A Classification and Evaluation of Approaches to Transitional Multicast
            draft-tsou-multicast-transition-taxonomy-01

Abstract

   A number of different contributions to the IETF make proposals in
   support of multicast during the transition from IPv4 to IPv6. This
   document provides a taxonomic framework to make it easier to see how
   the different proposals relate to each other. It analyzes the
   current work in progress in the light of this framework and draws a
   number of conclusions regarding how this work should move forward in
   the BEHAVE and SOFTWIRES Working Groups.

Table of Contents

1. Introduction

As evidenced by the (probably incomplete) set of references shown below, the handling of multicast during the transition from IPv4 to IPv6 is the focus of a considerable amount of activity. This has caused some difficulty within the BEHAVE and SOFTWIRES Working Groups, where the question has been how to select the drafts that should be Working Group work items, and which drafts would fit together and should be combined.

This document introduces a framework for classification of the subject matter of the various drafts dealing with multicast transition. It applies the framework to the drafts shown in the references, and draws conclusions from the results of the comparison.

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 RFC 2119 [RFC2119]. However, this document is purely descriptive, and thus uses no requirements language.

2. The Classification Framework

The classification framework proposed here consists of four dimensions: the scenario(s) considered, the multicast topologies considered, the translation techniques used, and whether multicast content is encapsulated (aside from the tunneling used during the PIM registration stage).

Scenarios
   The scenarios considered in the various multicast transition dicussions are either two-network (A, B) or three-network (A, B, C), where A, B, and C are multicast-capable networks.

   A and B always support different versions of IP, and in the three-network case C supports the same IP version as A. Thus with respect to IP version, the two-network cases can be classified as IPv4 receiving IPv6 (4 <- 6) and IPv6 receiving IPv4 (6 <- 4). The three-network cases can be classified as 4-6-4 or 6-4-6.

Multicast Topology
   This refers to whether a contribution considers source specific multicast (SSM), any-source multicast (ASM), or both.

Translation Technique
   Address translation is always required at points where the IP
   version changes, unless signalling and content are merely tunneled
   across an intervening network rather than making use of that
   network's native multicast capabilities. Different contributions
   go into different levels of detail regarding the translation
   schemes used for multicast source and group addresses. A basic
   distinction is between stateful and stateless methods.

Use of tunneling
   Some discussions deal only with translation. Others assume
   tunneling of the multicast content, which necessarily implies a
   three-network scenario. Translation is still required in this
   case, to map between the <source, group> address pair used in
   network B and the address pair used in networks A and C for the
   same multicast stream. Tunneling does save a translation step for
   multicast content reaching network A, compared with native
   transport across network B.

[3](). **Classification of the References According to the Framework**

   As an initial step, the following is a summary of the multicast
   references covered in this document. The ordering of these documents
   is that provided by the Datatracker search tool and is presumably
   related to the date version -00 of each document was submitted.

   1.   An IPv4 - IPv6 multicast translator [ID.venaas-mcast46]

        The initial version of this draft appeared in 2008, making it
        the original source of the ideas appearing in the other drafts
        described here. The basic scenario considered is two-network,
        though three-network examples are considered (e.g., IPv4 host
        sending to an IPv4 group through an IPv6 network). The document
        discusses both 4 <- 6 and 6 <- 4, but the specific translation
        mechanism proposed, embedding of IPv4 multicast addresses in
        IPv6, places limits on the 4 <- 6 scenario. Both ASM and SSM
        are considered. Alternatives are suggested for mapping of
        unicast source addresses. For ASM, the translator is assumed to
        be a rendezvous point on the IPv6 side. Tunneling is not
        discussed.

   2.   Framework for IPv4/IPv6 Multicast Translation
        [ID.venaas-v4v6mc-framework]

        The initial part of this paper looks at a set of two-network
        scenarios, both 4 <- 6 and 6 <- 4, with an additional variable
        being whether network A or network B (but not both at once) is

the Internet. For each scenario it considers the applicable
mechanisms, which are those considered in [ID.venaas-mcast46]
but with further considerations and suggestions in the more
difficult cases. For this later paper, [RFC6052] is available
as a resource for translation of IPv4 source addresses to IPv6.
Both ASM and SSM are considered, and tunneling is mentioned at a
couple of points but not expanded upon. The latter part of the
paper discusses addressing, routing, translation, and
application issues in the light of the scenario analyses.

3.   Multicast Proxy in IPv6/IPv4 Transition [ID.jiang-v4v6mc-proxy]

     This document considers the two-network scenarios 4 <- 6 and 6
     <- 4. It does not distinguish ASM or SSM. A multicast proxy
     forwards requests for multicast streams from network A to
     network B, receives the stream content, and forwards it directly
     to the receivers. An implicit mapping step occurs at the
     multicast proxy when a request from a network A receiver for a
     network A multicast stream is translated into a request from the
     multicast proxy for the corresponding network B multicast
     stream. The reverse stream-to-stream mapping is needed to
     ensure that the right streams from network B get forwarded to
     the right receivers. There is no discussion of the mechanism
     for achieving this mapping. There is no mention of tunneling.

4.   Multicast Extensions to DS-Lite Technique in Broadband
     Deployments [ID.qin-dslite-multicast]

     This document considers the three-network 4-6-4 scenario
     specifically as encountered when using the DS Lite transition
     mechanism. Network A is the network at the customer site.
     Network B is the IPv6 access network to which the customer site
     is attached. DS Lite uses IPv4-in-IPv6 tunneling to carry IPv4
     traffic across network B. The objective is to reduce the load on
     the AFTR at the border between network B and network C and make
     more efficient use of network B bandwidth through use of native
     network B multicast capability. The paper effectively restricts
     itself to SSM, although it discusses how to accommodate ASM
     where the sources are all in network C. It uses the embedded
     IPv4 address approach to translation, with specific dependency
     on [RFC6052] for source addresses and
     [ID.boucadair-64-multicast-address-format] for multicast group
     addresses.

5.   Multicast Considerations for Gateway-Initiated Dual-Stack Lite
     [ID.brockners-mcast-gi-ds-lite]

     Like the previous document, this one considers multicast for a

   4-6-4 tunneling scenario. It restricts itself to SSM.
   Translation is not discussed.

6.   Softwire Mesh Multicast [ID.xu-mesh-multicast]

   This document considers a three-network scenario where network B
   is an IP backbone. Both 4-6-4 and 6-4-6 scenarios are
   considered. No distinction is made between ASM and SSM. For
   the 4-6-4 scenario, the usual device of embedded IPv4 addresses
   is used. For the 6-4-6 scenario, additional signalling between
   the edge devices (AFBRs) is used to ensure that the multicast
   flows are matched up correctly. Tunneling is an essential
   feature of the discussion.

7.   IPv4-Embedded IPv6 Multicast Address Format
      [ID.boucadair-64-multicast-address-format]

   This document is written to standardize the format of the IPv4-
   embedded IPv6 multicast address. Effectively it addresses the
   two-network 6 <- 4 scenario or the three-network 4-6-4 scenario,
   although that is not explicitly stated. Both ASM and SSM
   addresses are considered. The proposal applies to both native
   and tunneled transport.

8.   IPv6 Multicast Using Native IPv4 Capabilities in a 6rd
      Deployment [ID.tsou-6rd-multicast]

   This document specifically addresses the three-network 6-4-6
   scenario, where network A is the customer site and network B is
   the access network to which the customer is attached. Neither
   ASM nor SSM is mentioned, although the proposal is applicable to
   both. The document relies on stateful mapping. Although
   described in conjunction with a tunneling mechanism (6rd), the
   proposal uses native transport for multicast content.

9.   Multicast Support for NAT64 [ID.sarikaya-mcast4nat64]

   This document considers the two-network 6 <- 4 scenario. ASM
   and SSM are discussed. IPv4 embedded addresses are used for
   translation, with explicit reference to
   [ID.boucadair-64-multicast-address-format]. The document
   assumes native transport rather than tunneling.

10.   A Generic Approach to Multicast Translation In Support of IPv6
      Transition [ID.tsou-translated-multicast]

   This document considers the generic three-network scenario, with
   intended applicability to both 4-6-4 and 6-4-6. Network A is

the customer site, network B is the access network to which the
customer is connected. (Note that the designations A, B, and C
used in the present document differ from those used in
[ID.tsou-translated-multicast].) ASM and SSM are both in the
intended scope. Translation uses the stateful mapping mechanism
described in [ID.tsou-6rd-multicast]. Tunneling of multicast
content is not considered.

11. A Generic Approach to Multicast tunneling In Support of IPv6
     Transition [ID.tsou-encapsulated-multicast]

     This document has the same scope as the previous document, with
     the one difference that it considers tunneled transport of
     multicast content.

12. IPv4/IPv6 Multicast Translation Framework
     [ID.lee-v4v6-mcast-fwk]

     This document is basically an enumeration of two-network
     scenarios, both 6 <- 4 and 4 <- 6, with comments on when they
     are likely to be encountered. Having become aware of
     [ID.venaas-v4v6mc-framework], the authors intend to retarget
     this work.

13. IPv4-IPv6 Multicast: Problem Statement and Use Cases
     [ID.jaclee-v4v6-mcast-ps]

     This document is a more in-depth examination of various
     scenarios that arise depending on the transition mechanisms
     deployed by the operator. It notes operational issues that can
     complicate the choice of transition techniques in particular
     deployments. One point raised is how to handle situations where
     there is a mixture of IPv6 and IPv4 terminals. A key constraint
     on the discussion is that the transition methods consider must
     accommodate a trend to single virtual connections between the
     customer site and the IP edge, as opposed to multiple service-
     specific connections. This is translated into a requirement to
     avoid use of private IPv4 addressing for the multicast
     sources(?).

14. Automatic IP Multicast Without Explicit Tunnels (AMT)
     [ID.mboned-auto-multicast]

     This document defines an architecture and protocol in the spirit
     of 6to4 ([RFC3056]), whereby isolated multicast sites can
     connect either to a multicast backbone or directly to each
     other. Access to each site is through an AMT Gateway. Access
     to the multicast backbone is through AMT Relays. The unicast

network serves as a link between these entities, through the
mechanism of encapsulation. Scalability issues are addressed
when necessary by supplementing this structure with permanent
tunnels between the the AMT sites and the multicast backbone.
Following discovery and handshakes using the protocol defined in
the AMT document, IGMP/MLD signalling is sent to the AMT
Gateways to establish listener relationships. AMT supports SSM
only from AMT sites, but allows them to receive ASM. From the
point of view of the network model used in this document, the
AMT site corresponds to network A, the intervening unicast
network to network B, and the multicast backbone to network C.

15. Multicast Support for Dual Stack Lite and 6RD
     [ID.sarikaya-dslite-multicast]

     This document considers the DS-Lite (4-6-4) and 6rd (6-4-6)
     scenarios. In both cases it proposes the simplest multicast
     strategy: encapsulation of multicast signalling in the one
     direction and multicast content in the other as unicast traffic.
     The customer equipment (DS-Lite B4, 6rd CE) is required to act
     as a proxy for IGMP or MLD respectively. The border element
     (DS-Lite AFTR, 6rd Border Relay) is required to act as an IGMP
     querier (respectively MLD querier). Otherwise network B is
     unaware of the multicast traffic.

16. Multicast Transition to IPv6 Only in Broadband Deployments
     [ID.tsou-multicast-transition-v6only]

     This document defines a transition path that assumes the
     technology of the CPE, access network, and multicast sources are
     all under operator control. The basic evolution of each of
     these three components is through dual stack to IPv6, but the
     timing is different for the different elements. In consequence
     of the suggested transition path, neither translation nor
     encapsulation is required at any stage.

   Table 1 summarizes the contents of the respective drafts according to
   the four dimensions presented in Section 2. The first column of the
   table identifies each draft by the sequence number given to it in the
   descriptions presented above.

+-------+-----------+--------------------+------------+------------+
| Draft | Scenarios | Translation | ASM/SSM | Tunneling |
| | | Mechanism | | |
+-------+-----------+--------------------+------------+------------+
1	4 <- 6	Specific IPv6	Both	No
		prefixes		
	6 <- 4	Embedded IPv4		

2	4 <- 6	Various ideas	Both	No
	6 <- 4	Embedded IPv4		
3	4 <- 6	ND	ND	No
	6 <- 4			
4	4-6-4	Embedded IPv4	SSM *	Yes
5	4-6-4	ND	SSM	Yes
6	4-6-4	Embedded IPv4	ND	Yes
	6-4-6	Additional signalling		
7	6 <- 4	Embedded IPv4	Both	ND
	4-6-4	ditto		
8	6-4-6	Stateful mapping	ND	No
9	6 <- 4	Embedded IPv4	Both	No
10	4-6-4	Stateful mapping	ND	No
	6-4-6	ditto		
11	4-6-4	Stateful mapping	ND	Yes
	6-4-6	ditto		
12	6 <- 4	ND	ND	ND
	4 <- 6			
13	All	General discussion	Some restriction	General discussion
14	4-6-4	Special prefixes plus supplementary protocol	SSM	Yes
	6-4-6			
15	4-6-4	Border element retains leaf node multicast state	SSM	Yes
	6-4-6	ditto		
16	4-4-4	None	ND	No

ND = Not discussed. * [ID.qin-dslite-multicast] considers the use of

ASM in network C only.

Table 1: Multicast Transition Draft Content Summarized

4. Conclusions

The above analysis leads to a number of suggestions for moving forward.

o For lack of time, a detailed comparison of
   [ID.venaas-v4v6mc-framework] and [ID.lee-v4v6-mcast-fwk] has not
   yet been carried out. This needs to be done to determine what the
   latter document adds to the discussion and the implications for
   the organization of documentation related to multicast transition.

o It is suggested that, because of the completeness of its coverage,
   [ID.venaas-v4v6mc-framework] should be the document on which all
   other multicast transition work should be based. However, details
   relating to embedding IPv4 addresses in IPv6 multicast addresses
   should be removed and left for another document. It is possible
   that other material may be subject to removal to a more
   specialized document, but this requires further thought. The
   document may belong formally to either BEHAVE or SOFTWIRES, but
   must be reviewed and approved by both Working Groups.

o The need for [ID.boucadair-64-multicast-address-format] is quite
   clear, and the document should become a BEHAVE work item.

o It should be possible to create a document considering the details
   of multicast transition in a tunneling environment, applicable to
   both 4-6-4 and 6-4-6. In particular, this document should make
   clear that translation is also applicable when tunneling, but to a
   lesser extent than in the case of native transport. It is
   expected that with this document and [ID.venaas-v4v6mc-framework]
   as starting points, [ID.qin-dslite-multicast] can discard some of
   its content and end up looking more like
   [ID.brockners-mcast-gi-ds-lite].

o [ID.xu-mesh-multicast] relates specifically to multicast
   transition in a backbone network. In this, it differs from the
   other documents reviewed, most of which assume that network B is
   an access network. Some content may duplicate material in other
   documents mentioned above and may thus be subject to removal, but
   there is definitely enough material for this to remain a separate
   document.

o The current priority is the 6 <- 4 or 4-6-4 scenario.
   Nevertheless, it is desirable that work on good solutions for the
   4 <- 6 or 6-4-6 scenario should continue, since this is a harder
   problem and may take some time. BEHAVE would be the natural
   location for this work.

The classification scheme proposed in this document did not capture
all of the contributions of the individual drafts. For instance,
[ID.sarikaya-mcast4nat64] has a number of useful details outside the
scope considered here. It may be possible to identify more problems
of general scope once work on the basics identified here is in hand.
It is also possible that a number of specialized documents like
[ID.xu-mesh-multicast] and [ID.brockners-mcast-gi-ds-lite] should be
written to cover specific scenarios.

5. Security Considerations

This document introduces no new security considerations.

6. IANA Considerations

This document introduces no IANA considerations.

7. References

7.1. Normative References

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

7.2. Informative References

[ID.boucadair-64-multicast-address-format]
           Boucadair, M., Qin, J., Lee, Y., Venaas, S., Li, X., and
           M. Xu, "IPv4-Embedded IPv6 Multicast Address Format (Work
           in progress)", February 2011.

[ID.brockners-mcast-gi-ds-lite]
           Brockners, F. and Y. Lee, "Multicast Considerations for
           Gateway-Initiated Dual-Stack Lite (Work in progress)",
           October 2010.

[ID.jaclee-v4v6-mcast-ps]
           Jacquenet, C., Boucadair, M., Lee, Y., Qin, J., and T.
           Tsou, "IPv4-IPv6 Multicast: Problem Statement and Use

                    Cases (Work in progress)", March 2011.

     [ID.jiang-v4v6mc-proxy]
                    Jiang, S. and D. Gu, "Multicast Proxy in IPv6/IPv4
                    Transition (Work in progress)", January 2011.

     [ID.lee-v4v6-mcast-fwk]
                    Lee, Y. and J. Qin, "IPv4/IPv6 Multicast Translation
                    Framework (Work in progress)", February 2011.

     [ID.mboned-auto-multicast]
                    Thaler, D., Talwar, M., Aggarwal, A., Vicisano, L., and T.
                    Pusateri, "Automatic IP Multicast Without Explicit Tunnels
                    (AMT) (Work in progress)", March 2010.

     [ID.qin-dslite-multicast]
                    Wang, Q., Qin, J., Sun, P., Boucadair, M., Jacquenet, C.,
                    and Y. Lee, "Multicast Extensions to DS-Lite Technique in
                    Broadband Deployments (Work in progress)", January 2011.

     [ID.sarikaya-dslite-multicast]
                    Sarikaya, B., "Multicast Support for Dual Stack Lite and
                    6RD (Work in progress)", March 2011.

     [ID.sarikaya-mcast4nat64]
                    Sarikaya, B., "Multicast Support for NAT64 (Work in
                    progress)", January 2011.

     [ID.tsou-6rd-multicast]
                    Tsou, T., Taylor, T., Zhou, C., and H. Ji, "IPv6 Multicast
                    Using Native IPv4 Capabilities in a 6rd Deployment (Work
                    in progress)", January 2011.

     [ID.tsou-encapsulated-multicast]
                    Tsou, T., Zhou, C., and H. Ji, "A Generic Approach to
                    Multicast tunneling In Support of IPv6 Transition (Work
                    in progress)", January 2011.

     [ID.tsou-multicast-transition-v6only]
                    Tsou, T. and T. Taylor, "Multicast Transition to IPv6 Only
                    in Broadband Deployments (Work in progress)", March 2010.

     [ID.tsou-translated-multicast]
                    Tsou, T., Taylor, T., Zhou, C., and H. Ji, "A Generic
                    Approach to Multicast Translation In Support of IPv6
                    Transition (Work in progress)", January 2011.

     [ID.venaas-mcast46]

               Venaas, S., Asaeda, H., SUZUKI, S., and T. Fujisaki, "An
               IPv4 - IPv6 multicast translator (Work in progress)",
               December 2010.

   [ID.venaas-v4v6mc-framework]
               Venaas, S., Li, X., and C. Bao, "Framework for IPv4/IPv6
               Multicast Translation (Work in progress)", December 2010.

   [ID.xu-mesh-multicast]
               Xu, M., Cui, Y., Yang, S., Metz, C., and G. Shepherd,
               "Softwire Mesh Multicast (Work in progress)",
               October 2010.

   [RFC3056]   Carpenter, B. and K. Moore, "Connection of IPv6 Domains
               via IPv4 Clouds", RFC 3056, February 2001.

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

Authors' Addresses

   Tina Tsou
   Huawei Technologies (USA)
   2330 Central Expressway
   Santa Clara, CA 95050
   USA

   Phone: +1 408 330 4424
   Email: tena@huawei.com
   URI: http://tinatsou.weebly.com/contact.html

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

   Phone:
   Email: cathyzhou@huawei.com

Tom Taylor
Huawei Technologies
1852 Lorraine Ave.t
Ottawa, Ontario   K1H 6Z8
Canada

Phone:
Email: tom111.taylor@bell.net