Network Working Group	F. Templin
Internet-Draft	Boeing Research & Technology
Intended status: Informational	May 05, 2011
Expires: November 06, 2011	

Operational Guidance for IPv6 Deployment in IPv4 Sites using ISATAP
draft-templin-v6ops-isops-00.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 satisfactory internal routing and addressing services for most applications. As more and more IPv6-only services are deployed in the Internet, however, end user devices within such sites will increasingly require at least basic IPv6 functionality for external access. It is also expected that more and more IPv6-only devices will be deployed within the site over time. This document therefore provides operational guidance for deployment of IPv6 within predominantly IPv4 sites using the Intra-Site Automatic Tunnel Addressing Protocol (ISATAP).

Status of this Memo

Copyright Notice

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). These cases and more are considered in RANGERS[RFC6139].

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. This document therefore provides operational guidance for using ISATAP to enable IPv6 services within predominantly IPv4 sites while causing no disruptions to existing IPv4 services.

2. Enabling IPv6 Services using ISATAP

Many existing sites within the Internet predominantly use IPv4-based services for their internal networking needs, but there is a growing requirement for enabling IPv6 services to support communications with IPv6-only correspondents. Smaller sites that wish to enable IPv6 typically arrange to obtain public IPv6 prefixes from an Internet Service Provider (ISP), where the prefixes may be either purely native or the near-native prefixes offered by 6rd [RFC5969]. Larger sites typically obtain provider independent IPv6 prefixes from an Internet registry and advertise the prefixes into the IPv6 routing system on their own behalf, i.e., they act as an ISP unto themselves. In either case, after obtaining IPv6 prefixes the site can automatically enable IPv6 services internally by configuring ISATAP.

The ISATAP service uses a Non-Broadcast, Multiple Access (NBMA) tunnel virtual interface model [RFC2491][RFC2529] based on IPv6-in-IPv4 encapsulation [RFC4213]. The service is further based on three basic node types known as advertising ISATAP routers, non-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). Non-advertising ISATAP routers configure their site-facing ISATAP interfaces as non-advertising router interfaces and obtain IPv6 addresses/prefixes via

autoconfiguration exchanges with advertising ISATAP routers. Finally, ISATAP hosts configure their site-facing ISATAP interfaces as simple host interfaces and also coordinate their autoconfiguration operations with advertising ISATAP routers.

Advertising ISATAP routers arrange to add their IPv4 addresses to the Potential Router List (PRL) within the site name service. The name service could be either the DNS or some other site-internal name resolution system, but the PRL should be published in such a way that ISATAP nodes can resolve the name "isatap.domainname" for the "domainname" suffix associated with their attached link. For example, if the domainname suffix associated with an ISATAP node's attached link is "example.com", then the name of the PRL for that link attachment point is "isatap.example.com". On the other hand, if the site name service is operating without a domainname suffix, then the name of the PRL is simply "isatap".

After the PRL is published, ISATAP nodes within the site will automatically discover advertising ISATAP routers and perform a Router Solicitation (RS) / Router Advertisement (RA) exchange to initiate Stateless Address AutoConfiguration (SLAAC), the Dynamic Host Configuration Protocol for IPv6 (DHCPv6) or both. The nodes can then use SLAAC-provided IPv6 addresses for basic IPv6 services and DHCPv6-provided IPv6 addresses/prefixes for fully-qualified IPv6 services.

3. SLAAC Services

Predominantly IPv4 sites can enable ISATAP SLAAC services for the purpose of providing basic IPv6 services to IPv4 hosts that need to communicate with IPv6-only correspondents. In order to provide a simple service that does not interact poorly with existing site topological arrangements, the site should not publish any ISATAP-provided IPv6 addresses that were configured using SLAAC within the site name service. Hence, ISATAP-provided SLAAC services are typically used primary for client-side operation. The following sections discuss operational considerations for enabling ISATAP SLAAC services within predominantly IPv4 sites.

3.1. ISATAP Router Behavior

Advertising ISATAP routers that support SLAAC services send RA messages in response to RS messages received on an advertising ISATAP interface. SLAAC services are enabled when advertising ISATAP routers advertise non-link-local IPv6 prefixes. When there are multiple advertising ISATAP routers, the routers can advertise the same IPv6 prefixes or a different set of IPv6 prefixes. For example, a first router may advertise 2001:db8:1::/64, a second may advertise 2001:db8:2::/64, etc.

The routers can further be configured to advertise different prefixes to different sets of hosts within the site (e.g., as identified by the host's IPv4 prefix) for the purpose of site partitioning. To discourage direct communications between ISATAP hosts using SLAAC-provided addresses, advertising ISATAP routers can send RAs that include Prefix Information Options (PIOs) with the (A, L) flags set to (1,0) [RFC4861].

3.2. ISATAP Host Behavior

ISATAP hosts resolve the PRL and send RS messages to obtain RA messages from an advertising ISATAP router. ISATAP routers that advertise prefixes for SLAAC purposes will typically advertise prefixes in PIOs with the (A, L) flags set to (1,0). In that case, the ISATAP host autoconfigures an address from the advertised IPv6 prefix and assigns the address to the ISATAP interface, but the host does not assign an IPv6 prefix to the ISATAP interface. Therefore, all IPv6 communications from the hosts will (initially) flow through the advertising ISATAP router. This arrangement prevents communication failure modes in which a pair of ISATAP hosts that use SLAAC are separated by a packet filtering gateway that would prevent direct communications via the tunneled IPv6 service.

3.3. Reference Operational Scenario

Figure 1 depicts a reference ISATAP network topology for allowing hosts within a predominantly IPv4 site to configure IPv6 services using ISATAP with SLAAC. The scenario shows two advertising ISATAP routers ('A', 'B'), two ISATAP hosts ('C', 'D'), and an ordinary IPv6 host ('E') outside of the site in a typical deployment configuration:

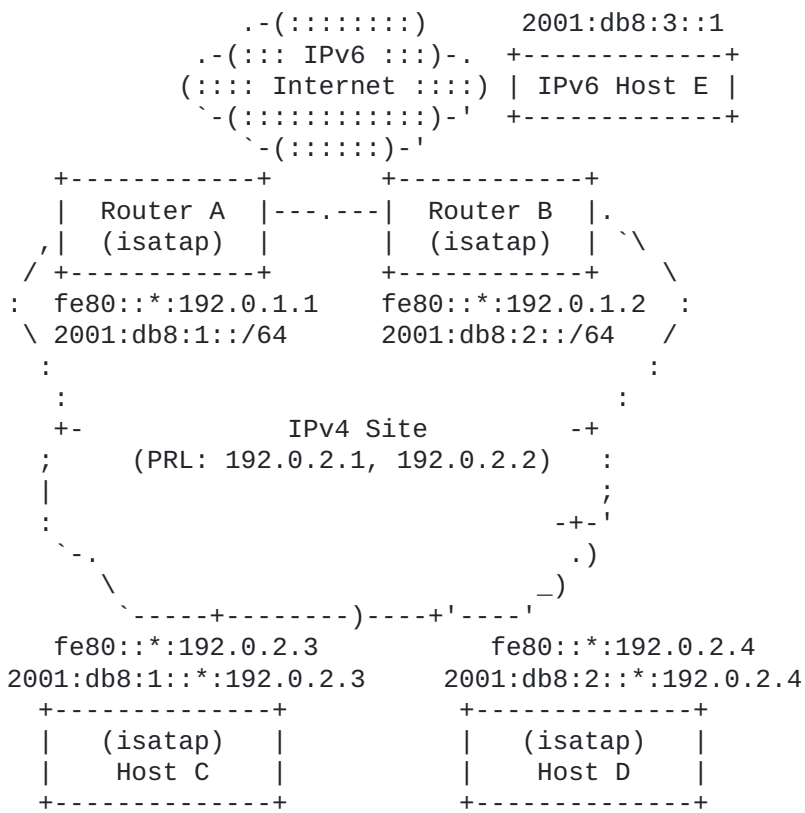

(* == "5efe")

In Figure 1, advertising ISATAP routers 'A' and 'B' within the IPv4 site connect to the IPv6 Internet. (Note that the routers may instead connect to the IPv6 Internet via a companion gateway as shown in Figure 2.) Advertising ISATAP router 'A' configures a site-interior IPv4 interface with address 192.0.2.1 and arranges to add

the address to the site's PRL. 'A' next configures an advertising ISATAP router interface with link-local IPv6 address fe80::5efe:192.0.2.1 over the IPv4 interface. In the same fashion, 'B' configures the IPv4 interface address 192.0.2.2, adds the address to the PRL, then configures its advertising ISATAP router interface with link-local address fe80::5efe:192.0.2.2.

ISATAP host 'C' connects to the site via an IPv4 interface with address 192.0.2.3, and also configures an ISATAP host interface with link-local address fe80::5efe:192.0.2.3 over the IPv4 interface. 'C' next resolves the PRL to discover the address 192.0.2.1 and performs an RS/RA exchange with 'A'. Based on the RA information, 'C' next configures a default IPv6 route with next-hop address fe80::5efe:192.0.2.1 via the ISATAP interface and processes the IPv6 prefix 2001:db8:1::/64 advertised in the PIO. When 'C' processes the prefix, it uses SLAAC to automatically configure the address 2001:db8:1::5efe:192.0.2.3. 'C' then assigns the address to the ISATAP interface, but does not assign the prefix itself to the interface if the 'L' bit in the PIO is 0.

In the same fashion, ISATAP host 'D' configures its IPv4 interface with address 192.0.2.4 and configures its ISATAP interface with link-local address fe80::5efe:192.0.2.4. 'D' next performs an RS/RA exchange with 'B', then uses SLAAC to autoconfigure the address 2001:db8:2::5efe:192.0.2.4.

Finally, IPv6 host 'E' connects to an IPv6 network outside of the site. 'E' 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 'C' has an IPv6 packet to send to host 'E', it prepares the packet with source address 2001:db8::5efe:192.0.2.3 and destination address 2001:db8:3::1. 'C' then uses IPv6-in-IPv4 encapsulation to forward the packet to router 'A', which in turn decapsulates the packet and forwards it into the public IPv6 Internet where it will be conveyed to 'E' via normal IPv6 routing. (Note that 'A' may "translate" the packet as it is forwarded across the site boundary such that it appears to come from a different source address than the one used by host 'C' within the site.) In the same fashion, host 'D' uses IPv6-in-IPv4 encapsulation via its default router 'B' to send IPv6 packets to IPv6 Internet hosts such as 'E'.

When host 'C' has an IPv6 packet to send to host 'D' (i.e., another ISATAP host within the site), it uses IPv6-in-IPv4 encapsulation to forward the packet to advertising ISATAP router 'A'. 'A' in turn conveys the packet to 'D' either directly or via 'B' as an intermediary. However, it is not expected that hosts 'C' and 'D' will normally use ISATAP services when communicating with each other within the site. Instead, they will continue to use legacy IPv4 services until a fully-qualified IPv6 intra-site service becomes available.

3.4. Loop Avoidance

In sites that provide IPv6 services through ISATAP with SLAAC as described in this section, advertising ISATAP routers must take

operational precautions to avoid routing loops. For example, with reference to Figure 1 an IPv6 packet that enters the site via advertising ISATAP router 'A' must not be allowed to exit the site via advertising ISATAP router 'B' based on an invalid SLAAC address.

As a simple mitigation, each advertising ISATAP router should drop any packets coming from the IPv6 Internet that would be forwarded back to the Internet via another advertising router. Additionally, each advertising ISATAP router should drop any encapsulated packets received from another advertising router that would be forwarded to the IPv6 Internet. (Note that IPv6 packets with link-local addresses are excluded from these checks, since they cannot be forwarded by an IPv6 router and may be necessary for router-to-router coordinations.) This corresponds to the mitigation documented in Section 3.2.3 of [I-D.ietf-v6ops-tunnel-loops], but other mitigations such as the tunnel endpoint verification checks listed in Section 3.1 of that document can also be employed.

Again with reference to Figure 1, when 'A' receives a packet coming from the IPv6 Internet with destination address 2001:db8:1::5efe:192.0.2.2, it drops the packet since the IPv4 address 192.0.2.2 corresponds to advertising ISATAP router 'B'. Similarly, when 'B' receives a packet coming from the tunnel with an IPv6 destination address that would cause the packet to be forwarded back out to the IPv6 Internet and with an IPv4 source address 192.0.2.1, it drops the packet since 192.0.2.1 corresponds to advertising ISATAP router 'A'.

4. DHCPv6 Services

Whether or not advertising ISATAP routers make basic IPv6 services available using SLAAC, they can also provide fully-qualified IPv6 services to ISATAP clients (i.e., both hosts and non-advertising ISATAP routers) using the Dynamic Host Configuration Protocol for IPv6 (DHCPv6). Any addresses/prefixes obtained via DHCPv6 are distinct from any IPv6 prefixes assigned to the ISATAP interface for SLAAC purposes, however. In this way, DHCPv6 addresses/prefixes are reached by viewing the ISATAP tunnel interface as a "transit" rather than viewing it as an ordinary IPv6 host interface.

ISATAP nodes employ the source address verification checks specified in Section 7.3 of [RFC5214] as a prerequisite for decapsulation of packets received on an ISATAP interface. In order to accommodate direct communications with hosts and non-advertising ISATAP routers that use DHCPv6, ISATAP nodes that support route optimization must employ an additional source address verification check. Namely, the node also considers the outer IPv4 source address correct for the inner IPv6 source address if:

* a forwarding table entry 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 following sections discuss operational considerations for enabling ISATAP DHCPv6 services within predominantly IPv4 sites.

4.1. ISATAP Router Behavior

Advertising ISATAP routers that support DHCPv6 services send RA messages in response to RS messages received on an advertising ISATAP interface. Advertising ISATAP routers also configure either a DHCPv6 relay or server function to service DHCPv6 requests received from other ISATAP nodes.

In many use case scenarios (e.g., small enterprise networks, 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, highly mobile MANETs, etc.), this might be impractical dues to scaling issues. When a proactive dynamic routing protocol cannot be used, 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 also acquire IPv6 prefixes, e.g., through the use of DHCPv6 Prefix Delegation [RFC3633] via an advertising router in the same fashion as described for host-based DHCPv6 stateful address autoconfiguration in Section 4.2. The advertising router in turn maintains IPv6 forwarding table entries that list the IPv4 address of the non-advertising router as the link-layer address of the next hop toward the delegated IPv6 prefixes.

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 ISATAP nodes on the link.

4.2. ISATAP Host Behavior

ISATAP hosts resolve the PRL and send RS messages to obtain RA messages from an advertising ISATAP router. Whether or not IPv6 prefixes for SLAAC are advertised, the host can acquire IPv6 addresses, e.g., through the use of DHCPv6 stateful address autoconfiguration [RFC3315]. 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.

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 IPv6 forwarding table entries that list the IPv4 address of the host as the link-layer address of the delegated IPv6 addresses.

4.3. Reference Operational Scenario

Figure 2 depicts a reference ISATAP network topology that uses DHCPv6. 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:

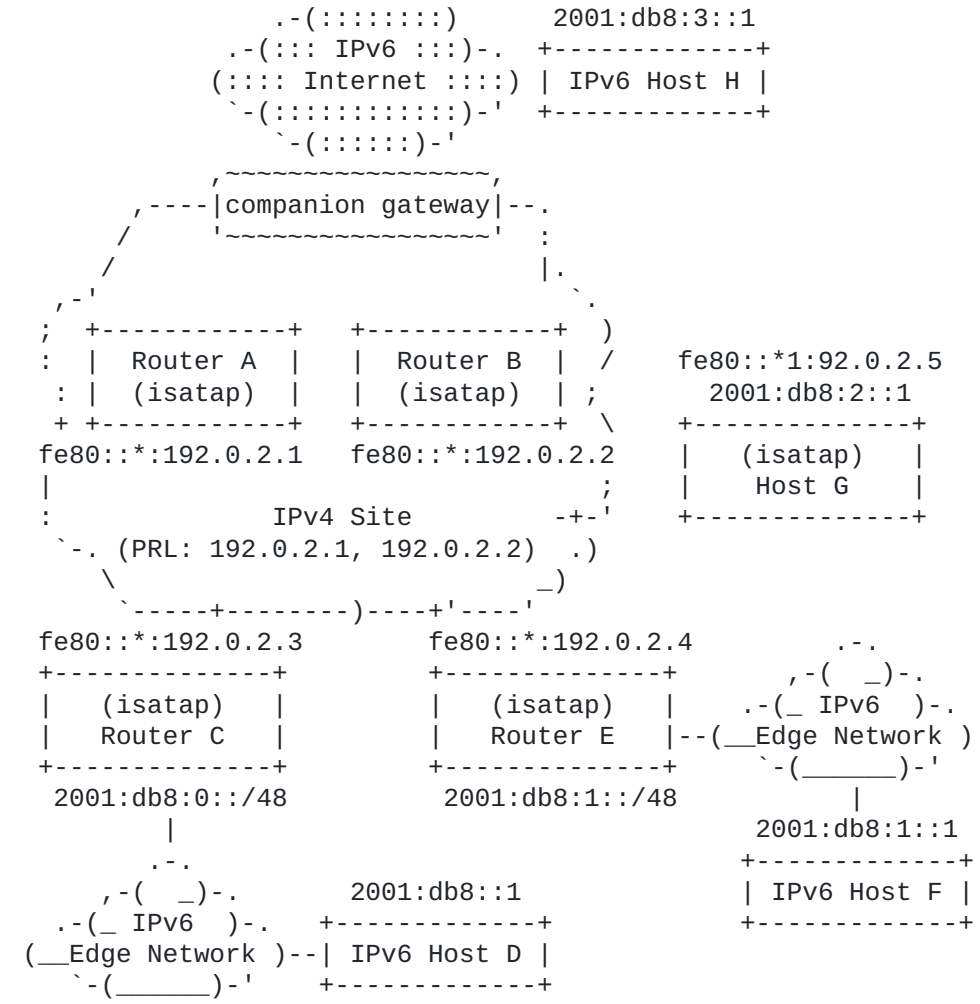

(* == "5efe")

In Figure 2, advertising ISATAP routers 'A' and 'B' within the IPv4 site connect to the IPv6 Internet via a companion gateway. (Note that the routers may instead connect to the IPv6 Internet directly as shown in Figure 1.) Advertising ISATAP router 'A' configures a provider network IPv4 interface with address 192.0.2.1 and arranges to add the address to the provider network PRL. 'A' next configures an advertising ISATAP router interface with link-local IPv6 address fe80::5efe:192.0.2.1 over the IPv4 interface. In the same fashion, advertising ISATAP router 'B' configures the IPv4 interface address 192.0.2.2, adds the address to the PRL, then configures the IPv6 ISATAP interface link-local address fe80::5efe:192.0.2.2.

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 192.0.2.3, but it does not add the IPv4 address to the site's PRL. 'C' next configures a non-advertising ISATAP router interface with link-local address fe80::5efe:192.0.2.3, then receives the IPv6 prefix 2001:db8::/48 through a DHCPv6 prefix delegation exchange via one of 'A' or 'B'. 'C' 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::1.

Non-advertising ISATAP router 'E' connects to the site, configures its ISATAP interface, receives a DHCPv6 prefix delegation, and engages in the IPv6 routing protocol the same as for 'C'. In particular, 'E' configures the IPv4 address 192.0.2.4, the ISATAP link-local address fe80::5efe:192.0.2.4, and the delegated IPv6 prefix 2001:db8:1::/48. 'E' finally 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 192.0.2.5, and also configures an ISATAP host interface with link-local address fe80::5efe:192.0.2.5 over the IPv4 interface. 'G' next performs an RS/RA exchange with 'B' to configure default IPv6 route with next-hop address fe80::5efe:192.0.2.2, then receives the IPv6 address 2001:db8:2::1 from a DHCPv6 address configuration exchange via 'B'. When 'G' receives the IPv6 address, it 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::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::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., a router that advertises "default"), 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.

4.4. Loop Avoidance

In a purely DHCPv6-based ISATAP deployment, no non-link-local IPv6 prefixes are assigned to ISATAP router interfaces. Therefore, 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 [I-D.ietf-v6ops-tunnel-loops].

Any routing loops introduced in the DHCPv6 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. Scaling Considerations

Figure 1 and Figure 2 depict ISATAP network topologies with only two advertising ISATAP routers within the site. In order to support larger numbers of ISATAP nodes, the site can deploy more advertising ISATAP routers to support load balancing and generally shortest-path routing.

Such an arrangement requires that the advertising ISATAP routers participate in an IPv6 routing protocol instance so that IPv6 addresses/prefixes can be mapped to the correct ISATAP router. The routing protocol instance can be configured as either a full mesh topology involving all advertising ISATAP routers, or as a partial mesh topology with each advertising ISATAP router associating with one or more companion gateways. Each such companion gateway would in turn participate in a full mesh between all companion gateways.

6. On-Demand Dynamic Routing

With respect to the reference operational scenarios depicted in Figure 2, 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 ISATAP redirection are specified in [I-D.templin-aero].

7. Site Partitioning Considerations

In common practice, site administrators often deploy packet filtering devices of various forms in order to divide the site into separate partitions. These devices may prevent IPv6-in-IPv4 encapsulated packets from traversing partition boundaries.

In order to avoid communication failures that may result from filtering, ISATAP clients (i.e., hosts and non-advertising routers)

should only enable the service after an initial reachability exchange with an advertising ISATAP router (e.g., in an initial RS/RA exchange). ISATAP client to client communications should therefore also only be used when the path between the clients is first tested in an initial reachability exchange.

8. Site Renumbering Considerations

Advertising ISATAP routers distribute IPv6 prefixes to ISATAP nodes within the site via DHCPv6 and/or SLAAC. If the site subsequently reconnects to a different ISP, however, the site must renumber to use addresses derived from the new IPv6 prefixes [RFC1900][RFC4192] [RFC5887].

For basic IPv6 services provided by SLAAC, site renumbering in the event of a change in an ISP-served IPv6 prefix entails a simple renumbering of IPv6 addresses and/or prefixes that are assigned to the ISATAP interfaces of hosts within the site. In some cases, filtering rules (e.g., within site border firewall filtering tables) may also require renumbering, but this operation can be automated and limited to only one or a few administrative "touch points". In order to renumber the ISATAP interfaces of hosts within the site using SLAAC, advertising ISATAP routers need only schedule the services offered by the old ISP for deprecation while beginning to advertise the IPv6 prefixes provided by the new ISP. ISATAP host interface address lifetimes will eventually expire, and the host will renumber its interfaces with addresses derived from the new prefixes.

For fully-qualified IPv6 services provided by DHCPv6, site renumbering in the event of a change in an ISP-served IPv6 prefix further entails locating and rewriting all IPv6 addresses in naming services, databases, configuration files, packet filtering rules, documentation, etc. If the site has published the IPv6 addresses of any site- internal nodes within the public Internet DNS system, then the corresponding resource records will also need to be updated during the renumbering operation. This can be accomplished via secure dynamic updates to the DNS.

9. Path MTU Considerations

IPv6-in-IPv4 encapsulation overhead effectively reduces the size of IPv6 packets that can traverse the tunnel in relation to the actual Maximum Transmission Unit (MTU) of the underlying IPv4 network path between the encapsulator and decapsulator. Two methods for accommodating IPv6 path MTU discovery over IPv6-in-IPv4 tunnels (i.e., the static and dynamic methods) are documented in Section 3.2 of [RFC4213].

The static method places a "safe" upper bound on the size of IPv6 packets permitted to enter the tunnel, however the method can be overly conservative when larger IPv4 path MTUs are available. The dynamic method can accommodate much larger IPv6 packet sizes in some cases, but can fail silently if the underlying IPv4 network path does not return the necessary error messages.

This document notes that sites that include well-managed IPv4 links, routers and other network middleboxes are candidates for use of the

dynamic MTU determination method, which may provide for a better operational IPv6 experience in the presence of IPv6-in-IPv4 tunnels.

10. Alternative Approaches

[RFC4554] proposes a use of VLANs for IPv4-IPv6 coexistence in enterprise networks. The ISATAP approach provides a more flexible and broadly-applicable alternative, and with fewer administrative touch points.

The tunnel broker service [RFC3053] uses point-to-point tunnels that require end users to establish an explicit administrative configuration of the tunnel far end, which may be outside of the administrative boundaries of the site.

6to4 [RFC3056] and Teredo [RFC4380] provide "last resort" unmanaged automatic tunneling services when no other means for IPv6 connectivity is available. These services are given lower priority when the ISATAP managed service and/or native IPv6 services are enabled.

IRON [RFC6179], RANGER [RFC5720], VET [RFC5558] and SEAL [RFC5320] are a tribute to those in all walks of life who serve with dignity and honor for the benefit of others.

11. IANA Considerations

This document has no IANA considerations.

12. 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 [I-D.ietf-v6ops-tunnel-loops].

13. Acknowledgments

The following are acknowledged for their insights that helped shape this work: Fred Baker, Brian Carpenter, Thomas Henderson, Philip Homburg, Lee Howard, Joel Jaeggli, Gabi Nakibly, Hemant Singh, Mark Smith, Ole Troan, Gunter Van de Velde, ...

14. References

14.1. Normative References

[RFC5214]	Templin, F., Gleeson, T. and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", RFC 5214, March 2008.
[RFC1918]	Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G. and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
[RFC4861]	Narten, T., Nordmark, E., Simpson, W. and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.
[RFC4213]	Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, October 2005.

[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.

14.2. Informative References

[RFC6139]	Russert, S., Fleischman, E. and F. Templin, "Routing and Addressing in Networks with Global Enterprise Recursion (RANGER) Scenarios", RFC 6139, February 2011.
[RFC1900]	Carpenter, B. and Y. Rekhter, "Renumbering Needs Work", RFC 1900, February 1996.
[RFC4192]	Baker, F., Lear, E. and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day", RFC 4192, September 2005.
[RFC5887]	Carpenter, B., Atkinson, R. and H. Flinck, "Renumbering Still Needs Work", RFC 5887, May 2010.
[RFC1687]	Fleischman, E., "A Large Corporate User's View of IPng", RFC 1687, August 1994.
[RFC5969]	Townsley, W. and O. Troan, "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification", RFC 5969, August 2010.
[RFC2491]	Armitage, G., Schulter, P., Jork, M. and G. Harter, "IPv6 over Non-Broadcast Multiple Access (NBMA) networks", RFC 2491, January 1999.
[RFC2529]	Carpenter, B. and C. Jung, "Transmission of IPv6 over IPv4 Domains without Explicit Tunnels", RFC 2529, March 1999.
[RFC4554]	Chown, T., "Use of VLANs for IPv4-IPv6 Coexistence in Enterprise Networks", RFC 4554, June 2006.
[RFC3053]	Durand, A., Fasano, P., Guardini, I. and D. Lento, "IPv6 Tunnel Broker", RFC 3053, January 2001.
[RFC3056]	Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, February 2001.
[RFC4380]	Huitema, C., "Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)", RFC 4380, February 2006.
[RFC5320]	Templin, F., "The Subnetwork Encapsulation and Adaptation Layer (SEAL)", RFC 5320, February 2010.
[RFC5558]	Templin, F., "Virtual Enterprise Traversal (VET)", RFC 5558, February 2010.
[RFC5720]	Templin, F., "Routing and Addressing in Networks with Global Enterprise Recursion (RANGER)", RFC 5720, February 2010.
[RFC6179]	Templin, F., "The Internet Routing Overlay Network (IRON)", RFC 6179, March 2011.
[I-D.ietf-v6ops-tunnel-loops]	Nakibly, G and F Templin, "Routing Loop Attack using IPv6 Automatic Tunnels: Problem Statement and Proposed Mitigations", Internet-Draft draft-ietf-v6ops-tunnel-loops-07, May 2011.
[I-D.templin-aero]	Templin, F, "Asymmetric Extended Route Optimization (AERO)", Internet-Draft draft-templin-aero-04, October 2011.

Author's Address

Fred L. Templin Templin Boeing Research & Technology P.O. Box 3707
MC 7L-49 Seattle, WA 98124 USA EMail: fltemplin@acm.org