CableLabs®
Technical Reports

Mapping of Address and Port (MAP) Technical Report

CL-TR-MAP-V01-160630

RELEASED

Notice

This CableLabs technical report is the result of a cooperative effort undertaken at the direction of Cable Television Laboratories, Inc. for the benefit of the cable industry and its customers. You may download, copy, distribute, and reference the documents herein only for the purpose of developing products or services in accordance with such documents, and educational use. Except as granted by CableLabs in a separate written license agreement, no license is granted to modify the documents herein (except via the Engineering Change process), or to use, copy, modify or distribute the documents for any other purpose.

This document may contain references to other documents not owned or controlled by CableLabs. Use and understanding of this document may require access to such other documents. Designing, manufacturing, distributing, using, selling, or servicing products, or providing services, based on this document may require intellectual property licenses from third parties for technology referenced in this document. To the extent this document contains or refers to documents of third parties, you agree to abide by the terms of any licenses associated with such third-party documents, including open source licenses, if any.

© Cable Television Laboratories, Inc. 2016

DISCLAIMER

This document is furnished on an "AS IS" basis and neither CableLabs nor its members provides any representation or warranty, express or implied, regarding the accuracy, completeness, noninfringement, or fitness for a particular purpose of this document, or any document referenced herein. Any use or reliance on the information or opinion in this document is at the risk of the user, and CableLabs and its members shall not be liable for any damage or injury incurred by any person arising out of the completeness, accuracy, or utility of any information or opinion contained in the document.

CableLabs reserves the right to revise this document for any reason including, but not limited to, changes in laws, regulations, or standards promulgated by various entities, technology advances, or changes in equipment design, manufacturing techniques, or operating procedures described, or referred to, herein.

This document is not to be construed to suggest that any company modify or change any of its products or procedures, nor does this document represent a commitment by CableLabs or any of its members to purchase any product whether or not it meets the characteristics described in the document. Unless granted in a separate written agreement from CableLabs, nothing contained herein shall be construed to confer any license or right to any intellectual property. This document is not to be construed as an endorsement of any product or company or as the adoption or promulgation of any guidelines, standards, or recommendations.

Document Status Sheet

Document Control Number:	CL-TR-MAP-V01-160630			
Document Title:	Mapping of Address and Port (MAP) Technical Report			
Revision History:	V01 – 06/30/2016			
Date:	June 30, 2016			
Status:	~~Work in Progress~~	~~Draft~~	**Released**	~~Closed~~
Distribution Restrictions:	~~Author Only~~	~~CL/Member~~	~~CL/ Member/ Vendor~~	**Public**

Trademarks

CableLabs® is a registered trademark of Cable Television Laboratories, Inc. Other CableLabs marks are listed at http://www.cablelabs.com/specs/certification/trademarks/. All other marks are the property of their respective owners.

Contents

Figures

Tables

This page left blank intentionally.

1 SCOPE

1.1 Introduction and Purpose

The continued reliance on IPv4 has driven a need for co-existence technologies that allow IPv4-only devices to continue to function on IPv6-only networks. Despite the Regional Internet Registries running out of assignable IPv4 address blocks and the high cost and lack of availability of blocks from secondary market brokers, adoption of IPv6 by consumer products has been slow. As a result, there remains a need to support IPv4-only devices on an IPv6-only network now and for the foreseeable future.

The need for technologies that extend the useful life of IPv4, or that provide for IPv4/IPv6 co-existence, is not new. Today, the most widely deployed co-existence technologies are NAT444 and Dual Stack-Lite. Thus far, these have served as the best options for extending the life of IPv4 as network operators continue to deploy IPv6 in their networks.

However, while these technologies solve some immediate problems, they are not without certain limitations. For example, the finite number of unique source ports affects the number of subscribers that can be associated with a single public IPv4 address. This affects both the scaling and the customer experience. Tunneled IPv6 traffic also has its limitations. The available system resources and the design of the system limits the number of stateful translations that a DS-Lite CGN (Carrier Grade NAT) can support. Ultimately, the tradeoff for deploying these technologies result in service-affecting network conditions that impact the user experience. Both latency and jitter are known to be impacted when deploying these address-sharing and stateful translation co-existence technologies.

Mapping of Address and Port (MAP) offers a compelling alternative to the aforementioned co-existence technologies for a number of reasons. MAP provides two mechanisms for processing IPv4 traffic originating from a subscriber home network, with each approach having benefits and drawbacks. IPv4 traffic is either encapsulated, as in the case of MAP-E, or translated, in the case of MAP-T, into IPv6 by devices at the customer edge and service provider network. MAP provides a solution that is stateless because the translation or encapsulation functions occurs at the Customer Edge using rules that are shared by both the CE and the Border Relay (BR). As such, the BR forwards IPv6 traffic that was translated or encapsulated from IPv4 statelessly. The lack of state information means that compared to NAT444 or DS-Lite, MAP provides a solution that better aligns with the scaling demands of large service providers. Ultimately, this characteristic enables a better user experience and is transparent to the subscriber.

This technical report explains the advantages of MAP, and discusses the differences between MAP-E and MAP-T modes of operation. Additionally, extensive information is included on the provisioning of MAP and deployment models for MAP in cable networks. This technical report also contains in-depth information on deployment considerations for service providers, as well as a technical comparison of MAP-E and MAP-T to aid in the selection process.

2 INFORMATIVE REFERENCES

This technical report uses the following informative references.

[eRouter] IPv4 and IPv6 eRouter Specification, CM-SP-eRouter-I18-160317, March 17, 2016, Cable Television Laboratories, Inc.

[RFC 1918] IETF RFC 1918, Address Allocation for Private Internets, Y. Rekhter, B. Moskowitz, D. Karrenberg, G.J. deGroot, E. Lear, February 1996.

[RFC 2473] IETF RFC 2473, Generic Packet Tunneling In IPv6 Specification, A. Conta, S. Deering, December 1998

[RFC 3315] IETF RFC 3315, Dynamic Host Configuration Protocol for IPv6 (DHCPv6), R. Droms, J. Bound, B. Volz, T. Lemon, C. Perkins, M. Carney, July 2003.

[RFC 3633] IETF RFC 3633, IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6, O. Troan, R. Droms, December 2003.

[RFC 6052] ETF RFC 6052, IPv6 Addressing of IPv4/IPv6 Translators, C. Bao, C. Huitema, M. Bagnulo, M. Boucadair, X. Li, October 2010.

[RFC 6127] IETF RFC 6127, IPv4 Run-Out and IPv4-IPv6 Co-Existence Scenarios, J. Arkko, M. Townsley, May 2011.

[RFC 6145] IETF RFC 6145, IP/ICMP Translation Algorithm, X. Li, C. Bao, F. Baker, April 2011.

[RFC 6333] IETF RFC 6333, Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion, A. Durand, R. Droms, J. Woodyatt, Y. Lee, August 2011.

[RFC 7597] IETF RFC 7597, Mapping of Address and Port With Encapsulation (MAP-E), O. Troan, W. Dec, X. Li, C. Bao, S. Matsushima, T. Murakami, T. Taylor, July 2015.

[RFC 7598] IETF RFC 7598, DHCPv6 Options for Configuration of Softwire Address and Port-Mapped Clients, T. Mrugalski, O. Troan, I. Farrer, S. Perreault, W. Dec, C. Bao, L. Yeh, X. Deng, July 2015.

[RFC 7599] IETF RFC 7599, Mapping of Address and Port Using Translation (MAP-T), X. Li, C. Bao, W. Dec, O. Troan, S. Matsushima, T. Murakami, July 2015.

[Softwire] Charter-ietf-softwire-05, https://datatracker.ietf.org/doc/charter-ietf-softwire/

[sRouter] Access Network Independent Standalone Router Specification, CL-SP-sRouter-I01-160317, March 17, 2016, Cable Television Laboratories, Inc.

3 TERMS AND DEFINITIONS

This document uses the following terms:

Anycast	A network addressing and routing protocol where messages are directed to a node that is the closest point in the network.
CE Router	Customer edge router. In a cable network, this typically means a standalone router connected to a cable modem.
Customer Edge	The demark point between the service provider network and the subscriber home network.
Deep Packet Inspection	A form of packet processing in which not only are the headers parsed, but the payload octets as well. Such parsing is used to identify the traffic more accurately than might be possible relying upon the headers alone.
Delegated Prefix	An IPv6 prefix assigned to the WAN interface of a customer edge device. It is used to assign IPv6 addresses to clients on the LAN network, or to sub-delegate smaller sized prefixes to internal routers behind the edge router.
Dual Stack-Lite	An IPv4/IPv6 co-existence technology that tunnels IPv4 traffic over IPv6.
eRouter	An eSAFE device that is implemented in conjunction with an embedded DOCSIS cable modem.
Full mesh network	A network topology in which each node on a network is connected to all other nodes in the network.
Hub and Spoke	A network topology in which all nodes are connected to a central hub. Also known as a star network.
MAP-E Agent	An application that acts on behalf of another device that supports MAP-E and fulfills the role of a BR for the purpose of enforcing configured MAP-E encapsulation rules.
Multicast	A one-to-many communication wherein a message is sent to a plurality of recipients, in which the traffic is replicated at layer 3.
NAPT	An IPv4 address extension technique in which multiple private host IP addresses and port numbers are mapped to a single public IP address.
NAT444	An IPv4 address extension technology that implies two IPv4 network address translations are completed before traffic is forwarded to the public Internet. Typically, this means translation is done at the customer edge device from the LAN to the WAN, and again at the Carrier Grade NAT deployed by the service provider.
Partial mesh network	A network topology in which all nodes are not interconnected, and only specific nodes are interconnected for communication and/or redundancy purposes.
Unicast	A one-to-one communication between a pair of end nodes.

4 ABBREVIATIONS AND ACRONYMS

This document uses the following abbreviations:

AS Autonomous System
BGP Border Gateway Protocol
BMR Base Mapping Rule
BR Border Relay
CE Customer Edge
CGN Carrier Grade NAT
CPE Customer Premise Equipment
DMR Default Mapping Rule
DNS Domain Name System
DPI Deep Packet Inspection
EA Embedded Address
ECMP Equal Cost Multi-Path
FMR Forward Mapping Rule
FQDN Fully Qualified Domain Name
GUA Global Unicast Address
IA_PD Identity Association for Prefix Delegation
ICMP Internet Control Message Protocol
IGP Interior Gateway Protocol
IPv4aaS IPv4 as a Service
MAP-E Mapping of Address and Port (Encapsulation)
MAP-T Mapping of Address and Port (Translation)
MTU Maximum Transmission Unit
NAPT Network Address and Port Translation
NAT Network Address Translation
NAT44 Network Address Translation 44
NAT444 Network Address Translation 444
OID Object Identifier
ORO Option Request Option (DHCP)
PMTUD Path Maximum Transmission Unit Discovery
PPPoE Point-to-Point Protocol over Ethernet
PSID Port Set Identifier
SIIT Stateless IP/ICMP Translation Algorithm
SP Service Provider

5 TECHNOLOGY OVERVIEW

5.1 Overview of MAP

Mapping of Address and Port (MAP) is an IPv6 transition technology that allows IPv4 addresses to be translated or encapsulated into IPv6 without the need for stateful translation on the Service Provider's network. This enables Service Providers to use simpler, more efficient IPv6-only networks to support IPv4-only devices on the subscriber's network without negatively influencing the user experience. By comparison, NAT444 merely preserves existing IPv4 functionality and does not provide useful progress towards the desired end goal of IPv6-only networking.

MAP defines two transport modes. The MAP-E mode uses encapsulation for transport of IPv4 traffic in an IPv6 header and follows encapsulation rules defined in [RFC 2473]. The MAP-T mode uses the stateless IP/ICMP translation algorithm (SIIT) to convert IPv4 headers into IPv6 headers and follows stateless translation rules defined in [RFC 6145].

MAP-E and MAP-T offer distinct advantages over other co-existence technologies, such as NAT444/NAT44 or Dual-Stack Lite. The primary advantage is that translation or de-encapsulation at the Border Relay (BR) is accomplished through stateless mapping of an IPv4 address into an IPv6 prefix. The critical difference between the MAP methodology and traditional translation technologies, such as NAT44, is that MAP assigns a port range to each of the CPEs that share the same public IPv4 address. This allows for unique representation of the address and port range combination that is then translated or encapsulated in the IPv6 header. In this way, IPv6 route aggregation logic is used to direct packets in a stateless fashion, rather than a relying upon per-subscriber or per-connection criteria to determine state.

MAP provides several other advantages over the aforementioned co-existence technologies. For instance, MAP back office and provisioning requirements are closely aligned with those of a native IPv6 deployment. For advanced CE capabilities, unlike competing technologies, additional protocols are not required between the CE and service provider NAT function. For example, Port Control Protocol (PCP) is required by DS-Lite and NAT444 in order to support dynamic and static port forwarding. MAP also provides a nearly stateless mode of operation that avoids the resource intensive creation and maintenance of stateful translations.

5.2 MAP Domains

A MAP domain can be defined as one or more Border Relay (BR) nodes that operate within a defined administrative boundary (the domain) that may coincide with a geographical region or market that contains a population of Customer Edge (CE) devices. CE devices with MAP capabilities encapsulate or translate the IPv4 traffic that egresses the subscriber's home network. Using pre-defined and provisioned mapping rules, the BR forwards this traffic to the IPv4 Internet. Public IPv4 traffic destined for the customer network is translated or encapsulated in the same way at the BR, and is again converted to IPv4 at the ingress path of the CE.

MAP domains are defined through a set of standardized parameters that are shared between the BR and a given population of CE devices assigned to the domain. These parameters are configured on the BR via the CLI, XML schema or some other means. For CE devices, configurations may be established statically, via a downloaded configuration file, or through option encodings conveyed via DHCPv6 during address acquisition.

All nodes in the MAP domain share the following set of common parameters:

- One or more Mapping Rules consisting of:
  - **Rule IPv6 prefix**: The parent IPv6 prefix assigned to a particular mapping rule.
  - **Rule IPv4 prefix**: The pool of available IPv4 addresses associated with a particular mapping rule. The parameter is conveyed as a single IPv4 prefix and a prefix length.
  - **Embedded Address (EA) bits**: A set of bits in the end-user's assigned IPv6 prefix (offset from the Rule IPv6 prefix) which represent their assigned IPv4 address. When multiple subscribers share an IPv4 address, these bits include an identifier for the unique set of layer-4 ports that a subscriber may use.
  - **PSID offset**: Provides the means to define the beginning of the assigned layer-4 port range.

- A transport mode defined IPv6 address or prefix used to convey a domain's BR:
  - **MAP-E BR Address**: An IPv6 address representing the tunnel termination point of the MAP-E BR.
  - **MAP-T BR Prefix**: An IPv6 prefix assigned to a MAP-T BR whose unicast addresses represents all possible IPv4 destinations. The Default Mapping Rule (DMR) may be used to describe this prefix.

The usage of these parameters is explained in further detail in this document.

As noted previously, a MAP domain can contain one or more mapping rules. The mapping rule selected by a CE to determine its dedicated IPv4 addressing or shared address and port set assignment is called the Basic Mapping Rule (BMR). This rule is identified from the list of mapping rules and an IPv6 user prefix supplied during the provisioning process. Although multiple BMRs are possible, a BMR can only match a single IPv6 user prefix. In contrast to the CE, the BR uses all domain defined mapping rules to support the forwarding of traffic to and from the CE.

5.3 IPv4 as a Service (IPv4aaS)

IPv4 as a Service (IPv4aaS) is the use of IPv6 as a transport to carry IPv4 communications. In the context of this document, a CE accomplishes IPv4 communication without the assignment of a public IPv4 address on its WAN for use with IPv4 Network Address Translation (NAT). The customer premise CE router is provisioned with an IPv6 address for its WAN and IPv6 delegated prefix for its LAN. The LAN of the customer premise continues to support dual-stack operation where the LAN IPv4 address utilizes [RFC 1918] address blocks and the IPv6 prefix is subdivided and assigned in the manner described in [eRouter] and [sRouter].

MAP is the key enabling technology for IPv4aaS for both of the following modes of operation:

- One-to-one IPv4 address translation, which is the allocation of a single, dedicated IPv4 address to a customer premise with no IPv4 address sharing across subscribers. As with dual stack deployments, the CE router will employ NAPT to share the IPv4 address across devices. In this mode, MAP is used solely as a mechanism to transport IPv4 traffic through the IPv6 domain to a centralized IPv4 forwarding agent.
- Many-to-one IPv4 address translation, which is the multiplexing of a single publicly routable IPv4 address across one or more customer premises. The multiplexing of several subscribers with a relatively smaller number of public IPv4 addresses is typically referred to as Carrier Grade NAT (CGNAT or CGN). IPv4aaS in this context utilizes IPv6 as a transport to carry the customer premise IPv4 traffic to an IPv4 forwarding agent in conjunction with the use of CGN. CGN in this case is when a customer premise has been allocated a subset of available transport protocol layer ports mapped to a single publicly routable IPv4 address in order to uniquely identify that customer's flow. More than one public IPv4 address can be used in order to increase scaling, but there is no direct correlation between the number of subscribers and the number of public IPv4 addresses.

While both of the above approaches are valid modes of operation, the use of one-to-one translation does not allow adopters of IPv4aaS to maximize their current IPv4 allocations. As available pools of IPv4 address space deplete, the ability to multiplex IPv4 using CGN techniques is critical, until such time as IPv6 emerges as the dominant protocol.

6 REFERENCE ARCHITECTURE

Figure 1 is a graphical representation of the MAP reference architecture.

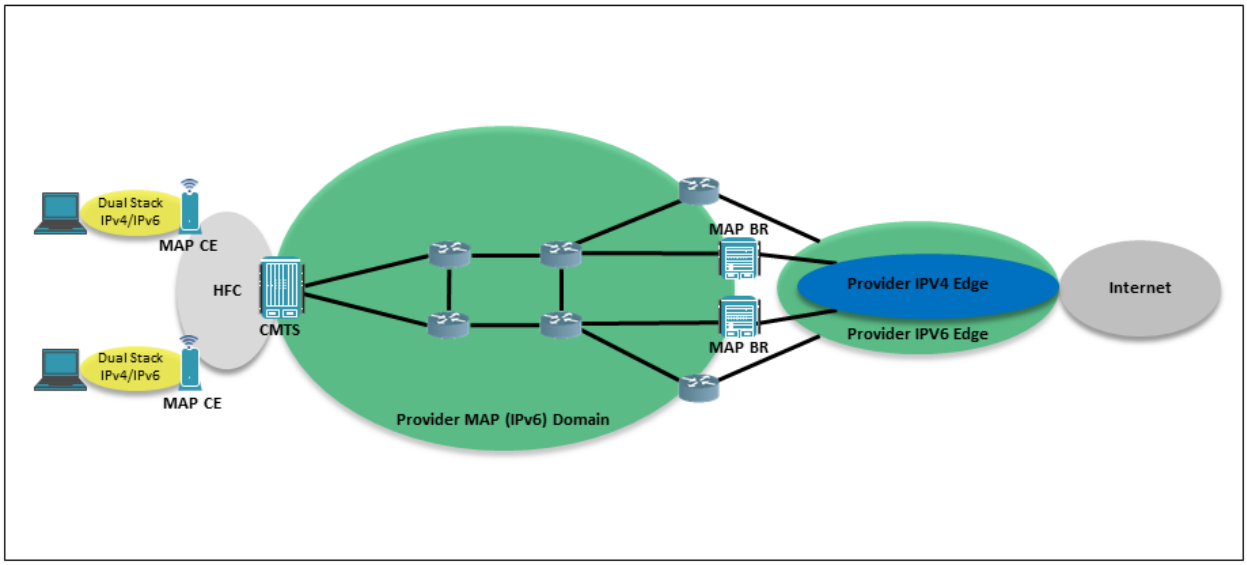

Figure 1 - MAP Reference Architecture

6.1 Border Relay

The BR typically resides at the northbound edge of the Service Provider network. However, placement of the BR at the northbound edge is not always optimal. Please refer to Section 7 for other considerations related to deployment of the BR for variable sized MAP domains.

The primary function of the BR is to route traffic to and from the public Internet. For packets that are natively IPv6, these packets are routed to the public IPv6 Internet. For packets that are IPv4 and have been encapsulated or translated into IPv6 at the customer edge, these packets are translated or de-capsulated at the BR and then routed to the public IPv4 Internet.

IPv4 internet traffic that is destined for the SP's MAP subscribers is routed through the BR. The BR will then encapsulate or translate these packets using MAP-E or MAP-T algorithms so that they can be transported as IPv6 packets over the SP's IPv6-only network. When these packets are routed to the customer edge device, they are translated or de-capsulated back to IPv4 before being routed to the subscriber's LAN.

BRs residing in the same MAP domain share the same set of domain parameters as the CE devices that are provisioned in that domain. A MAP domain may contain several Border Relays to increase scalability and provide high availability. Unlike the extensive clustering requirements of stateful approaches, MAP uses the well-established strategies of *anycast* (MAP-E) or IPv6 *prefix origination* (MAP-T) to extend the per-domain BR function across multiple nodes. These two approaches are described in further detail in the sections that follow.

6.1.1 Anycast

Anycast is a networking technique that leverages routing protocols to determine the closest interface or destination for packets that are sent to the anycast address. An anycast address is an IP address that is typically assigned to multiple interfaces on multiple nodes. The main advantage of anycast is that it enables the network to make routing decisions, such that packets destined for the anycast address are directed to the nearest device, as determined by router routing tables.

Since MAP-E BRs are identified by a specific IPv6 address, anycast is a viable mechanism for routing and scalability within the MAP-E BR infrastructure. This is not the case for MAP-T, however, as MAP-T BRs are

identifiable only through the BR prefix. As a consequence, MAP-T BRs utilize a translation technique in conjunction with a dynamic routing protocol for advertising themselves in the MAP-T domain. This mechanism is described in the next section of this document.

IPv6 anycast addressing simplifies the deployment and scaling of the MAP-E BR infrastructure. MAP-E BRs that are geographically dispersed can collectively advertise or announce themselves using a dynamic routing protocol, such as IGP or BGP. This approach allows MAP-E BRs to be deployed in a redundant fashion while allowing BR capacity to be augmented as capacity demands increase over time.

One or more MAP-E BR IPv6 addresses can be used when configuring MAP-E CEs either via DHCPv6 options or through SNMP MIB2 varbinds expressed as TLV 202.11 encodings present in the cable modem configuration file. For added flexibility, a DNS FQDN can also be used for one or more MAP-E BR IPv6 addresses to further simplify MAP-E BR discovery. The use of FQDNs for MAP-E is discussed in further detail in Appendix I of this document.

The use of IPv6 anycast and DNS FQDNs are both considered ideal techniques for MAP-E BR discovery given the stateless nature of MAP-E. This, of course, assumes that each MAP-E BR and CE share common rule sets for determining BR reachability.

6.1.2 BR Prefix Origination

The BR prefix represents the IPv4 destinations that are outside of a MAP-T domain. Using the algorithms defined in [RFC 6052], any IPv4 address may be embedded in the BR prefix to form a unique IPv6 address.

MAP-T BRs can collectively advertise themselves by announcing the BR prefix through either an IGP or BGP. The BR prefix is announced by all BRs in a given MAP-T domain. As with anycast, this allows for the same one-to-one-of-many relationship between a MAP-T CE and a set of MAP-T BRs.

The [RFC 6052] based translation technique also provides the benefit of accommodating additional granularity for both link and BR load balancing. Since MAP-T traffic consists of TCP/UDP IPv6 traffic, it supports 5-Tuple ECMP (flow based load balancing) both in-transit or across a set of BRs.

6.2 Customer Edge

Customer edge devices that support MAP enable the subscriber to connect IPv4-only or dual stack clients to the LAN, and allow them to communicate over an SP's IPv6-only networks, as if they were forwarding native IPv6.

In cable networks, the most common edge devices are the eRouter (a home router with an embedded cable modem), or a cable modem with a standalone router connected to it. A router or home gateway supporting MAP-E or MAP-T converts IPv4 packets originating on the LAN into IPv6 packets forwarded on the WAN. Packets destined for the public Internet are forwarded to a BR, where IPv4 packets are translated or de-encapsulated back to IPv4 before being routed to the public IPv4 Internet. CE devices residing in the same MAP domain as the BR share the same set of domain parameters.

6.3 Provisioning Methodologies

As defined in [eRouter], there are two mechanisms for configuring or provisioning MAP parameters on eRouter devices: DHCP option encodings and TLV configuration. Mechanisms for configuring standalone home gateways for MAP are out of scope for this technical report.

An eRouter supports both MAP-E and MAP-T provisioning using either provisioning mechanism, but cannot support the configuration of both modes simultaneously.

6.3.1 Provisioning of MAP via DHCPv6 Option Encodings

[RFC 7598] defines DHCPv6 options for configuring MAP-E and MAP-T clients. These DHCPv6 options provide critical information, such as SP IPv6 prefix, SP IPv4 prefix, BR IPv6 address, and unique rule sets required for participation in the provisioned MAP domain.

eRouters supporting MAP configuration via DHCP must request either the MAP-E Container Option (94) or the MAP-T Container Option (95) in the DHCPv6 Option Request Option (1) [RFC 3315] encoded in the DHCPv6

Solicit messages sent during provisioning. DHCPv6 servers responding to these requests will provide critical information for MAP-E or MAP-T configuration for the MAP domain. At a minimum, the following parameters are provided:

- S46 Rule Option (89)
- S46 Port Parameters Option (93)
- S46 BR Option (90) – MAP-E only
- S46 DMR Option (91) – MAP-T only

Examples of the proper DHCPv6 encodings for the MAP options are as follows:

- ***MAP-E***

  Client

-OPTION_ORO(6)

  -OPTION_S46_CONT_MAPE (94)

  Server

-OPTION_S46_CONT_MAPE (94)

  -OPTION_S46_RULE (89)

    -OPTION_S46_PORTPARAMS (93)

  -OPTION_S46_BR (90)

- ***MAP-T***

  Client

-OPTION_ORO(6)

  -OPTION_S46_CONT_MAPT (95)

  Server

-OPTION_S46_CONT_MAPT (95)

  -OPTION_S46_RULE (89)

    -OPTION_S46_PORTPARAMS (93)

  -OPTION_S46_DMR (91)

6.3.2 Provisioning of MAP via Configuration File Varbinds

[eRouter] defines MIBs for the provisioning of MAP-E and MAP-T parameters on eRouter devices. These MIBs are derived from a recent TR-181 data model and are defined in [eRouter]. This provisioning model does not include the end-user IPv6 prefix, which is acquired through DHCPv6 provisioning via assignment of an IA_PD.

The required data objects for MAP-E and MAP-T provisioning via TLV configuration file are shown in Table 1 and Table 2.

Attributes required to complete the activation of MAP are required as part of DOCSIS provisioning via the cable modem configuration file and are used in combination with the end-user IPv6 prefix to derive the full MAP configuration. The derived MAP configuration determines the specific MAP forwarding and potential NAPT behavior of the CE. For a detailed example of how MAP CE configuration parameters are determined, refer to Appendix I.

In the event that TLV encodings for both MAP-E and MAP-T are present in the configuration file, the eRouter must ignore these parameters and provision without any version of MAP being enabled.

6.3.2.1 MAP-E Explicit Provisioning

The data objects that appear in Table 1 are the required attributes for MAP-E TLV provisioning.

Table 1 - Required MIB Objects for TLV202.11 Provisioning of MAP-E

Object Name	Index	TruthValue	Value
clabMAPEnable		TruthValue	1
clabMAPDomainEnable	1	TruthValue	1
clabMAPDomainTransportMode	1	Integer	1
clabMAPDomainBRIPv6Prefix	1	InetAddressIPv6	IPv6 address
clabMAPDomainBRIPv6PrefixLen	1	InetAddressPrefixLength	/128
clabMAPDomainPSIDOffset	1	Unsigned32	0..16 (4 - default)
clabMAPDomainRowStatus	1	RowStatus	CreateAndGo(4)
clabMAPDomainRuleEnable	1,1	TruthValue	1
clabMAPDomainRuleIPv6Prefix	1,1	InetAddressIPv6	Rule IPv6 Prefix
clabMAPDomainRuleIPv6PrefixLen	1,1	InetAddressPrefixLength	Rule IPv6 Prefix Length
clabMAPDomainRuleIPv4Prefix	1,1	InetAddressIPv4	Rule IPv4 Prefix
clabMAPDomainRuleIPv4PrefixLen	1,1	InetAddressPrefixLength	Rule IPv4 Prefix Length
clabMAPDomainRuleEABitsLength	1,1	Unsigned32	0..48 (0 - default)
clabMAPDomainRuleIsFMR	1,1	TruthValue	0 - No (default); 1 - Yes
clabMAPDomainRuleRowStatus	1,1	RowStatus	CreateAndGo(4)
clabMAPDomainIfEnable	1,1	TruthValue	1
clabMAPDomainIfAlias	1,1	SnmpAdminString	Domain Interface handle
clabMAPDomainIfName	1,1	SnmpAdminString	Domain tunnel Interface Name
clabMAPDomainIfLowerLayers	1,1	SnmpAdminString	IfIndex values list
clabMAPDomainIfRowStatus	1,1	RowStatus	CreateAndGo(4)

6.3.2.2 MAP-T Explicit Provisioning

The data objects that appear in Table 2 are the required attributes for MAP-T TLV provisioning.

Table 2 - Required MIB Objects for TLV202.11 Provisioning of MAP-T

Object Name	Index	TruthValue	Value
clabMAPEnable		TruthValue	1
clabMAPDomainEnable	1	TruthValue	1
clabMAPDomainTransportMode	1	Integer	2
clabMAPDomainBRIPv6Prefix	1	InetAddressIPv6	BR IPv6 Prefix
clabMAPDomainBRIPv6PrefixLen	1	InetAddressPrefixLength	BR Prefix length
clabMAPDomainPSIDOffset	1	Unsigned32	0..16 (4 - default)

Object Name	Index	TruthValue	Value
clabMAPDomainRowStatus	1	RowStatus	CreateAndGo(4)
clabMAPDomainRuleEnable	1,1	TruthValue	1
clabMAPDomainRuleIPv6Prefix	1,1	InetAddressIPv6	Rule IPv6 Prefix
clabMAPDomainRuleIPv6PrefixLen	1,1	InetAddressPrefixLength	Rule IPv6 Prefix Length
clabMAPDomainRuleIPv4Prefix	1,1	InetAddressIPv4	Rule IPv4 Prefix
clabMAPDomainRuleIPv4PrefixLen	1,1	InetAddressPrefixLength	Rule IPv4 Prefix Length
clabMAPDomainRuleEABitsLength	1,1	Unsigned32	0..48 (0 - default)
clabMAPDomainRuleIsFMR	1,1	TruthValue	0 - No (default); 1 - Yes
clabMAPDomainRuleRowStatus	1,1	RowStatus	CreateAndGo(4)
clabMAPDomainIfEnable	1,1	TruthValue	1
clabMAPDomainIfAlias	1,1	SnmpAdminString	Domain If handle
clabMAPDomainIfName	1,1	SnmpAdminString	Domain tunnel If Name
clabMAPDomainIfLowerLayers	1,1	SnmpAdminString	IfIndex values list
clabMAPDomainIfRowStatus	1,1	RowStatus	CreateAndGo(4)

6.4 MAP Forwarding Behavior and Mapping Rules

Mapping rules determine the forwarding behavior to destinations both within, and external to, a particular MAP domain. The forwarding behaviors are defined by the following two MAP topology models.

6.4.1 Hub and Spoke (Default)

In the model shown in Figure 2, all MAP traffic must traverse the BR. This includes MAP traffic between two CEs belonging to the same MAP domain.

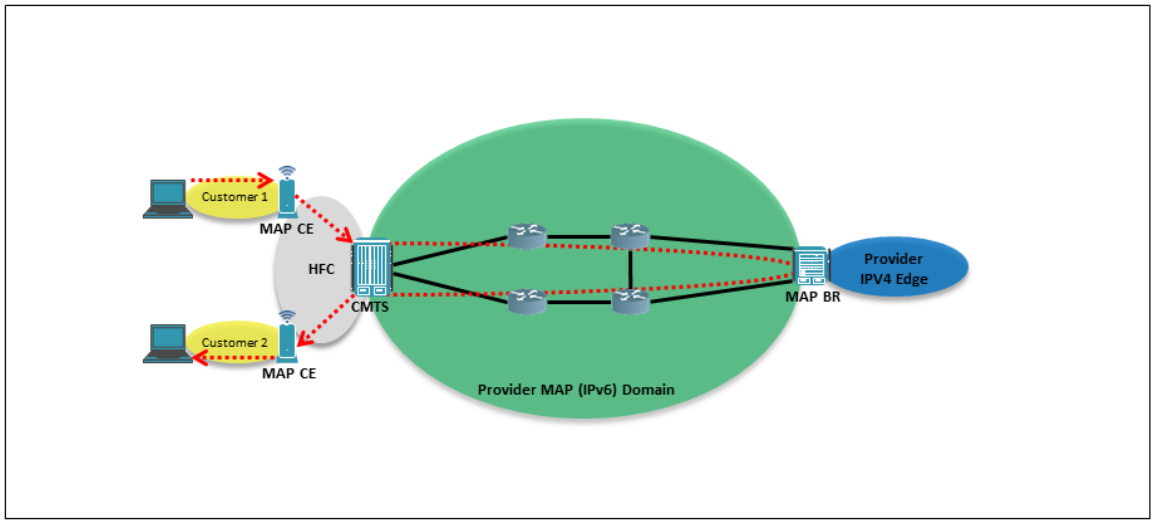

Figure 2 - MAP Hub and Spoke Model

6.4.2 Mesh (Optimization)

Mapping rules can be configured to permit MAP traffic to be forwarded between CEs without an intervening BR. This is accomplished through the definition of a Forward Mapping Rule (FMR). For traffic within a particular MAP domain, the Base Mapping Rule (BMR) may be configured to be interpreted as an FMR in order to support mesh mode operation.

These two models may be combined to support a hybrid environment where traffic within a MAP domain is mesh mode forwarded, while IPv4 destinations outside the MAP domain are forwarded through the BR as shown in Figure 3.

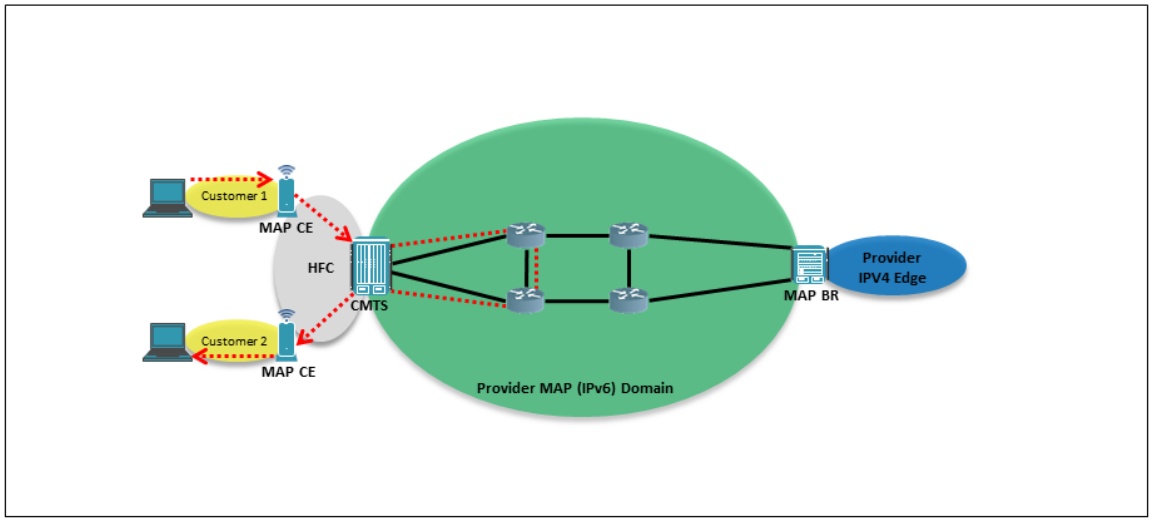

Figure 3 - MAP Mesh Mode Operation Model

Any LAN originated packet arriving at the subscriber's MAP CE whose target IPv4 address is contained within an FMR enabled mapping rule will be forwarded directly to the target MAP CE IPv6 address. This is possible because mesh mode leverages the characteristic that each MAP CE contains a unique MAP IPv6 address that can be algorithmically matched to a target IPv4 address and set of layer-4 ports. This allows both MAP-E and MAP-T to support mesh operation when a mapping rule is configured as an FMR.

7 DEPLOYMENT CONSIDERATIONS

7.1 BR Placement Considerations

When designing a MAP deployment, a Service Provider must consider the tradeoffs associated with the placement of the BR within the topology. This section provides an analysis of long versus short layer-3 path distances between a MAP BR and MAP CE.

7.1.1 BR Deployed at the Edge of a Large MAP Network Domain

Large MAP domains are characterized by BRs that are placed at or near the IPv4 Internet Autonomous Systems (AS) peering boundary and at a distance from the MAP CEs. This type of BR placement is well suited to the IPv4aaS deployment paradigm, in that it attempts to limit the extent of IPv4 deployment.

In this deployment model, a MAP path through a BR can steer traffic away from a particular IPv4 destination. Figure 4 illustrates one such scenario where the MAP path is not always the most efficient path.

Customer traffic destined for **AS1** will take the path through the best connected BR. The path through **MAP BR #1** happens to be the most efficient path to **AS1**. Traffic from the same customer with a destination of **AS2** will also take the path through the best-connected BR. In this second case, a better path to **AS2** can be found by taking the path through **MAP BR #2**.

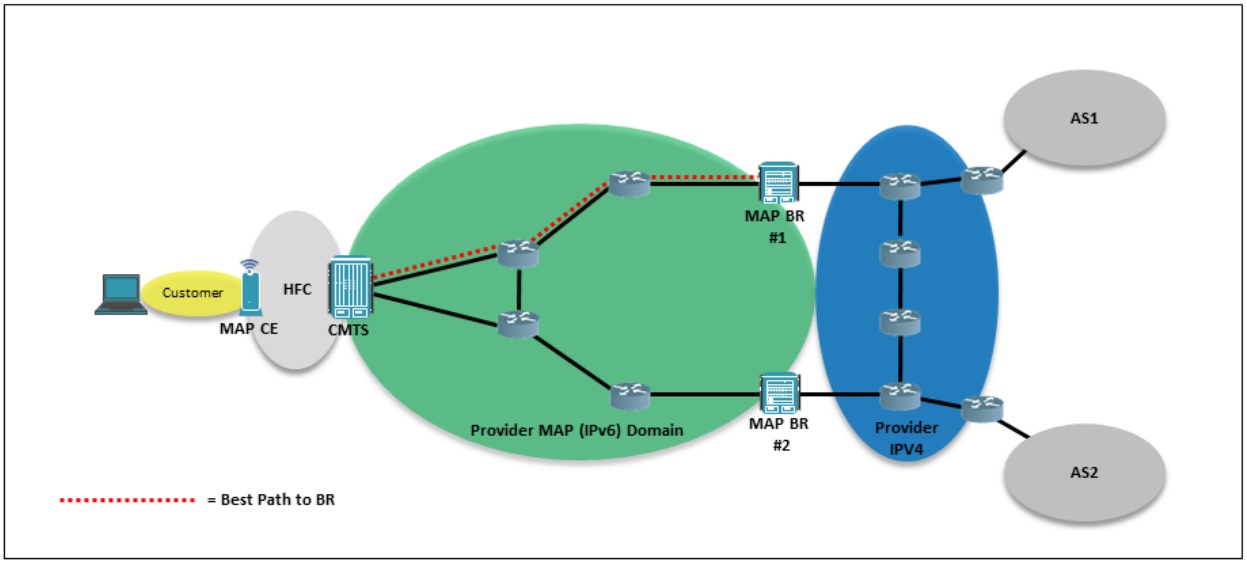

Figure 4 - Border Relay Edge Deployment for Large MAP Domain

In a hub and spoke environment, placing the BRs at a distance from the MAP CEs may also increase the path for peer-to-peer customer traffic within the provider's network. The resultant deleterious effects on legacy IPv4 performance for traffic between two customers in the same MAP network domain may serve as encouragement for migration to an IPv6-based solution.

7.1.2 BR Deployed at the Edge of a Small MAP Network Domain

Small MAP domains are characterized by BRs that are placed near the MAP CE devices. In a hub and spoke environment, placing the BR function near the MAP CEs will result in optimal path selection for MAP CE traffic with IPv4 destinations outside of the MAP domain. In this deployment model, the path taken to the best-connected BR has very little impact on the IPv4 traffic engineering policies for IPv4 Internet destinations. This deployment strategy also provides the benefit of optimizing the forwarding path for IPv4 peer-to-peer applications and service provider IPv4 applications residing between the BR and MAP subscribers.

The downside to this approach is an increase in the number of BRs required. Additionally, minimizing the distance between the BR and MAP CE reduces the effectiveness of BR function aggregation. Another cost is the impact on transitioning IPv4 to a small, residual protocol on the Service Provider's network. Decreasing the MAP IPv6 domain size comes at the expense of a larger IPv4 footprint and the associated operational demands. Figure 5 illustrates the topics covered in this section.

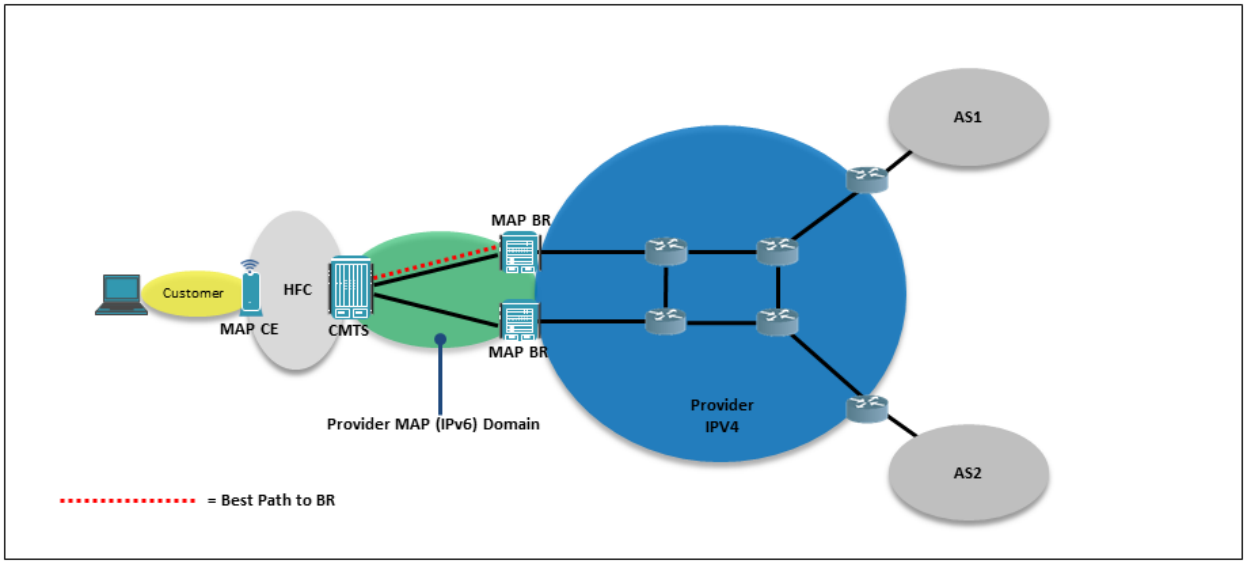

Figure 5 - Border Relay Edge Deployment for Small MAP Domains

7.1.3 Forwarding Path Optimization for CDN and Other Internal Service Provider Resources

For IPv4 resources residing on the service provider network, it is advantageous to place the BR close to the subscribers. If an IPv4 addressed internal resource such as CDN is deployed between the BR and the MAP subscriber, suboptimal routing paths will result as all IPv4 traffic requires traversal through the BR. As the distance between the BR and MAP subscriber and CDN increases, the efficiency of the path taken decreases.

In lieu of being able to place a BR close to the subscriber, there are two mechanisms for avoiding such a suboptimal path. In the case of MAP-T, an FMR or DMR based IPv4-mapped IPv6 address may be assigned to allow IPv4 addressed applications direct access to the CDN resource. In the case of MAP-E, an FMR may be used in conjunction with a MAP-E agent installed on the CDN servers addressed with a IPv4-mapped IPv6 address. In both cases, when the IPv6 addresses are advertised via a routing protocol the most optimal path will be utilized.

7.2 Provisioning Considerations

7.2.1 IPv6 End-user Prefix and DHCPv6 Prefix Delegation

[RFC 3633] defines a capability that can be used by a requesting router to set the IPv6 prefix field to zero (0) and the prefix-length field to a user-defined value. Routers that request a prefix length that is not equal to the expected IPv6 End-user prefix can break the MAP algorithms that utilize the EA bits. Therefore, a MAP CE should not request an IPv6 end-user prefix that is not equal to the prefix length obtained from the BMR.

When using MAP DHCPv6 provisioning, the server should enforce a prefix length compatible with the defined EA bit location regardless of the DHCPv6 prefix length hint. This policy can be applied when a MAP CE's DHCP client includes an ORO with options 94 or 95.

7.2.2 Explicitly Provisioned PSID Assignment

MAP supports the creation of a BMR where the EA length is zero and assignment of the PSID and PSID length occurs via the provisioning system. Using this approach, the IPv4 and optional transport layer port range is not

derived from the CE assigned IPv6 prefix. In theory, such a deployment can provide for a more efficient allocation of IPv4 addresses as un-provisioned systems will not consume IPv4 addresses and/or transport protocol ports. In practice, this capability limits the operational scalability of the MAP deployment and increases the complexity of the BR and provisioning system implementations. Given these constraints, we recommend avoiding this implementation strategy.

7.3 Fragmentation and Path MTU

The encapsulation or translation between the IPv4 and IPv6 address families directly affects how MAP BRs and CEs handle packet fragmentation, the forwarding of fragments, and Path MTU Discovery. These characteristics dictate several design considerations to ensure the reliability of the MAP domain.

7.3.1 BR Fragment Forwarding

In a "Many-to-One IPv4 Address Translation" environment, an IPv4 packet that is destined for the MAP subscriber and arriving at a BR from outside the MAP domain cannot be forwarded if it does not contain the destination IPv4 address and port number (or the identifier field for selected ICMP traffic). In the case of an IP fragment, this information is only contained in the initial fragment. A BR must either reassemble the packet, or if possible use cache-and-forwarding in order to perform the encapsulation or translation function properly.

Both techniques present a scenario where the BR does not operate in a stateless fashion. An environment with excessive legitimate and illegitimate fragment volume can impact the performance and stability of a BR. An example of an illegitimate fragment would be a DoS attack. Policing and other remediation tactics may be used to limit the impact of excessive fragment volumes.

7.3.2 MAP Domain Path MTU

Current broadband cable Internet offerings employ DOCSIS infrastructure that supports an IP MTU of 1500 bytes. This is consistent with most prevalent Internet IP MTU of 1500 bytes. Therefore, mechanisms that are in-place to handle native IPv4 or IPv6 packets larger than 1500 bytes are typically not needed. When a MAP CE or BR performs encapsulation or translation on a given IPv4 packet, it will increase its size (typically 40 bytes for MAP-E and 20 bytes for MAP-T) with a resulting packet that will not fit into a DOCSIS frame (DOCSIS 3.0 version or earlier).

An operator must consider the path MTU for all possible paths in a MAP domain. Operators that cannot depend on a well-managed MTU (such as those that utilize third-party networks) should base the MAP BR and CE tunnel MTU around the IPv6 minimum MTU of 1280 as calculated below:

- MAP-E: 1280 – 40 (encapsulation overhead) = 1240
- MAP-T: 1280 – 20 (translation overhead) – 8 (accommodate IPv6 fragment extension header) = 1252

By basing the tunnel MTU on the 1280 byte value, an operator is guaranteed that the packet will not require fragmentation as it traverses the MAP domain. This outcome is important, as fragmentation is not a permitted function on an intermediate IPv6 router.

Operators that can guarantee a minimum IP path MTU of 1500 bytes can set the MAP BR and CE tunnel MTU as follows:

- MAP-E: 1500 – 40 (encapsulation overhead) = 1460
- MAP-T: 1500 – 20 (translation overhead) – 8 (accommodate IPv6 fragment extension header) = 1472

Setting the appropriate MAP path MTU is also important in support of Path Maximum Transmission Unit Discovery (PMTUD), which may not operate properly if allowed to traverse the MAP domain.

In a BR implementation that is capable of looking deep enough into the layer 4 headers to adjust TCP MSS, a considerable amount of headaches may be eliminated by adjusting it downwards in the manner often done in an L2TP/PPPoE DSL environment. This **needs** to happen on the end that first sees a large TCP packet, so the BR end is fine for web surfing. Optimally, however, TCP MSS rewriting should happen on both the CE router and the BR. On the CE end, an ICMPv4 packet-too-big message upon receipt of a packet that the CE knows is too big to encapsulate

will suffice, since it can be reasonably assumed that there will be no ICMP filters between the CE router and the customer device.

7.3.3 BR Packet Fragmentation

The overhead associated with encapsulation or translation from IPv4, within or converted to an IPv6 header, may require the BR to perform packet fragmentation. A MAP destined IPv4 packet arriving at a BR from outside the MAP domain is required to be fragmented by the BR when all of the following conditions are met:

- The packet arrives with the DF bit set to zero,
- The packet is too large to transmit over the MAP domain PATH MTU.

Once again, traffic that must be fragmented, regardless of whether it is legitimate traffic, can lead to BR performance and stability issues.

7.4 Logging Considerations for Law Enforcement Requests

In many countries a Service Provider has a legal obligation to track subscribers by their assigned IP address. Traditional carrier deployments accommodate this requirement by logging the customer to IP address mapping as it is assigned through DHCP or other address provisioning mechanisms. When responding to subpoenas from law enforcement agencies and internal abuse reporting, a Service Provider can use the address log and the applicable time and date range to identify the correct subscriber.

MAP address provisioning systems provide IPv6 addresses and prefixes and do not explicitly assign an IPv4 address to a specific subscriber. Furthermore, when implementing many-to-one IPv4 address translation, the provider can no longer map a single subscriber to a single IPv4 address. Given these two factors, the address provisioning log cannot be used on its own to match reported IPv4 Internet traffic with the originating subscriber. Since this represents a departure from LE's traditional "map an IP address to a user" modus operandi, and indeed they may not have taken the time to preserve evidence of which ports were in use for traffic of interest, this will represent a substantial educational burden on the part of the service provider.

The combination of IPv6 prefix to IPv4 address mapping, as well as port-mapping algorithms, allows MAP to avoid additional logging requirements often associated with other Carrier Grade NAT solutions. A direct mapping of stored DHCPv6 (or other IPv6 prefix provisioning mechanism) prefix assignments to MAP-originated IPv4 Internet sessions is possible through programmatic techniques. Prior to the deployment of MAP, a software tool should be deployed with the Service Provider's configured MAP domain rules which supports the calculation of subscriber address mapping. Throughout the life of a MAP deployment, a Service Provider must synchronize the configured MAP network rules with those stored in the software tool. The software must also store MAP domain rules and their timelines in service in order to support historical queries of inactive mapping rules.

By following these recommendations, a provider can accommodate law enforcement and abuse requests while limiting the logging activity to the IPv6 end-user prefix allocation. Appendix II provides an example of associating a MAP subscriber's assigned IPv6 end-user prefix with an IPv4 Internet session.

7.5 MAP Domain Deployment Options

Several address allocation, MAP rules, and domain layout options are available to Service Providers in the planning stage of a MAP deployment. There are several factors a network operator must consider when planning a deployment including: The availability of contiguous IPv4 address blocks, MAP provisioning system capabilities, level of effort to deploy and manage, and alignment with existing aggregation schemes and traffic engineering requirements.

7.5.1 Single MAP Domain With a Single Mapping Rule

Some operators may have the flexibility to choose this option, which is the easiest to deploy and manage. The following items should be evaluated when considering this deployment strategy:

- The provider has a block of contiguous IPv4 addresses large enough to accommodate the number of users and the selected sharing ratio within the MAP domain.
- Customer prefixes contained in the Rule IPv6 prefix should be deployed as MAP users in large enough numbers to obtain a desired IPv4 address sharing ratio. This type of deployment should not be considered when the conversion of customers to MAP is going to occur over a long period of time.

The example below presents a single domain single mapping rule scenario that supports 65,536 users with a sharing ratio of 256 users per IPv4 address. In this scenario each customer receives a /60 IPv6 end-user prefix. The provisioning and BR configuration overhead is minimal as it consists of a single mapping rule and either a BR address or prefix/DMR depending on whether MAP-E or MAP-T mode is configured.

Map Domain Definition:

MAP Rule 1

- Rule IPv6 Prefix: 2001:db8::/44
- Rule IPv4 Prefix 192.0.2.0/24
- EA Length: 16

BR address or BR prefix

7.5.2 Single MAP Domain With Multiple Mapping Rules

Network operators with a requirement to deploy a single MAP domain with multiple blocks of IPv4 addresses can do so by defining multiple mapping rules. The resulting size of the provisioning and BR configuration will be significantly larger than the first deployment model example. Additionally, the MAP CEs will need to support the matching of their end-user IPv6 prefix to the correct Rule IPv6 prefix so they can identify the mapping rule that will be applied as a BMR.

The following example supports the same number of users and sharing ratio as the first example. Another shared characteristic is that the cumulative list of Rule IPv6 and Rule IPv4 definitions result in the 192.0.2.0/24 and 2001:db8::/44 prefixes used in the first example. For any given MAP domain, the use of parent prefixes is not a requirement. Furthermore, the prefix length, EA length, and other definable attributes do not need to match across the defined MAP rules.

Map Domain Definition:

MAP Rule 1

- Rule IPv6 Prefix: 2001:db8::/46
- Rule IPv4 Prefix 192.0.2.0/26
- EA Length: 14

MAP Rule 2

- Rule IPv6 Prefix: 2001:db8:4::/46
- Rule IPv4 Prefix 192.0.0.0/26
- EA Length: 14

MAP Rule 3

- Rule IPv6 Prefix: 2001:db8:8::/46
- Rule IPv4 Prefix 192.0.32.0/26
- EA Length: 14

MAP Rule 4

- Rule IPv6 Prefix: 2001:db8:c::/47
- Rule IPv4 Prefix 192.0.48.0/27
- EA Length: 13

MAP Rule 5

- Rule IPv6 Prefix: 2001:db8:e::/47
- Rule IPv4 Prefix 192.0.64.0/27
- EA Length: 13

BR address or BR prefix

7.5.3 Multiple MAP Domains With One or More Mapping Rules

The last MAP deployment option can support multiple blocks of IPv4 addresses with or without multiple MAP rules. The deployment of multiple domains requires the setting of a unique per domain BR address or BR prefix. In this environment, a particular MAP domain can be dedicated to single BR or a group of BRs. Barring implementation-specific restrictions, multiple domains may be configured on a per BR basis.

The example that follows defines multiple MAP domains in order to provide the same capabilities and prefix allocation configuration as the second example.

Map Domain Definition A:

MAP Rule 1

- Rule IPv6 Prefix: 2001:db8::/46
- Rule IPv4 Prefix 192.0.2.0/26
- EA Length: 14

BR address or BR prefix A

Map Domain Definition B: MAP Rule 1 • Rule IPv6 Prefix: 2001:db8:4::/46 • Rule IPv4 Prefix 192.0.0/26 • EA Length: 14 BR address or BR prefix B
Map Domain Definition C: MAP Rule 1 • Rule IPv6 Prefix: 2001:db8:8::/46 • Rule IPv4 Prefix 192.0.32.0/26 • EA Length: 14 BR address or BR prefix C
Map Domain Definition D: MAP Rule 1 • Rule IPv6 Prefix: 2001:db8:c::/47 • Rule IPv4 Prefix 192.0.48.0/27 • EA Length: 13 MAP Rule 2 • Rule IPv6 Prefix: 2001:db8:e::/47 • Rule IPv4 Prefix 192.0.64.0/27 • EA Length: 13 BR address or BR prefix D

8 COMPARISON OF MAP-E AND MAP-T

8.1 Technical Characteristics

The following table provides a summary of the technical characteristics of MAP-E and MAP-T, followed by further details.

Table 3 - MAP-E and MAP-T Technical Characteristics Summary

Characteristic/Capability	MAP-E	MAP-T
BR Reachability	BR IPv6 Address	BR Prefix Address
Layer-3 Transparency	Full	Partial
Header Overhead	~ 40 Bytes	~ 20 Bytes
Native IPv4-mapped IPv6 addresses	Requires MAP-E agent	Supported
Native edge filtering and classification	Requires DPI	Supported
CE implementation complexity	Simple	Moderate
Traffic Engineering	Not Supported	Supported
Equal Cost Multi-Path (ECMP)	1-Tuple	5-Tuple

- BR Reachability: This component defines the addressing methodology used to reach a BR within a Hub and Spoke deployment.
  - MAP-E: In this environment, MAP IPv4 traffic is encapsulated within an IPv6 header which defines the source and destination for a particular [Softwire] (BR and MAP CE). Figure 6 illustrates this concept.
  - MAP-T: The [Softwire] technique utilized in this scenario is header translation between IPv4 and IPv6. The IPv4 address for a destination outside the MAP domain is embedded in the BR prefix so that it is not lost during the header translation process. Figure 7 illustrates this concept.
- Layer-3 Transparency: A measure of the ability to retain all the values stored in the IPv4 header within a given MAP IPv6 domain transport mode.
  - MAP-E: Full transparency can be achieved as long as the operator or encapsulation implementation does not explicitly change the IPv4 header values.
  - MAP-T: A high level of layer-3 transparency is achieved by mapping most IPv4 header values to functionally equivalent IPv6 header values. The following IPv4 header values are not supported:
    - The combination of DF=1 and MF=1
    - IP Options
- Header Overhead: The number of bytes required to support encapsulation or translation and its impact to available payload.
  - MAP-E: Overhead is roughly 40 bytes
  - MAP-T: Overhead is roughly 20 bytes
- Native IPv4-mapped IPv6 Addresses: Provides the capability for a MAP subscriber IPv4 application to access an IPv6 based service directly without traversing the BR.
  - MAP-E: Requires a MAP-E agent installed on the application server or a MAP-E capable network node to provide this capability. This approach also requires an FMR to be provisioned on the MAP CE.
  - MAP-T: Support is built in and does not require agents to be installed on the target application servers. If the IPv4-mapped IPv6 address is selected from the provisioned DMR, no FMR is required.

- Native Edge Filtering and Classification: The ability to perform 5-Tuple packet filtering and QoS classification on MAP traffic at the service provider's customer edge network.
  - MAP-E: The ability to ignore the IPv6 header and process the encapsulated IPv4 header and transport header is typically not supported on service provider CMTS, router, and switching platforms. Therefore, classification and filtering based on 5-tuple values is typically not possible.
  - MAP-T: Traffic can be QoS classified and filtered on any platform that supports these functions in IPv6.
- CE implementation complexity: A measure of the level effort required to develop a MAP-E or MAP-T CE. Facilities such as provisioning and NAPT functions that are equivalent on the two transport modes are considered simple to implement and are the baseline for the evaluation.
  - MAP-E: No additional module development beyond the MAP baseline as [RFC 2473] based tunnel is available on most Linux based CE platforms.
  - MAP-T: The stateless translation module must be developed on any platform wishing to implement MAP-T.
- Traffic Engineering: The ability to alter the path taken through a service provider network for any particular MAP IPv4 destination that must traverse a BR.
  - MAP-E: For non-mesh MAP environments, the destination address is always the BR address. Traffic engineering for a particular destination is not possible.
  - MAP-T: Traffic engineering is possible as more specific IPv4-mapped IPv6 prefixes may be injected into a MAP domain for any BR prefix.
- ECMP: The ability to load balance MAP traffic over two or more equally weighted "best path" links.
  - MAP-E: The lack of 5-tuple visibility limits the load balancing algorithm to the CE MAP IPv6 address.
  - MAP-T: Per flow load balancing is supported on any network that supports 5-tuple IPv6 ECMP.

8.2 Packet Walk Comparison

The difference between MAP-E and MAP-T is the mechanism by which traffic is transported over an IPv6 network domain. This traffic can consist of MAP traffic that stays within the MAP domain such as CE to CE, or traffic that leaves the domain with the help of a BR. Figure 6 and Figure 7 below demonstrate the difference between a MAP CE performing MAP-E encapsulation and MAP-T stateless translation in a hub and spoke environment. In this scenario a workstation with the IPv4 address of 192.168.1.101 transmits a packet to a server outside the MAP domain with an address of 1.1.1.1. The differences between MAP-E and MAP-T packet processing within the CE is illustrated in the delta between the two versions of packet #3 in Figure 6 and Figure 7.

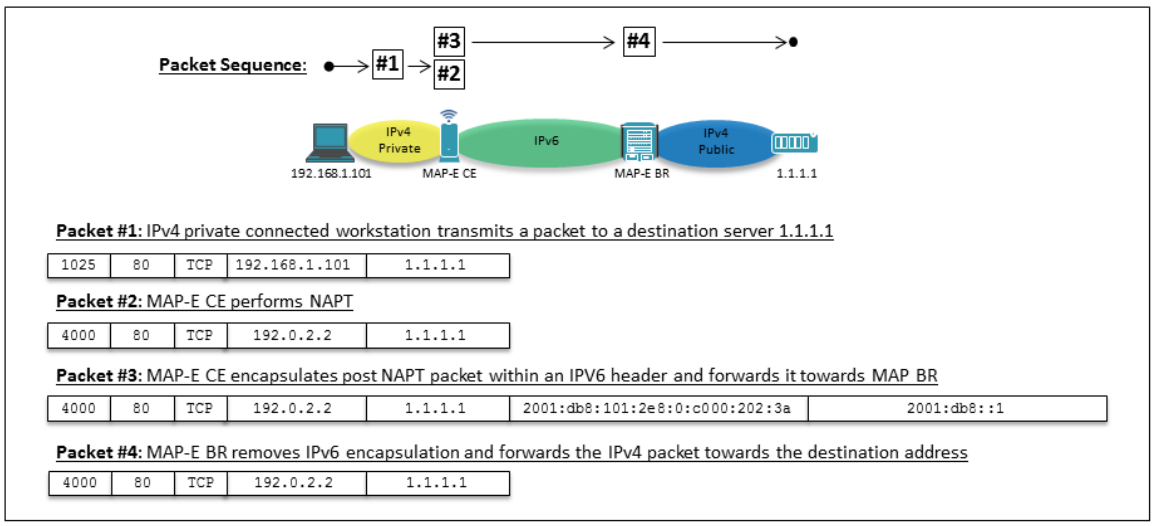

Figure 6 - MAP-E Encapsulation Function

When the MAP-E encapsulation function is performed an IPv6 header is applied to packet #2 with a source address of the CE MAP address (2001:db8:101:2e8:0:c000:202:3a) and the destination address of the BR (2001:db8::1).

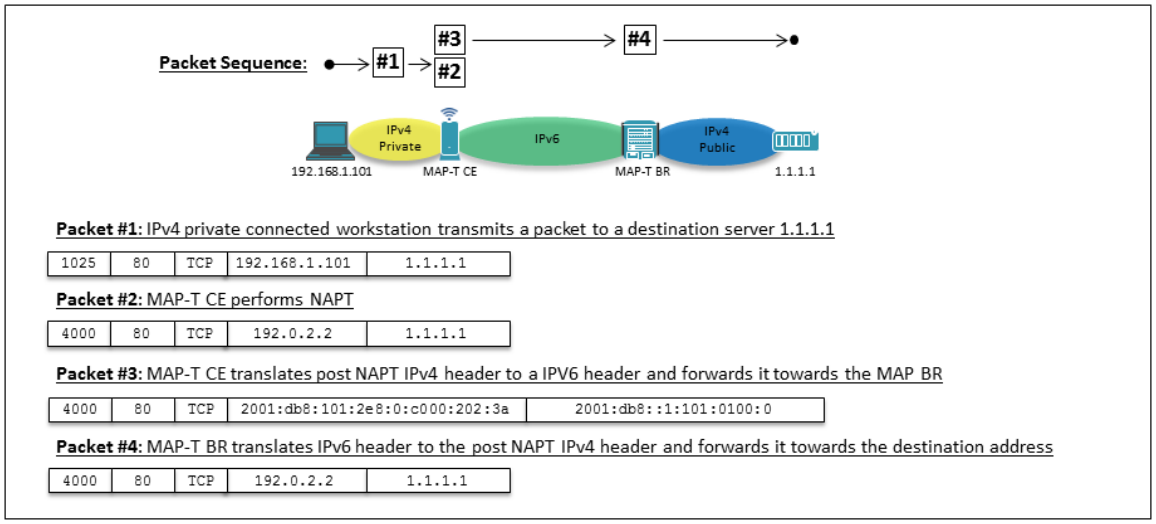

Figure 7 - MAP-T Translation Function

When the MAP-T stateless translation function is performed the IPv4 header is replaced with an IPv6 header. The source is changed to the CE MAP address (2001:db8:101:2e8:0:c000:202:3a) while the destination is the combination of the BR prefix (2001:db8::/64) and the destination IPv4 address which produces a [RFC 6145] complaint address (2001:db8::1:101:0100:0).

Appendix I Using the Base Mapping Rule to Configure a MAP CE

A MAP domain contains one or more mapping rules that govern the encapsulation or translation behavior of MAP nodes. During the address and MAP provisioning process, a MAP CE uses the mapping rules and the assigned end-user IPv6 prefix(es) to determine its IPv4 address and optionally its layer-4 port allocation (in support of "Many-to-one IPv4 sharing" scheme). A rule used in this fashion is defined as a Base Mapping Rule (BMR).

This section provides an example of BMR processing by a MAP CE and the resulting configuration. This example does not cover the "Explicitly Provisioned PSID Assignment" scenario described in Section 7.2.2.

The MAP domain definition for this example is a follows:

- Mapping Rule #1
  - Rule IPv4 Prefix: 192.0.2.0/24
  - Rule IPv6 Prefix: 2001:db8:9500::/40
  - EA Bits Length: 16
- PSID Offset: 6

Using the domain definition and the mapping rule contained within the values may be derived:

1. End-user IPv6 prefix length: A Rule IPv6 Prefix length of 40-bits and an EA length of 16-bits results in an End-user IPv6 prefix of 56-bits. Therefore the IPv6 address provisioning system will need to allocate a unique /56 prefix to each MAP CE from the parent Rule IPv6 prefix of 2001:db8:9500::/40.
2. PSID bits and "Many-to-One IPv4" sharing ratio: All possible IPv4 addresses for Rule IPv4 Prefix value of 192.0.2.0/24 can be represented by an 8-bit index for the 256 possible IPv4 addresses. Subtracting these 8-bits from the EA bits length value of 16 results in 8-bits left for PSID identification. Once again, since 8-bits can represent 256 values, the PSID attribute can now support 256 MAP CEs (or customers) for each IPv4 address.

In summary the following items have been calculated from the example domain mapping rule:

- End-user IPv6 prefix length: 56-bits
- EA Bits
  - IPv4 address index bits: The 8 high order bits following the first 40-bits of the Rule IPv6 prefix
  - PSID bits: The 8-bits following the IPv4 address index bits

The provisioning sequence provides the MAP domain parameters listed in the beginning of this section and an end-user IPv6 prefix of 2001:db8:9511:2200::/56. The mapping rule for this example can be processed as a BMR as the end-user IPv6 prefix is the correct prefix length and is a child prefix of 2001:db8:9500::/40.

Since the EA bits contain bits representing a PSID value, both the shared IPv4 address and PSID value must be derived through the following steps:

1. Shared IPv4 address: This value is provided by the IPv4 address index bits contained in the 8 high order bits of the End-user IPv6 prefix EA which are highlighted in green:

   2001:db8:95**11**:2200::/56

   The value 0x11 represents the 17th address of 192.0.2.17 from the Rule IPv4 prefix of 192.0.2.0/24.

2. PSID: This value is provided by the PSID bits contained in the 8 low order bits of the End-user IPv6 prefix EA which are highlighted in red:

   2001:db8:9511:**22**00::/56

The PSID value in this example is 0x22 (34).

Given the PSID value it is now possible to determine the set of usable source TCP/UDP ports and ICMP identifier values defined by the PSID. The MAP CE will be restricted to this set of values when assigning a source port during the NAPT process. The set of usable ports values are determined in the following fashion.

1. Apply the PSID value of 0x22 (0b00100010) to the 16-bit representation of all possible UDP/TCP port values. The PSID bits follow the high order bits defined by the PSID offset. The example domain definition uses a PSID offset of 6 resulting in the following PSID value application:

   00000000100010**00**

   Where:
   PSID Offset Bits = The 6 initial high order bits
   PSID Bits = The next 8 bits (the length of the PSID bits)
   Contiguous Port Bits = The 2 low order bits representing 4 contiguous ports

Using the following rules the set of usable ports can be determined:

- The PSID Offset Bits can be any value as long as at least one bit is set
- The PSID Bits value must remain fixed
- The Contiguous Port Bits can be any value

These rules will result in the following port set allocation characteristics:

- The CE is allocated 252 ports which are split across 63 ranges of 4 ports each
- The first usable port range is 1160-1163 (0000010010001000 - 0000010010001011)
- The last usable port range is 64648-64651 (1111110010001000 - 1111110010001011)

Lastly, the MAP CE will construct the MAP IPv6 address that originates and terminates the MAP function for the applicable BMR. This address is contained with the End-user IPv6 prefix of 2001:db8:9511:2200::/56 and is constructed as follows:

2001:db8:9511:2200:0000:c000:0211:0022

Where:
64 high order bits: The End-user IPv6 prefix right padded with 0

16 bits: Set to the value 0
32 bits: Set to the shared IPv4 address of 192.0.2.0.17
16 bits: Set to the PSID value of 0x22 (34) and left padded with 0

Section 6 of [RFC 7597] defines the MAP IPv6 address format including instances where the End-user IPv6 prefix length is larger than 64.

Appendix II Resolving IPv4 Traffic to Originating IPv6 Prefix

This section provides an example of the algorithms used to associate a particular IPv4 flow with the originating MAP subscriber. A hypothetical law enforcement tracking request for a TCP session originating from IP address 192.0.2.15 and a source port of 1478 is used in this example. This scenario is based on a MAP domain with the following parameters:

- Rule IPv6 Prefix: 2001:db8:9500::/40
- Rule IPv4 Prefix: 192.0.2.0/24
- EA Bits: 16 (sharing ratio of 256 to 1)
- PSID Offset: 6

The techniques used in this example are for illustration purposes only. A software implementation can utilize existing IPv6 libraries and simple bit arithmetic to produce the same output.

The first step is to derive the PSID value associated with the reported TCP source port of 1478. This is required as this example is a many-to-one IPv4 address translation scenario. This value is the first of two components that are stored in the EA bits. The following steps are required to complete the conversion:

1. The TCP source port value of 1478 is represented in the context of 16-bit integer as follows:

   ```
   0000010111000110
   ```

2. The PSID length in bits value is calculated in the following fashion:

   ```
   EA Bits - ( 32 - Rule IPv4 Prefix Length )  = PSID Length
      16   -  (          (32 - 24)           )  = 8
   ```

3. It is now possible to identify the PSID value from the 16-bit representation of the source TCP port. Starting at the high order bit, we can now use the PSID offset value of 6 to determine where the PSID bits begin. Using the value calculated in step #2 we can determine where the PSID bits end. We apply this logic in following fashion:

   00000101110001**10**

   Where:
   PSID Offset Bits = The 6 (PSID offset value) initial high order bits
   PSID Bits = The next 8 (PSID length) bits

   The value stored in the PSID bits of 01110001 which gives us a PSID value of 113 (0x71)

The next step is to determine the second value store in the EA bits, which is the IPv4 address mapping component. The following steps are used to determine this value:

1. The length in bits for the IPv4 address mapping is determined in the following fashion:

   ```
   EA Bits – PSID Length = IPv4 address mapping bit length
     16    -      8      = 8
   ```

2. Next, the reported IPv4 address of 192.0.2.15 is converted to the corresponding index value for the IPv4 address mapping bits. The IPv4 address mapping bit length value provides the range of possible index values. The example Rule IPv4 prefix of 192.0.2.0/24 provides a range of index values starting at index value 0 for 192.0.2.0 and ending in value 255 for 192.0.2.255. Therefore, the IP address 192.0.2.15 has an index value of 15.

At this point of the process, the steps above have provided the following values:

- PSID Length = 8 bits
- PSID Value = 113 (0x71)
- Address Mapping Length = 8 bits
- Address Mapping Index value = 15 (0x0f)

Using this information, the Rule IPv6 prefix can be converted to the IPv6 end-user prefix associated with originating subscriber site used in the example. This is done by appending the EA bits following the first 40 high order bits of the rule IPv6 prefix as illustrated in the following steps:

- Convert the Rule IPv6 prefix to its 128-bit integer representation. This example will utilize hex notation for presentation purposes:

  0x20010db895000000000000000000000

- Append the 8-bit PSID value to the 8-bit Address Mapping Index value to form the 16-bit EA bits value:

  0x0f71

- The 16-bit EA value is appended to the first 40 (Rule IPv6 prefix length) high order bits of the Rule IPv6 prefix as follows:

  0x20010db8950f71000000000000000000

- Convert the 128-bit integer representation back to IPv6 prefix notation with a resulting 56-bit prefix length (Rule IPv6 prefix length + EA length in bits):

  2001:db8:950f:7100::/56

The service provider can now search through the DHCPv6 (or other address provisioning mechanism) logs to identify the subscriber associated with 2001:db8:950f:7100::/56 for the date and time reported in the law enforcement request.

Appendix III How to Build a Virtual MAP Environment

This section provides a tutorial on building a virtual MAP-E or MAP-T environment using VirtualBox. The assumed audience for this tutorial is IP network engineers who are comfortable with VirtualBox, Linux, and traditional CLI based network configuration. Two approaches may be taken for deploying the virtual environment. The environment may be deployed manually by using the step-by-step instructions provided in this section. Alternatively, an OVF 1.0 VirtualBox appliance package is available for either MAP-E or MAP-T allowing for the rapid deployment of the entire environment. The tutorial also provides coverage of how MAP domain parameters are generated given a set of Service Provider defined constraints. The virtual environment utilizes DHCPv6 provisioning, as a virtual MAP supporting TLV 202 provisioning is not available as of the document publication date.

III.1 Hardware and Software Requirements for Virtualization Platform

The instructions for building out a complete environment require a single host running VirtualBox. Alternative hypervisors may be used provided the implementer understands how to convert the referenced VirtualBox images to the desired hypervisor image standard. When using a VirtualBox host the following system requirements apply:

- Core I5 or better processor with virtualization support enabled in the BIOS
- Minimum 8GB RAM
- 10GB of disk space available to virtual hosts
- Supported Operating Systems
  - Windows 7 64-bit, Windows 8/8.1 64-bit, Windows 10 64-bit
  - MAC OS X 10.8 or better
  - Ubuntu 14.04 64-bit or other comparable Linux distribution
  - VirtualBox 5.0.10 and above

III.1.1 Virtual Environment Capabilities

The MAP virtual environment is a self-contained environment designed to allow for the observation of MAP in actions. The environment's use of the 2001:db8::/32 documentation prefix prevents clients from reaching the IPv6 internet. Furthermore, Internet edge connectivity is emulated using a RFC1918 addressed (10.0.2.32/30) VirtualBox NAT interface, which adds an additional layer of NAT and which lacks IPv6 reachability. Both of these components can be modified assuming:

- Sufficient GUA space is available to the environment
- Sufficient public IPv4 space and connectivity is available to the virtualization host edge facing interface and sufficient public addressing to accommodate the defined Rule IPv4 prefix.

Elimination of the NAT edge and the direct attachment of the environment to the Internet is not recommended as the deployed components are not hardened and may pose a security risk.

III.1.2 MAP Virtual Network Environment

The virtual MAP environment is built by deploying the following components into VirtualBox:

- MAP CE: An OpenWRT image that supports both MAP-E and MAP-T.
- Access Router: A router that provides access layer functionality to the MAP CE function. In an MSO network a CMTS would typically provide this function. The access router also provides connectivity to and from the provisioning server.
- Provisioning Server: An Ubuntu 14.04 server that provides DHCPv6, DNS, and BGP route server functionality. The route server is required to inject applicable IA_PD routes on the access router using DHCPv6.

- MAP BR: The BR function implemented on a Vyatta ASAMAP open source implementation developed by Maakazu Asama.
- Customer System: Any desktop operating system deployable in a VirtualBox. The example utilizes Xubuntu 14.04 desktop operating system

The diagram below is a high-level representation of the MAP virtual network environment.

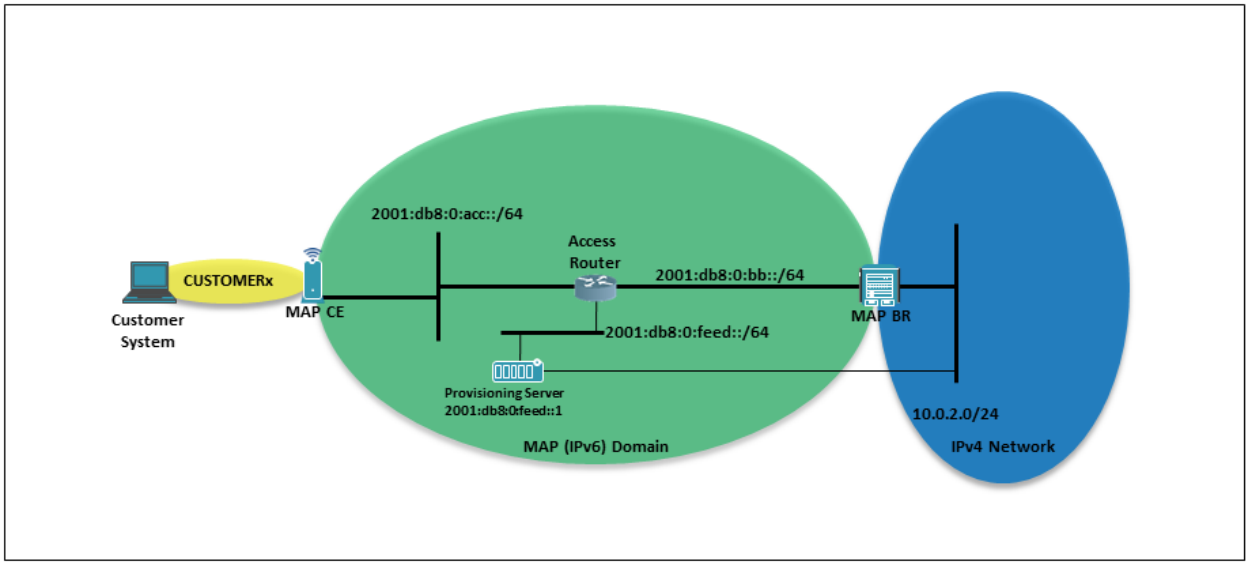

Figure 8 - MAP Virtual Network Environment

The following table provides a mapping between infrastructure addressing and corresponding VirtualBox network name. The "Edge" network refers to a VirtualBox "NAT network" interface type that is created prior to the deployment of the virtual MAP environment.

Table 4 - VirtualBox Network Definitions

Network Name	Network Type	Prefix Allocation	Description
ACCESS	Internal Network	2001:db8:0:acc::/64	MAP CE WAN attachment network
PROVISION	Internal Network	2001:db8:0:feed::/64	DHCPv6 and DNS server network
CORE	Internal Network	2001:db8:0:bb::/64	BR MAP Domain Facing Network
NatNetwork	NAT Network	10.0.2.0/24	NAT interface network with Internet connectivity
VirtualBox Host – Only Ethernet Adapter	Host-Only Network	192.168.56.0/24	Provides out-of-band access to provisioning server, access router, and MAP BR.
CUSTOMERx	Internal Network	IPv4: Customer RFC1918 IPv6: Provider issued prefix	Customer LAN with user defined RFC1918 network and an IPv6 prefix that was provisioned by the service provider. A unique network name must be created for each customer. For example, the first customer site should be named CUSTOMER1

The following table provides the mapping for the deployed network components and their virtual network points of attachment.

Table 5 - Interface Mappings

Node Name	Interface to Virtual Network Mapping
MAP CE	Adapter 1:CUSTOMER1 Adapter 2: ACCESS

Node Name	Interface to Virtual Network Mapping
Provisioning Server	Adapter 1: PROVISION Adapter 2: EDGE Adapter 3: VirtualBox Host-Only (optional)
Access Router	Adapter 1: CORE Adapter 2: ACCESS Adapter 3: PROVISION Adapter 4: VirtualBox Host-Only (optional)
MAP BR	Adapter 1: EDGE Adapter 2: CORE Adapter 3: VirtualBox Host-Only (optional)
Customer System	Adapter 1: CUSTOMER1 Adapter 2: VirtualBox Host-Only (optional)

III.1.3 Configuring the Environment

The MAP Virtual environment may be deployed using a fully configured VirtualBox environment import or a more manual step-by-step configuration approach. The resulting virtual environments will differ in how they handle the routing of the IA_PD assigned to a particular MAP-CE. The step-by-step approach requires a manually entered static route for each IA_PD assignment to the corresponding MAP-CE. The fully configured VirtualBox environment includes a custom script that uses ExaBGP to automatically inject the IA_PD route whenever it is assigned to the MAP-CE.

III.1.3.1 Generating MAP Domain Parameters for the Tutorial Virtual Environment

A MAP domain is defined by a unique set of mapping rules and a BR address or prefix. The virtual environment MAP domain described in this tutorial will employ the simplest case consisting of a single mapping rule. A mapping rule requires the following attributes:

- Rule IPv4 prefix
- Rule IPv6 prefix
- EA bit length
- PSID offset

Generating a mapping rule requires that a base set of constraints and associated attribute are defined. The values chosen are generally arbitrary with the goal of being easy to implement in a virtual environment running on a laptop. The attribute values are as follows:

- **Rule IPv4 prefix 10.0.2.32/30:** A host or prefix that will act as the public translated address or addresses that may be shared across multiple MAP CEs. The MAP tutorial virtual environment allocates four private IPv4 addresses (10.0.2.32 – 10.0.2.35) in support of the VirtualBox NAT interface. In a real world deployment this address space would be public IPv4 address space as opposed to [RFC 1918].
- **Sharing Ratio of 64-to-1:** The ratio of customers to a single IPv4 address.
- **End-user IPv6 Prefix length /60:** Each customer will obtain a /60 IPv6 prefix allocation for MAP and CE LAN interface.
- **A PSID offset value 2:** Ports (0-16383) will be excluded from all PSID assignments.

With the initial values determined, the rest of the mapping rule attributes are derived to form a complete mapping rule set. The following list demonstrates the evolution of a mapping rule from two attributes to a full set of attributes:

1. Rule IPv4 Prefix and PSID offset values have been defined resulting in a partial mapping rule of:

   Rule IPv4 Prefix: 10.0.2.32/30

PSID offset: 2

2. The following steps are used to derive the MAP rule EA bits length
   a. The Rule IPv4 prefix of 10.0.2.32/30 requires a 2-bit allocation within the defined EA bit length (as a /30 can represent for address values).
   b. The defined 64-1 ratio requires 64 unique PSID values which require an additional 6-bit allocation within the defined EA bit length.
   c. The EA length of 8-bits is obtained by adding the bits required to map the address (2-bits) and PSID (6-bits).

   The EA length value is now added to the existing mapping rule resulting in the following:

   Rule IPv4 Prefix: 10.0.2.32/30
   EA Length: 8
   PSID Offset: 2

3. The prefix length of the Rule IPv6 prefix is *calculated* by subtracting the EA length value from the end-user IPv6 prefix length value. For example: **60** (end-user IPv6 length) **- 8** (EA length) **= 52** (Rule IPv6 prefix length). The calculated rule IPv6 prefix length is applied to an arbitrarily selected available prefix of 2001:db8:ca6::/52 to be used as the Rule IPv6 prefix. This results in the complete mapping rule of:

   Rule IPv4 Prefix: 10.0.2.32/30
   Rule IPv6 Prefix: 2001:db8:ca6::/52
   EA Length: 8
   PSID Offset: 2

With the mapping rule for the MAP domain established the BR reachability attribute must be defined. The unused prefix of 2001:db8:0:def::/64 from the virtual lab allocated space of 2001:db8::/32 is used to provide the BR address or BR prefix allocation (depending on which version of the lab MAP-E/MAP-T is implemented). The following attribute values are defined:

- MAP-E BR address: 2001:db8:0:def::1
- MAP-T BR prefix: 2001:db8:0:def::/64

III.1.3.2 Generating the ISC DHCP MAP Configuration

Given a full set of MAP domain parameters the requirements in [RFC 7598] are used to generate the DHCPv6 options for MAP-E or MAP-T. For learning purposes this document provides a breakdown of the components that make up a container option. Cisco provides an online tool at http://map46.cisco.com/MAP.php which generates these values automatically. The end of this section provides the full ISC DHCP configuration for both transport modes of MAP. The virtual lab is intended to operate in one of the two supported transport modes and should not use both configurations simultaneously. For clarity purposes, the sections supporting the route injection scripts are omitted in the examples.

III.1.3.2.1 MAP-E

The first step in constructing a MAP-E container option is to generate a single S46 Rule Option (since the lab MAP domain contains a single mapping rule) with its embedded S46 Port Parameters Option. The S46 Rule Option is constructed in the two steps that follow:

1. Construct S46 Rule Option with the following fields and values:
   - option-code: 00:59 (2 octets)
     - Explanation: IANA designated value of 89 for S46 Rule Option
   - option-length: 00:17 (2 octets)
     - Explanation: The length in number of octets of the option excluding the option-code and option-length and including the S46 Port Parameters Option. The calculated value is 23 (00:17) which represents 17 octets for the base S46 Rule Option (option-code and option-

length excluded) + 8 octets for the S46 Port Parameters Option (option-code and option-length included). The S46 Port Parameters Option will be constructed in step #2.

- flags: 00 (1 octet)
  - Explanation: The virtual environment reflects a flags field with no value.
- ea-len: 08 (1 octet)
  - Explanation: The EA length of 8 bits defined during the mapping rule generation.
- prefix4-len: 1e (1 octet)
  - Explanation: The Rule IPv4 prefix length is 30 bits as specified by the mapping rule.
- ipv4-prefix: 0a:00:02:20 (4 octets)
  - Explanation: The Rule IPv4 prefix of 10.0.2.32 represented in hexadecimal format.
- prefix6-len: 34 (1 octet)
  - Explanation: The Rule IPv6 prefix length is 52 as specified by the mapping rule.
- ipv6-prefix: 20:01:0d:b8:0c:a6:00 (7 octets)
  - Explanation: The Rule IPv6 prefix of 2001:db8:ca6::. This field is right padded with an additional 4 bit set to 0 to make this field evenly divisible by 8 as is required by [RFC 7598].

Resulting S46 Rule Option octet sequence (without S46 Port Parameters Option):

```
00:59:00:17:00:08:1e:0a:00:02:20:34:20:01:0d:b8:0c:a6:00
```

2. Construct the S46 Port Parameters Option that is contained within the S46 Rule Option:
   - option-code: 00:5d (2 octets)
     - Explanation: IANA designated value of 93 for S46 Port Parameters Option
   - option-length: 00:04 (2 octets)
     - Explanation: The length in number of octets of the option excluding the option-code and option length. The calculated value is 4.
   - offset: 02 (1 octet)
     - Explanation: The PSID offset as specified by the mapping rule.
   - PSID-len: 00 (1 octet)
     - Explanation: The virtual environment does not implement an explicitly provisioned PSID value. The PSID value is derived from the assigned End-user IPv6 prefix. The value of this field is 0 to indicate this type of posture.
   - PSID: 00:00 (2 octets)
     - Explanation: Since the PSID-len field is 0, this field has been set to the value of 0 and is not processed.

   Resulting S46 Port Parameters Option octet sequence:

   ```
   00:5d:00:04:02:00:00:00
   ```

The next option to be constructed contains the BR address information for the defined MAP-E domain.

3. Construct the S46 BR Option:
   - option-code: 00:5a (2 octets)
     - IANA designated value of 90 for S46 BR Option
   - option-length: 00:10 (2 octets)
     - Explanation: The length in number of octets of the option excluding the option-code and option length. The calculated value is 16.
   - br-ipv6-address: 20:01:0d:b8:00:00:0d:ef:00:00:00:00:00:00:00:01 (16 octets)
     - Explanation: The MAP-E BR address of 2001:db8:0:def::1

   Resulting S46 BR Option octet sequence:

   00:5a:00:10:20:01:0d:b8:00:00:0d:ef:00:00:00:00:00:00:00:01

4. The octet sequence for the encapsulated options within the S46 MAP-E Container Option is generated by concatenating the results of steps 1,2, and 3 (in the order listed):

```
00:59:00:17:00:08:1e:0a:00:02:20:34:20:01:0d:b8:0c:a6:00:00:5d:00:04:02:00:00:00:
00:5a:00:10:20:01:0d:b8:00:00:0d:ef:00:00:00:00:00:00:00:01
```

The resulting ISC DHCP DHCPv6 configuration file is displayed below.

```
# dhcpd6.conf

# A MAP-E DHCPv6 Server Configuration

authoritative;
default-lease-time 3600;
max-lease-time 86400;

# S46 MAP-E Container Option is utilized
option dhcp6.map-option code 94 = string;

#
subnet6 2001:db8:0:feed::/64 {
  range6 2001:db8:0:feed::a 2001:db8:0:feed::a;
}

subnet6 2001:db8:0:acc::/64 {
  range6 2001:db8:0:acc::100 2001:db8:0:acc::1ff;
  option dhcp6.name-servers 2001:db8:0:feed::1000;

  option dhcp6.map-option
    00:59:00:17:00:08:1e:0a:00:02:20:34:20:01:0d:b8:0c:a6:00:    # S46 Rule
    00:5d:00:04:02:00:00:00:                                     # S46 PortParams
    00:5a:00:10:20:01:0d:b8:00:00:0d:ef:00:00:00:00:00:00:00:01; # S46 BR

  # Prefix delegation configuration for subnet 2001:db8:0:acc::/64.
  # 256 /60 prefixes are assignable to MAP CEs.
  prefix6 2001:db8:ca6:: 2001:db8:ca6:ff0:: /60;
}
```

III.1.3.2.2 MAP-T

The first step in constructing a MAP-T container option is to generate a single S46 Rule Option (since the lab MAP domain contains a single mapping rule) with its embedded S46 Port Parameters Option. The S46 Rule Option is constructed in the two steps that follow:

1. Construct S46 Rule Option with the following fields and values:
    - option-code: 00:59 (2 octets)
        - Explanation: IANA designated value of 89 for S46 Rule Option
    - option-length: 00:17 (2 octets)
        - Explanation: The length in number of octets of the option excluding the option-code and option-length and including the S46 Port Parameters Option. The calculated value is 23 (00:17) which represents 17 octets for the base S46 Rule Option (option-code and option-length excluded) + 8 octets for the S46 Port Parameters Option (option-code and option-length included). The S46 Port Parameters Option will be constructed in step #2.
    - flags: 00 (1 octet)
        - Explanation: The virtual environment reflects a flags field with no value.
    - ea-len: 08 (1 octet)
        - Explanation: The EA length of 8 bits defined during the mapping rule generation.
    - prefix4-len: 1e (1 octet)
        - Explanation: The Rule IPv4 prefix length is 30 bits as specified by the mapping rule.

- ipv4-prefix: 0a:00:02:20 (4 octets)
  - Explanation: The Rule IPv4 prefix of 10.0.2.32 represented in hexadecimal format.
- prefix6-len: 34 (1 octet)
  - Explanation: The Rule IPv6 prefix length is 52 as specified by the mapping rule.
- ipv6-prefix: 20:01:0d:b8:0c:a6:00 (7 octets)
  - Explanation: The Rule IPv6 prefix of 2001:db8:ca6::. This field is right padded with an additional 4 bit set to 0 to make this field evenly divisible by 8 as is required by [RFC 7598].

Resulting S46 Rule Option octet sequence (without S46 Port Parameters Option):

00:59:00:17:00:08:1e:0a:00:02:20:34:20:01:0d:b8:0c:a6:00

2. Construct the S46 Port Parameters Option that is contained within the S46 Rule Option:
   - option-code: 00:5d (2 octets)
     - Explanation: IANA designated value of 93 for S46 Port Parameters Option
   - option-length: 00:04 (2 octets)
     - Explanation: The length in number of octets of the option excluding the option-code and option length. The calculated value is 4.
   - offset: 02 (1 octet)
     - Explanation: The PSID offset as specified by the mapping rule.
   - PSID-len: 00 (1 octet)
     - Explanation: The virtual environment does not implement an explicitly provisioned PSID value. The PSID value is derived from the assigned End-user IPv6 prefix. The value of this field is 0 to indicate this type of posture.
   - PSID: 00:00 (2 octets)
     - Explanation: Since the PSID-len field is 0 this field has been set to the value of 0 and is not processed.

   Resulting S46 Port Parameters Option octet sequence:

   ```
   00:5d:00:04:02:00:00:00
   ```

The next option to be constructed contains the BR prefix (DMR) information for the defined MAP-T domain.

3. Construct the S46 DMR Option:
   - option-code: 00:5b (2 octets)
     - IANA designated value of 91 for S46 DMR Option
   - option-length: 00:09 (2 octets)
     - Explanation: The length in number of octets of the option excluding the option-code and option length. The calculated value is 16.
   - dmr-prefix-len: 40 (1 octet)
     - Explanation: The BR prefix of length of 64 as specified by the MAP-T domain parameters.
   - dmr-ipv6-prefix: 20:01:0d:b8:00:00:0d:ef (8 octets)
     - Explanation: The MAP-T BR prefix of 2001:db8:0:def::/64

   Resulting S46 DMR Option octet sequence:

   ```
   00:5b:00:09:40:20:01:0d:b8:00:00:0d:ef
   ```

   The octet sequence for the encapsulated options within the S46 MAP-T Container Option is generated by concatenating the results of steps 1, 2, and 3 (in the order listed):

   ```
   00:59:00:17:00:08:1e:0a:00:02:20:34:20:01:0d:b8:0c:a6:00:00:5d:00:04:02:00:00:00:
   ```

```
        00:5b:00:09:40:20:01:0d:b8:00:00:0d:ef
```

The resulting ISC DHCP DHCPv6 configuration file is displayed below.

```
# dhcpd6.conf

# A MAP-T DHCPv6 Server Configuration

authoritative;
default-lease-time 3600;
max-lease-time 86400;

# S46 MAP-T Container Option is utilized
option dhcp6.map-option code 95 = string;

#

subnet6 2001:db8:0:feed::/64 {
 range6 2001:db8:0:feed::a 2001:db8:0:feed::a;
}

subnet6 2001:db8:0:acc::/64 {
  range6 2001:db8:0:acc::100 2001:db8:0:acc::1ff;
  option dhcp6.name-servers 2001:db8:0:feed::1000;

  option dhcp6.map-option
    00:59:00:17:00:08:1e:0a:00:02:20:34:20:01:0d:b8:0c:a6:00: # S46 Rule
    00:5d:00:04:02:00:00:00:                                  # S46 PortParams
    00:5b:00:10:40:20:01:0d:b8:00:00:0d:ef;                   # S46 DMR

  # Prefix delegation configuration for subnet 2001:db8:0:acc::/64.
  # 256 /60 prefixes are assignable to MAP CEs.
  prefix6 2001:db8:ca6:: 2001:db8:ca6:ff0:: /60;
}
```

III.1.3.3 Installing and Configuring the Virtual Appliances

This section will describe the method used to install and configure each of the virtual appliances/servers within the virtual environment.

III.1.3.3.1 Provisioning Server

The virtual lab environment provisioning server utilizes the Ubuntu Server 14.04.3 LTS 64-bit Linux distribution. It is available for download at http://www.ubuntu.com/download/server. The server should be installed into VirtualBox with the following configuration posture:

- Name: provisioning
- Type: Linux
- Version: Ubuntu (64 bit)
- Memory size: 512MB
- Drive Size: 8.00 GB
- Hard drive file type: VDI
- Storage on physical hard drive: Dynamically allocated
- Network Interfaces: Per the "Interface to Virtual Network Mapping" table the interfaces should be configured as follows:
    - Adapter 1

    - Attached to: Internal Network
    - Name: PROVISION
- Adapter 2:
    - Attached to: NAT
- Adapter 3:
    - Attached to: Host-only Adapter
    - Name: VirtualBox Host-Only Ethernet Adapter

The server should be installed and configured as follows:

- Network Configuration
    - eth0
        - IPv6 address: 2001:db8:0:feed::1000/64
    - eth1
        - IPv4 address: dhcp
    - eth2
        - IPv4 address: dhcp
    - Routes
        - IPv4: dhcp on interface eth1 provides a default out the NAT interface
        - IPv6: 2001:db8:ca6::/52 via 2001:db8:0:feed::1 dev eth0
            - Provides reachability to customer prefixes
        - IPv6: 2001:db8:0:acc::/64 via 2001:db8:0:feed::1 dev eth0
            - Provides reachability to MAP CE WAN interfaces
    - DNS Server
        - Server should be configured as a recursive server for the environment IA-NA and IA_PD prefixes.
    - DHCPv6 Server
        - Server should be configured with the example dhcpd6.conf for the applicable MAP transport mode.
- Required Packages
    - bind9
    - isc-dhcp-server

III.1.3.3.2 Access Router

The virtual lab environment access router utilizes the ASAMAP Vyatta router distribution. It is available for download at http://www.ubuntu.com/download/server. The access router image should be imported into VirtualBox with the following configuration posture:

- Name: access (Access Router)
- Type: Linux
- Version: Debian (32 bit)
- Memory size: 512MB
- Drive Size: 2.00 GB
- Hard drive file type: VDI
- Storage on physical hard drive: Dynamically allocated
- Network Interfaces: Per the "Interface to Virtual Network Mapping" table the interfaces should be configured as follows:
    - Adapter 1
        - Attached to: Internal Network

    - Name: CORE
- Adapter 2
    - Attached to: Internal Network
    - Name: ACCESS
- Adapter 3
    - Attached to: Internal Network
    - Name: PROVISION
- Adapter 4:
    - Attached to: Host-only Adapter
    - Name: VirtualBox Host-Only Ethernet Adapter

The following represents the configuration posture the access router that does not use the automated route injection scheme for the MAP CE IA_PDs.

- Access Router Configuration
    - eth0
        - IPv6 address: 2001:db8:0:bb::2/64
        - Vyatta syntax

```
set interfaces ethernet eth0 address 2001:db8:0:bb::2/64
```

    - eth1
        - IPv6 address: 2001:db8:0:acc::1/64
        - The router will also provide Router Advertisements
        - Vyatta syntax

```
set interfaces ethernet eth1 address 2001:db8:0:acc::1/64
set interfaces ethernet eth1 ipv6 router-advert send-advert true
set interfaces ethernet eth1 ipv6 router-advert prefix 2001:db8:0:acc::/64 autonomous-flag false
set interfaces ethernet eth1 ipv6 router-advert managed-flag true
set interfaces ethernet eth1 ipv6 router-advert other-config-flag true
```

    - eth2
        - IPv6 address: 2001:db8:0:feed::1/64
        - Vyatta syntax

```
set interfaces ethernet eth2 address 2001:db8:0:feed::1/64
```

    - eth3
        - IPv4 address: dhcp
        - Vyatta syntax

```
set interfaces ethernet eth3 address dhcp
```

    - DHCPv6 Relay Vyatta syntax

```
set service dhcpv6-relay listen-interface eth1 address ::
set service dhcpv6-relay listen-port 547
set service dhcpv6-relay upstream-interface eth2 address 2001:db8:0:feed::1000
```

    - Routes
        - IPv4: None required
        - IPv6:
            - Default Route

- Vyatta syntax

```
set protocols static route6 ::/0 next-hop 2001:db8:0:bb::1
```

- Static routes for each provisioned MAP CE IA_PD must be configured manually for environments lacking the automatic route injection process. For example, a MAP CE with an IA-NA of 2001:db8:0:acc::100 and an IA_PD of 2001:db8:ca6::/60 would result in the following Vyatta configuration statement:

```
set protocols static route6 2001:db8:ca6::/60 next-hop 2001:db8:0:acc::100
```

III.1.3.3.3 MAP-E BR

The virtual lab environment MAP-E BR utilizes the ASAMAP Vyatta router distribution. It is available for download at http://www.ubuntu.com/download/server. The MAP-E BR image should be imported into VirtualBox with the following configuration posture:

- Name: MAP-E BR
- Type: Linux
- Version: Debian (32 bit)
- Memory size: 512MB
- Drive Size: 2.00 GB
- Hard drive file type: VDI
- Storage on physical hard drive: Dynamically allocated
- Network Interfaces: Per the "Interface to Virtual Network Mapping" table the interfaces should be configured as follows:
  - Adapter 1
    - Attached to: NAT Network
    - Name: EDGE
  - Adapter 2
    - Attached to: Internal Network
    - Name: CORE
  - Adapter 3:
    - Attached to: Host-only Adapter
    - Name: VirtualBox Host-Only Ethernet Adapter

The following represents the configuration posture of the access router that does not use the automated route injection scheme for the MAP CE IA_PDs.

- MAP-E BR Configuration
  - eth0
    - IPv4 address: 10.0.2.31/24
    - Proxy-ARP will be enabled for the MAP Rule IPv4 prefix addresses

    Vyatta syntax

```
set interfaces ethernet eth0 address 10.2.0.31/24
set interfaces ethernet eth0 ip enable-proxy-arp
```

  - eth1
    - IPv6 address: 2001:db8:0:bb::1/64

        - Vyatta syntax

```
set interfaces ethernet eth1 address 2001:db8:0:bb::1/64
```

    - eth2
        - IPv4 address: dhcp
        - Vyatta syntax

```
set interfaces ethernet eth3 address dhcp
```

    - MAP-E Vyatta configuration syntax

```
set interfaces map map0 br-address 2001:db8:0:def::1/64
set interfaces map map0 default-forwarding-mode encapsulation
set interfaces map map0 default-forwarding-rule true
set interfaces map map0 role br
set interfaces map map0 ipv6-fragment-size 1500
set interfaces map map0 rule 1 ea-length 8
set interfaces map map0 rule 1 ipv4-prefix 10.0.2.32/30
set interfaces map map0 rule 1 ipv6-prefix 2001:db8:ca6::/52
set interfaces map map0 rule 1 psid-offset 2
set protocols static interface-route 10.0.2.32/30 next-hop-interface map0
```

    - Routes
        - IPv4:
            - Default Route
                - Vyatta syntax

```
set protocols static route 0.0.0.0/0 next-hop 10.0.2.1
```

        - IPv6:
            - Default Route
                - Given the non-GUA addressing, a default route has no practical use. Environments utilizing GUA IA_PD assignments may add a default route if applicable.
            - Static Routes

```
set protocols static route6 2001:db8:ca6::/52 next-hop 2001:db8:0:bb::2
```

III.1.3.3.4 MAP-T BR

The virtual lab environment MAP-T BR utilizes the ASAMAP Vyatta router distribution. It is available for download at http://www.ubuntu.com/download/server. The MAP-T BR image should be imported into VirtualBox with the following configuration posture:

- Name: MAP-T BR
- Type: Linux
- Version: Debian (32 bit)
- Memory size: 512MB
- Drive Size: 2.00 GB
- Hard drive file type: VDI

- Storage on physical hard drive: Dynamically allocated
- Network Interfaces: Per the "Interface to Virtual Network Mapping" table the interfaces should be configured as follows:
  - Adapter 1
    - Attached to: NAT Network
    - Name: EDGE
  - Adapter 2
    - Attached to: Internal Network
    - Name: CORE
  - Adapter 3:
    - Attached to: Host-only Adapter
    - Name: VirtualBox Host-Only Ethernet Adapter

The following represents the configuration posture of the access router that does not use the automated route injection scheme for the MAP CE IA_PDs.

- MAP-T BR Configuration
  - eth0
    - IPv4 address: 10.0.2.31/24
    - Proxy-ARP will be enabled for the MAP Rule IPv4 prefix addresses
    - Vyatta syntax

```
set interfaces ethernet eth0 address 10.2.0.31/24
set interfaces ethernet eth0 ip enable-proxy-arp
set interfaces ethernet eth0 description EDGE
```

  - eth1
    - IPv6 address: 2001:db8:0:bb::1/64
    - Vyatta syntax

```
set interfaces ethernet eth1 address 2001:db8:0:bb::1/64
set interfaces ethernet eth1 description CORE
```

  - eth2
    - IPv4 address: dhcp
    - Vyatta syntax

```
set interfaces ethernet eth3 address dhcp
set interfaces ethernet eth3 description HostOnly
```

  - MAP-E Vyatta configuration syntax

```
set interfaces map map0 br-address 2001:db8:0:def::/64
set interfaces map map0 default-forwarding-mode translation
set interfaces map map0 default-forwarding-rule true
set interfaces map map0 role br
set interfaces map map0 ipv6-fragment-size 1500
set interfaces map map0 rule 1 ea-length 8
set interfaces map map0 rule 1 ipv4-prefix 10.0.2.32/30
set interfaces map map0 rule 1 ipv6-prefix 2001:db8:ca6::/52
set interfaces map map0 rule 1 psid-offset 2
```

```
set protocols static interface-route 10.0.2.32/30 next-hop-interface map0
```

- Routes
  - IPv4:
    - Default Route
      - Vyatta syntax

```
set protocols static route 0.0.0.0/0 next-hop 10.0.2.1
```

  - IPv6:
    - Default Route
      - Given the non-GUA addressing a default route has no practical use. Environments that utilize GUA IA_PD assignments may add a default route if applicable.
    - Static Routes

```
set protocols static route6 2001:db8:ca6::/52 next-hop 2001:db8:0:bb::2
```

III.1.3.3.5 MAP-CE

The virtual lab environment MAP-CE utilizes the OpenWrt Linux distribution. It is available for download at http://www.ubuntu.com/download/server. The OpenWrt router should be imported into VirtualBox with the following configuration posture:

- Name: MAP-CE
- Type: Linux
- Version: Other (64 bit)
- Memory size: 64MB
- Drive Size: NA as it is set by image import
- Hard drive file type: VMDK
- Disk Controller type: IDE
- Network Interfaces: Per the "Interface to Virtual Network Mapping" table the interfaces should be configured as follows:
  - Adapter 1
    - Attached to: Internal Network
    - Name: CUSTOMER1
  - Adapter 2:
    - Attached to: Internal Network
    - Name: ACCESS

The MAP CE should be installed and configured as follows:

- Network Configuration
  - "wan"
    - All configuration options are commented out
  - "wan6"
    - DHCPv6 must be enabled

- Packages
  - map
  - map-t

III.1.3.3.6 Customer System

The virtual lab environment can utilize any VirtualBox compatible operating system. The system should be configured with the following VirtualBox network posture:

- Name: Customer 1
- Network Interfaces: Per the "Interface to Virtual Network Mapping" table, the interfaces should be configured as follows:
  - Adapter 1
    - Attached to: Internal Network
    - Name: CUSTOMER1
  - Adapter 2:
    - Attached to: Host-only Adapter
    - Name: VirtualBox Host-Only Ethernet Adapter

The customer system should be installed and configured as follows:

- Network Configuration
  - MAP CE Facing Interface
    - IPv4: DHCP
    - IPv6: SLAAC
  - Host-Only Facing Interface
    - IPv4: DHCP
    - IPv6: N/A

III.1.4 Installing the Environment Using Pre-Packaged Appliances

For users who do not wish to build a virtual MAP-E or MAP-T environment from scratch, fully configured VirtualBox appliance modules have been created for ease in configuring both MAP-E and MAP-T test environments. The following instructions can be used to download and install both MAP-E and MAP-T environments quite simply.

Before beginning the download or installation process, be sure that the server being used to host the VirtualBox appliances complies with the system requirements described in Section III.1.3.3.3.

1. Install Oracle VM VirtualBox Manager software on the server. It is recommended that version 5.0.10 or above be used to host the services.
2. Download the MAP-E or MAP-T appliance modules from the hosted web site:
   a. MapE_Appliance.ova
   b. MapT_Appliance.ova
3. After successful download of the appliance modules, launch VirtualBox on the server.
4. When the VirtualBox UI appears, select 'File', then 'Import Appliance'.
5. A pop up window will appear asking to select which appliance to import. Click the folder icon at the right. Select the MAP-E or MAP-T appliance module that you wish to install. Navigate to the correct directory to locate the correct appliance module as required. Highlight the file name and click 'Open'.

6. The directory path and file name will populate on the import window. Click 'Next' to begin the installation process.
7. An 'Appliance Settings' window will appear. Click the 'Import' button to continue the installation process.
8. In the left panel of the VirtualBox UI you will see the following modules appear: MAP CE, Provisioning Server, Access Router, and MAP BR.
9. Install each of the appliance modules in the following order by double clicking on each icon. Launching the module will cause a terminal window to open. When the login prompt appears in each window that is an indication that the module was installed successfully. It is not necessary to login to any of these processes to do any further configurations. They will remain running until they are shut down.
    a. Provisioning Server
    b. Access router
    c. BR
    d. MAP CE

    Note: The modules will not start if the networks bound to the "NatNetwork" and "VirtualBox Host-Only Ethernet Adapter" are not defined in the VirtualBox preferences. Table 4 provides additional information on the network definitions and their bindings.
10. Once the MAP CE module installation has completed, you can confirm proper functionality and connectivity by entering the following commands in the MAP CE window at the 'root@OpenWrt' prompt:
    a. ip –4 route (displays WAN gateway 192.168.1.0 address)
    b. ifconfig eth1 (displays IPv6 link local and globally routable MAP IPv6 address for WAN)
    c. ifconfig br-lan (displays IPv6 link local and globally routable MAP IPv6 address for WAN and RFC 1918 addresses for LAN)
    d. Ping 8.8.8.8 to confirm connectivity to the public DNS server and internet
11. Optionally, you can install your own client desktop (Linux, Windows, etc.) to use for connecting to the public internet and for generating traffic from the MAP CE. This module is not included as a VirtualBox appliance in the MAP-E or MAP-T bundles. The use of client desktops is at the discretion of the end user and is out of scope for this document.

    Note: All NICs associated with the VirtualBox should be set to promiscuous mode. This can be done by choosing 'Settings' from the VirtualBox Manager UI and selecting 'Network'. Each configured adapter should be set to 'Allow All'.
12. Coverage of the module implementation is beyond the scope of this document. However, access to the modules can be obtained using the following credentials:
    a. Provisioning Server - Username: administrator; Password: password
    b. All Vyatta routers including the BR - Username: vyatta; Password: vyatta
    c. MAP CE - Username: root; there is no password set on the MAP CE.
13. Graceful shut down of each module is recommended. The following methods should be used to ensure a graceful shut down of each process:
    a. For the Provisioning Server, MAP BR and Access Router, select 'Machine' from the window menu and choose 'ACPI Shutdown'.
    b. For the MAP CE, type 'halt' at the command prompt.

Appendix IV DNS Behavior

IV.1 Introduction

The exhaustion of IPv4 introduces a challenge as it relates to dual stack customer premise environments, which are likely to be widespread for years to come.

Host operating systems will continue to operate in dual stack with the expectation that off customer premise network communications be predominantly IPv6-only. The use of IPv4 ideally would be limited to customer premise network. The reality of the situation is with dual stack customer premise local area networks there inevitably will be applications and services that lag in their support of IPv6. These applications will unfortunately find that their network-based communications, which are IPv4-only, will be carried to carrier grade network address translation (CGN) devices that are located elsewhere in the service provider network. The focus of this document is DNS and not the overarching protocol flows for the multitude of applications and services.

Interaction with DNS plays a critical role in how applications behave and perform. The transportation of IPv4 DNS queries over an IPv6 infrastructure may have a significant impact on the applications, service, customer, and ultimately the service provider infrastructure.

A customer premise must have IPv6 connectivity if IPv6 is to be the transport for IPv4. The same IPv6 transport is also used to support native IPv6 based communications to and from the Internet. The IPv6 connectivity provided today can and should be leveraged to support improved communications for IPv4 hosts, nodes, applications, and services that reside on the customer premise LAN.

IV.2 Topology

The topology for a customer premise network where IPv4 translation or encapsulation is being used to support legacy IPv4-only communications is as follows:

- Dual stack customer premise LAN
- [RFC 1918] addressing is utilized for IPv4
- Globally routable IPv6 is delegated and assigned to the customer premises router
- IPv6 configuration options including DNS server IPv6 addresses are assigned and configured to the customer premises router

The technology used to enable IPv4 at the customer premise may vary. Examples of popular technologies include GRE over IPv6 and Mapping of Address and Port (MAP). For the purposes of this document, Internet facing IPv4 communications are assumed to traverse a CGN.

IV.3 DNS Resolver Behavior

Today customer premise LANs typically leverage one or more of the following. For both IPv6 and IPv4 one more of the following apply:

- All DNS queries are sent by default to the customer premises router, typically 192.168.0.1 for example. The router may implement a DNS proxy or simply forward to an alternate DNS server for resolution.
- All hosts on the customer LAN are assigned the publicly routable IPv4 address of DNS servers on the Internet, which may include DNS recursive name server supported by the service provider.
- Finally, hosts on the customer LAN may also utilize a DNS server that is administered locally. This DNS server may resolve local names and addresses and will forward DNS queries it cannot satisfy using one or more of the approaches listed above.

When legacy IPv4 communications require support from the customer LAN, if DNS over IPv4 is used today most of the above would result in one of the following:

- Added round trip for DNS over IPv4 over the IPv4aaS infrastructure

- The DNS query forwarded through the CGN

Since the customer premise router is assumed to have IPv6 connectivity and minimally a DNS server IPv6 address, an optimization should be utilized to address the potential performance impact to DNS.

1. Hosts must prefer the use of DNS over IPv6 for all query types, including A RR queries. This assumes that a host is dual stack and will issue a DNS query over IPv6 for an A and/or AAAA RR.
2. The locally administered DNS server must forward customer LAN DNS queries using one or more of the following, regardless of the transport of receipt for the query:
    a. If the locally administered DNS server forwards DNS queries over IPv4, the query must be forwarded to the LAN IPv4 address of the customer premises router.
    b. If the locally administered DNS server is able and configured to support IPv6, the DNS query must be forwarded over IPv6 and not to the LAN IPv4 address of the customer premises router. The destination DNS server IPv6 address can be learned or discovered in a number of ways, which is out of scope for this document.
    c. To avoid using IPv4aaS for DNS queries, the locally administered DNS server should not support DNS recursion for any query types. DNS recursion relies on authoritative DNS server configurations, which is controlled by third parties, to determine the desired IP transport for DNS queries.
3. DNS queries over IPv4 that are sent directly to the customer premises router must be forwarded by the customer premises router using DNS over IPv6 to the IPv6 DNS servers configured on the customer premises router.

All original attributes of the original DNS query must remain intact and as-is. The resolver receiving the DNS query must forward the same, over IPv6 to the target DNS server configured on the customer premises router.

The DNS servers performing the DNS recursion function must have native IPv4 and IPv6 connectivity so they may contact authoritative servers of either address family. The specification of the behavior for recursive DNS server is out of scope for this document.

IV.4 Summary

Adhering to the above will help ensure that DNS queries over IPv4 are not unnecessarily translated and, more importantly, are not subject to lengthy round trip communications to reach DNS servers over IPv6. Further, the use of DNS servers over IPv6 that are closer in proximity to the customer will help to ensure that geo-location accuracy remains intact.

Appendix V Acknowledgements

We wish to thank the following authors contributing directly to this document:

John Jason Brzozowski – Comcast

Jordan Gottlieb – Charter Communications

John Berg – CableLabs/Charter Communications

Kirk Erichsen – Time Warner Cable

Steve Burroughs - CableLabs