

Network Working Group
Request for Comments: 4214
Category: Experimental

F. Templin
Nokia
T. Gleeson
Cisco Systems K.K.
M. Talwar
D. Thaler
Microsoft Corporation
October 2005

Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)

Status of This Memo

This memo defines an Experimental Protocol for the Internet community. It does not specify an Internet standard of any kind. Discussion and suggestions for improvement are requested. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2005).

IESG Note

The content of this RFC was at one time considered by the IETF, and therefore it may resemble a current IETF work in progress or a published IETF work. 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 RFC should exercise caution in evaluating its value for implementation and deployment. See RFC 3932 for more information.

Abstract

The Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) connects IPv6 hosts/routers over IPv4 networks. ISATAP views the IPv4 network as a link layer for IPv6 and views other nodes on the network as potential IPv6 hosts/routers. ISATAP supports an automatic tunneling abstraction similar to the Non-Broadcast Multiple Access (NBMA) model.

1. Introduction

This document specifies a simple mechanism called the Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) that connects IPv6 hosts/routers over IPv4 networks. Dual-stack (IPv6/IPv4) nodes use ISATAP to automatically tunnel IPv6 packets in IPv4, i.e., ISATAP views the IPv4 network as a link layer for IPv6 and views other nodes on the network as potential IPv6 hosts/routers.

ISATAP enables automatic tunneling whether global or private IPv4 addresses are used, and presents a Non-Broadcast Multiple Access (NBMA) abstraction similar to [RFC2491][RFC2492][RFC2529][RFC3056].

The main objectives of this document are to: 1) describe the domain of applicability, 2) specify addressing requirements, 3) specify automatic tunneling using ISATAP, 4) specify the operation of IPv6 Neighbor Discovery over ISATAP interfaces, and 5) discuss Site Administration, Security, and IANA considerations.

2. Requirements

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 [BCP14].

This document also uses internal conceptual variables to describe protocol behavior and external variables that an implementation must allow system administrators to change. The specific variable names, how their values change, and how their settings influence protocol behavior are provided in order to demonstrate protocol behavior. An implementation is not required to have them in the exact form described here, as long as its external behavior is consistent with that described in this document.

3. Terminology

The terminology of [RFC2460][RFC2461] applies to this document. The following additional terms are defined:

ISATAP node:

A node that implements the specifications in this document.

ISATAP interface:

An ISATAP node's Non-Broadcast Multi-Access (NBMA) IPv6 interface, used for automatic tunneling of IPv6 packets in IPv4.

ISATAP interface identifier:

An IPv6 interface identifier with an embedded IPv4 address constructed as specified in Section 6.1.

ISATAP address:

An IPv6 unicast address that matches an on-link prefix on an ISATAP interface of the node, and that includes an ISATAP interface identifier.

locator:

An IPv4 address-to-interface mapping; i.e., a node's IPv4 address and its associated interface.

locator set:

A set of locators associated with an ISATAP interface. Each locator in the set belongs to the same site.

4. Domain of Applicability

The domain of applicability for this technical specification is automatic tunneling of IPv6 packets in IPv4 for ISATAP nodes within sites that observe the security considerations found in this document, including host-to-router, router-to-host, and host-to-host automatic tunneling in certain enterprise networks and 3GPP/3GPP2 wireless operator networks. (Other scenarios with a sufficient trust basis ensured by the mechanisms specified in this document also fall within this domain of applicability.)

Extensions to the above domain of applicability (e.g., by combining the mechanisms in this document with those in other technical specifications) are out of the scope of this document.

5. Node Requirements

ISATAP nodes observe the common functionality requirements for IPv6 nodes found in [NODEREQ] and the requirements for dual IP layer operation found in ([MECH], Section 2). They also implement the additional features specified in this document.

6. Addressing Requirements

6.1. ISATAP Interface Identifiers

ISATAP interface identifiers are constructed in Modified EUI-64 format ([RFC3513], Section 2.5.1 and Appendix A) by concatenating the 24-bit IANA OUI (00-00-5E), the 8-bit hexadecimal value 0xFE, and a 32-bit IPv4 address in network byte order as follows:

0	1	3	6
0	5	1	3
+-----+			
0000000ug00000000	0101111011111110	
+-----+			

When the IPv4 address is known to be globally unique, the "u" bit (universal/local) is set to 1; otherwise, the "u" bit is set to 0. "g" is the individual/group bit, and "m" are the bits of the IPv4 address.

6.2. ISATAP Interface Address Configuration

Each ISATAP interface configures a set of locators consisting of IPv4 address-to-interface mappings from a single site; i.e., an ISATAP interface's locator set MUST NOT span multiple sites.

When an IPv4 address is removed from an interface, the corresponding locator SHOULD be removed from its associated locator set(s). When a new IPv4 address is assigned to an interface, the corresponding locator MAY be added to the appropriate locator set(s).

ISATAP interfaces form ISATAP interface identifiers from IPv4 addresses in their locator set and use them to create link-local ISATAP addresses ([RFC2462], Section 5.3).

6.3. Multicast/Anycast

It is not possible to assume the general availability of wide-area IPv4 multicast, so (unlike 6over4 [RFC2529]) ISATAP must assume that its underlying IPv4 carrier network only has unicast capability. Support for IPv6 multicast over ISATAP interfaces is not described in this document.

Similarly, support for Reserved IPv6 Subnet Anycast Addresses is not described in this document.

7. Automatic Tunneling

ISATAP interfaces use the basic tunneling mechanisms specified in ([MECH], Section 3). The following sub-sections describe additional specifications.

7.1. Encapsulation

ISATAP addresses are mapped to a link-layer address by a static computation; i.e., the last four octets are treated as an IPv4 address.

7.2. Handling ICMPv4 Errors

ISATAP interfaces SHOULD process ARP failures and persistent ICMPv4 errors as link-specific information indicating that a path to a neighbor may have failed ([RFC2461], Section 7.3.3).

7.3. Decapsulation

The specification in ([MECH], Section 3.6) is used. Additionally, when an ISATAP node receives an IPv4 protocol 41 datagram that does not belong to a configured tunnel interface, it determines whether the packet's IPv4 destination address and arrival interface match a locator configured in an ISATAP interface's locator set.

If an ISATAP interface that configures a matching locator is found, the decapsulator MUST verify that the packet's IPv4 source address is correct for the encapsulated IPv6 source address. The IPv4 source address is correct if:

- the IPv6 source address is an ISATAP address that embeds the IPv4 source address in its interface identifier, or
- the IPv4 source address is a member of the Potential Router List (see Section 8.1).

Packets for which the IPv4 source address is incorrect for this ISATAP interface are checked to determine whether they belong to another tunnel interface.

7.4. Link-Local Addresses

ISATAP interfaces use link-local addresses constructed as specified in Section 6 of this document.

7.5. Neighbor Discovery over Tunnels

ISATAP interfaces use the specifications for neighbor discovery found in the following section of this document.

8. Neighbor Discovery for ISATAP Interfaces

ISATAP interfaces use the neighbor discovery mechanisms specified in [RFC2461]. The following sub-sections describe specifications that are also implemented.

8.1. Conceptual Model of a Host

To the list of Conceptual Data Structures ([RFC2461], Section 5.1), ISATAP interfaces add the following:

Potential Router List (PRL)

A set of entries about potential routers; used to support router and prefix discovery. Each entry ("PRL(i)") has an associated timer ("TIMER(i)"), and an IPv4 address ("V4ADDR(i)") that represents a router's advertising ISATAP interface.

8.2. Router and Prefix Discovery - Router Specification

Advertising ISATAP interfaces send Solicited Router Advertisement messages as specified in ([RFC2461], Section 6.2.6) except that the messages are sent directly to the soliciting node; i.e., they might not be received by other nodes on the link.

8.3. Router and Prefix Discovery - Host Specification

The Host Specification in ([RFC2461], Section 6.3) is used. The following sub-sections describe specifications added by ISATAP interfaces.

8.3.1. Host Variables

To the list of host variables ([RFC2461], Section 6.3.2), ISATAP interfaces add the following:

PrlRefreshInterval

Time in seconds between successive refreshments of the PRL after initialization. The designated value of all ones (0xffffffff) represents infinity.
Default: 3600 seconds

MinRouterSolicitInterval

Minimum time in seconds between successive solicitations of the same advertising ISATAP interface. The designated value of all ones (0xffffffff) represents infinity.

8.3.2. Potential Router List Initialization

ISATAP nodes initialize an ISATAP interface's PRL with IPv4 addresses discovered via manual configuration, a DNS Fully Qualified Domain Name (FQDN) [STD13], a DHCPv4 option, a DHCPv4 vendor-specific option, or an unspecified alternate method. FQDNs are established via manual configuration or an unspecified alternate method. FQDNs are resolved into IPv4 addresses through a static host file lookup,

querying the DNS service, querying a site-specific name service, or with an unspecified alternate method.

After initializing an ISATAP interface's PRL, the node sets a timer for the interface to `PrlRefreshInterval` seconds and re-initializes the interface's PRL as specified above when the timer expires. When an FQDN is used, and when it is resolved via a service that includes TTLs with the IPv4 addresses returned (e.g., DNS 'A' resource records [STD13]), the timer SHOULD be set to the minimum of `PrlRefreshInterval` and the minimum TTL returned. (Zero-valued TTLs are interpreted to mean that the PRL is re-initialized before each Router Solicitation event; see Section 8.3.4.)

8.3.3. Processing Received Router Advertisements

To the list of checks for validating Router Advertisement messages ([RFC2461], Section 6.1.1), ISATAP interfaces add the following:

- IP Source Address is a link-local ISATAP address that embeds `V4ADDR(i)` for some `PRL(i)`.

Valid Router Advertisements received on an ISATAP interface are processed as specified in ([RFC2461], Section 6.3.4).

8.3.4. Sending Router Solicitations

To the list of events after which Router Solicitation messages may be sent ([RFC2461], Section 6.3.7), ISATAP interfaces add the following:

- `TIMER(i)` for some `PRL(i)` expires.

Since unsolicited Router Advertisements may be incomplete and/or absent, ISATAP nodes MAY schedule periodic Router Solicitation events for certain `PRL(i)`s by setting the corresponding `TIMER(i)`.

When periodic Router Solicitation events are scheduled, the node SHOULD set `TIMER(i)` so that the next event will refresh remaining lifetimes stored for `PRL(i)` before they expire, including the Router Lifetime, Valid Lifetimes received in Prefix Information Options, and Route Lifetimes received in Route Information Options [DEFAULT]. `TIMER(i)` MUST be set to no less than `MinRouterSolicitInterval` seconds where `MinRouterSolicitInterval` is configurable for the node, or for a specific `PRL(i)`, with a conservative default value (e.g., 2 minutes).

When `TIMER(i)` expires, the node sends Router Solicitation messages as specified in ([RFC2461], Section 6.3.7) except that the messages are sent directly to `PRL(i)`; i.e., they might not be received by other routers. While the node continues to require periodic Router

Solicitation events for PRL(i), and while PRL(i) continues to act as a router, the node resets TIMER(i) after each expiration event as described above.

8.4. Neighbor Unreachability Detection

Hosts SHOULD perform Neighbor Unreachability Detection ([RFC2461], Section 7.3). Routers MAY perform neighbor unreachability detection, but this might not scale in all environments.

After address resolution, hosts SHOULD perform an initial reachability confirmation by sending Neighbor Solicitation messages and receiving a Neighbor Advertisement message. Routers MAY perform this initial reachability confirmation, but this might not scale in all environments.

9. Site Administration Considerations

Site administrators maintain a Potential Router List (PRL) of IPv4 addresses representing advertising ISATAP interfaces of routers.

The PRL is commonly maintained as an FQDN for the ISATAP service in the site's name service (see Section 8.3.2). There are no mandatory rules for the selection of the FQDN, but site administrators are encouraged to use the convention "isatap.domainname" (e.g., isatap.example.com).

When the site's name service includes TTLs with the IPv4 addresses returned, site administrators SHOULD configure the TTLs with conservative values to minimize control traffic.

10. Security Considerations

Implementors should be aware that, in addition to possible attacks against IPv6, security attacks against IPv4 must also be considered. Use of IP security at both IPv4 and IPv6 levels should nevertheless be avoided, for efficiency reasons. For example, if IPv6 is running encrypted, encryption of IPv4 would be redundant unless traffic analysis is felt to be a threat. If IPv6 is running authenticated, then authentication of IPv4 will add little. Conversely, IPv4 security will not protect IPv6 traffic once it leaves the ISATAP domain. Therefore, implementing IPv6 security is required even if IPv4 security is available.

The threats associated with IPv6 Neighbor Discovery are described in [RFC3756].

There is a possible spoofing attack in which spurious ip-protocol-41 packets are injected into an ISATAP link from outside. Since an ISATAP link spans an entire IPv4 site, restricting access to the link can be achieved by restricting access to the site; i.e., by having site border routers implement IPv4 ingress filtering and ip-protocol-41 filtering.

Another possible spoofing attack involves spurious ip-protocol-41 packets injected from within an ISATAP link by a node pretending to be a router. The Potential Router List (PRL) provides a list of IPv4 addresses representing advertising ISATAP interfaces of routers that hosts use in filtering decisions. Site administrators should ensure that the PRL is kept up to date, and that the resolution mechanism (see Section 9) cannot be subverted.

The use of temporary addresses [RFC3041] and Cryptographically Generated Addresses [CGA] on ISATAP interfaces is outside the scope of this specification.

11. IANA Considerations

The IANA has specified the format for Modified EUI-64 address construction ([RFC3513], Appendix A) in the IANA Ethernet Address Block. The text in Appendix A of this document has been offered as an example specification. The current version of the IANA registry for Ether Types can be accessed at:

<http://www.iana.org/assignments/ethernet-numbers>

12. Acknowledgements

The ideas in this document are not original, and the authors acknowledge the original architects. Portions of this work were sponsored through SRI International internal projects and government contracts. Government sponsors include Monica Farah-Stapleton and Russell Langan (U.S. Army CECOM ASE0), and Dr. Allen Moshfegh (U.S. Office of Naval Research). SRI International sponsors include Dr. Mike Frankel, J. Peter Marcotullio, Lou Rodriguez, and Dr. Ambatipudi Sastry.

The following are acknowledged for providing peer review input: Jim Bound, Rich Draves, Cyndi Jung, Ambatipudi Sastry, Aaron Schrader, Ole Troan, and Vlad Yasevich.

The following are acknowledged for their significant contributions: Alain Durand, Hannu Flinck, Jason Goldschmidt, Nathan Lutchansky, Karen Nielsen, Mohan Parthasarathy, Chirayu Patel, Art Shelest, Markku Savela, Pekka Savola, Margaret Wasserman, and Brian Zill.

The authors acknowledge the work of Quang Nguyen on "Virtual Ethernet", under the guidance of Dr. Lixia Zhang, that proposed very similar ideas to those that appear in this document. This work was first brought to the authors' attention on September 20, 2002.

Appendix A. Modified EUI-64 Addresses in the IANA Ethernet Address Block

Modified EUI-64 addresses ([RFC3513], Section 2.5.1 and Appendix A) in the IANA Ethernet Address Block are formed by concatenating the 24-bit IANA OUI (00-00-5E) with a 40-bit extension identifier and inverting the "u" bit; i.e., the "u" bit is set to one (1) to indicate universal scope and set to zero (0) to indicate local scope.

Modified EUI-64 addresses have the following appearance in memory (bits transmitted right-to-left within octets, octets transmitted left-to-right):

0		23		63
	OUI		extension identifier	
000000ug00000000 01011110xx				

When the first two octets of the extension identifier encode the hexadecimal value 0xFFFFE, the remainder of the extension identifier encodes a 24-bit vendor-supplied id as follows:

0		23		39		63
	OUI		0xFFFFE		vendor-supplied id	
000000ug00000000 0101111011111111 11111110xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx						

When the first octet of the extension identifier encodes the hexadecimal value 0xFE, the remainder of the extension identifier encodes a 32-bit IPv4 address as follows:

0		23		31		63
	OUI		0xFE		IPv4 address	
000000ug00000000 0101111011111110 xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx						

Normative References

- [BCP14] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [STD13] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [RFC2461] Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.

- [RFC2462] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.
- [RFC3513] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", RFC 3513, April 2003.
- [MECH] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, October 2005.

Informative References

- [RFC2491] Armitage, G., Schulter, P., Jork, M., and G. Harter, "IPv6 over Non-Broadcast Multiple Access (NBMA) networks", RFC 2491, January 1999.
- [RFC2492] Armitage, G., Schulter, P., and M. Jork, "IPv6 over ATM Networks", RFC 2492, January 1999.
- [RFC2529] Carpenter, B. and C. Jung, "Transmission of IPv6 over IPv4 Domains without Explicit Tunnels", RFC 2529, March 1999.
- [RFC3041] Narten, T. and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 3041, January 2001.
- [RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, February 2001.
- [RFC3756] Nikander, P., Kempf, J., and E. Nordmark, "IPv6 Neighbor Discovery (ND) Trust Models and Threats", RFC 3756, May 2004.
- [CGA] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, March 2005.
- [DEFAULT] Draves, R. and D. Thaler, "Default Router Preferences and More-Specific Routes", Work in Progress, December 2003.
- [NODEREQ] Loughney, J., Ed., "IPv6 Node Requirements", Work in Progress, May 2004.

Authors' Addresses

Fred L. Templin
Nokia
313 Fairchild Drive
Mountain View, CA 94110
US

EMail: fltemplin@acm.org

Tim Gleeson
Cisco Systems K.K.
Shinjuku Mitsui Building
2-1-1 Nishishinjuku, Shinjuku-ku
Tokyo 163-0409
Japan

EMail: tgleeson@cisco.com

Mohit Talwar
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052-6399
US

Phone: +1 425 705 3131
EMail: mohitt@microsoft.com

Dave Thaler
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052-6399
US

Phone: +1 425 703 8835
EMail: dthaler@microsoft.com

Full Copyright Statement

Copyright (C) The Internet Society (2005).

This document is subject to the rights, licenses and restrictions contained in BCP 78 and at www.rfc-editor.org/copyright.html, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

