MAP Encapsulation (MAP-E) - specification
draft-mdt-softwire-map-encapsulation-00

Abstract

   This document specifies the "Mapping of Address and Port" (MAP)
   encapsulation based solution (MAP-E) with an automatic tunneling
   mechanism for providing IPv4 connectivity service to end users over a
   service provider's IPv6 network.

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 30, 2012.

Copyright Notice

Table of Contents

[1](#). **Introduction**

   MAP-E is a protocol mechanism of the &ldque;Mapping of Address and
   Port" (MAP) encapsulation based solution to deploy IPv4 to sites via
   a service provider's (SP's) IPv6 network with the automatic tunneling
   mechanism (IPv4-in-IPv6). Similar to Dual-Stack Lite
   [I-D.ietf-softwire-dual-stack-lite], MAP-E is designed to allow IPv4
   traffic to be delivered over an IPv6 network without the direct
   provisioning of IPv4 addresses. Like 6rd [RFC5969], MAP-E is
   operated in a fully stateless manner within the SP network.

   MAP-E relies on IPv6 and is designed to deliver production-quality
   dual-stack service while allowing IPv4 to be phased out within the SP
   network. The phasing out of IPv4 within the SP network is
   independent of whether the end user disables IPv4 service or not.
   Further, &ldque;Greenfield&ldque; IPv6-only networks may use MAP-E in
   order to deliver IPv4 to sites via the IPv6 network in a way that
   does not require protocol translation between IPv4 and IPv6.

   MAP-E utilizes an algorithmic mapping, defined in MAP
   [I-D.mdt-softwire-mapping-address-and-port], between the IPv6 and
   IPv4 addresses that are assigned for use within the SP network. This
   mapping can provide automatic determination of IPv6 tunnel endpoints
   from IPv4 destination addresses, allowing the stateless operation of
   MAP-E. MAP-E views the IPv6 network as a link layer for IPv4 and
   supports an automatic tunneling abstraction similar to the Non-
   Broadcast Multiple Access (NBMA) [RFC2491] model.

   The MAP algorithmic mapping is also used to automatically provision
   IPv4 addresses and allocating a set of non-overlapping ports for each
   MAP-E CE. The "SP-facing" (i.e., "WAN") side of the MAP-E CE,
   operate as native IPv6 interface with no need for IPv4 operation or
   support. On the "end-user-facing" (i.e., "LAN") side of a CE, IPv6
   and IPv4 might be implemented as for any native dual-stack service
   delivered by the SP.

   A MAP-E domain consists of MAP-E Customer Edge (CE) routers and one
   or more MAP-E Border Relays (BRs). IPv4 packets encapsulated by
   MAP-E follow the IPv6 routing topology within the SP network between
   CEs and among CEs and BRs. CE to CE traffic is direct, while BRs are
   traversed only for IPv4 packets that are destined to or are arriving
   from outside a given MAP-E domain. As MAP-E is stateless, BRs may be
   reached using anycast for failover and resiliency.

   MAP-E does not require any stateful NAPT [RFC3022] functions at the
   BRs or elsewhere within the SP network. Instead, MAP-E allows for
   sharing of IPv4 addresses among multiple sites by automatically
   allocating a set of non-overlapping ports for each CE as part of the

stateless mapping function. It is expected that the CE will, in
turn, perform local IPv4 Network Address and Port Translation (NAPT)
[RFC3022] functions for the site as is commonly performed today,
except avoiding ports outside of the allocated port set. Although
MAP-E is designed primarily to support IPv4 deployment to a customer
site (such as a residential home network) by an SP, it can equally be
applied to an individual host acting as a CE router.

2. 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].

3. Terminology

MAP-E: Mapping of Address and Port - Encapsulation
                       mode. MAP-E utilizes a simple IPv4-in-IPv6
                       tunneling along with the MAP extensions for
                       mapping between IPv4 and IPv6 defined in MAP
                       [I-D.mdt-softwire-mapping-address-and-port] and
                       this draft.

MAP-E domain (Domain): A set of MAP-E CEs and BRs connected to the
                       same virtual MAP-E link. A service provider
                       may deploy MAP-E with a single MAP-E domain, or
                       may utilize multiple MAP-E domains. Each
                       domain requires a separate MAP-E rule set.

MAP-E Border Relay (BR): A MAP-E enabled router managed by the
                       service provider at the edge of a MAP-E domain.
                       A Border Relay router has at least one of each
                       of the following: an IPv6-enabled interface, a
                       MAP-E virtual interface acting as an endpoint
                       for the MAP-E IPv4 in IPv6 tunnel, and an IPv4
                       interface connected to the native IPv4 network.
                       A MAP-E BR may also be referred to simply as a
                       "BR" within the context of MAP-E.

MAP-E Customer Edge (CE): A device functioning as a Customer Edge
                       router in a MAP-E deployment. In a residential
                       broadband deployment, this type of device is
                       sometimes referred to as a "Residential
                       Gateway" (RG) or "Customer Premises Equipment"
                       (CPE). A typical MAP-E CE serving a
                       residential site has one WAN side interface,

one or more LAN side interfaces, and a MAP-E
virtual interface. A MAP-E CE may also be
referred to simply as a "CE" within the context
of MAP-E.

Shared IPv4 address: An IPv4 address that is shared among multiple
nodes. Each node has a separate part of the
transport layer port space.

MAP-E Rule: A MAP rule defining the mapping relationship
for a given MAP-E domain between IPv4 and IPv6,
defined in MAP
[I-D.mdt-softwire-mapping-address-and-port]

4. MAP-E Configuration

The IPv4 prefix, IPv4 address or shared IPv4 address for use at a
customer site is automatically obtained based on BMR defined in MAP
[I-D.mdt-softwire-mapping-address-and-port] from the IPv6 prefix
delegated to the site.

For a given MAP-E domain, the BR and CE MUST be configured with a set
of mapping rules (BMR, FMR and DMR) defined in
[I-D.mdt-softwire-mapping-address-and-port] . The configured values
for these elements MUST be consistent for all CEs and BRs within a
given MAP-E domain.

The configuration elements in the set of mapping rules (BMR, FMR and
DMR) may be provisioned via IPv6 DHCP defined in
[I-D.mdt-softwire-map-dhcp-option] or manually.

The only remaining provisioning information in order to enable MAP-E
is an IPv6 prefix. This IPv6 prefix is configured as part of
obtaining IPv6 Internet access (i.e., configured via SLAAC, DHCPv6,
DHCPv6 PD, manual or otherwise).

5. MAP-E Node Behavior

5.1. Provisioning of MAP-E BR

The MAP-E BR needs to be provisioned with information for the MAP-E
domain or domains it is expected to handle, along with any necessary
routing processes. For each MAP-E domain, the BR will have the
following parameters:

o The MAP Domain IPv4 and IPv6 prefix, and their lengths (Basic

Mapping Rule)

o The MAP EA-bits (CE index), including IPv4 suffix, length and any port-range (including any excluded ports and the port number continuity parameter)

o The BR prefix and its length (Default Mapping Rule)

o The subnet ID

A BR when configured for BMR, FMR and DMR, and performs the following functions:

o Configures the IPv4/IPv6 stateless encapsulation parameters (BMR, FMR and DMR)

Based on the above configuration, the IPv4-in-IPv6 encapsulation function can be performed by the BR.

o Derive IPv4 address along with any applicable port-range from IPv4-translatable address (BMR)

o Derive IPv4-translatable address from IPv4 address and port number (FMR)

5.2. Packet Forwarding Behavior on MAP-E BR

(a) BR reception of an IPv4 packet

Step 1   BR looks up an appropriate mapping rule (FMR)
                       with a specific Domain IPv4 prefix which has
                       the longest match with an IPv4 destination
                       address in the received IPv4 packet. If the
                       FMR is not found, the received packet should be
                       discarded. If the length of Domain IPv4 prefix
                       plus EA-bits associated with the FMR does not
                       exceed 32 bits, BR proceeds to step 2. If the
                       length exceeds 32 bits, BR checks that the
                       received packet contains a complete IPv4
                       datagram. If the packet is fragmented, BR
                       should reassemble the packet. Once BR can
                       obtain the complete IPv4 datagram, BR proceeds
                       to step 2 as though the datagram has been
                       received in a single packet.

Step 2   BR generates a CE IPv6 address from the IPv4
                          destination address or the IPv4 destination
                          address and the destination port based on the
                          FMR found in step 1. If the CE IPv6 address
                          can be successfully generated, BR encapsulates
                          the IPv4 packet in IPv6 and forwards the IPv6
                          packet via the IPv6 interface. If the length
                          of the IPv6 encapsulated packet exceeds the MTU
                          of the IPv6 interface, the fragmentation should
                          be done in IPv6.

   (b) BR reception of an IPv6 packet

   Step 1   If the received IPv6 packet is fragmented, the
                          reassembly should be done in IPv6 at first.
                          Once BR obtains a complete IPv6 packet, BR
                          looks up an appropriate mapping rule (BMR) with
                          a specific Domain IPv6 prefix which has the
                          longest match with an IPv6 source address in
                          the received IPv6 packet. If the BMR rule is
                          not found, the received IPv6 packet should be
                          discarded. BR derives a CE IPv6 address from
                          the IPv4 source address or the IPv4 source
                          address and the source port in the encapsulated
                          IPv4 packet based on the BMR. If the CE IPv6
                          address is eqaul to the IPv6 source address in
                          the received IPv6 packet, BR decapsulates the
                          IPv4 packet and then forward it via the IPv4
                          interface.

5.3. Provisioning of MAP-E CE

   A MAP-E CE requires the following parameters for provisioning:

   o The MAP Domain IPv4 and IPv6 prefix, and their lengths (Basic
   Mapping Rule)

   o The MAP EA-bits (CE index), including IPv4 suffix, length and any
   port-range (including any excluded ports and the port number
   continuity parameter)

   o The BR prefix and its length (Default Mapping Rule)

   A MAP-E CE that receives a MAP DHCP option
   [I-D.mdt-softwire-map-dhcp-option] for BMR, FMR and DMR and performs
   the following (MAP initialization) functions:

   o Configures the NAT44 port-range mapping function parameters (BMR)

   o Configures the IPv4/IPv6 stateless encapsulation parameters (BMR,
   FMR and DMR) Based on the above configuration, the IPv4/IPv6
   encapsulation function can be performed in CE.

   o Derives IPv4 address along with any applicable port-range from
   IPv4-translatable address (BMR)

   o Derives IPv4-translatable address from IPv4 address (FMR)

5.4. Packet Forwarding Behavior on MAP-E CE

   (a) CE reception of an IPv4 packet

   Step 1   CE looks up an appropriate mapping rule (FMR)
                         with a specific Doamin IPv4 prefix which has
                         the longest match with an IPv4 destination
                         address in the received IPv4 packet. If the
                         FMR is found, the length of Domain IPv4 prefix
                         plus EA-bits must be checked. If the length
                         does not exceeds 32 bits, CE proceeds to step
                         2. If the length exceeds 32 bits, CE checks
                         that the received IPv4 packet contains a
                         complete IPv4 datagram. If the packet is
                         fragmented, CE should reassemble the packet.
                         Once CE can obtain the complete IPv4 datagram,
                         CE proceeds to step 2 as though the datagram
                         has been received in a single packet. If the
                         FMR is not found, CE proceeds to step 2.

    Step 2   If the FMR is found in step 1, CE derives a
                         IPv6 destination address from the IPv4
                         destination address or the IPv4 destination
                         address and the destination port based on the
                         FMR. If the IPv6 destination address can be
                         derived successfully, CE encapsulates the IPv4
                         packet in IPv6 whose destination address is set
                         to the derived IPv6 address. If the FMR is not
                         found in step 1, CE uses the DMR and then CE
                         encapsulates the IPv4 packet in IPv6 whose
                         destination address is set to the BR IPv6
                         address. Then CE forwards the IPv6 packet via
                         IPv6 interface. If the length of the IPv6
                         packet exceeds the MTU of the IPv6 interface,
                         the fragmentation should be done in IPv6.
                         Moreover, if using IPv4 shared address, a
                         Datagram ID in the received IPv4 header must be
                         over-written before encapsulating the IPv4
                         packet in IPv6. In case of shared IPv4

                         address, the Datagram ID must be unique among
                         CEs sharing the same IPv4 address. Hence, CE
                         should assign the unique value and set this
                         value to the datagram ID in IPv4 header. This
                         value may be generated from the port-range
                         assigned to the CE to keep the uniqueness among
                         CEs sharing same IPv4 address.

   (b) CE reception of an IPv6 packet

   Step 1   If the received IPv6 packet is fragmented, the
                         reassembly should be done in IPv6 at first.
                         Once CE obtains a complete IPv6 packet, CE
                         looks up an appropriate mapping rule (BMR) with
                         a specific Domain IPv6 prefix which has the
                         longest match with an IPv6 source address in
                         the recieved IPv6 packet. If the BMR is found,
                         the CE derives a CE IPv6 address from the IPv4
                         source address or the IPv4 source address and
                         the source port based on the BMR and then
                         checks that the IPv6 source address of the
                         received IPv6 packet is matched to it. If the
                         BMR is not found, CE checks that the IPv6
                         source address is matched to the BR IPv6
                         address. In case of success, the CE can
                         decapsulate the IPv4 packet and forward it via
                         the IPv4 interface.

6. Deriving IPv6 address from IPv4

6.1. Deriving IPv6 address from IPv4 Address and Port Number at the BR

   IPv6 Source Address and Source Port Number:

   At the BR, the IPv6 source address MUST be set to the BR IPv6 address
   as per DMR MAP [I-D.mdt-softwire-mapping-address-and-port]. The
   source Layer 4 port number MUST be unchanged.

   IPv6 Destination Address and Destination Port Number:

   At the BR, the IPv6 destination address (IPv4-translatable address)
   MUST be derived from the IPv4 destination address and the destination
   port number per FMR MAP [I-D.mdt-softwire-mapping-address-and-port].
   The destination Layer 4 port number MUST be unchanged.

6.2. Deriving IPv6 address from IPv4 Address and Port Number at the CE

IPv6 Source Address and Source Port Number:

At the CE, the IPv6 source address (IPv4-translatable address) MUST
be derived from the IPv4 source address as per BMR MAP
[I-D.mdt-softwire-mapping-address-and-port]. The source port number
MUST be unchanged.

IPv6 Destination Address and Destination Port Number:

At the CE, if Forwarding Mapping Rules (FMRs) are enabled, the IPv4
packet MUST be checked to see if the IPv4 destination address matches
the FMR. If matching, the IPv6 destination address (IPv4-converted
address) MUST be derived from the IPv4 destination address and the
destination port number as per FMR. Otherwise, the IPv6 destination
address MUST be set to the BR IPv6 address per DMR MAP
[I-D.mdt-softwire-mapping-address-and-port]. The destination port
number MUST be unchanged.

7. Encapsulation and Fragmentation Considerations

Maximum transmission unit (MTU) and fragmentation issues for IPv4 in
IPv6 tunneling are discussed in detail in Section 7.2 of [RFC2473].
MAP-E's scope is limited to a service provider network. IPv6 Path
MTU discovery MAY be used to adjust the MTU of the tunnel as
described in Section 7.2 of [RFC2473], or the MAP-E Tunnel MTU might
be explicitly configured.

The use of an anycast source address could lead to any ICMP error
message generated on the path being sent to a different BR.
Therefore, using dynamic tunnel MTU Section 7.2 of [RFC2473] is
subject to IPv6 Path MTU blackholes.

Multiple BRs using the same anycast source address could send
fragmented packets to the same MAP-E CE at the same time. If the
fragmented packets from different BRs happen to use the same fragment
ID, incorrect reassembly might occur. For this reason, a BR using an
anycast source address MUST NOT fragment the IPv6 encapsulated packet
unless BR's having identical rules are required to use disjoint
ranges of fragment ID.

If the MTU is well-managed such that the IPv6 MTU on the CE WAN side
interface is set so that no fragmentation occurs within the boundary
of the SP, then the MAP-E Tunnel MTU should be set to the known IPv6
MTU minus the size of the encapsulating IPv6 header (40 bytes). For
example, if the IPv6 MTU is known to be 1500 bytes, the MAP-E Tunnel

MTU might be set to 1460 bytes. Absent more specific information,
the MAP-E Tunnel MTU SHOULD default to 1280 bytes.

Alternatively, if BR's having identical rule are required to use
disjoint ranges of fragment ID, a BR using an anycast source address
SHOULD fragment the IPv6 encapsulated packet correctly.

For MAP-E domain traversal, IPv4 packets are encapsulated in IPv6
packets whose Next header is set to 4 (i.e. IPv4). If fragmentation
of IPv6 packets is needed, it is performed according to [RFC2460].
Absent more specific information, the path MTU of a MAP-E Domain has
to be set to 1280 [RFC2460].

In domains where IPv4 addresses are not shared, IPv6 destinations are
derived from IPv4 addresses alone. Thus, each IPv4 packet can be
encapsulated and decapsulated independently of each other. MAP-E
processing is completely stateless.

On the other hand, in domains where IPv4 addresses are shared, BR's
and CE's can have to encapsulate IPv4 packets whose IPv6 destinations
depend on destination ports. Precautions are needed, due to the fact
that the destination port of a fragmented datagram is available only
in its first fragment. A sufficient precaution consists in
reassembling each datagram received in multiple packets, and to treat
it as though it would have been received in single packet. This
function is such that MAP-E is in this case stateful at the IP layer.
(This is common with DS-lite and NAT64/DNS64 which, in addition, are
stateful at the transport layer.) At Domain entrance, this ensures
that all pieces of all received IPv4 datagrams go to the right IPv6
destinations.

Another peculiarity of shared IPv4 addresses is that, without
precaution, a destination could simultaneously receive from different
sources fragmented datagrams that have the same Datagram ID (the
Identification field of [RFC0791]. This would disturb the reassembly
process. To eliminate this risk, CE MUST rewrite the datagram ID to
an unique value among CEs having same shared IPv4 address upon
sending the packets over MAP-E tunnel. This value SHOULD be
generated locally within the port-range assigned to a given CE. Note
that replacing a Datagram ID in an IPv4 header implies an update of
its Header-checksum fieald, by adding to it the one's complement
difference between the old and the new values.

8. Packet Forwarding Considerations

8.1. Mesh model

   Basically, MAP-E should allow the mesh model in order for all CEs to
   communicate each others directly. If one mapping rules is applied to
   a given MAP-E domain, all CEs can communicate each others directly.
   If multiple mapping rules are applied to a given MAP-E domain, or if
   multiple MAP-E domains are existed, CE can communicate each others
   directly only if all CEs know all mapping rules. When a CE receives
   an IPv4 packet from its LAN side, the CE looks up a mapping rule
   corresponding to an IPv4 destination address in the received IPv4
   packet. If the corresponding mapping rule is found, CE can
   communicate to another CE directly based on the mapping rule defined
   as Forwarding mapping rule (FMR) in
   [I-D.mdt-softwire-mapping-address-and-port] . If the corresponding
   mapping rule is not found, CE must forward the packet to a given BR.

8.2. Hub & Spoke model

   In order to allow the mesh topology so that all CEs can communicate
   each others directly, all CE should know all mapping rules applied to
   a given MAP-E domain or MAP-E domains. However, if a CE knows only
   subset of mapping rules applied to a given MAP-E domain or MAP-E
   domains, a CE can not communicate to some CEs due to the lack of
   mapping rules. In this case, an IPv4 packet toward to these CEs must
   be forwarded to a given BR. In order to achieve the hub & spoke mode
   fully, Forwarding mapping rule (FMR) defined in
   [I-D.mdt-softwire-mapping-address-and-port] should be disabled. In
   this case, all CEs do not look up the mapping rules upon receiving an
   IPv4 packet from its LAN side and then CE must encapsulate the IPv4
   packet with IPv6 whose destination must be a given BR.

9. NAT Considerations

   NAT44 should be implemented in CPE which has MAP-E CE function. The
   NAT44 must conform that best current practice documented in
   [RFC4787], [RFC5508] and [RFC5382]. When there are restricted
   available port numbers in a given MAP-E CE, the NAT44 must restrict
   mapping ports within the port-set.

10. ICMP Considerations

   ICMP message should be supported in MAP-E domain. Hence, the NAT44
   in MAP-E CE must implement the behavior for ICMP message conforming
   to the best current practice documented in [RFC5508].

   If a MAP-E CE receives an ICMP message having ICMP identifier field

   in ICMP header, NAT44 in the MAP-E CE must rewrite this field to a
   specific value assigned from the port-set. BR and other CEs must
   handle this field similar to the port number in tcp/udp header upon
   receiving the ICMP message with ICMP identifier field.

   If a MAP-E BR and CE receives an ICMP error message without ICMP
   identifier field for some errors that is detected inside a IPv6
   tunnel, a MAP-E BR and CE should replay the ICMP error message to the
   original source. This behavior should be implemented conforming to
   the section 8 of [RFC2473]. The MAP-E BR and CE obtain the origianl
   IPv6 tunnel packet storing in ICMP payload and then decapsulate IPv4
   packet. Finally the MAP-E BR and CE generate a new ICMP error
   message from the decapsulated IPv4 packet and then forward it.

   If a MAP-E BR receives an ICMP error message on its IPv4 interface,
   the MAP-E BR should replay the ICMP message to an appropriate MAP-E
   CE. If IPv4 address is not shared, the MAP-E BR generates a CE IPv6
   address from the IPv4 destination address in the ICMP error message
   and encapsulates the ICMP message in IPv6. If IPv4 address is
   shared, the MAP-E BR derives an original IPv4 packet from the ICMP
   payload and generates a CE IPv6 address from the source address and
   the source port in the original IPv4 packet. If the MAP-E BR can
   generate the CE IPv6 address, the MAP-E BR encapsulates the ICMP
   error message in IPv6 and then forward it to its IPv6 interface.

11. Security Considerations

   Spoofing attacks: With consistency checks between IPv4 and IPv6
                           sources that are performed on IPv4/IPv6 packets
                           received by BR's and CE's (Section 5), MAP-E
                           does not introduce any opportunity for spoofing
                           attack that would not pre-exist in IPv6.

   Denial-of-service attacks: In MAP-E domains where IPv4 addresses are
                           shared, the fact that IPv4 datagram reassembly
                           may be necessary introduces an opportunity for
                           DOS attacks. This is inherent to address
                           sharing, and is common with other address
                           sharing approaches such as DS-lite and NAT64/
                           DNS64. The best protection against such
                           attacks is to accelerate IPv6 enablement in
                           both clients and servers so that, where MAP-E
                           is supported, it is less and less used.

Routing-loop attacks: This attack may exist in some automatic-
                       tunneling scenarios are documented in
                       [I-D.ietf-v6ops-tunnel-loops].   They cannot
                       exist with MAP-E because each BRs checks that
                       the IPv6 source address of a received IPv6
                       packet is a CE address.

Attacks facilitated by restricted port set: From hosts that are not
                       subject to ingress filtering of [RFC2827], some
                       attacks are possible by intervening with faked
                       packets during ongoing transport connections
                       ([RFC4953], [RFC5961], [RFC6056].   The attacks
                       depend on guessing which ports are currently
                       used by target hosts. Using unrestricted port
                       set which mean that are IPv6 is exactly
                       preferable. To avoid this attacks using
                       restricted port set, NAT44 filtering behavior
                       SHOULD be "Address-Dependent Filtering".

12. IANA Consideration

This document makes no request of IANA.

13. Acknowledgements

This draft is based on original idea described in
[I-D.despres-softwire-sam]. The authors would like to thank Remi
Despres, Mark Townsley, Wojciech Dec and Olivier Vautrin.

14. References

14.1. Normative References

[I-D.mdt-softwire-map-dhcp-option]
          Mrugalski, T., Boucadair, M., Troan, O., Deng, X., and C.
          Bao, "DHCPv6 Options for Mapping of Address and Port",
          draft-mdt-softwire-map-dhcp-option-01 (work in progress),
          October 2011.

[I-D.mdt-softwire-mapping-address-and-port]
          Troan, O., "Mapping of Address and Port (MAP)",
          draft-mdt-softwire-mapping-address-and-port-01 (work in
          progress), October 2011.

[RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791,

                September 1981.

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

    [RFC2460]   Deering, S. and R. Hinden, "Internet Protocol, Version 6
                (IPv6) Specification", RFC 2460, December 1998.

    [RFC2491]   Armitage, G., Schulter, P., Jork, M., and G. Harter, "IPv6
                over Non-Broadcast Multiple Access (NBMA) networks",
                RFC 2491, January 1999.

    [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.

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

14.2. Informative References

    [I-D.despres-softwire-sam]
                Despres, R., "Stateless Address Mapping (SAM) - a
                Simplified Mesh-Softwire Model",
                draft-despres-softwire-sam-01 (work in progress),
                July 2010.

    [I-D.ietf-softwire-dual-stack-lite]
                Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-
                Stack Lite Broadband Deployments Following IPv4
                Exhaustion", draft-ietf-softwire-dual-stack-lite-11 (work
                in progress), May 2011.

    [I-D.ietf-v6ops-tunnel-loops]
                Nakibly, G. and F. Templin, "Routing Loop Attack using
                IPv6 Automatic Tunnels: Problem Statement and Proposed
                Mitigations", draft-ietf-v6ops-tunnel-loops-07 (work in
                progress), May 2011.

    [I-D.operators-softwire-stateless-4v6-motivation]
                Boucadair, M., Matsushima, S., Lee, Y., Bonness, O.,
                Borges, I., and G. Chen, "Motivations for Stateless IPv4
                over IPv6 Migration Solutions",
                draft-operators-softwire-stateless-4v6-motivation-02 (work
                in progress), June 2011.

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

   [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering:
              Defeating Denial of Service Attacks which employ IP Source
              Address Spoofing", BCP 38, RFC 2827, May 2000.

   [RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network
              Address Translator (Traditional NAT)", RFC 3022,
              January 2001.

   [RFC4787] Audet, F. and C. Jennings, "Network Address Translation
              (NAT) Behavioral Requirements for Unicast UDP", BCP 127,
              RFC 4787, January 2007.

   [RFC4953] Touch, J., "Defending TCP Against Spoofing Attacks",
              RFC 4953, July 2007.

   [RFC5382] Guha, S., Biswas, K., Ford, B., Sivakumar, S., and P.
              Srisuresh, "NAT Behavioral Requirements for TCP", BCP 142,
              RFC 5382, October 2008.

   [RFC5508] Srisuresh, P., Ford, B., Sivakumar, S., and S. Guha, "NAT
              Behavioral Requirements for ICMP", BCP 148, RFC 5508,
              April 2009.

   [RFC5961] Ramaiah, A., Stewart, R., and M. Dalal, "Improving TCP's
              Robustness to Blind In-Window Attacks", RFC 5961,
              August 2010.

   [RFC5969] Townsley, W. and O. Troan, "IPv6 Rapid Deployment on IPv4
              Infrastructures (6rd) -- Protocol Specification",
              RFC 5969, August 2010.

   [RFC6056] Larsen, M. and F. Gont, "Recommendations for Transport-
              Protocol Port Randomization", BCP 156, RFC 6056,
              January 2011.

Authors' Addresses

   Tetsuya Murakami (editor)
   IP Infusion
   1188 East Arques Avenue
   Sunnyvale
   USA

   Email: tetsuya@ipinfusion.com

   Ole Troan
   cisco
   Oslo
   Norway

   Email: ot@cisco.com

   Satoru Matsushima
   SoftBank
   1-9-1 Higashi-Shinbashi, Munato-ku
   Tokyo
   Japan

   Email: satoru.matsushima@tm.softbank.co.jp