Introduction to the Survey of IPv4 Addresses in
Currently Deployed IETF Standards Track and Experimental Documents

Status of this Memo

   This memo provides information for the Internet community. It does
   not specify an Internet standard of any kind. Distribution of this
   memo is unlimited.

Copyright Notice

Abstract

   This document is a general overview and introduction to the v6ops
   IETF workgroup project of documenting all usage of IPv4 addresses in
   IETF standards track and experimental RFCs. It is broken into seven
   documents conforming to the current IETF areas. It also describes
   the methodology used during documentation, which types of RFCs have
   been documented, and provides a concatenated summary of results.

Table of Contents

1.0. Introduction

   This document is the introduction to a document set aiming to
   document all usage of IPv4 addresses in IETF standards. In an effort
   to have the information in a manageable form, it has been broken into
   7 documents, conforming to the current IETF areas (Application [1],
   Internet [2], Operations and Management [3], Routing [4], Security
   [5], Sub-IP [6], and Transport [7]). It also describes the
   methodology used during documentation, which types of RFCs that have
   been documented, and provides a concatenated summary of results.

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
   managed, with varying degrees of success, by changes to the network
   layer and to routing protocols. (Although largely ignored in the
   changes to network layer and routing protocols, the tremendous
   advances in computational hardware during the past two decades have
   been of significant benefit in management of scaling problems
   encountered thus far.)

   The first "modern" transition to the network layer occurred during
   the early 1980's, moving from the Network Control Protocol (NCP) to
   IPv4. This culminated in the famous "flag day" of January 1, 1983.
   IP Version 4 originally specified an 8 bit network and 24 bit host
   addresses, as documented in RFC 760. A year later, IPv4 was updated
   in RFC 791 to include the famous A, B, C, D, and E class system.

   Networks were growing in such a way that it was clear that a
   convention 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 insufficiently robust to
   scale with the growth of the Internet. The first version of BGP was
   documented in 1989 in RFC 1105.

Yet another scaling issue, exhaustion of the class B address space
became apparent in the early 1990s. The growth and commercialization
of the Internet stimulated organisations requesting IP addresses in
alarming numbers. By May of 1992, over 45% of the Class B space had
been 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 temporarily circumvented the problem of address space
exhaustion, but had a 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
circumvented the problem for the present, but they continue to pose
potential scaling issues.

Growth in the population of Internet hosts since the mid-1980s would
have long overwhelmed the IPv4 address space if industry had not
supplied a circumvention in the form of Network Address Translators
(NATs). To do this, the Internet has watered down 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 debate 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. An Observation on the Classification of Standards

It has become clear during the course of this investigation that
there has been little management of the status of standards over the
years. Some attempt has been made by the introduction of the
classification of standards into Full, Draft, Proposed, Experimental,
and Historic. However, there has not been a concerted effort to
actively manage the classification for older standards. Standards
are only classified as Historic when either a newer version of the
protocol is deployed and it is randomly noticed that an RFC describes
a long dead protocol, or a serious flaw is discovered in a protocol.
Another issue is the status of Proposed Standards. Since this is the
entry level position for protocols entering the standards process,
many old protocols or non-implemented protocols linger in this status
indefinitely. This problem also exists for Experimental RFCs.
Similarly, the problem exists for the Best Current Practices (BCP)
and For You Information (FYI) series of documents.

To exemplify this point, there are 61 Full Standards, only 4 of which have been reclassified to Historic. There are 65 Draft Standards, 611 Proposed Standards, and 150 Experimental RFCs, of which only 66 have been reclassified as Historic. That is a rate of less than 8%. It should be obvious that in the more that 30 years of protocol development and documentation, there should be at least as many (if not a majority of) protocols that have been retired compared to the ones that are currently active.

Please note that there is occasionally some confusion of the meaning of a "Historic" classification. It does NOT necessarily mean that the protocol is not being used. A good example of this concept is the Routing Information Protocol(RIP) version 1. There are many thousands of sites using this protocol even though it has Historic status. There are potentially hundreds of otherwise classified RFC's that should be reclassified.

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 and BCP RFCs are not addressed. RFCs that have been obsoleted by either newer versions or because they have transitioned through the standards process are not covered. RFCs which have been classified as Historic are also not included.

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 25,000 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 produce as complete a document set as possible, but it is likely that some subtle (or perhaps not so subtle) dependence was missed. The author encourages those familiar (designers, implementers or anyone who has an intimate knowledge) with any protocol to review the appropriate sections and make comments.

3.0. Summary of Results

   In the initial survey of RFCs, 173 positives were identified out of a
   total of 877, broken down as follows:

         Standards: 30 out of 68 or 44.12%
         Draft Standards: 16 out of 68 or 23.53%
         Proposed Standards: 98 out of 597 or 16.42%
         Experimental RFCs: 29 out of 144 or 20.14%

   Of those identified, many require no action because they document
   outdated and unused protocols, while others are active document
   protocols being updated by the appropriate working groups (SNMP MIBs
   for example).

   Additionally, there are many instances of standards that should be
   updated but do not cause any operational impact (STD 3/RFCs 1122 and
   1123 for example) if they are not updated.

   In this statistical survey, a positive is defined as a RFC containing
   an IPv4 dependency, regardless of context.

3.1. Application Area Specifications

   In the initial survey of RFCs, 34 positives were identified out of a
   total of 257, broken down as follows:

         Standards: 1 out of 20 or 5.00%
         Draft Standards: 4 out of 25 or 16.00%
         Proposed Standards: 19 out of 155 or 12.26%
         Experimental RFCs: 10 out of 57 or 17.54%

   For more information, please look at [1].

3.2. Internet Area Specifications

   In the initial survey of RFCs, 52 positives were identified out of a
   total of 186, broken down as follows:

         Standards: 17 out of 24 or 70.83%
         Draft Standards: 6 out of 20 or 30.00%
         Proposed Standards: 22 out of 111 or 19.91%
         Experimental RFCs: 7 out of 31 or 22.58%

   For more information, please look at [2].

3.3. Operations and Management Area Specifications

   In the initial survey of RFCs, 36 positives were identified out of a
   total of 153, broken down as follows:

         Standards: 6 out of 15 or 40.00%
         Draft Standards: 4 out of 15 or 26.67%
         Proposed Standards: 26 out of 112 or 23.21%
         Experimental RFCs: 0 out of 11 or 0.00%

   For more information, please look at [3].

3.4. Routing Area Specifications

   In the initial survey of RFCs, 23 positives were identified out of a
   total of 46, broken down as follows:

         Standards: 3 out of 3 or 100.00%
         Draft Standards: 1 out of 3 or 33.33%
         Proposed Standards: 13 out of 29 or 44.83%
         Experimental RFCs: 6 out of 11 or 54.54%

   For more information, please look at [4].

3.5. Security Area Specifications

   In the initial survey of RFCs, 4 positives were identified out of a
   total of 124, broken down as follows:

         Standards: 0 out of   1 or 0.00%
         Draft Standards: 1 out of   3 or 33.33%
         Proposed Standards: 1 out of 102 or 0.98%
         Experimental RFCs: 2 out of 18 or 11.11%

   For more information, please look at [5].

3.6. Sub-IP Area Specifications

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

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

   For information about the Sub-IP Area standards, please look at [6].

3.7. Transport Area Specifications

   In the initial survey of RFCs, 24 positives were identified out of a
   total of 104, broken down as follows:

              Standards: 3 out of 5 or 60.00%
              Draft Standards: 0 out of 2 or 0.00%
              Proposed Standards: 17 out of 82 or 20.73%
              Experimental RFCs: 4 out of 15 or 26.67%

   For more information, please look at [7].

4.0. Discussion of "Long Term" Stability of Addresses on Protocols

   In attempting this analysis, it was determined that a full scale
   analysis is well beyond the scope of this document. Instead, a short
   discussion is presented on how such a framework might be established.

   A suggested approach would be to do an analysis of protocols based on
   their overall function, similar (but not strictly) to the OSI network
   reference model. It might be more appropriate to frame the
   discussion in terms of the different Areas of the IETF.

   The problem is fundamental to the overall architecture of the
   Internet and its future. One of the stated goals of the IPng (now
   IPv6) was "automatic" and "easy" address renumbering. An additional
   goal is "stateless autoconfiguration." To these ends, a substantial
   amount of work has gone into the development of such protocols as
   DHCP and Dynamic DNS. This goes against the Internet age-old "end-
   to-end principle."

   Most protocol designs implicitly count on certain underlying
   principles that currently exist in the network. For example, the
   design of packet switched networks allows upper level protocols to
   ignore the underlying stability of packet routes. When paths change
   in the network, the higher level protocols are typically unaware and
   uncaring. This works well since whether the packet goes A-B-C-D-E-F
   or A-B-X-Y-Z-E-F is of little consequence.

   In a world where endpoints (i.e., A and F in the example above)
   change at a "rapid" rate, a new model for protocol developers should
   be considered. It seems that a logical development would be a change
   in the operation of the Transport layer protocols. The current model
   is essentially a choice between TCP and UDP, neither of which
   provides any mechanism for an orderly handoff of the connection if
   and when the network endpoint (IP) addresses change. Perhaps a third

major transport layer protocol should be developed, or perhaps
updated TCP and UDP specifications that include this function might
be a better solution.

There are many, many variables that would need to go into a
successful development of such a protocol. Some issues to consider
are: timing principles; overlap periods as an endpoint moves from
address A, to addresses A and B (answers to both), to only B; delays
due to the recalculation of routing paths, etc...

5.0. Security Considerations

This memo examines the IPv6-readiness of specifications; this does
not have security considerations in itself.

6.0. Acknowledgements

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

The editor, Andreas Bergstrom, would like to thank Pekka Savola for
guidance and collection of comments for the editing of this document.
He would further like to thank Alan E. Beard, Jim Bound, Brian
Carpenter, and Itojun for valuable feedback on many points of this
document.

7.0. References

7.1. Normative References

   [1] Sofia, R. and P. Nesser, II, "Survey of IPv4 Addresses in
        Currently Deployed IETF Application Area Standards", RFC 3795,
        June 2004.

   [2] Mickles, C., Editor and P. Nesser, II, "Survey of IPv4 Addresses
        in Currently Deployed IETF Internet Area Standards", RFC 3790,
        June 2004.

   [3] Nesser, II, P. and A. Bergstrom, Editor, "Survey of IPv4
        Addresses in Currently Deployed IETF Operations & Management
        Area Standards", RFC 3796, June 2004.

   [4] Olvera, C. and P. Nesser, II, "Survey of IPv4 Addresses in
        Currently Deployed IETF Routing Area Standards", RFC 3791, June
        2004.

   [5] Nesser, II, P. and A. Bergstrom, Editor, "Survey of IPv4
        Addresses in Currently Deployed IETF Security Area Standards",
        RFC 3792, June 2004.

   [6] Nesser, II, P. and A. Bergstrom, Editor, "Survey of IPv4
        Addresses in Currently Deployed IETF Sub-IP Area Standards", RFC
        3793, June 2004.

   [7] Nesser, II, P. and A. Bergstrom, Editor, "Survey of IPv4
        Addresses in Currently Deployed IETF Transport Area Standards",
        RFC 3794, June 2004.

8.0. Authors' Addresses

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

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

   Phone: +1 425 481 4303
   Fax: +1 425 482 9721
   EMail: phil@nesser.com

   Andreas Bergstrom, Editor
   Ostfold University College
   Rute 503 Buer
   N-1766 Halden
   Norway

   EMail: andreas.bergstrom@hiof.no

9.0. Full Copyright Statement

Intellectual Property

Acknowledgement