

## Routing Aspects Of IPv6 Transition

Status of this memo

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

Abstract

This document gives an overview of the routing aspects of the IPv6 transition. It is based on the protocols defined in the document "Transition Mechanisms for IPv6 Hosts and Routers" [1]. Readers should be familiar with the transition mechanisms before reading this document.

The proposals contained in this document are based on the work of the Ngtrans working group.

### 1. TERMINOLOGY

This paper uses the following terminology:

- node - a protocol module that implements IPv4 or IPv6.
- router - a node that forwards packets not explicitly addressed to itself.
- host - any node that is not a router.
- border router - a router that forwards packets across routing domain boundaries.
- link - a communication facility or medium over which nodes can communicate at the link layer, i.e., the layer immediately below internet layer.
- interface - a node's attachment to a link.
- address - an network layer identifier for an interface or a group of interfaces.

neighbors - nodes attached to the same link.

routing domain - a collection of routers which coordinate routing knowledge using a single routing protocol.

routing region (or just "region") - a collection of routers interconnected by a single internet protocol (e.g. IPv6) and coordinating their routing knowledge using routing protocols from a single internet protocol stack. A routing region may be a superset of a routing domain.

tunneling - encapsulation of protocol A within protocol B, such that A treats B as though it were a datalink layer.

reachability information - information describing the set of reachable destinations that can be used for packet forwarding decisions.

routing information - same as reachability information.

address prefix - the high-order bits in an address.

routing prefix - address prefix that expresses destinations which have addresses with the matching address prefixes. It is used by routers to advertise what systems they are capable of reaching.

route leaking - advertisement of network layer reachability information across routing region boundaries.

## 2. ISSUES AND OUTLINE

This document gives an overview of the routing aspects of IPv4 to IPv6 transition. The approach outlined here is designed to be compatible with the existing mechanisms for IPv6 transition [1].

During an extended IPv4-to-IPv6 transition period, IPv6-based systems must coexist with the installed base of IPv4 systems. In such a dual internetworking protocol environment, both IPv4 and IPv6 routing infrastructure will be present. Initially, deployed IPv6-capable domains might not be globally interconnected via IPv6-capable internet infrastructure and therefore may need to communicate across IPv4-only routing regions. In order to achieve dynamic routing in such a mixed environment, there need to be mechanisms to globally distribute IPv6 network layer reachability information between dispersed IPv6 routing regions. The same techniques can be used in later stages of IPv4-to-IPv6 transition to route IPv4 packets between isolated IPv4-only routing region over IPv6 infrastructure.

The IPng transition provides a dual-IP-layer transition, augmented by use of encapsulation where necessary and appropriate. Routing issues related to this transition include:

- (1) Routing for IPv4 packets
- (2) Routing for IPv6 packets
  - (2a) IPv6 packets with IPv6-native addresses
  - (2b) IPv6 packets with IPv4-compatible addresses
- (3) Operation of manually configured static tunnels
- (4) Operation of automatic encapsulation
  - (4a) Locating encapsulators
  - (4b) Ensuring that routing is consistent with encapsulation

Basic mechanisms required to accomplish these goals include: (i) Dual-IP-layer Route Computation; (ii) Manual configuration of point-to-point tunnels; and (iii) Route leaking to support automatic encapsulation.

The basic mechanism for routing of IPv4 and IPv6 involves dual-IP-layer routing. This implies that routes are separately calculated for IPv4 addresses and for IPv6 addressing. This is discussed in more detail in section 3.1.

Tunnels (either IPv4 over IPv6, or IPv6 over IPv4) may be manually configured. For example, in the early stages of transition this may be used to allow two IPv6 domains to interact over an IPv4 infrastructure. Manually configured static tunnels are treated as if they were a normal data link. This is discussed in more detail in section 3.2.

Use of automatic encapsulation, where the IPv4 tunnel endpoint address is determined from the IPv4 address embedded in the IPv4-compatible destination address of IPv6 packet, requires consistency of routes between IPv4 and IPv6 routing domains for destinations using IPv4-compatible addresses. For example, consider a packet which starts off as an IPv6 packet, but then is encapsulated in an IPv4 packet in the middle of its path from source to destination. This packet must locate an encapsulator at the correct part of its path. Also, this packet has to follow a consistent route for the entire path from source to destination. This is discussed in more detail in section 3.3.

The mechanisms for tunneling IPv6 over IPv4 are defined in the transition mechanisms specification [1].

### 3. MORE DETAIL OF BASIC APPROACHES

#### 3.1 Basic Dual-IP-layer Operation

In the basic dual-IP-layer transition scheme, routers may independently support IPv4 and IPv6 routing. Other parts of the transition, such as DNS support, and selection by the source host of which packet format to transmit (IPv4 or IPv6) are discussed in [1]. Forwarding of IPv4 packets is based on routes learned through running IPv4-specific routing protocols. Similarly, forwarding of IPv6 packets (including IPv6-packets with IPv4-compatible addresses) is based on routes learned through running IPv6-specific routing protocols. This implies that separate instances of routing protocols are used for IPv4 and for IPv6 (although note that this could consist of two instances of OSPF and/or two instances of RIP, since both OSPF and RIP are capable of supporting both IPv4 and IPv6 routing).

A minor enhancement would be to use an single instance of an integrated routing protocol to support routing for both IPv4 and IPv6. At the time that this is written there is no protocol which has yet been enhanced to support this. This minor enhancement does not change the basic dual-IP-layer nature of the transition.

For initial testing of IPv6 with IPv4-compatible addresses, it may be useful to allow forwarding of IPv6 packets without running any IPv6-compatible routing protocol. In this case, a dual (IPv4 and IPv6) router could run routing protocols for IPv4 only. It then forwards IPv4 packets based on routes learned from IPv4 routing protocols. Also, it forwards IPv6 packets with an IPv4-compatible destination address based on the route for the associated IPv4 address. There are a couple of drawbacks with this approach: (i) It does not specifically allow for routing of IPv6 packets via IPv6-capable routers while avoiding and routing around IPv4-only routers; (ii) It does not produce routes for "non-compatible" IPv6 addresses. With this method the routing protocol does not tell the router whether neighboring routers are IPv6-compatible. However, neighbor discovery may be used to determine this. Then if an IPv6 packet needs to be forwarded to an IPv4-only router it can be encapsulated to the destination host.

#### 3.2 Manually Configured Static Tunnels

Tunneling techniques are already widely deployed for bridging non-IP network layer protocols (e.g. AppleTalk, CLNP, IPX) over IPv4 routed infrastructure. IPv4 tunneling is an encapsulation of arbitrary packets inside IPv4 datagrams that are forwarded over IPv4 infrastructure between tunnel endpoints. For a tunneled protocol, a tunnel appears as a single-hop link (i.e. routers that establish a

tunnel over a network layer infrastructure can inter-operate over the tunnel as if it were a one-hop, point-to-point link). Once a tunnel is established, routers at the tunnel endpoints can establish routing adjacencies and exchange routing information. Describing the protocols for performing encapsulation is outside the scope of this paper (see [1]). Static point-to-point tunnels may also be established between a host and a router, or between two hosts. Again, each manually configured point-to-point tunnel is treated as if it was a simple point-to-point link.

### 3.3 Automatic Tunnels

Automatic tunneling may be used when both the sending and destination nodes are connected by IPv4 routing. In order for automatic tunneling to work, both nodes must be assigned IPv4-compatible IPv6 addresses. Automatic tunneling can be especially useful where either source or destination hosts (or both) do not have any adjacent IPv6-capable router. Note that by "adjacent router", this includes routers which are logically adjacent by virtue of a manually configured point-to-point tunnel (which is treated as if it is a simple point-to-point link).

With automatic tunneling, the resulting IPv4 packet is forwarded by IPv4 routers as a normal IPv4 packet, using IPv4 routes learned from routing protocols. There are therefore no special issues related to IPv4 routing in this case. There are however routing issues relating to how IPv6 routing works in a manner which is compatible with automatic tunneling, and how tunnel endpoint addresses are selected during the encapsulation process. Automatic tunneling is useful from a source host to the destination host, from a source host to a router, and from a router to the destination host. Mechanisms for automatic tunneling from a router to another router are not currently defined.

#### 3.3.1 Host to Host Automatic Tunneling

If both source and destination hosts make use of IPv4-compatible IPv6 addresses, then it is possible for automatic tunneling to be used for the entire path from the source host to the destination host. In this case, the IPv6 packet is encapsulated in an IPv4 packet by the source host, and is forwarded by routers as an IPv4 packet all the way to the destination host. This allows initial deployment of IPv6-capable hosts to be done prior to the update of any routers.

A source host may make use of Host to Host automatic tunneling provided that the following are both true:

- the source address is an IPv4-compatible IPv6 address.
- the destination address is an IPv4-compatible IPv6 address.
- the source host does know of one or more neighboring IPv4-capable routers, or the source and destination are on the same subnet.

If all of these requirements are true, then the source host may encapsulate the IPv6 packet in an IPv4 packet, using a source IPv4 address which is extracted from the associated source IPv6 address, and using a destination IPv4 address which is extracted from the associated destination IPv6 address.

Where host to host automatic tunneling is used, the packet is forwarded as a normal IPv4 packet for its entire path, and is decapsulated (i.e., the IPv4 header is removed) only by the destination host.

### 3.3.2 Host to Router Configured Default Tunneling

In some cases "configured default" tunneling may be used to encapsulate the IPv6 packet for transmission from the source host to an IPv6-backbone. However, this requires that the source host be configured with an IPv4 address to use for tunneling to the backbone.

Configured default tunneling is particularly useful if the source host does not know of any local IPv6-capable router (implying that the packet cannot be forwarded as a normal IPv6 packet directly over the link layer), and when the destination host does not have an IPv4-compatible IPv6 address (implying that host to host tunneling cannot be used).

Host to router configured default tunneling may optionally also be used even when the host does know of a local IPv6 router. In this case it is a policy decision whether the host prefers to send a native IPv6 packet to the IPv6-capable router or prefers to send an encapsulated packet to the configured tunnel endpoint.

Similarly host to router default configured tunneling may be used even when the destination address is an IPv4-compatible IPv6 address. In this case for example a policy decision may be made to prefer tunneling for part of the path and native IPv6 for part of the path, or alternatively to use tunneling for the entire path from source host to destination host.

A source host may make use of host to router configured default tunneling provided that ALL of the following are true:

- the source address is an IPv4-compatible IPv6 address.
- the source host does know of one or more neighboring IPv4-capable routers
- the source host has been configured with an IPv4 address of an dual router which can serve as the tunnel endpoint.

If all of these requirements are true, then the source host may encapsulate the IPv6 packet in an IPv4 packet, using a source IPv4 address which is extracted from the associated source IPv6 address, and using a destination IPv4 address which corresponds to the configured address of the dual router which is serving as the tunnel endpoint.

When host to router configured default tunneling is used, the packet is forwarded as a normal IPv4 packet from the source host to the dual router serving as tunnel endpoint, is decapsulated by the dual router, and is then forwarded as a normal IPv6 packet by the tunnel endpoint.

#### 3.3.2.1 Routing to the Endpoint for the Configured Default Tunnel

The dual router which is serving as the end point of the host to router configured default tunnel must advertise reachability into IPv4 routing sufficient to cause the encapsulated packet to be forwarded to it.

The simplest approach is for a single IPv4 address to be assigned for use as a tunnel endpoint. One or more dual routers, which have connectivity to the IPv6 backbone and which are capable of serving as tunnel endpoint, advertise a host route to this address into IPv4 routing in the IPv4-only region. Each dual host in the associated IPv4-only region is configured with the address of this tunnel endpoint and selects a route to this address for forwarding encapsulated packet to a tunnel end point (for example, the nearest tunnel end point, based on whatever metric(s) the local routing protocol is using).

Finally, in some cases there may be some reason for specific hosts to prefer one of several tunnel endpoints, while allowing all potential tunnel endpoints to serve as backups in case the preferred endpoint is not reachable. In this case, each dual router with IPv6 backbone connectivity which is serving as potential tunnel endpoint is given a unique IPv4 address taken from a single IPv4 address block (where the IPv4 address block is assigned either to the organization administering the IPv4-only region, or to the organization

administering the local part of the IPv6 backbone). In the likely case that there are much less than 250 such dual routers serving as tunnel endpoints, we suggest using multiple IPv4 addresses selected from a single 24-bit IPv4 address prefix for this purpose. Each dual router then advertises two routes into the IPv4 region: A host route corresponding to the tunnel endpoint address specifically assigned to it, and also a standard (prefix) route to the associated IPv4 address block. Each dual host in the IPv4-only region is configured with a tunnel endpoint address which corresponds to the preferred tunnel endpoint for it to use. If the associated dual router is operating, then the packet will be delivered to it based upon the host route that it is advertising into the IPv4-only region. However, if the associated dual router is down, but some other dual router serving as a potential tunnel endpoint is operating, then the packet will be delivered to the nearest operating tunnel endpoint.

### 3.3.3 Router to Host Automatic Tunneling

In some cases the source host may have direct connectivity to one or more IPv6-capable routers, but the destination host might not have direct connectivity to any IPv6-capable router. In this case, provided that the destination host has an IPv4-compatible IPv6 address, normal IPv6 forwarding may be used for part of the packet's path, and router to host tunneling may be used to get the packet from an encapsulating dual router to the destination host.

In this case, the hard part is the IPv6 routing required to deliver the IPv6 packet from the source host to the encapsulating router. For this to happen, the encapsulating router has to advertise reachability for the appropriate IPv4-compatible IPv6 addresses into the IPv6 routing region. With this approach, all IPv6 packets (including those with IPv4-compatible addresses) are routed using routes calculated from native IPv6 routing. This implies that encapsulating routers need to advertise into IPv6 routing specific route entries corresponding to any IPv4-compatible IPv6 addresses that belong to dual hosts which can be reached in an neighboring IPv4-only region. This requires manual configuration of the encapsulating routers to control which routes are to be injected into IPv6 routing protocols. Nodes in the IPv6 routing region would use such a route to forward IPv6 packets along the routed path toward the router that injected (leaked) the route, at which point packets are encapsulated and forwarded to the destination host using normal IPv4 routing.

Depending upon the extent of the IPv4-only and dual routing regions, the leaking of routes may be relatively simple or may be more complex. For example, consider a dual Internet backbone, connected via one or two dual routers to an IPv4-only stub routing domain. In



this case, it is likely that there is already one summary address prefix which is being advertised into the Internet backbone in order to summarize IPv4 reachability to the stub domain. In such a case, the border routers would be configured to announce the IPv4 address prefix into the IPv4 routing within the backbone, and also announce the corresponding IPv4-compatible IPv6 address prefix into IPv6 routing within the backbone.

A more difficult case involves the border between a major Internet backbone which is IPv4-only, and a major Internet backbone which supports both IPv4 and IPv6. In this case, it requires that either (i) the entire IPv4 routing table be fed into IPv6 routing in the dual routing domain (implying a doubling of the size of the routing tables in the dual domain); or (ii) Manual configuration is required to determine which of the addresses contained in the Internet routing table include one or more IPv6-capable systems, and only these addresses be advertised into IPv6 routing in the dual domain.

#### 3.3.4 Example of How Automatic Tunnels May be Combined

Clearly tunneling is useful only if communication can be achieved in both directions. However, different forms of tunneling may be used in each direction, depending upon the local environment, the form of address of the two hosts which are exchanging IPv6 packets, and the policies in use.

Table 1 summarizes the form of tunneling that will result given each possible combination of host capabilities, and given one possible set of policy decisions. This table is derived directly from the requirements for automatic tunneling discussed above.

The example in table 1 uses a specific set of policy decisions: It is assumed in table 1 that the source host will transmit a native IPv6 where possible in preference over encapsulation. It is also assumed that where tunneling is needed, host to host tunneling will be preferred over host to router tunneling. Other combinations are therefore possible if other policies are used.

Due to a specific policy choice, the default sending rules in [1] may not be followed.

Note that IPv6-capable hosts which do not have any local IPv6 router must be given an IPv4-compatible v6 address in order to make use of their IPv6 capabilities. Thus, there are no entries for IPv6-capable hosts which have an incompatible IPv6 address and which also do not have any connectivity to any local IPv6 router. In fact, such hosts could communicate with other IPv6 hosts on the same local network without the use of a router. However, since this document focuses on

routing and router implications of IPv6 transition, direct communication between two hosts on the same local network without any intervening router is outside the scope of this document.

Also, table 1 does not consider manually configured point-to-point tunnels. Such tunnels are treated as if they were normal point-to-point links. Thus any two IPv6-capable devices which have a manually configured tunnel between them may be considered to be directly connected.

| Host A                                  | Host B                                  | Result   |
|---|---|--|
| v4-compatible addr.<br>no local v6 rtr. | v4-compatible addr.<br>no local v6 rtr. | host to host tunneling<br>in both directions   |
| v4-compatible addr.<br>no local v6 rtr. | v4-compatible addr.<br>local v6 rtr.    | A->B: host to host tunnel<br>B->A: v6 forwarding plus<br>rtr->host tunnel                        |
| v4-compatible addr.<br>no local v6 rtr. | incompatible addr.<br>local v6 rtr.     | A->B: host to rtr tunnel<br>plus v6 forwarding<br>B->A: v6 forwarding plus<br>rtr to host tunnel |
| v4-compatible addr.<br>local v6 rtr.    | v4-compatible addr.<br>local v6 rtr.    | end to end native v6<br>in both directions   |
| v4-compatible addr.<br>local v6 rtr.    | incompatible addr.<br>local v6 rtr.     | end to end native v6<br>in both directions   |
| incompatible addr.<br>local v6 rtr.     | incompatible addr.<br>local v6 rtr.     | end to end native v6<br>in both directions   |

Table 1: Summary of Automatic Tunneling Combinations

### 3.3.5 Example

Figure 2 illustrates an example network with two regions A and B. Region A is dual, meaning that the routers within region A are capable of forwarding both IPv4 and IPv6. Region B is IPv4-only, implying that the routers within region B are capable of routing only IPv4. The illustrated routers R1 through R4 are dual. The illustrated routers r5 through r9 are IPv4-only. Also assume that hosts H3 through H8 are dual. Thus H7 and H8 have been upgraded to be IPv6-capable, even though they exist in a region in which the routers are not IPv6-capable. However, host h1 and h2 are IPv4-only.

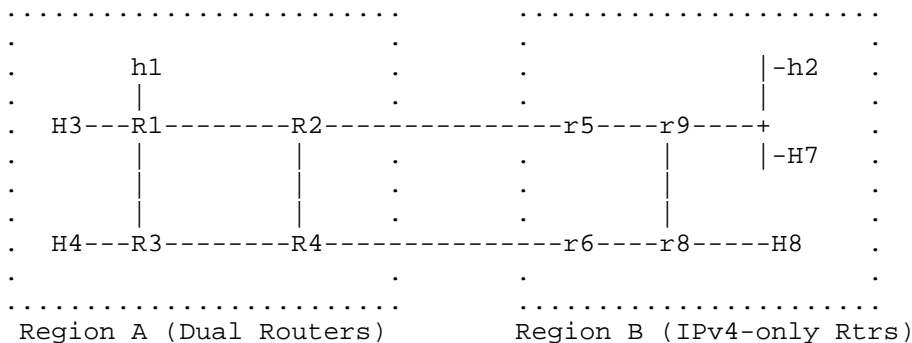


Figure 2: Example of Automatic Tunneling

Consider a packet from h1 to H8. In this case, since h1 is IPv4-only, it will send an IPv4 packet. This packet will traverse regions A and B as a normal IPv4 packet for the entire path. Routing will take place using normal IPv4 routing methods, with no change from the operation of the current IPv4 Internet (modulo normal advances in the operation of IPv4, of course). Similarly, consider a return packet from H8 to h1. Here again H8 will transmit an IPv4 packet, which will be forwarded as a normal IPv4 packet for the entire path.

Consider a packet from H3 to H8. In this case, since H8 is in an IPv4-only routing domain, we can assume that H8 uses an IPv4-compatible IPv6 address. Since both source and destination are IPv6-capable, H3 may transmit an IPv6 packet destined to H8. The packet will be forwarded as far as R2 (or R4) as an IPv6 packet.

Router R2 (or R4) will then encapsulate the full IPv6 packet in an IPv4 header for delivery to H8. In this case it is necessary for routing of IPv6 within region A to be capable of delivering this packet correctly to R2 (or R4). As explained in section 3.3, routers R2 and R4 may inject routes to IPv4-compatible IPv6 addresses into the IPv6 routing used within region A corresponding to the routes which are available via IPv4 routing within region B.

Consider a return packet from H8 to H3. Again, since both source and destination are IPv6-capable, a IPv6 packet may be transmitted by H8. However, since H8 does not have any direct connectivity to an IPv6-capable router, H8 must make use of an automatic tunnel. Which form of automatic tunnel will be used depends upon the type of address assigned to H3.

If H3 is assigned an IPv4-compatible address, then the requirements specified in section 3.3.1 will all be satisfied. In this case host H8 may encapsulate the full IPv6 packet in an IPv4 header using a source IPv4 address extracted from the IPv6 address of H8, and using a destination IPv4 address extracted from the IPv6 address of H3.

If H3 has an IPv6-only address, then it is not possible for H8 to extract an IPv4 address to use as the destination tunnel address from the IPv6 address of H3. In this case H8 must use host to router tunneling, as specified in section 3.3.2. In this case one or both of R2 and R4 must have been configured with a tunnel endpoint IPv4 address (R2 and R4 may use either the same address or different addresses for this purpose). R2 and/or R4 therefore advertise reachability to the tunnel endpoint address to r5 and r6 (respectively), which advertise this reachability information into region B. Also, H8 must have been configured to know which tunnel endpoint address to use for host to router tunneling. This will result in the IPv6 packet, encapsulated in an IPv4 header, to be transmitted as far as the border router R2 or R4. The border router will then strip off the IPv4 header, and forward the remaining IPv6 packet as a normal IPv6 packet using the normal IPv6 routing used in region A.

#### 4. SECURITY CONSIDERATIONS

Use of tunneling may violate firewalls of underlying routing infrastructure.

No other security issues are discussed in this paper.

#### 5. REFERENCES

- [1] Gilligan, B. and E. Nordmark. Transition Mechanisms for IPv6 Hosts and Routers, Sun Microsystems, RFC 1933, April 1996.

#### 6. AUTHORS' ADDRESSES

Ross Callon  
Cascade Communications Co.  
5 Carlisle Road  
Westford, MA 01886  
email: rcallon@casc.com

Dimitry Haskin  
Bay Networks, Inc.  
2 Federal Street  
Billerica, MA 01821  
email: dhaskin@baynetworks.com