

Softwire Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 10, 2015

I. Farrer
Q. Sun
Deutsche Telekom AG
March 9, 2015

Multiple Tunnel Endpoints on Border Router
draft-farrer-softwire-br-multiendpoints-00

Abstract

This document defines a mechanism to enable an IPv4-over-IPv6 Softwire Border Router to support multiple tunnel endpoint on one BR instance. This feature can be used to group users based on IPv6 tunnel endpoint addresses and achieve high availability.

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](#)].

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 September 10, 2015.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents

Internet-Draft

BR Multi-endpoints

March 2015

carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Changes to BR's Behavior	3
3.	Changes to Initiator's Behavior	4
4.	Operational Considerations	4
5.	Security Considerations	4
6.	IANA Considerations	5
7.	Acknowledgements	5
8.	References	5
8.1.	Normative References	5
8.2.	Informative References	5
	Authors' Addresses	6

[1.](#) Introduction

The Softwire WG has developed a number of IPv4-over-IPv6 transition mechanisms that utilize a hub-and-spoke network architecture. Although the schema for configuring an BR varies according to the mechanism being implemented (using either a per-subscriber rule or an algorithmic mapping rule), in their current shape all mechanisms only allow for a single IPv6 address to be used as the BR's IPv6 tunnel address. This address is provisioned to all Softwire Initiators to use as the destination address for their IPv4-in-IPv6 tunneled traffic and is used by the BR as the source address for encapsulated traffic being sent to Softwire initiators.

The inherent limitation in having only a single IPv6 tunnel endpoint address available for the BR is that it is not possible to differentiate client's tunneled traffic based on BR's address in the encapsulating IPv6 packet. However, by implementing multiple IPv6 tunnel endpoint addresses, the BR can support different classes of users, grouped by their tunnel endpoint address. Grouping clients based on a common tunnel endpoint address makes it simple for intermediate IPv6 network elements to identify client's traffic group by examining the encapsulating IPv6 header, e.g. so that traffic forwarding policies can be applied.

It also allows for flexible, anycast based geographical resilience models where each BR supports a primary group of users and a secondary group, differentiated by the tunnel endpoint address. Traffic is flexibly routed through auto-routing protocols and Equal-

Cost Multipath (ECMP).

This document describes a method that enables one Border Router to serve multiple groups of clients. The BR's mapping/binding table must hold an additional "BR source IPv6 address" field for each Softwire Initiator it is configured to support.

This mechanism can be applied to lw4over6 [[I-D.ietf-softwire-lw4over6](#)], MAP-E [[I-D.ietf-softwire-map](#)] and MAP-T [[I-D.ietf-softwire-map-t](#)].

DISCUSSION - Is this necessary for MAP BRs, or can this already be supported?

[2.](#) Changes to BR's Behavior

Existing BRs implementing lw4o6, MAP-E or MAP-T are provisioned with a set of rules defining packet processing behaviour. The rule/binding table on the Border Router only contains the mapping between the IPv6 and IPv4 address and source ports for the Softwire Initiators. In this mechanism, the rule/binding table is extended to include the IPv6 tunnel address(es) configured on the BR as another field. The Softwire Initiators' IPv6-IPv4 mapping rules are then linked to the related BR's IPv6 tunnel addresses. As such, it is possible for one BR to serve multiple groups of Softwire Initiators, independently from each other.

On receiving an IPv6 packet, the BR MUST use both the source and the destination IPv6 addresses as input parameters to search for a matching entry in the mapping rule table, instead of only using the source IPv6 address/prefix information. If a successful match is made, the encapsulated/translated IPv4 packet is then verified as documented in related Softwire mechanisms.

When the BR receives a packet from the IPv4 Internet, it looks up for the matching entry using the destination IPv4 address and port

number. The BR MUST also retrieve the associated BR's IPv6 address to use for the encapsulating packet's source IPv6 address. Depending on the implemented mechanism, the mapping rule for constructing the destination IPv6 address will need to be retrieved as normal.

The rest of encapsulation/decapsulation/translation process is aligned with the related mechanisms.

[3.](#) Changes to Initiator's Behavior

The Software Initiator's behavior is identical to that in the related mechanisms. That is, when it receives an IPv4-in-IPv6 packet it checks the source and destination addresses against its configuration. The source address of the packet MUST match the BR's tunnel endpoint address configured on the client.

[4.](#) Operational Considerations

Border Routers need to be provisioned with multiple sets of tunnel endpoint IPv6 addresses, IPv4-IPv6 mapping rules for Software Initiators and routable IPv4 prefixes. The provisioning mechanisms could include NETCONF, TR69 or out-of-band static configuration. This mechanism is out of scope for this document.

BRs implementing this mechanism can be deployed using IPv6 anycast to achieve high availability. Since multiple IPv6 anycast addresses can be configured on the BR as tunnel endpoint addresses, a BR is able to serve one primary domain while serving other domains as backup. The BR advertises the IPv6 anycast prefix(es), as well as the routable IPv4 prefix(es). ECMP can be used to leverage for stateless load-balancing across multiple BRs.

However, as the reachable IPv4 customer prefixes are being advertised by all instances serving that domain simultaneously, IPv4 traffic which ingresses the network will, by default, use the cluster which has the lowest routing metric to the ingress point in the network. This may result in different paths for egress and ingress traffic.

Whilst stateless and per-subscriber state software mechanisms (described in [[RFC6269](#)]) don't require the ingress/egress traffic paths to be symmetrical, it may be desirable for an operator to engineer this way for effective capacity planning. The exact mechanism for achieving this will be dependant on the network's topology and how the operator utilizes equal-cost multipath based load balancing.

NOTE: One possible other consideration is that as there is an additional lookup action that needs to be carried out for packets at the BR, there may be a packet processing overhead.

[5.](#) Security Considerations

TBD

Farrer & Sun

Expires September 10, 2015

[Page 4]

Internet-Draft

BR Multi-endpoints

March 2015

[6.](#) IANA Considerations

This document does not include an IANA request.

[7.](#) Acknowledgements

The authors would like to thank Madhusudhan Vadde for contributions to this work.

[8.](#) References

[8.1.](#) Normative References

[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.

[8.2.](#) Informative References

[I-D.ietf-software-lw4over6]

Cui, Y., Qiong, Q., Boucadair, M., Tsou, T., Lee, Y., and I. Farrer, "Lightweight 4over6: An Extension to the DS-Lite Architecture", [draft-ietf-softwire-lw4over6-13](#) (work in progress), November 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)", [draft-ietf-softwire-map-13](#) (work in progress), March 2015.

[I-D.ietf-softwire-map-dhcp]

Mrugalski, T., Troan, O., Farrer, I., Perreault, S., Dec, W., Bao, C., Yeh, L., and X. Deng, "DHCPv6 Options for configuration of Softwire Address and Port Mapped Clients", [draft-ietf-softwire-map-dhcp-12](#) (work in progress), March 2015.

[I-D.ietf-softwire-map-t]

Li, X., Bao, C., Dec, W., Troan, O., Matsushima, S., and T. Murakami, "Mapping of Address and Port using Translation (MAP-T)", [draft-ietf-softwire-map-t-08](#) (work in progress), December 2014.

[RFC6269] Ford, M., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", [RFC 6269](#), June 2011.

Authors' Addresses

Ian Farrer
Deutsche Telekom AG
CTO-ATI, Landgrabenweg 151
Bonn, NRW 53227
Germany

Email: ian.farrer@telekom.de

Qi Sun
Deutsche Telekom AG
CTO-ATI, Landgrabenweg 151
Bonn, NRW 53227
Germany

Email: qui.sun@external.telekom.de