

Mapping of Address and Port

Stateless IPv4 over IPv6

[draft-mdt-softwire-mapping-address-and-port](#)

Congxiao Bao, Mohamed Boucadair, Gang Chen, Maoke Chen, Wojciech Dec, Xiaohong Deng, Remi Despres, Jouni Korhonen, Xing Li, Satoru Matsushima, Tomasz Mrugalski, Tetsuya Murakami, Jacni Qin, Qiong Sun, Ole Troan, Tina Tsou, Dan Wing, Leaf Yeh, Jan Zorz

Anglo-EU Translation Guide

What the British say	What the British mean	What others understand
I hear what you say	I disagree and do not want to discuss it further	He accepts my point of view
With the greatest respect...	I think you are an idiot	He is listening to me
That's not bad	That's good	That's poor
That is a very brave proposal	You are insane	He thinks I have courage
Quite good	A bit disappointing	Quite good
I would suggest...	Do it or be prepared to justify yourself	Think about the idea, but do what you like
Oh, incidentally/ by the way	The primary purpose of our discussion is...	That is not very important
I was a bit disappointed that	I am annoyed that	It doesn't really matter
Very interesting	That is clearly nonsense	They are impressed
I'll bear it in mind	I've forgotten it already	They will probably do it
I'm sure it's my fault	It's your fault	Why do they think it was their fault?
You must come for dinner	It's not an invitation, I'm just being polite	I will get an invitation soon
I almost agree	I don't agree at all	He's not far from agreement
I only have a few minor comments	Please re-write completely	He has found a few typos
Could we consider some other options	I don't like your idea	They have not yet decided

MAP is:

- a building block. Describes how to map between an IPv4 address, prefix or IPv4 address and port into and IPv6 address.
- Used by various protocol mechanisms (4rd-
{E,T,U}, dIVI-PD, stateless 4over6...)
- Provisioned by the MAP DHCPv6 option

Stateful IPv6 over IPv4

- L2TP
- 6PE/6VPE
- BGP tunneling
- LISP
- GRE
- Configured Tunnels RFC1933

Stateless IPv6 over IPv4

- Automatic tunnels RFC1933
- 6over4
- 6to4
- ISATAP
- 6rd
- Teredo


Stateful IPv4 over IPv6

- DS-Lite
- L2TP
- GRE
- LISP
- DS-Lite with A+P
- IPv4 over DS-lite


Stateless IPv4 over IPv6

- 4rd-E
 - Stateless 4over6
 - SA46T-AS
 - dIVI
 - 4rd-T
 - 4rd-U
 - dIVI-pd
 - SD-NAT
- MAP(A+P)**

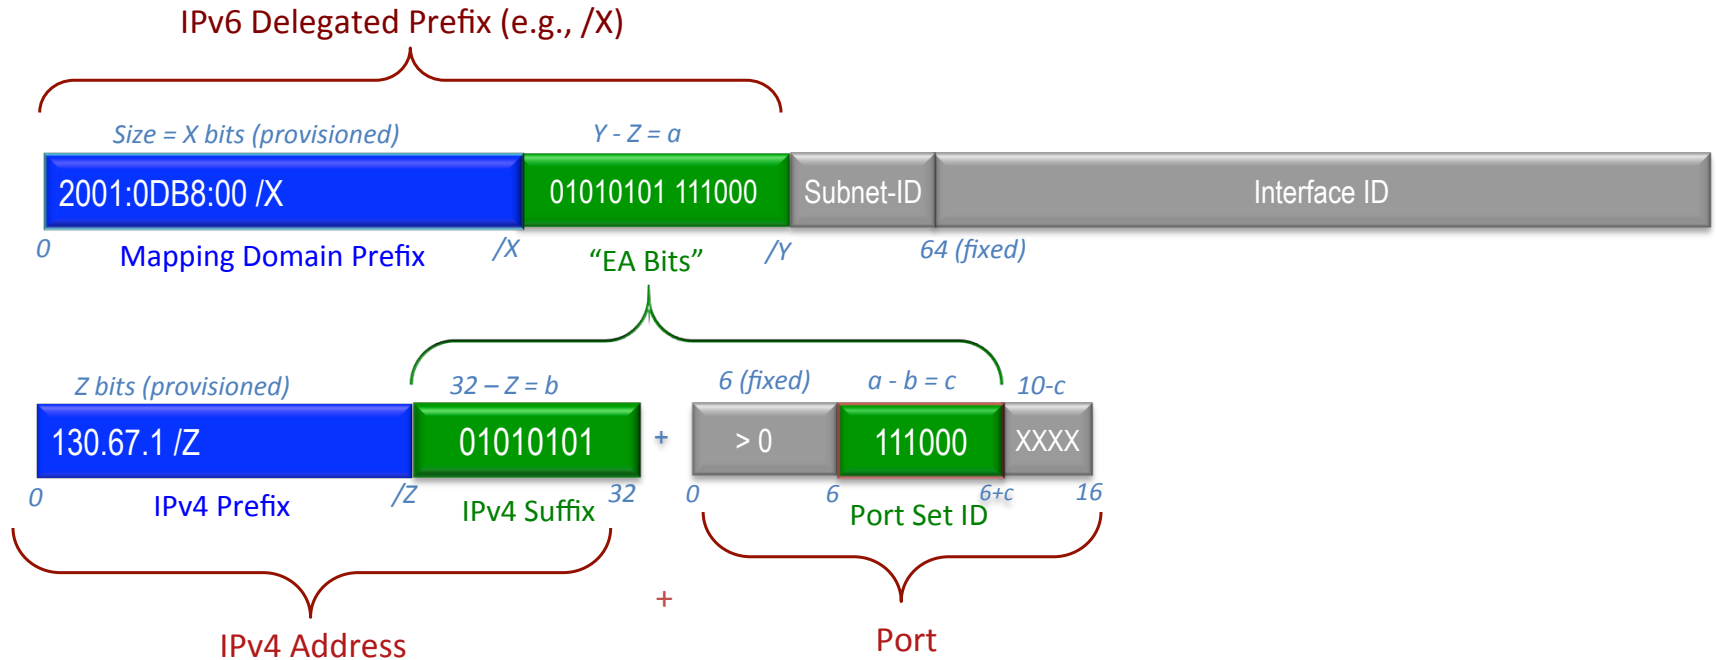
Dual Stack Lite (DS-Lite)



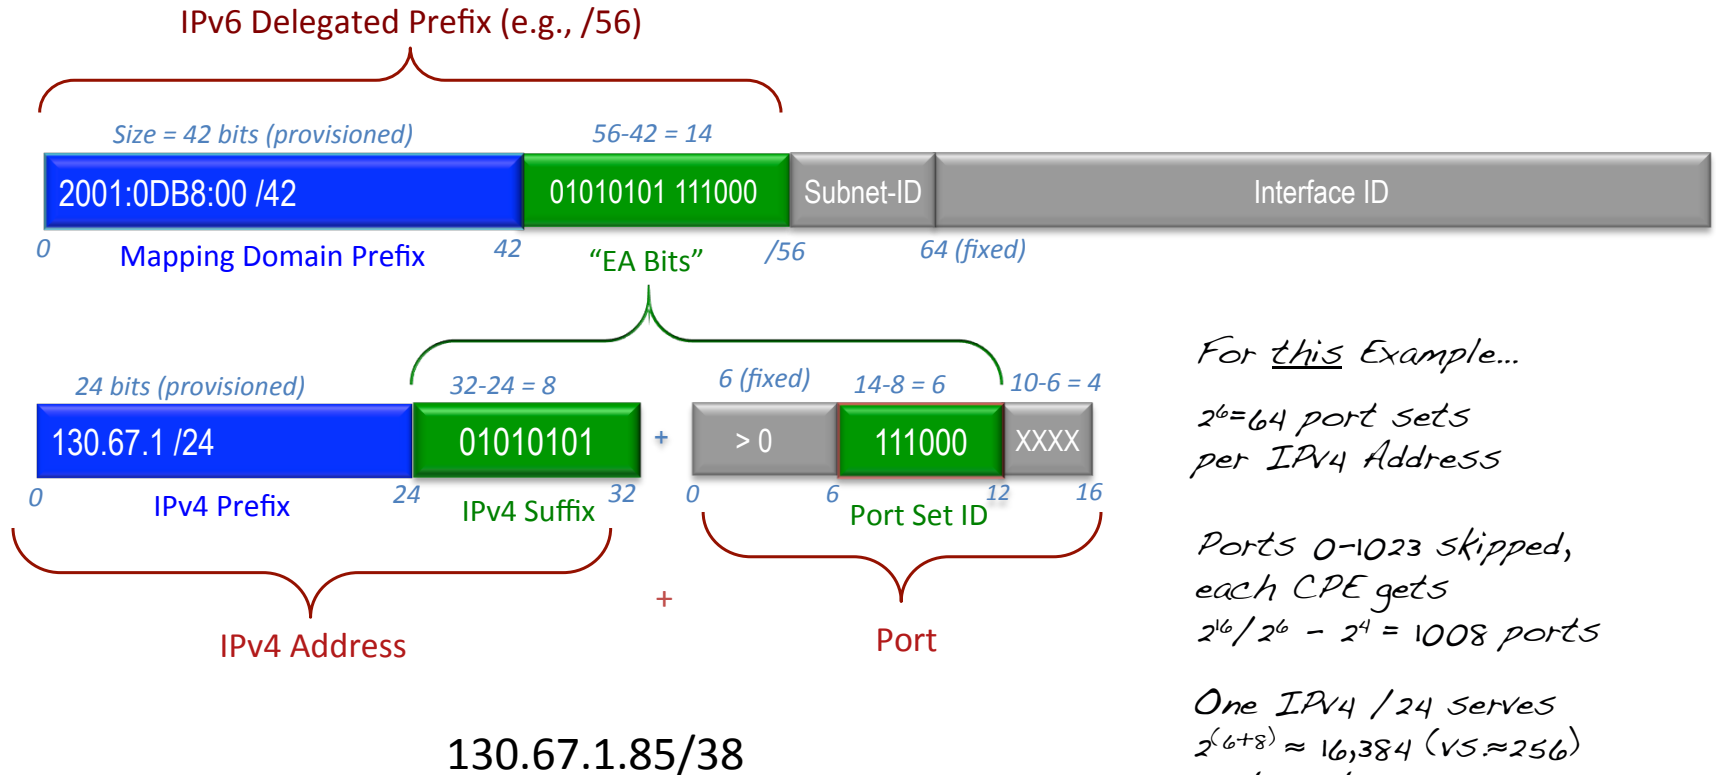
IPv6 Residual Deployment (4rd)



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
 $2^6 / 2^6 - 2^4 = 1008$ ports

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

MAP consensus:

- Choice of port mapping algorithm. “infix” / generalized modulus algorithm and bit representation.
- IPv6 prefix format and encoding of IPv4 address, prefix or shared IPv4 address bits.
- A MAP domain can have multiple mapping rules (accommodate multiple IPv4 subnets)

MAP - open issues

#1: Granularity of port set size

- If the requirement is: “must support differentiated sharing ratio within a single shared IPv4 address/IPv4 subnet within a single mapping rule”
 - Solution is Max PSID (see Remi’s slide)
 - *Tradeoff: Destination spray*
- Alternative: Fixed port set size per IPv4 subnet. One mapping rule == IPv4 subnet and sharing ratio

#2: Checksum (translation)

- If the requirement is: “A translator should generate packets with a valid L4 checksum” and “this should be done without modifying the IPv4 packet”
 - Solution is: Make the IPv6 header checksum neutral by embedding a Checksum Preserver field (16 bits) in the IPv6 addresses
 - *Tradeoff: Destination spray. Source address and destination address will vary depending on changes in IPv6 header*
- Alternative: Rewrite the L4 checksum

#3: Destination spray

- Encapsulated or translated traffic is not anchored to a single “tunnel end-point” address.
- How to determine if received traffic is native or MAP?
 - Solution: Add a V-octet in the Interface-identifier
 - *Tradeoff: Incurs penalty for native IPv6 traffic (V-octet filter), can we be guaranteed that no valid interface-id (or longer prefix) doesn't overlap with the V-octet?*
- Alternative: Reserve /64 (for translation with IPv4 prefix, or for BR) or /128 for MAP traffic.

#4: Interface-id

- Must include the complete IPv4 DA in the IPv6 address for out of domain translated traffic.
- If the requirement is “the interface-id should be globally unique, should be possible to parse (for features / classifiers) without knowing the mapping rule”:
 - Solution is: include full IPv4 address, PSID and L fields.
 - *Tradeoffs: Must redundant information in prefix and interface-id be enforced?*
- Alternative: Just include the IPv4 address as specified in RFC6052

#5: Multiple BR prefixes

- If the requirement is: “The same MAP DHCPv6 option must be given to MAP CEs”
 - The solution is: For mapping rules to include the BR IPv6 address, and the CEs do “source based routing” to pick correct exit point.
 - *Tradeoff: More complex mapping rules, more complex CE implementation*
- Alternative: If CEs are required to use different BR addresses, provision with different DHCPv6 options.

Next steps:

- Sense of the room with regards to open issues
- Adopt MAP drafts as working group items
- Converge the various encapsulation and translation protocol mechanisms and adopt
- Drink beer.