Network Working Group   C. Huitema
Request for Comments: 3750   Microsoft
Category: Informational   R. Austein
                                                                   ISC
                                                           S. Satapati
                                                    Cisco Systems, Inc.
                                                         R. van der Pol
                                                            NLnet Labs
                                                            April 2004

                 Unmanaged Networks IPv6 Transition Scenarios

Status of this Memo

Copyright Notice

Abstract

   This document defines the scenarios in which IPv6 transition
   mechanisms are to be used in unmanaged networks. In order to
   evaluate the suitability of these mechanisms, we need to define the
   scenarios in which these mechanisms have to be used. One specific
   scope is the "unmanaged network", which typically corresponds to a
   home or small office network. The scenarios are specific to a single
   subnet, and are defined in terms of IP connectivity supported by the
   gateway and the Internet Service Provider (ISP). We first examine
   the generic requirements of four classes of applications: local,
   client, peer to peer and server. Then, for each scenario, we infer
   transition requirements by analyzing the needs for smooth migration
   of applications from IPv4 to IPv6.

Table of Contents

1. Introduction

   In order to evaluate the suitability of transition mechanisms from
   IPv4 [RFC791] to IPv6 [RFC2460], we need to define the environment or
   scope in which these mechanisms have to be used. One specific scope
   is the "unmanaged networks", which typically correspond to home
   networks or small office networks.

   This document studies the requirement posed by various transition
   scenarios, and is organized in to four main sections. Section 2
   defines the topology that we are considering. Section 3 presents the
   four classes of applications that we consider for unmanaged networks:
   local applications, client applications, peer-to-peer applications,
   and server applications. Section 4 studies the requirements of these
   four classes of applications. Section 5 analyses how these
   requirements translate into four configurations that we expect to
   encounter during IPv6 deployment: gateways which do not provide IPv6,
   dual-stack gateways connected to dual-stack ISPs, dual-stack gateways
   connected to IPv4-only ISPs, and IPv6-capable gateways connected to
   IPv6-only ISPs. While these four configurations are certainly not an
   exhaustive list of possible configurations, we believe that they
   represent the common cases for unmanaged networks.

2. Topology

   The typical unmanaged network is composed of a single subnet,
   connected to the Internet through a single Internet Service Provider
   (ISP) connection. Several hosts may be connected to the subnet:

```
        +------+
        | Host +--+
        +------+  |
                  |
        +------+  |
        | Host +--+                          +--------------
        +------+  |                          |
             :    |                  +-----+ |
             :  +---------+          |     | |
             +--+ Gateway +------| ISP | Internet
             :  +---------+          |     | |
             :                    +-----+ |
        +------+  |                          |
        | Host +--+                          +--------------
        +------+  |
                  |
        +------+  |
        | Host +--+
        +------+
```

Between the subnet and the ISP access link is a gateway, which may or
may not perform NAT and firewall functions. When the gateway
performs NAT functions [RFC3022], it generally allocates private IPv4
addresses to the local hosts [RFC1918]. A key point of this
configuration is that the gateway is typically not "managed". In
most cases, it is a simple "appliance" that incorporates some static
policies. There are many cases in which the gateway is procured and
configured by the ISP.

Note that there are also some cases in which we find two gateways
back to back, one managed by the ISP and the other added by the owner
of the unmanaged network. They are not covered in this memo because
most of them either require some management, or the gateway added by
the user can function as an L2 switch.

The access link between the unmanaged network and the ISP might be
either a static, permanent connection or a dynamic connection such as
a dial-up or ISDN line.

In a degenerate case, an unmanaged network might consist of a single
host, directly connected to an ISP.

There are some cases in which the "gateway" is replaced by a layer-2
bridge. In such deployments, the hosts have direct access to the ISP
service. In order to avoid lengthy developments, we will treat these
cases as if the gateway was not present, i.e., as if each host was
connected directly to the ISP.

Our definition of unmanaged networks explicitly exclude networks
composed of multiple subnets. We will readily admit that some home
networks and some small business networks contain multiple subnets,
but in the current state of the technology, these multiple subnet
networks are not "unmanaged": some competent administrator has to
explicitly configure the routers. We will thus concentrate on single
subnet networks, where no such competent operator is expected.

3. Applications

   Users may use or wish to use the unmanaged network services in four
   types of applications: local, client, servers and peer-to-peers.
   These applications may or may not run easily on today's networks
   (some do, some don't).

3.1. Local Applications

   "Local applications" are only meant to involve the hosts that are
   part of the unmanaged network. Typical examples would be file
   sharing or printer sharing.

   Local applications work effectively in IPv4 unmanaged networks, even
   when the gateway performs NAT or firewall functions. In fact,
   firewall services at the gateway are often deemed desirable, as they
   isolate the local applications from interference by Internet users.

3.2. Client Applications

   "Client applications" are those that involve a client on the
   unmanaged network and a server at a remote location. Typical
   examples would be accessing a web server from a client inside the
   unmanaged network, or reading and sending e-mail with the help of a
   server outside the unmanaged network.

   Client applications tend to work correctly in IPv4 unmanaged
   networks, even when the gateway performs NAT or firewall functions:
   these translation and firewall functions are designed precisely to
   enable client applications.

3.3. Peer-to-Peer Applications

   There are really two kinds of "peer-to-peer" applications: ones which
   only involve hosts on the unmanaged network, and ones which involve
   both one or more hosts on the unmanaged network and one or more hosts
   outside the unmanaged network. We will only consider the latter kind
   of peer-to-peer applications, since the former can be considered a
   subset of the kind of local applications discussed in section 3.1.

   Peer-to-peer applications often don't work well in unmanaged IPv4
   networks. Application developers often have to enlist the help of a
   "relay server", in effect restructuring the peer-to-peer connection
   into a pair of back-to-back client/server connections.

3.4. Server Applications

   "Server applications" involve running a server in the unmanaged
   network for use by other parties outside the network. Typical
   examples would be running a web server or an e-mail server on one of
   the hosts inside the unmanaged network.

   Deploying these servers in most unmanaged IPv4 networks requires some
   special programming of the NAT or firewall [RFC2993], and is more
   complex when the NAT only publishes a small number of global IP

addresses and relies on "port translation". In the common case in
which the NAT manages exactly one global IP address and relies on
"port translation", a given external port can only be used by one
internal server.

Deploying servers usually requires providing each server with a
stable DNS name, and associating a global IPv4 address with that
name, whether the address be that of the server itself or that of the
router acting as a firewall or NAT. Since updating DNS is a
management task, it falls somewhat outside the scope of an unmanaged
network. On the other hand, it is also possible to use out-of-band
techniques (such as cut-and-paste into an instant message system) to
pass around the address of the target server.

4. Application Requirements of an IPv6 Unmanaged Network

   As we transition to IPv6, we must meet the requirements of the
   various applications, which we can summarize in the following way:
   applications that worked well with IPv4 should continue working well
   during the transition; it should be possible to use IPv6 to deploy
   new applications that are currently hard to deploy in IPv4 networks;
   and the deployment of these IPv6 applications should be simple and
   easy to manage, but the solutions should also be robust and secure.

   The application requirements for IPv6 Unmanaged Networks fall into
   three general categories: connectivity, naming, and security.
   Connectivity issues include the provision of IPv6 addresses and their
   quality: do hosts need global addresses, should these addresses be
   stable or, more precisely, what should the expected lifetimes of
   these addresses be? Naming issues include the management of names
   for the hosts: do hosts need DNS names, and is inverse name
   resolution [DNSINADDR] a requirement? Security issues include
   possible restriction to connectivity, privacy concerns and, generally
   speaking, the security of the applications.

4.1. Requirements of Local Applications

   Local applications require local connectivity. They must continue to
   work even if the unmanaged network is isolated from the Internet.

   Local applications typically use ad hoc naming systems. Many of
   these systems are proprietary; an example of a standard system is the
   service location protocol (SLP) [RFC2608].

   The security of local applications will usually be enhanced if these
   applications can be effectively isolated from the global Internet.

4.2. Requirements of Client Applications

   Client applications require global connectivity. In an IPv6 network,
   we would expect the client to use a global IPv6 address, which will
   have to remain stable for the duration of the client-server session.

   Client applications typically use the domain name system to locate
   servers. In an IPv6 network, the client must be able to locate a DNS
   resolver.

   Many servers try to look up a DNS name associated with the IP address
   of the client. In an IPv4 network, this IP address will often be
   allocated by the Internet service provider to the gateway, and the
   corresponding PTR record will be maintained by the ISP. In many
   cases, these PTR records are perfunctory, derived in an algorithmic
   fashion from the IPv4 address; the main information that they contain
   is the domain name of the ISP. Whether or not an equivalent function
   should be provided in an IPv6 network is unclear.

4.2.1. Privacy Requirement of Client Applications

   It is debatable whether the IPv6 networking service should be
   engineered to enhance the privacy of the clients, and specifically
   whether support for RFC 3041 [RFC3041] should be required. RFC 3041
   enables hosts to pick IPv6 addresses in which the host identifier is
   randomized; this was designed to make sure that the IPv6 addresses
   and the host identifier cannot be used to track the Internet
   connections of a device's owner.

   Many observe that randomizing the host identifier portion of the
   address is only a half measure. If the unmanaged network address
   prefix remains constant, the randomization only hides which host in
   the unmanaged network originates a given connection, e.g., the
   children's computer versus their parents'. This would place the
   privacy rating of such connections on a par with that of IPv4
   connections originating from an unmanaged network in which a NAT
   manages a static IPv4 address; in both cases, the IPv4 address or the
   IPv6 prefix can be used to identify the unmanaged network, e.g., the
   specific home from which the connection originated.

   However, randomization of the host identifier does provide benefits.
   First, if some of the hosts in the unmanaged network are mobile, the
   randomization destroys any correlation between the addresses used at
   various locations: the addresses alone could not be used to determine
   whether a given connection originates from the same laptop moving
   from work to home, or used on the road. Second, the randomization
   removes any information that could be extracted from a hardwired host
   identifier; for example, it will prevent outsiders from correlating a

serial number with a specific brand of expensive electronic
equipment, and to use this information for planning marketing
campaigns or possibly burglary attempts.

Randomization of the addresses is not sufficient to guarantee
privacy. Usage can be tracked by a variety of other means, from
application level "cookies" to complex techniques involving data
mining and traffic analysis. However, we should not make a bad
situation worse. Other attacks to privacy may be possible, but this
is not a reason to enable additional tracking through IPv6 addresses.

Randomization of the host identifier has some costs: the address
management in hosts is more complex for the hosts, reverse DNS
services are harder to provide, and the gateway may have to maintain
a larger cache of neighbor addresses; however, experience from
existing implementation shows that these costs are not overwhelming.
Given the limited benefits, it would be unreasonable to require that
all hosts use privacy addresses; however, given the limited costs, it
is reasonable to require that all unmanaged networks allow use of
privacy addresses by those hosts that choose to do so.

4.3. Requirements of Peer-to-Peer Applications

Peer-to-peer applications require global connectivity. In an IPv6
network, we would expect the peers to use a global IPv6 address,
which will have to remain stable for the duration of the peer-to-peer
session.

There are multiple aspects to the security of peer-to-peer
applications, many of which relate to the security of the rendezvous
system. If we assume that the peers have been able to safely
exchange their IPv6 addresses, the main security requirement is the
capability to safely exchange data between the peers without
interference by third parties.

Private conversations by one of the authors with developers of peer-
to-peer applications suggest that many individuals would be willing
to consider an "IPv6-only" model if they can get two guarantees:

1) That there is no regression from IPv4, i.e., that all customers
   who could participate in a peer-to-peer application using IPv4 can
   also be reached by IPv6.

2) That IPv6 provides a solution for at least some of their hard
   problems, e.g., enabling peers located behind an IPv4 NAT to
   participate in a peer-to-peer application.

Requiring IPv6 connectivity for a popular peer-to-peer application
could create what economists refer to as a "network effect", which in
turn could significantly speed up the deployment of IPv6.

4.4. Requirements of Server Applications

Server applications require global connectivity, which in an IPv6
network implies global addresses. In an IPv4 network utilizing a
NAT, for each service provided by a server, the NAT has to be
configured to forward packets sent to that service to the server that
offers the service.

Server applications normally rely on the publication of the server's
address in the DNS. This, in turn, requires that the server be
provisioned with a "global DNS name".

The DNS entries for the server will have to be updated, preferably in
real time, if the server's address changes. In practice, updating
the DNS can be slow, which implies that server applications will have
a better chance of being deployed if the IPv6 addresses remain
stable.

The security of server applications depends mostly on the correctness
of the server, and also on the absence of collateral effects: many
incidents occur when the opening of a server on the Internet
inadvertently enables remote access to some other services on the
same host.

5. Stages of IPv6 Deployment

We expect the deployment of IPv6 to proceed from an initial state in
which there is little or no deployment, to a final stage in which we
might retire the IPv4 infrastructure. We expect this process to
stretch over many years; we also expect it to not be synchronized, as
different parties involved will deploy IPv6 at different paces.

In order to get some clarity, we distinguish three entities involved
in the transition of an unmanaged network: the ISP (possibly
including ISP consumer premise equipment (CPE)), the home gateway,
and the hosts (computers and appliances). Each can support IPv4-
only, both IPv4 and IPv6, or IPv6-only. That gives us 27
possibilities. We describe the most important cases. We will assume
that in all cases the hosts are a combination of IPv4-only, dual
stack, and (perhaps) IPv6-only hosts.

The cases we will consider are:

A) a gateway that does not provide IPv6 at all;
B) a dual-stack gateway connected to a dual stack ISP;
C) a dual stack gateway connected to an IPV4-only ISP; and
D) a gateway connected to an IPv6-only ISP

In most of these cases, we will assume that the gateway includes a
NAT: we realize that this is not always the case, but we submit that
it is common enough that we have to deal with it; furthermore, we
believe that the non-NAT variants of these cases map fairly closely
to this same set of cases. In fact, we can consider three non-NAT
variants: directly connected host; gateway acting as a bridge; and
gateway acting as a non-NAT IP router.

The cases of directly connected hosts are, in effect, variants of
cases B, C, and D, in which the host can use all solutions available
to gateways: case B if the ISP is dual stack, case C if the ISP only
provides IPv4 connectivity, and case D if the ISP only provides IPv6
connectivity.

In the cases where the gateway is a bridge, the hosts are, in effect,
directly connected to the ISP, and for all practical matter, behave
as directly connected hosts.

The case where the gateway is an IP router but not a NAT will be
treated as small variants in the analysis of case A, B, C, and D.

5.1. Case A, Host Deployment of IPv6 Applications

In this case, the gateway doesn't provide IPv6; the ISP may or may
not provide IPv6, but this is not relevant since the non-upgraded
gateway would prevent the hosts from using the ISP service. Some
hosts will try to get IPv6 connectivity in order to run applications
that require IPv6, or work better with IPv6. The hosts, in this
case, will have to handle the IPv6 transition mechanisms on their
own.

There are two variations of this case, depending on the type of
service implemented by the gateway. In many cases, the gateway is a
direct obstacle to the deployment of IPv6, but a gateway which is
some form of bridge-mode CPE or which is a plain (neither filtering
nor NAT) router does not really fall into this category.

5.1.1. Application Support in Case A

The focus of Case A is to enable communication between a host on the
unmanaged network and some IPv6-only hosts outside of the network.

The primary focus in the immediate future, i.e., for the early
adopters of IPv6, will be peer-to-peer applications. However, as
IPv6 deployment progresses, we will likely find a situation where
some networks have IPv6-only services deployed, at which point we
would like case A client applications to be able to access those
services.

Local applications are not a primary focus of Case A. At this stage,
we expect all clients in the unmanaged network to have either IPv4
only or dual stack support. Local applications can continue working
using IPv4.

Server applications are also not a primary focus of Case A. Server
applications require DNS support, which is difficult to engineer for
clients located behind a NAT, which is likely to be present in this
case. Besides, server applications presently cater mostly to IPv4
clients; putting up an IPv6-only server is not very attractive.

In contrast, peer-to-peer applications are probably both attractive
and easy to deploy: they are deployed in a coordinated fashion as
part of a peer-to-peer network, which means that hosts can all
receive some form of an IPv6 upgrade; they often provide their own
naming infrastructure, in which case they are not dependent on DNS
services.

5.1.2. Addresses and Connectivity in Case A

We saw in 5.1.1 that the likely motivation for deployment of IPv6
connectivity in hosts in case A is a desire to use peer-to-peer and
client IPv6 applications. These applications require that all
participating nodes get some form of IPv6 connectivity, i.e., at
least one globally reachable IPv6 address.

If the local gateway provides global IPv4 addresses to the local
hosts, then these hosts can individually exercise the mechanisms
described in case C, "IPv6 connectivity without provider support."
If the local gateway implements a NAT function, another type of
mechanism is needed. The mechanism to provide connectivity to peers
behind NAT should be easy to deploy, and light weight; it will have
to involve tunneling over a protocol that can easily traverse NAT,
either TCP or preferably UDP, as tunneling over TCP can result in
poor performance in cases of time-outs and retransmissions. If
servers are needed, these servers will, in practice, have to be
deployed as part of the "support infrastructure" for the peer-to-peer
network or for an IPv6-based service; economic reality implies that
the cost of running these servers should be as low as possible.

5.1.3. Naming Services in Case A

   At this phase of IPv6 deployment, hosts in the unmanaged domain have
   access to DNS services over IPv4 through the existing gateway. DNS
   resolvers are supposed to serve AAAA records, even if they only
   implement IPv4; the local hosts should thus be able to obtain the
   IPv6 addresses of IPv6-only servers.

   Reverse lookup is difficult to provide for hosts on the unmanaged
   network if the gateway is not upgraded. This is a potential issue
   for client applications. Some servers require a reverse lookup as
   part of accepting a client's connection, and may require that the
   direct lookup of the corresponding name matches the IPv6 address of
   the client. There is thus a requirement to provide either a reverse
   lookup solution, or to make sure that IPv6 servers do not require
   reverse lookup.

5.2. Case B, IPv6 Connectivity with Provider Support

   In this case, the ISP and gateway are both dual stack. The gateway
   can use native IPv6 connectivity to the ISP and can use an IPv6
   prefix allocated by the ISP.

5.2.1. Application Support in Case B

   If the ISP and the gateway are dual-stack, client applications,
   peer-to-peer applications, and server applications can all be enabled
   easily on the unmanaged network.

   We expect the unmanaged network to include three kinds of hosts:
   IPv4 only, IPv6-only, and dual stack. Obviously, dual stack hosts
   can interact easily with either IPv4 only hosts or IPv6-only hosts,
   but an IPv4 only host and an IPv6-only host cannot communicate
   without a third party performing some kind of translation service.
   Our analysis concludes that unmanaged networks should not have to
   provide such translation services.

   The argument for providing translation services is that their
   availability would accelerate the deployment of IPv6-only devices,
   and thus the transition to IPv6. This is, however, a dubious
   argument since it can also be argued that the availability of these
   translation services will reduce the pressure to provide IPv6 at all,
   and to just continue fielding IPv4-only devices. The remaining
   pressure to provide IPv6 connectivity would just be the difference in
   "quality of service" between a translated exchange and a native
   interconnect.

The argument against translation service is the difficulty of
providing these services for all applications, compared to the
relative ease of installing dual stack solutions in an unmanaged
network. Translation services can be provided either by application
relays, such as HTTP proxies, or by network level services, such as
NAT-PT [RFC2766]. Application relays pose several operational
problems: first, one must develop relays for all applications;
second, one must develop a management infrastructure to provision the
host with the addresses of the relays; in addition, the application
may have to be modified if one wants to use the relay selectively,
e.g., only when direct connection is not available. Network level
translation poses similar problems: in practice, network level
actions must be complemented by "application layer gateways" that
will rewrite references to IP addresses in the protocol, and while
these relays are not necessary for every application, they are
necessary for enough applications to make any sort of generalized
translation quite problematic; hosts may need to be parameterized to
use the translation service, and designing the right algorithm to
decide when to translate DNS requests has proven very difficult.

Not assuming translation services in the network appears to be both
more practical and more robust. If the market requirement for a new
device requires that it interact with both IPv4 and IPv6 hosts, we
may expect the manufacturers of these devices to program them with a
dual stack capability; in particular, we expect general purpose
systems, such as personal computers, to be effectively dual-stack.
The only devices that are expected to be capable of only supporting
IPv6 are those designed for specific applications, which do not
require interoperation with IPv4-only systems. We also observe that
providing both IPv4 and IPv6 connectivity in an unmanaged network is
not particularly difficult: we have a fair amount of experience using
IPv4 in unmanaged networks in parallel with other protocols, such as
IPX.

5.2.2. Addresses and Connectivity in Case B

   In Case B, the upgraded gateway will act as an IPv6 router; it will
   continue providing the IPv4 connectivity, perhaps using NAT. Nodes
   in the local network will typically obtain:

      - IPv4 addresses (from or via the gateway),
      - IPv6 link local addresses, and
      - IPv6 global addresses.

   In some networks, NAT will not be in use and the local hosts will
   actually obtain global IPv4 addresses. We will not elaborate on
   this, as the availability of global IPv4 addresses does not bring any
   additional complexity to the transition mechanisms.

To enable this scenario, the gateway needs to use a mechanism to
obtain a global IPv6 address prefix from the ISP, and advertise this
address prefix to the hosts in the unmanaged network; several
solutions will be assessed in a companion memo [EVAL].

5.2.3. Naming Services in Case B

In case B, hosts in the unmanaged domain have access to DNS services
through the gateway. As the gateway and the ISP both support IPv4
and IPv6, these services may be accessible by the IPv4-only hosts
using IPv4, by the IPv6-only hosts using IPv6, and by the dual stack
hosts using either. Currently, IPv4 only hosts usually discover the
IPv4 address of the local DNS resolver using DHCP; there must be a
way for IPv6-only hosts to discover the IPv6 address of the DNS
resolver.

There must be a way to resolve the name of local hosts to their IPv4
or IPv6 addresses. Typing auto-configured IPv6 addresses in a
configuration file is impractical; this implies either some form of
dynamic registration of IPv6 addresses in the local service, or a
dynamic address discovery mechanism. Possible solutions will be
compared in the evaluation draft [EVAL].

The requirement to support server applications in the unmanaged
network implies a requirement to publish the IPv6 addresses of local
servers in the DNS. There are multiple solutions, including domain
name delegation. If efficient reverse lookup functions are to be
provided, delegation of a fraction of the ip6.arpa tree is also
required.

The response to a DNS request should not depend on the protocol by
which the request is transported: dual-stack hosts may use either
IPv4 or IPv6 to contact the local resolver, the choice of IPv4 or
IPv6 may be random, and the value of the response should not depend
on a random event.

DNS transition issues in a dual IPv4/IPv6 network are discussed in
[DNSOPV6].

5.3. Case C, IPv6 Connectivity without Provider Support

In this case, the gateway is dual stack, but the ISP is not. The
gateway has been upgraded and offers both IPv4 and IPv6 connectivity
to hosts. It cannot rely on the ISP for IPv6 connectivity, because
the ISP does not yet offer ISP connectivity.

5.3.1. Application Support in Case C

   Application support in case C should be identical to that of case B.

5.3.2. Addresses and Connectivity in Case C

   The upgraded gateway will behave as an IPv6 router; it will continue
   providing the IPv4 connectivity, perhaps using NAT. Nodes in the
   local network will obtain:

      - IPv4 addresses (from or via the gateway),
      - IPv6 link local addresses,
      - IPv6 global addresses.

   There are two ways to bring immediate IPv6 connectivity on top of an
   IPv4 only infrastructure: automatic tunnels, e.g., provided by the
   6TO4 technology [RFC3056], or configured tunnels. Both technologies
   have advantages and limitations, which will be studied in another
   document.

   There will be some cases where the local hosts actually obtain global
   IPv4 addresses. We will not discuss this scenario, as it does not
   make the use of transition technology harder, or more complex. Case
   A has already examined how hosts could obtain IPv6 connectivity
   individually.

5.3.3. Naming Services in Case C

   The local naming requirements in case C are identical to the local
   naming requirements of case B, with two differences: delegation of
   domain names, and management of reverse lookup queries.

   A delegation of some domain name is required in order to publish the
   IPv6 addresses of servers in the DNS.

   A specific mechanism for handling reverse lookup queries will be
   required if the gateway uses a dynamic mechanism, such as 6to4, to
   obtain a prefix independently of any IPv6 ISP.

5.4. Case D, ISP Stops Providing Native IPv4 Connectivity

   In this case, the ISP is IPv6-only, so the gateway loses IPv4
   connectivity, and is faced with an IPv6-only service provider. The
   gateway itself is dual stack, and the unmanaged network includes IPv4
   only, IPv6-only, and dual stack hosts. Any interaction between hosts
   in the unmanaged network and IPv4 hosts on the Internet will require
   the provision of some inter-protocol services by the ISP.

5.4.1. Application Support in Case D

   At this phase of the transition, IPv6 hosts can participate in all
   types of applications with other IPv6 hosts. IPv4 hosts in the
   unmanaged network will be able to perform local applications with
   IPv4 or dual stack local hosts.

   As in case B, we will assume that IPv6-only hosts will not interact
   with IPv4-only hosts, either local or remote. We must however assume
   that IPv4-only hosts and dual stack hosts will want to interact with
   IPv4 services available on the Internet: the inability to do so would
   place the IPv6-only provider at a great commercial disadvantage
   compared to other Internet service providers.

   There are three possible ways that an ISP can provide hosts in the
   unmanaged network with access to IPv4 applications: by using a set of
   application relays, by providing an address translation service, or
   by providing IPv4-over-IPv6 tunnels. Our analysis concludes that a
   tunnel service seems to be vastly preferable.

   We already mentioned the drawbacks of the application gateway
   approach when analyzing case B: it is necessary to provide relays for
   all applications, to develop a way to provision the hosts with the
   addresses of these relays, and to modify the applications so that
   they will only use the relays when needed. We also observe that in
   an IPv6-only ISP, the application relays would only be accessible
   over IPv6, and would thus not be accessible by the "legacy" IPv4-only
   hosts. The application relay approach is thus not very attractive.

   Providing a network address and protocol translation service between
   IPv6 and IPv4 would also have many drawbacks. As in case B, it will
   have to be complemented by "application layer gateways" that will
   rewrite references to IP addresses in the protocol; hosts may need to
   be parameterized to use the translation service, and we would have to
   solve DNS issues. The network level protocol translation service
   doesn't appear to be very desirable.

   The preferable alternative to application relays and network address
   translation is the provision of an IPv4-over-IPv6 service.

5.4.2. Addresses and Connectivity in Case D

   The ISP assigns an IPv6 prefix to the unmanaged network, so hosts
   have a global IPv6 address and use it for global IPv6 connectivity.
   This will require delegation of an IPv6 address prefix, as
   investigated in case C.

To enable IPv4 hosts and dual stack hosts accessibility to remote
IPv4 services, the ISP must provide the gateway with at least one
IPv4 address, using some form of IPv4-over-IPv6 tunneling. Once such
addresses have been provided, the gateway effectively acquires dual-
stack connectivity; for hosts inside the unmanaged network, this will
be indistinguishable from the IPv4 connectivity obtained in case B or
C.

5.4.3. Naming Services in Case D

The loss of IPv4 connectivity has a direct impact on the provision of
naming services. In many IPv4 unmanaged networks, hosts obtain their
DNS configuration parameters from the local gateway, typically
through the DHCP service. If the same mode of operation is desired
in case D, the gateway will have to be provisioned with the address
of a DNS resolver and with other DNS parameters, and this
provisioning will have to use IPv6 mechanisms. Another consequence
is that the DNS service in the gateway will only be able to use IPv6
connectivity to resolve queries; if local hosts perform DNS
resolution autonomously, they will have the same restriction.

On the surface, this seems to indicate that the local hosts will only
be able to resolve names if the domain servers are accessible through
an IPv6 address documented in an AAAA record. However, the DNS
services are just one case of "IPv4 servers accessed by IPv6 hosts":
it should be possible to simply send queries through the IPv4
connectivity services to reach the IPv4 only servers.

The gateway should be able to act as a recursive DNS name server for
the remaining IPv4 only hosts.

6. Security Considerations

Security considerations are discussed as part of the applications'
requirements. They include:

- the guarantee that local applications are only used locally,
- the protection of the privacy of clients
- the requirement that peer-to-peer connections are only used by
  authorized peers
- the requirement that tunneling protocols used for IPv6 access over
  IPv4 be designed for secure use
- the related requirement that servers in the infrastructure
  supporting transition scenarios be designed so as to not be
  vulnerable to abuse.

The security solutions currently used in IPv4 networks include a combination of firewall functions in the gateway, authentication and authorization functions in the applications, encryption and authentication services provided by IP security, Transport Layer Security and application specific services, and host-based security products, such as anti-virus software and host firewalls. The applicability of these tools in IPv6 unmanaged networks will be studied in a another document.

7. Acknowledgements

   This document has benefited from the comments of the members of the IETF V6OPS working group, and from extensive reviews by Chris Fischer, Tony Hain, Kurt Erik Lindqvist, Erik Nordmark, Pekka Savola, and Margaret Wasserman.

8. References

8.1. Normative References

   [RFC791]   Postel, J., "Internet Protocol", STD 5, RFC 791,
                September 1981.

   [RFC2460]   Deering, S. and R. Hinden, "Internet Protocol, Version 6
                (IPv6) Specification", RFC 2460, December 1998.

8.2. Informative References

   [EVAL]   Evaluation of Transition Mechanisms for Unmanaged
                Networks, Work in Progress.

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

   [RFC2608]   Guttman, E., Perkins, C., Veizades, J. and M. Day,
                "Service Location Protocol, Version 2", RFC 2608, June
                1999.

   [RFC3056]   Carpenter, B. and K. Moore, "Connection of IPv6 Domains
                via IPv4 Clouds", RFC 3056, February 2001.

   [RFC3022]   Srisuresh, P. and K. Egevang. "Traditional IP Network
                Address Translator (Traditional NAT)", RFC 3022, January
                2001.

   [RFC2993]   Hain, T., "Architectural Implications of NAT", RFC 2993,
                November 2000.

   [RFC3041]   Narten, T. and R. Draves, "Privacy Extensions for
               Stateless Address Autoconfiguration in IPv6", RFC 3041,
               January 2001.

   [RFC2766]   Tsirtsis, G. and P. Srisuresh, "Network Address
               Translation - Protocol Translation (NAT-PT)", RFC 2766,
               February 2000.

   [DNSOPV6]   Durand, A., Ihren, J. and P. Savola, "Operational
               Considerations and Issues with IPv6 DNS", Work in
               Progress.

   [DNSINADDR] Senie, D., "Requiring DNS IN-ADDR Mapping", Work in
               Progress.

9. Authors' Addresses

   Christian Huitema
   Microsoft Corporation
   One Microsoft Way
   Redmond, WA 98052-6399

   EMail: huitema@microsoft.com

   Rob Austein
   Internet Systems Consortium
   950 Charter Street
   Redwood City, CA 94063
   USA

   EMail: sra@isc.org

   Suresh Satapati
   Cisco Systems, Inc.
   San Jose, CA 95134
   USA

   EMail: satapati@cisco.com

   Ronald van der Pol
   NLnet Labs
   Kruislaan 419
   1098 VA Amsterdam
   NL

   EMail: Ronald.vanderPol@nlnetlabs.nl

10. Full Copyright Statement

Intellectual Property

Acknowledgement