Provider Independent IPv6 Addressing Architecture

Status of this Memo

This document is an Internet-Draft. 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. Internet-Drafts may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet-Drafts as reference material or to cite them other than as a ``working draft'' or ``work in progress.''

To learn the current status of any Internet-Draft, please check the 1id-abstracts.txt listing contained in the Internet-Drafts Shadow Directories on ds.internic.net, nic.nordu.net, ftp.nisc.sri.com, or munnari.oz.au.

Abstract

An IPv6 addressing architecture which maximize the provider independence is described.

With flatly routed IPv4 addresses, one can subscribe to multiple providers and change providers at will without a lot of efforts. But, IPv6 packets will be hierarchically routed and their addresses will have hierarchical structure, whose higher part is determined by the network provider.

By separating a 128 bit IPv6 address into 64 bit ILOC (Internet LOCator), first 4 bytes of which is flat routable provider part and the rest 4 bytes of which is hierarchical intra-provider part, and 64 bit IID (Internet ID), which is not routable but globally unique, it is possible to preserve some of the provider independence of IPv4.

The architecture can also identify geographical location of providers.

Introduction

IPv4 addresses do not contain provider dependent information. Thus, with IPv4 addresses, we can select and change providers without reassigning IP addresses.

On the other hand, IPv6 addresses contain provider dependent part for routing aggregation.

If host identification is controlled by providers, it is difficult to change providers or have multiple provider. It is likely that subscribers are tuned into the provider they choose and tends to promote provider monopoly.

The problem, obviously, is in provider-dependent hierarchical routing.

And, the solution, obviously, is not to do provider-dependent hierarchical routing. This memo proposes an address assignment scheme for IPv6 which does flat routing for providers.

Routing below providers is, of course, hierarchical so that enough scalability is assured to support 10^12 networks and 10^15 hosts.

A mechanism to identify small geographical locations is also included to have geographically-near-optimal and least-costly routing with proper route selection.

Assignment Plan for IPv6 address

The 16 byte IPv6 address is divided into two fields: 8 byte ILOC (Internet LOCator) and 8 byte IID (Internet ID).

ILOC is further divided into three sub fields: 4 byte "Provider ID", 2 byte "Subscriber ID" and 2 byte "Subnet ID".

"Provider ID" and "Subscriber ID" together is called "Provider dependent part".

Provider dependent part is supplied by providers and dynamically reconfigurable at system boot time or even during operation. It is expected that routers to providers announce provider dependent part information.

IID is a globally unique ID of a host or a multicast group to identify the host or the multicast group. While an IID uniquely identifies a single host, a host may have multiple IIDs. But, within a lifetime of some connection or reservation such as for TCP or flow, the same IID should be used regardless of the routing changes.

IID is supplied by subscribers. The configuration may be automatic. But it is expected that renumbering is necessary not so often, in

general, only when a host is purchased or the host is moved to different suborganization of the provider. Host specific information such as IP address to host name mapping is looked up only through IID. For autoconfiguration, it must be possible to derive some IIDs from MAC addresses. As 2^28<10^15, MAC addresses to support 10^15 hosts must be longer than 48 bits (actually, a lot longer than that), IID is designed to have 64 bits.

Subnet ID is also supplied by subscribers and identifies a subnet within a subscribers LAN. The configuration may be automatic through nearest routers. Renumbering is necessary when a location of a host is changed to a different subnet.

Network layer identification of a host is done through IID just like the current IPv4.

Routing is controlled purely by the ILOC part of IPv6 address, which is 8byte aligned and, thus, is more efficient than schemes using full 16 bytes.

For the deliverly between adjacent routers or a router and a host, which is within a single link layer, only the identification is necessary. So, for IPv6 ARP equivalent, IID is used and ILOC is not consulted.

Users can change providers at will just by disconnecting one of its external routers and connect it to a new provider, and ILOC part will be automatically reconfigured. But, it is still necessary to update information of DNS, which is further explained in the "DNS interaction" section.

A host of a subscriber belonging to multiple providers may have multiple provider dependent parts.

Different interfaces of a host is, in general, distinguished by ILOC part.

Assignment Plan for Provider ID, Subscriber ID and Subnet ID

4 byte provider ID uniquely identifies a single provider in the Internet, while a provider may have multiple provider IDs.

4 byte provider ID combined with 2 byte subscriber ID uniquely identifies a subscriber in the Internet.

Provider ID of 0 is reserved for subscriber local routing.

2 byte subnet ID uniquely identifies a subnet in a subscribers

network.

A large subscriber having more than 65536 subnets will have multiple subscriber IDs from a provider.

A large provider having more than 65536 subscriber IDs or having some geographical constraints, as explained in the next section, will have multiple provider IDs.

Avoiding Treework

A goal of IPv6 is to avoid treework and make it real network.

With usual provider based addressing, users within a single provider share the same routing prefix so that it is difficult to use better route outside the provider. That is, complex configuration is necessary to break the routing hierarchy, which does not scale. It is unlikely that providers welcome such configuration only to allow subscribers pay less to the provider.

So, in this proposal, to identify small geographical locations, a provider ID should not cover an area of 100Km radius. That is, a large provider must use different provider IDs for hosts located more than 200Km apart. The distance is measured only for IP entry point of providers, so that PPP service through non-IP public data network for 1,000Km-distant user is allowed.

Thus, it becomes naturally possible to favor local links outside of the centralized provider, if local IXes are available.

Only about 17,000 IDs are necessary to cover the surface of the Earth. Inter-planetary communication is NOT considered here.

Assignment Plan for IIDs

There are two kinds of IIDs, structured and unstructured.

Structured IIDs are maintained by IANA and is used to lookup IANA maintained database of in-addr.arpa. equivalent of IPv6.

Unstructured IIDs are directly derived from globally unique MAC addresses of hosts and useful for autoconfiguration.

The most significant bit of structured IID is 0.

First three bytes of structured IID are assigned from IANA to country NICs.

Each country NIC uses the next three bytes for independent subscribers.

A subscriber use the last two bytes for the internal use.

If the lease significant bit of IID is 0, it is for unicasting. Otherwise, the IID identifies multicast. An IID of all 1's except of the first bit is for broadcasting.

The most significant bit of unstructured IID is 1. For the current 48 bit MAC addresses maintained by IEEE, 16 bits prefix of "1000000000000000" is added to form the IID. Rest of the unstructured IID space is reserved for the future use.

DNS Interaction

While IPv6 addresses are mostly autoconfigurable, it is still necessary to use DNS to advertise IPv6 addresses all over the Internet.

As IIDs are considered to be rather static, they can be stored in DNS just as the current IPv4 addresses.

On the other hand, if a subscriber adds or changes a provider, it is necessary to reflect the change to DNS. To do so without a lot of pain, the DNS lookup of ILOC should be indirect. That is, each DNS node of a host should not directly have ILOC but have a pointer to some node. The node, which is pointed by a lot of hosts within the same subscriber, then, have (multiple) ILOC information of the subscriber. Thus, it is possible to quickly change a provider without much difficulty. Of course, some statically configured raw addresses, which is necessary for the minimal operation of DNS itself, will still need to be reconfigured.

Supporting 10^12 networks and 10^15 hosts

How the requirement to support 10^12 networks and 10^15 hosts can be satisfied?

First, how routing between 10^12 networks is possible? 10^3 hosts within a subnet can easily be identified by the IID. Thus, (Provider ID, Subscriber ID, Subnet ID) must identify 10^12 networks. It is not unnatural that a provider, in average, supports at least 10^3 subscribers. It can also be safely assumed that subnet ID can identify 10^1 subnets. Thus, Provider ID is required to identify 10^8 hosts. Considering that the requirement contains 10^2 safety factor, the least significant byte of the Provider ID are reserved for the factor. The remaining 3 bytes (2^24>10^7) are much more than

enough to identify 10^6 providers. It is assumed that by the time we support 10^12 networks, flat routing of 10^6 providers is not a problem at all. The reserved lowest byte of provider ID may also be used for two-level routing among providers.

Next, how identification of 10^15 hosts is possible? 10^3 hosts of a subscriber can easily be identified by the last two bytes of IID. For the remaining 10^12 factor, IANA and country NICs are expected to manage the upper 6 bytes (for about 2*10^24 hosts) densely enough. Thus, it is feasible that more than 10^12 networks can be identified with the 6 bytes.

Conclusion

As the 16 byte address space is so large, it is possible to use it wisely to enjoy full provider independence, including provider change without renumbering and long distance provider selection by provider IDs.

Acknowledgement

The work is inspired by the concept of locator/EID of NIMROD and PIP. Josh Osborne helped the Author to clarify the geographical identification feature of the proposal.

References

(to be provided)

Security Considerations

(to be provided)

Author's Addresses

Masataka Ohta
Computer Center
Tokyo Institute of Technology
2-12-1, O-okayama, Meguro-ku
Tokyo 152, JAPAN

Phone: +81-3-5734-3299
Fax: +81-3-5734-3415
EMail: mohta@necom830.cc.titech.ac.jp