

Network Working Group	N. Ward	
Internet-Draft	Braintrust Ltd.	
Intended status: Standards Track	March 03, 2009	
Expires: September 4, 2009		

[TOC](#)

## 6to4 Qualification

**draft-nward-6to4-qualification-00**

### Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

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>.

This Internet-Draft will expire on September 4, 2009.

### Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

### Abstract

A deployment problem exists with existing self-configuring 6to4 implementations making often incorrect assumptions about the state of their IPv4 network connectivity.

This document describes the problem, and proposes a qualification mechanism by which nodes can validate that their connectivity to the global IPv6 network is suitable for use with the 6to4 protocol.

## Requirements Language

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 [RFC 2119 \(Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," March 1997.\)](#) [RFC2119].

---

## Table of Contents

- [1.](#) Introduction
  - [2.](#) Test Nodes
    - [2.1.](#) Symmetric Test Node
    - [2.2.](#) Asymmetric Test Node
  - [3.](#) Qualification
    - [3.1.](#) Initial Interface Configuration
    - [3.2.](#) Symmetric Test
    - [3.3.](#) Asymmetric Test
    - [3.4.](#) Full Configuration
  - [4.](#) Examples
    - [4.1.](#) Symmetric Test Pass
    - [4.2.](#) Symmetric Test Failure
    - [4.3.](#) Asymmetric Test Pass
    - [4.4.](#) Asymmetric Test Failure
  - [5.](#) IANA Considerations
    - [5.1.](#) Symmetric Test Prefix
    - [5.2.](#) Asymmetric Test Prefix
  - [6.](#) Security Considerations
  - [7.](#) Normative References
  - [§](#) Author's Address
- 

## 1. Introduction

[TOC](#)

6to4 as described in [RFC3056 \(Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds," February 2001.\)](#) [RFC3056] and [RFC3068 \(Huitema, C., "An Anycast Prefix for 6to4 Relay Routers," June 2001.\)](#) [RFC3068] is a commonly used method for connecting IPv6 capable networks without IPv6 transit to the global IPv6 network using automatic tunneling over IPv4, and anycasted relay routers. This document refers to a 6to4 router without other IPv6 transit as a 6to4 node.

Many implementations of 6to4 are automatically configured. The assumption is made that if a 6to4 node has a public (i.e. globally routable) IPv4 address, it is capable of 6to4, and the 6to4 interface

and default route via the well known 6to4 relay address is automatically configured. Some of these implementations exist in end systems, and some implementations exist in mass market CPE devices. In reality, however, a public IPv4 address may be subject to IPv4 NAT, or some kind of firewalling which restricts the ability for a 6to4 node to exchange 6to4 packets with other 6to4 nodes and 6to4 relay routers. In this case, 6to4 should not be used. This document defines a mechanism for a 6to4 node to qualify the suitability of its IPv4 network connection for use with 6to4.

---

## 2. Test Nodes

[TOC](#)

There are two new special nodes that MUST connect to the 6to4 network, and respond to ICMPv6 echo-request messages in certain ways. Anyone MAY run these test nodes, their usage scope MAY be limited by filtering advertisements for TBD1 and TBD2.

---

### 2.1. Symmetric Test Node

[TOC](#)

The symmetric test node is a typical node with a 6to4 interface, which MUST have 2002::/16 routed to its 6to4 interface.

The host MUST have connectivity to the global IPv4 network. The TBD prefix MUST be advertised to any networks for which this host is to provide a 6to4 qualification symmetric testing service.

Connectivity to the IPv6 network is not required. Routes to subnets of 2002::/16 MUST NOT exist on the host.

A packet filter MAY be configured to allow only IP protocol 41 (0x29) to and from TBD1(IPv4). A packet filter MAY be configured to prevent any other packets to and from other addresses in the TBD1 prefix. TBD1 SHOULD be routed to the host, and the host MUST null-route this prefix if so.

The host MUST have an IPv4 interface with the address TBD1(IPv4).

TBD1(IPv4) MUST be routed to the host.

The host MUST have a 6to4 interface configured with the IPv6 address TBD1(6to4) and the local IPv4 address TBD1(IPv4).

The host MUST reply to ICMPv6 echo request messages received on the 6to4 interface.

---

[TOC](#)

## 2.2. Asymmetric Test Node

The asymmetric test node is a typical node with a 6to4 interface, which MUST have 2002::/16 routed to its 6to4 interface.

The host MUST have connectivity to the global IPv4 network, and a 6to4-capable IPv4 address.

The host MUST have connectivity to an IPv6 network, including typically at least one 6to4 relay. Routes to subnets of 2002::/16 MUST NOT exist on the host. TBD2 must be advertised to networks with 6to4 relays for which this host is to provide a 6to4 qualification asymmetric testing service.

The host MUST have an IPv6 interface configured with the address TBD2(1).

TBD2 SHOULD be routed to the host, and the host MUST null-route this prefix if so.

The host MUST have a 6to4 interface configured with the IPv6 address corresponding to a locally assigned IPv4 address. The host MUST NOT use 192.88.99.1 as the local IPv4 address for the 6to4 interface.

The host MUST reply to ICMPv6 echo request messages destined to TBD2(1) when the source IPv6 address is in the 2002::/16 prefix.

---

## 3. Qualification

[TOC](#)

Two similar tests are performed by sending ICMPv6 echo request message encapsulated in IPv4, and waiting for a response. This can be trivially implemented in software on most platforms as a low priority process, as it does not require a large amount of processing power.

Sections 3.1 through 3.4 are the 4 steps in the qualification process, and SHALL be performed in order. The qualification process SHOULD be run at regular intervals to ensure reliable 6to4 connectivity.

---

### 3.1. Initial Interface Configuration

[TOC](#)

When a new IPv4 address becomes available that is suspected of being suitable for 6to4 use (candidate IPv4 address), the 6to4 interface is configured with the appropriate address based on the candidate IPv4 address, with a 128-bit prefix length, and the following IPv6 routes:

Destination	Next-hop or connected interface
2002:C058:6301::/48	Connected via 6to4 interface
TBD1(6to4)	Connected via 6to4 interface
TBD2	Next-hop via 2002:C058:6301::

---

### 3.2. Symmetric Test

[TOC](#)

The 6to4 node sends an ICMPv6 echo request message to TBD1(6to4) from its 6to4 IPv6 address. The 6to4 encapsulation will have an IPv4 destination address of TBD1(IPv4), and a source address of the candidate IPv4 address.

A symmetric test node responds to this ICMPv6 message normally.

The 6to4 node receives an ICMPv6 echo response message encapsulated in IPv4 with the candidate IPv4 address as the destination, and TBD1(IPv4) as the source.

If the 6to4 node does not receive this message within a timeout the candidate IPv4 address SHALL be considered unusable and this test run SHALL cease. Failure at this stage MAY mean an IPv4 firewall is in place.

If the 6to4 node receives this message, qualification testing SHOULD proceed to the asymmetric test phase.

---

### 3.3. Asymmetric Test

[TOC](#)

The 6to4 node sends an ICMPv6 echo request message to TBD2 from its 6to4 IPv6 address. The 6to4 encapsulation will have an IPv4 destination address of 192.88.99.1, and a source address of the candidate IPv4 address.

An asymmetric test node responds to this ICMPv6 message normally.

The 6to4 node receives an ICMPv6 echo response message encapsulated in IPv4. The encapsulation MUST have the candidate IPv4 address as the destination address, and an IPv4 address that is not 192.88.99.1 as the source address.

If the 6to4 node does not receive this message within a timeout the candidate IPv4 address SHALL be considered unusable and this test run SHALL cease. Failure at this stage MAY mean a stateful IPv4 firewall and possibly IPv4 NAT is in place.

If the 6to4 node receives this message, the candidate IPv4 address SHOULD be considered usable, and the interface SHOULD be configured as per the Full Configuration section.

---

### 3.4. Full Configuration

[TOC](#)

Once a 6to4 node has an IPv4 address suitable for 6to4 use, the following routes are installed:

Destination	Next-hop or connected interface
2002::/16	Connected via 6to4 interface
::/0	Next-hop via 2002:C058:6301::


The Symmetric Test and Asymmetric Test phases are performed periodically. If either of these tests fail the 2002::/16 route is withdrawn.

## 4. Examples

[TOC](#)


The following are diagrams of example test scenarios.

### 4.1. Symmetric Test Pass

[TOC](#)


- 1) 6to4 node sends ICMPv6 echo request message to TBD1(6to4).
- 2) Symmetric test node responds, and response is received by 6to4 node.


### 4.2. Symmetric Test Failure

[TOC](#)


- 1) 6to4 node sends ICMPv6 echo request message to TBD1(6to4).
- 2) Symmetric test node responds, and response is not received by 6to4 node because an intermediary firewall restriction.


[TOC](#)

### 4.3. Asymmetric Test Pass



- 1) 6to4 node sends ICMPv6 echo request message to TBD2 via 6to4 relay.
- 2) 6to4 relay decapsulates IPv6 packet and forwards to closest asymmetric test node.
- 3) Asymmetric test node responds to ICMPv6 message over 6to4 with source address of its local IPv4 interface, and response is received by 6to4 node.

### 4.4. Asymmetric Test Failure

[TOC](#)

- 1) 6to4 node sends ICMPv6 echo request message to TBD2 via 6to4 relay.

- 2) 6to4 relay decapsulates IPv6 packet and forwards to closest asymmetric test node.
- 3) Asymmetric test node responds to ICMPv6 message over 6to4 with source address of its local IPv4 interface. A stateful firewall prevents the packet reaching 6to4 node.

---

## 5. IANA Considerations

[TOC](#)

This document requests the assignment of two prefixes:

---

### 5.1. Symmetric Test Prefix

[TOC](#)

A 24-bit IPv4 prefix, TBD1. Only one IPv4 address is used, however 24 bits is likely to be widely accepted in BGP peering sessions.

Note to editor: TBD1 is used in two ways in this document, TBD1(IPv4) and TBD1(6to4):

- TBD1(IPv4) is an IPv4 address in this prefix with 1 in the host part, ie AAA.BBB.CCC.1/32
- TBD1(6to4) is an IPv6 address in the 6to4 prefix corresponding to TBD1(6to4) with 1 in the host part, ie. 2002:AABB:CCDD::1

---

### 5.2. Asymmetric Test Prefix

[TOC](#)

A 32-bit IPv6 prefix, TBD2. Only one IPv6 address is used, however 32 bits is likely to be widely accepted in BGP peering sessions.

Note to editor: TBD2(1) refers to an address in this prefix with 1 in the host part, ie. AAAA:BBBB::1.

---

## 6. Security Considerations

[TOC](#)

Unknown at this time.

---

## 7. Normative References

[TOC](#)

- |           |  |
|-----------|--|
| [RFC2119] | <a href="#">Bradner, S.</a> , " <a href="#">Key words for use in RFCs to Indicate Requirement Levels</a> ," BCP 14, RFC 2119, March 1997 ( <a href="#">TXT</a> , <a href="#">HTML</a> , <a href="#">XML</a> ). |
|-----------|--|



[RFC3056]	Carpenter, B. and K. Moore, " <a href="#">Connection of IPv6 Domains via IPv4 Clouds</a> ," RFC 3056, February 2001 ( <a href="#">TXT</a> ).
[RFC3068]	Huitema, C., " <a href="#">An Anycast Prefix for 6to4 Relay Routers</a> ," RFC 3068, June 2001 ( <a href="#">TXT</a> ).

---

## Author's Address

[TOC](#)

	Nathan Ward
	Braintrust Ltd.
	Level 1, 206 Symonds St.
	Auckland, 1010
	NZ
Phone:	+64-21-431675
Email:	<a href="mailto:nward@braintrust.co.nz">nward@braintrust.co.nz</a>