HP-UX IPv6 Transition Mechanisms

White Paper

Table of Contents

1. Executive Summary

The Internet has evolved as a critical business communication medium; the current Internet Protocol (IPv4), which has been amazingly resilient so far, will not be able to support the continuous growth of the Internet indefinitely. The new emerging markets of nomadic personal computing devices, IP enabled devices and networked home entertainment require unique addressing, improved mobility and security. IPv4, though retrofitted to meet some of the needs, is fast reaching its limit.

The Internet Engineering Task Force (IETF) has designed IPv6, the successor to IPv4, to meet the growing demands of the Internet and Enterprise Networks. IPv6 was designed to improve upon the scalability, security, mobility, ease of configuration and management capabilities of IPv4. Some of the key advantages of IPv6 are:

- Expanded Routing and Addressing: IPv6 increases the address size from 32 bits to 128 bits, thus exponentially increasing the number of available unique addresses. It also supports multiple levels of addressing hierarchy.
- Mandatory Security: Security extension headers are part of the base IPv6 protocol and hence a mandatory part of every IPv6 implementation. Together with the increased address space, it allows to provide end-to-end security.
- Natural Mobility support: The improved option support in IPv6, together with stateless address auto-configuration, routing headers, security headers and a new type of address called the "anycast" address contribute to the natural design of mobility for IPv6 nodes.
- Integrated QoS support: The Flow Label and Traffic Class fields are part of the IPv6 header and can be used for prioritization and reservation of network traffic.

Even though IPv6 provides many improvements over IPv4, IPv6 adoption is going to be gradual due to the vast installed base of IPv4. The Internet is too large for any kind of controlled upgrade. IPv6 will have to be deployed in a highly diffusive and incremental manner with minor interdependencies. The newly deployed IPv6 nodes should also be able to interoperate with existing IPv4 nodes. The IETF has designed many transition mechanisms to help achieve these goals.

This document describes the transition mechanisms currently supported by HP-UX and how these mechanisms can be used to deploy IPv6 with no impact to the existing IPv4 infrastructure.

2. Overview

Large-scale migration to any new technology is challenging and time consuming. IPv6 is no exception; its deployment is expected to be gradual. As IPv6 is being deployed, there will be a lengthy transition period during which IPv4 and IPv6 protocols will need to coexist. The IETF developed a number of transition mechanisms that will facilitate IPv6 deployment without impacting the existing installed base. The main goals of these transition mechanisms are to allow newly deployed IPv6 nodes to interoperate with existing IPv4 nodes and allow isolated IPv6 nodes to communicate with each other using the existing IPv4 infrastructure.

HP-UX (11iv1, 11iv2) supports the following transition mechanisms:

- Dual stack
- Tunneling
- 6to4

The transition mechanisms dual stack and tunneling are specified in [3] and 6to4 is specified in [4].

3. Terminology

IPv4-only node: A host or router that supports only IPv4. An IPv4-only node does not understand IPv6. The installed base of IPv4 hosts and routers existing before the transition begins are IPv4-only nodes.

IPv6/IPv4 node: A host or router that supports both IPv4 and IPv6. Those nodes are also called dual stack nodes.

IPv6-only node: A host or router that supports IPv6 only, and does not support IPv4.

IPv6 node: Any host or router that supports IPv6. IPv6/IPv4 and IPv6-only nodes are both IPv6 nodes.

IPv4 node: Any host or router that supports IPv4. IPv6/IPv4 and IPv4-only nodes are both IPv4 nodes.

4. Dual Stack

Dual stack (or dual IP layer) is a technique that provides complete support for both IPv4 and IPv6 protocols in a single node, thus allowing the node to support both existing IPv4 applications and new IPv6 applications.

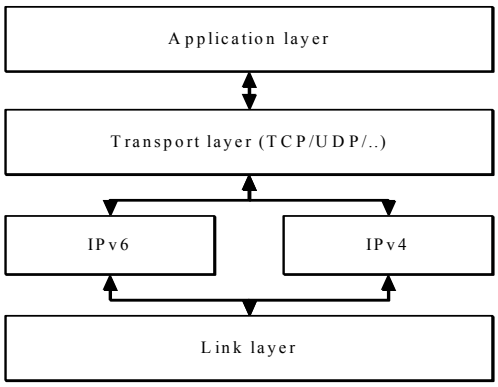

Figure 1: Dual-stack

Figure 1 illustrates the protocol stack in a dual stack node. The common transport layer components support two network layer components (IPv4 and IPv6), allowing applications on the node to communicate with both IPv4 and IPv6 applications. A dual stack node should be configured with both IPv4 and IPv6 addresses.

Dual stack is the most straightforward technique that allows interoperation between IPv6 applications and IPv4 applications. IPv6-aware applications such as telnet, ftp, DNS, Web Servers deployed on dual stack nodes can service existing IPv4-only nodes and new IPv6 nodes avoiding the need to run two separate applications for the same purpose.

IPv6-aware applications running on dual stack nodes communicate with IPv4 applications running on IPv4-only nodes using a special IPv6 address called an IPv4-mapped IPv6 address [1]. The IPv4-mapped IPv6 address with the format "::ffff:a.b.c.d" contains the IPv4 address in the low-order 32 bits.

When an IPv4 client application sends a request to an IPv6 server application bound to a wildcard address, the dual stack node presents the client address as a 128 bit IPv4-mapped IPv6 address to the IPv6 application. This allows the IPv6 application to service the request as a regular IPv6 request. While transmitting the response, the dual stack node will properly interpret the IPv4-mapped IPv6 address as an IPv4 address and the packets exchanged between the two nodes will be IPv4 packets.

5. Tunneling

Tunneling is a key transition mechanism strategy that allows isolated IPv6 nodes or networks to communicate with each other using a virtual link created over the existing IPv4 infrastructure. The IPv6 nodes that act as tunnel endpoints should be dual stack nodes. The tunnel entry-point node encapsulates each IPv6 packet with an IPv4 header and transmits the encapsulated IPv4 packet using the existing IPv4 routing infrastructure. The tunnel exit-point node receives the encapsulated packet, decapsulates it, and passes it to its final destination.

Tunnel configuration can be either manual (configured) or automatic, they primarily differ in how the IPv4 tunnel end-point is determined.

The manually configured tunnels are point-to-point tunnels; the tunnel configuration is done on both the tunnel entry-point node and exit-point node. The encapsulating node determines the IPv4 tunnel endpoint from the configuration information.

The automatic tunnels are point-to-multipoint tunnels, and they require no manual configuration. They use a special IPv6 address with an embedded IPv4 tunnel endpoint address. 6to4 is an example of automatic tunneling mechanism.

Tunneling can be used in different scenarios, as discussed in the following sections.

5.1 Host-to-Host

In a host-to-host configuration, the tunnel endpoints are IPv6/IPv4 (dual stack) hosts that are interconnected by an IPv4 infrastructure. This tunnel configuration enables newly deployed isolated IPv6/IPv4 hosts to communicate with each other by creating an IPv6-in-IPv4 tunnel. In this case, the tunnel spans the entire end-to-end path that the packet takes.

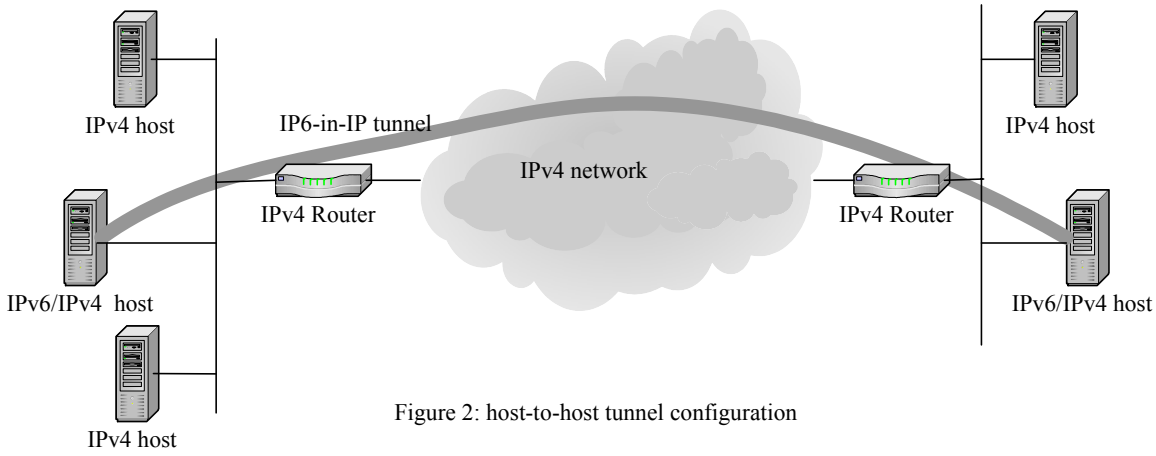

Figure 2: host-to-host tunnel configuration

5.2 Router-to-Router

In a router-to-router tunnel configuration, IPv6/IPv4 routers interconnected by an IPv4 infrastructure can tunnel IPv6 packets between themselves. In this case, the tunnel spans one segment of the end-to-end path the packet takes.

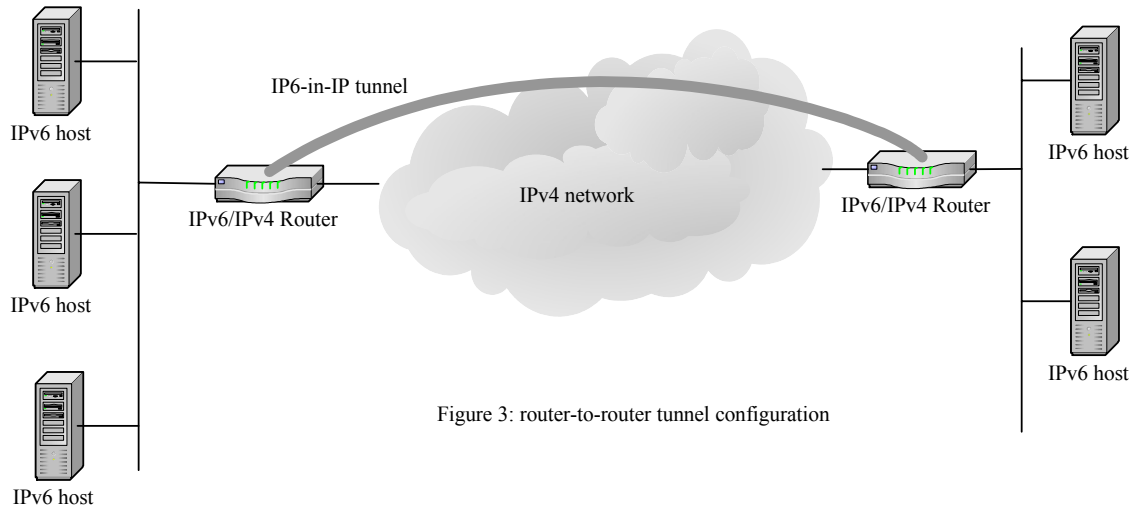

Figure 3: router-to-router tunnel configuration

The router-to-router tunnel configuration is useful to connect two IPv6 domains separated by IPv4 network.

5.3 Host-to-Router and Router-to-Host

In a host-to-router configuration, the tunnel spans the first segment of the packet's end-to-end path. In the router-to-host configuration, the tunnel spans the last segment of the packet's end-to-end path.

The host-to-router and router-to-host tunnel configuration is useful when a whole site needs access to a service running on a dual stack host in an IPv4 network.

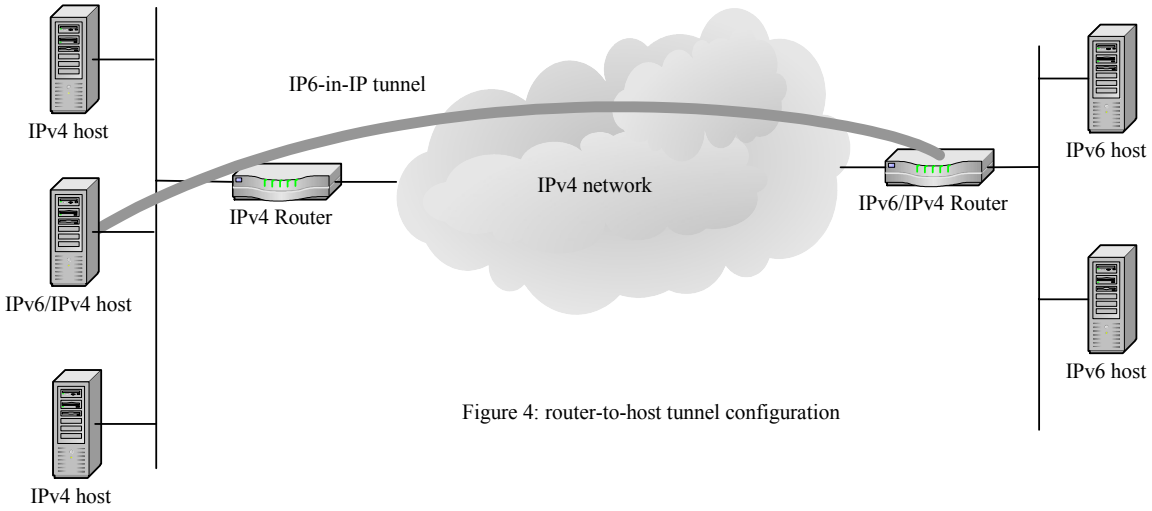

Figure 4: router-to-host tunnel configuration

6. 6to4

6to4 is an automatic router-to-router tunneling mechanism that can be used to provide connectivity between isolated IPv6 sites or hosts across the IPv4 infrastructure and with IPv6-only sites via relay routers. Each 6to4 site should have a site border 6to4 (dual stack) router, which is one endpoint of the 6to4 automatic tunnel.

6to4 further defines an address assignment scheme that allows a site to obtain an unique externally routable prefix if the site has at least one globally unique IPv4 address. The Internet Assigned Number Authority has assigned the unique prefix 2002::/16 for 6to4 mechanism. The site border dual stack router should have at least one global IPv4 address, a 6to4 prefix can be generated by concatenating the 2002:: prefix to the global IPv4 address. For example, if the dual stack router has an IPv4 address 15.13.136.1, then its 6to4 prefix will be 2002:0f0d:8801::/48. The embedded IPv4 address will be used as the tunnel endpoint address by the 6to4 mechanism.

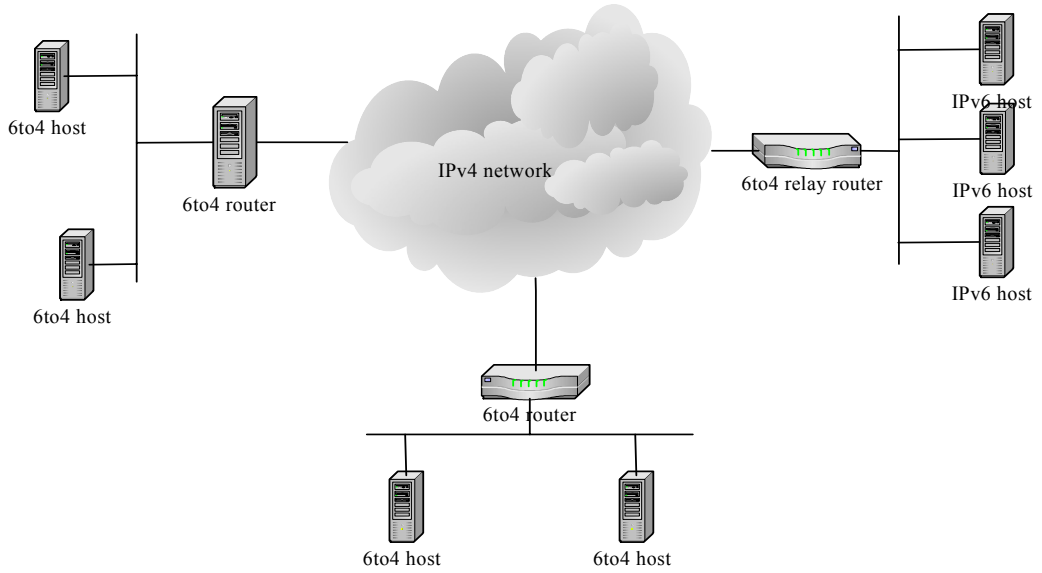

Figure 5: 6to4 network topology

Figure 5 illustrates a 6to4 network topology. 6to4 topology consists of one or more 6to4 hosts in a 6to4 site, a 6to4 border router on the site that has at least one IPv4 connection to the Internet and a 6to4 relay router that is used to connect to native IPv6 site. A 6to4 relay router is different from a 6to4 router in that it has a 6to4 pseudo interface and at least one IPv6 interface.

Within the 6to4 site, the 6to4 router advertises the 6to4 prefix derived from its global IPv4 address. The 6to4 hosts in the site autoconfigure IPv6 addresses using the advertised 6to4 prefix and create a default route to the advertising 6to4 router. The 6to4 hosts within the 6to4 site communicate using IPv6 routing protocols.

IPv6 communication with 6to4 hosts in a different 6to4 site is sent to the 6to4 router using the default route. The 6to4 router will encapsulate the IPv6 packets in an IPv4 header. The IPv4 source address will be the IPv4 address of the 6to4 router and the destination IPv4 address will be the IPv4 address embedded in the destination 6to4 IPv6 address. The encapsulated packet will be sent to the destination 6to4 router using IPv4 routing infrastructure. The destination 6to4 router will decapsulate the packet and will forward the IPv6 packet to its final destination.

In case of communication with native IPv6 hosts, the 6to4 router will send the encapsulated packet to the 6to4 relay router using its default route pointing to the 6to4 relay router.

7. Transition from IPv4 to IPv6

To prepare for a transition from IPv4 to IPv6, the following steps are recommended:

- Develop IP-version independent applications: Existing and new applications should be made protocol independent without specific dependencies on either IPv4 or IPv6. The applications should use version independent data structures and version independent APIs as specified in [2].
- Upgrade the DNS Infrastructure in the organization: IPv6 addresses are stored in DNS using the new record type AAAA. The DNS infrastructure must be upgraded to support AAAA records along with the support for A records.
- Upgrade hosts to dual stack: Hosts must be upgraded to dual stack. This allows the hosts to support both existing IPv4 applications and new version independent applications.
- Upgrade routers to dual stack: Routers must be upgraded to dual stack to support IPv6 routing and transition mechanisms.

8. Deployment Scenario

An organization is expanding its operations and starting a new branch office. The organization is short of globally unique IPv4 addresses and prefers not to use private IPv4 addresses. The organization decides to use IPv6 in the new branch office to take advantage of new IPv6 capabilities and to be in the forefront of new technology.

The main requirement of the organization is its new branch office should be able to communicate with services available in the IPv4 networked main office and other branch offices. To facilitate the communication, the organization is willing to upgrade some nodes in the main office and branch office as long as the upgrade does not disrupt existing operations.

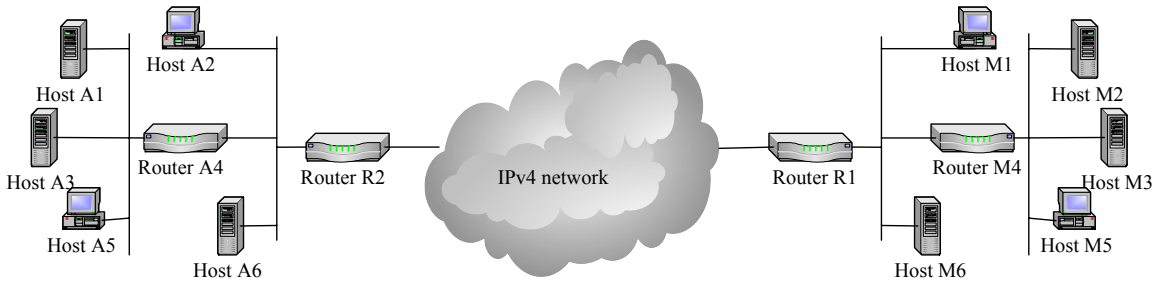

Figure 6: Initial Organization topology

Figure 6 illustrates the initial topology of an organization consisting of a main office and branch office A connected by IPv4.

The organization should perform the following pre-deployment tasks before setting up the new branch office B with native IPv6 support:

- Acquire an IPv6 prefix: The organization should acquire global unicast prefix from its service provider.
- DNS Infrastructure: The organization should upgrade its DNS Infrastructure to support IPv4 A records and IPv6 AAAA records.
- Required applications upgrade: The organization should identify the existing office applications that must be made IPv6-aware so they can communicate with existing peer IPv4 applications in the main office and other branch office and with IPv6 applications in new branch office.
- Required nodes upgrade: The organization should identify which nodes must be upgraded to dual stack in the existing infrastructure to facilitate communication between the new branch office and existing offices.

For example, assume that the employees in the new branch office need access to the payroll application running on server M3 in main office, accounting applications on servers A1 and A3 in branch office A.

To successfully communicate,

- The payroll application and accounting application must be made IPv6-aware.
- The server M3 in the main office, servers A1 and A3 in branch office A should be upgraded to dual stack.
- The router A4 should be upgraded to dual stack.

Figure 7 illustrates the topology of the organization after setting up the new branch office and the required upgrades in the existing infrastructure.

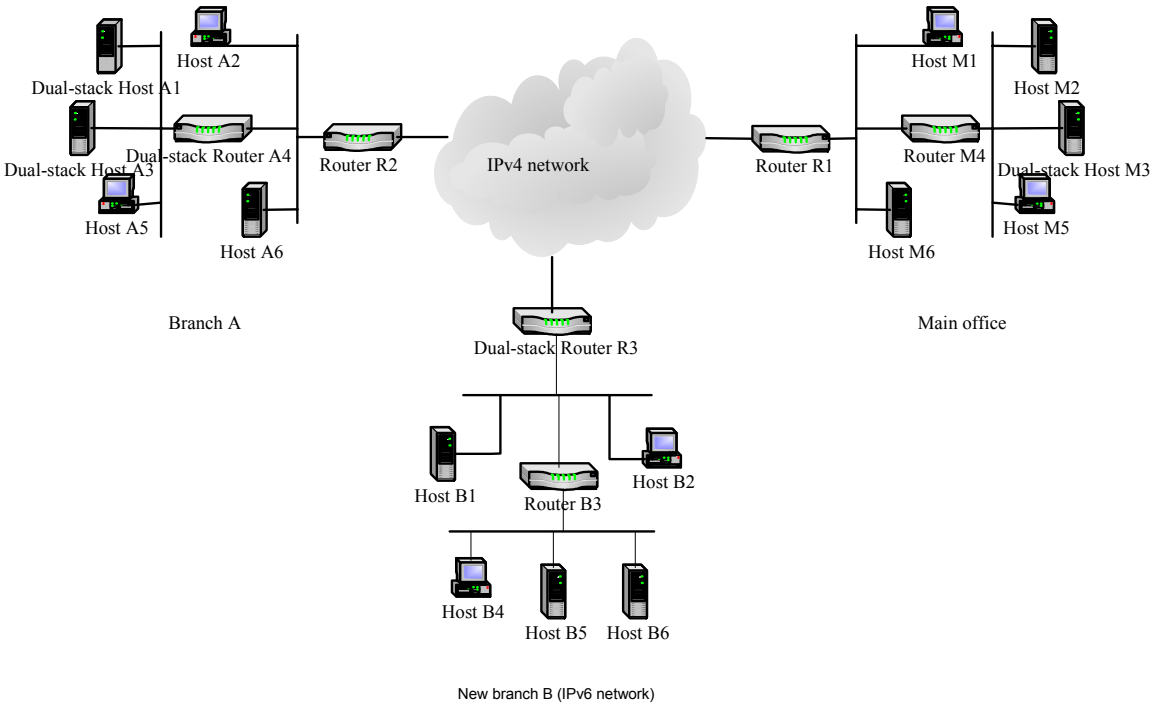

Figure 7: Organization topology with new branch office

8.1 Access to Payroll Application in the Main Office

To enable employees in branch office B to access the payroll application running on server M3 in the main office, a router-to-host tunnel configuration can be used. A router–to-host configured tunnel is setup between router R3 and server M3. M3 is the endpoint for the tunnel from R3 to M3 and R3 is the endpoint for tunnel from M3 to R3.

The communication between hosts in branch office B and server M3 is described below:

1. A host in branch office B sends an IPv6 packet with server M3 as the destination address.
2. The router R3 receives the outgoing IPv6 packet. Based on the tunnel information, it encapsulates the IPv6 packet in an IPv4 packet with the source address set to IPv4 address of R3 and destination address set to the IPv4 address of server M3.
3. The encapsulated packet is forwarded to server M3 using the IPv4 routing infrastructure. When M3 receives the packet, it decapsulates it and since it is also the final destination of the IPv6 packet, it sends the packet up to the application.
4. When server M3 sends the response back, the IPv6 packet is encapsulated in IPv4 packet with the source address set to the IPv4 address of M3 and the destination address set to IPv4 address of R3.
5. The encapsulated packet is received by router R3. R3 decapsulates the packet and forwards the IPv6 packet to its final destination.

8.2 Access to Accounting Application in Branch Office A

To enable employees in the branch office B to access the accounting application running on server A1 and A3, a router-to-router tunnel configuration between router R3 and router A4 can be used. The router A4 is the endpoint for the tunnel from R3 to A4 and R3 is the endpoint for the tunnel from A4 to R3.

The communication between nodes in branch office B and server A1 and A3 happens as described in the previous section. The routers A4 and R3 will do the encapsulation and decapsulation.

9. Related Documents

[1] R. Hinden, S. Deering, “IP version 6 Addressing Architecture”, RFC 2373, July 1998.

[2] R. Gilligan, S. Thomson, J. Bound, J. McCann, W. Stevens, "Basic Socket Interface Extensions for IPv6," RFC 3493, February 2003.

[3] E. Nordmark, R. E. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers," RFC 2893, August 2000.

[4] Carpenter, B and Moore, K, “Connection of IPv6 Domains via IPv4 Clouds”, RFC 3056, February 2001.

[5] HP-UX IPv6 Porting guide.