

Independent Submission
Request for Comments: 5579
Category: Informational
ISSN: 2070-1721

F. Templin, Ed.
Boeing Research & Technology
February 2010

Transmission of IPv4 Packets over Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) Interfaces

Abstract

The Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) specifies a Non-Broadcast, Multiple Access (NBMA) interface type for the transmission of IPv6 packets over IPv4 networks using automatic IPv6-in-IPv4 encapsulation. The original specifications make no provisions for the encapsulation and transmission of IPv4 packets, however. This document specifies a method for transmitting IPv4 packets over ISATAP interfaces.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This is a contribution to the RFC Series, independently of any other RFC stream. The RFC Editor has chosen to publish this document at its discretion and makes no statement about its value for implementation or deployment. Documents approved for publication by the RFC Editor are not a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc5579>.

IESG Note

This RFC is not a candidate for any level of Internet Standard. The IETF disclaims any knowledge of the fitness of this RFC for any purpose and in particular notes that the decision to publish is not based on IETF review for such things as security, congestion control, or inappropriate interaction with deployed protocols. The RFC Editor has chosen to publish this document at its discretion. Readers of this document should exercise caution in evaluating its value for implementation and deployment. See RFC 3932 for more information.

Copyright Notice

Copyright (c) 2010 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 carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

1. Introduction	3
2. Terminology	3
3. ISATAP Interface Model	3
4. ISATAP Interface MTU	4
5. IPv6 Operation	4
6. IPv4 Operation	4
6.1. ISATAP IPv4 Address Configuration	4
6.2. IPv4 Route Configuration	5
6.3. ISATAP Interface Determination	5
6.4. Next-Hop Resolution	5
6.5. Encapsulation and Transmission	6
6.6. IPv4 Multicast Mapping	6
6.7. Recursive Encapsulation Avoidance	7
7. Security Considerations	7
8. Acknowledgements	7
9. References	7
9.1. Normative References	7
9.2. Informative References	8
Appendix A. Encapsulation Avoidance	9

1. Introduction

The Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) [RFC5214] specifies a Non-Broadcast, Multiple Access (NBMA) interface type for the transmission of IPv6 packets over IPv4 networks using automatic IPv6-in-IPv4 encapsulation. ISATAP interfaces therefore typically configure IPv6 addresses and prefixes, but they do not configure IPv4 addresses and prefixes. In typical implementations and deployments, an ISATAP interface therefore appears as an ordinary interface configured for IPv6 operation but with a null IPv4 configuration. This places an unnecessary limitation on the ISATAP domain of applicability.

ISATAP interfaces perform automatic IPv6-in-IPv4 encapsulation over a virtual IPv6 overlay that spans a region within a connected IPv4 routing topology (i.e., a "site") comprising ordinary IPv4 routers. ISATAP interfaces configure IPv6 link-local addresses that encapsulate an IPv4 address assigned to an underlying IPv4 interface within the IPv6 link-local prefix "fe80::/10", as specified in Sections 6 and 7 of [RFC5214]. ISATAP interfaces may additionally configure IPv6 addresses from a non-link-local IPv6 prefix in exactly the same fashion. As a result, [RFC5214] extends the basic transition mechanisms specified in [RFC4213].

This document specifies mechanisms and operational practices that enable automatic IPv4-in-IPv4 encapsulation over ISATAP interfaces in the same manner as for IPv6-in-IPv4 encapsulation. As a result, this document also extends the IPv4-in-IPv4 tunneling mechanisms specified in [RFC2003]. These mechanisms are useful in a wide variety of enterprise network scenarios, e.g., as discussed in the RANGER [RANGER] and VET [VET] proposals.

The following sections specify IPv4 operation over ISATAP interfaces. A working knowledge of the IPv4 and IPv6 protocols ([RFC0791] and [RFC2460]), IPv4-in-IPv4 encapsulation [RFC2003], and IPv6-in-IPv4 encapsulation ([RFC4213] and [RFC5214]) is assumed.

2. Terminology

The keywords MUST, MUST NOT, REQUIRED, SHALL, SHALL NOT, SHOULD, SHOULD NOT, RECOMMENDED, MAY, and OPTIONAL, when they appear in this document, are to be interpreted as described in [RFC2119].

3. ISATAP Interface Model

ISATAP interfaces use automatic IPv6-in-IPv4 encapsulation to span a region within a connected IPv4 routing topology (i.e., a "site") in a single IPv6 hop. That is to say that the site comprises border nodes

with ISATAP interfaces that forward IPv6-in-IPv4 packets across the site in a single IPv6 hop, and ordinary IPv4 routers between the border nodes that decrement the Time to Live (TTL) in the outer IPv4 header. Border nodes that configure ISATAP interfaces within the same site therefore see each other as single-hop neighbors.

ISATAP interfaces are configured over one or more of the node's underlying IPv4 interfaces connected to the site. These underlying IPv4 interfaces configure site- or greater-scoped IPv4 addresses, and the underlying IPv4 interfaces of two "neighboring" ISATAP interfaces may be separated by many IPv4 hops within the site.

This specification simply extends the ISATAP interface model to also support IPv4-in-IPv4 encapsulation. When IPv4-in-IPv4 encapsulation is used, the ISATAP interface spans exactly the same underlying site as for IPv6-in-IPv4 encapsulation.

4. ISATAP Interface MTU

ISATAP interface MTU considerations are exactly as specified in Section 3.2 of [RFC4213] and apply equally for both IPv6 and IPv4 operation.

5. IPv6 Operation

IPv6 operations over ISATAP interfaces are exactly as specified in [RFC5214].

6. IPv4 Operation

The following sections specify IPv4 operation over ISATAP interfaces:

6.1. ISATAP IPv4 Address Configuration

As for IPv6 operation, IPv4 operation requires that all ISATAP interfaces connected to the same site configure a unique Layer 3 IPv4 address ("L3ADDR") taken from an IPv4 prefix for the site. L3ADDR is used for next-hop determination, but it may also be used as the source address for packets that originate from the ISATAP interface itself.

When a unique "name" for the ISATAP site is required (e.g., to distinguish it from other ISATAP sites), L3ADDR is taken from a public IPv4 prefix. Otherwise, it may be taken from a link-local-scoped IPv4 prefix (e.g., 169.254/16 [RFC3927]).

Methods for ensuring L3ADDR uniqueness include dynamic allocation using DHCP [RFC2131], manual configuration, etc.

6.2. IPv4 Route Configuration

As for any IPv4 interface, IPv4 Forwarding Information Base (FIB) entries (i.e., IPv4 routes) are configured over ISATAP interfaces via either administrative or dynamic mechanisms.

Next-hop addresses in FIB entries configured over an ISATAP interface correspond to the L3ADDR assigned to the ISATAP interface of a neighbor.

6.3. ISATAP Interface Determination

When the node's IPv4 layer has a packet to send, it performs an IPv4 FIB lookup to determine the outgoing ISATAP interface and the next-hop L3ADDR. The node then checks the packet length against the MTU configured on the ISATAP interface.

If the packet is no larger than the MTU, the node admits it into the ISATAP interface without fragmentation. If the packet is larger than the MTU, the node examines the "Don't Fragment (DF)" flag in the IPv4 header. If DF=1, it drops the packet and returns an ICMPv4 "fragmentation needed" message to the original source [RFC1191]; otherwise, it fragments the packet using IPv4 fragmentation and admits each fragment into the ISATAP interface.

6.4. Next-Hop Determination and Address Mapping

As for ISATAP for IPv6, ISATAP for IPv4 requires a means for determining the L3ADDR of neighbors on the ISATAP interface that can provide a next-hop toward the final destination. Next-hop determination for default routes is through the Potential Router List (PRL) discovery procedures specified in Section 8.3.2 of [RFC5214]. Next-hop determination methods for more-specific routes include forwarding initial packets via a default router until a redirect is received, name service lookup (e.g., as described in [VET]), etc.

After a next-hop L3ADDR is discovered, the node admits IPv4 packets/fragments into the ISATAP interface and maps the next-hop L3ADDR into a next-hop Layer 2 address ("L2ADDR"), which in reality is the IPv4 address of an underlying interface of the ISATAP neighbor that may be many IPv4 hops away.

Address mapping for IPv4 differs from the IPv6 version in that no algorithmic mapping between L3ADDR and L2ADDR is possible. ISATAP for IPv4 therefore requires an L3ADDR->L2ADDR address mapping method that is coordinated on a per-site basis such that all nodes in the site follow the same convention. Examples include name service lookup (e.g., as described in [VET]), static mapping table lookup,

etc.

The node next performs an IPv4 FIB lookup on the next-hop L2ADDR to determine the correct underlying IPv4 interface. If the FIB lookup fails, the node drops the packet and returns an ICMPv4 "Destination Unreachable" message to the original source [RFC0792]; otherwise, it encapsulates the packet and submits it to the IPv4 layer as described below.

6.5. Encapsulation and Transmission

After performing the IPv4 FIB lookup on the next-hop L2ADDR, the node encapsulates the packet as specified in [RFC2003] with the IPv4 address of the underlying interface as the outer IPv4 source address and the next-hop L2ADDR as the outer IPv4 destination address. The node sets the DF flag in the outer IPv4 header according to Section 3.2 of [RFC4213]. The node also sets the IP protocol field in the outer IPv4 header to 4 (i.e., ip-protocol-4).

The node then submits the encapsulated packet to the IPv4 layer. The IPv4 layer fragments the packet if necessary, then forwards each fragment to the underlying IPv4 interface. The underlying IPv4 interface then performs address resolution on the outer IPv4 destination address (i.e., the next-hop L2ADDR) and submits the packet for transmission on the underlying link layer.

6.6. IPv4 Multicast Mapping

In many aspects, ISATAP is simply a unicast-only derivative of "6over4" [RFC2529]. For various reasons, however, ISATAP has seen practical wide-scale deployment while the 6over4 approach has been silently carried forward through ongoing research efforts. This specification extends the ISATAP interface model to support IPv4 multicast mapping in a manner that exactly parallels IPv6 multicast mapping in 6over4 (see [RFC2529], Section 6). Indeed, the approach might more aptly be named "4over4" were it not for the fact that the name "ISATAP" has already become ingrained in the widely published terminology.

IPv4 multicast mapping is available only on ISATAP interfaces configured over sites that support IPv4 multicast. For such sites, an IPv4 packet sent on an ISATAP interface with a multicast destination address DST MUST be encapsulated for transmission on an underlying IPv4 interface to the IPv4 multicast address of Organization-Local Scope using the mapping below. The mapped address SHOULD be taken from the block 239.193.0.0/16, a different sub-block of the Organization-Local Scope address block, or -- if none of those are available -- from the expansion blocks defined in [RFC2365].

Note that when they are formed using the expansion blocks, they use only a /16-sized block.

```
+-----+-----+-----+-----+
| 239   | OLS   | DST2  | DST3   |
+-----+-----+-----+-----+
```

DST2, DST3 Last two bytes of IPv4 multicast address.

OLS From the configured Organization-Local
Scope address block. SHOULD be 193;
see [RFC2365] for details.

Figure 1: ISATAPv4 Multicast Mapping

No new IANA registration procedures are required for the above.

6.7. Recursive Encapsulation Avoidance

The node must take care in managing its IPv4 FIB table entries in order to avoid looping through recursive encapsulations.

7. Security Considerations

The security considerations specified in [RFC2003] apply equally to this document. The security considerations specified in ISATAP [RFC5214] and 6over4 [RFC2529] also apply, with the exception that ip-protocol-4 encapsulation is used instead of ip-protocol-41.

Updated tunnel security considerations are found in [SECURITY].

8. Acknowledgements

This work extends the ISATAP interface model, which has evolved through the insights of many contributors over the course of many decades. Special thanks to Brian Carpenter and Jari Arkko for their helpful review input.

9. References

9.1. Normative References

- [RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.
- [RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, September 1981.

- [RFC1191] Mogul, J. and S. Deering, "Path MTU discovery", RFC 1191, November 1990.
- [RFC2003] Perkins, C., "IP Encapsulation within IP", RFC 2003, October 1996.
- [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.
- [RFC2529] Carpenter, B. and C. Jung, "Transmission of IPv6 over IPv4 Domains without Explicit Tunnels", RFC 2529, March 1999.
- [RFC3927] Cheshire, S., Aboba, B., and E. Guttman, "Dynamic Configuration of IPv4 Link-Local Addresses", RFC 3927, May 2005.
- [RFC4213] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, October 2005.
- [RFC5214] Templin, F., Gleeson, T., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", RFC 5214, March 2008.

9.2. Informative References

- [SECURITY] Hoagland, J., Krishnan, S., and D. Thaler, "Security Concerns With IP Tunneling", Work in Progress, October 2008.
- [VET] Templin, F., "Virtual Enterprise Traversal (VET)", RFC 5558, February 2010.
- [RANGER] Templin, F., "Routing and Addressing in Networks with Global Enterprise Recursion (RANGER)", RFC 5720, February 2010.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.
- [RFC2365] Meyer, D., "Administratively Scoped IP Multicast", BCP 23, RFC 2365, July 1998.

Appendix A. Encapsulation Avoidance

In some instances, an ISATAP interface may be configured over a site in which the L3ADDRs and L2ADDRs of all ISATAP neighbors are also known to be routable within the underlying site. In that case, the ISATAP interface MAY avoid encapsulation and submit the unencapsulated packets directly to the IPv4 layer. Note however that this approach does not provide for the use of indirection afforded through encapsulation.

Author's Address

Fred L. Templin (editor)
Boeing Research & Technology
P.O. Box 3707 MC 7L-49
Seattle, WA 98124
USA

EMail: fltemplin@acm.org

