Softwire WG
Internet-Draft
Intended status: Standards Track
Expires: July 22, 2013

M. Boucadair
France Telecom
I. Farrer
Deutsche Telekom

January 18, 2013

Unified IPv4-in-IPv6 Softwire CPE
draft-bfmk-softwire-unified-cpe-02

Abstract

Transporting IPv4 packets encapsulated in IPv6 is a common solution to the problem of IPv4 service continuity over IPv6-only provider networks. A number of differing functional approaches have been developed for this, each having their own specific characteristics. As these approaches share a similar functional architecture and use the same data plane mechanisms, this memo describes a specification whereby a single CPE can interwork with all of the standardized and proposed approaches to providing encapsulated IPv4 in IPv6 services.

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 RFC 2119 [RFC2119].

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on July 22, 2013.

Copyright Notice

Table of Contents

1. Introduction

IPv4 service continuity is one of the major technical challenges which must be considered during IPv6 migration. Over the past few years, a number of different approaches have been developed to assist with this problem. These approaches, or modes, exist in order to meet the particular deployment, scaling, addressing and other requirements of different service provider's networks. Section 2 of this document describes these approaches in more detail.

A common feature shared between all of the differing modes is the integration of softwire tunnel end-point functionality into the CPE router. Due to this inherent data plane similarity, a single CPE may be capable of supporting several different approaches. Users may also wish to configure a specific mode of operation.

A service provider's network may also have more than one mode enabled in order to support diverse CPE client functionality, during migration between modes or where services require specific supporting softwire architectures.

For softwire based services to be successfully established, it is essential that the customer end-node, the service provider end-node and provisioning systems are able to indicate their capabilities and preferred mode of operation.

This memo describes the logic required by both the CPE tunnel end-node and the service provider's provisioning infrastructure so that softwire services can be provided in mixed-mode environments.

1.1. Rationale

The following rationale has been adopted for this document:

(1) Describe the functionality of each the different solution modes and provide clear distinctions between them
(2) Simplify solution migration paths: Define unified CPE behavior, allowing for smooth migration between the different modes
(3) Deterministic CPE co-existence behavior: Specify the behavior when several modes co-exist in the CPE
(4) Deterministic service provider co-existence behavior: Specify the behavior when several modes co-exist in the service providers network
(5) Re-usability: Maximize the re-use of existing functional blocks including tunnel end-points, port restricted NAPT44, forwarding behavior, etc.

(6) Solution agnostic: Adopt neutral terminology and avoid (as far as possible) overloading the document with solution-specific terms
(7) Flexibility: Allow operators to compile CPE software only for the mode(s) necessary for their chosen deployment context(s)
(8) Simplicity: Provide a model that allows operators to only implement the specific mode(s) that they require without the additional complexity of unneeded modes.

2. IPv4 Service Continuity Architectures: A 'Big Picture' Overview

The solutions which have been proposed within the Softwire WG can be categorized into three main functional approaches, differentiated by the amount and type of state that the service provider needs to maintain within their network:

(1) Full stateful approach (DS-Lite, [RFC6333]): Requires per-session state to be maintained in the Service Provider's network.
(2) Binding approach (e.g., Lightweight 4over6 (Lw4o6) [I-D.cui-softwire-b4-translated-ds-lite][I-D.ietf-softwire-public-4over6] or MAP 1:1 [I-D.ietf-softwire-map]): Requires a single per-subscriber state (or a few) to be maintained in the Service Provider's network.
(3) Full stateless approach (MAP, [I-D.ietf-softwire-map]): Does not require per-session or per-subscriber state to be maintained in the Service Provider's network.

All these approaches share a similar architecture, with a tunnel endpoint located in the CPE and a remote tunnel endpoint. All use IPv6 as the transport protocol for the delivery of an IPv4 connectivity service using an IPv4-in-IPv6 encapsulation scheme [RFC2473].

Several cases can be envisaged:

1. The CPE is complied to support only one mode: No issue is raised by this case.
2. The CPE supports several modes but only one mode is explicitly configured: No issue is raised by this case.
3. The CPE supports several modes but no mode is explicitly enabled: the CPE will need additional triggers to decide which mode to activate.
4. The CPE supports several modes and several modes are configured: the CPE will need additional triggers to decide which mode to activate.

As this document describes a provisioning profile whereby a single CPE could be capable of supporting any, or multiple modes, the

customer should not be required to have any knowledge of the capabilities and configuration of their CPE, or of their service provider's network.

The service provider, however, may have only a single mode enabled, or may have multiple modes, but with one preferred mode. For this reason, it is necessary to approach the configuration of CPEs from the standpoint of the service provider's network capabilities.

2.1. Functional Elements

The functional elements for each of the solution modes are listed in Table 1:

Mode	Customer side	Network side
DS-Lite	B4	AFTR
Lw4o6	lwB4	lwAFTR
MAP	MAP CE	MAP BR

Table 1: Functional Elements

Table 2 describes each functional element:

Functional Element	Description
B4	An IPv4-in-IPv6 tunnel endpoint; the B4 creates a tunnel to a pre-configured remote tunnel endpoint.
AFTR	Provides both an IPv4-in-IPv6 tunnel endpoint and a NAT44 function implemented in the same node.
lwB4	A B4 which supports port-restricted IPv4 addresses. An lwB4 MAY also provide a NAT44 function.
lwAFTR	An IPv4-in-IPv6 tunnel endpoint which maintains per-subscriber address binding. Unlike the AFTR, it MUST NOT perform a NAPT44 function.
MAP CE	A B4 which supports port-restricted IPv4 addresses. It MAY be co-located with a NAT44. A MAP CE forwards IPv4-in-IPv6 packets using provisioned mapping rules to derive the remote tunnel endpoint.
MAP BR	An IPv4-in-IPv6 tunnel endpoint. A MAP BR forwards IPv4-in-IPv6 packets following pre-configured mapping rules.

Table 2: Required Element Functionality

Table 3 identifies features required by the customer end-node.

Functional Element	IPv4-in-IPv6 tunnel endpoint	Port-restricted IPv4	Port-restricted NAT44
B4	Yes	N/A	No
lwB4	Yes	Yes	Optional
MAP-E CE	Yes	Yes	Optional

Table 3: Supported Features

2.2. Required Provisoning Information

Table 4 identifies the provisioning information required for each solution mode.

Mode	Provisioning Information
DS-Lite	Remote IPv4-in-IPv6 Tunnel Endpoint Address
Lw4o6	Remote IPv4-in-IPv6 Tunnel Endpoint Address
	IPv4 Address
	Port Set
MAP-E	Mapping Rules
	MAP Domain Parameters

Table 4: Provisioning Information

Note: MAP Mapping Rules are translated into the following configuration parameters: Set of remote IPv4-in-IPv6 tunnel endpoint addresses, IPv4 address and port set.

Note: Required provisioning information for each mode may also be represented as follows:

DS-Lite: - Remote IPv4-in-IPv6 Tunnel Endpoint

Lw4o6: - DS-Lite set of provisioning information
- IPv4 address
- Port set

MAP-E: - Lw4o6 set of provisioning information

- Forwarding mapping rules

3. Unified Softwire CPE Behaviour

This section specifies a unified CPE behavior capable of supporting any one, or combination of, the three modes.

3.1. IPv4 Address Functional Requirements

The following two requirements must be met by the functional elements:

Full IPv4 Address Assingment All the aforementioned modes MUST be designed to allow either a full or a shared IPv4 address to be assigned to a customer end-node. DS-Lite and MAP-E fulfill this requirement. With minor changes, the [I-D.cui-softwire-b4-translated-ds-lite] specification can be updated to assign full IPv4 addresses.

Customer End-Node NAT A NAT function within the customer end-node is not required for DS-Lite, while it is optional for both MAP-E and Lw4o6. When NAT is enabled for MAP-E or Lw4o6, the customer end-node NAT MUST be able to restrict the external translated source ports to the set of ports that it has been provisioned with.

3.2. Generic CPE Bootstrapping Logic

The generic provisioning logic is designed to meet the following requirements:

- When several service continuity modes are supported by the same CPE, it MUST be possible to configure a single mode for use.

- For each network attachment, the end-node MUST NOT activate more than one mode.

- The CPE MAY be configured by a user or via remote device management means (e.g., DHCP, TR-069).

- A network which supports one or several modes MUST return valid configuration data enabling requesting devices to unambiguously select a single mode to use for attachment.

- A CPE which supports only one mode or it is configured to enable only mode MUST ignore any configuration parameter which is not required for the mode it supports.

This section sketches a generic algorithm to be followed by a CPE supporting one or more of the modes listed above. Based on the retrieved information, the CPE will determine which mode to activate.

(1) If a given mode is enabled (DS-Lite, Lw4o6 or MAP-E), the CPE MUST be configured with the required provisioning information listed in Table 4. If all of the required information is not available locally, the CPE MUST use available provisioning means (e.g., DHCP) to retrieve the missing configuration data.

(2) If the CPE supports several modes, but no mode is explicitly enabled, the CPE MUST use available provisioning means (e.g., DHCP) to retrieve available configuration parameters and use the availability of individual parameters to ascertain which functional mode to configure:

  (2.1) If only a Remote IPv4-in-IPv6 Tunnel Endpoint is received, the CPE MUST proceed as follows:
    (2.1.1) IPv4-in-IPv6 tunnel endpoint initialization is defined in [RFC6333].
    (2.1.2) Outbound IPv4 packets are forwarded to the next hop as specified in Section 3.4.

  (2.2) If a Remote IPv4-in-IPv6 Tunnel Endpoint, an IPv4 Address and optionally a Port Set are received, the CPE MUST behave as follows:
    (2.2.1) IPv4-in-IPv6 tunnel endpoint initialization is similar to the B4 [RFC6333].
    (2.2.2) When NAPT44 is required (e.g., because the CPE is a router), a NAPT44 module is enabled.
    (2.2.3) The tunnel endpoint address is selected from the native IPv6 addresses configured on the CPE. No particular considerations are required to be taken into account to generate the Interface Identifier.
    (2.2.4) When a port set is provisioned, the external source ports MUST be restricted to the provisioned set of ports.
    (2.2.5) After translation, outbound IPv4 packets are forwarded to the next hop as specified in Section 3.4.

  (2.3) If Mapping Rule(s) are received, the CPE MUST behave as follows:

(2.3.1) IPv4-in-IPv6 tunnel endpoint initialization is similar to the B4 [RFC6333].

(2.3.2) The tunnel endpoint is assigned with an IPv6 address which includes an IPv4 address. The MAP Interface Identifier is based on the format specified in Section 2.2 of [RFC6052].

(2.3.3) When NAPT44 is required (e.g., because the CPE is a router), a NAPT44 module is enabled.

(2.3.4) When a port set is provisioned, the external source port MUST be restricted to the provisioned set of ports.

(2.3.5) After translation, outbound IPv4 packets then forwarded to the next hop as specified in Section 3.4.

3.3. Customer Side DHCP Based Provisioning

[DISCUSSION NOTE:

1. This section will be updated to reflect the consensus from DHC WG.

2. As it is proposed that OPTION_MAP would be used for all new softwire provisioning, should we rename OPTION_MAP to OPTION_SW (incl. the associated sub-options)?]

]

DHCP-based configuration SHOULD be implemented by the customer end-node using the following two DHCP options:

OPTION_AFTR_NAME [RFC6334] Provides the FQDN for the remote IPv4-in-IPv6 tunnel end-point.

OPTION_MAP [I-D.ietf-softwire-map-dhcp] Provides IPv4-related configuration for binding mode and/or mapping rules for stateless mode (including MAP parameters such as offset, domain prefix, etc.). OPTION_MAP_BIND is a sub-option used to convey an IPv4 address (for example, encoded as an IPv4-mapped IPv6 address [RFC4291]). This address is used when binding mode is enabled. The receipt of OPTION_MAP_BIND is an implicit indication to the customer side device to operate in binding, rather than stateless mode.

The customer end-node uses the DHCP Option Request Option (ORO) to request either one or both of these options depending on which modes

it is capable of and configured to support.

The DHCP option(s) sent in the response allow the service provider to inform the customer end-node which operating mode to enable.

The following table shows the different DHCP options (and sub-options) that the service provider can supply in a response.

DHCP Option	Stateful Mode	Binding Mode	Stateless Mode
OPTION_AFTR_NAME	Yes	Yes	Optional
OPTION_MAP_BIND	No	Yes	No
OPTION_MAP_RULE	No	No	Yes
OPTION_MAP_PORTPARAMS	No	Optional	Optional

Table 5: DHCP Option Provisioning Matrix

The customer side device MUST interpret the received DHCP configuration parameters according to the logic defined in Section 3.2:

- If only OPTION_AFTR_NAME is received, then the device MUST operate in stateful mode
- If both OPTION_AFTR_NAME and OPTION_MAP_BIND are received then the device MUST operate in binding mode
- If one or more OPTION_MAP_RULE options are received, then the customer side device MUST operate in stateless mode
- If both OPTION_AFTR_NAME and OPTION_MAP_RULE(s) are received, then the customer side device MUST operate as a MAP CE. OPTION_AFTR_NAME provides the FQDN of the MAP BR.
- If OPTION_MAP_PORTPARAMS is received as a sub-option to either OPTION_MAP_BIND or OPTION_MAP_RULE, then NAPT44 MUST be configured using the supplied port-set for external translated source ports.

From the service providers side, the following rule MUST be followed:

- The DHCP server MUST NOT send both OPTION_MAP_BIND and OPTION_MAP_RULE in a single OPTION_MAP response.

3.4. Forwarding Action by the Customer End-Node

For all modes, the longest prefix match algorithm MUST be enforced to forward outbound IPv4 packets.

Specifically, this algorithm will:

- Always return the address of the AFTR for the DS-Lite mode.
- Always return the address of the lwAFTR for the binding mode.
- Return the next hop according to the pre-configured mapping rules for the stateless mode (i.e., MAP-E).

4. Security Considerations

Security considerations discussed in Section 7 of [I-D.ietf-softwire-stateless-4v6-motivation] and Section 11 of [RFC6333] should be taken into account.

5. IANA Considerations

This document does not require any action from IANA.

6. Acknowledgements

Many thanks to T. Tsou, S. Perrault, S. Sivakumar, O. Troan, W. Dec, M. Chen, for their review and comments.

Special thanks to S. Krishnan for the suggestions and guidance.

7. References

7.1. Normative References

[I-D.cui-softwire-b4-translated-ds-lite]
Cui, Y., Sun, Q., Boucadair, M., Tsou, T., Lee, Y., and I. Farrer, "Lightweight 4over6: An Extension to the DS-Lite Architecture", draft-cui-softwire-b4-translated-ds-lite-09 (work in progress), October 2012.

[I-D.ietf-softwire-map]
Troan, O., Dec, W., Li, X., Bao, C., Matsushima, S., and T. Murakami, "Mapping of Address and Port with

Encapsulation (MAP)", draft-ietf-softwire-map-02 (work in progress), September 2012.

[I-D.ietf-softwire-map-dhcp]
Mrugalski, T., Troan, O., Bao, C., Dec, W., and L. Yeh, "DHCPv6 Options for Mapping of Address and Port", draft-ietf-softwire-map-dhcp-01 (work in progress), August 2012.

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

[RFC2473] Conta, A. and S. Deering, "Generic Packet Tunneling in IPv6 Specification", RFC 2473, December 1998.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.

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

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

7.2. Informative References

[I-D.ietf-softwire-public-4over6]
Cui, Y., Wu, J., Wu, P., Vautrin, O., and Y. Lee, "Public IPv4 over IPv6 Access Network", draft-ietf-softwire-public-4over6-04 (work in progress), October 2012.

[I-D.ietf-softwire-stateless-4v6-motivation]
Boucadair, M., Matsushima, S., Lee, Y., Bonness, O., Borges, I., and G. Chen, "Motivations for Carrier-side Stateless IPv4 over IPv6 Migration Solutions", draft-ietf-softwire-stateless-4v6-motivation-05 (work in progress), November 2012.

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

Authors' Addresses

Mohamed Boucadair
France Telecom
Rennes
France

Email: mohamed.boucadair@orange.com

Ian Farrer
Deutsche Telekom
Germany

Email: ian.farrer@telekom.de