

INTERNET DRAFT
<draft-ietf-ngtrans-unmanscope-02.txt>
November 1, 2002
Expires May 1, 2002

C. Huitema
Microsoft
R. Austein
Bourgeois Dilettant
R. van der Pol
NLnet Labs

Unmanaged Networks Transition Scope

Status of this memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Abstract

In order to evaluate the suitability of transition mechanisms, 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.

1 Introduction

In order to evaluate the suitability of transition mechanisms, 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.

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 are connected to the subnet:

```


+-----+
| Host +--+

```

Huitema et al.

[Page 1]

INTERNET DRAFT Unmanaged Networks Transition Scope March 11, 2002



Between the subnet and the ISP access link is a gateway, which may or may not perform NAT and firewall function. A key point of this configuration is that the gateway is typically not "managed". In most cases, it is a simple "appliance", which incorporates some static policies. There are however many cases in which the gateway is procured and configured by the ISP, and there are also some common cases in which we find to back to back gateways, one managed by the ISP and the other added by the owner of the unmanaged network.

The access link between the unmanaged network and the ISP can be either static, i.e. a permanent connection, or dynamically established, i.e. a dial-up or ISDN connection.

In a degenerate case, an unmanaged network can be constituted of a single host, directly connected to an ISP.

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 network: their status vary.

3.1 Local applications

Local applications are meant to only involve the hosts that are part

of the unmanaged network. Typical examples are the sharing of file or printers.

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

Huitema et al.

[Page 2]

INTERNET DRAFT Unmanaged Networks Transition Scope March 11, 2002

[3.2](#) **Client applications**

Client applications are those that involve a client on the unmanaged network and a server at a remote location. A typical example is 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.

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

[3.3](#) **Peer-to-peer applications**

There are two kinds of peer-to-peer applications, the "local peer-to-peer" that only involve hosts on the unmanaged network, and the "remote peer-to-peer" that involve both hosts on the unmanaged network and hosts outside the network. We will only consider here the "remote peer-to-peer" applications, as the local peer-to-peer applications are a subset of the "local applications."

Peer-to-peer applications are a restricted subset of the server applications, in which the services are only meant to be used by well identified peers outside the unmanaged network. These applications are often facilitated by a server outside the unmanaged networks. Examples of a peer-to-peer application would be a video-conference over IP, facilitated by a SIP server, or a distributed game application, facilitated by a "game lobby".

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", to effectively restructure the peer-to-peer connection in two 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. 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, 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 the servers with a

Huitema et al.

[Page 3]

INTERNET DRAFT Unmanaged Networks Transition Scope March 11, 2002

stable DNS name, and associating the global IPv4 address of the nat/firewall with that name. Since updating DNS is a management task, it somewhat falls outside the scope of an unmanaged network.

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: the applications that used to work 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; the deployment of these IPv6 applications should be simple and easy.

The application requirements are expressed in mostly three dimensions: connectivity, naming, and security. Connectivity issues include the provision of IPv6 addresses and their quality: do host need a global scope address, should this address be stable, or more precisely what should be the expected lifetime of the address. Naming issues include the management of names for the hosts: do hosts need a DNS-name, is inverse name resolution 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 working 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 standard system is the service location protocol (SLP).

The security of local applications is 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 server.

Many servers try to look up a DNS name associated to 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 most cases, these PTR records are perfunctory, derived in an algorithmic

Huitema et al.

[Page 4]

INTERNET DRAFT Unmanaged Networks Transition Scope March 11, 2002

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.

An important security issues with client applications is, maintaining the privacy of the client. Many potential users of IPv6 observe that the current NAT configuration provides some level of privacy: all communications to outside servers appear to come from the single IPv4 address of the NAT; the source IPv4 address does not identify a specific host inside the unmanaged network. Users migrating to IPv6 should not experience a "privacy regression"; it should be possible to prevent external servers from using the source IPv6 addresses of successive connections to link these connections to a specific internal host.

*** There is a debate as to whether this privacy requirement should be maintained. An opinion is that this is just a random by-product of using a NAT that should not be turned in a requirement. However, privacy advocates will probably insist on the "no regression" requirement. ***

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 between client and server.

Peer-to-peer applications often use ad hoc naming systems, sometimes derived from an "instant messaging" service. Many of these systems are proprietary; an example of standard system is the session

initiation protocol (SIP). In these systems, the peers register their presence to a "rendezvous" server, using a name specific to the service; the case of SIP, they would use a SIP URL, of the form "sip:user@example.com". A peer to peer session typically starts by an exchange of synchronization messages through the rendezvous servers, during which the peers exchange the addresses that will be used for the 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 with developers of peer-to-peer applications showed that many 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 that could participate in a peer-to-peer application using IPv4 can

Huitema et al.

[Page 5]

INTERNET DRAFT Unmanaged Networks Transition Scope March 11, 2002

also be reached by IPv6.

2) That IPv6 provides a solution for at least some of their hard problems, i.e. 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.

Server applications normally rely on the publication of the server's address in the DNS. This, in turns, 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 is slow, which implies that server applications will have a better chance of being deployed if the IPv6 addresses remain stable for a long period.

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

The deployment of IPv6 over time is expected 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 several years; we also expect it to not be synchronized, as different parties involved will deploy IPv6 at different pace. In order to get some clarity, we distinguish three entities involved in the transition of an unmanaged network: the ISP (possibly including ISP 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 consider that in all cases the hosts are a combination of IPv4-only, dual stack and IPv6-only hosts.

The cases we will consider are:

- A) Gateway does not provide IPv6
- B) ISP and gateway are dual stack
- C) Gateway is IPv6 capable, dual stack, ISP is not

Huitema et al.

[Page 6]

INTERNET DRAFT Unmanaged Networks Transition Scope March 11, 2002

- D) ISP is IPv6-only

The case where the ISP is IPv6 capable, but the gateway is not is similar to case A.

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.

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. Besides, server applications, at this stage, cater mostly to IPv4 clients; putting up an IPv6-only server is not very attractive.

In contrast, peer-to-peer applications are 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 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 a primary motivation for the deployment of IPv6 connectivity in hosts is participation to peer-to-peer applications, and also to IPv6-only client applications. These applications require that all participating nodes get some form of IPv6 connectivity, i.e. at least one globally reachable IPv6 address. The mechanism to provide connectivity to peers behind NAT should be easy to deploy, and light weight; it will have to involve tunneling over UDP, as this is the practical way to traverse a NAT. If servers are

Huitema et al.

[Page 7]

INTERNET DRAFT Unmanaged Networks Transition Scope March 11, 2002

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 hard to provide 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 either provide a reverse lookup solution, or 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 dual stack. The gateway can use native IPv6 connectivity to the ISP and 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. Our first requirement is that IPv6-only hosts should be "upward compatible" with IPv4 only hosts, i.e. be able to access all the services and applications that are available to IPv4 only hosts, and then to be able to access additional IPv6-only services. As a result, the application support for IPv6-only hosts should be at least the following:

- support for local applications with IPv6-only and dual stack host, and at a minimum possibility to initiate associations with local IPv4 hosts.
- support for client applications with both IPv4 and IPv6 servers, with the condition that the IPv4 server has a global IPv4 address.
- support for peer-to-peer applications with IPv6 hosts.
- provision of server applications to IPv6 clients.

Huitema et al.

[Page 8]

INTERNET DRAFT Unmanaged Networks Transition Scope March 11, 2002

The paragraph is brand new text that the authors have not have not discussed fully, and may not entirely agree:

-- Whether to provide "backward compatibility" or not is debatable. Backward compatibility would allow IPv4 only clients to use IPv6-only services, but it can be perceived as a disincentive to the IPv6 deployment, as there would then be less reason to upgrade the IPv4 only hosts. The case for backward compatibility depends on the type of application. There is a strong incentive to enable interoperability for local applications, e.g. to let the existing IPv4 only hosts use the services a brand new IPv6-only printer.

There is also a reasonable case for letting IPv4 only clients access remote IPv6 services, if it can be achieved simply. There is a much lesser case for letting these clients participate in peer-to-peer application, as these applications tend to behave poorly in the presence of translation or port mapping services; there is also not much of a case for letting local IPv4 servers publish their service on the IPv6 Internet. --

Another opinion is:

-- I think that there is a case to be made that using translation tools boosts the previously mentioned network effect and makes it work in our favor; the somewhat crippled nature of the translation tools is in fact also a good thing, because it gives people an incentive to throw them away once they realize that they are broken. I realize that this is heresy, so be it. --

5.2.2 Addresses and connectivity in Case B

In Case B, the upgraded gateway will behave as an IPv6 router; it will continue providing the IPv4 connectivity of a non-upgraded NAT. Nodes in the local network will obtain:

- IPv4 natted addresses,
- IPv6 link local addresses,
- IPv6 global addresses.

The hosts could also obtain IPv6 site local addresses, if the gateway advertises a site local prefix. This is as debatable: site local addresses provide some isolation to site local application from network connectivity events and network based attacks; however, managing non unique addresses can be problematic if some local hosts are multi-homed, as is common with VPN connections.

To enable this scenario, the gateway need to use a mechanism obtain a global 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](#)].

In order to meet the "backward compatibility" requirement, some form

Huitema et al.

[Page 9]

INTERNET DRAFT Unmanaged Networks Transition Scope March 11, 2002

of translation service must be used to enable connectivity between local IPv4 only and IPv6-only hosts.

To use the translation service, the IPv6 hosts will have to send packets to a "mapped address" that represents the IPv4 host in the IPv6 address space. Several possibilities are evaluated in a companion document.

The IPv4 hosts can use the translation service in two ways. If the dialog is initiated by the IPv6 host, the IPv4 host simply has to respond to the incoming packets, reversing source and destination addresses. In order to initiate a dialog, the IPv4 host will have to learn the "mapped IPv4 address" of the IPv6 host; this requires either some form of explicit configuration, in which the IPv6 host acquires an IPv4 address from the gateway and publishes it in name services or in local applications; or some form of implicit translation, in which the gateway intercepts the name services and replaces a published IPv6 address by an automatically provisioned IPv4 address. Implicit translation is problematic, since it cannot work in presence of secure name services; however, the hosts that need translation most are also those that are least likely to support secure name services.

5.2.3 Naming services in Case B

At this phase of IPv6 deployment, 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 discover the IPv4 address of the local DNS server using DHCP; there must be a way for IPv6-only hosts to discover the IPv6 address of the DNS server.

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.

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 variations of domain name delegation. If we want to provide efficient reverse lookup functions, delegation of a fraction of the ip6.arpa tree is also required.

If "backward compatibility" of local applications is required, the local IPv4 only hosts will have to obtain the IPv4 address at which an IPv6-only local host can be contacted; an address translation relay may then be used to provide connectivity. This implies that

the local DNS service must be able to associate a local IPv4 address with the local IPv6-only nodes, and that this association must be coordinated with the translation function.

If "backward compatibility" of client applications is required, the local IPv4 only hosts will also have to obtain the IPv4 address at which an IPv6-only remote host can be contacted.

The response to a DNS request should not depend of the protocol in which the request is transported: dual-stack hosts may indifferently use IPv4 or IPv6 to contact the local resolver; the choice of IPv4 or IPv6 will be random; the value of the response should not depend of a random event.

5.3 Case C, IPv6 connectivity without provider support

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

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 of non-upgraded NAT. Nodes in the local network will obtain:

- IPv4 natted addresses,
- IPv6 link local addresses,
- IPv6 global addresses.

The clients could also obtain IPv6 site local addresses, if the gateway advertises a site local prefix. Whether gateways should actually do that is debatable.

There are two ways to bring immediate IPv6 connectivity on top of an IPv4 only infrastructure: automatic tunnels provided by the [6T04] technology, or configured tunnels. Both technologies have advantages and limitations, which will be studied in a companion document.

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

Huitema et al.

[Page 11]

IPv6 addresses of servers in the DNS. As the ISP does not provide support for IPv6 in case C, the delegation mechanism will have to be provided independently of the IP connectivity mechanism.

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.

5.4.1 Application support in Case D

At this phase of the transition, IPv6 hosts can perform all types of applications with other IPv6 hosts. The specific support requirements concern the amount of support for the remaining IPv4 hosts, and the access by IPv6 hosts to IPv4 services. The following proposition seems to be a good compromise:

- IPv4 only hosts should be able to use the services available on the IPv4 internet, i.e. they should not lose services that were previously available.
- It is desirable that IPv4 only clients be able to access IPv6 servers available on the IPv6 Internet.
- IPv6-only hosts should also be able to use the services available on the IPv4 internet, as they were in case B.

The support of local applications should be the same as in Case B.

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. The gateway may provide [RFC 1918](#) addresses to the hosts, but they don't have global IPv4 connectivity anymore. The unmanaged network does not have any global IPv4 addresses anymore, so translation tools that require those cannot be used.

We have a requirement that IPv6-only hosts continue to be able to access previously available IPv4 services. Since the IPv6-only hosts by definition cannot use IPv4, they will require some Internet based translation service, which may be provided by the IPv6 ISP or by a third party. The IPv6-only host will have to obtain the IPv6 service representation of the IPv4 address of the server, either by using a

lookup service or by combining a learned IPv6 prefix and the IPv4 address of the server.

Huitema et al.

[Page 12]

We have a requirement that IPv4 only hosts be able to use IPv6-only services provided on the global Internet. This is a requirement already covered in case B; the proposed solution requires some form of IPv4 to IPv6 translation in the gateway.

We also have a requirement that IPv4 hosts be able to continue using IPv4 services that were previously available. A possibility is to combine a local translation service provided by the gateway with the Internet based translation service used by IPv6-only hosts, effectively using back to back translations: from local IPv4 to global IPv6 in the gateway, and from global IPv6 to global IPv4 in the Internet based service.

5.4.3 Naming services in Case D

The loss of IPv4 connectivity has a direct impact on the provision of naming services. An obvious consequence is the gateway will have to be provisioned with the address of a DNS server and with other DNS parameters, and that 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 a 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 address translation services to reach the IPv4 only servers.

The gateway should be able to act as a "DNS proxy" 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.

7 IANA Considerations

This memo does not include any request to IANA.

8 Copyright

The following copyright notice is copied from [RFC 2026](#) [Bradner, 1996], [Section 10.4](#), and describes the applicable copyright for this document.

Copyright (C) The Internet Society July 12, 2001. All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assignees.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

9 Intellectual Property

The following notice is copied from [RFC 2026](#) [Bradner, 1996], [Section 10.4](#), and describes the position of the IETF concerning intellectual property claims made against this document.

The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use other technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in [BCP-11](#). Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification

can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any

Huitema et al.

[Page 14]

copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

10 Acknowledgements

This draft has benefited from extensive reviews by Tony Hain, Suresh K Satapati, and Margaret Wasserman.

11 References

[EVAL] Evaluation of Transition Mechanisms for Unmanaged Networks, work in progress.

12 Authors' Addresses

Christian Huitema
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052-6399
Email: huitema@microsoft.com

Rob Austein
Email: sra@hactrn.net

Ronald van der Pol
Email: Ronald.vanderPol@surfnet.nl

Table of Contents:

1	Introduction	1
2	Topology	1
3	Applications	2
3.1	Local applications	2
3.2	Client applications	3
3.3	Peer-to-peer applications	3
3.4	Server applications	3
4	Application requirements of an IPv6 unmanaged network	4
4.1	Requirements of local applications	4
4.2	Requirements of client applications	4
4.3	Requirements of peer-to-peer applications	5
4.4	Requirements of server applications	6
5	Stages of IPv6 deployment	6
5.1	Case A, host deployment of IPv6 applications	7
5.1.1	Application support in Case A	7
5.1.2	Addresses and connectivity in Case A	7
5.1.3	Naming services in Case A	8
5.2	Case B, IPv6 connectivity with provider support	8
5.2.1	Application support in Case B	8
5.2.2	Addresses and connectivity in Case B	9
5.2.3	Naming services in Case B	10
5.3	Case C, IPv6 connectivity without provider support	11
5.3.1	Application support in Case C	11
5.3.2	Addresses and connectivity in Case C	11
5.3.3	Naming services in Case C	11
5.4	Case D, ISP stops providing native IPv4 connectivity	12
5.4.1	Application support in Case D	12
5.4.2	Addresses and connectivity in Case D	12
5.4.3	Naming services in Case D	13
6	Security Considerations	13
7	IANA Considerations	13
8	Copyright	13
9	Intellectual Property	14
10	Acknowledgements	15
11	References	15
12	Authors' Addresses	15

