IPng Transition
Internet Draft

Dan Harrington
Digital Equipment Corp.
Quaizar Vohra
University of New Hampshire
November 1996

Limiting the role of IPv4-compatible Addresses in IPv6

draft-harrington-ngtrans-v4comp-00.txt

Abstract

This draft presents a proposal to limit IPv4-compatible IPv6 addresses to tunnelling interfaces in the transition from IPv4 to IPv6. The reasons and context for restricting the usage in this manner will be presented.

Table Of Contents

1. Introduction

IPv4-compatible addresses are designed to ease the transition of IPv4 to IPv6, by utilizing the readily available IPv4 address space and protocols to provide IPv6 connectivity. They currently serve two roles, both related to tunnelling:

- To allow isolated IPv6 nodes to come up on the Internet and communicate with other IPv6 nodes via automatic tunneling, which requires a minimal amount of configuration.

- Identifying an IPv6 router's next-hop interface address over a manually configured tunnel.

These tasks both require implementations to treat an IPv4 tunnel as a pseudo-NBMA link, where ::/96 is treated as an on-link IPv6 prefix for the tunnel interface. In this model, all IPv4-compatible addresses are on-link to the tunnel interface and the IPv4 Internet forms one large link layer, in which address resolution is a trivial function. Manually configured tunnels are used with static routes to IPv6 prefixes, where the next-hop is an IPv4-compatible address on the link. While this link type does not use the standard link-local prefix of FE80:: or Neighbor Discovery protocols, it does have its own characteristics and rules [V6TUNNELS]. Conceptually, then, it can be seen that IPv6 packets using IPv4-compatible addresses could be treated as using a special type of link-local address, and the Hop Limit could be set to a value of 1 with no dire consequences.

The current Transition Mechanisms specification [RFC1933], however, also include a provision to allow an IPv4-compatible address to be assigned to an interface for native IPv6 communications, with all the requirements of Neighbor Discovery. It is this usage which we wish to prohibit, for the sake of reduced complexity and increased interoperability.

2. Architectural and Philosophical Issues

Although IPv4 and IPv6 represent different network protocols, IPv4 addresses can be represented as IPv6 addresses. However, they still define an IPv4 endpoint, that is, an interface on a link connected to an IPv4 network, using IPv4 protocols. Using them in multiple fashions, for both IPv4 and IPv6 packets on a given interface as well as for tunnelling, can and will lead to interoperability problems, as has been reported on the NGTRANS mailing list [NGLIST]. This dual usage also leads to unnecessary implementation complexity; for example, the source address selection algorithm should not permit the use of an IPv4-compatible address (as source or destination) with a global IPv6 address (as destination or source).

As mentioned above, the encapsulation of IPv6 packets in IPv4 packets essentially uses the IPv4 network as a specialized media type. The "Generic Packet Tunneling in IPv6" [V6TUNNELS] specification gives the mechanism by which one protocol may be run over another. In keeping with the general IP philosophy of an

address being associated with a particular interface [RFC1122], it should be held that a tunnel interface is not merely an abstraction, but a "real" interface to a specific media type, with its own rules and behaviours.

Finally, restricting the usage of IPv4-compatible addresses will simplify the definition, implementation, and usage of this address form, and smooth the IPv4 to IPv6 transition. Simple, clear definitions are easy to explain; special cases and asterisks are not. If IPv6 is to be widely accepted and deployed, the training and educational aspects of the architecture must not be ignored.

3. Isolated Hosts

Two interpretations of the term "Isolated Host" have been proposed in the course of discussing IPv4-compatible address usage. Both are presented below, and hopefully these definitions can be clarified, and consensus reached, through further discussion.

3.1. Class 1 Isolated Nodes

The first interpretation of an isolated host is a host which does not have an on-link IPv6 router, and which thus must encapsulate all packets to off-link destinations. But this node is connected to an IPv6-capable Internet Service Provider (ISP) and thus has a provider based IPv6 address [RFC1897][V6PROVIDER], which we will refer to as PBA. This PBA is assigned to the tunnel interface and is used as source address in outgoing packets. The node has a manually configured tunnel to an ISP router. This PBA is based upon the ISP's prefix and the IPV4 address of the IPv4 interface through which the encapsulated packets get forwarded to the ISP. Note that the IPv4-compatible might be usable as the link-local address in a routing protocol, but this is yet to be determined.

So this isolated node has global IPv6 connectivity via the ISP. This isolated node has a default IPv6 route (::/0) with the ISP router as next-hop, which may be identifed by an IPv4-compatible address. Examples of this class of isolated node can be found on the current 6-bone. [6BONE]

3.2. Class 2 Isolated Nodes

The second form of isolated nodes are those nodes which are not connected to an IPv6-capable ISP, i.e. they don't have a PBA. All they have is an IPv4-compatible address and they communicate with other IPv6 nodes which have IPv4-compatible addresses using end-to-end automatic tunneling. This requires that the destination node also has an IPv4-compatible address, and implies that the packet will make a single hop (i.e. the IPv6 packet will not be

forwarded).

In the evolution of the 6bone, the second class of host is not represented. It remains to be seen how common this type of host will be as IPv6 is deployed commercially. For these nodes to

communicate with other IPv6 nodes on the Internet, the remote IPv6 system must have automatic tunneling enabled on every IPv6 node on the Internet. At some point in transition, when the IPv4 address space is exhausted, new IPv6 nodes will not be able to get IPv4-compatible addresses to do automatic tunneling. These nodes will only have PBAs and would not be able to communicate with class 2 isolated nodes. So while this class of system represents a simple configuration, it can be seen that from the beginning these nodes may only be able to communicate with a subset of the IPv6 network, and the percentage of unreachable hosts will likely increase over time. Also, the extensive use of IPv4-compatible addresses for communications between IPv6 systems will exercise the IPv4 routing infrastructure, without promoting the use of IPv6 hierarchical routing, thereby taxing an overburdened service without any gain in operational experience in the new technology.

4. Other Issues

One important issue is whether IPv4-compatible addresses should be assigned to all physical interfaces having IPv4 addresses. We believe that this is not a good idea as it creates several problems without being a solution for any existing problem. There are other issues to consider as well.

Another disadvantage is that IPv4-compatible addresses will have to be treated specially in name services like DNS and DHCP, with duplication of data and potential operational confusion resulting.

4.1. Host Issues

Hosts may have to deal with multiple mechanisms for obtaining addresses, and support dual address lifetime (or lease) constructs. While DHCP is commonly used to obtain IPv4 addresses, DHCPv6 does not support the assignment of IPv4-compatible addresses, and thus the server will not recognize such addresses as belonging to any given client. [DHCPv6]

Also, assigning an IPv4-compatible address to the interface on which IPv4 is running may not be generally possible. For example, an IPv4 host using SLIP could support an IPv6 implementation using tunnelling, but not a native interface. There may be other examples of media types which support one protocol but not the other.

4.2. Router Issues

In addition to the issues presented above, which focus largely on the impact to IPv6 hosts, there are various concerns related to dual IPv6/IPv4 routers. In the current RFC 1933 model, dual protocol routers at the borders of IPv6 islands may be called upon to perform

routing of packets using IPv4-compatible source and destination addresses. There are several reasons why this is not a good idea:

- While encapsulation of IPv6 packets in IPv4 tunnels will be a necessary function of dual IPv4/IPv6 routers, it would be best

to reduce the need for this function by having the originating host use automatic tunnelling.

- The routers may have greater memory requirements than otherwise. See the draft "IPv6 Routing Table Issues" [RJA] for details.

5. Acknowledgements

The authors wish to thank Pedro Roque, Jim Bound, Ran Atkinson, Bill Lenharth and Matt Thomas for their input and consideration, as well as the growing community of IPv6 developers.

6. References

[RFC1933] R. Gilligan, E. Nordmark, "Transition Mechanisms for IPv6 Hosts and Routers", RFC 1933, April 1996.

[V6TUNNELS] A. Conta, S. Deering, "Generic Packet Tunneling in IPv6", <draft-ietf-ipngwg-ipv6-tunnel-04.txt>, Work in Progress, October 1996.

[RFC1122] R. Braden, "Requirements for Internet Hosts - Communication Layers", RFC 1122, October 1989.

[NGLIST] Interoperability problem described on ngtrans mailing list, Wednesday March 13, 1996.

[RFC1897] R. Hinden, "IPv6 Testing Address Allocation", RFC 1897, January 1996.

[V6PROVIDER] Y. Rekhter et al, "An IPv6 Provider-Based Unicast Address Format", <draft-ietf-ipngwg-unicast-addr-fmt-04.txt>, Work in Progress, March 1996.

[6BONE] http://www-cnr.lbl.gov/6bone

[DHCPv6] J. Bound, C. Perkins, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", <draft-ietf-dhc-dhcpv6-07.txt>, Work in Progress, August 1996.

[RJA] R. Atkinson, "IPv6 Routing Table Size Issues", <draft-ietf-ipngwg-ipv6-routing-00.txt>, Work in Progress, October 1996.

7. Author's Addresses

Dan Harrington
P.O. Box 81W
W. Townsend, MA

Quaizar Vohra
Interoperability Lab
7 Leavitt Lane

University of New Hampshire
Durham, NH 03824
qv@iol.unh.edu