Network Working Group
Internet-Draft
Intended status: Experimental
Expires: November 10, 2012

F. Templin
Boeing Research & Technology
May 09, 2012

ISATAP Updates
draft-templin-isupdate-04.txt

Abstract

Many end user sites in the Internet today still have predominantly IPv4 internal infrastructures. These sites range in size from small home/office networks to large corporate enterprise networks, but share the commonality that IPv4 continues to provide operational internal routing and addressing services for most applications. As more and more IPv6-only services are deployed, however, end user devices within such sites will increasingly require at least basic IPv6 functionality. This document therefore discusses updates to the Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) to better accommodate these needs.

Table of Contents

1. Introduction

End user sites in the Internet today currently use IPv4 routing and addressing internally for core operating functions such as web browsing, filesharing, network printing, e-mail, teleconferencing and numerous other site-internal networking services. Such sites typically have an abundance of public or private IPv4 addresses for internal networking, and are separated from the public Internet by firewalls, packet filtering gateways, proxies, address translators and other site border demarcation devices. To date, such sites have had little incentive to enable IPv6 services internally [RFC1687].

End-user sites that currently use IPv4 services internally come in endless sizes and varieties. For example, a home network behind a Network Address Translator (NAT) may consist of a single link supporting a few laptops, printers etc. As a larger example, a small business may consist of one or a few offices with several networks connecting considerably larger numbers of computers, routers, handheld devices, printers, faxes, etc. Moving further up the scale, large banks, restaurants, major retailers, large corporations, etc. may consist of hundreds or thousands of branches worldwide that are tied together in a complex global enterprise network. Additional examples include personal-area networks, mobile vehicular networks, disaster relief networks, tactical military networks, and various forms of Mobile Ad-hoc Networks (MANETs), etc.

With the proliferation of IPv6 devices in the public Internet, however, existing IPv4 sites will increasingly require a means for enabling IPv6 services so that hosts within the site can communicate with IPv6-only correspondents. Such services must be deployable with minimal configuration, and in a fashion that will not cause disruptions to existing IPv4 services. The Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) [RFC5214] provides a simple-to-use service that sites can deploy in the near term to meet these requirements, as discussed in [I-D.templin-v6ops-isops]. However, the ISATAP base specification has several fundamental limitations that make for a "spartan" IPv6 service. This document discusses the motivations for new functionality followed by the updates and operational practices necessary to provide a more fully-functioned service.

2. Motivation

The base ISATAP specification does not support stateful address configuration nor prefix delegation (e.g., via DHCPv6 [RFC3315][RFC3633]) on ISATAP interfaces. Instead, the base specification requires a special IPv6 address format in which a

node's site-internal IPv4 address is embedded literally within the interface identifier of its public IPv6 address. This exposes the site-internal IPv4 address structure to IPv6 networks and correspondents outside of the site. Furthermore, static linkage of the node's site-internal IPv4 address to its public IPv6 address limits the node's ability to renumber its IPv4 address without also deprecating the IPv6 address. These limitations may be more of a concern in some ISATAP deployments than others, but can be obviated by address configuration methods that support non-ISATAP interface identifiers.

The ISATAP base specification further does not support router-to-router tunneling, i.e., it permits only router-to-host and host-to-host tunneling. In practical terms, the ISATAP base specification therefore does not allow for deployment of "stub" IPv6-only networks inside of a parent site. Examples include an IPv6-only bluetooth network of embedded devices, a laptop user's personal-area network, an IPv6-only fileshare workgroup, etc. Without updates to the ISATAP base specification, these limitations could only be addressed by a site-wide native IPv6 deployment, which the site may not be prepared to finance or support in the near term.

Finally, the base specification provides no means for address selection preference of IPv4 over ISATAP for communications within the same site. Although this need could be addressed in the future by a DHCP option [I-D.ietf-6man-addr-select-opt], it may be necessary or preferable in some environments for ISATAP clients to discover address selection preferences only from the information advertised by ISATAP routers. This document therefore specifies updates to the base specification to address these needs.

3. ISATAP Updates

The basic ISATAP model supports two basic node types - namely, advertising ISATAP routers and ISATAP hosts. Advertising ISATAP routers configure their site-facing ISATAP interfaces as advertising router interfaces (see: [RFC4861], Section 6.2.2). ISATAP hosts configure their site-facing ISATAP interfaces as simple host interfaces and also coordinate their autoconfiguration operations with advertising ISATAP routers.

This document introduces a third node type known as "non-advertising ISATAP routers". Non-advertising ISATAP routers configure their site-facing ISATAP interfaces as non-advertising router interfaces and obtain IPv6 addresses/prefixes via manual or automatic configuration arrangements with advertising ISATAP routers. Non-advertising ISATAP routers connect IPv6 networks to the ISATAP link,

and can therefore support a router-to-router tunneling mode not supported under the base specification.

To support this router-to-router tunneling (and also to support the assignment of native IPv6 addresses on ISATAP interfaces) ISATAP nodes add an update to the existing source address verification checks specified in Section 7.3 of [RFC5214]. Namely, the node also considers the outer IPv4 source address correct for the inner IPv6 source address if:

- a stateful address mapping exists that lists the packet's IPv4 source address as the link-layer address corresponding to the inner IPv6 source address via the ISATAP interface.

The basic ISATAP model further does not specify any IPv6 multicast mappings. This precludes the use of services such as DHCPv6 which require a link-scoped IPv6 multicasting service. To support DHCPv6 services, ISATAP hosts and non-advertising ISATAP routers that observe this specification map the IPv6 "All_DHCP_Relay_Agents_and_Servers" link-scoped multicast address to the IPv4 address of an advertising ISATAP router that advertises availability of the DHCPv6 service. The advertising ISATAP router in turn configures a DHCPv6 server or relay function, and accepts DHCPv6 messages sent by clients using this mapping. The advertising router also maintains a stateful address mapping that lists the IPv4 address of the client as the link-layer address of any delegated IPv6 addresses or prefixes.

Finally, this document updates the address selection policies of the base specification as follows. For communications between two nodes whose IPv6 addresses are covered by the same IPv6 prefix advertised in Router Advertisements (RAs) on an ISATAP interface, prefer IPv4 over IPv6 if the L bit in the Prefix Information Option (PIO) is set to 0.

Using these updates, a much richer ISATAP service model is made possible. The following sections describe the new modes of operation that are enabled by the updates.

4. Advanced IPv6 Services Enabled by Updates

Whether or not advertising ISATAP routers make stateless IPv6 services available using StateLess Address AutoConfiguration (SLAAC), they can also provide advanced IPv6 services to ISATAP clients (i.e., both hosts and non-advertising ISATAP routers) using the updates specified in this document. Any addresses/prefixes obtained via the advanced (stateful) services are distinct from any IPv6 prefixes

advertised on the ISATAP interface for SLAAC purposes, however.

The following sections discuss operational considerations for enabling ISATAP DHCPv6 services within predominantly IPv4 sites.

4.1. Advertising ISATAP Router Behavior

Advertising ISATAP routers that support DHCPv6 services send IPv6-in-IPv4 encapsulated RA messages that advertise availability of the service in response to IPv6-in-IPv4 encapsulated Router Solicitation (RS) messages received on an advertising ISATAP interface. They also configure either a DHCPv6 relay or server function to service DHCPv6 requests received from ISATAP clients.

4.2. ISATAP Host Behavior

ISATAP hosts send RS messages to obtain RA messages from an advertising ISATAP router. When the DHCPv6 service is available, the host can acquire IPv6 addresses through the use of DHCPv6 stateful address autoconfiguration [RFC3315] whether or not IPv6 prefixes for SLAAC are advertised. To acquire addresses, the host performs standard DHCPv6 exchanges while mapping the IPv6 "All_DHCP_Relay_Agents_and_Servers" link-scoped multicast address to the IPv4 address of an advertising ISATAP router that supports the DHCPv6 service.

After the host receives IPv6 addresses, it assigns them to its ISATAP interface and forwards any of its outbound IPv6 packets via the advertising router as a default router. The advertising router in turn maintains stateful address mappings that list the IPv4 address of the host as the link-layer address of the delegated IPv6 addresses. Note that IPv6 addresses acquired from DHCPv6 therefore need not be ISATAP addresses, i.e., even though the addresses are assigned to the ISATAP interface.

4.3. Non-Advertising ISATAP Router Behavior

Non-advertising ISATAP routers send RS messages to obtain RA messages from an advertising ISATAP router, i.e., they act as "hosts" on their non-advertising ISATAP interfaces. Non-advertising ISATAP routers can acquire IPv6 prefixes through the use of DHCPv6 Prefix Delegation [RFC3633] via an advertising router that supports DHCPv6 services in the same fashion as described above for host-based address autoconfiguration. The advertising router in turn maintains stateful address mappings that list the IPv4 address of the non-advertising router as the link-layer address of the next hop toward the delegated IPv6 prefixes.

In many use case scenarios (e.g., small enterprise networks, small and stable MANETs, etc.), advertising and non-advertising ISATAP routers can engage in a proactive dynamic IPv6 routing protocol (e.g., OSPFv3, RIPng, etc.) over their ISATAP interfaces so that IPv6 routing/forwarding tables can be populated and standard IPv6 forwarding between ISATAP routers can be used. In other scenarios (e.g., large enterprise networks, large and dynamic MANETs, etc.), this might be impractical due to scaling issues.

After the non-advertising ISATAP router acquires IPv6 prefixes, it can sub-delegate them to routers and links within its attached IPv6 edge networks, then can forward any outbound IPv6 packets coming from its edge networks via other nodes on the ISATAP link.

4.4. Reference Operational Scenario

Figure 1 depicts a reference ISATAP network topology enabled by the updated ISATAP services specified in this document. The scenario shows two advertising ISATAP routers ('A', 'B'), two non-advertising ISATAP routers ('C', 'E'), an ISATAP host ('G'), and three ordinary IPv6 hosts ('D', 'F', 'H') in a typical deployment configuration:

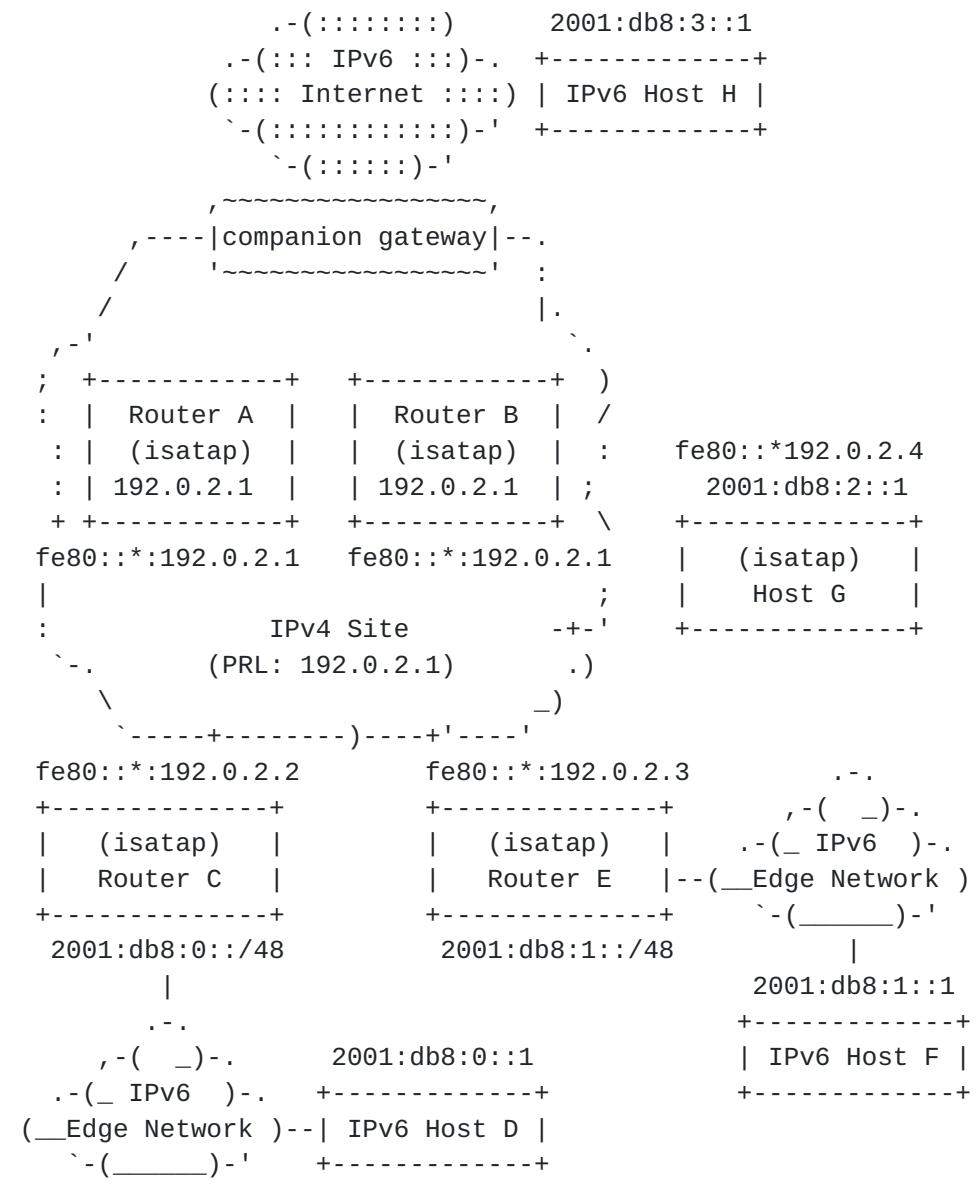

Figure 1: Reference ISATAP Network Topology

In Figure 1, advertising ISATAP routers 'A' and 'B' within the IPv4 site provide DHCPv6 services and connect to the IPv6 Internet either directly or via a companion gateway. The advertising routers both configure the IPv4 anycast address 192.0.2.1 on a site-interior IPv4 interface, and configure an advertising ISATAP interface with link-local ISATAP address fe80::5efe:192.0.2.1. The site administrator then places the single IPv4 address 192.0.2.1 in the Potential Router List (PRL) for the site. 'A' and 'B' then both advertise the anycast address/prefix into the site's IPv4 routing system so that ISATAP clients can locate the router that is topologically closest. (Note:

advertising ISATAP routers can instead use individual IPv4 unicast addresses instead of a shared IPv4 anycast address. In that case, the PRL may contain multiple IPv4 addresses of advertising routers.)

Non-advertising ISATAP router 'C' connects to one or more IPv6 edge networks and also connects to the site via an IPv4 interface with address 192.0.2.2. 'C' next configures a non-advertising ISATAP router interface with link-local ISATAP address fe80::5efe:192.0.2.2, then discovers router 'A' via an RS/RA exchange. 'C' next receives the IPv6 prefix 2001:db8:0::/48 through a DHCPv6 prefix delegation exchange via 'A', then engages in an IPv6 routing protocol over its ISATAP interface and announces the delegated IPv6 prefix. 'C' finally sub-delegates the prefix to its attached edge networks, where IPv6 host 'D' autoconfigures the address 2001:db8:0::1.

Non-advertising ISATAP router 'E' connects to the site, configures its ISATAP interface, performs an RS/RA exchange, receives a DHCPv6 prefix delegation, and engages in the IPv6 routing protocol the same as for 'C'. In particular, 'E' configures the IPv4 address 192.0.2.3 and the link-local ISATAP address fe80::5efe:192.0.2.3. 'E' then receives the delegated IPv6 prefix 2001:db8:1::/48 and sub-delegates the prefix to its attached edge networks, where IPv6 host 'F' autoconfigures IPv6 address 2001:db8:1::1.

ISATAP host 'G' connects to the site via an IPv4 interface with address 192.0.2.4, and also configures an ISATAP host interface with link-local ISATAP address fe80::5efe:192.0.2.4 over the IPv4 interface. 'G' next performs an RS/RA exchange to discover 'B" and configures a default IPv6 route with next-hop address fe80::5efe: 192.0.2.1. 'G' then receives the IPv6 address 2001:db8:2::1 via a DHCPv6 address configuration exchange via 'B'; it then assigns the address to the ISATAP interface but does not assign a non-link-local IPv6 prefix to the interface.

Finally, IPv6 host 'H' connects to an IPv6 network outside of the ISATAP domain. 'H' configures its IPv6 interface in a manner specific to its attached IPv6 link, and autoconfigures the IPv6 address 2001:db8:3::1.

Following this autoconfiguration, when host 'D' has an IPv6 packet to send to host 'F', it prepares the packet with source address 2001: db8:0::1 and destination address 2001:db8:1::1, then sends the packet into the edge network where IPv6 forwarding will eventually convey it to router 'C'. 'C' then uses IPv6-in-IPv4 encapsulation to forward the packet to router 'E', since it has discovered a route to 2001: db8:1::/48 with next hop 'E' via dynamic routing over the ISATAP interface. Router 'E' finally sends the packet into the edge network where IPv6 forwarding will eventually convey it to host 'F'.

In a second scenario, when 'D' has a packet to send to ISATAP host 'G', it prepares the packet with source address 2001:db8:0::1 and destination address 2001:db8:2::1, then sends the packet into the edge network where it will eventually be forwarded to router 'C' the same as above. 'C' then uses IPv6-in-IPv4 encapsulation to forward the packet to router 'A' (i.e., 'C's default router), which in turn forwards the packet to 'G'. Note that this operation entails two hops across the ISATAP link (i.e., one from 'C' to 'A', and a second from 'A' to 'G'). If 'G' also participates in the dynamic IPv6 routing protocol, however, 'C' could instead forward the packet directly to 'G' without involving 'A'.

In a third scenario, when 'D' has a packet to send to host 'H' in the IPv6 Internet, the packet is forwarded to 'C' the same as above. 'C' then forwards the packet to 'A', which forwards the packet into the IPv6 Internet.

In a final scenario, when 'G' has a packet to send to host 'H' in the IPv6 Internet, the packet is forwarded directly to 'B', which forwards the packet into the IPv6 Internet.

4.5. Site Administration Guidance

Site administrators configure advertising ISATAP routers that also support the DHCPv6 relay/server function to send RA messages with the M flag set to 1 as an indication to clients that the stateful DHCPv6 address autoconfiguration services are available. If stateless DHCPv6 services are also available, the RA messages also set the O flag to 1.

Gateways and packet filtering devices of various forms are often deployed in order to divide the site into separate partitions. Although the purely stateful model does not involve the advertisement of non-link-local IPv6 prefixes on ISATAP interfaces, alignment of IPv6 prefixes used for stateful address assignment with IPv4 site partitions is still recommended so that ISATAP clients can prefer native IPv4 communications over ISATAP IPv6 services for correspondents within their contiguous IPv4 partition.

For example, if the site is assigned the aggregate prefix 2001:db8:0::/48, then the site administrators can assign the more-specific prefixes 2001:db8:0:0::/64, 2001:db8:0:1::/64, 2001:db8:0:2::/64, etc. to the different IPv4 partitions within the site. The administrators can then institute a policy that prefers native IPv4 addresses for communications between clients covered by the same /64 prefix.

Site administrators can implement this policy implicitly by

configuring advertising ISATAP routers to advertise each /64 prefix with both the A and L flags set to 0 as an indication that IPv4 should be preferred over IPv6 destinations that configure addresses from the same prefix. Site administrators can instead (or in addition) implement address selection policy rules [RFC3484] through explicit configurations in each ISATAP client.

For example, each ISATAP client associated with the prefix 2001:db8:0:0::/64 can add the prefix to its address selection policy table with a lower precedence than the prefix ::ffff:0:0/96. In this way, IPv4 addresses are preferred over IPv6 addresses from within the same /64 prefix. The prefix could be added to each ISATAP client either manually, or through an automated service such as a DHCP option [I-D.ietf-6man-addr-select-opt]. In this way, clients will use IPv4 communications to reach correspondents within the same IPv4 site partition, and will use IPv6 communications to reach correspondents in other partitions and/or outside of the site.

When the PRL includes an anycast address, the client may be directed to a first DHCPv6 relay/server in initial message exchanges and to a different relay/server in subsequent exchanges. In order to address this uncertainty, site administrators should configure DHCPv6 servers to include a Server Unicast option so that clients can remain associated with the same server that was reached during the initial exchange. (Alternatively, the administrator could arrange for the site's DHCPv6 servers to maintain a distributed database of client bindings.)

Finally, site administrators should configure ISATAP routers to not send ICMPv6 Redirect messages to inform a source client of a better next hop toward the destination unless there is strong assurance that the client and the next hop are within the same IPv4 site partition.

4.6. On-Demand Dynamic Routing

With respect to the reference operational scenarios depicted in Figure 1, there may be use cases in which a proactive dynamic IPv6 routing protocol cannot be used. For example, in large enterprise network deployments it would be impractical for all ISATAP routers to engage in a common routing protocol instance due to scaling considerations.

In those cases, an on-demand routing capability can be enabled in which ISATAP nodes send initial packets via an advertising ISATAP router and receive redirection messages back. For example, when a non-advertising ISATAP router 'C' has a packet to send to a host located behind non-advertising ISATAP router 'E', it can send the initial packets via advertising router 'A' which will return

redirection messages to inform 'C' that 'E' is a better first hop. Protocol details for this redirection procedure (including a means for detecting whether the direct path is usable) are specified in [I-D.templin-aero].

4.7. Loop Avoidance

When no advertising ISATAP routers advertise IPv6 prefixes for SLAAC purposes, no non-link-local IPv6 prefixes are assigned to ISATAP router interfaces. In that case, an ISATAP router cannot mistake another router for an ISATAP host due to an address that matches an on-link prefix. This corresponds to the mitigation documented in Section 3.2.4 of [RFC6324].

Any routing loops introduced in the stateful scenario would therefore be due to a misconfiguration in IPv6 routing the same as for any IPv6 router, and hence are out of scope for this document.

5. Manual Configuration

In addition to any SLAAC and/or DHCPv6 services, when the updates in this document are employed site administrators can use manual configuration to assign non-ISATAP IPv6 addresses to the ISATAP interfaces of client end systems. Site administrators can also use manual configuration to assign IPv6 prefixes to non-advertising ISATAP routers instead of (or in addition to) using DHCPv6 prefix delegation.

The IPv6 prefixes used for manual configuration must be distinct from any prefixes used for SLAAC, however they may overlap with the prefixes used for DHCPv6 as long as there is administrative assurance that the same IPv6 addresses/prefixes will not be delegated by both DHCPv6 and manual configuration. The manual configuration scenarios and routing considerations are otherwise the same as discussed in Section 4.

When manually configured IPv6 addresses/prefixes are used, the prefixes must be covered by a shorter IPv6 prefix advertised into the IPv6 routing system by one or more advertising ISATAP routers. The advertising routers must further maintain stateful address mappings that associate the addresses/prefixes with the ISATAP clients to which the addresses/prefixes are delegated, i.e., the same as for DHCPv6.

6. IANA Considerations

This document has no IANA considerations.

7. Security Considerations

In addition to the security considerations documented in [RFC5214], sites that use ISATAP should take care to ensure that no routing loops are enabled [RFC6324]. Additional security concerns with IP tunneling are documented in [RFC6169].

8. Acknowledgments

The following are acknowledged for their insights that helped shape this work: Dmitry Anipko, Fred Baker, Ron Bonica, Brian Carpenter, Remi Despres, Thomas Henderson, Philip Homburg, Lee Howard, Ray Hunter, Joel Jaeggli, John Mann, Gabi Nakibly, Christoper Palmer, Hemant Singh, Mark Smith, Dave Thaler, Ole Troan, and Gunter Van de Velde.

9. References

9.1. Normative References

[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.

[RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.

[RFC5214] Templin, F., Gleeson, T., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", RFC 5214, March 2008.

9.2. Informative References

[I-D.ietf-6man-addr-select-opt] Matsumoto, A., Fujisaki, T., Kato, J., and T. Chown, "Distributing Address Selection Policy using DHCPv6",

draft-ietf-6man-addr-select-opt-03 (work in progress), February 2012.

[I-D.templin-aero] Templin, F., "Asymmetric Extended Route Optimization (AERO)", draft-templin-aero-08 (work in progress), February 2012.

[I-D.templin-v6ops-isops] Templin, F., "Operational Guidance for IPv6 Deployment in IPv4 Sites using ISATAP", draft-templin-v6ops-isops-16 (work in progress), May 2012.

[RFC1687] Fleischman, E., "A Large Corporate User's View of IPng", RFC 1687, August 1994.

[RFC3484] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, February 2003.

[RFC5720] Templin, F., "Routing and Addressing in Networks with Global Enterprise Recursion (RANGER)", RFC 5720, February 2010.

[RFC6169] Krishnan, S., Thaler, D., and J. Hoagland, "Security Concerns with IP Tunneling", RFC 6169, April 2011.

[RFC6324] Nakibly, G. and F. Templin, "Routing Loop Attack Using IPv6 Automatic Tunnels: Problem Statement and Proposed Mitigations", RFC 6324, August 2011.

Author's Address

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

Email: fltemplin@acm.org