Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: December 25, 2015

C. Zhou
Huawei Technologies
T. Taylor
PT Taylor Consulting
Q. Sun
China Telecom
M. Boucadair
France Telecom

June 23, 2015

Attribute-Value Pairs For Provisioning Customer Equipment Supporting IPv4-Over-IPv6 Transitional Solutions
draft-ietf-dime-4over6-provisioning-03

Abstract

During the transition from IPv4 to IPv6, customer equipment may have to support one of the various transition methods that have been defined for carrying IPv4 packets over IPv6. This document enumerates the information that needs to be provisioned on a customer edge router to support a list of transition techniques based on tunneling IPv4 in IPv6, with a view to defining reusable components for a reasonable transition path between these techniques. To the extent that the provisioning is done dynamically, AAA support is needed to provide the information to the network server responsible for passing the information to the customer equipment. This document specifies Diameter (RFC 6733) attribute-value pairs to be used for that purpose.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

A number of transition techniques have been defined to allow IPv4 packets to pass between hosts and IPv4 networks over an intervening IPv6 network while minimizing the number of public IPv4 addresses that need to be consumed by the hosts. Different operators will deploy different technologies, and sometimes one operator will use more than one technology, depending on what is supported by the available equipment and upon other factors both technical and economic.

Each technique requires the provisioning of some subscriber-specific information on the customer edge device. The provisioning may be by DHCPv6 [RFC3315] or by some other method. This document is indifferent to the specific provisioning technique used, but assumes a deployment in which that information is managed by AAA (Authentication, Authorization, and Accounting) servers. It further assumes that this information is delivered to intermediate network nodes for onward provisioning using the Diameter protocol [RFC6733].

As described below, in the particular case where the Lightweight IPv4 Over IPv6 (Lw4o6) [I-D.ietf-softwire-lw4over6] transition method has been deployed, per-subscriber-site information almost identical to that passed to the subscriber site [I-D.ietf-softwire-map-dhcp] also needs to be delivered to the border router serving that site. The Diameter protocol may be used for this purpose too.

This document analyzes the information required to configure the customer edge equipment for the following set of transition methods:

o Dual-Stack Lite (DS-Lite) [RFC6333],

o Lightweight IPv4 Over IPv6 (LW4over6) [I-D.ietf-softwire-lw4over6], and

o Mapping of Address and Port with Encapsulation (MAP-E) [I-D.ietf-softwire-map].

[I-D.ietf-softwire-dslite-multicast] specifies a generic solution for delivery of IPv4 multicast services to IPv4 clients over an IPv6 multicast network. The solution was developed with DS-Lite in mind

but it is however not limited to DS-Lite. As such, it applies also for LW4over6 and MAP-E. This document analyzes the information required to configure the customer edge equipment for the support of multicast in the context of DS-Lite, MAP, and LW4over6 in particular.

On the basis of those analyses it specifies a number of attribute-value pairs (AVPs) to allow the necessary subscriber-site-specific configuration information to be carried in Diameter.

This document doesn't specify any new commands or Application-Ids and that the AVPs could be used for any Diameter application suitable for provisioning.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

The abbreviation "CE" denotes the equipment at the customer edge that terminates the customer end of an IPv6 transitional tunnel. This will usually be a router, but could be a host directly connected to the network.

The term "tunnel source address" is used to denote the IPv6 source address used in the outer header of packets sent from the CE through an LW4over6 transitional tunnel to the border router.

2. Description of the Parameters Required By Each Transition Method

This section reviews the parameters that need to be provisioned for each of the transition methods listed above. This enumeration provides the justification for the AVPs defined in the next section.

A means is required to indicate which transition method(s) a given subscriber is allowed to use. The approach taken in this document is to specify Grouped AVPs specific to LW4over6 and MAP-E. The operator can control which of these two transition methods a given subscriber uses by ensuring that AAA passes only the Grouped AVP relevant to that method. A Grouped AVP is unnecessary for Dual-Stack Lite, since AAA has only to provide the Fully Qualified Domain Name (FQDN) of the DS-Lite Address Family Transition Router (AFTR) (see Section 2.1). Hence when no Grouped AVP is provided either for LW4over6 or MAP-E and only the AFTR's FQDN is present, this indicates that the subscriber equipment will use the Dual-Stack Lite transition method. Provisioning of multicast is an orthogonal activity, since it is independent of the transition method.

2.1. Parameters For Dual-Stack Lite (DS-Lite)

DS-Lite is documented in [RFC6333]. The Basic Bridging BroadBand (B4) element at the customer premises needs to discover the IPv6 address of the AFTR (border router). For the reasons discussed in Section 3.2, the AAA server provision the B4 element with the AFTR's Fully Qualified Domain Name (FQDN) that is passed to a B4's IP resolution library. The AFTR's FQDN is contained in the Border-Router-Name AVP (see Section 3.2).

The B4 element could also be configured with the IPv4 address of the B4 interface facing the tunnel, with valid values from 192.0.0.2 to 192.0.0.7 and the default value of 192.0.0.2 in the absence of provisioning. Provisioning such information through AAA is problematic because it is most likely used in a case where multiple B4 instances occupy the same device. This document therefore assumes that the B4 interface address is determined by other means than AAA (implementation-dependent or static assignment).

2.2. Lightweight IPv4 Over IPv6 (LW4over6)

Light Weight IPv4 Over IPv6 (LW4over6) is documented in [I-D.ietf-softwire-lw4over6]. LW4over6 requires four items to be provisioned to the customer equipment:

- o IPv6 address of the border router.
- o IPv6 prefix used by the CE to construct the tunnel source address. In the terminology of [I-D.ietf-softwire-lw4over6], this is the IPv6 Binding Prefix.
- o an IPv4 address to be used on the external side of the CE; and
- o if the IPv4 address is shared, a specification of the port set the subscriber site is allowed to use. Please see the description in Section 2.3. For LW4over6, all three of the parameters 'a', 'k', and PSID described in that section are required. The default value of the offset parameter 'a' is 0.

As discussed in Section 4 of [I-D.ietf-softwire-lw4over6], it is necessary to synchronize this configuration with corresponding per-subscriber configuration at the border router. The border router information consists of the same public IPv4 address and port set parameters that are passed to the CE, bound together with the full /128 IPv6 address (not just the Binding Prefix) configured as the tunnel source address at the CE.

2.3. Port Set Specification

When an external IPv4 address is shared, LW4over6 and MAP-E restrict the CE to use of a subset of all available ports on the external side. Both transition methods use the algorithm defined in Appendix B of [I-D.ietf-softwire-map] to derive the values of the port numbers in the port set. This algorithm features three parameters, describing the positioning and value of the Port Set Identifier (PSID) within each port number of the generated set:

- an offset 'a' from the beginning of the port number to the first bit of the PSID;
- the length 'k' of the PSID within the port number, in bits; and
- the value of the PSID itself.

2.4. Mapping of Address and Port with Encapsulation (MAP-E)

Mapping of Address and Port with Encapsulation (MAP-E) is described in [I-D.ietf-softwire-map]. MAP-E requires the provisioning of the following per-subscriber information at the customer edge device:

- the IPv6 address of one or more border routers, or in MAP-E terminology, MAP border relays.
- the unique End-user IPv6 prefix for the customer edge device. This may be provided by AAA or acquired by other means.
- the Basic Mapping Rule for the customer edge device. This includes the following parameters:
  * the rule IPv6 prefix and length;
  * the rule IPv4 prefix and length. A prefix length of 0 indicates that the entire IPv4 address or prefix is coded in the Extended Address (EA) bits of the End-user IPv6 prefix rather than in the mapping rule.
  * the number of EA bits included in the End-user IPv6 prefix;
  * port set parameters giving the set of ports the CE is allowed to use when the IPv4 address is shared. Please see the description of these parameters in Section 2.3. At a minimum, the offset parameter 'a' is required. For MAP-E this has the default value 6. The parameters 'k' and PSID are needed if they cannot be derived from the mapping rule information and

the EA bits (final case of Section 5.2 of [I-D.ietf-softwire-map]).

- whether the device is to operate in mesh or hub-and-spoke mode;

- in mesh mode only, zero or more Forwarding Mapping Rules, described by the same set of parameters as the Basic Mapping Rule;

As indicated in Section 5, bullet 1 of the MAP-E document, a MAP CE can be provisioned with multiple End-user IPv6 prefixes, each associated with its own Basic Mapping Rule. This does not change the basic requirement for representation of the corresponding information in the form of Diameter AVPs, but adds a potential requirement for multiple instances of this information to be present in the Diameter message, differing in the value of the End-user IPv6 prefix (in contrast to the Forward Mapping Rule instances).

The border router needs to be configured with the superset of the Mapping Rules passed to the customer sites it serves. Since this requirement does not require direct coordination with CE configuration in the way LW4over6 does, it is out of scope of the present document. However, the AVPs defined here may be useful if a separate Diameter application is used to configure the border router.

2.5. Parameters For Multicast

[I-D.ietf-softwire-dslite-multicast] specifies a generic solution for delivery of IPv4 multicast services to IPv4 clients over an IPv6 multicast network. The solution can be in particular deployed in a DS-Lite context, but is also adaptable to LW4over6 and MAP-E. For example, [I-D.ietf-softwire-multicast-prefix-option] specifies how DHCPv6 [RFC3315] can be used to provision multicast-related information. The following lists the multicast-related information that needs to be provisioned:

- ASM-mPrefix64: the IPv6 multicast prefix to be used to synthesize the IPv4-embedded IPv6 addresses of the multicast groups in the Any-Source Multicast (ASM) mode. This is achieved by concatenating the ASM-mPrefix64 and a IPv4 multicast address; the Pv4 multicast address is inserted in the last 32 bits of the IPv4-embedded IPv6 multicast address.

- SSM-mPrefix64: the IPv6 multicast prefix to be used to synthesize the IPv4-embedded IPv6 addresses of the multicast groups in the Source-Specific Multicast (SSM, [RFC4607]) mode. This is achieved by concatenating the SSM-mPrefix64 and a IPv4 multicast address; the Pv4 multicast address is inserted in the last 32 bits of the IPv4-embedded IPv6 multicast address.

o uPrefix64: the IPv6 unicast prefix to be used in SSM mode for constructing the IPv4-embedded IPv6 addresses representing the IPv4 multicast sources in the IPv6 domain. uPrefix64 may also be used to extract the IPv4 address from the received multicast data flows. The address mapping follows the guidelines documented in [RFC6052].

2.6. Summary and Discussion

It appears that two items are common to the different transition methods and the corresponding AVPs to carry them can be reused:

o a representation of the IPv6 address of a border router;

o A set of prefixes for delivery of multicast services to IPv4 clients over an IPv6 multicast network.

[RFC6519] sets a precedent for representation of the IPv6 address of a border router as an FQDN. This can be dereferenced to one or more IP addresses by the provisioning system before being passed to the customer equipment, or left as an FQDN as it as in [RFC6334].

The remaining requirements are transition-method-specific:

o for LW4over6, a representation of a binding between (1) either the IPv6 Binding Prefix or a full /128 IPv6 address, (2) a public IPv4 address, and (3) (if the IPv4 address is shared) a port set identifier;

o for MAP-E, a representation of the unique End-user IPv6 prefix for the CE, if not provided by other means;

o for MAP-E, a representation of a Mapping Rule;

o for MAP-E, an indication of whether mesh mode or hub-and-spoke mode is to be used.

3. Attribute-Value Pair Definitions

This section provides the specifications for the AVPs needed to meet the requirements summarized in Section 2.6. Within the context of their usage, all of these AVPs MUST have the M bit set and the V bit cleared.

3.1. IP-Prefix-Length AVP

The IP-Prefix-Length AVP (AVP code TBD00) is of type Unsigned32. It provides the length of an IPv4 or IPv6 prefix. Valid values are from 0 to 32 for IPv4, and from 0 to 128 for IPv6. Tighter limits are given below for particular contexts of use of this AVP.

NOTE: The IP-Prefix-Length AVP is only relevant when associated with an IP-Address AVP in a Grouped AVP.

3.2. Border-Router-Name AVP

Following on the precedent set by [RFC6334] and [RFC6519], this document identifies a border router using an FQDN rather than an address. The Border-Router-Name AVP (AVP Code TBD01) is of type OctetString. The FQDN encoding MUST follow the Name Syntax defined in [RFC1035][RFC1123][RFC2181] and are represented in ASCII form.

3.3. 64-Multicast-Attributes AVP

The 64-Multicast-Attributes AVP (AVP Code TBD02) is of type Grouped. It contains the multicast-related IPv6 prefixes needed for providing IPv4 multicast over IPv6 using DS-Lite, MAP-E, or LW4over6, as mentioned in Section 2.5.

The syntax is shown in Figure 1.

```
        64-Multicast-Attributes  ::= < AVP Header: TBD02 >
                       [ ASM-mPrefix64 ]
                       [ SSM-mPrefix64 ]
                       [ Delegated-IPv6-Prefix ]
                      *[ AVP ]
```

Figure 1: 64-Multicast-Attributes AVP

64-Multicast-Attributes AVP MUST include at least the ASM-mPrefix64 AVP or the SSM-mPrefix64 AVP.

Both the ASM-mPrefix64 AVP and the SSM-mPrefix64 AVP MAY be present.

The Delegated-IPv6-Prefix AVP MUST be present when the SSM-mPrefix64 AVP is present. The Delegated-IPv6-Prefix AVP MAY be present when the ASM-mPrefix64 AVP is present.

3.3.1. ASM-mPrefix64 AVP

The ASM-mPrefix64 AVP (AVP Code TBD03) conveys the value of ASM_mPrefix64 as mentioned in Section 2.5. The ASM-mPrefix64 AVP is of type Grouped, as shown in Figure 2.

```
          ASM-mPrefix64  ::= < AVP Header: TBD03 >
                            { IP-Address }
                            { IP-Prefix-Length }
                           *[ AVP ]
```

Figure 2: ASM-mPrefix64 AVP

IP-Address (AVP code 518) is defined in [RFC5777] and is of type Address. Within the ASM-mPrefix64 AVP, it provides the value of an IPv6 prefix. The AddressType field in IP-Address MUST have value 2 (IPv6). The conveyed multicast IPv6 prefix MUST belong to the ASM range. Unused bits in IP-Address beyond the actual prefix MUST be set to zeroes by the sender and ignored by the receiver.

The IP-Prefix-Length AVP (AVP code TBD00) provides the actual length of the prefix contained in the IP-Address AVP. Within the ASM-mPrefix64 AVP, valid values of the IP-Prefix-Length AVP are from 24 to 96.

3.3.2. SSM-mPrefix64 AVP

The SSM-mPrefix64 AVP (AVP Code TBD04) conveys the value of SSM_mPrefix64 as mentioned in Section 2.5. The SSM-mPrefix64 AVP is of type Grouped, as shown in Figure 3.

```
          SSM-mPrefix64  ::= < AVP Header: TBD04 >
                            { IP-Address }
                            { IP-Prefix-Length }
                           *[ AVP ]
```

Figure 3: SSM-mPrefix64 AVP

IP-Address (AVP code 518) provides the value of an IPv6 prefix. The AddressType field in IP-Address MUST have value 2 (IPv6). The conveyed multicast IPv6 prefix MUST belong to the SSM range. Unused bits in IP-Address beyond the actual prefix MUST be set to zeroes by the sender and ignored by the receiver.

The IP-Prefix-Length AVP (AVP code TBD00) provides the actual length of the prefix contained in the IP-Address AVP. With regard to prefix length, note that Section 6 of [RFC3306] requires that bits 33-95 of an SSM address in the FF3x range be set to zero, meaning that the

prefix length for an SSM prefix is effectively 96. However, Section 1 of [RFC4607] suggests that the lower limit of 32 bits be preserved to allow potential future use of bits 33-95. Hence applications SHOULD accept prefix lengths between 32 and 96 inclusive.

3.3.3. Delegated-IPv6-Prefix AVP As uPrefix64

Within the 64-Multicast-Attributes AVP, the Delegated-IPv6-Prefix AVP (AVP Code 123) conveys the value of uPrefix64, a unicast IPv6 prefix, as mentioned in Section 2.5. The Delegated-IPv6-Prefix AVP is defined in [RFC4818]. As specified by [RFC6052], the value in the Prefix-Length field MUST be one of 32, 48, 56, 64 or 96.

3.4. Tunnel-Source-Pref-Or-Addr AVP

The Tunnel-Source-Pref-Or-Addr AVP (AVP Code TBD05) conveys either the IPv6 Binding Prefix or the tunnel source address on the CE, as described in Section 2.2. The Tunnel-Source-Pref-Or-Addr AVP is of type Grouped, with syntax as shown in Figure 4. The Tunnel-Source-Pref-Or-Addr AVP MUST contain either the Delegated-IPv6-Prefix AVP or the Tunnel-Source-IPv6-Address AVP, not both.

```
        Tunnel-Source-Pref-Or-Addr  ::= < AVP Header: TBD05 >
                       [ Delegated-IPv6-Prefix ]
                       [ Tunnel-Source-IPv6-Address ]
                      *[ AVP ]
```

Figure 4: Tunnel-Source-Pref-Or-Addr AVP

This AVP is defined separately from the LW4over6-Binding AVP (which includes it) to provide flexibility in the transport of the tunnel source address from the provisioning system to AAA while also supporting the provision of a complete binding to the LW4over6 border router.

3.4.1. Delegated-IPv6-Prefix As the IPv6 Binding Prefix

The Delegated-IPv6-Prefix AVP (AVP code 123) is of type Octetstring, and is defined in [RFC4818]. Within the Tunnel-Source-Pref-Or-Addr AVP, it conveys the IPv6 Binding Prefix assigned to the CE. Valid values in the Prefix-Length field are from 0 to 128 (full address), although a more restricted range is obviously more reasonable.

3.4.2. Tunnel-Source-IPv6-Address AVP

The Tunnel-Source-IPv6-Address AVP (AVP code TBD06) is of type Address. It provides the address that the CE has assigned to its end of an LW4over6 tunnel. The AddressType field in this AVP MUST be set to 2 (IPv6).

3.5. Port-Set-Identifier

The Port-Set-Identifier AVP (AVP Code TBD07) is a structured OctetString with four octets of data, hence a total AVP length of 12. The description of the structure which follows refers to refers to the parameters described in Section 2.3.

- o The first (high-order) octet is the Offset field. It is interpreted as an 8-bit unsigned integer giving the offset 'a' from the beginning of a port number to the beginning of the port set identifier (PSID) to which that port belongs. Valid values are from 0 to 15.

- o The next octet, the PSIDLength, is also interpreted as an 8-bit unsigned integer and gives the length 'k' in bits of the port set identifier (PSID). Valid values are from 0 to (16 - a). A value of 0 indicates that the PSID is not present (probable case for MAP-E, see Section 2.4), and the PSIDValue field MUST be ignored.

- o The final two octets contain the PSIDValue field. They give the value of the PSID itself, right-justified within the field. That is, the value of the PSID occupies the 'k' lowest-order bits of the PSIDValue field.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|    Offset     |    Length     |         PSID  Value           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

3.6. LW4over6-Binding AVP

The LW4over6-Binding AVP (AVP Code TBD08) is of type Grouped. It contains the elements of configuration that constitute the binding between an LW4over6 tunnel and IPv4 packets sent through that tunnel, as described in Section 2.2.

```
                    LW4over6-Binding  ::= < AVP Header: TBD08 >
                             { Tunnel-Source-Pref-Or-Addr }
                             { LW4over6-External-IPv4-Addr }
                             [ Port-Set-Identifier ]
                            *[ AVP ]
```

Figure 5

The Tunnel-Source-Pref-Or-Addr AVP is defined in Section 3.4 and provides either the Binding Prefix or the full IPv6 tunnel source address.

The LW4over6-External-IPv4-Addr AVP is defined in Section 3.6.1.

The Port-Set-Identifier AVP is defined in Section 3.5. It identifies the specific set of ports assigned to the LW4over6 tunnel, when the IPv4 address is being shared.

3.6.1. LW4over6-External-IPv4-Addr AVP

The LW4over6-External-IPv4-Addr AVP (AVP Code TBD09) uses the Address derived data format defined in Section 4.3.1 of [RFC6733]. It provides the CE's external IPv4 address within the LW4over6 tunnel associated with the given binding. The AddressType field MUST be set to 1 (IPv4), and the total length of the AVP MUST be 14 octets.

3.7. MAP-E-Attributes

The MAP-E-Attributes AVP (AVP Code TBD10) is of type Grouped. It contains the configuration data identified in Section 2.4 for all of the mapping rules (Basic and Forwarding) in a single MAP domain. Multiple instances of this AVP will be present if the CE belongs to multiple MAP domains.

```
                    MAP-E-Attributes  ::= < AVP Header: TBD06 >
                           1*{ Border-Router-Name }
                           1*{ MAP-Mapping-Rule }
                             [ MAP-Mesh-Mode ]
                             [ Delegated-IPv6-Prefix ]
                            *[ AVP ]
```

Figure 6

The Border-Router-Name AVP is defined in Section 3.2. It provides the FQDN of a MAP border relay at the edge of the MAP domain to which the containing MAP-E-Attributes AVP relates. At least one instance of this AVP MUST be present.

The MAP-Mapping-Rule AVP is defined in Section 3.9. At least one instance of this AVP MUST be present. If the MAP-E domain supports mesh mode (indicated by the presence of the MAP-Mesh-Mode AVP), additional MAP-Mapping-Rule instances MAY be present. If the MAP-E domain is operating in hub-and-spoke mode, additional MAP-Mapping-Rule instances MUST NOT be present.

The MAP-Mesh-Mode AVP is defined in Section 3.8. The absence of the mesh mode indicator attribute indicates that the CE is required to operate in hub-and-spoke mode.

The Delegated-IPv6-Prefix AVP (AVP Code 123) provides the End-user IPv6 prefix assigned to the CE for the MAP domain to which the containing MAP-E-Attributes AVP relates. The AVP is defined in [RFC4818]. Valid values of the Prefix-Length field range from 0 to 128.

The Delegated-IPv6-Prefix AVP is optional because, depending on deployment, the End-user IPv6 prefix may be provided by AAA or by other means. If multiple instances of the MAP-E-Attributes AVP containing the Delegated-IPv6-Prefix AVP are present, each instance of the latter MUST have a different value.

3.8. MAP-Mesh-Mode

The MAP-Mesh-Mode AVP (AVP Code TBD11) is of type Enumerated and indicates whether the CE has to operate in mesh mode or hub-and-spoke when using MAP-E. The following values are supported:

0 MESH

1 HUB_AND_SPOKE

The absence of the mesh mode indicator attribute indicates that the CE is required to operate in hub-and-spoke mode.

3.9. MAP-Mapping-Rule

The MAP-Mapping-Rule AVP (AVP Code TBD12) is of type Grouped, and is used only in conjunction with MAP-based transition methods. Mapping rules are required both by the MAP border relay and by the CE. The components of the MAP-Mapping-Rule AVP provide the contents of a mapping rule as described in Section 2.4.

The syntax of the MAP-Mapping-Rule AVP is as follows:

```
            MAP-Mapping-Rule  ::= < AVP Header: TBD12 >
                                 { Rule-IPv4-Addr-Or-Prefix }
                                 { Rule-IPv6-Prefix    }
                                 { EA-Field-Length     }
                                 { Port-Set-Identifier }
                                *[ AVP ]
```

Figure 7

The Rule-IPv4-Addr-Or-Prefix, Rule-IPv6-Prefix, EA-Field-Length, and Port-Set-Identifier AVPs MUST all be present.

The Port-Set-Identifier AVP provides information to identify the specific set of ports assigned to the CE. For more information see Section 2.4 and Section 2.3. The Port-Set-Identifier AVP is defined in Section 3.5.

3.9.1. Rule-IPv4-Addr-Or-Prefix AVP

The Rule-IPv4-Addr-Or-Prefix AVP (AVP Code TBD13) conveys the rule IPv4 prefix and length as described in Section 2.4. The Rule-IPv4-Addr-Or-Prefix AVP is of type Grouped, as shown in Figure 8.

```
             Rule-IPv4-Addr-Or-Prefix  ::= < AVP Header: TBD13 >
                            { IP-Address }
                            { IP-Prefix-Length }
                           *[ AVP ]
```

Figure 8: Rule-IPv4-Addr-Or-Prefix AVP

IP-Address (AVP code 518) is defined in [RFC5777] and is of type Address. Within the Rule-IPv4-Addr-Or-Prefix AVP, it provides the value of a unicast IPv4 address or prefix. The AddressType field in IP-Address MUST have value 1 (IPv4). Unused bits in IP-Address beyond the actual prefix MUST be set to zeroes by the sender and ignored by the receiver.

The IP-Prefix-Length AVP (AVP code TBD00) provides the actual length of the prefix contained in the IP-Address AVP. Within the Rule-IPv4-Addr-Or-Prefix AVP, valid values of the IP-Prefix-Length AVP are from 0 to 32 (full address), based on the different cases identified in Section 5.2 of [I-D.ietf-softwire-map].

3.9.2. Rule-IPv6-Prefix AVP

The Rule-IPv6-Prefix AVP (AVP Code TBD14) conveys the rule IPv6 prefix and length as described in Section 2.4. The Rule-IPv6-Prefix AVP is of type Grouped, as shown in Figure 9.

```
             Rule-IPv6-Prefix  ::= < AVP Header: TBD14 >
                         { IP-Address }
                         { IP-Prefix-Length }
                        *[ AVP ]
```

Figure 9: Rule-IPv6-Prefix AVP

IP-Address (AVP code 518) is defined in [RFC5777] and is of type Address. Within the Rule-IPv6-Prefix AVP, it provides the value of a unicast IPv6 prefix. The AddressType field in IP- Address MUST have value 2 (IPv6). Unused bits in IP-Address beyond the actual prefix MUST be set to zeroes by the sender and ignored by the receiver.

The IP-Prefix-Length AVP (AVP code TBD00) provides the actual length of the prefix contained in the IP-Address AVP. Within the Rule-IPv6-Prefix AVP, the minimum valid prefix length is 0. The maximum value is bounded by the length of the End-user IPv6 prefix associated with the mapping rule, if present in the form of the Delegated-IPv6-Prefix AVP in the enclosing MAP-E-Attributes AVP. Otherwise the maximum value is 128.

3.9.3. EA-Field-Length AVP

The EA-Field-Length AVP (AVP Code TBD15) is of type Unsigned32. Valid values range from 0 to 48. See Section 5.2 of [I-D.ietf-softwire-map] for a description of the use of this parameter in deriving IPv4 address and port number configuration.

4. Attribute Value Pair flag rules

Attribute Name	AVP Code	Section Defined	Value Type	AVP flag rules: MUST	AVP flag rules: MUST NOT
IP-Prefix-Length	TBD00	3.1	Unsigned32		V
Border-Router-Name	TBD01	3.2	OctetString		V
64-Multicast-Attributes	TBD02	3.3	Grouped		V
ASM-mPrefix64	TBD03	3.3.1	Grouped		V
SSM-mPrefix64	TBD04	3.3.2	Grouped		V
Tunnel-Source-Pref-Or-Addr	TBD05	3.4	Grouped		V
Tunnel-Source-IPv6-Address	TBD06	3.4.2	Address		V
Port-Set-Identifier	TBD07	3.5	OctetString		V
LW4over6-Binding	TBD08	3.6	Grouped		V
LW4over6-External-IPv4-Addr	TBD09	3.6.1	Address		V
MAP-E-Attributes	TBD10	3.7	Grouped		V
MAP-Mesh-Mode	TBD11	3.8	Enumerated		V
MAP-Mapping-Rule	TBD12	3.9	Grouped		V
Rule-IPv4-Addr-Or-Prefix	TBD13	3.9.1	Grouped		V
Rule-IPv6-Prefix	TBD14	3.9.2	Grouped		V
EA-Field-Length	TBD15	3.9.3	Unsigned32		V

As described in the Diameter base protocol [RFC6733], the M-bit usage for a given AVP in a given command may be defined by the application.

5. Acknowledgements

Huawei Technologies funded Tom Taylor's work on earlier versions of this document.

Special thanks to Lionel Morand for the detailed review.

6. IANA Considerations

This memo requests to IANA to register the following Diameter AVP codes:

Code	Attribute Name	Reference
TBD00	IP-Prefix-Length	This document
TBD01	Border-Router-Name	This document
TBD02	64-Multicast-Attributes	This document
TBD03	ASM-mPrefix64	This document
TBD04	SSM-mPrefix64	This document
TBD05	Tunnel-Source-Pref-Or-Addr	This document
TBD06	Tunnel-Source-IPv6-Address	This document
TBD07	Port-Set-Identifier	This document
TBD08	LW4over6-Binding	This document
TBD09	LW4over6-External-IPv4-Addr	This document
TBD10	MAP-E-Attributes	This document
TBD11	MAP-Mesh-Mode	This document
TBD12	MAP-Mapping-Rule	This document
TBD13	Rule-IPv4-Addr-Or-Prefix	This document
TBD14	Rule-IPv6-Prefix	This document
TBD15	EA-Field-Length	This document

Table 1

7. Security Considerations

The AVPs defined in this document face two threats, both dependent on man-in-the-middle attacks on the Diameter delivery path. The more serious threat is denial of service through modification of the AVP contents leading to misconfiguration. The lesser threat is disclosure of subscriber addresses allowing the attacker to track subscriber activity.

Diameter security is currently provided on a hop-by-hop basis (see Section 2.2 of [RFC6733]). The Diameter end-to-end security problem has not been solved, so man-in-the-middle attacks on Diameter peers

along the path are possible. The present document does not propose to solve that general problem, but simply warn that it exists.

8. References

8.1. Normative References

[I-D.ietf-softwire-lw4over6]
Cui, Y., Sun, Q., Boucadair, M., Tsou, T., Lee, Y., and I. Farrer, "Lightweight 4over6: An Extension to the DS-Lite Architecture (work in progress)", March 2014.

[I-D.ietf-softwire-map]
Troan, O., Dec, W., Li, X., Bao, C., Matsushima, S., Murakami, T., and T. Taylor, "Mapping of Address and Port with Encapsulation (MAP) (work in progress)", January 2014.

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.

[RFC1123] Braden, R., "Requirements for Internet Hosts - Application and Support", STD 3, RFC 1123, October 1989.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2181] Elz, R. and R. Bush, "Clarifications to the DNS Specification", RFC 2181, July 1997.

[RFC3306] Haberman, B. and D. Thaler, "Unicast-Prefix-based IPv6 Multicast Addresses", RFC 3306, August 2002.

[RFC4818] Salowey, J. and R. Droms, "RADIUS Delegated-IPv6-Prefix Attribute", RFC 4818, April 2007.

[RFC5777] Korhonen, J., Tschofenig, H., Arumaithurai, M., Jones, M., and A. Lior, "Traffic Classification and Quality of Service (QoS) Attributes for Diameter", RFC 5777, February 2010.

[RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, August 2011.

[RFC6733] Fajardo, V., Arkko, J., Loughney, J., and G. Zorn, "Diameter Base Protocol", RFC 6733, October 2012.

8.2. Informative References

[I-D.ietf-softwire-dslite-multicast]
Qin, J., Boucadair, M., Jacquenet, C., Lee, Y., and Q. Wang, "Delivery of IPv4 Multicast Services to IPv4 Clients over an IPv6 Multicast Network", draft-ietf-softwire-dslite-multicast-09 (work in progress), March 2015.

[I-D.ietf-softwire-map-dhcp]
Mrugalski, T., Troan, O., Farrer, I., Perrault, S., Dec, W., Bao, C., Yeh, L., and X. Deng, "DHCPv6 Options for configuration of Softwire Address and Port Mapped Clients (Work in progress)", March 2014.

[I-D.ietf-softwire-multicast-prefix-option]
Boucadair, M., Qin, J., Tsou, T., and X. Deng, "DHCPv6 Option for IPv4-Embedded Multicast and Unicast IPv6 Prefixes", draft-ietf-softwire-multicast-prefix-option-08 (work in progress), March 2015.

[RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.

[RFC4607] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", RFC 4607, August 2006.

[RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", RFC 6052, October 2010.

[RFC6334] Hankins, D. and T. Mrugalski, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite", RFC 6334, August 2011.

[RFC6519] Maglione, R. and A. Durand, "RADIUS Extensions for Dual-Stack Lite", RFC 6519, February 2012.

Authors' Addresses

Cathy Zhou
Huawei Technologies
Bantian, Longgang District
Shenzhen 518129
P.R. China

Email: cathy.zhou@huawei.com

Tom Taylor
PT Taylor Consulting
Ottawa
Canada

Email: tom.taylor.stds@gmail.com

Qiong Sun
China Telecom
P.R.China

Phone: 86 10 58552936
Email: sunqiong@ctbri.com.cn

Mohamed Boucadair
France Telecom
Rennes 35000
France

Email: mohamed.boucadair@orange.com