

Evolving the Internet Addressing Architecture

James Kempf

DoCoMo Communications Laboratories USA

kempf@docomolabs-usa.com

Abstract-- This paper summarizes discussion from the open architecture meeting on the Internet addressing architecture held by the IAB at IETF 57 in July, 2003 [1].

The basic Internet addressing architecture remains the same today as when it was first developed 30 years ago. Device interfaces are assigned fixed numerical addresses serving three basic purposes:

1. As an endpoint identifier, uniquely identifying a communication endpoint
2. As a locator, identifying where in the network topology a device interface is located,
3. As a forwarding identifier, allowing intermediaries to pass a packet to a destination device interface.

As a practical matter, 1 and 3 tend to be functionally active, while 2 is an attribute of the address related to 3. These functions are different and there is tension between the constraints of routing implied by 2 and 3 and the requirements for end to end transport session identification implied by 1.

In IPv4 there are two types of addresses:

1. Unique structured addresses for global Internet use,
2. Private, reusable addresses for local use.

Addresses of type 2 were originally added for two purposes:

1. For devices that never, ever are connected to the global Internet
2. In order to alleviate a perceived pending shortage of IPv4 address space on the global Internet.

As it turned out, private addresses have proven useful for other purposes independently of any lack of global connectivity or shortage of address space.

In IPv6, there are three types of addresses:

1. Unique structured addresses for global Internet use,
2. Site local addresses for scoped local use,
3. Link local addresses for use within a single subnet.

Address types 1 and 2 are essentially the same in both IPv4 and IPv6, though 2 is expressed somewhat differently in IPv6.

Global addresses can be used as locators and forwarding identifiers anywhere in the globally addressable Internet. Private, reusable addresses and site local addresses can be used as locators and forwarding identifiers only within a private addressing realm. The borders of a private addressing realm are flexible, and depend on the configuration of routers at the border. Border routers drop private addresses, enforcing the scope of addressability. Link local addresses can only be used on the local subnet, and are typically dropped by the first hop router.

The IPv4 addressing architecture customarily assumes a 1:1 mapping between an address and a device interface in a single addressing realm (either global or private). This mapping is not enforced by the architecture, however. The IPv6 architecture is exploring having multiple addresses per device interface, potentially from multiple realms. At a minimum, every device interface in IPv6 requires at least one link local

address for performing subnet configuration, and may have more. In addition, an interface can be configured (potentially automatically) with one and/or several global and/or site local addresses. The exact implications of this flexibility in practice are not entirely clear, and it is also not clear what issues it addresses.

Recent challenges for the architecture are posed by wireless devices and multi-interface devices. A wireless device interface is a challenge because its point of attachment to the Internet changes in time, requiring the address-as-locator and address-as-forwarding-identifier to change. However, if an open transport connection remains on the device when it moves its point of attachment, the address-as-endpoint-identifier must remain the same. Thus, a wedge is driven between the endpoint identifier function of addresses and the locator/forwarding identifier functions, as the address changes in time. Existing IP mobility protocols such as Mobile IP solve some part of the problem by utilizing two addresses, one as endpoint identifier and one as locator/forwarding identifier. This solution is challenged by security solutions such as IPsec which embeds an expectation that the locator/forwarding identifier address is the same as the endpoint identifier address.

A multi-interface device is a challenge because the device-as-endpoint can be identified through any one of a number of possible addresses, again driving a wedge between the endpoint identifier function and locator/forwarding identifier functions, this time in topological, network space. For traditional, wired multi-interface devices in IPv4, each interface was customarily treated as a separate endpoint identifier for all practical purposes, and assigned a different host name in the DNS. However, particularly for wireless devices, the mapping of logical traffic flows to multiple interfaces either serially or in parallel has the potential to increase robustness and performance, but exactly how to maintain endpoint identification across such mapping is an open question. Mobile IP has similar difficulties with multi-interface devices.

Recently, the IETF has begun to re-examine the decision for including addresses of type 2 (locally scoped addresses) in the IPv6 addressing architecture. There has been discussion of deprecating such addresses as a specific class of address. The re-examination is complicated by the multiple uses to which people put private addresses. Any decision to deprecate site local addressing in IPv6 must therefore be accompanied by an alternative that provides support for these, and there are discussions on alternatives that preserve the filtering at flexible site borders without specifying an address type for that purpose.

Reference

- [1] <http://www.ietf.org/proceedings/03jul/index.html> (iab)