     Operational Guidance for IPv6 Deployment in IPv4 Sites using ISATAP
                     draft-templin-v6ops-isops-15.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. 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

publication of this document. Please review these documents
carefully, as they describe your rights and restrictions with respect
to this document. Code Components extracted from this document must
include Simplified BSD License text as described in Section 4.e of
the Trust Legal Provisions and are provided without warranty as
described in the Simplified BSD License.

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
   discussed 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. The
   terminology of ISATAP (see: [RFC5214], Section 3) applies also to
   this document.

2. **Enabling IPv6 Services using ISATAP**

   Existing sites within the Internet will soon need to enable IPv6
   services. 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

Internet Service Provider (ISP) unto themselves. Smaller sites that
wish to enable IPv6 can arrange to obtain public IPv6 prefixes from
an ISP, where the prefixes may be either purely native or the near-
native prefixes offered by 6rd [RFC5969]. Alternatively, the site
can obtain prefixes independently of an ISP e.g., via a tunnel broker
[RFC3053], by using one of its public IPv4 addresses to form a 6to4
prefix [RFC3056][RFC3068], etc. (Note however that experience shows
that the 6to4 method has some problems in current deployments that
can lead to connectivity failures [RFC6343].) In any 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 encapsulation format can further
use Differentiated Service (DS) [RFC2983] and Explicit Congestion
Notification (ECN) [RFC3168] mapping between the inner and outer IP
headers to ensure expected per-hop behavior within well-managed
sites.

The ISATAP service is based on two node types known as advertising
ISATAP routers and ISATAP hosts. (A third node type known as non-
advertising ISATAP routers is defined in [I-D.templin-isupdate] but
out of scope for this document.) Each node may further have multiple
ISATAP interfaces (i.e., one interface for each site), and may act as
an advertising ISATAP router on some of those interfaces and a simple
ISATAP host on others. Hence, the node type is considered on a per-
interface basis.

Advertising ISATAP routers configure their ISATAP interfaces as
advertising router interfaces (see: [RFC4861], Section 6.2.2).
ISATAP hosts configure their ISATAP interfaces as simple host
interfaces and also coordinate their autoconfiguration operations
with advertising ISATAP routers. In this sense, advertising ISATAP
routers are "servers" while ISATAP hosts are "clients" in the service
model.

Advertising ISATAP routers arrange to add their IPv4 address to the
site's Potential Router List (PRL) so that ISATAP clients can
discover them, as discussed in Sections 8.3.2 and 9 of [RFC5214].
Alternatively, site administrators could include IPv4 anycast
addresses in the PRL and assign each such address to multiple
advertising ISATAP routers. In that case, IPv4 routing within the
site would direct the ISATAP client to the nearest advertising ISATAP
router.

After the PRL is published, ISATAP clients within the site can
automatically perform unicast IPv6 Neighbor Discovery Router

Solicitation (RS) / Router Advertisement (RA) exchanges with
advertising ISATAP routers using IPv6-in-IPv4 encapsulation
[RFC4861][RFC5214]. In the exchange, the IPv4 source address of the
RS and the destination address of the RA are an IPv4 address of the
client, while the IPv4 destination address of the RS and the source
address of the RA are an IPv4 address of the server found in the PRL.
Similarly, the IPv6 source address of the RS is a link-local ISATAP
address that embeds the client's IPv4 address, while the source
address of the RA is a link-local ISATAP address that embeds the
server's IPv4 address. (The destination addresses of the RS and RA
may be either the neighbor's link-local ISATAP address or a link-
scoped multicast address depending on the implementation.)

Following router discovery, ISATAP clients can configure and assign
IPv6 addresses and/or prefixes using Stateless Address
AutoConfiguration (SLAAC) [RFC4862][RFC5214]. While out of scope for
this document, use of the Dynamic Host Configuration Protocol for
IPv6 (DHCPv6) [RFC3315] is also possible when necessary updates to
the ISATAP base specification are implemented [I-D.templin-isupdate].

3. SLAAC Services

Predominantly IPv4 sites can enable SLAAC services for ISATAP clients
that need to communicate with IPv6 correspondents. SLAAC services
are enabled using either the "shared" or "individual" prefix model.
In the shared prefix model, all advertising ISATAP routers advertise
a common prefix (e.g., 2001:db8::/64) to ISATAP clients within the
site. In the individual prefix model, advertising ISATAP router
advertise individual prefixes (e.g., 2001:db8:0:1::/64, 2001:db8:0:
2::/64, 2001:db8:0:3::/64, etc.) to ISATAP clients within partition
of the site. Note that combinations of the shared and individual
prefix models are also possible, in which some of the site's ISATAP
routers advertise shared prefixes and others advertise individual
prefixes.

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

3.1. Advertising 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 in Prefix Information
Options (PIOs) with the A flag set to 1[RFC4861]. When there are
multiple advertising ISATAP routers, the routers can advertise a
shared IPv6 prefix or individual IPv6 prefixes.

3.2. **ISATAP Host Behavior**

   ISATAP hosts resolve the PRL and send RS messages to obtain RA
   messages from an advertising ISATAP router. When the host receives
   RA messages, it uses SLAAC to configure IPv6 addresses from any
   advertised prefixes with the A flag set to 1 as specified in
   [RFC4862][RFC5214] then assigns the addresses to the ISATAP
   interface. The host also assigns any of the advertised prefixes with
   the L flag set to 1 to the ISATAP interface. (Note that the IPv6
   link-local prefix fe80::/64 is always considered on-link on an ISATAP
   interface.)

3.3. **Reference Operational Scenario - Shared Prefix Model**

   Figure 1 depicts an example ISATAP network topology for allowing
   hosts within a predominantly IPv4 site to configure ISATAP services
   using SLAAC with the shared prefix model. The example 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. In this model, routers 'A' and 'B' both
   advertise the same (shared) IPv6 prefix 2001:db8::/64 into the IPv6
   routing system, and also advertise the prefix to ISATAP clients
   within the site for SLAAC purposes.

```
                       .-(:::::::::)       2001:db8:1::1
                    .-(::: IPv6 :::)-.  +-------------+
                   (::::: Internet :::) | IPv6 Host E |
                    `-(:::::::::::::)-'  +-------------+
                      `-(:::::::)-'
                 ,~~~~~~~~~~~~~~~~,
             ,----|companion gateway|--.
            /      '~~~~~~~~~~~~~~~~'   :
           /                           |.
        ,-'                             `.
      ;   +------------+   +------------+  )
      :   |  Router A  |   |  Router B  | /
       :  |  (isatap)  |   |  (isatap)  | :
       :  | 192.0.2.1  |   | 192.0.2.1  | ;
       + +------------+    +------------+  \
       fe80::*:192.0.2.1  fe80::*:192.0.2.1
       | 2001:db8::/64       2001:db8::/64   |
       |                                  ;
       :              IPv4 Site        -+-'
        `-.          (PRL: 192.0.2.1)      .)
          \                               _)
           `-----+--------)----+'----'
        fe80::*:192.0.2.18          fe80::*:192.0.2.34
      2001:db8::*:192.0.2.18     2001:db8::*:192.0.2.34
        +--------------+             +--------------+
        |   (isatap)   |             |   (isatap)   |
        |    Host C    |             |    Host D    |
        +--------------+             +--------------+
```

   (* == "5efe")

    Figure 1: Example ISATAP Network Topology using Shared Prefix Model

   With reference to Figure 1, advertising ISATAP routers 'A' and 'B'
   within the IPv4 site connect to the IPv6 Internet either directly or
   via a companion gateway. The routers advertise the shared prefix
   2001:db8::/64 into the IPv6 Internet routing system either as a
   singleton /64 or as part of a shorter aggregated IPv6 prefix if the
   routing system will not accept prefixes as long as a /64. For the
   purpose of this example, we also assume that the IPv4 site is
   configured within multiple IPv4 subnets - each with an IPv4 prefix
   length of /28.

   Advertising ISATAP routers 'A' and 'B' both configure the IPv4
   anycast address 192.0.2.1 on a site-interior IPv4 interface, then
   configure an advertising ISATAP router interface for the site 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

site's PRL. '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 also use individual IPv4 unicast addresses instead
of, or in addition to, a shared IPv4 anycast address. In that case,
the PRL will contain multiple IPv4 addresses of advertising routers -
some of which may be anycast and others unicast.)

ISATAP host 'C' connects to the site via an IPv4 interface with
address 192.0.2.18/28, and also configures an ISATAP host interface
with link-local ISATAP address fe80::5efe:192.0.2.18 over the IPv4
interface. 'C' next resolves the PRL, and sends an RS message to the
IPv4 address 192.0.2.1, where IPv4 routing will direct it to the
closest of either 'A' or 'B'. Assuming 'A' is closest, 'C' receives
an RA from 'A' then 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::/64 advertised in the PIO. If the A flag
is set in the PIO, 'C' uses SLAAC to automatically configure the IPv6
address 2001:db8::5efe:192.0.2.18 (i.e., an address with an ISATAP
interface identifier) and assigns it to the ISATAP interface. If the
L flag is set, 'C' also assigns the prefix 2001:db8::/64 to the
ISATAP interface, and the IPv6 address becomes a true ISATAP address.

In the same fashion, ISATAP host 'D' configures its IPv4 interface
with address 192.0.2.34/28 and configures its ISATAP interface with
link-local ISATAP address fe80::5efe:192.0.2.34. 'D' next performs
an RS/RA exchange that is serviced by 'B', then uses SLAAC to
autoconfigure the address 2001:db8::5efe:192.0.2.34 and a default
IPv6 route with next-hop address fe80::5efe:192.0.2.1. 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:db8:1::1.

Following this autoconfiguration, when host 'C' inside the site has
an IPv6 packet to send to host 'E' outside the site, it prepares the
packet with source address 2001:db8::5efe:192.0.2.18 and destination
address 2001:db8:1::1. 'C' then uses IPv6-in-IPv4 encapsulation to
forward the packet to the IPv4 address 192.0.2.1 which will be
directed to 'A' based on IPv4 routing. 'A' in turn decapsulates the
packet and forwards it into the public IPv6 Internet where it will be
conveyed to 'E' via normal IPv6 routing. 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 'E' outside the site sends IPv6 packets to ISATAP host 'C'
inside the site, the IPv6 routing system may direct the packet to
either of 'A' or 'B'. If the site is not partitioned internally, the
router that receives the packet can use ISATAP to statelessly forward

the packet directly to 'C'. If the site may be partitioned
internally, however, the packet must first be forwarded to 'C's
serving router based on IPv6 routing information. This implies that,
in a partitioned site, the advertising ISATAP routers must connect
within a full or partial mesh of IPv6 links, and must either run a
dynamic IPv6 routing protocol or configure static routes so that
incoming IPv6 packets can be forwarded to the correct serving router.

In this example, 'A' can configure the IPv6 route 2001:db8::5efe:
192.0.2.32/124 with the IPv6 address of the next hop toward 'B' in
the mesh network as the next hop, and 'B' can configure the IPv6
route 2001:db8::5efe:192.0.2.16/124 with the IPv6 address of the next
hop toward 'A' as the next hop. (Notice that the /124 prefixes
properly cover the /28 prefix of the IPv4 address that is embedded
within the IPv6 address.) In that case, when 'A' receives a packet
from the IPv6 Internet with destination address 2001:db8::5efe:
192.0.2.34, it first forwards the packet toward 'B' over an IPv6 mesh
link. 'B' in turn uses ISATAP to forward the packet into the site,
where IPv4 routing will direct it to 'D'. In the same fashion, when
'B' receives a packet from the IPv6 Internet with destination address
2001:db8::5efe:192.0.2.18, it first forwards the packet toward 'A'
over an IPv6 mesh link. 'A' then uses ISATAP to forward the packet
into the site, where IPv4 routing will direct it to 'C'.

Finally, when host 'C' inside the site connects to host 'D' inside
the site, it has the option of using the native IPv4 service or the
ISATAP IPv6-in-IPv4 encapsulation service. When there is operational
assurance that IPv4 services between the two hosts are available, the
hosts may be better served to continue to use legacy IPv4 services in
order to avoid encapsulation overhead and to avoid any IPv4
protocol-41 filtering middleboxes that may be in the path. If 'C'
and 'D' may be in different IPv4 network partitions, however, IPv6-
in-IPv4 encapsulation should be used with one or both of routers 'A'
and 'B' serving as intermediate gateways.

3.4. Reference Operational Scenario - Individual Prefix Model

Figure 2 depicts an example ISATAP network topology for allowing
hosts within a predominantly IPv4 site to configure ISATAP services
using SLAAC with the individual prefix model. The example 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. In the figure, ISATAP routers 'A' and 'B'
both advertise different prefixes taken from the aggregated prefix
2001:db8::/48, with 'A' advertising 2001:db8:0:1::/64 and 'B'
advertising 2001:db8:0:2::/64.

```
                       .-(:::::::::)       2001:db8:1::1
                    .-(::: IPv6 :::)-.  +-------------+
                   (::::: Internet ::::) | IPv6 Host E |
                    `-(:::::::::::::)-'  +-------------+
                      `-(:::::::)-'
                  ,~~~~~~~~~~~~~~~~~,
              ,----|companion gateway|--.
             /      '~~~~~~~~~~~~~~~~~'   :
            /                            |.
        ,-'                                `.
       ;   +------------+   +------------+  )
       :   |  Router A  |   |  Router B  |  /
        :  |  (isatap)  |   |  (isatap)  |  :
        :  | 192.0.2.1  |   | 192.0.2.1  |  ;
        + +------------+   +------------+  \
         fe80::*:192.0.2.1   fe80::*:192.0.2.1
         2001:db8:0:1::/64   2001:db8:0:2::/64
         |                               ;
         :              IPv4 Site        -+-'
          `-.        (PRL: 192.0.2.1)      .)
            \                             _)
             `-----+--------)----+'----'
         fe80::*:192.0.2.18         fe80::*:192.0.2.34
        2001:db8:0:1::*:192.0.2.18  2001:db8:0:2::*:192.0.2.34
          +--------------+          +--------------+
          |   (isatap)   |          |   (isatap)   |
          |    Host C    |          |    Host D    |
          +--------------+          +--------------+

     (* == "5efe")
```

       Figure 2: Example ISATAP Network Topology using Individual Prefix
                                   Model

   With reference to Figure 2, advertising ISATAP routers 'A' and 'B'
   within the IPv4 site connect to the IPv6 Internet either directly or
   via a companion gateway. Router 'A' advertises the individual prefix
   2001:db8:0:1::/64 into the IPv6 Internet routing system, and router
   'B' advertises the individual prefix 2001:db8:0:2::/64. The routers
   could instead both advertise a shorter shared prefix such as 2001:
   db8::/48 into the IPv6 routing system, but in that case they would
   need to configure a mesh of IPv6 links between themselves in the same
   fashion as described for the shared prefix model in Section 3.4. For
   the purpose of this example, we also assume that the IPv4 site is
   configured within multiple IPv4 subnets - each with an IPv4 prefix
   length of /28.

   Advertising ISATAP routers 'A' and 'B' both configure the IPv4

anycast address 192.0.2.1 on a site-interior IPv4 interface, then
configure an advertising ISATAP router interface for the site 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
site's PRL. '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 also use individual IPv4 unicast addresses instead
of, or in addition to, a shared IPv4 anycast address. In that case,
the PRL will contain multiple IPv4 addresses of advertising routers -
some of which may be anycast and others unicast.)

ISATAP host 'C' connects to the site via an IPv4 interface with
address 192.0.2.18/28, and also configures an ISATAP host interface
with link-local ISATAP address fe80::5efe:192.0.2.18 over the IPv4
interface. 'C' next resolves the PRL, and sends an RS message to the
IPv4 address 192.0.2.1, where IPv4 routing will direct it to the
closest of either 'A' or 'B'. Assuming 'A' is closest, 'C' receives
an RA from 'A' then 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:db8:0:1:/64 advertised in the PIO. If the A
flag is set in the PIO, 'C' uses SLAAC to automatically configure the
IPv6 address 2001:db8:0:1::5efe:192.0.2.18 (i.e., an address with an
ISATAP interface identifier) and assigns it to the ISATAP interface.
If the L flag is set, 'C' also assigns the prefix 2001:db8:0:1::/64
to the ISATAP interface, and the IPv6 address becomes a true ISATAP
address.

In the same fashion, ISATAP host 'D' configures its IPv4 interface
with address 192.0.2.34/28 and configures its ISATAP interface with
link-local ISATAP address fe80::5efe:192.0.2.34. 'D' next performs
an RS/RA exchange that is serviced by 'B', then uses SLAAC to
autoconfigure the address 2001:db8:0:2::5efe:192.0.2.34 and a default
IPv6 route with next-hop address fe80::5efe:192.0.2.1. 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:db8:1::1.

Following this autoconfiguration, when host 'C' inside the site has
an IPv6 packet to send to host 'E' outside the site, it prepares the
packet with source address 2001:db8::5efe:192.0.2.18 and destination
address 2001:db8:1::1. 'C' then uses IPv6-in-IPv4 encapsulation to
forward the packet to the IPv4 address 192.0.2.1 which will be
directed to 'A' based on IPv4 routing. 'A' in turn decapsulates the
packet and forwards it into the public IPv6 Internet where it will be
conveyed to 'E' via normal IPv6 routing. 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 'E' outside the site sends IPv6 packets to ISATAP host 'C'
   inside the site, the IPv6 routing system will direct the packet to
   'A' since 'A' advertises the individual prefix that matches 'C's
   destination address. 'A' can then use ISATAP to statelessly forward
   the packet directly to 'C'. If 'A' and 'B' both advertise the shared
   shorter prefix 2001:db8::/48 into the IPv6 routing system, however
   packets coming from 'E' may be directed to either 'A' or 'B'. In
   that case, the advertising ISATAP routers must connect within a full
   or partial mesh of IPv6 links the same as for the shared prefix
   model, and must either run a dynamic IPv6 routing protocol or
   configure static routes so that incoming IPv6 packets can be
   forwarded to the correct serving router.

   In this example, 'A' can configure the IPv6 route 2001:db8:0:2::/64
   with the IPv6 address of the next hop toward 'B' in the mesh network
   as the next hop, and 'B' can configure the IPv6 route 2001:db8:
   0.1::/64 with the IPv6 address of the next hop toward 'A' as the next
   hop. Then, when 'A' receives a packet from the IPv6 Internet with
   destination address 2001:db8:0:2::5efe:192.0.2.34, it first forwards
   the packet toward 'B' over an IPv6 mesh link. 'B' in turn uses
   ISATAP to forward the packet into the site, where IPv4 routing will
   direct it to 'D'. In the same fashion, when 'B' receives a packet
   from the IPv6 Internet with destination address 2001:db8:0:1::5efe:
   192.0.2.18, it first forwards the packet toward 'A' over an IPv6 mesh
   link. 'A' then uses ISATAP to forward the packet into the site,
   where IPv4 routing will direct it to 'C'.

   Finally, when host 'C' inside the site connects to host 'D' inside
   the site, it has the option of using the native IPv4 service or the
   ISATAP IPv6-in-IPv4 encapsulation service. When there is operational
   assurance that IPv4 services between the two hosts are available, the
   hosts may be better served to continue to use legacy IPv4 services in
   order to avoid encapsulation overhead and to avoid any IPv4
   protocol-41 filtering middleboxes that may be in the path. If 'C'
   and 'D' may be in different IPv4 network partitions, however, IPv6-
   in-IPv4 encapsulation should be used with one or both of routers 'A'
   and 'B' serving as intermediate gateways.

3.5. SLAAC Site Administration Guidance

   In common practice, firewalls, gateways and packet filtering devices
   of various forms are often deployed in order to divide the site into
   separate partitions. In both the shared and individual prefix models
   described above, the entire site can be represented by the aggregate
   IPv6 prefix assigned to the site, while each site partition can be
   represented by "sliver" IPv6 prefixes taken from the aggregate. In
   order to provide a simple service that does not interact poorly with
   the site topology, site administrators should therefore institute an

address plan to align IPv6 sliver prefixes with IPv4 site partition
boundaries.

For example, in the shared prefix model in Section 3.3, the aggregate
prefix is 2001:db8::/64, and the sliver prefixes are 2001:db8::5efe:
192.0.2.0/124, 2001:db8::5efe:192.0.2.16/124, 2001:db8::5efe:
192.0.2.32/124, etc. In the individual prefix model in Section 3.4,
the aggregate prefix is 2001:db8::/48 and the sliver prefixes are
2001:db8:0:0::/64, 2001:db8:0:1::/64, 2001:db8:0:2::/64, etc.

When individual prefixes are used, site administrators can configure
advertising ISATAP routers to advertise different individual prefixes
to different sets of clients, e.g., based on the client's IPv4 subnet
prefix such that the IPv6 prefixes are congruent with the IPv4
addressing plan. (For example, administrators can configure each
advertising ISATAP router to provide services only to certain sets of
ISATAP clients through inbound IPv6 Access Control List (ACL) entries
that match the IPv4 subnet prefix embedded in the ISATAP interface
identifier of the IPv6 source address). When a shared prefix is
used, site administrators instead configure the ISATAP routers to
advertise the shared prefix to all clients.

Advertising ISATAP routers can advertise prefixes with the (A, L)
flags set to (1,0) so that ISATAP clients will use SLAAC to
autoconfigure IPv6 addresses with ISATAP interface identifiers from
the prefixes and assign them to the receiving ISATAP interface, but
they will not assign the prefix itself to the ISATAP interface. In
that case, the advertising router must assign the sliver prefix for
the site partition to the advertising ISATAP interface. In this way,
the advertising router considers the addresses covered by the sliver
prefix as true ISATAP addresses, but the ISATAP clients themselves do
not. This configuration enables a hub-and-spokes architecture which
in some cases may be augmented by route optimization based on the
receipt of ICMPv6 Redirects.

Site administrators can implement address selection policy rules
[RFC3484] through explicit configurations in each ISATAP client.
Site administrators implement this policy by configuring address
selection policy rules [RFC3484] in each ISATAP client in order to
give preference to IPv4 destination addresses over destination
addresses derived from one of the client's IPv6 sliver prefixes.

For example, site administrators can configure each ISATAP client
associated with a sliver prefix such as 2001:db8::5efe:192.0.2.64/124
to 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
sliver. 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] discovered by the client, e.g., using
Stateless DHCPv6 [RFC3736]. 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.

It should be noted that sliver prefixes longer than /64 cannot be
advertised for SLAAC purposes. Also, sliver prefixes longer than /64
do not allow for interface identifier rewriting by address
translators. These factors may favor the individual prefix model in
some deployment scenarios, while the flexibility afforded by the
shared prefix model may be more desirable in others. Additionally,
if the network is small then the shared prefix model works well. If
the network is large, however, a better alternative may be to deploy
separate ISATAP routers in each region and have each advertise their
own individual prefix.

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.

3.6. Loop Avoidance

In sites that provide IPv6 services through ISATAP with SLAAC as
described in this section, site administrators must take operational
precautions to avoid routing loops. For example, each advertising
ISATAP router should drop any incoming IPv6 packets that would be
forwarded back to itself via another of the site's advertising
routers. Additionally, each advertising ISATAP router should drop
any encapsulated packets received from another advertising router
that would be forwarded back to that same advertising router. This
corresponds to the mitigation documented in Section 3.2.3 of
[RFC6324], but other mitigations specified in that document can also
be employed.

Note that IPv6 packets with link-local ISATAP addresses are exempt
from these checks, since they cannot be forwarded by an IPv6 router
and may be necessary for router-to-router coordinations.

3.7. Interface Identifier Compatibility Considerations

[RFC5214] Section 6.1 specifies the setting of the "u" bit in the
Modified EUI-64 interface identifier format used by ISATAP.
Implementations that comply with the specification set the "u" bit to
1 when the IPv4 address is known to be globally unique, however some
legacy implementations unconditionally set the "u" bit to 0.

Implementations interpret the ISATAP interface identifier only within
the link to which the corresponding ISATAP prefix is assigned, hence
the value of the "u" bit is interpreted only within the context of an
on-link prefix and not within a global context. Implementers are
responsible for ensuring that their products are interoperable,
therefore implementations must make provisions for ensuring "u" bit
compatibility for intra-link communications.

Site administrators should accordingly configure access control list
entries and other literal representations of ISATAP interface
identifiers such that both values of the "u" bit are accepted. For
example, if the site administrator configures an access control list
entry that matches the prefix "fe80::0000:5efe:192.0.2.0/124" they
should also configure a companion list entry that matches the prefix
"fe80::0200:5efe:192.0.2.0/124.

4. Manual Configuration

When no autoconfiguration services are available (e.g., if there are
no advertising ISATAP routers present), site administrators can use
manual configuration to assign IPv6 addresses with ISATAP interface
identifiers to the ISATAP interfaces of clients. Otherwise, site
administrators should avoid manual configurations that would in any
way invalidate the assumptions of the autoconfiguration service. For
example, manually configured addresses may not be automatically
renumbered during a site-wide renumbering event, which could
subsequently result in communication failures.

5. Scaling Considerations

Section 3 depicts ISATAP network topologies with only two advertising
ISATAP routers within the site. In order to support larger numbers
of ISATAP clients (and/or multiple site partitions), 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. Site Renumbering Considerations

   Advertising ISATAP routers distribute IPv6 prefixes to ISATAP clients
   within the site. 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 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 clients 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 clients within the site
   using SLAAC, advertising ISATAP routers need only schedule the
   services offered by the old ISP for deprecation and begin to
   advertise the IPv6 prefixes provided by the new ISP. ISATAP client
   interface address lifetimes will eventually expire, and the host will
   renumber its interfaces with addresses derived from the new prefixes.
   ISATAP clients should also eventually remove any deprecated SLAAC
   prefixes from their address selection policy tables, but this action
   is not time-critical.

   Finally, 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.

7. 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.
The dynamic MTU determination method can potentially also present a
larger MTU to IPv6 correspondents outside of the site, since IPv6
path MTU discovery is considered robust even over the wide area in
the public IPv6 Internet.

8. 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][RFC3068] 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.

6rd [RFC5969] enables a stateless prefix delegation capability based
on IPv4-embedded IPv6 prefixes, whereas ISATAP enables a stateful
prefix delegation capability based on native IPv6 prefixes.

IRON [RFC6179], RANGER [RFC5720], VET [RFC5558] and SEAL [RFC5320]
were developed as the "next-generation" of ISATAP and extend to a
wide variety of use cases [RFC6139]. However, these technologies are
not yet widely implemented or deployed.

9. IANA Considerations

This document has no IANA considerations.

10. 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].

11. Acknowledgments

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

12. References

12.1. Normative References

   [RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and
              E. Lear, "Address Allocation for Private Internets",
              BCP 5, RFC 1918, February 1996.

   [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.

   [RFC3736] Droms, R., "Stateless Dynamic Host Configuration Protocol
              (DHCP) Service for IPv6", RFC 3736, April 2004.

   [RFC4213] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms
              for IPv6 Hosts and Routers", RFC 4213, October 2005.

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

   [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless
              Address Autoconfiguration", RFC 4862, September 2007.

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

12.2. 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-isupdate]
              Templin, F., "ISATAP Updates", draft-templin-isupdate-01
              (work in progress), July 2011.

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

   [RFC1900] Carpenter, B. and Y. Rekhter, "Renumbering Needs Work",
              RFC 1900, February 1996.

   [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.

   [RFC2983] Black, D., "Differentiated Services and Tunnels",
              RFC 2983, October 2000.

   [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.

   [RFC3068] Huitema, C., "An Anycast Prefix for 6to4 Relay Routers",
              RFC 3068, June 2001.

   [RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition
              of Explicit Congestion Notification (ECN) to IP",
              RFC 3168, September 2001.

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

   [RFC4192] Baker, F., Lear, E., and R. Droms, "Procedures for
              Renumbering an IPv6 Network without a Flag Day", RFC 4192,
              September 2005.

   [RFC4380] Huitema, C., "Teredo: Tunneling IPv6 over UDP through
              Network Address Translations (NATs)", RFC 4380,
              February 2006.

   [RFC4554] Chown, T., "Use of VLANs for IPv4-IPv6 Coexistence in
              Enterprise Networks", RFC 4554, June 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.

   [RFC5887] Carpenter, B., Atkinson, R., and H. Flinck, "Renumbering
              Still Needs Work", RFC 5887, May 2010.

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

   [RFC6139] Russert, S., Fleischman, E., and F. Templin, "Routing and
              Addressing in Networks with Global Enterprise Recursion
              (RANGER) Scenarios", RFC 6139, February 2011.

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

   [RFC6179] Templin, F., "The Internet Routing Overlay Network
              (IRON)", RFC 6179, March 2011.

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

   [RFC6343] Carpenter, B., "Advisory Guidelines for 6to4 Deployment",
              RFC 6343, 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