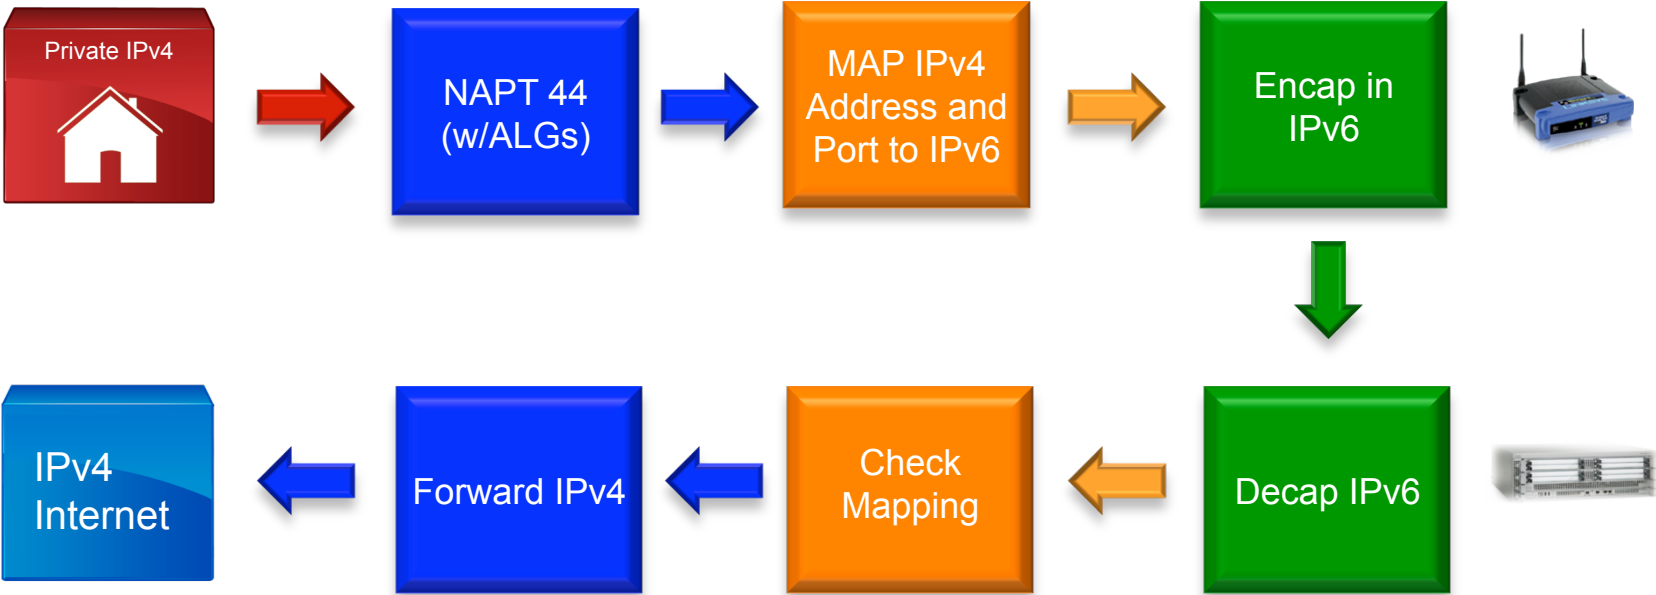
- You must have deployed IPv6 to use any of this!
- MAP uses the power of IPv6 routing to deliver IPv4 without CGN!
- LW46 lighter than DS-Lite, both heavier than MAP
- Commercial products are arriving now, open source CPE code available



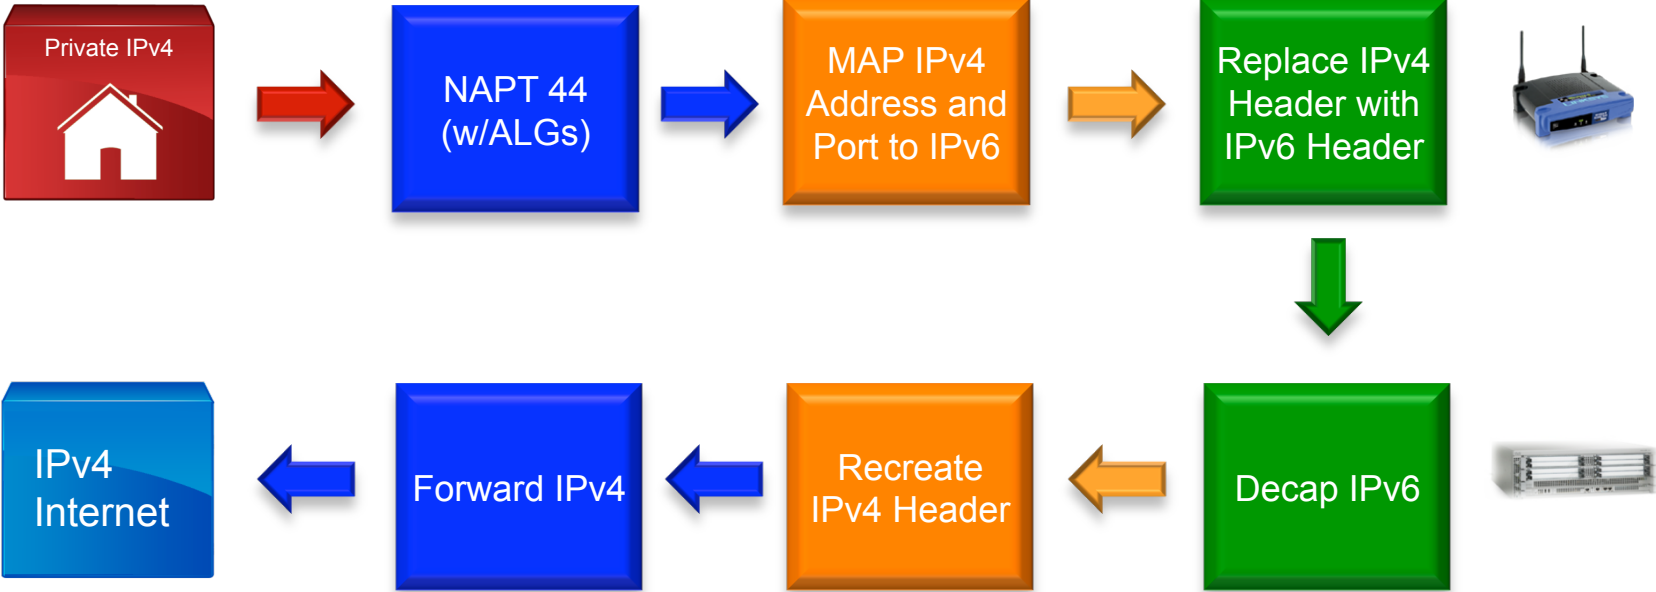
Thank you.



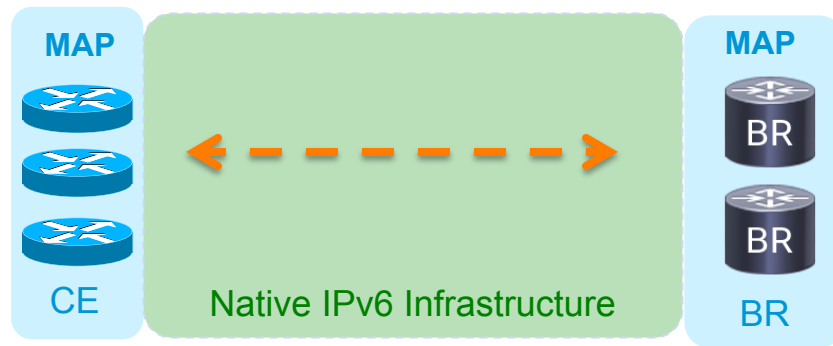
Forwarding (Encapsulation, MAP-E)





Forwarding (Translation, MAP-T)



Encapsulation or Translation – Boils down to 20 bytes



OR



MAP-E

MAP-T