

Survey of IPv4 Addresses in Currently Deployed
IETF General Area Standards

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

Status of this Memo

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/lid-abstracts.txt>

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

Abstract

This document seeks to document all usage of IPv4 addresses in currently deployed IETF General Area documented standards. In order to successfully transition from an all IPv4 Internet to an all IPv6 Internet, many interim steps will be taken. One of these steps is the evolution of current protocols that have IPv4 dependencies. It is hoped that these protocols (and their implementations) will be redesigned to be network address independent, but failing that will at least dually support IPv4 and IPv6. To this end, all Standards (Full, Draft, and Proposed) as well as Experimental RFCs will be surveyed and any dependencies will be documented.

[1.0](#) Introduction

This work began as a megolithic document [draft-ietf-ngtrans-ipv4survey-XX.txt](#). In an effort to rework the information into a more manageable form, it has been broken into 8 documents conforming to the current IETF areas (Application, General, Internet, Management & Operations, Routing, Security, Sub-IP and Transport).

1.1 Short Historical Perspective

There are many challenges that face the Internet Engineering community. The foremost of these challenges has been the scaling issue. How to grow a network that was envisioned to handle thousands of hosts to one that will handle tens of millions of networks with billions of hosts. Over the years this scaling problem has been overcome with changes to the network layer and to routing protocols. (Ignoring the tremendous advances in computational hardware)

The first "modern" transition to the network layer occurred in during the early 1980's from the Network Control Protocol (NCP) to IPv4. This culminated in the famous "flag day" of January 1, 1983. This version of IP was documented in [RFC 760](#). This was a version of IP with 8 bit network and 24 bit host addresses. A year later IP was updated in [RFC 791](#) to include the famous A, B, C, D, & E class system.

Networks were growing in such a way that it was clear that a need for breaking networks into smaller pieces was needed. In October of 1984 [RFC 917](#) was published formalizing the practice of subnetting.

By the late 1980's it was clear that the current exterior routing protocol used by the Internet (EGP) was not sufficient to scale with the growth of the Internet. The first version of BGP was documented in 1989 in [RFC 1105](#).

The next scaling issues to become apparent in the early 1990's was the exhaustion of the Class B address space. The growth and commercialization of the Internet had organizations requesting IP addresses in alarming numbers. In May of 1992 over 45% of the Class B space was allocated. In early 1993 [RFC 1466](#) was published directing assignment of blocks of Class C's be given out instead of Class B's. This solved the problem of address space exhaustion but had significant impact of the routing infrastructure.

The number of entries in the "core" routing tables began to grow exponentially as a result of [RFC 1466](#). This led to the implementation of BGP4 and CIDR prefix addressing. This may have solved the problem for the present but there are still potential scaling issues.

Current Internet growth would have long overwhelmed the current address space if industry didn't supply a solution in Network Address Translators (NATs). To do this the Internet has sacrificed the underlying "End-to-End" principle.

In the early 1990's the IETF was aware of these potential problems and began a long design process to create a successor to IPv4 that would address these issues. The outcome of that process was IPv6.

The purpose of this document is not to discuss the merits or problems of

IPv6. That is a debate that is still ongoing and will eventually be decided on how well the IETF defines transition mechanisms and how industry accepts the solution. The question is not "should," but "when."

[1.2](#) A Brief Aside

Throughout this document there are discussions on how protocols might be updated to support IPv6 addresses. Although current thinking is that IPv6 should suffice as the dominant network layer protocol for the lifetime of the author, it is not unreasonable to contemplate further upgrade to IP. Work done by the IRTF Interplanetary Internet Working Group shows one idea of far reaching thinking. It may be a reasonable idea (or may not) to consider designing protocols in such a way that they can be either IP version aware or independent. This idea must be balanced against issues of simplicity and performance. Therefore it is recommended that protocol designer keep this issue in mind in future designs.

Just as a reminder, remember the words of Jon Postel:

"Be conservative in what you send; be liberal in what you accept from others."

[2.0](#) Methodology

To perform this study each class of IETF standards are investigated in order of maturity: Full, Draft, and Proposed, as well as Experimental. Informational RFC are not addressed. RFCs that have been obsoleted by either newer versions or as they have transitioned through the standards process are not covered.

Please note that a side effect of this choice of methodology is that some protocols that are defined by a series of RFC's that are of different levels of standards maturity are covered in different spots in the document. Likewise other natural groupings (i.e. MIBs, SMTP extensions, IP over FOO, PPP, DNS, etc.) could easily be imagined.

[2.1](#) Scope

The procedure used in this investigation is an exhaustive reading of the applicable RFC's. This task involves reading approximately 25000 pages of protocol specifications. To compound this, it was more than a process of simple reading. It was necessary to attempt to understand the purpose and functionality of each protocol in order to make a proper determination of IPv4 reliability. The author has made every effort to make this effort and the resulting document as complete as possible, but it is likely that some subtle (or perhaps not so subtle) dependence was missed. The author encourage those familiar (designers, implementers or anyone who has an

intimate knowledge) with any protocol to review the appropriate sections and make comments.

[2.2](#) Document Organization

The rest of the document sections are described below.

Sections [3](#), [4](#), [5](#), and [6](#) each describe the raw analysis of Full, Draft, and Proposed Standards, and Experimental RFCs. Each RFC is discussed in its turn starting with [RFC 1](#) and ending with [RFC 3247](#). The comments for each RFC is "raw" in nature. That is, each RFC is discussed in a vacuum and problems or issues discussed do not "look ahead" to see if the problems have already been fixed.

[Section 7](#) is an analysis of the data presented in Sections [3](#), [4](#), [5](#), and [6](#). It is here that all of the results are considered as a whole and the problems that have been resolved in later RFCs are correlated.

[3.0](#) Full Standards

Full Internet Standards (most commonly simply referred to as "Standards") are fully mature protocol specification that are widely implemented and used throughout the Internet.

[3.1](#) [RFC 2600](#) INTERNET OFFICIAL PROTOCOL STANDARDS

Although this is classified as a standard, it does not describe a protocol. It is a list of assigned protocol numbers and therefore has no IPv6 transition issues.

[3.2](#) [RFC 1700](#) Assigned Numbers

Although this is classified as a standard, it does not describe a protocol. It is a list of assigned protocol numbers and therefore has no IPv6 transition issues.

[3.3](#) Host Requirements. [RFC1122](#), [RFC1123](#)

[3.3.1](#) [RFC 1122](#)

[RFC 1122](#) defines requirements for Internet hosts. In [Section 1.1.3](#) (Internet Protocol Suite), the section on layer 3 (Internet Layer) mandates hosts implement IP, but does not specify a version requirement.

[Section 3](#) is devoted to a discussion of IP, ICMP, and IGMP and is riddled with specific IPv4 requirements. Any IPv6 only host would be non-compliant with [RFC 1122](#). Some of the most obvious problems are shown below.

In section "3.2.1.1 Version Number" it says: 'A datagram whose version number is not 4 MUST be silently discarded.'

In section "3.2.1.3 Addressing" is clearly out of date even with the current addressing of IPv4 addresses.

A new version of [RFC 1122](#) should be written. It should either be IPv6 focused (as the current [RFC 1122](#) is IPv4 focused) and be labeled as such (i.e. "IPv6 Host Requirements") or should be written generically with appropriate external links. The later would be difficult since the document is meant to be self-contained, so the former is a more likely solution.

[3.3.2 RFC 1123](#)

[Section 2.1](#) "Host Names and Numbers" makes references to applications making explicit references to "dotted decimal" notation and the form "#.#.#.#"

[Section 5.2.17](#) "Domain Literals:" says "A mailer MUST be able to accept and parse an Internet domain literal whose content ("dtext"; see [RFC-822](#)) is a dotted decimal host address."

There are also many references to IP addresses scattered throughout the document that make no reference to the format or version of the addresses.

Most of this document (as well as [RFC 1122](#)) is a series of references and consolidation of data from numerous RFCs. The few references to IPv4 addresses might not warrant a rewrite. However the technology of the Internet has changed substantially in the last 11 years. With a requirement of rewriting [RFC 1122](#) it makes sense to update this document for other reasons, not IPv6 related.

[RFC 2181](#) is considered to be an "Update" to [RFC1123](#) but is only related to DNS issues and does not fix the problems pointed out here.

[3.4 RFC 1221](#) Host Access Protocol specification

There are no IPv4 dependencies in this protocol.

[4.0](#) Draft Standards

Draft Standards represent the penultimate standard level in the IETF. A protocol can only achieve draft standard when there are multiple, independent, interoperable implementations. Draft Standards are usually quite mature and widely used.

[5.0](#) Proposed Standards

Proposed Standards are introductory level documents. There are no requirements for even a single implementation. In many cases Proposed are never implemented or advanced in the IETF standards process. They therefore are often just proposed ideas that are presented to the Internet community. Sometimes flaws are exposed or they are one of many competing solutions to problems. In these later cases, no discussion is presented as it would not serve the purpose of this discussion.

[5.1](#) [RFC 1812](#) Requirements for IP Version 4 Routers

This document is only intended to describe requirements for IPv4 implementations in routers and is not discussed in this document.

[6.0](#) Experimental RFCs

Experimental RFCs typically define protocols that do not have widescale implementation or usage on the Internet. They are often propriety in nature or used in limited arenas. They are documented to the Internet community in order to allow potential interoperability or some other potential useful scenario. In a few cases they are presented as alternatives to the mainstream solution to an acknowledged problem.

[7.0](#) Summary of Results

In the initial survey of RFCs 0 positives were identified out of a total of 7, broken down as follows:

Standards	0 of 6 or 0.00%
Draft Standards	0 of 0
Proposed Standards	0 of 1 or 0.00%
Experimental RFCs	0 of 0

Of those identified many require no action because they document outdated and unused protocols, while others are document protocols that are actively being updated by the appropriate working groups. Additionally there are many instances of standards that SHOULD be updated but do not cause any operational impact if they are not

updated. The remaining instances are documented below.

The author has attempted to organize the results in a format that allows easy reference to other protocol designers. The following recommendations uses the documented terms "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" described in [RFC 2119](#). They should only be interpreted in the context of [RFC 2119](#) when they appear in all caps. That is, the word "should" in the previous SHOULD NOT be interpreted as in [RFC 2119](#).

The assignment of these terms has been based entirely on the authors perceived needs for updates and should not be taken as an official statement.

[7.1](#) Standards

[7.2](#) Draft Standards

[7.3](#) Proposed Standards

[7.4](#) Experimental RFCs

[8.0](#) Acknowledgements

The author would like to acknowledge the support of the Internet Society in the research and production of this document. Additionally the author would like to thanks his partner in all ways, Wendy M. Nesser.

[9.0](#) Authors Address

Please contact the author with any questions, comments or suggestions at:

Philip J. Nesser II
Principal
Nesser & Nesser Consulting
[13501](#) 100th Ave NE, #5202
Kirkland, WA 98034

Email: phil@nesser.com

Phone: +1 425 481 4303

Fax: +1 425 48