

# Mapping of Address and Port using Translation (MAP-T)

E. Jordan Gottlieb

Network Engineering and Architecture

[jordan.gottlieb@charter.com](mailto:jordan.gottlieb@charter.com).

# What is MAP-T?

## One of two transport modes of MAP

- Mapping of Address and Port with Encapsulation (MAP-E) is the other transport mode

## IPv4/IPv6 stateless translation

- Used to support IPv4 network to IPv4 connectivity over an IPv6 network

## Supports N:1 IPv4 address sharing

- Competing technology in the CGN/LSN solution space
- Different than other CGN/LSN solutions as state remains on the subscriber's router
- Subscriber router is assigned a public IPv4 address (logically) with a unique set of layer-4 ports

# Interest by Vertical

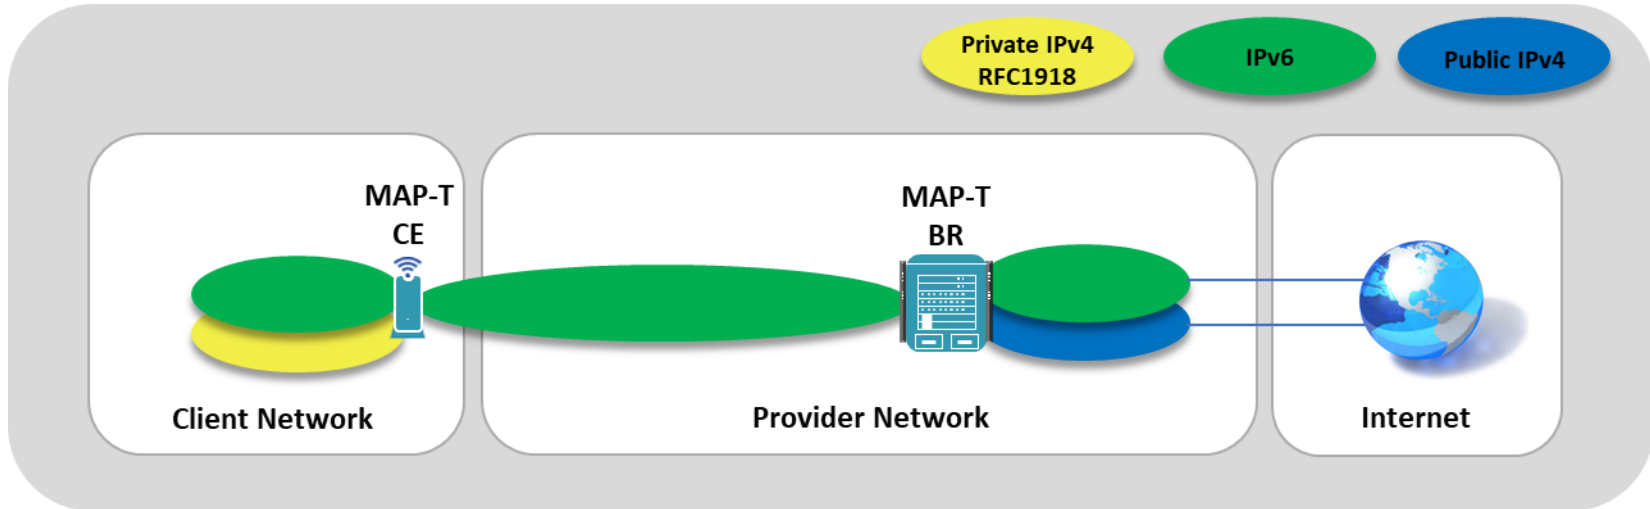
## Content Provider

- Shared IPv4 address versus dedicated addressing per subscriber
  - Licensing
  - Reputation
  - Possibly geolocation (depends on implementation)
- Layer-4 port multiplexing scale (not 16-Bit)
- Peer-to-peer traffic
  - Triangular routing
  - Implication for UPnP/IGD and static port forwarding
  - UDP hole punching techniques are not impacted
- IPv6 strategy

## Service Provider

- Potential for adoption

# MAP-T Components in Context



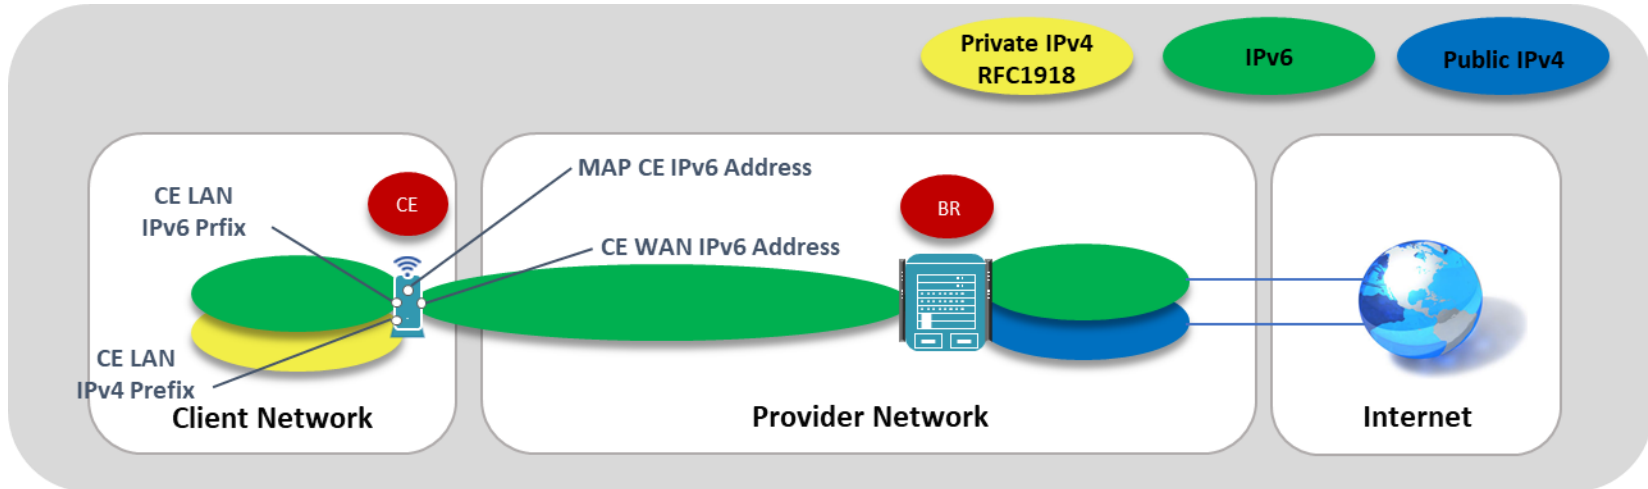
## MAP-T Customer Edge (MAP-T CE)

- A home gateway (wireless router, cable modem + router, etc.)
- Provides IPv4/IPv6 stateless translation as well as native IPv6 forwarding

## MAP-T Border Relay (MAP-T BR)

- Provider side component which supports the IPv4/IPv6 stateless translation

# MAP-T Subscriber Addressing



## IPv4 MAP-T CE Addressing

- Subscriber assigned IPv4 prefix on LAN/WLAN segment (RFC1918)

## IPv6 CE Addressing

- Provider facing interface is an IPv6 address that is independent of MAP-T
- MAP-T address is derived from assigned end-user prefix
- LAN/WLAN prefix is allocated from assigned end-user prefix

# MAP-T Terminology

## Mapping Rule

- Associates IPv4 prefix to an IPv6 prefix
  - The IPv4 prefix is a prefix used for the “IPv4 NAT Pool” or shared IPv4 address(s)
  - The IPv6 prefix is the aggregate of the prefixes assigned to individual subscribers
- Used by MAP-T CE to determine its IPv4 address and optionally the layer-4 port allocation
- Allows a MAP-T BR to support IPv4/IPv6 stateless translation

## Base Mapping Rule

- The mapping rule(s) used by a MAP-T CPE given a list of mapping rules

## Default Mapping Rule (DMR) or BR Prefix

- An IPv6 prefix used to embed IPv4 destination addresses (RFC 6052)
- Similar to IPv4 default route

# MAP-T Terminology

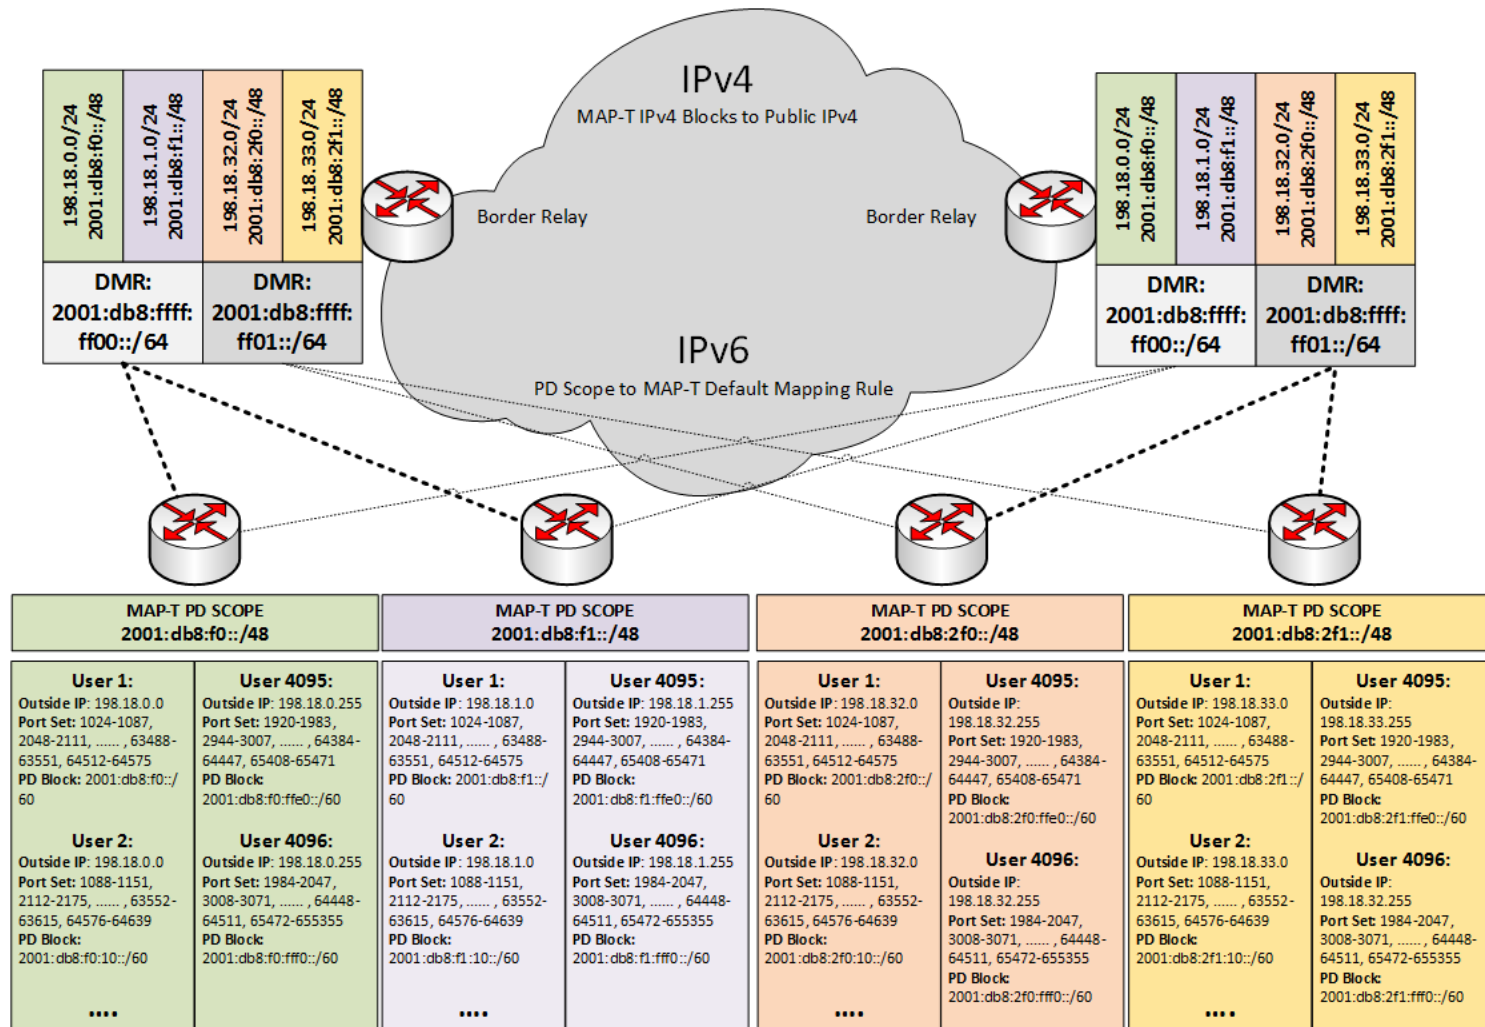
## MAP Domain

- A mechanism to group a set of MAP-T CEs with a set of MAP-T BRs

## Embedded Address (EA) bits


- A set of bits within a IPv6 user prefix which identify a unique or shared IPv4 address and a index into the associated set of layer-4 ports
  - Mapping rule provides the offset information for these bits

# MAP-T Terminology in Context





# MAP-T in Action – Subscriber to Internet



**Packet #1:** IPv4 private connected workstation transmits a packet to a destination server 1.1.1.1

1025	80	TCP	192.168.1.101	1.1.1.1
------	----	-----	---------------	---------

**Packet #2:** MAP-T CE performs NAPT

4000	80	TCP	192.0.2.2	1.1.1.1
------	----	-----	-----------	---------


**Packet #3:** MAP-T CE translates post NAPT IPv4 header to a IPV6 header and forwards it towards the MAP BR

4000	80	TCP	2001:db8:101:2e8:0:c000:202:3a	2001:db8::1:101:0100:0
------	----	-----	--------------------------------	------------------------

**Packet #4:** MAP-T BR translates IPv6 header to the post NAPT IPv4 header and forwards it towards the destination address

4000	80	TCP	192.0.2.2	1.1.1.1
------	----	-----	-----------	---------

# MAP-T in Action – Internet to Subscriber

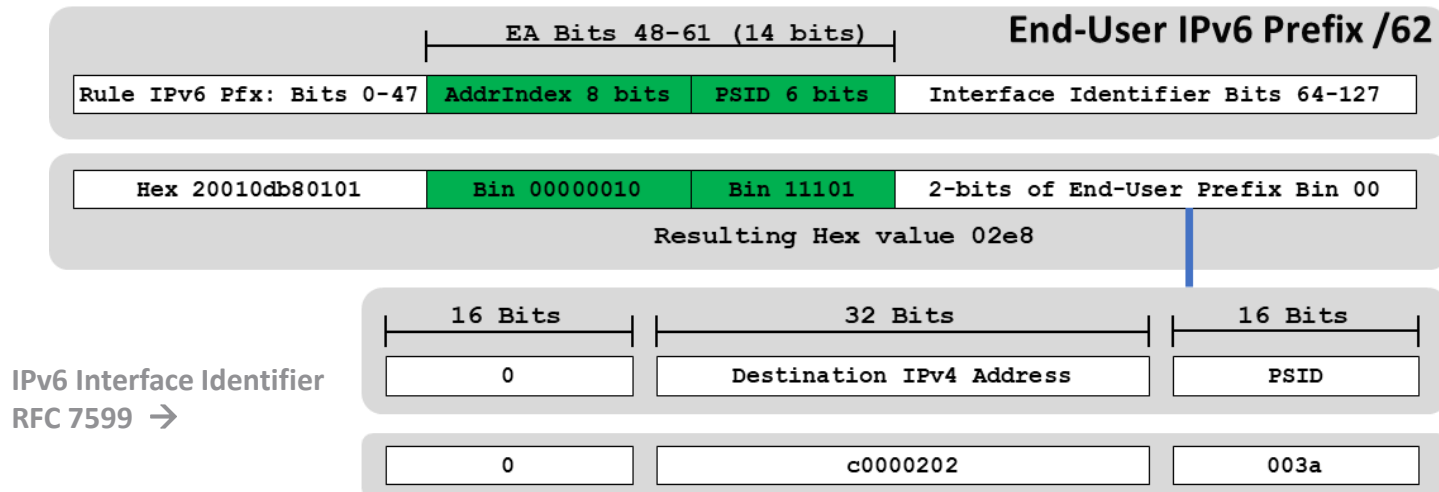


- Incoming packet src=1.1.1.1:80 dst=192.0.2.2:4000
- BR Prefix: 2001:db8::/64
- Mapping Rule
  - Rule IPv6 Prefix: 2001:db8:101::/48
  - Rule IPv4 Prefix: 192.0.2.0/24
  - EA Length: 14
- Source stateless translation – Easy!
  - Embed 1.1.1.1 in BR Prefix 2001:db8::/64 = 2001:db8::1:101:0100:0
- Result

80	4000	TCP	2001:db8::1:101:0100:0	Destination
----	------	-----	------------------------	-------------

# MAP-T in Action – Internet to Subscriber

- Destination stateless translation (Translate to MAP-T CE address)
  - Find mapping rule by matching longest match against Rule IPv4 prefix
  - EA bits contain index value to destination IPv4 and index value to port set which contains the destination port (PSID)
    - Destination port 4000 belongs to PSID index value 58 (hex 3a)



- Result

80	4000	TCP	2001:db8::1:101:0100:0	2001:db8:101:2e8::c000:202:3a
----	------	-----	------------------------	-------------------------------

# MAP-T Benefits

## Nearly stateless border relay

- Fragment forwarding

## Nearly seamless alignment with IPv6 deployment

- No additional logging overhead
- DHCPv6 provisioning delta is option 95
- Supports existing 5-tuple use cases
  - Classification
  - ECMP
  - Filtering

## Simple integration with existing deployment

- No need for separate forwarding logic (routing instance/VPN/policy-routing)
- Asymmetrical forwarding is supported (no pinning required)

## IPv4-Mapped IPv6 Address

- Resources such as CDN can be IPv6 addressed and accessed by IPv4 clients without traversing the BR

# Transition Technology Comparison

	MAP-T	NAT444	MAP-E lw4over6	DS-Lite
Subscriber Assigned Address Scalability	IPv6	RFC6598	IPv6	IPv6
Asymmetrical Forwarding	Yes	No	Yes	Yes
5-Tuple Visibility	Yes	Yes	No	No
CPE Software Required	Yes	No	Yes	Yes
Address Allocation and NAT Efficiency	Fixed	Dynamic	Fixed Dynamic	Dynamic
Support for IPv4 client access of IPv6 resources without gateway	Yes	No	No	No
Encapsulation Overhead	28 Bytes	None	40 Bytes	40 Bytes
Provisioning integration required with NAT or gateway/relay	No	Yes - PCP	No Yes	Yes - PCP

# MAP-T at Charter Communications

## MAP-T CPE

- Implementations from 5 commercial vendors
  - Hardware acceleration from 3
- Support for base mapping rule from 16 mapping rules

## MAP-T BR

- Implementation from 3 commercial vendors
- 150 Gbps per 1U

## Provisioning

- Existing provisioning platform extended to support provisioning multiple mapping rules

## Early Field Trial Underway

# Resources

MAP-T RFC – RFC 7599

Cablelabs MAP Technical Report:

- <https://apps.cablelabs.com/specification/14954>

MAP Python Package:

- <https://github.com/ejordangottlieb/pyswmap>
  - Used for test traffic generation and validation
  - Basis for regulatory web service which resolves a public IPv4 address and port to a particular subscriber

# Questions/Comments

