Network Working Group
Request for Comments: 3791
Category: Informational

C. Olvera
Consulintel
P. Nesser, II
Nesser & Nesser Consulting

June 2004

Survey of IPv4 Addresses in Currently Deployed IETF Routing Area Standards Track and Experimental Documents

Status of this Memo

Copyright Notice

Abstract

This investigation work seeks to document all usage of IPv4 addresses in currently deployed IETF Routing 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.

Table of Contents

1. Introduction

This work aims to document all usage of IPv4 addresses in currently deployed IETF Routing Area documented standards. Also, throughout this document there are discussions on how routing protocols might be updated to support IPv6 addresses.

This material was originally presented within a single document, but in an effort to have the information in a manageable form, it has subsequently been split into 7 documents conforming to the current IETF main areas (Application [2], Internet [3], Operations & Management [4], Routing [this document], Security [5], Sub-IP [6] and Transport [7]).

The general overview, methodology used during documentation and scope of the investigation for the whole 7 documents can be found in the introduction of this set of documents [1].

It is important to mention that to perform this study the following classes of IETF standards are investigated: Full, Draft, and Proposed, as well as Experimental. Informational, BCP and Historic RFCs are not addressed. RFCs that have been obsoleted by either newer versions or as they have transitioned through the standards process are also not covered.

2. Document Organization

The main Sections of this document are described below.

Sections 3, 4, 5, and 6 each describe the raw analysis of Full, Draft, Proposed Standards and Experimental RFCs. Each RFC is discussed in its turn starting with RFC 1 and ending (around) RFC 3100. The comments for each RFC are "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. 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 1722 (STD 57) RIP Version 2 Protocol Applicability Statement

RIPv2 is only intended for IPv4 networks.

3.2. RFC 2328 (STD 54) OSPF Version 2

This RFC defines a protocol for IPv4 routing. It is highly assumptive about address formats being IPv4 in nature.

3.3. RFC 2453 (STD 56) RIP Version 2

RIPv2 is only intended for IPv4 networks.

4. 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.

4.1. RFC 1771 A Border Gateway Protocol 4 (BGP-4)

This RFC defines a protocol used for exchange of IPv4 routing information and does not support IPv6 as is defined.

4.2. RFC 1772 Application of the Border Gateway Protocol in the Internet

This RFC is a discussion of the use of BGP-4 on the Internet.

4.3. RFC 3392 Capabilities Advertisement with BGP-4

Although the protocol enhancements have no IPv4 dependencies, the base protocol, BGP-4, is IPv4 only.

5. 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 1195 Use of OSI IS-IS for routing in TCP/IP and dual environments

This document specifies a protocol for the exchange of IPv4 routing information.

5.2. RFC 1370 Applicability Statement for OSPF

This document discusses a version of OSPF that is limited to IPv4.

5.3. RFC 1397 Default Route Advertisement In BGP2 and BGP3 Version of The Border Gateway Protocol

BGP2 and BGP3 are both deprecated and therefore are not discussed in this document.

5.4. RFC 1478 An Architecture for Inter-Domain Policy Routing

The architecture described in this document has no IPv4 dependencies.

5.5. RFC 1479 Inter-Domain Policy Routing Protocol Specification: Version 1 (IDPR)

There are no IPv4 dependencies in this protocol.

5.6. RFC 1517 Applicability Statement for the Implementation of Classless Inter-Domain Routing (CIDR)

This document deals exclusively with IPv4 addressing issue.

5.7. RFC 1518 An Architecture for IP Address Allocation with CIDR

This document deals exclusively with IPv4 addressing issue.

5.8. RFC 1519 Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy

This document deals exclusively with IPv4 addressing issue.

5.9. RFC 1582 Extensions to RIP to Support Demand Circuits

This protocol is an extension to a protocol for exchanging IPv4 routing information.

5.10. RFC 1584 Multicast Extensions to OSPF

This document defines the use of IPv4 multicast to an IPv4 only routing protocol.

5.11. RFC 1793 Extending OSPF to Support Demand Circuits

There are no IPv4 dependencies in this protocol other than the fact that it is a new functionality for a routing protocol that only supports IPv4 networks.

5.12. RFC 1997 BGP Communities Attribute

Although the protocol enhancements have no IPv4 dependencies, the base protocol, BGP-4, is IPv4 only.

5.13. RFC 2080 RIPng for IPv6

This RFC documents a protocol for exchanging IPv6 routing information and is not discussed in this document.

5.14. RFC 2091 Triggered Extensions to RIP to Support Demand Circuits

This RFC defines an enhancement for an IPv4 routing protocol and while it has no IPv4 dependencies it is inherently limited to IPv4.

5.15. RFC 2338 Virtual Router Redundancy Protocol (VRRP)

This protocol is IPv4 specific, there are numerous references to 32-bit IP addresses.

5.16. RFC 2370 The OSPF Opaque LSA Option

There are no IPv4 dependencies in this protocol other than the fact that it is a new functionality for a routing protocol that only supports IPv4 networks.

5.17. RFC 2439 BGP Route Flap Damping

The protocol enhancements have no IPv4 dependencies, even though the base protocol, BGP-4, is IPv4 only routing protocol.

5.18. RFC 2545 Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing

This RFC documents IPv6 routing methods and is not discussed in this document.

5.19. RFC 2740 OSPF for IPv6

This document defines an IPv6 specific protocol and is not discussed in this document.

5.20. RFC 2784 Generic Routing Encapsulation (GRE)

This protocol is only defined for IPv4. The document states in the Appendix:

- o IPv6 as Delivery and/or Payload Protocol

This specification describes the intersection of GRE currently deployed by multiple vendors. IPv6 as delivery and/or payload protocol is not included.

5.21. RFC 2796 BGP Route Reflection - An Alternative to Full Mesh IBGP

Although the protocol enhancements have no IPv4 dependencies, the base protocol, BGP-4, is IPv4 only routing protocol. This specification updates but does not obsolete RFC 1966.

5.22. RFC 2858 Multiprotocol Extensions for BGP-4

In the Abstract:

Currently BGP-4 is capable of carrying routing information only for IPv4. This document defines extensions to BGP-4 to enable it to carry routing information for multiple Network Layer protocols (e.g., IPv6, IPX, etc...). The extensions are backward compatible - a router that supports the extensions can interoperate with a router that doesn’t support the extensions.

The document is therefore not examined further in this document.

5.23. RFC 2890 Key and Sequence Number Extensions to GRE

There are no IPv4 dependencies in this protocol.

5.24. RFC 2894 Router Renumbering for IPv6

The RFC defines an IPv6 only document and is not concerned in this survey.

5.25. RFC 2918 Route Refresh Capability for BGP-4

Although the protocol enhancements have no IPv4 dependencies, the base protocol, BGP-4, is IPv4 only routing protocol.

5.26. RFC 3065 Autonomous System Confederations for BGP

Although the protocol enhancements have no IPv4 dependencies, the base protocol, BGP-4, is IPv4 only routing protocol.

5.27. RFC 3101 The OSPF Not-So-Stubby Area (NSSA) Option

This document defines an extension to an IPv4 routing protocol.

5.28. RFC 3107 Carrying Label Information in BGP-4

There are no IPv4 dependencies in this protocol.

5.29. RFC 3122 Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification

This is an IPv6 related document and is not discussed in this document.

6. Experimental RFCs

Experimental RFCs typically define protocols that do not have wide scale 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.

6.1. RFC 1075 Distance Vector Multicast Routing Protocol (DVMRP)

This document defines a protocol for IPv4 multicast routing.

6.2. RFC 1383 An Experiment in DNS Based IP Routing

This proposal is IPv4 limited:

This record is designed for easy general purpose extensions in the DNS, and its content is a text string. The RX record will contain three fields: A record identifier, A cost indicator, and An IP address.

The three strings will be separated by a single comma. An example of record would thus be:

domain	type	record	value
*.27.32.192.in-addr.arpa	IP	TXT	RX, 10, 10.0.0.7

which means that for all hosts whose IP address starts by the three octets "192.32.27" the IP host "10.0.0.7" can be used as a gateway, and that the preference value is 10.

6.3. RFC 1476 RAP: Internet Route Access Protocol

This document defines an IPv7 routing protocol and has been abandoned by the IETF as a feasible design. It is not considered in this document.

6.4. RFC 1765 OSPF Database Overflow

There are no IPv4 dependencies in this protocol other than the fact that it is a new functionality for a routing protocol that only supports IPv4 networks.

6.5. RFC 1863 A BGP/IDRP Route Server alternative to a full mesh routing

This protocol is both IPv4 and IPv6 aware and needs no changes.

6.6. RFC 1966 BGP Route Reflection An alternative to full mesh IBGP

Although the protocol enhancements have no IPv4 dependencies, the base protocol, BGP-4, is IPv4 only routing protocol. This specification has been updated by RFC 2796.

6.7. RFC 2189 Core Based Trees (CBT version 2) Multicast Routing

The document specifies a protocol that depends on IPv4 multicast. There are many packet formats defined that show IPv4 usage.

6.8. RFC 2201 Core Based Trees (CBT) Multicast Routing Architecture

See previous Section for the IPv4 limitation in this protocol.

6.9. RFC 2337 Intra-LIS IP multicast among routers over ATM using Sparse Mode PIM

This protocol is designed for IPv4 multicast.

6.10. RFC 2362 Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification

This protocol is both IPv4 and IPv6 aware and needs no changes.

6.11. RFC 2676 QoS Routing Mechanisms and OSPF Extensions

There are IPv4 dependencies in this protocol. It requires the use of the IPv4 TOS header field.

7. Summary of Results

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%

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 authors have attempted to organize the results in a format that allows easy reference to other protocol designers. The assignment of statements has been based entirely on the authors perceived needs for updates and should not be taken as an official statement.

7.1. Standards

7.1.1. STD 57 RIP Version 2 Protocol Applicability Statement (RFC 1722)

This problem has been fixed by RFC 2081, RIPng Protocol Applicability Statement.

7.1.2. STD 54 OSPF Version 2 (RFC 2328)

This problem has been fixed by RFC 2740, OSPF for IPv6.

7.1.3. STD 56 RIP Version 2 (RFC 2453)

This problem has been fixed by RFC 2080, RIPng for IPv6.

7.2. Draft Standards

7.2.1. Border Gateway Protocol 4 (RFC 1771)

This problem has been fixed in RFC 2858 Multiprotocol Extensions for BGP-4, RFC 2545 Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing, and in [8].

RFC 2858 extends BGP to support multi-protocol extensions that allows routing information for other address families to be exchanged. RFC 2545 further extends RFC 2858 for full support of exchanging IPv6 routing information and additionally clarifies support of the extended BGP-4 protocol using TCP+IPv6 as a transport mechanism. RFC 1771, 2858 & 2545 must be supported in order to provide full IPv6 support.

Note also that all the BGP extensions analyzed previously in this memo function without changes with the updated version of BGP-4.

7.3. Proposed Standards

7.3.1. Use of OSI IS-IS for routing in TCP/IP and dual environments (RFC 1195)

This problem is being addressed by the IS-IS WG [9].

7.3.2. Applicability Statement for OSPFv2 (RFC 1370)

This problem has been resolved in RFC 2740, OSPF for IPv6.

7.3.3. Applicability of CIDR (RFC 1517)

The contents of this specification has been treated in various IPv6 addressing architecture RFCs, see RFC 3513 & 3587.

7.3.4. CIDR Architecture (RFC 1518)

The contents of this specification has been treated in various IPv6 addressing architecture RFCs, see RFC 3513 & 3587.

7.3.5. Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy (RFC 1519)

The contents of this specification has been treated in various IPv6 addressing architecture RFCs, see RFC 3513 & 3587.

7.3.6. RIP Extensions for Demand Circuits (RFC 1582)

This problem has been addressed in RFC 2080, RIPng for IPv6.

7.3.7. OSPF Multicast Extensions (RFC 1584)

This functionality has been covered in RFC 2740, OSPF for IPv6.

7.3.8. OSPF For Demand Circuits (RFC 1793)

This functionality has been covered in RFC 2740, OSPF for IPv6.

7.3.9. RIP Triggered Extensions for Demand Circuits (RFC 2091)

This functionality is provided in RFC 2080, RIPng for IPv6.

7.3.10. Virtual Router Redundancy Protocol (VRRP)(RFC 2338)

The problems identified are being addressed by the VRRP WG [10].

7.3.11. OSPF Opaque LSA Option (RFC 2370)

This problem has been fixed by RFC 2740, OSPF for IPv6. Opaque options support is an inbuilt functionality in OSPFv3.

7.3.12. Generic Routing Encapsulation (GRE)(RFC 2784)

Even though GRE tunneling over IPv6 has been implemented and used, its use has not been formally specified. Clarifications are required.

7.3.13. OSPF NSSA Option (RFC 3101)

This functionality has been covered in RFC 2740, OSPF for IPv6.

7.4. Experimental RFCs

7.4.1. Distance Vector Multicast Routing Protocol (RFC 1075)

This protocol is a routing protocol for IPv4 multicast routing. It is no longer in use and need not be redefined.

7.4.2. An Experiment in DNS Based IP Routing (RFC 1383)

This protocol relies on IPv4 DNS RR, but is no longer relevant has never seen much use; no action is necessary.

7.4.3. Core Based Trees (CBT version 2) Multicast Routing (RFC 2189)

This protocol relies on IPv4 IGMP Multicast and a new protocol standard may be produced. However, the multicast routing protocol has never been in much use and is no longer relevant; no action is necessary.

7.4.4. Core Based Trees (CBT) Multicast Routing Architecture (RFC 2201)

See previous Section for the limitation in this protocol.

7.4.5. Intra-LIS IP multicast among routers over ATM using Sparse Mode PIM (RFC 2337)

This protocol is designed for IPv4 multicast. However, Intra-LIS IP multicast among routers over ATM is not believed to be relevant anymore. A new mechanism may be defined for IPv6 multicast.

7.4.6. QoS Routing Mechanisms and OSPF Extensions (RFC 2676)

QoS extensions for OSPF were never used for OSPFv2, and there seems to be little need for them in OSPFv3.

However, if necessary, an update to this document could simply define the use of the IPv6 Traffic Class field since it is defined to be exactly the same as the IPv4 TOS field.

8. Security Considerations

This document examines the IPv6-readiness of routing specification; this does not have security considerations in itself.

9. Acknowledgements

The original author, Philip J. Nesser II, would like to acknowledge the support of the Internet Society in the research and production of this document.

He also would like to thanks his partner in all ways, Wendy M. Nesser.

Cesar Olvera would like to thanks Pekka Savola for an extended guidance and comments for the edition of this document, and Jordi Palet for his support and reviews.

Additionally, he would further like to thank Andreas Bergstrom, Brian Carpenter, Jeff Haas, Vishwas Manral, Gabriela Medina, Venkata Naidu, Jeff Parker and Curtis Villamizar for valuable feedback.

10. References

10.1. Normative References

[1] Nesser, II, P. and A. Bergstrom, Editor, "Introduction to the Survey of IPv4 Addresses in Currently Deployed IETF Standards", RFC 3789, June 2004.

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

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

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

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

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

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

10.2. Informative References

[8] Chen, E. and J. Yuan, "AS-wide Unique BGP Identifier for BGP-4", Work in Progress, December 2003.

[9] Hopps, C., "Routing IPv6 with IS-IS", Work in Progress, January 2003.

[10] Hinden, R., "Virtual Router Redundancy Protocol for IPv6", Work in Progress, February 2004.

11. Authors' Addresses

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

Cesar Olvera Morales
Researcher
Consulintel
San Jose Artesano, 1
28108 - Alcobendas
Madrid, Spain

Phone: +34 91 151 81 99
Fax: +34 91 151 81 98
EMail: cesar.olvera@consulintel.es

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

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

12. Full Copyright Statement

Intellectual Property

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.