IETF
Internet Draft
January 1999

B. Carpenter
K. Moore

# Connection of IPv6 Domains via IPv4 Clouds without Explicit Tunnels



Abstract

draft-ietf-ngtrans-6to4-00.txt

This memo specifies an optional mechanism for assigning a unique IPv6 address prefix to any site that currently has at least one globally unique IPv4 address, and describes scenarios for using such a prefix during the co-existence phase of IPv4 to IPv6 transition.

The motivation for this method is to allow isolated IPv6 domains, attached to an IPv4 network which has no native IPv6 support, to communicate with other such IPv6 domains with minimal manual configuration. Effectively it treats the IPv4 network as a virtual link layer. It also automatically provides 80 bits of globally unique IPv6 address space to any site with at least one globally unique IPv4 address. If combined with a Network Address Translator (NAT), it allows the NAT to provide a globally-unique and globally-routable IPv6 address to each of its client hosts.

## Status of this Memo

Table of Contents:

## 1. Introduction

<DISCLAIMER> This version has been released at the request of the NGTRANS chairs prior to the interim meeting. It has known deficiencies (MTU text is wrong, multicast text is confused, address selection needs work, scenario for 6to4/native IPv6 interworking needs much expansion, various nits).</DISCLAIMER>

This memo specifies an optional mechanism for assigning a unique IPv6 address prefix to any site that currently has at least one globally unique IPv4 address, and describes scenarios for using such a prefix during the co-existence phase of IPv4 to IPv6 transition. Note that these scenarios are only part of the total picture of transition to IPv6, in addition to the mechanisms in [RFC 1933].

The motivation for this method is to allow isolated IPv6 domains, attached to a wide area network which has no native IPv6 support, to communicate with other such IPv6 domains with minimal manual configuration. Effectively it treats the wide area IPv4 network as a virtual link layer.

IPv6 domains connected using this method do not require IPv4-compatible addresses or configured tunnels. In this way IPv6 gains considerable independence of the underlying wide area network and can step over many hops of IPv4 subnets. The abbreviated name of this mechanism is 6to4 (not to be confused with [6OVER4]). The 6to4 mechanism is implemented entirely in boundary routers, without host modifications.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

## 2. IPv6 Prefix Allocation

Suppose that a subscriber site has at least one valid, globally unique 32-bit IPv4 address, referred to in this document as V4ADDR. This address MUST be duly allocated to the site by an address registry (possibly via a service provider) and it MUST NOT be a private address [RFC 1918].

The IANA has permanently assigned one 13-bit IPv6 Top Level Aggregator (TLA) identifier under the IPv6 Format Prefix 001 [AARCH, AGGR], referred to in this document as TLA624. Its numeric value is 0x0010.

[*** this assignment remains to be made and may change ***]

The subscriber site is then deemed to have the following IPv6 address prefix, without any further assignment procedures being necessary:

- Prefix length: 48 bits
- Format prefix: 001
- TLA value: TLA624
- NLA value: V4ADDR

This is illustrated as follows:

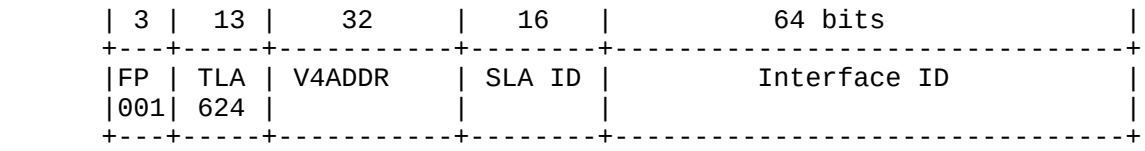


Thus, this prefix has exactly the same format as normal prefixes assigned according to [AGGR]. Within the subscriber site it can be used for automated address assignment and discovery according to the normal mechanisms such as [CONF, DISC]. No changes are required in IPv6 host software. If the subscriber site is not yet running native IPv6, but is running IPv4 multicast, this "6 to 4" address prefix can be used in conjunction with the "6 over 4" mechanism [6OVER4]. Thus isolated IPv6 hosts within isolated IPv6 domains can communicate by using "6 over 4" to a boundary router and "6 to 4" over the wide area.

## 3. Maximum Transmission Unit

The default MTU size for IPv6 packets sent to an IPv4 domain is 1480 octets. This size may be varied by a Router Advertisement [DISC] containing an MTU option which specifies a different MTU, or by manual configuration of each node.

Note that if by chance the IPv6 MTU size proves to be too large for some intermediate IPv4 subnet, IPv4 fragmentation will ensue. While undesirable, this is not disastrous. However, the IPv4 "do not fragment" bit MUST NOT be set in the encapsulating IPv4 header.

## 4. Frame Format

IPv6 packets are transmitted in IPv4 packets [RFC 791] with an IPv4 protocol type of 41, the same as has been assigned [RFC 1933] for IPv6 packets that are tunneled inside of IPv4 frames. The IPv4 header contains the Destination and Source IPv4 addresses. One or both of these will be identical to the V4ADDR field of an IPv6 prefix formed as specified above (see section 6 for more details). The IPv4 packet body contains the IPv6 header and payload.

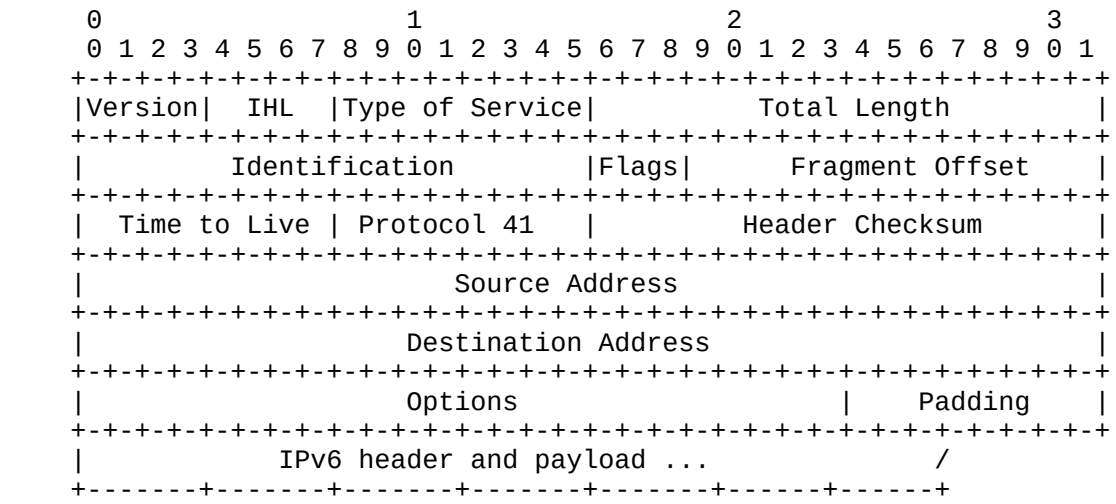


If there are IPv4 options, then padding SHOULD be added to the IPv4 header such that the IPv6 header starts on a boundary that is a 32-bit offset from the end of the datalink header.

The IPv4 Time to Live will be set as normal [RFC 791], as will the encapsulated IPv6 hop limit [IPv6].

## 5. Multicast and Anycast

Nothing prevents IPv6 multicast packets being sent to or sourced from a 6to4 site. However, since unicast routing for 6to4 has some peculiarities discusssed in the next section, a multicast treee that covers both 6to4 and non-6to4 sites is likely to have a sub-optimal topology. If it has a single root in the 6to4 address space, the multicast packets are likely to traverse large regions of the IPv4 network as well as corresponding regions of the IPv6 network. If the tree has multiple roots in the 6to4 address space, 6to4 encapsulation of the same multicast packet will take place multiple times.

The allocated anycast address space [ANYCAST] is compatible with TLA624 prefixes.

## 6. Scenarios, scaling, and transition to normal prefixes

The typical deployment scenario for 6to4 is for use between a number of sites, each of which has at least one connection to the global IPv4 Internet. There is no requirement that the sites all connect to the same Internet service provider. Thus any of the sites is able to send IPv4 packets to any of the others. By definition, each site has an IPv6 prefix in the format defined in Section 2. It will therefore create DNS records for these addresses. For example, site A which owns IPv4 address 192.1.2.3 will create DNS records with the IPv6 prefix {FP=001, TLA=TLA624, NLA=192.1.2.3}/48. Site B which owns

address 9.254.253.252 will create DNS records with the IPv6 prefix {FP=001, TLA=TLA624, NLA=9.254.253.252}/48.

Suppose an IPv6 host on site B queries the DNS entry for a host on site A, and the DNS returns multiple IPv6 prefixes. If the host picks the 6to4 prefix according to the normal rules for multiple prefixes, it will simply send packets to an IPv6 address formed with the prefix {FP=001, TLA=TLA624, NLA=192.1.2.3}/48. It is essential that they are sourced from the prefix {FP=001, TLA=TLA624, NLA=9.254.253.252}/48.

[*** Query - how does the source prefix get specified? ***]

[*** Note - intend to insert an A6 record exmple here ***]

The only change to standard IPv6 routing is that the egress router on each 6to4 site MUST include the sending rule:

```
if the destination address of an IPv6 packet is
  {FP=001, TLA=TLA624}/16
  then
      if the NLA field is an IPv4 address assigned to this site
        then queue the packet for local IPv6 forwarding
        else encapsulate the packet in IPv4 as in Section 3
             with destination address set to the NLA value V4ADDR;
             queue the packet for global IPv4 forwarding.
```

A simple decapsulation rule for incoming IPv4 packets with protocol type 41 is also required:

```
Apply any security checks (see Section 8)
Remove the IPv4 header
Submit the packet to local IPv6 routing.
```

In this scenario, no IPv4 routing information is imported into IPv6 routing (nor vice versa). The above special sending rule is the only contamination of IPv6 forwarding, and it occurs only at egress routers.

Any IPv6 router willing to act as a relay from native IPv6 to the 6to4 address space advertises a route to {FP=001, TLA=TLA624}/16. Within a 6to4 site, this prefix will normally be handled by the default IPv6 router.

In this scenario, any number of 6to4 sites can interoperate with no prior agreement, no tunnel setup, and no special requirements from the IPv4 service. All that is required is the appropriate DNS entries and the special sending rule configured in the egress router. This router SHOULD also generate the appropriate IPv6 prefix announcements [CONF, DISC].

Note that 6to4 only requires one unique IPv4 address per participating site. It is RECOMMENDED that in any case each site should use only one IPv4 address, and that should be the address of its egress router. Note that this router may well also be a firewall and/or an IPv4 network address translator (NAT). This does not affect the 6to4 mechanism. In particular, using 6to4 in conjunction with an

IPv4 NAT offers the site concerned an extra 80 bits of globally unique address space, automatically and free of charge, behind the IPv4 address of the NAT.

Because of the lack of setup and prior agreement, and the distributed deployment model, there are believed to be no particular scaling issues with the 6to4 mechanism.

Sites which are multihomed on IPv4 MAY extend the 6to4 scenario by using a TLA624 prefix for each IPv4 egress router, thereby automatically obtaining a degree of IPv6 multihoming.

Sites which have at least one native IPv6 egress, in addition to a 6to4 egress, will therefore have at least one IPv6 prefix which is not a TLA624 prefix. Such sites' DNS entries will reflect this. If two such sites need to interoperate, whether the 6to4 route or the native route will be used depends on the priorities of the DNS entries. These priorities are an operational choice by which a site can decide whether it wants to receive its IPv6 traffic in preference via 6to4 or via the native connection.

When a site acquires a native IPv6 connection it MUST NOT advertise its TLA624 prefix on that connection, and service providers MUST filter out and discard any TLA624 prefix advertisements longer than /16.

If these rules are followed, then a site can migrate from using 6to4 to using native IPv6 connections over a long period of co-existence, with no need to stop 6to4 until it has ceased to be used.

There is nothing to stop the above scenario being deployed within a private corporate network as part of its internal transition to IPv6; the corporate IPv4 backbone would serve as the virtual link layer for individual corporate sites using TLA624 prefixes. In this case the V4ADDR could even be a private IPv4 address [RFC 1918] as long as it was unique within the private network and the corresponding DNS record was never advertised outside.

## 7. IANA considerations

No assignments by the IANA are required except the special TLA value TLA624 = 0x0010. [*** value to be confirmed ***]

## 8. Security considerations

Implementors should be aware that, in addition to posssible attacks against IPv6, security attacks against IPv4 must also be considered. Use of IP security at both IPv4 and IPv6 levels should nevertheless be avoided, for efficiency reasons. For example, if IPv6 is running encrypted, encryption of IPv4 would be redundant except if traffic analysis is felt to be a threat. If IPv6 is running authenticated,

then authentication of IPv4 will add little. Conversely, IPv4 security will not protect IPv6 traffic once it leaves the 6to4 domain. Therefore, implementing IPv6 security is required even if IPv4 security is available.

By default, 6to4 traffic will be accepted and decapsulated from any source from which regular IPv4 traffic is accepted. If this is for any reason felt to be a security risk (for example, if IPv6 spoofing is felt to be more likley than IPv4 spoofing), then additional source-based packet filtering could be applied. A possible plausibility check is whether the encapsulating IPv4 address is consistent with the encapsulated TLA624 address. If this check is applied, exceptions to it must be configured to admit traffic from relay routers (Section 6). TLA624 traffic must also be excepted from checks applied to prevent spoofing of "6 over 4" traffic [6OVER4].

## Acknowledgements

The basic idea presented above is probably not original, and we have had invaluable comments from members of the NGTRANS working group. Some text has been copied from [6OVER4].

## References

[AARCH] Hinden, R., and S. Deering, "IP Version 6 Addressing Architecture", RFC 2373

[AGGR] Hinden., R, O'Dell, M., and Deering, S., "An IPv6 Aggregatable Global Unicast Address Format", RFC 2374

[CONF] Thomson, S., and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462

[DISC] Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461

[IPV6] Deering, S., and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460

[6OVER4] Carpenter, B., and Jung., C. "Transmission of IPv6 over IPv4 Domains without Explicit Tunnels", draft-ietf-ipngwg-6over4-02.txt (work in progress).

[ANYCAST] Johnson, D. and Deering, S., Reserved IPv6 Subnet Anycast Addresses, draft-ietf-ipngwg-resv-anycast-01.txt (work in progress).

[RFC 791] Postel, J., "Internet Protocol", RFC 791

[RFC 1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., de Groot, G., Lear, E., "Address Allocation for Private Internets", RFC 1918

[RFC 1933] Transition Mechanisms for IPv6 Hosts and Routers. R. Gilligan & E. Nordmark, RFC 1933

[RFC 2119] Key words for use in RFCs to Indicate Requirement Levels. S. Bradner, RFC 2119

## Authors' Addresses

Brian E. Carpenter
IBM United Kingdom Laboratories
MP 185, Hursley Park
Winchester, Hampshire SO21 2JN, UK

Email: brian@hursley.ibm.com

Keith Moore
Innovative Computing Laboratory
University of Tennessee
104 Ayres Hall
Knoxville TN 37996, USA

Email: moore@cs.utk.edu

## Intellectual Property

PLACEHOLDER for full IETF IPR Statement if needed.

## Full Copyright Statement