

Network Working Group
Request for Comments: 4779
Category: Informational

S. Asadullah
A. Ahmed
C. Popoviciu
Cisco Systems
P. Savola
CSC/FUNET
J. Palet
Consulintel
January 2007

ISP IPv6 Deployment Scenarios in Broadband Access Networks

Status of This Memo

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

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

This document provides a detailed description of IPv6 deployment and integration methods and scenarios in today's Service Provider (SP) Broadband (BB) networks in coexistence with deployed IPv4 services. Cable/HFC, BB Ethernet, xDSL, and WLAN are the main BB technologies that are currently deployed, and discussed in this document. The emerging Broadband Power Line Communications (PLC/BPL) access technology is also discussed for completeness. In this document we will discuss main components of IPv6 BB networks, their differences from IPv4 BB networks, and how IPv6 is deployed and integrated in each of these networks using tunneling mechanisms and native IPv6.

Table of Contents

1.	Introduction	4
2.	Common Terminology	5
3.	Core/Backbone Network	5
3.1.	Layer 2 Access Provider Network	5
3.2.	Layer 3 Access Provider Network	6
4.	Tunneling Overview	7
4.1.	Access over Tunnels - Customers with Public IPv4 Addresses	7
4.2.	Access over Tunnels - Customers with Private IPv4 Addresses	8
4.3.	Transition a Portion of the IPv4 Infrastructure	8
5.	Broadband Cable Networks	9
5.1.	Broadband Cable Network Elements	9
5.2.	Deploying IPv6 in Cable Networks	10
5.2.1.	Deploying IPv6 in a Bridged CMTS Network	12
5.2.2.	Deploying IPv6 in a Routed CMTS Network	14
5.2.3.	IPv6 Multicast	23
5.2.4.	IPv6 QoS	24
5.2.5.	IPv6 Security Considerations	24
5.2.6.	IPv6 Network Management	25
6.	Broadband DSL Networks	26
6.1.	DSL Network Elements	26
6.2.	Deploying IPv6 in IPv4 DSL Networks	28
6.2.1.	Point-to-Point Model	29
6.2.2.	PPP Terminated Aggregation (PTA) Model	30
6.2.3.	L2TPv2 Access Aggregation (LAA) Model	33
6.2.4.	Hybrid Model for IPv4 and IPv6 Service	36
6.3.	IPv6 Multicast	38
6.3.1.	ASM-Based Deployments	39
6.3.2.	SSM-Based Deployments	39
6.4.	IPv6 QoS	40
6.5.	IPv6 Security Considerations	41
6.6.	IPv6 Network Management	42
7.	Broadband Ethernet Networks	42
7.1.	Ethernet Access Network Elements	42
7.2.	Deploying IPv6 in IPv4 Broadband Ethernet Networks	43
7.2.1.	Point-to-Point Model	44
7.2.2.	PPP Terminated Aggregation (PTA) Model	46
7.2.3.	L2TPv2 Access Aggregation (LAA) Model	48
7.2.4.	Hybrid Model for IPv4 and IPv6 Service	50
7.3.	IPv6 Multicast	52
7.4.	IPv6 QoS	53
7.5.	IPv6 Security Considerations	54
7.6.	IPv6 Network Management	55

8.	Wireless LAN	55
8.1.	WLAN Deployment Scenarios	55
8.1.1.	Layer 2 NAP with Layer 3 termination at NSP Edge Router	56
8.1.2.	Layer 3 Aware NAP with Layer 3 Termination at Access Router	59
8.1.3.	PPP-Based Model	61
8.2.	IPv6 Multicast	63
8.3.	IPv6 QoS	65
8.4.	IPv6 Security Considerations	65
8.5.	IPv6 Network Management	67
9.	Broadband Power Line Communications (PLC)	67
9.1.	PLC/BPL Access Network Elements	68
9.2.	Deploying IPv6 in IPv4 PLC/BPL	69
9.2.1.	IPv6 Related Infrastructure Changes	69
9.2.2.	Addressing	69
9.2.3.	Routing	70
9.3.	IPv6 Multicast	71
9.4.	IPv6 QoS	71
9.5.	IPv6 Security Considerations	71
9.6.	IPv6 Network Management	71
10.	Gap Analysis	71
11.	Security Considerations	74
12.	Acknowledgements	74
13.	References	74
13.1.	Normative References	74
13.2.	Informative References	76

1. Introduction

This document presents the options available in deploying IPv6 services in the access portion of a BB Service Provider (SP) network - namely Cable/HFC, BB Ethernet, xDSL, WLAN, and PLC/BPL.

This document briefly discusses the other elements of a provider network as well. It provides different viable IPv6 deployment and integration techniques, and models for each of the above-mentioned BB technologies individually. The example list is not exhaustive, but it tries to be representative.

This document analyzes how all the important components of current IPv4-based Cable/HFC, BB Ethernet, xDSL, WLAN, and PLC/BPL networks will behave when IPv6 is integrated and deployed.

The following important pieces are discussed:

- A. Available tunneling options
- B. Devices that would have to be upgraded to support IPv6
- C. Available IPv6 address assignment techniques and their use
- D. Possible IPv6 Routing options and their use
- E. IPv6 unicast and multicast packet transmission
- F. Required IPv6 Quality of Service (QoS) parameters
- G. Required IPv6 Security parameters
- H. Required IPv6 Network Management parameters

It is important to note that the addressing rules provided throughout this document represent an example that follows the current assignment policies and recommendations of the registries. However, they can be adapted to the network and business model needs of the ISPs.

The scope of the document is to advise on the ways of upgrading an existing infrastructure to support IPv6 services. The recommendation to upgrade a device to dual stack does not stop an SP from adding a new device to its network to perform the necessary IPv6 functions discussed. The costs involved with such an approach could be offset by lower impact on the existing IPv4 services.

2. Common Terminology

BB: Broadband

CPE: Customer Premise Equipment

GWR: Gateway Router

ISP: Internet Service Provider

NAP: Network Access Provider

NSP: Network Service Provider

QoS: Quality of Service

SP: Service Provider

3. Core/Backbone Network

This section intends to briefly discuss some important elements of a provider network tied to the deployment of IPv6. A more detailed description of the core network is provided in other documents [RFC4029].

There are two types of networks identified in the Broadband deployments:

- A. Access Provider Network: This network provides the broadband access and aggregates the subscribers. The subscriber traffic is handed over to the Service Provider at Layer 2 or 3.
- B. Service Provider Network: This network provides Intranet and Internet IP connectivity for the subscribers.

The Service Provider network structure beyond the Edge Routers that interface with the Access provider is beyond the scope of this document.

3.1. Layer 2 Access Provider Network

The Access Provider can deploy a Layer 2 network and perform no routing of the subscriber traffic to the SP. The devices that support each specific access technology are aggregated into a highly redundant, resilient, and scalable Layer 2 core. The network core can involve various technologies such as Ethernet, Asynchronous Transfer Mode (ATM), etc. The Service Provider Edge Router connects to the Access Provider core.

This type of network may be transparent to the Layer 3 protocol. Some possible changes may come with the intent of supporting IPv6 provisioning mechanisms, as well as filtering and monitoring IPv6 traffic based on Layer 2 information such as IPv6 Ether Type Protocol ID (0x86DD) or IPv6 multicast specific Media Access Control (MAC) addresses (33:33:xx:xx:xx:xx).

3.2. Layer 3 Access Provider Network

The Access Provider can choose to terminate the Layer 2 domain and route the IP traffic to the Service Provider network. Access Routers are used to aggregate the subscriber traffic and route it over a Layer 3 core to the SP Edge Routers. In this case, the impact of the IPv6 deployment is significant.

The case studies in this document discuss only the relevant network elements of such a network: Customer Premise Equipment, Access Router, and Edge Router. In real networks, the link between the Access Router and the Edge Router involves other routers that are part of the aggregation and the core layer of the Access Provider network.

The Access Provider can forward the IPv6 traffic through its Layer 3 core in three possible ways:

- A. IPv6 Tunneling: As a temporary solution, the Access Provider can choose to use a tunneling mechanism to forward the subscriber IPv6 traffic to the Service Provider Edge Router. This approach has the least impact on the Access Provider network; however, as the number of users increase and the amount of IPv6 traffic grows, the ISP will have to evolve to one of the scenarios listed below.
- B. Native IPv6 Deployment: The Access Provider routers are upgraded to support IPv6 and can become dual stack. In a dual-stack network, an IPv6 Interior Gateway Protocol (IGP), such as OSPFv3 [RFC2740] or IS-IS [ISISv6], is enabled. RFC 4029 [RFC4029] discusses the IGP selection options with their benefits and drawbacks.
- C. MPLS 6PE Deployment [6PE]: If the Access Provider is running MPLS in its IPv4 core, it could use 6PE to forward IPv6 traffic over it. In this case, only a subset of routers close to the edge of the network need to be IPv6 aware. With this approach, BGP becomes important in order to support 6PE.

The 6PE approach has the advantage of having minimal impact on the Access Provider network. Fewer devices need to be upgraded and

configured while the MPLS core continues to switch the traffic, unaware that it transports both IPv4 and IPv6. 6PE should be leveraged only if MPLS is already deployed in the network. At the time of writing this document, a major disadvantage of the 6PE solution is that it does not support multicast IPv6 traffic.

The native approach has the advantage of supporting IPv6 multicast traffic, but it may imply a significant impact on the IPv4 operational network in terms of software configuration and possibly hardware upgrade.

More detailed Core Network deployment recommendations are discussed in other documents [RFC4029]. The handling of IPv6 traffic in the Core of the Access Provider Network will not be discussed for the remainder of this document.

4. Tunneling Overview

If SPs are not able to deploy native IPv6, they might use tunneling-based transition mechanisms to start an IPv6 service offering, and move to native IPv6 deployment at a later time.

Several tunneling mechanisms were developed specifically to transport IPv6 over existing IPv4 infrastructures. Several of them have been standardized and their use depends on the existing SP IPv4 network and the structure of the IPv6 service. The requirements for the most appropriate mechanisms are described in [v6tc] with more updates to follow. Deploying IPv6 using tunneling techniques can imply as little changes to the network as upgrading software on tunnel end points. A Service Provider could use tunneling to deploy IPv6 in the following scenarios:

4.1. Access over Tunnels - Customers with Public IPv4 Addresses

If the customer is a residential user, it can initiate the tunnel directly from the IPv6 capable host to a tunnel termination router located in the NAP or ISP network. The tunnel type used should be decided by the SP, but it should take into consideration its availability on commonly used software running on the host machine. Of the many tunneling mechanisms developed, such as IPv6 Tunnel Broker [RFC3053], Connection of IPv6 Domains via IPv4 Clouds [RFC3056], Generic Packet Tunneling in IPv6 [RFC2473], ISATAP [RFC4214], Basic Transition Mechanisms for IPv6 Hosts and Routers [RFC4213], and Transmission of IPv6 over IPv4 Domains without Explicit Tunnels [RFC2529], some are more popular than the others. At the time of writing this document, the IETF Softwire Working Group was tasked with standardizing a single tunneling protocol [Softwire] for this application.

If the end customer has a GWR installed, then it could be used to originate the tunnel, thus offering native IPv6 access to multiple hosts on the customer network. In this case, the GWR would need to be upgraded to dual stack in order to support IPv6. The GWR can be owned by the customer or by the SP.

4.2. Access over Tunnels - Customers with Private IPv4 Addresses

If the end customer receives a private IPv4 address and needs to initiate a tunnel through Network Address Translation (NAT), techniques like 6to4 may not work since they rely on public IPv4 address. In this case, unless the existing GWRs support protocol-41-forwarding [Protocol41], the end user might have to use tunnels that can operate through NATs (such as Teredo [RFC4380]). Most GWRs support protocol-41-forwarding, which means that hosts can initiate the tunnels - in which case the GWR is not affected by the IPv6 service.

The customer has the option to initiate the tunnel from the device (GWR) that performs the NAT functionality, similar to the GWR scenario discussed in Section 4.1. This will imply hardware replacement or software upgrade and a native IPv6 environment behind the GWR.

It is also worth observing that initiating an IPv6 tunnel over IPv4 through already established IPv4 IPsec sessions would provide a certain level of security to the IPv6 traffic.

4.3. Transition a Portion of the IPv4 Infrastructure

Tunnels can be used to transport the IPv6 traffic across a defined segment of the network. As an example, the customer might connect natively to the Network Access Provider, where a tunnel is used to transit the traffic over IPv4 to the ISP. In this case, the tunnel choice depends on its capabilities (for example, whether or not it supports multicast), routing protocols used (there are several types that can transport Layer 2 messages, such as GRE [RFC2784], L2TPv3 [RFC3931], or pseudowire), manageability, and scalability (dynamic versus static tunnels).

This scenario implies that the access portion of the network has been upgraded to support dual stack, so the savings provided by tunneling in this scenario are very small compared with the previous two scenarios. Depending on the number of sites requiring the service, and considering the expenses required to manage the tunnels (some tunnels are static while others are dynamic [DynamicTunnel]) in this case, the SPs might find the native approach worth the additional investments.

In all the scenarios listed above, the tunnel selection process should consider the IPv6 multicast forwarding capabilities if such service is planned. As an example, 6to4 tunnels do not support IPv6 multicast traffic.

The operation, capabilities, and deployment of various tunnel types have been discussed extensively in the documents referenced earlier as well as in [RFC4213] and [RFC3904]. Details of a tunnel-based deployment are offered in the next section of this document, which discusses the case of Cable Access, where the current Data Over Cable Service Interface Specification (DOCSIS 2.0) [RF-Interface] and prior specifications do not provide support for native IPv6 access. Although Sections 6, 7, 8, and 9 focus on a native IPv6 deployments over DSL, Fiber to the Home (FTTH), wireless, and PLC/BPL and because this approach is fully supported today, tunnel-based solutions are also possible in these cases based on the guidelines of this section and some of the recommendations provided in Section 5.

5. Broadband Cable Networks

This section describes the infrastructure that exists today in cable networks providing BB services to the home. It also describes IPv6 deployment options in these cable networks.

DOCSIS standardizes and documents the operation of data over cable networks. DOCSIS 2.0 and prior specifications have limitations that do not allow for a smooth implementation of native IPv6 transport. Some of these limitations are discussed in this section. For this reason, the IPv6 deployment scenarios discussed in this section for the existing cable networks are tunnel based. The tunneling examples presented here could also be applied to the other BB technologies described in Sections 6, 7, 8, and 9.

5.1. Broadband Cable Network Elements

Broadband cable networks are capable of transporting IP traffic to/from users to provide high speed Internet access and Voice over IP (VoIP) services. The mechanism for transporting IP traffic over cable networks is outlined in the DOCSIS specification [RF-Interface].

Here are some of the key elements of a cable network:

Cable (HFC) Plant: Hybrid Fiber Coaxial plant, used as the underlying transport

CMTS: Cable Modem Termination System (can be a Layer 2 bridging or Layer 3 routing CMTS)

GWR: Residential Gateway Router (provides Layer 3 services to hosts)

Host: PC, notebook, etc., which is connected to the CM or GWR

CM: Cable Modem

ER: Edge Router

MSO: Multiple Service Operator

Data Over Cable Service Interface Specification (DOCSIS): Standards defining how data should be carried over cable networks

Figure 5.1 illustrates the key elements of a Cable Network.




Figure 5.1

5.2. Deploying IPv6 in Cable Networks

One of the motivators for an MSO to deploy IPv6 over its cable network is to ease management burdens. IPv6 can be enabled on the CM, CMTS, and ER for management purposes. Currently portions of the cable infrastructure use IPv4 address space [RFC1918]; however, there is a finite number of those. Thus, IPv6 could have utility in the cable space implemented on the management plane initially and focused

on the data plane for end-user services later. For more details on using IPv6 for management in cable networks, please refer to Section 5.6.1.

There are two different deployment modes in current cable networks: a bridged CMTS environment and a routed CMTS environment. IPv6 can be deployed in both of these environments.

1. Bridged CMTS Network

In this scenario, both the CM and CMTS bridge all data traffic. Traffic to/from host devices is forwarded through the cable network to the ER. The ER then routes traffic through the ISP network to the Internet. The CM and CMTS support a certain degree of Layer 3 functionality for management purposes.

2. Routed CMTS Network

In a routed network, the CMTS forwards IP traffic to/from hosts based on Layer 3 information using the IP source/destination address. The CM acts as a Layer 2 bridge for forwarding data traffic and supports some Layer 3 functionality for management purposes.

Some of the factors that hinder deployment of native IPv6 in current routed and bridged cable networks include:

- A. Changes need to be made to the DOCSIS specification [RF-Interface] to include support for IPv6 on the CM and CMTS. This is imperative for deploying native IPv6 over cable networks.
- B. Problems with IPv6 Neighbor Discovery (ND) on CM and CMTS. In IPv4, these devices rely on Internet Group Multicast Protocol (IGMP) join messages to track membership of hosts that are part of a particular IP multicast group. In order to support ND, a multicast-based process, the CM and CMTS will need to support IGMPv3/Multicast Listener Discovery Version 2 (MLDv2) or v1 snooping.
- C. Classification of IPv6 traffic in the upstream and downstream direction. The CM and CMTS will need to support classification of IPv6 packets in order to give them the appropriate priority and QoS. Service providers that wish to deploy QoS mechanisms also have to support classification of IPv6 traffic.

Due to the above mentioned limitations in deployed cable networks, at the time of writing this document, the only option available for cable operators is to use tunneling techniques in order to transport IPv6 traffic over their current IPv4 infrastructure. The following

sections will cover tunneling and native IPv6 deployment scenarios in more detail.

5.2.1. Deploying IPv6 in a Bridged CMTS Network

In IPv4, the CM and CMTS act as Layer 2 bridges and forward all data traffic to/from the hosts and the ER. The hosts use the ER as their Layer 3 next hop. If there is a GWR behind the CM it can act as a next hop for all hosts and forward data traffic to/from the ER.

When deploying IPv6 in this environment, the CM and CMTS will continue to act as bridging devices in order to keep the transition smooth and reduce operational complexity. The CM and CMTS will need to bridge IPv6 unicast and multicast packets to/from the ER and the hosts. If there is a GWR connected to the CM, it will need to forward IPv6 unicast and multicast traffic to/from the ER.

IPv6 can be deployed in a bridged CMTS network either natively or via tunneling. This section discusses the native deployment model. The tunneling model is similar to ones described in Sections 5.2.2.1 and 5.2.2.2.

Figure 5.2.1 illustrates the IPv6 deployment scenario.




Figure 5.2.1

5.2.1.1. IPv6 Related Infrastructure Changes

In this scenario, the CM and the CMTS bridge all data traffic so they will need to support bridging of native IPv6 unicast and multicast traffic. The following devices have to be upgraded to dual stack: Host, GWR, and ER.

5.2.1.2. Addressing

The proposed architecture for IPv6 deployment includes two components that must be provisioned: the CM and the host. Additionally if there is a GWR connected to the CM, it will also need to be provisioned. The host or the GWR use the ER as their Layer 3 next hop.

5.2.1.2.1. IP Addressing for CM

The CM will be provisioned in the same way as in currently deployed cable networks, using an IPv4 address on the cable interface connected to the MSO network for management functions. During the initialization phase, it will obtain its IPv4 address using Dynamic Host Configuration Protocol (DHCPv4), and download a DOCSIS configuration file identified by the DHCPv4 server.

5.2.1.2.2. IP Addressing for Hosts

If there is no GWR connected to the CM, the host behind the CM will get a /64 prefix via stateless auto-configuration or DHCPv6.

If using stateless auto-configuration, the host listens for routing advertisements (RAs) from the ER. The RAs contain the /64 prefix assigned to the segment. Upon receipt of an RA, the host constructs its IPv6 address by combining the prefix in the RA (/64) and a unique identifier (e.g., its modified EUI-64 (64-bit Extended Unique Identifier) format interface ID).

If DHCPv6 is used to obtain an IPv6 address, it will work in much the same way as DHCPv4 works today. The DHCPv6 messages exchanged between the host and the DHCPv6 server are bridged by the CM and the CMTS.

5.2.1.2.3. IP Addressing for GWR

The GWR can use stateless auto-configuration (RA) to obtain an address for its upstream interface, the link between itself and the ER. This step is followed by a request via DHCP-PD (Prefix Delegation) for a prefix shorter than /64, typically /48 [RFC3177], which in turn is divided into /64s and assigned to its downstream interfaces connecting to the hosts.

5.2.1.3. Data Forwarding

The CM and CMTS must be able to bridge native IPv6 unicast and multicast traffic. The CMTS must provide IP connectivity between hosts attached to CMs, and must do so in a way that meets the expectation of Ethernet-attached customer equipment. In order to do that, the CM and CMTS must forward Neighbor Discovery (ND) packets between ER and the hosts attached to the CM.

Communication between hosts behind different CMs is always forwarded through the CMTS. IPv6 communication between the different sites relies on multicast IPv6 ND [RFC2461] frames being forwarded correctly by the CM and the CMTS.

In order to support IPv6 multicast applications across DOCSIS cable networks, the CM and bridging CMTS need to support IGMPv3/MLDv2 or v1 snooping. MLD is almost identical to IGMP in IPv4, only the name and numbers are changed. MLDv2 is identical to IGMPv3 and also supports ASM (Any-Source Multicast) and SSM (Source-Specific Multicast) service models. Implementation work on CM/CMTS should be minimal because the only significant difference between IPv4 IGMPv3 and IPv6 MLDv2 is the longer addresses in the protocol.

5.2.1.4. Routing

The hosts install a default route that points to the ER or the GWR. No routing protocols are needed on these devices, which generally have limited resources. If there is a GWR present, it will also use static default route to the ER.

The ER runs an IGP such as OSPFv3 or IS-IS. The connected prefixes have to be redistributed. If DHCP-PD is used, with every delegated prefix a static route is installed by the ER. For this reason, the static routes must also be redistributed. Prefix summarization should be done at the ER.

5.2.2. Deploying IPv6 in a Routed CMTS Network

In an IPv4/IPv6 routed CMTS network, the CM still acts as a Layer 2 device and bridges all data traffic between its Ethernet interface and cable interface connected to the cable operator network. The CMTS acts as a Layer 3 router and may also include the ER functionality. The hosts and the GWR use the CMTS as their Layer 3 next hop.

When deploying IPv6, the CMTS/ER will need to either tunnel IPv6 traffic or natively support IPv6.

There are five possible deployment scenarios for IPv6 in a routed CMTS network:

1. IPv4 Cable (HFC) Network

In this scenario, the cable network, including the CM and CMTS, remain IPv4 devices. The host and ER are upgraded to dual stack. This is the easiest way for a cable operator to provide IPv6 service, as no changes are made to the cable network.

2. IPv4 Cable (HFC) Network, GWR at Customer Site

In this case, the cable network, including the CM and CMTS, remain IPv4 devices. The host, GWR, and ER are upgraded to dual stack. This scenario is also easy to deploy since the cable operator just needs to add GWR at the customer site.

3. Dual-stacked Cable (HFC) Network, CM, and CMTS Support IPv6

In this scenario, the CMTS is upgraded to dual stack to support IPv4 and IPv6. Since the CMTS supports IPv6, it can act as an ER as well. The CM will act as a Layer 2 bridge, but will need to bridge IPv6 unicast and multicast traffic. This scenario is not easy to deploy since it requires changes to the DOCSIS specification. The CM and CMTS may require hardware and software upgrades to support IPv6.

4. Dual-stacked Cable (HFC) Network, Standalone GWR, and CMTS Support IPv6

In this scenario there is a stand-alone GWR connected to the CM. Since the IPv6 functionality exists on the GWR, the CM does not need to be dual stack. The CMTS is upgraded to dual stack and it can incorporate the ER functionality. This scenario may also require hardware and software changes on the GWR and CMTS.

5. Dual-stacked Cable (HFC) Network, Embedded GWR/CM, and CMTS Support IPv6

In this scenario, the CM and GWR functionality exists on a single device, which needs to be upgraded to dual stack. The CMTS will also need to be upgraded to a dual-stack device. This scenario is also difficult to deploy in existing cable network since it requires changes on the Embedded GWR/CM and the CMTS.

The DOCSIS specification will also need to be modified to allow native IPv6 support on the Embedded GWR/CM.

5.2.2.1. IPv4 Cable Network, Host, and ER Upgraded to Dual Stack

This is one of the most cost-effective ways for a cable operator to offer IPv6 services to its customers. Since the cable network remains IPv4, there is relatively minimal cost involved in turning up IPv6 service. All IPv6 traffic is exchanged between the hosts and the ER.

Figure 5.2.2.1 illustrates this deployment scenario.

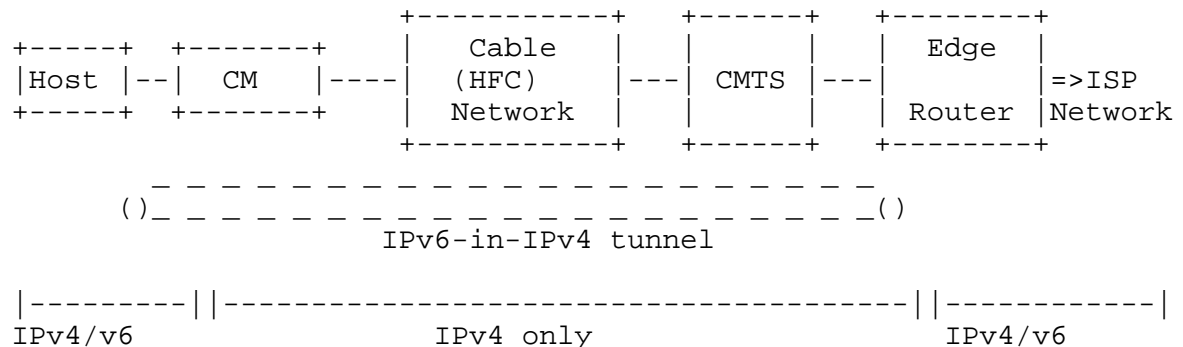


Figure 5.2.2.1

5.2.2.1.1. IPv6 Related Infrastructure Changes

In this scenario, the CM and the CMTS will only need to support IPv4, so no changes need to be made to them or the cable network. The following devices have to be upgraded to dual stack: Host and ER.

5.2.2.1.2. Addressing

The only device that needs to be assigned an IPv6 address at the customer site is the host. Host address assignment can be done in multiple ways. Depending on the tunneling mechanism used, it could be automatic or might require manual configuration.

The host still receives an IPv4 address using DHCPv4, which works the same way in currently deployed cable networks. In order to get IPv6 connectivity, host devices will also need an IPv6 address and a means to communicate with the ER.

5.2.2.1.3. Data Forwarding

All IPv6 traffic will be sent to/from the ER and the host device. In order to transport IPv6 packets over the cable operator IPv4 network, the host and the ER will need to use one of the available IPv6 in IPv4 tunneling mechanisms.

The host will use its IPv4 address to source the tunnel to the ER. All IPv6 traffic will be forwarded to the ER, encapsulated in IPv4 packets. The intermediate IPv4 nodes will forward this traffic as regular IPv4 packets. The ER will need to terminate the tunnel and/or provide other IPv6 services.

5.2.2.1.4. Routing

Routing configuration on the host will vary depending on the tunneling technique used. In some cases, a default or static route might be needed to forward traffic to the next hop.

The ER runs an IGP such as OSPFv3 or ISIS.

5.2.2.2. IPv4 Cable Network, Host, GWR and ER Upgraded to Dual Stack

The cable operator can provide IPv6 services to its customers, in this scenario, by adding a GWR behind the CM. Since the GWR will facilitate all IPv6 traffic between the host and the ER, the cable network, including the CM and CMTS, does not need to support IPv6, and can remain as IPv4 devices.

Figure 5.2.2.2 illustrates this deployment scenario.




Figure 5.2.2.2

5.2.2.2.1. IPv6 Related Infrastructure Changes

In this scenario, the CM and the CMTS will only need to support IPv4, so no changes need to be made to them or the cable network. The following devices have to be upgraded to dual stack: Host, GWR, and ER.

5.2.2.2.2. Addressing

The only devices that need to be assigned an IPv6 address at customer site are the host and GWR. IPv6 address assignment can be done statically at the GWR downstream interface. The GWR will send out RA messages on its downstream interface, which will be used by the hosts to auto-configure themselves with an IPv6 address. The GWR can also configure its upstream interface using RA messages from the ER and use DHCP-PD for requesting a /48 [RFC3177] prefix from the ER. This /48 prefix will be used to configure /64s on hosts connected to the GWR downstream interfaces. The uplink to the ISP network is configured with a /64 prefix as well.

The GWR still receives a global IPv4 address on its upstream interface using DHCPv4, which works the same way in currently deployed cable networks. In order to get IPv6 connectivity to the Internet, the GWR will need to communicate with the ER.

5.2.2.2.3. Data Forwarding

All IPv6 traffic will be sent to/from the ER and the GWR, which will forward IPv6 traffic to/from the host. In order to transport IPv6 packets over the cable operator IPv4 network, the GWR and the ER will need to use one of the available IPv6 in IPv4 tunneling mechanisms. All IPv6 traffic will need to go through the tunnel, once it comes up.

The GWR will use its IPv4 address to source the tunnel to the ER. The tunnel endpoint will be the IPv4 address of the ER. All IPv6 traffic will be forwarded to the ER, encapsulated in IPv4 packets. The intermediate IPv4 nodes will forward this traffic as regular IPv4 packets. In case of 6to4 tunneling, the ER will need to support 6to4 relay functionality in order to provide IPv6 Internet connectivity to the GWR, and hence, the hosts connected to the GWR.

5.2.2.2.4. Routing

Depending on the tunneling technique used, additional configuration might be needed on the GWR and the ER. If the ER is also providing a 6to4 relay service then a default route will need to be added to the GWR pointing to the ER, for all non-6to4 traffic.

If using manual tunneling, the GWR and ER can use static routing or an IGP such as RIPng [RFC2080]. The RIPng updates can be transported over a manual tunnel, which does not work when using 6to4 tunneling since it does not support multicast.

Customer routes can be carried to the ER using RIPng updates. The ER can advertise these routes in its IGP. Prefix summarization should be done at the ER.

If DHCP-PD is used for address assignment, a static route is automatically installed on the ER for each delegated /48 prefix. The static routes need to be redistributed into the IGP at the ER, so there is no need for a routing protocol between the ER and the GWR.

The ER runs an IGP such as OSPFv3 or ISIS.

5.2.2.3. Dual-Stacked Cable (HFC) Network, CM, and CMTS Support IPv6

In this scenario the cable operator can offer native IPv6 services to its customers since the cable network, including the CMTS, supports IPv6. The ER functionality can be included in the CMTS or it can exist on a separate router connected to the CMTS upstream interface. The CM will need to bridge IPv6 unicast and multicast traffic.

Figure 5.2.2.3 illustrates this deployment scenario.




Figure 5.2.2.3

5.2.2.3.1. IPv6 Related Infrastructure Changes

Since the CM still acts as a Layer 2 bridge, it does not need to be dual stack. The CM will need to support bridging of IPv6 unicast and multicast traffic and IGMPv3/MLDv2 or v1 snooping, which requires changes in the DOCSIS specification. In this scenario, the following devices have to be upgraded to dual stack: Host and CMTS/ER.

5.2.2.3.2. Addressing

In cable networks today, the CM receives a private IPv4 address using DHCPv4 for management purposes. In an IPv6 environment, the CM will continue to use an IPv4 address for management purposes. The cable operator can also choose to assign an IPv6 address to the CM for management, but the CM will have to be upgraded to support this functionality.

IPv6 address assignment for the CM and host can be done via DHCP or stateless auto-configuration. If the CM uses an IPv4 address for management, it will use DHCPv4 for its address assignment and the CMTS will need to act as a DHCPv4 relay agent. If the CM uses an IPv6 address for management, it can use DHCPv6, with the CMTS acting as a DHCPv6 relay agent, or the CMTS can be statically configured with a /64 prefix and it can send out RA messages out the cable interface. The CMTS connected to the cable interface can use the RA messages to auto-configure themselves with an IPv6 address. All CMTS connected to the cable interface will be in the same subnet.

The hosts can receive their IPv6 address via DHCPv6 or stateless auto-configuration. With DHCPv6, the CMTS may need to act as a DHCPv6 relay agent and forward DHCP messages between the hosts and the DHCP server. With stateless auto-configuration, the CMTS will be configured with multiple /64 prefixes and send out RA messages to the hosts. If the CMTS is not also acting as an ER, the RA messages will come from the ER connected to the CMTS upstream interface. The CMTS will need to forward the RA messages downstream or act as an ND proxy.

5.2.2.3.3. Data Forwarding

All IPv6 traffic will be sent to/from the CMTS and hosts. Data forwarding will work the same way it works in currently deployed cable networks. The CMTS will forward IPv6 traffic to/from hosts based on the IP source/destination address.

5.2.2.3.4. Routing

No routing protocols are needed between the CMTS and the host since the CM and host are directly connected to the CMTS cable interface. Since the CMTS supports IPv6, hosts will use the CMTS as their Layer 3 next hop.

If the CMTS is also acting as an ER, it runs an IGP such as OSPFv3 or IS-IS.

5.2.2.4. Dual-Stacked Cable (HFC) Network, Stand-Alone GWR, and CMTS Support IPv6

In this case, the cable operator can offer IPv6 services to its customers by adding a GWR between the CM and the host. The GWR will facilitate IPv6 communication between the host and the CMTS/ER. The CMTS will be upgraded to dual stack to support IPv6 and can act as an ER as well. The CM will act as a bridge for forwarding data traffic and does not need to support IPv6.

This scenario is similar to the case described in Section 5.2.2.2. The only difference in this case is that the ER functionality exists on the CMTS instead of on a separate router in the cable operator network.

Figure 5.2.2.4 illustrates this deployment scenario.




Figure 5.2.2.4

5.2.2.4.1. IPv6 Related Infrastructure Changes

Since the CM still acts as a Layer 2 bridge, it does not need to be dual stack, nor does it need to support IPv6. In this scenario, the following devices have to be upgraded to dual stack: Host, GWR, and CMTS/ER.

5.2.2.4.2. Addressing

The CM will still receive a private IPv4 address using DHCPv4, which works the same way in existing cable networks. The CMTS will act as a DHCPv4 relay agent.

The address assignment for the host and GWR happens in a similar manner as described in Section 5.2.2.2.2.

5.2.2.4.3. Data Forwarding

Data forwarding between the host and CMTS/ER is facilitated by the GWR and happens in a similar manner as described in Section 5.2.2.2.3.

5.2.2.4.4. Routing

In this case, routing is very similar to the case described in Section 5.2.2.2.4. Since the CMTS now incorporates the ER functionality, it will need to run an IGP such as OSPFv3 or IS-IS.

5.2.2.5. Dual-Stacked Cable (HFC) Network, Embedded GWR/CM, and CMTS Support IPv6

In this scenario, the cable operator can offer native IPv6 services to its customers since the cable network, including the CM/Embedded GWR and CMTS, supports IPv6. The ER functionality can be included in the CMTS or it can exist on a separate router connected to the CMTS upstream interface. The CM/Embedded GWR acts as a Layer 3 device.

Figure 5.2.2.5 illustrates this deployment scenario.

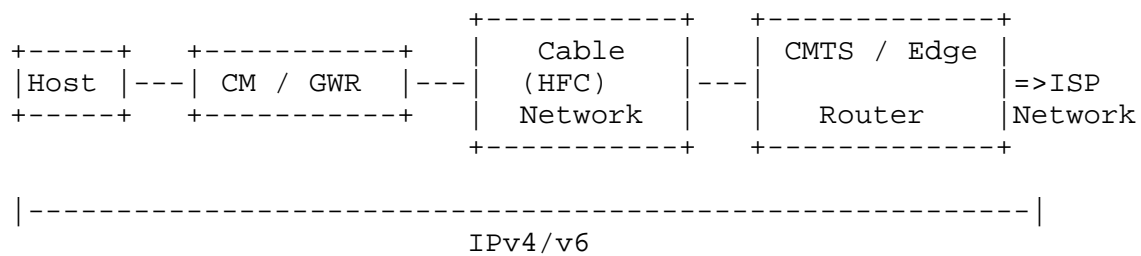


Figure 5.2.2.5

5.2.2.5.1. IPv6 Related Infrastructure Changes

Since the CM/GWR acts as a Layer 3 device, IPv6 can be deployed end-to-end. In this scenario, the following devices have to be upgraded to dual stack: Host, CM/GWR, and CMTS/ER.

5.2.2.5.2. Addressing

Since the CM/GWR is dual stack, it can receive an IPv4 or IPv6 address using DHCP for management purposes. As the GWR functionality is embedded in the CM, it will need an IPv6 address for forwarding data traffic. IPv6 address assignment for the CM/GWR and host can be done via DHCPv6 or DHCP-PD.

If using DHCPv6, the CMTS will need to act as a DHCPv6 relay agent. The host and CM/GWR will receive IPv6 addresses from pools of /64 prefixes configured on the DHCPv6 server. The CMTS will need to glean pertinent information from the DHCP Offer messages, sent from the DHCP server to the DHCP clients (host and CM/GWR), much like it does today in DHCPv4. All CM/GWR connected to the same cable interface on the CMTS belong to the same management /64 prefix. The hosts connected to the same cable interface on the CMTS may belong to different /64 customer prefixes, as the CMTS may have multiple /64 prefixes configured under its cable interfaces.

It is also possible to use DHCP-PD for an IPv6 address assignment. In this case, the CM/GWR will use stateless auto-configuration to assign an IPv6 address to its upstream interface using the /64 prefix sent by the CMTS/ER in an RA message. Once the CM/GWR assigns an IPv6 address to its upstream interface, it will request a /48 [RFC3177] prefix from the CMTS/ER and chop this /48 prefix into /64s for assigning IPv6 addresses to hosts. The uplink to the ISP network is configured with a /64 prefix as well.

5.2.2.5.3. Data Forwarding

The host will use the CM/GWR as the Layer 3 next hop. The CM/GWR will forward all IPv6 traffic to/from the CMTS/ER and hosts. The CMTS/ER will forward IPv6 traffic to/from hosts based on the IP source/destination address.

5.2.2.5.4. Routing

The CM/GWR can use a static default route pointing to the CMTS/ER or it can run a routing protocol such as RIPng or OSPFv3 between itself and the CMTS. Customer routes from behind the CM/GWR can be carried to the CMTS using routing updates.

If DHCP-PD is used for address assignment, a static route is automatically installed on the CMTS/ER for each delegated /48 prefix. The static routes need to be redistributed into the IGP at the CMTS/ER so there is no need for a routing protocol between the CMTS/ER and the GWR.

If the CMTS is also acting as an ER, it runs an IGP such as OSPFv3 or IS-IS.

5.2.3. IPv6 Multicast

In order to support IPv6 multicast applications across DOCSIS cable networks, the CM and bridging CMTS will need to support IGMPv3/MLDv2 or v1 snooping. MLD is almost identical to IGMP in IPv4, only the

name and numbers are changed. MLDv2 is almost identical to IGMPv3 and also supports ASM (Any-Source Multicast) and SSM (Source-Specific Multicast) service models.

SSM is more suited for deployments where the SP intends to provide paid content to the users (video or audio). These types of services are expected to be of primary interest. Moreover, the simplicity of the SSM model often overrides the scalability issues that would be resolved in an ASM model. ASM is, however, an option that is discussed in Section 6.3.1. The Layer 3 CM, GWR, and Layer 3 routed CMTS/ER will need to be enabled with PIM-SSM, which requires the definition and support for IGMPv3/MLDv1 or v2 snooping, in order to track join/leave messages from the hosts. Another option would be for the Layer 3 CM or GWR to support MLD proxy routing. The Layer 3 next hop for the hosts needs to support MLD.

Refer to Section 6.3 for more IPv6 multicast details.

5.2.4. IPv6 QoS

IPv6 will not change or add any queuing/scheduling functionality already existing in DOCSIS specifications. But the QoS mechanisms on the CMTS and CM would need to be IPv6 capable. This includes support for IPv6 classifiers, so that data traffic to/from host devices can be classified appropriately into different service flows and be assigned appropriate priority. Appropriate classification criteria would need to be implemented for unicast and multicast traffic.

Traffic classification and marking should be done at the CM for upstream traffic and the CMTS/ER for downstream traffic, in order to support the various types of services: data, voice, and video. The same IPv4 QoS concepts and methodologies should be applied for IPv6 as well.

It is important to note that when traffic is encrypted end-to-end, the traversed network devices will not have access to many of the packet fields used for classification purposes. In these cases, routers will most likely place the packets in the default classes. The QoS design should take into consideration this scenario and try to use mainly IP header fields for classification purposes.

5.2.5. IPv6 Security Considerations

Security in a DOCSIS cable network is provided using Baseline Privacy Plus (BPI+). The only part that is dependent on IP addresses is encrypted multicast. Semantically, multicast encryption would work the same way in an IPv6 environment as in the IPv4 network. However,

appropriate enhancements will be needed in the DOCSIS specification to support encrypted IPv6 multicast.

There are limited changes that have to be done for hosts in order to enhance security. The privacy extensions [RFC3041] for auto-configuration should be used by the hosts. IPv6 firewall functions could be enabled, if available on the host or GWR.

The ISP provides security against attacks that come from its own subscribers, but it could also implement security services that protect its subscribers from attacks sourced from the outside of its network. Such services do not apply at the access level of the network discussed here.

The CMTS/ER should protect the ISP network and the other subscribers against attacks by one of its own customers. For this reason Unicast Reverse Path Forwarding (uRPF) [RFC3704] and Access Control Lists (ACLs) should be used on all interfaces facing subscribers. Filtering should be implemented with regard for the operational requirements of IPv6 [IPv6-Security].

The CMTS/ER should protect its processing resources against floods of valid customer control traffic such as: Router and Neighbor Solicitations, and MLD Requests.

All other security features used with the IPv4 service should be similarly applied to IPv6 as well.

5.2.6. IPv6 Network Management

IPv6 can have many applications in cable networks. MSOs can initially implement IPv6 on the control plane and use it to manage the thousands of devices connected to the CMTS. This would be a good way to introduce IPv6 in a cable network. Later, the MSO can extend IPv6 to the data plane and use it to carry customer traffic as well as management traffic.

5.2.6.1. Using IPv6 for Management in Cable Networks

IPv6 can be enabled in a cable network for management of devices like CM, CMTS, and ER. With the rollout of advanced services like VoIP and Video-over-IP, MSOs are looking for ways to manage the large number of devices connected to the CMTS. In IPv4, an RFC1918 address is assigned to these devices for management purposes. Since there is a finite number of RFC1918 addresses available, it is becoming difficult for MSOs to manage these devices.

By using IPv6 for management purposes, MSOs can scale their network management systems to meet their needs. The CMTS/ER can be configured with a /64 management prefix that is shared among all CMs connected to the CMTS cable interface. Addressing for the CMs can be done via stateless auto-configuration or DHCPv6. Once the CMs receive a /64 prefix, they can configure themselves with an IPv6 address.

If there are devices behind the CM that need to be managed by the MSO, another /64 prefix can be defined on the CMTS/ER. These devices can also use stateless auto-configuration to assign themselves an IPv6 address.

Traffic sourced from or destined to the management prefix should not cross the MSO's network boundaries.

In this scenario, IPv6 will only be used for managing devices on the cable network. The CM will no longer require an IPv4 address for management as described in DOCSIS 3.0 [DOCSIS3.0-Reqs].

5.2.6.2. Updates to MIB Modules/Standards to Support IPv6

The current DOCSIS, PacketCable, and CableHome MIB modules are already designed to support IPv6 objects. In this case, IPv6 will neither add nor change any of the functionality of these MIB modules. The Textual Convention used to represent Structure of Management Information Version 2 (SMIv2) objects representing IP addresses was updated [RFC4001] and a new Textual Convention InetAddressType was added to identify the type of the IP address used for IP address objects in MIB modules.

There are some exceptions; the MIB modules that might need to add IPv6 support are defined in the DOCSIS 3.0 OSSI specification [DOCSIS3.0-OSSI].

6. Broadband DSL Networks

This section describes the IPv6 deployment options in today's high-speed DSL networks.

6.1. DSL Network Elements

Digital Subscriber Line (DSL) broadband services provide users with IP connectivity over the existing twisted-pair telephone lines called the local-loop. A wide range of bandwidth offerings are available depending on the quality of the line and the distance between the Customer Premise Equipment and the DSL Access Multiplexer (DSLAM).

The following network elements are typical of a DSL network:

DSL Modem: It can be a stand-alone device, be incorporated in the host, incorporate router functionalities, and also have the capability to act as a CPE router.

Customer Premise Router (CPR): It is used to provide Layer 3 services for customer premise networks. It is usually used to provide firewalling functions and segment broadcast domains for a small business.

DSL Access Multiplexer (DSLAM): It terminates multiple twisted-pair telephone lines and provides aggregation to BRAS.

Broadband Remote Access Server (BRAS): It aggregates or terminates multiple Permanent Virtual Circuits (PVCs) corresponding to the subscriber DSL circuits.

Edge Router (ER): It provides the Layer 3 interface to the ISP network.

Figure 6.1 depicts all the network elements mentioned.

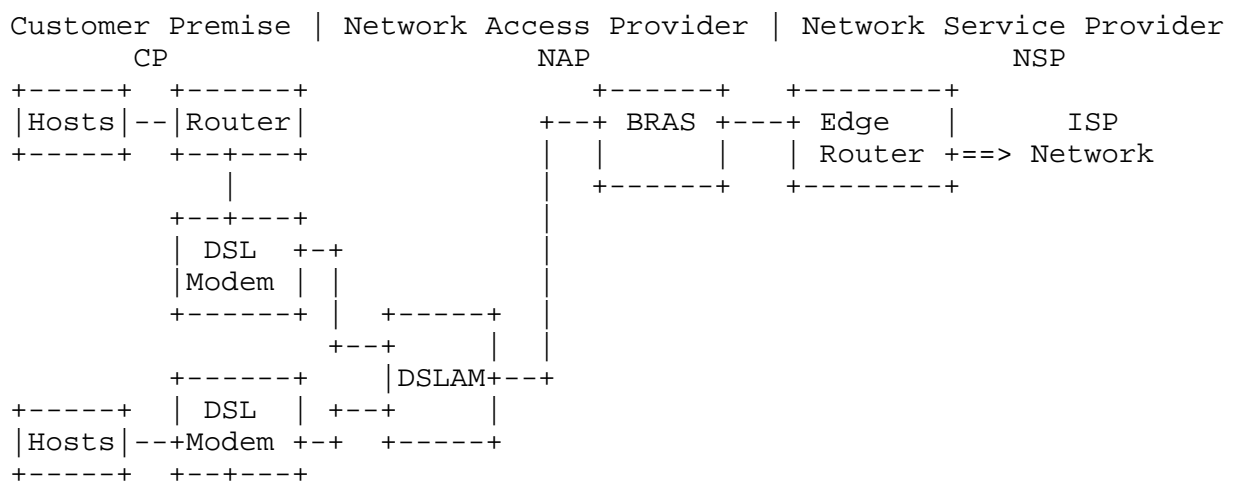


Figure 6.1

6.2. Deploying IPv6 in IPv4 DSL Networks

There are three main design approaches to providing IPv4 connectivity over a DSL infrastructure:

1. Point-to-Point Model: Each subscriber connects to the DSLAM over a twisted pair and is provided with a unique PVC that links it to the service provider. The PVCs can be terminated at the BRAS or at the Edge Router. This type of design is not very scalable if the PVCs are not terminated as close as possible to the DSLAM (at the BRAS). In this case, a large number of Layer 2 circuits has to be maintained over a significant portion of the network. The Layer 2 domains can be terminated at the ER in three ways:
 - A. In a common bridge group with a virtual interface that routes traffic out.
 - B. By enabling a Routed Bridged Encapsulation feature, all users could be part of the same subnet. This is the most common deployment approach of IPv4 over DSL but it might not be the best choice in IPv6 where address availability is not an issue.
 - C. By terminating the PVC at Layer 3, each PVC has its own prefix. This is the approach that seems more suitable for IPv6 and is presented in Section 6.2.1.

None of these ways requires that the CPE (DSL modem) be upgraded.

2. PPP Terminated Aggregation (PTA) Model: PPP sessions are opened between each subscriber and the BRAS. The BRAS terminates the PPP sessions and provides Layer 3 connectivity between the subscriber and the ISP. This model is presented in Section 6.2.2.
3. Layer 2 Tunneling Protocol (L2TP) Access Aggregation (LAA) Model: PPP sessions are opened between each subscriber and the ISP Edge Router. The BRAS tunnels the subscriber PPP sessions to the ISP by encapsulating them into L2TPv2 [RFC2661] tunnels. This model is presented in Section 6.2.3.

In aggregation models, the BRAS terminates the subscriber PVCs and aggregates their connections before providing access to the ISP.

In order to maintain the deployment concepts and business models proven and used with existing revenue generating IPv4 services, the IPv6 deployment will match the IPv4 one. This approach is presented

in Sections 6.2.1 - 6.2.3 that describe current IPv4 over DSL broadband access deployments. Under certain circumstances where new service types or service needs justify it, IPv4 and IPv6 network logical architectures could be different as described in Section 6.2.4.

6.2.1. Point-to-Point Model

In this scenario, the Ethernet frames from the Host or the Customer Premise Router are bridged over the PVC assigned to the subscriber.

Figure 6.2.1 describes the protocol architecture of this model.




Figure 6.2.1

6.2.1.1. IPv6 Related Infrastructure Changes

In this scenario, the DSL modem and the entire NAP is Layer 3 unaware, so no changes are needed to support IPv6. The following devices have to be upgraded to dual stack: Host, Customer Router (if present), and Edge Router.

6.2.1.2. Addressing

The Hosts or the Customer Routers have the Edge Router as their Layer 3 next hop.

If there is no Customer Router, all the hosts on the subscriber site belong to the same /64 subnet that is statically configured on the Edge Router for that subscriber PVC. The hosts can use stateless auto-configuration or stateful DHCPv6-based configuration to acquire an address via the Edge Router.

However, as manual configuration for each customer is a provisioning challenge, implementers are encouraged to develop mechanism(s) that automatically map the PVC (or some other customer-specific information) to an IPv6 subnet prefix, and advertise the customer-specific prefix to all the customers with minimal configuration.

If a Customer Router is present:

- A. It is statically configured with an address on the /64 subnet between itself and the Edge Router, and with /64 prefixes on the interfaces connecting the hosts on the customer site. This is not a desired provisioning method being expensive and difficult to manage.
- B. It can use its link-local address to communicate with the ER. It can also dynamically acquire, through stateless auto-configuration, the prefix for the link between itself and the ER. The later option allows it to contact a remote DHCPv6 server, if needed. This step is followed by a request via DHCP-PD for a prefix shorter than /64 that, in turn, is divided in /64s and assigned to its downstream interfaces.

The Edge Router has a /64 prefix configured for each subscriber PVC. Each PVC should be enabled to relay DHCPv6 requests from the subscribers to DHCPv6 servers in the ISP network. The PVCs providing access for subscribers that use DHCP-PD as well, have to be enabled to support the feature. The uplink to the ISP network is configured with a /64 prefix as well.

The prefixes used for subscriber links and the ones delegated via DHCP-PD should be planned in a manner that allows as much summarization as possible at the Edge Router.

Other information of interest to the host, such as DNS, is provided through stateful DHCPv6 [RFC3315] and stateless DHCPv6 [RFC3736].

6.2.1.3. Routing

The CPE devices are configured with a default route that points to the Edge Router. No routing protocols are needed on these devices, which generally have limited resources.

The Edge Router runs the IPv6 IGP used in the NSP: OSPFv3 or IS-IS. The connected prefixes have to be redistributed. If DHCP-PD is used, with every delegated prefix a static route is installed by the Edge Router. For this reason, the static routes must also be redistributed. Prefix summarization should be done at the Edge Router.

6.2.2. PPP Terminated Aggregation (PTA) Model

The PTA architecture relies on PPP-based protocols (PPPoA [RFC2364] and PPPoE [RFC2516]). The PPP sessions are initiated by Customer Premise Equipment and are terminated at the BRAS. The BRAS

authorizes the session, authenticates the subscriber, and provides an IP address on behalf of the ISP. The BRAS then does Layer 3 routing of the subscriber traffic to the NSP Edge Router.

When the NSP is also the NAP, the BRAS and NSP Edge Router could be the same piece of equipment and provide the above mentioned functionality.

There are two types of PPP encapsulations that can be leveraged with this model:

A. Connection using PPPoA




Figure 6.2.2.1

The PPP sessions are initiated by the Customer Premise Equipment. The BRAS authenticates the subscriber against a local or a remote database. Once the session is established, the BRAS provides an address and maybe a DNS server to the user; this information is acquired from the subscriber profile or from a DHCP server.

This solution scales better than the Point-to-Point, but since there is only one PPP session per ATM PVC, the subscriber can choose a single ISP service at a time.

B. Connection using PPPoE

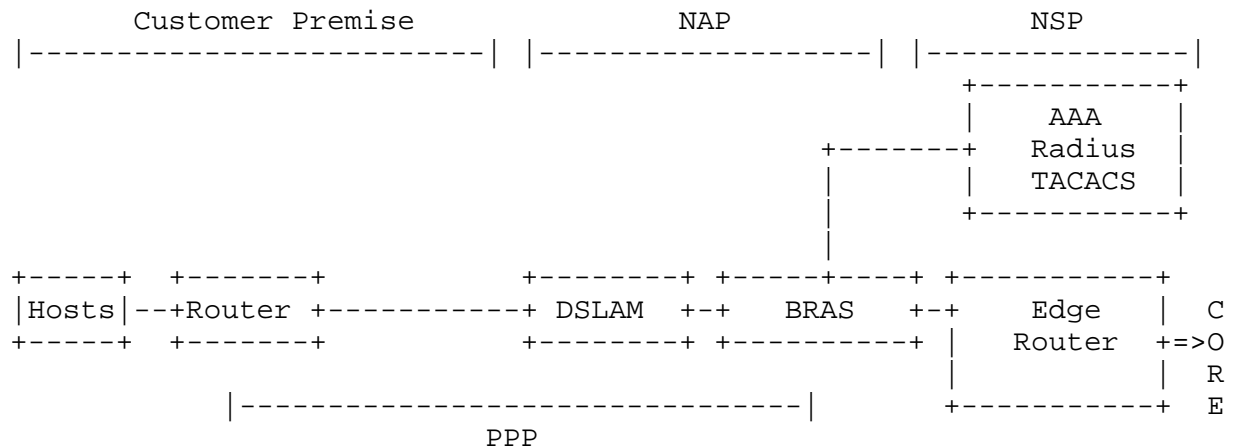


Figure 6.2.2.2

The operation of PPPoE is similar to PPPoA with the exception that with PPPoE multiple sessions can be supported over the same PVC, thus allowing the subscriber to connect to multiple services at the same time. The hosts can initiate the PPPoE sessions as well. It is important to remember that the PPPoE encapsulation reduces the IP MTU available for the customer traffic due to additional headers.

The network design and operation of the PTA model is the same, regardless of the PPP encapsulation type used.

6.2.2.1. IPv6 Related Infrastructure Changes

In this scenario the BRAS is Layer 3 aware and it has to be upgraded to support IPv6. Since the BRAS terminates the PPP sessions it has to support the implementation of these PPP protocols with IPv6. The following devices have to be upgraded to dual stack: Host, Customer Router (if present), BRAS, and Edge Router.

6.2.2.2. Addressing

The BRAS terminates the PPP sessions and provides the subscriber with an IPv6 address from the defined pool for that profile. The subscriber profile for authorization and authentication can be located on the BRAS or on an Authentication, Authorization, and Accounting (AAA) server. The Hosts or the Customer Routers have the BRAS as their Layer 3 next hop.

The PPP session can be initiated by a host or by a Customer Router. In the latter case, once the session is established with the BRAS and an address is negotiated for the uplink to the BRAS, DHCP-PD can be used to acquire prefixes for the Customer Router other interfaces.

The BRAS has to be enabled to support DHCP-PD and to relay the DHCPv6 requests of the hosts on the subscriber sites.

The BRAS has /64 prefixes configured on the link to the Edge router. The Edge Router links are also configured with /64 prefixes to provide connectivity to the rest of the ISP network.

The prefixes used for subscribers and the ones delegated via DHCP-PD should be planned in a manner that allows maximum summarization at the BRAS.

Other information of interest to the host, such as DNS, is provided through stateful [RFC3315] and stateless [RFC3736] DHCPv6.

6.2.2.3. Routing

The CPE devices are configured with a default route that points to the BRAS router. No routing protocols are needed on these devices, which generally have limited resources.

The BRAS runs an IGP to the Edge Router: OSPFv3 or IS-IS. Since the addresses assigned to the PPP sessions are represented as connected host routes, connected prefixes have to be redistributed. If DHCP-PD is used, with every delegated prefix a static route is installed by the Edge Router. For this reason, the static routes must also be redistributed. Prefix summarization should be done at the BRAS.

The Edge Router is running the IGP used in the ISP network: OSPFv3 or IS-IS.

A separation between the routing domains of the ISP and the Access Provider is recommended if they are managed independently. Controlled redistribution will be needed between the Access Provider IGP and the ISP IGP.

6.2.3. L2TPv2 Access Aggregation (LAA) Model

In the LAA model, the BRAS forwards the CPE initiated session to the ISP over an L2TPv2 tunnel established between the BRAS and the Edge Router. In this case, the authentication, authorization, and subscriber configuration are performed by the ISP itself. There are two types of PPP encapsulations that can be leveraged with this model:

A. Connection via PPPoA

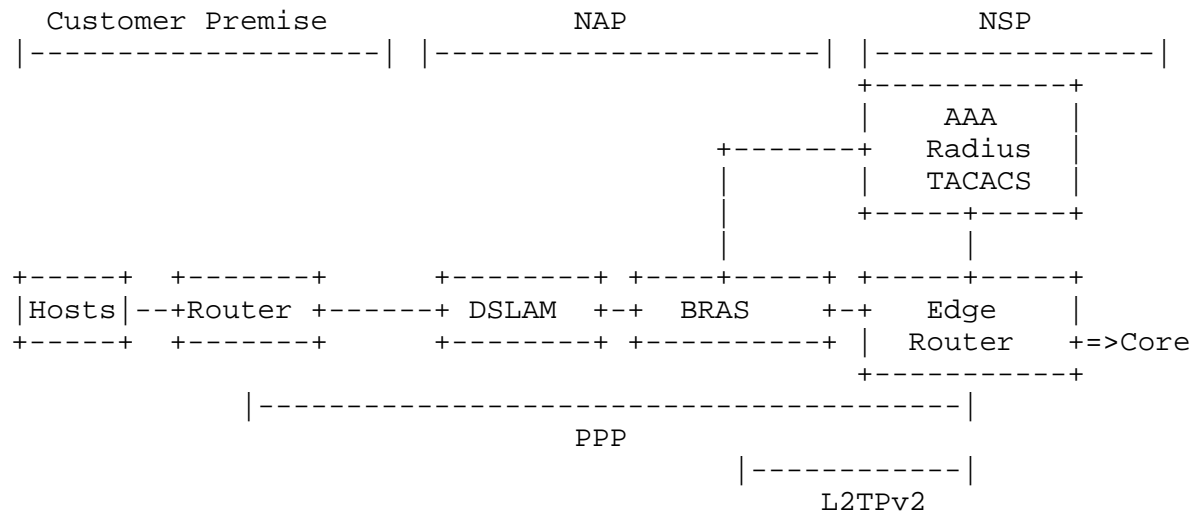


Figure 6.2.3.1

B. Connection via PPPoE

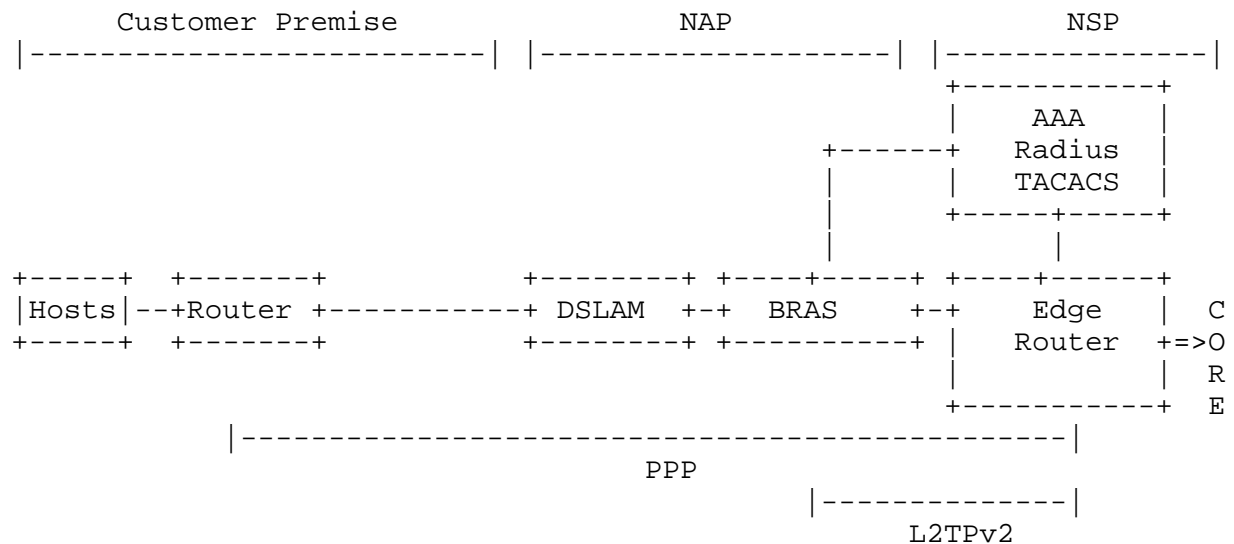


Figure 6.2.3.2

The network design and operation of the PTA model is the same, regardless of the PPP encapsulation type used.

6.2.3.1. IPv6 Related Infrastructure Changes

In this scenario, the BRAS is forwarding the PPP sessions initiated by the subscriber over the L2TPv2 tunnel established to the L2TP Network Server (LNS), the aggregation point in the ISP network. The L2TPv2 tunnel between the L2TP Access Concentrator (LAC) and LNS can run over IPv6 or IPv4. These capabilities have to be supported on the BRAS. The following devices have to be upgraded to dual stack: Host, Customer Router, and Edge Router. If the tunnel is set up over IPv6, then the BRAS must be upgraded to dual stack.

6.2.3.2. Addressing

The Edge Router terminates the PPP sessions and provides the subscriber with an IPv6 address from the defined pool for that profile. The subscriber profile for authorization and authentication can be located on the Edge Router or on an AAA server. The Hosts or the Customer Routers have the Edge Router as their Layer 3 next hop.

The PPP session can be initiated by a host or by a Customer Router. In the latter case, once the session is established with the Edge Router, DHCP-PD can be used to acquire prefixes for the Customer Router interfaces. The Edge Router has to be enabled to support DHCP-PD and to relay the DHCPv6 requests generated by the hosts on the subscriber sites.

The BRAS has a /64 prefix configured on the link to the Edge Router. The Edge Router links are also configured with /64 prefixes to provide connectivity to the rest of the ISP network. Other information of interest to the host, such as DNS, is provided through stateful [RFC3315] and stateless [RFC3736] DHCPv6.

It is important to note here a significant difference between this deployment for IPv6 versus IPv4. In the case of IPv4, the customer router or CPE can end up on any Edge Router (acting as LNS), where the assumption is that there are at least two of them for redundancy purposes. Once authenticated, the customer will be given an address from the IP pool of the ER (LNS) it connected to. This allows the ERs (LNSs) to aggregate the addresses handed out to the customers. In the case of IPv6, an important constraint that likely will be enforced is that the customer should keep its own address, regardless of the ER (LNS) it connects to. This could significantly reduce the prefix aggregation capabilities of the ER (LNS). This is different than the current IPv4 deployment where addressing is dynamic in nature, and the same user can get different addresses depending on the LNS it ends up connecting to.

One possible solution is to ensure that a given BRAS will always connect to the same ER (LNS) unless that LNS is down. This means that customers from a given prefix range will always be connected to the same ER (primary, if up, or secondary, if not). Each ER (LNS) can carry summary statements in their routing protocol configuration for the prefixes for which they are the primary ER (LNS), as well as for the ones for which they are the secondary. This way the prefixes will be summarized any time they become "active" on the ER (LNS).

6.2.3.3. Routing

The CPE devices are configured with a default route that points to the Edge Router that terminates the PPP sessions. No routing protocols are needed on these devices, which generally have limited resources.

The BRAS runs an IPv6 IGP to the Edge Router: OSPFv3 or IS-IS. Different processes should be used if the NAP and the NSP are managed by different organizations. In this case, controlled redistribution should be enabled between the two domains.

The Edge Router is running the IPv6 IGP used in the ISP network: OSPFv3 or IS-IS.

6.2.4. Hybrid Model for IPv4 and IPv6 Service

It was recommended throughout this section that the IPv6 service implementation should map the existing IPv4 one. This approach simplifies manageability and minimizes training needed for personnel operating the network. In certain circumstances such mapping is not feasible. This typically becomes the case when a Service Provider plans to expand its service offering with the new IPv6 deployed infrastructure. If this new service is not well supported in a network design such as the one used for IPv4, then a different design might be used for IPv6.

An example of such circumstances is that of a provider using an LAA design for its IPv4 services. In this case all the PPP sessions are bundled and tunneled across the entire NAP infrastructure which is made of multiple BRAS routers, aggregation routers etc. The end point of these tunnels is the ISP Edge Router. If the provider decides to offer multicast services over such a design, it will face the problem of NAP resources being over utilized. The multicast traffic can be replicated only at the end of the tunnels by the Edge Router and the copies for all the subscribers are carried over the entire NAP.

A Modified Point-to-Point (as described in Section 6.2.4.2) or PTA model is more suitable to support multicast services because the packet replication can be done closer to the destination at the BRAS. Such topology saves NAP resources.

In this sense, IPv6 deployment can be viewed as an opportunity to build an infrastructure that might better support the expansion of services. In this case, an SP using the LAA design for its IPv4 services might choose a modified Point-to-Point or PTA design for IPv6.

6.2.4.1. IPv4 in LAA Model and IPv6 in PTA Model

The coexistence of the two PPP-based models, PTA and LAA, is relatively straightforward. The PPP sessions are terminated on different network devices for the IPv4 and IPv6 services. The PPP sessions for the existing IPv4 service deployed in an LAA model are terminated on the Edge Router. The PPP sessions for the new IPv6 service deployed in a PTA model are terminated on the BRAS.

The logical design for IPv6 and IPv4 in this hybrid model is presented in Figure 6.2.4.1.

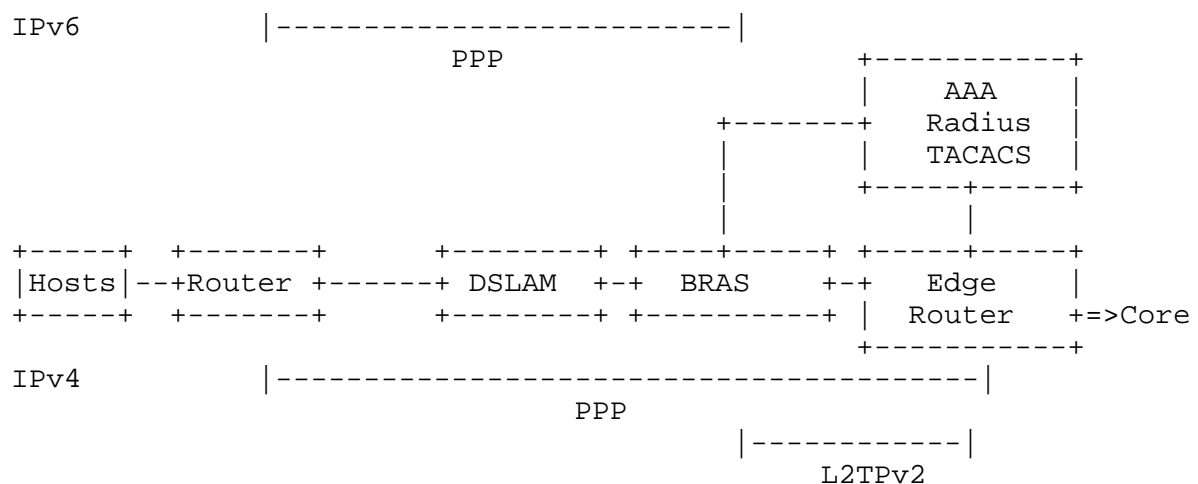


Figure 6.2.4.1

6.2.4.2. IPv4 in LAA Model and IPv6 in Modified Point-to-Point Model

In this particular scenario the Point-to-Point model used for the IPv6 service is a modified version of the model described in section 6.2.1.

For the IPv4 service in the LAA model, the PVCs are terminated on the BRAS and PPP sessions are terminated on the Edge Router (LNS). For IPv6 service in the Point-to-Point model, the PVCs are terminated at the Edge Router as described in Section 6.2.1. In this hybrid model, the Point-to-Point link could be terminated on the BRAS, a NAP-owned device. The IPv6 traffic is then routed through the NAP network to the NSP. In order to have this hybrid model, the BRAS has to be upgraded to a dual-stack router. The functionalities of the Edge Router, as described in Section 6.2.1, are now implemented on the BRAS.

The other aspect of this deployment model is the fact that the BRAS has to be capable of distinguishing between the IPv4 PPP traffic that has to be bridged across the L2TPv2 tunnel and the IPv6 packets that have to be routed to the NSP. The IPv6 Routing and Bridging Encapsulation (RBE) has to be enabled on all interfaces with PVCs supporting both IPv4 and IPv6 services in this hybrid design.

The logical design for IPv6 and IPv4 in this hybrid model is presented in Figure 6.2.4.2.

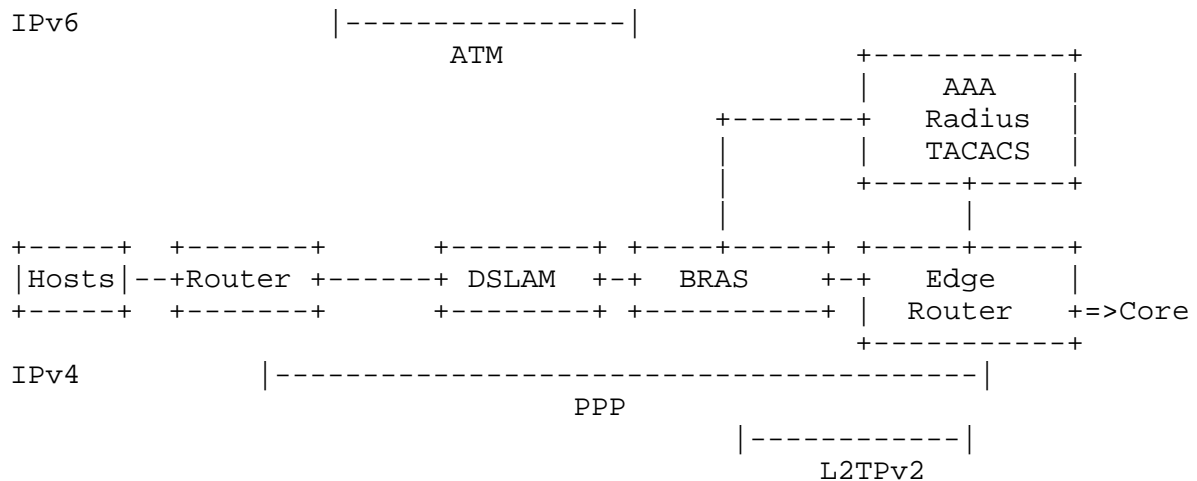


Figure 6.2.4.2

6.3. IPv6 Multicast

The deployment of IPv6 multicast services relies on MLD, identical to IGMP in IPv4 and on PIM for routing. ASM (Any Source Multicast) and SSM (Single Source Multicast) service models operate almost the same as in IPv4. Both have the same benefits and disadvantages as in IPv4. Nevertheless, the larger address space and the scoped address architecture provide major benefits for multicast IPv6. Through RFC 3306, the large address space provides the means to assign global

multicast group addresses to organizations or users that were assigned unicast prefixes. It is a significant improvement with respect to the IPv4 GLOP mechanism [RFC3180].

This facilitates the deployment of multicast services. The discussion of this section applies to all the multicast sections in the document.

6.3.1. ASM-Based Deployments

Any Source Multicast (ASM) is useful for Service Providers that intend to support the forwarding of multicast traffic of their customers. It is based on the Protocol Independent Multicast - Sparse Mode (PIM-SM) protocol and it is more complex to manage because of the use of Rendezvous Points (RPs). With IPv6, static RP and Bootstrap Router [BSR] can be used for RP-to-group mapping similar to IPv4. Additionally, the larger IPv6 address space allows for building up of group addresses that incorporate the address of the RP. This RP-to-group mapping mechanism is called Embedded RP and is specific to IPv6.

In inter-domain deployments, Multicast Source Discovery Protocol (MSDP) [RFC3618] is an important element of IPv4 PIM-SM deployments. MSDP is meant to be a solution for the exchange of source registration information between RPs in different domains. This solution was intended to be temporary. This is one of the reasons why it was decided not to implement MSDP in IPv6 [IPv6-Multicast].

For multicast reachability across domains, Embedded RP can be used. As Embedded RP provides roughly the same capabilities as MSDP, but in a slightly different way, the best management practices for ASM multicast with embedded RP still remain to be developed.

6.3.2. SSM-Based Deployments

Based on PIM-SSM, the Source-Specific Multicast deployments do not need an RP or related protocols (such as BSR or MSDP), but rely on the listeners to know the source of the multicast traffic they plan to receive. The lack of RP makes SSM not only simpler to operate, but also robust; it is not impacted by RP failures or inter-domain constraints. It also has a higher level of security (no RP to be targeted by attacks). For more discussions on the topic of IPv6 multicast, see [IPv6-Multicast].

The typical multicast service offered for residential and very small businesses is video/audio streaming, where the subscriber joins a multicast group and receives the content. This type of service model is well supported through PIM-SSM which is very simple and easy to

manage. PIM-SSM has to be enabled throughout the SP network. MLDv2 is required for PIM-SSM support. Vendors can choose to implement features that allow routers to map MLDv1 group joins to predefined sources.

Subscribers might use a set-top box that is responsible for the control piece of the multicast service (does group joins/leaves). The subscriber hosts can also join desired multicast groups as long as they are enabled to support MLDv1 or MLDv2. If a customer premise router is used, then it has to be enabled to support MLDv1 and MLDv2 in order to process the requests of the hosts. It has to be enabled to support PIM-SSM in order to send PIM joins/leaves up to its Layer 3 next hop whether it is the BRAS or the Edge Router. When enabling this functionality on a CPR, its limited resources should be taken into consideration. Another option would be for the CPR to support MLD proxy routing.

The router that is the Layer 3 next hop for the subscriber (BRAS in the PTA model or the Edge Router in the LAA and Point-to-Point model) has to be enabled to support MLDv1 and MLDv2 in order to process the requests coming from subscribers without CPRs. It has to be enabled for PIM-SSM in order to receive joins/leaves from customer routers and send joins/leaves to the next hop towards the multicast source (Edge Router or the NSP core).

MLD authentication, authorization and accounting are usually configured on the Edge Router in order to enable the ISP to control the subscriber access of the service and do billing for the content provided. Alternative mechanisms that would support these functions should be investigated further.

6.4. IPv6 QoS

The QoS configuration is particularly relevant on the router that represents the Layer 3 next hop for the subscriber (BRAS in the PTA model or the Edge Router in the LAA and Point-to-Point model) in order to manage resources shared amongst multiple subscribers, possibly with various service level agreements.

In the DSL infrastructure, it is expected that there is already a level of traffic policing and shaping implemented for IPv4 connectivity. This is implemented throughout the NAP and is beyond the scope of this document.

On the BRAS or the Edge Router, the subscriber-facing interfaces have to be configured to police the inbound customer traffic and shape the traffic outbound to the customer based on the service level agreements (SLAs). Traffic classification and marking should also be

done on the router closest (at Layer 3) to the subscriber in order to support the various types of customer traffic (data, voice, and video) and to optimally use the infrastructure resources. Each provider (NAP, NSP) could implement their own QoS policies and services so that reclassification and marking might be performed at the boundary between the NAP and the NSP, in order to make sure the traffic is properly handled by the ISP. The same IPv4 QoS concepts and methodologies should be applied with IPv6 as well.

It is important to note that when traffic is encrypted end-to-end, the traversed network devices will not have access to many of the packet fields used for classification purposes. In these cases, routers will most likely place the packets in the default classes. The QoS design should take into consideration this scenario and try to use mainly IP header fields for classification purposes.

6.5. IPv6 Security Considerations

There are limited changes that have to be done for CPEs in order to enhance security. The privacy extensions for auto-configuration [RFC3041] should be used by the hosts. ISPs can track the prefixes it assigns to subscribers relatively easily. If, however, the ISPs are required by regulations to track their users at a /128 address level, the privacy extensions may be implemented in parallel with network management tools that could provide traceability of the hosts. IPv6 firewall functions should be enabled on the hosts or CPR, if present.

The ISP provides security against attacks that come from its own subscribers but it could also implement security services that protect its subscribers from attacks sourced from the outside of its network. Such services do not apply at the access level of the network discussed here.

The device that is the Layer 3 next hop for the subscribers (BRAS or Edge Router) should protect the network and the other subscribers against attacks by one of the provider customers. For this reason, uRPF and ACLs should be used on all interfaces facing subscribers. Filtering should be implemented with regard for the operational requirements of IPv6 [IPv6-Security].

The BRAS and the Edge Router should protect their processing resources against floods of valid customer control traffic such as: Router and Neighbor Solicitations, and MLD Requests. Rate limiting should be implemented on all subscriber-facing interfaces. The emphasis should be placed on multicast-type traffic, as it is most often used by the IPv6 control plane.

All other security features used with the IPv4 service should be similarly applied to IPv6 as well.

6.6. IPv6 Network Management

The necessary instrumentation (such as MIB modules, NetFlow Records, etc.) should be available for IPv6.

Usually, NSPs manage the edge routers by SNMP. The SNMP transport can be done over IPv4 if all managed devices have connectivity over both IPv4 and IPv6. This would imply the smallest changes to the existing network management practices and processes. Transport over IPv6 could also be implemented, and it might become necessary if IPv6 only islands are present in the network. The management applications may be running on hosts belonging to the NSP core network domain. Network Management Applications should handle IPv6 in a similar fashion to IPv4; however, they should also support features specific to IPv6 (such as neighbor monitoring).

In some cases, service providers manage equipment located on customers' LANs. The management of equipment at customers' LANs is out of scope of this memo.

7. Broadband Ethernet Networks

This section describes the IPv6 deployment options in currently deployed Broadband Ethernet Access Networks.

7.1. Ethernet Access Network Elements

In environments that support the infrastructure deploying RJ-45 or fiber (Fiber to the Home (FTTH) service) to subscribers, 10/100 Mbps Ethernet broadband services can be provided. Such services are generally available in metropolitan areas in multi-tenant buildings where an Ethernet infrastructure can be deployed in a cost-effective manner. In such environments, Metro-Ethernet services can be used to provide aggregation and uplink to a Service Provider.

The following network elements are typical of an Ethernet network:

Access Switch: It is used as a Layer 2 access device for subscribers.

Customer Premise Router: It is used to provide Layer 3 services for customer premise networks.

Aggregation Ethernet Switches: Aggregates multiple subscribers.

Broadband Remote Access Server (BRAS)

Edge Router (ER)

- B. PPP Terminated Aggregation (PTA) Model: PPP sessions are opened between each subscriber and the BRAS. The BRAS terminates the PPP sessions and provides Layer 3 connectivity between the subscriber and the ISP.

This model is presented in Section 7.2.2.

- C. L2TPv2 Access Aggregation (LAA) Model: PPP sessions are opened between each subscriber and the ISP termination devices. The BRAS tunnels the subscriber PPP sessions to the ISP by encapsulating them into L2TPv2 tunnels.

This model is presented in Section 7.2.3.

In aggregation models the BRAS terminates the subscriber VLANs and aggregates their connections before providing access to the ISP.

In order to maintain the deployment concepts and business models proven and used with existing revenue generating IPv4 services, the IPv6 deployment will match the IPv4 one. This approach is presented in Sections 7.2.1 - 7.2.3 that describe currently deployed IPv4 over Ethernet broadband access deployments. Under certain circumstances where new service types or service needs justify it, IPv4 and IPv6 network architectures could be different as described in Section 7.2.4.

7.2.1. Point-to-Point Model

In this scenario, the Ethernet frames from the Host or the Customer Premise Router are bridged over the VLAN assigned to the subscriber.

Figure 7.2.1 describes the protocol architecture of this model.




Figure 7.2.1

7.2.1.1. IPv6 Related Infrastructure Changes

In this scenario, the Access Switch is on the customer site and the entire NAP is Layer 3 unaware, so no changes are needed to support IPv6. The following devices have to be upgraded to dual stack: Host, Customer Router, and Edge Router.

The Access switches might need upgrades to support certain IPv6-related features such as MLD Snooping.

7.2.1.2. Addressing

The Hosts or the Customer Routers have the Edge Router as their Layer 3 next hop. If there is no Customer Router all the hosts on the subscriber site belong to the same /64 subnet that is statically configured on the Edge Router for that subscriber VLAN. The hosts can use stateless auto-configuration or stateful DHCPv6-based configuration to acquire an address via the Edge Router.

However, as manual configuration for each customer is a provisioning challenge, implementations are encouraged to develop mechanism(s) that automatically map the VLAN (or some other customer-specific information) to an IPv6 subnet prefix, and advertise the customer-specific prefix to all the customers with minimal configuration.

If a Customer Router is present:

- A. It is statically configured with an address on the /64 subnet between itself and the Edge Router, and with /64 prefixes on the interfaces connecting the hosts on the customer site. This is not a desired provisioning method, being expensive and difficult to manage.
- B. It can use its link-local address to communicate with the ER. It can also dynamically acquire, through stateless auto-configuration, the address for the link between itself and the ER. This step is followed by a request via DHCP-PD for a prefix shorter than /64 that in turn is divided in /64s and assigned to its interfaces connecting the hosts on the customer site.

The Edge Router has a /64 prefix configured for each subscriber VLAN. Each VLAN should be enabled to relay DHCPv6 requests from the subscribers to DHCPv6 servers in the ISP network. The VLANs providing access for subscribers that use DHCP-PD have to be enabled to support the feature. The uplink to the ISP network is configured with a /64 prefix as well.

The prefixes used for subscriber links and the ones delegated via DHCP-PD should be planned in a manner that allows as much summarization as possible at the Edge Router.

Other information of interest to the host, such as DNS, is provided through stateful [RFC3315] and stateless [RFC3736] DHCPv6.

7.2.1.3. Routing

The CPE devices are configured with a default route that points to the Edge Router. No routing protocols are needed on these devices, which generally have limited resources.

The Edge Router runs the IPv6 IGP used in the NSP: OSPFv3 or IS-IS. The connected prefixes have to be redistributed. If DHCP-PD is used, with every delegated prefix a static route is installed by the Edge Router. For this reason, the static routes must also be redistributed. Prefix summarization should be done at the Edge Router.

7.2.2. PPP Terminated Aggregation (PTA) Model

The PTA architecture relies on PPP-based protocols (PPPoE). The PPP sessions are initiated by Customer Premise Equipment and are terminated at the BRAS. The BRAS authorizes the session, authenticates the subscriber, and provides an IP address on behalf of the ISP. The BRAS then does Layer 3 routing of the subscriber traffic to the NSP Edge Router.

When the NSP is also the NAP, the BRAS and NSP Edge Router could be the same piece of equipment and provide the above mentioned functionality.

The PPPoE logical diagram in an Ethernet Broadband Network is shown in Fig 7.2.2.1.

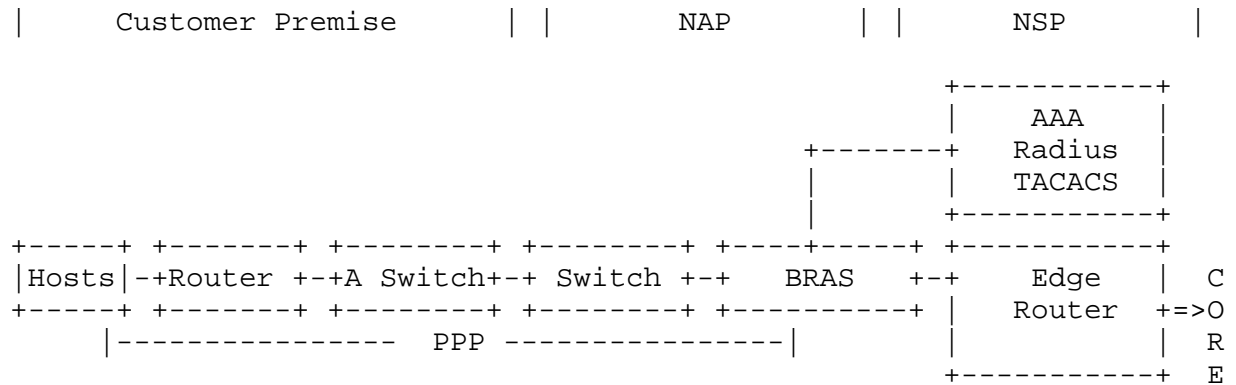


Figure 7.2.2.1

The PPP sessions are initiated by the Customer Premise Equipment (Host or Router). The BRAS authenticates the subscriber against a local or remote database. Once the session is established, the BRAS provides an address and maybe a DNS server to the user; this information is acquired from the subscriber profile or a DHCP server.

This model allows for multiple PPPoE sessions to be supported over the same VLAN, thus allowing the subscriber to connect to multiple services at the same time. The hosts can initiate the PPPoE sessions as well. It is important to remember that the PPPoE encapsulation reduces the IP MTU available for the customer traffic.

7.2.2.2. IPv6 Related Infrastructure Changes

In this scenario, the BRAS is Layer 3 aware and has to be upgraded to support IPv6. Since the BRAS terminates the PPP sessions, it has to support PPPoE with IPv6. The following devices have to be upgraded to dual stack: Host, Customer Router (if present), BRAS and Edge Router.

7.2.2.2. Addressing

The BRAS terminates the PPP sessions and provides the subscriber with an IPv6 address from the defined pool for that profile. The subscriber profile for authorization and authentication can be located on the BRAS, or on an AAA server. The Hosts or the Customer Routers have the BRAS as their Layer 3 next hop.

The PPP session can be initiated by a host or by a Customer Router. In the latter case, once the session is established with the BRAS, DHCP-PD can be used to acquire prefixes for the Customer Router interfaces. The BRAS has to be enabled to support DHCP-PD and to relay the DHCPv6 requests of the hosts on the subscriber sites.

The BRAS has a /64 prefix configured on the link facing the Edge router. The Edge Router links are also configured with /64 prefixes to provide connectivity to the rest of the ISP network.

The prefixes used for subscribers and the ones delegated via DHCP-PD should be planned in a manner that allows maximum summarization at the BRAS.

Other information of interest to the host, such as DNS, is provided through stateful [RFC3315] and stateless [RFC3736] DHCPv6.

7.2.2.3. Routing

The CPE devices are configured with a default route that points to the BRAS router. No routing protocols are needed on these devices, which generally have limited resources.

The BRAS runs an IGP to the Edge Router: OSPFv3 or IS-IS. Since the addresses assigned to the PPP sessions are represented as connected host routes, connected prefixes have to be redistributed. If DHCP-PD is used, with every delegated prefix a static route is installed by the BRAS. For this reason, the static routes must also be redistributed. Prefix summarization should be done at the BRAS.

The Edge Router is running the IGP used in the ISP network: OSPFv3 or IS-IS. A separation between the routing domains of the ISP and the Access Provider is recommended if they are managed independently. Controlled redistribution will be needed between the Access Provider IGP and the ISP IGP.

7.2.3. L2TPv2 Access Aggregation (LAA) Model

In the LAA model, the BRAS forwards the CPE initiated session to the ISP over an L2TPv2 tunnel established between the BRAS and the Edge Router. In this case, the authentication, authorization, and subscriber configuration are performed by the ISP itself.

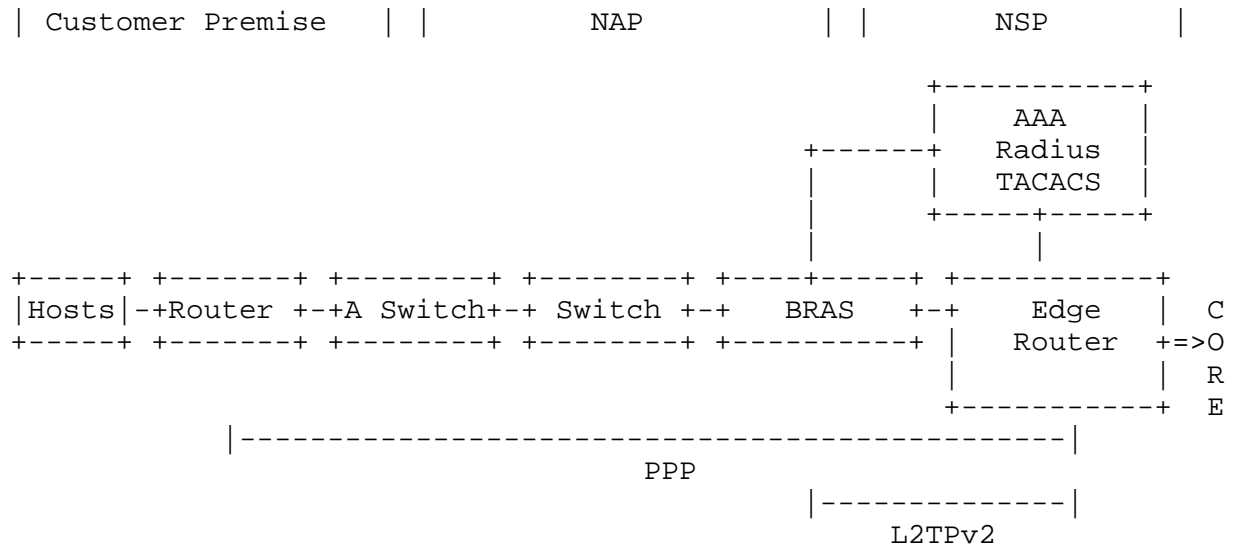


Figure 7.2.3.1

7.2.3.1. IPv6 Related Infrastructure Changes

In this scenario, the BRAS is Layer 3 aware and has to be upgraded to support IPv6. The PPP sessions initiated by the subscriber are forwarded over the L2TPv2 tunnel to the aggregation point in the ISP network. The BRAS (LAC) can aggregate IPv6 PPP sessions and tunnel them to the LNS using L2TPv2. The L2TPv2 tunnel between the LAC and LNS could run over IPv6 or IPv4. These capabilities have to be supported on the BRAS. The following devices have to be upgraded to dual stack: Host, Customer Router (if present), BRAS and Edge Router.

7.2.3.2. Addressing

The Edge Router terminates the PPP sessions and provides the subscriber with an IPv6 address from the defined pool for that profile. The subscriber profile for authorization and authentication can be located on the Edge Router or on an AAA server. The Hosts or the Customer Routers have the Edge Router as their Layer 3 next hop.

The PPP session can be initiated by a host or by a Customer Router. In the latter case, once the session is established with the Edge Router and an IPv6 address is assigned to the Customer Router by the Edge Router, DHCP-PD can be used to acquire prefixes for the Customer Router other interfaces. The Edge Router has to be enabled to support DHCP-PD and to relay the DHCPv6 requests of the hosts on the subscriber sites. The uplink to the ISP network is configured with a /64 prefix as well.

The BRAS has a /64 prefix configured on the link to the Edge Router. The Edge Router links are also configured with /64 prefixes to provide connectivity to the rest of the ISP network.

Other information of interest to the host, such as DNS, is provided through stateful [RFC3315] and stateless [RFC3736] DHCPv6.

The address assignment and prefix summarization issues discussed in Section 6.2.3.2 are relevant in the same way for this media access type as well.

7.2.3.3. Routing

The CPE devices are configured with a default route that points to the Edge Router that terminates the PPP sessions. No routing protocols are needed on these devices, which have limited resources.

The BRAS runs an IPv6 IGP to the Edge Router: OSPFv3 or IS-IS. Different processes should be used if the NAP and the NSP are managed by different organizations. In this case, controlled redistribution should be enabled between the two domains.

The Edge Router is running the IPv6 IGP used in the ISP network: OSPFv3 or IS-IS.

7.2.4. Hybrid Model for IPv4 and IPv6 Service

It was recommended throughout this section that the IPv6 service implementation should map the existing IPv4 one. This approach simplifies manageability and minimizes training needed for personnel operating the network. In certain circumstances, such mapping is not feasible. This typically becomes the case when a Service Provider plans to expand its service offering with the new IPv6 deployed infrastructure. If this new service is not well supported in a network design such as the one used for IPv4, then a different design might be used for IPv6.

An example of such circumstances is that of a provider using an LAA design for its IPv4 services. In this case, all the PPP sessions are bundled and tunneled across the entire NAP infrastructure, which is made of multiple BRAS routers, aggregation routers, etc. The end point of these tunnels is the ISP Edge Router. If the SP decides to offer multicast services over such a design, it will face the problem of NAP resources being over-utilized. The multicast traffic can be replicated only at the end of the tunnels by the Edge Router, and the copies for all the subscribers are carried over the entire NAP.

A Modified Point-to-Point (see Section 7.2.4.2) or a PTA model is more suitable to support multicast services because the packet replication can be done closer to the destination at the BRAS. Such a topology saves NAP resources.

In this sense, IPv6 deployments can be viewed as an opportunity to build an infrastructure that can better support the expansion of services. In this case, an SP using the LAA design for its IPv4 services might choose a modified Point-to-Point or PTA design for IPv6.

7.2.4.1. IPv4 in LAA Model and IPv6 in PTA Model

The coexistence of the two PPP-based models, PTA and LAA, is relatively straightforward. It is a straightforward overlap of the two deployment models. The PPP sessions are terminated on different network devices for the IPv4 and IPv6 services. The PPP sessions for the existing IPv4 service deployed in an LAA model are terminated on the Edge Router. The PPP sessions for the new IPv6 service deployed in a PTA model are terminated on the BRAS.

The logical design for IPv6 and IPv4 in this hybrid model is presented in Figure 7.2.4.1.

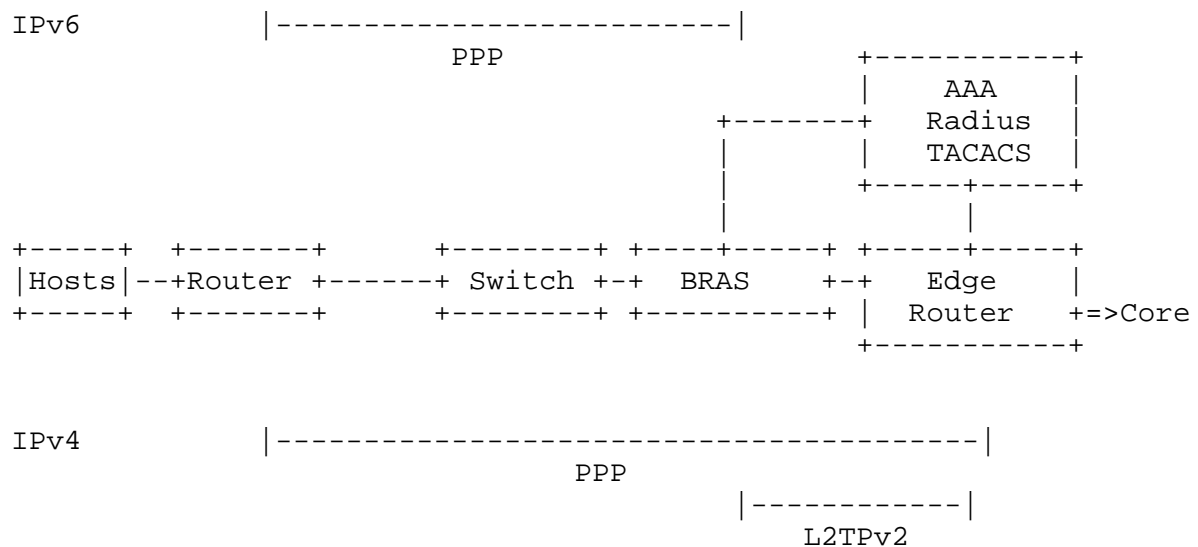


Figure 7.2.4.1

7.2.4.2. IPv4 in LAA Model and IPv6 in Modified Point-to-Point Model

The coexistence of the modified Point-to-Point and the LAA models implies a few specific changes.

For the IPv4 service in LAA model, the VLANs are terminated on the BRAS, and PPP sessions are terminated on the Edge Router (LNS). For the IPv6 service in the Point-to-Point model, the VLANs are terminated at the Edge Router as described in Section 6.2.1. In this hybrid model, the Point-to-Point link could be terminated on the BRAS, a NAP-owned device. The IPv6 traffic is then routed through the NAP network to the NSP. In order to have this hybrid model, the BRAS has to be upgraded to a dual-stack router. The functionalities of the Edge Router, as described in Section 6.2.1, are now implemented on the BRAS.

The logical design for IPv6 and IPv4 in this hybrid model is in Figure 7.2.4.2.




Figure 7.2.4.2

7.3. IPv6 Multicast

The typical multicast services offered for residential and very small businesses are video/audio streaming where the subscriber joins a multicast group and receives the content. This type of service model is well supported through PIM-SSM, which is very simple and easy to manage. PIM-SSM has to be enabled throughout the ISP network. MLDv2 is required for PIM-SSM support. Vendors can choose to implement features that allow routers to map MLDv1 group joins to predefined sources.

Subscribers might use a set-top box that is responsible for the control piece of the multicast service (does group joins/leaves). The subscriber hosts can also join desired multicast groups as long as they are enabled to support MLDv1 or MLDv2. If a CPR is used, then it has to be enabled to support MLDv1 and MLDv2 in order to process the requests of the hosts. It has to be enabled to support PIM-SSM in order to send PIM joins/leaves up to its Layer 3 next hop whether it is the BRAS or the Edge Router. When enabling this functionality on a CPR, its limited resources should be taken into consideration. Another option would be for the CPR to support MLD proxy routing. MLD snooping or similar Layer 2 multicast-related protocols could be enabled on the NAP switches.

The router that is the Layer 3 next hop for the subscriber (BRAS in the PTA model or the Edge Router in the LAA and Point-to-Point model) has to be enabled to support MLDv1 and MLDv2 in order to process the requests coming from subscribers without CPRs. It has to be enabled for PIM-SSM in order to receive joins/leaves from customer routers and send joins/leaves to the next hop towards the multicast source (Edge Router or the NSP core).

MLD authentication, authorization, and accounting are usually configured on the edge router in order to enable the ISP to control the subscriber access of the service and do billing for the content provided. Alternative mechanisms that would support these functions should be investigated further.

Please refer to section 6.3 for more IPv6 multicast details.

7.4. IPv6 QoS

The QoS configuration is particularly relevant on the router that represents the Layer 3 next hop for the subscriber (BRAS in the PTA model or the Edge Router in the LAA and Point-to-Point model) in order to manage resources shared amongst multiple subscribers, possibly with various service level agreements.

On the BRAS or the Edge Router, the subscriber-facing interfaces have to be configured to police the inbound customer traffic and shape the traffic outbound to the customer based on the SLAs. Traffic classification and marking should also be done on the router closest (at Layer 3) to the subscriber in order to support the various types of customer traffic: data, voice, video, and to optimally use the network resources. This infrastructure offers a very good opportunity to leverage the QoS capabilities of Layer 2 devices. Diffserv-based QoS used for IPv4 should be expanded to IPv6.

Each provider (NAP, NSP) could implement their own QoS policies and services so that reclassification and marking might be performed at the boundary between the NAP and the NSP, in order to make sure the traffic is properly handled by the ISP. The same IPv4 QoS concepts and methodologies should be applied for the IPv6 as well.

It is important to note that when traffic is encrypted end-to-end, the traversed network devices will not have access to many of the packet fields used for classification purposes. In these cases, routers will most likely place the packets in the default classes. The QoS design should take into consideration this scenario and try to use mainly IP header fields for classification purposes.

7.5. IPv6 Security Considerations

There are limited changes that have to be done for CPEs in order to enhance security. The privacy extensions [RFC3041] for auto-configuration should be used by the hosts with the same considerations for host traceability as discussed in Section 6.5. IPv6 firewall functions should be enabled on the hosts or Customer Premise Router, if present.

The ISP provides security against attacks that come from its own subscribers, but it could also implement security services that protect its subscribers from attacks sourced from outside its network. Such services do not apply at the access level of the network discussed here.

If any Layer 2 filters for Ethertypes are in place, the NAP must permit the IPv6 Ethertype (0X86DD).

The device that is the Layer 3 next hop for the subscribers (BRAS Edge Router) should protect the network and the other subscribers against attacks by one of the provider customers. For this reason uRPF and ACLs should be used on all interfaces facing subscribers. Filtering should be implemented with regard for the operational requirements of IPv6 [IPv6-Security].

The BRAS and the Edge Router should protect their processing resources against floods of valid customer control traffic such as: Router and Neighbor Solicitations, and MLD Requests. Rate limiting should be implemented on all subscriber-facing interfaces. The emphasis should be placed on multicast-type traffic, as it is most often used by the IPv6 control plane.

All other security features used with the IPv4 service should be similarly applied to IPv6 as well.

7.6. IPv6 Network Management

The necessary instrumentation (such as MIB modules, NetFlow Records, etc.) should be available for IPv6.

Usually, NSPs manage the edge routers by SNMP. The SNMP transport can be done over IPv4 if all managed devices have connectivity over both IPv4 and IPv6. This would imply the smallest changes to the existing network management practices and processes. Transport over IPv6 could also be implemented and it might become necessary if IPv6 only islands are present in the network. The management applications may be running on hosts belonging to the NSP core network domain. Network Management Applications should handle IPv6 in a similar fashion to IPv4; however, they should also support features specific to IPv6 such as neighbor monitoring.

In some cases, service providers manage equipment located on customers' LANs.

8. Wireless LAN

This section provides a detailed description of IPv6 deployment and integration methods in currently deployed wireless LAN (WLAN) infrastructure.

8.1. WLAN Deployment Scenarios

WLAN enables subscribers to connect to the Internet from various locations without the restriction of staying indoors. WLAN is standardized by IEEE 802.11a/b/g.

Figure 8.1 describes the current WLAN architecture.

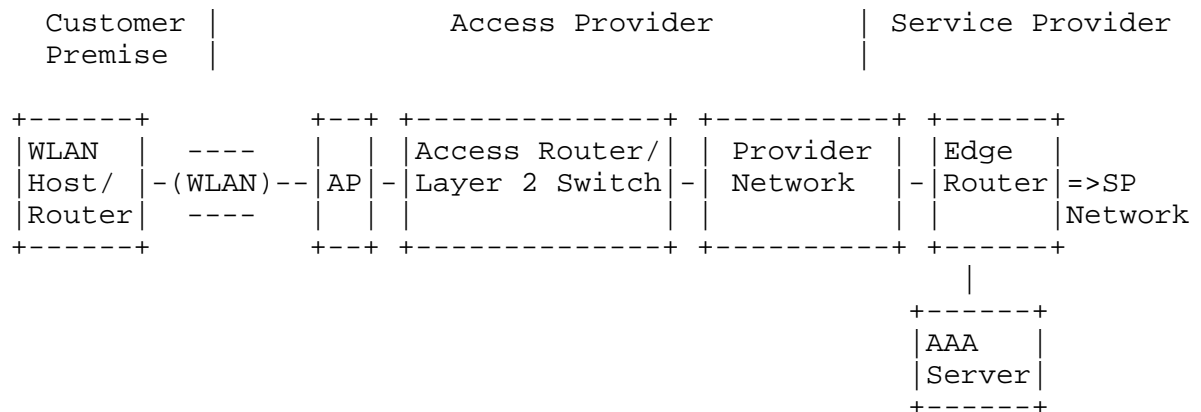


Figure 8.1

The host should have a wireless Network Interface Card (NIC) in order to connect to a WLAN network. WLAN is a flat broadcast network and works in a similar fashion as Ethernet. When a host initiates a connection, it is authenticated by the AAA server located at the SP network. All the authentication parameters (username, password, etc.) are forwarded by the Access Point (AP) to the AAA server. The AAA server authenticates the host; once successfully authenticated, the host can send data packets. The AP is located near the host and acts as a bridge. The AP forwards all the packets coming to/from host to the Edge Router. The underlying connection between the AP and Edge Router could be based on any access layer technology such as HFC/Cable, FTTH, xDSL, etc.

WLANs operate within limited areas known as WiFi Hot Spots. While users are present in the area covered by the WLAN range, they can be connected to the Internet given they have a wireless NIC and required configuration settings in their devices (notebook PCs, PDAs, etc.). Once the user initiates the connection, the IP address is assigned by the SP using DHCPv4. In most of the cases, SP assigns a limited number of public IP addresses to its customers. When the user disconnects the connection and moves to a new WiFi hot spot, the above-mentioned process of authentication, address assignment, and accessing the Internet is repeated.

There are IPv4 deployments where customers can use WLAN routers to connect over wireless to their service provider. These deployment types do not fit in the typical Hot Spot concept, but rather they serve fixed customers. For this reason, this section discusses the WLAN router options as well. In this case, the ISP provides a public IP address and the WLAN Router assigns private addresses [RFC1918] to all WLAN users. The WLAN Router provides NAT functionality while WLAN users access the Internet.

While deploying IPv6 in the above-mentioned WLAN architecture, there are three possible scenarios as discussed below.

- A. Layer 2 NAP with Layer 3 termination at NSP Edge Router
- B. Layer 3 aware NAP with Layer 3 termination at Access Router
- C. PPP-Based Model

8.1.1. Layer 2 NAP with Layer 3 termination at NSP Edge Router

When a Layer 2 switch is present between AP and Edge Router, the AP and Layer 2 switch continues to work as a bridge, forwarding IPv4 and IPv6 packets from WLAN Host/Router to Edge Router and vice versa.

When initiating the connection, the WLAN Host is authenticated by the AAA server located at the SP network. All the parameters related to authentication (username, password, etc.) are forwarded by the AP to the AAA server. The AAA server authenticates the WLAN Hosts, and once the WLAN Host is authenticated and associated successfully with the WLAN AP, it acquires an IPv6 address. Note that the initiation and authentication process is the same as used in IPv4.

Figure 8.1.1 describes the WLAN architecture when a Layer 2 Switch is located between AP and Edge Router.




Figure 8.1.1

8.1.1.1. IPv6 Related Infrastructure Changes

IPv6 will be deployed in this scenario by upgrading the following devices to dual stack: WLAN Host, WLAN Router (if present), and Edge Router.

8.1.1.2. Addressing

When a customer WLAN Router is not present, the WLAN Host has two possible options to get an IPv6 address via the Edge Router.

- A. The WLAN Host can get the IPv6 address from an Edge Router using stateless auto-configuration [RFC2462]. All hosts on the WLAN belong to the same /64 subnet that is statically configured on the Edge Router. The IPv6 WLAN Host may use stateless DHCPv6 for obtaining other information of interest such as DNS, etc.
- B. The IPv6 WLAN Host can use DHCPv6 [RFC3315] to get an IPv6 address from the DHCPv6 server. In this case, the DHCPv6 server would be located in the SP core network, and the Edge Router would simply act as a DHCP Relay Agent. This option is similar

to what is done today in case of DHCPv4. It is important to note that host implementation of stateful auto-configuration is rather limited at this time, and this should be considered if choosing this address assignment option.

When a customer WLAN Router is present, the WLAN Host has two possible options as well for acquiring IPv6 address.

- A. The WLAN Router may be assigned a prefix between /48 and /64 [RFC3177] depending on the SP policy and customer requirements. If the WLAN Router has multiple networks connected to its interfaces, the network administrator will have to configure the /64 prefixes to the WLAN Router interfaces connecting the WLAN Hosts on the customer site. The WLAN Hosts connected to these interfaces can automatically configure themselves using stateless auto-configuration.
- B. The WLAN Router can use its link-local address to communicate with the ER. It can also dynamically acquire through stateless auto-configuration the address for the link between itself and the ER. This step is followed by a request via DHCP-PD for a prefix shorter than /64 that, in turn, is divided in /64s and assigned to its interfaces connecting the hosts on the customer site.

In this option, the WLAN Router would act as a requesting router and the Edge Router would act as a delegating router. Once the prefix is received by the WLAN Router, it assigns /64 prefixes to each of its interfaces connecting the WLAN Hosts on the customer site. The WLAN Hosts connected to these interfaces can automatically configure themselves using stateless auto-configuration. The uplink to the ISP network is configured with a /64 prefix as well.

Usually it is easier for the SPs to stay with the DHCP-PD and stateless auto-configuration model and point the clients to a central server for DNS/domain information, proxy configurations, etc. Using this model, the SP could change prefixes on the fly, and the WLAN Router would simply pull the newest prefix based on the valid/preferred lifetime.

The prefixes used for subscriber links and the ones delegated via DHCP-PD should be planned in a manner that allows maximum summarization at the Edge Router.

Other information of interest to the host, such as DNS, is provided through stateful [RFC3315] and stateless [RFC3736] DHCPv6.

8.1.1.3. Routing

The WLAN Host/Router is configured with a default route that points to the Edge Router. No routing protocols are needed on these devices, which generally have limited resources.

The Edge Router runs the IGP used in the SP network such as OSPFv3 or IS-IS for IPv6. The connected prefixes have to be redistributed. Prefix summarization should be done at the Edge Router. When DHCP-PD is used, the IGP has to redistribute the static routes installed during the process of prefix delegation.

8.1.2. Layer 3 Aware NAP with Layer 3 Termination at Access Router

When an Access Router is present between the AP and Edge Router, the AP continues to work as a bridge, bridging IPv4 and IPv6 packets from WLAN Host/Router to Access Router and vice versa. The Access Router could be part of the SP network or owned by a separate Access Provider.

When the WLAN Host initiates the connection, the AAA authentication and association process with WLAN AP will be similar, as explained in Section 8.1.1.

Figure 8.1.2 describes the WLAN architecture when the Access Router is located between the AP and Edge Router.




Figure 8.1.2

8.1.2.1. IPv6 Related Infrastructure Changes

IPv6 is deployed in this scenario by upgrading the following devices to dual stack: WLAN Host, WLAN Router (if present), Access Router, and Edge Router.

8.1.2.2. Addressing

There are three possible options in this scenario for IPv6 address assignment:

- A. The Edge Router interface facing towards the Access Router is statically configured with a /64 prefix. The Access Router receives/ configures a /64 prefix on its interface facing towards the Edge Router through stateless auto-configuration. The network administrator will have to configure the /64 prefixes to the Access Router interface facing toward the customer premise. The WLAN Host/Router connected to this interface can automatically configure itself using stateless auto-configuration.
- B. This option uses DHCPv6 [RFC3315] for IPv6 prefix assignments to the WLAN Host/Router. There is no use of DHCP PD or stateless auto-configuration in this option. The DHCPv6 server can be located on the Access Router, the Edge Router, or somewhere in the SP network. In this case, depending on where the DHCPv6 server is located, the Access Router or the Edge Router would relay the DHCPv6 requests.
- C. It can use its link-local address to communicate with the ER. It can also dynamically acquire through stateless auto-configuration the address for the link between itself and the ER. This step is followed by a request via DHCP-PD for a prefix shorter than /64 that, in turn, is divided in /64s and assigned to its interfaces connecting the hosts on the customer site.

In this option, the Access Router would act as a requesting router, and the Edge Router would act as a delegating router. Once the prefix is received by the Access Router, it assigns /64 prefixes to each of its interfaces connecting the WLAN Host/Router on the customer site. The WLAN Host/Router connected to these interfaces can automatically configure itself using stateless auto-configuration. The uplink to the ISP network is configured with a /64 prefix as well.

It is easier for the SPs to stay with the DHCP PD and stateless auto-configuration model and point the clients to a central server for DNS/domain information, proxy configurations, and others. Using this model, the provider could change prefixes on the fly, and the Access Router would simply pull the newest prefix based on the valid/preferred lifetime.

As mentioned before, the prefixes used for subscriber links and the ones delegated via DHCP-PD should be planned in a manner that allows the maximum summarization possible at the Edge Router. Other information of interest to the host, such as DNS, is provided through stateful [RFC3315] and stateless [RFC3736] DHCPv6.

8.1.2.3. Routing

The WLAN Host/Router is configured with a default route that points to the Access Router. No routing protocols are needed on these devices, which generally have limited resources.

If the Access Router is owned by an Access Provider, then the Access Router can have a default route, pointing towards the SP Edge Router. The Edge Router runs the IGP used in the SP network such as OSPFv3 or IS-IS for IPv6. The connected prefixes have to be redistributed. If DHCP-PD is used, with every delegated prefix a static route is installed by the Edge Router. For this reason the static routes must be redistributed. Prefix summarization should be done at the Edge Router.

If the Access Router is owned by the SP, then the Access Router will also run IPv6 IGP, and will be part of the SP IPv6 routing domain (OSPFv3 or IS-IS). The connected prefixes have to be redistributed. If DHCP-PD is used, with every delegated prefix a static route is installed by the Access Router. For this reason, the static routes must be redistributed. Prefix summarization should be done at the Access Router.

8.1.3. PPP-Based Model

PPP Terminated Aggregation (PTA) and L2TPv2 Access Aggregation (LAA) models, as discussed in Sections 6.2.2 and 6.2.3, respectively, can also be deployed in IPv6 WLAN environment.

8.1.3.1. PTA Model in IPv6 WLAN Environment

While deploying the PTA model in IPv6 WLAN environment, the Access Router is Layer 3 aware and it has to be upgraded to support IPv6. Since the Access Router terminates the PPP sessions initiated by the WLAN Host/Router, it has to support PPPoE with IPv6.

Figure 8.1.3.1 describes the PTA Model in IPv6 WLAN environment.




Figure 8.1.3.1

8.1.3.1.1. IPv6 Related Infrastructure Changes

IPv6 is deployed in this scenario by upgrading the following devices to dual stack: WLAN Host, WLAN Router (if present), Access Router, and Edge Router.

8.1.3.1.2. Addressing

The addressing techniques described in Section 6.2.2.2 apply to the IPv6 WLAN PTA scenario as well.

8.1.3.1.3. Routing

The routing techniques described in Section 6.2.2.3 apply to the IPv6 WLAN PTA scenario as well.

8.1.3.2. LAA Model in IPv6 WLAN Environment

While deploying the LAA model in IPv6 WLAN environment, the Access Router is Layer 3 aware and has to be upgraded to support IPv6. The PPP sessions initiated by the WLAN Host/Router are forwarded over the L2TPv2 tunnel to the aggregation point in the SP network. The Access Router must have the capability to support L2TPv2 for IPv6.

Figure 8.1.3.2 describes the LAA Model in IPv6 WLAN environment.

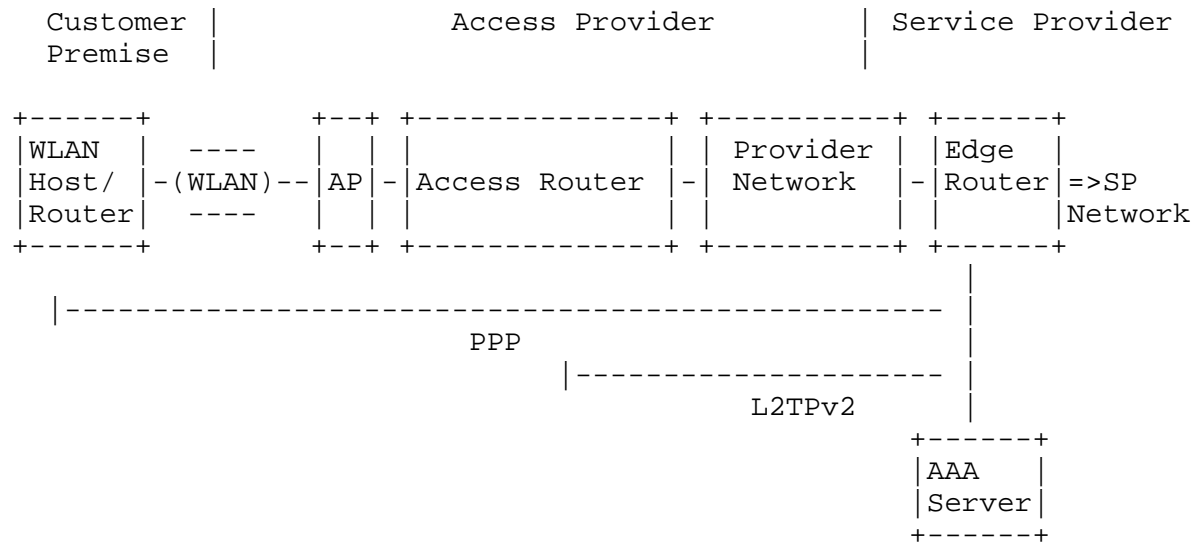


Figure 8.1.3.2

8.1.3.2.1. IPv6 Related Infrastructure Changes

IPv6 is deployed in this scenario by upgrading the following devices to dual stack: WLAN Host, WLAN Router (if present), Access Router, and Edge Router.

8.1.3.2.2. Addressing

The addressing techniques described in Section 6.2.3.2 apply to the IPv6 WLAN LAA scenario as well.

8.1.3.2.3. Routing

The routing techniques described in Section 6.2.3.3 apply to the IPv6 WLAN LAA scenario as well.

8.2. IPv6 Multicast

The typical multicast services offered are video/audio streaming where the IPv6 WLAN Host joins a multicast group and receives the content. This type of service model is well supported through PIM-SSM, which is enabled throughout the SP network. MLDv2 is required for PIM-SSM support. Vendors can choose to implement features that allow routers to map MLDv1 group joins to predefined sources.

It is important to note that in the shared wireless environments, multicast can have a significant bandwidth impact. For this reason, the bandwidth allocated to multicast traffic should be limited and fixed, based on the overall capacity of the wireless specification used in 802.11a, 802.11b, or 802.11g.

The IPv6 WLAN Hosts can also join desired multicast groups as long as they are enabled to support MLDv1 or MLDv2. If WLAN/Access Routers are used, then they have to be enabled to support MLDv1 and MLDv2 in order to process the requests of the IPv6 WLAN Hosts. The WLAN/Access Router also needs to be enabled to support PIM-SSM in order to send PIM joins up to the Edge Router. When enabling this functionality on a WLAN/Access Router, its limited resources should be taken into consideration. Another option would be for the WLAN/Access Router to support MLD proxy routing.

The Edge Router has to be enabled to support MLDv1 and MLDv2 in order to process the requests coming from the IPv6 WLAN Host or WLAN/Access Router (if present). The Edge Router has also needs to be enabled for PIM-SSM in order to receive joins from IPv6 WLAN Hosts or WLAN/Access Router (if present), and send joins towards the SP core.

MLD authentication, authorization, and accounting are usually configured on the Edge Router in order to enable the SP to do billing for the content services provided. Further investigation should be made in finding alternative mechanisms that would support these functions.

Concerns have been raised in the past related to running IPv6 multicast over WLAN links. Potentially these are the same kind of issues when running any Layer 3 protocol over a WLAN link that has a high loss-to-signal ratio, where certain frames that are multicast based are dropped when settings are not adjusted properly. For instance, this behavior is similar to an IGMP host membership report, when done on a WLAN link with a high loss-to-signal ratio and high interference.

This problem is inherited by WLAN that can impact both IPv4 and IPv6 multicast packets; it is not specific to IPv6 multicast.

While deploying WLAN (IPv4 or IPv6), one should adjust their broadcast/multicast settings if they are in danger of dropping application dependent frames. These problems are usually caused when the AP is placed too far (not following the distance limitations), high interference, etc. These issues may impact a real multicast application such as streaming video or basic operation of IPv6 if the frames were dropped. Basic IPv6 communications uses functions such as Duplicate Address Detection (DAD), Router and Neighbor

Solicitations (RS, NS), Router and Neighbor Advertisement (RA, NA), etc., which could be impacted by the above mentioned issues as these frames are Layer 2 Ethernet multicast frames.

Please refer to Section 6.3 for more IPv6 multicast details.

8.3. IPv6 QoS

Today, QoS is done outside of the WiFi domain, but it is nevertheless important to the overall deployment.

The QoS configuration is particularly relevant on the Edge Router in order to manage resources shared amongst multiple subscribers possibly with various service level agreements (SLAs). However, the WLAN Host/Router and Access Router could also be configured for QoS. This includes support for appropriate classification criteria, which would need to be implemented for IPv6 unicast and multicast traffic.

On the Edge Router, the subscriber-facing interfaces have to be configured to police the inbound customer traffic and shape the traffic outbound to the customer, based on the SLA. Traffic classification and marking should also be done on the Edge Router in order to support the various types of customer traffic: data, voice, and video. The same IPv4 QoS concepts and methodologies should be applied for the IPv6 as well.

It is important to note that when traffic is encrypted end-to-end, the traversed network devices will not have access to many of the packet fields used for classification purposes. In these cases, routers will most likely place the packets in the default classes. The QoS design should take into consideration this scenario and try to use mainly IP header fields for classification purposes.

8.4. IPv6 Security Considerations

There are limited changes that have to be done for WLAN the Host/Router in order to enhance security. The privacy extensions [RFC3041] for auto-configuration should be used by the hosts with the same consideration for host traceability as described in Section 6.5. IPv6 firewall functions should be enabled on the WLAN Host/Router, if present.

The ISP provides security against attacks that come from its own subscribers, but it could also implement security services that protect its subscribers from attacks sourced from outside its network. Such services do not apply at the access level of the network discussed here.

If the host authentication at hotspots is done using a web-based authentication system, then the level of security would depend on the particular implementation. User credentials should never be sent as clear text via HTTP. Secure HTTP (HTTPS) should be used between the web browser and authentication server. The authentication server could use RADIUS and LDAP services at the back end.

Authentication is an important aspect of securing WLAN networks prior to implementing Layer 3 security policies. For example, this would help avoid threats to the ND or stateless auto-configuration processes. 802.1x [IEEE8021X] provides the means to secure the network access; however, the many types of EAP (PEAP, EAP-TLS, EAP-TTLS, EAP-FAST, and LEAP) and the capabilities of the hosts to support some of the features might make it difficult to implement a comprehensive and consistent policy.

The 802.11i [IEEE80211i] amendment has many components, the most obvious of which are the two new data-confidentiality protocols, Temporal Key Integrity Protocol (TKIP) and Counter-Mode/CBC-MAC Protocol (CCMP). 802.11i also uses 802.1x's key-distribution system to control access to the network. Because 802.11 handles unicast and broadcast traffic differently, each traffic type has different security concerns. With several data-confidentiality protocols and the key distribution, 802.11i includes a negotiation process for selecting the correct confidentiality protocol and key system for each traffic type. Other features introduced include key caching and pre-authentication.

The 802.11i amendment is a step forward in wireless security. The amendment adds stronger encryption, authentication, and key management strategies that could make wireless data and systems more secure.

If any Layer 2 filters for Ethertypes are in place, the NAP must permit the IPv6 Ethertype (0X86DD).

The device that is the Layer 3 next hop for the subscribers (Access or Edge Router) should protect the network and the other subscribers against attacks by one of the provider customers. For this reason uRPF and ACLs should be used on all interfaces facing subscribers. Filtering should be implemented with regard for the operational requirements of IPv6 [IPv6-Security].

The Access and the Edge Router should protect their processing resources against floods of valid customer control traffic such as: RS, NS, and MLD Requests. Rate limiting should be implemented on all

subscriber-facing interfaces. The emphasis should be placed on multicast-type traffic, as it is most often used by the IPv6 control plane.

8.5. IPv6 Network Management

The necessary instrumentation (such as MIB modules, NetFlow Records, etc) should be available for IPv6.

Usually, NSPs manage the edge routers by SNMP. The SNMP transport can be done over IPv4 if all managed devices have connectivity over both IPv4 and IPv6. This would imply the smallest changes to the existing network management practices and processes. Transport over IPv6 could also be implemented and it might become necessary if IPv6 only islands are present in the network. The management applications may be running on hosts belonging to the NSP core network domain. Network Management Applications should handle IPv6 in a similar fashion to IPv4; however, they should also support features specific to IPv6 (such as neighbor monitoring).

In some cases, service providers manage equipment located on customers' LANs.

9. Broadband Power Line Communications (PLC)

This section describes the IPv6 deployment in Power Line Communications (PLC) Access Networks. There may be other choices, but it seems that this is the best model to follow. Lessons learnt from cable, Ethernet, and even WLAN access networks may be applicable also.

Power Line Communications are also often called Broadband Power Line (BPL) and sometimes even Power Line Telecommunications (PLT).

PLC/BPL can be used for providing, with today's technology, up to 200Mbps (total, upstream+downstream) by means of the power grid. The coverage is often the last half mile (typical distance from the medium-to-low voltage transformer to the customer premise meter) and, of course, as an in-home network (which is out of the scope of this document).

The bandwidth in a given PLC/BPL segment is shared among all the customers connected to that segment (often the customers connected to the same medium-to-low voltage transformer). The number of customers can vary depending on different factors, such as distances and even countries (from a few customers, just 5-6, up to 100-150).

PLC/BPL could also be used in the medium voltage network (often configured as Metropolitan Area Networks), but this is also out of the scope of this document, as it will be part of the core network, not the access one.

9.1. PLC/BPL Access Network Elements

This section describes the different elements commonly used in PLC/BPL access networks.

Head End (HE): Router that connects the PLC/BPL access network (the power grid), located at the medium-to-low voltage transformer, to the core network. The HE PLC/BPL interface appears to each customer as a single virtual interface, all of them sharing the same physical media.

Repeater (RPT): A device that may be required in some circumstances to improve the signal on the PLC/BPL. This may be the case if there are many customers in the same segment or building. It is often a bridge, but it could also be a router if, for example, there is a lot of peer-to-peer traffic in a building and due to the master-slave nature of the PLC/BPL technology, is required to improve the performance within that segment. For simplicity within this document, the RPT will always be considered a transparent Layer 2 bridge, so it may or may not be present (from the Layer 3 point of view).

Customer Premise Equipment (CPE): Modem (internal to the host), modem/bridge (BCPE), router (RCPE), or any combination among those (i.e., modem+bridge/router), located at the customer premise.

Edge Router (ER)

Figure 9.1 depicts all the network elements indicated above.

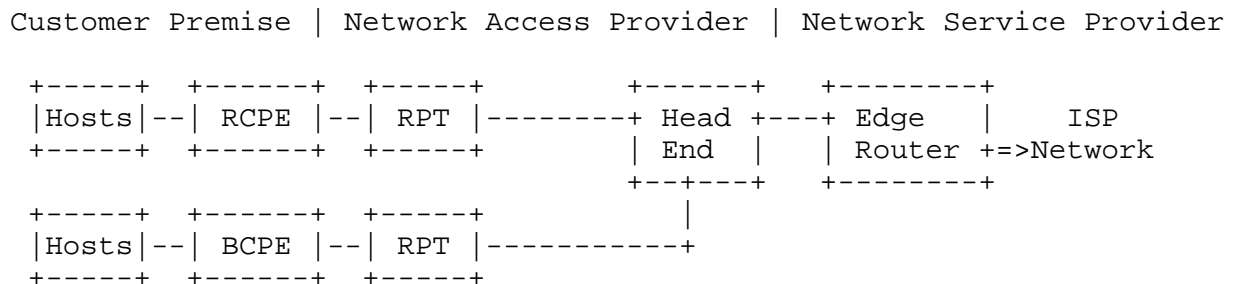


Figure 9.1

The logical topology and design of PLC/BPL is very similar to Ethernet Broadband Networks as discussed in Section 7. IP connectivity is typically provided in a Point-to-Point model, as described in Section 7.2.1

9.2. Deploying IPv6 in IPv4 PLC/BPL

The most simplistic and efficient model, considering the nature of the PLC/BPL networks, is to see the network as a point-to-point, one to each customer. Even if several customers share the same physical media, the traffic is not visible among them because each one uses different channels, which are, in addition, encrypted by means of 3DES.

In order to maintain the deployment concepts and business models proven and used with existing revenue-generating IPv4 services, the IPv6 deployment will match the IPv4 one. Under certain circumstances where new service types or service needs justify it, IPv4 and IPv6 network architectures could be different. Both approaches are very similar to those already described for the Ethernet case.

9.2.1. IPv6 Related Infrastructure Changes

In this scenario, only the RPT is Layer 3 unaware, but the other devices have to be upgraded to dual stack Hosts, RCPE, Head End, and Edge Router.

9.2.2. Addressing

The Hosts or the RCPEs have the HE as their Layer 3 next hop.

If there is no RCPE, but instead a BCPE, all the hosts on the subscriber site belong to the same /64 subnet that is statically configured on the HE. The hosts can use stateless auto-configuration or stateful DHCPv6-based configuration to acquire an address via the HE.

If an RCPE is present:

- A. It is statically configured with an address on the /64 subnet between itself and the HE, and with /64 prefixes on the interfaces connecting the hosts on the customer site. This is not a desired provisioning method, being expensive and difficult to manage.
- B. It can use its link-local address to communicate with the HE. It can also dynamically acquire through stateless auto-configuration the address for the link between itself and the HE. This step is

followed by a request via DHCP-PD for a prefix shorter than /64 (typically /48 [RFC3177]) that, in turn, is divided in /64s and assigned to its interfaces connecting the hosts on the customer site. This should be the preferred provisioning method, being cheaper and easier to manage.

The Edge Router needs to have a prefix, considering that each customer in general will receive a /48 prefix, and that each HE will accommodate customers. Consequently, each HE will require $n \times /48$ prefixes.

It could be possible to use a kind of Hierarchical Prefix Delegation to automatically provision the required prefixes and fully auto-configure the HEs, and consequently reduce the network setup, operation, and maintenance cost.

The prefixes used for subscriber links and the ones delegated via DHCP-PD should be planned in a manner that allows as much summarization as possible at the Edge Router.

Other information of interest to the host, such as DNS, is provided through stateful [RFC3315] and stateless [RFC3736] DHCPv6.

9.2.3. Routing

If no routers are used on the customer premise, the HE can simply be configured with a default route that points to the Edge Router. If a router is used on the customer premise (RCPE), then the HE could also run an IGP (such as OSPFv3, IS-IS or even RIPng) to the ER. The connected prefixes should be redistributed. If DHCP-PD is used, with every delegated prefix a static route is installed by the HE. For this reason, the static routes must also be redistributed. Prefix summarization should be done at the HE.

The RCPE requires only a default route pointing to the HE. No routing protocols are needed on these devices, which generally have limited resources.

The Edge Router runs the IPv6 IGP used in the NSP: OSPFv3 or IS-IS. The connected prefixes have to be redistributed, as well as any routing protocols (other than the ones used on the ER) that might be used between the HE and the ER.

9.3. IPv6 Multicast

The considerations regarding IPv6 Multicast for Ethernet are also applicable here, in general, assuming the nature of PLC/BPL is a shared media. If a lot of Multicast is expected, it may be worth considering using RPT which are Layer 3 aware. In that case, one extra layer of Hierarchical DHCP-PD could be considered, in order to facilitate the deployment, operation, and maintenance of the network.

9.4. IPv6 QoS

The considerations introduced for QoS in Ethernet are also applicable here. PLC/BPL networks support QoS, which basically is the same whether the transport is IPv4 or IPv6. It is necessary to understand that there are specific network characteristics, such as the variability that may be introduced by electrical noise, towards which the PLC/BPL network will automatically self-adapt.

9.5. IPv6 Security Considerations

There are no differences in terms of security considerations if compared with the Ethernet case.

9.6. IPv6 Network Management

The issues related to IPv6 Network Management in PLC networks should be similar to those discussed for Broadband Ethernet Networks in Section 7.6. Note that there may be a need to define MIB modules for PLC networks and interfaces, but this is not necessarily related to IPv6 management.

10. Gap Analysis

Several aspects of deploying IPv6 over SP Broadband networks were highlighted in this document, aspects that require additional work in order to facilitate native deployments, as summarized below:

- A. As mentioned in section 5, changes will need to be made to the DOCSIS specification in order for SPs to deploy native IPv6 over cable networks. The CM and CMTS will both need to support IPv6 natively in order to forward IPv6 unicast and multicast traffic. This is required for IPv6 Neighbor Discovery to work over DOCSIS cable networks. Additional classifiers need to be added to the DOCSIS specification in order to classify IPv6 traffic at the CM and CMTS in order to provide QoS. These issues are addressed in a recent proposal made to Cable Labs for DOCSIS 3.0 [DOCSIS3.0-Reqs].

- B. Section 6 stated that current RBE-based IPv4 deployment might not be the best approach for IPv6, where the addressing space available gives the SP the opportunity to separate the users on different subnets. The differences between IPv4 RBE and IPv6 RBE were highlighted in Section 6. If, however, support and reason are found for a deployment similar to IPv4 RBE, then the environment becomes NBMA and the new feature should observe RFC2491 recommendations.
- C. Section 6 discussed the constraints imposed on an LAA-based IPv6 deployment by the fact that it is expected that the subscribers keep their assigned prefix, regardless of LNS. A deployment approach was proposed that would maintain the addressing schemes contiguous and offers prefix summarization opportunities. The topic could be further investigated for other solutions or improvements.
- D. Sections 6 and 7 pointed out the limitations (previously documented in [IPv6-Multicast]) in deploying inter-domain ASM; however, SSM-based services seem more likely at this time. For such SSM-based services of content delivery (video or audio), mechanisms are needed to facilitate the billing and management of listeners. The currently available feature of MLD AAA is suggested; however, other methods or mechanisms might be developed and proposed.
- E. In relation to Section 8, concerns have been raised related to running IPv6 multicast over WLAN links. Potentially, these are the same kind of issues when running any Layer 3 protocol over a WLAN link that has a high loss-to-signal ratio; certain frames that are multicast based are dropped when settings are not adjusted properly. For instance this behavior is similar to an IGMP host membership report, when done on a WLAN link with high loss-to-signal ratio and high interference. This problem is inherited by WLAN that can impact both IPv4 and IPv6 multicast packets; it is not specific to IPv6 multicast.
- F. The privacy extensions were mentioned as a popular means to provide some form of host security. ISPs can track relatively easily the prefixes assigned to subscribers. If, however, the ISPs are required by regulations to track their users at host address level, the privacy extensions [RFC3041] can be implemented only in parallel with network management tools that could provide traceability of the hosts. Mechanisms should be defined to implement this aspect of user management.

- G. Tunnels are an effective way to avoid deployment dependencies on the IPv6 support on platforms that are out of the SP control (GWRs or CPEs) or over technologies that did not standardize the IPv6 support yet (cable). They can be used in the following ways:
- i. Tunnels directly to the CPE or GWR with public or private IPv4 addresses.
 - ii. Tunnels directly to hosts with public or private IPv4 addresses. Recommendations on the exact tunneling mechanisms that can/should be used for last-mile access need to be investigated further and should be addressed by the IETF Softwire Working Group.
- H. Through its larger address space, IPv6 allows SPs to assign fixed, globally routable prefixes to the links connecting each subscriber.
- This approach changes the provisioning methodologies that were used for IPv4. Static configuration of the IPv6 addresses for all these links on the Edge Routers or Access Routers might not be a scalable option. New provisioning mechanisms or features might need to be developed in order to deal with this issue, such as automatic mapping of VLAN IDs/PVCs (or other customer-specific information) to IPv6 prefixes.
- I. New deployment models are emerging for the Layer 2 portion of the NAP where individual VLANs are not dedicated to each subscriber. This approach allows Layer 2 switches to aggregate more than 4096 users. MAC Forced Forwarding [RFC4562] is an example of such an implementation, where a broadcast domain is turned into an NBMA-like environment by forwarding the frames based on both Source and Destination MAC addresses. Since these models are being adopted by the field, the implications of deploying IPv6 in such environments need to be further investigated.
- J. The deployment of IPv6 in continuously evolving access service models raises some issues that may need further investigation. Examples of such topics are [AUTO-CONFIG]:
- i. Network Service Selection & Authentication (NSSA) mechanisms working in association with stateless auto-configuration. As an example, NSSA relevant information, such as ISP preference, passwords, or profile ID, can be sent by hosts with the RS [RFC4191].

- ii. Providing additional information in Router Advertisements to help access nodes with prefix selection in multi-ISP/multi-homed environments.

Solutions to some of these topics range from making a media access capable of supporting native IPv6 (cable) to improving operational aspects of native IPv6 deployments.

11. Security Considerations

Please refer to the individual "IPv6 Security Considerations" technology sections for details.

12. Acknowledgements

We would like to thank Brian Carpenter, Patrick Grossetete, Toerless Eckert, Madhu Sudan, Shannon McFarland, Benoit Lourdelet, and Fred Baker for their valuable comments. The authors would like to acknowledge the structure and information guidance provided by the work of Mickles, et al., on "Transition Scenarios for ISP Networks" [ISP-CASES].

13. References

13.1. Normative References

- [RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
- [RFC2080] Malkin, G. and R. Minnear, "RIPng for IPv6", RFC 2080, January 1997.
- [RFC2364] Gross, G., Kaycee, M., Lin, A., Malis, A., and J. Stephens, "PPP Over AAL5", RFC 2364, July 1998.
- [RFC2461] Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.
- [RFC2462] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.
- [RFC2473] Conta, A. and S. Deering, "Generic Packet Tunneling in IPv6 Specification", RFC 2473, December 1998.

- [RFC2516] Mamakos, L., Lidl, K., Evarts, J., Carrel, D., Simone, D., and R. Wheeler, "A Method for Transmitting PPP Over Ethernet (PPPoE)", RFC 2516, February 1999.
- [RFC2529] Carpenter, B. and C. Jung, "Transmission of IPv6 over IPv4 Domains without Explicit Tunnels", RFC 2529, March 1999.
- [RFC2661] Townsley, W., Valencia, A., Rubens, A., Pall, G., Zorn, G., and B. Palter, "Layer Two Tunneling Protocol "L2TP"", RFC 2661, August 1999.
- [RFC2740] Coltun, R., Ferguson, D., and J. Moy, "OSPF for IPv6", RFC 2740, December 1999.
- [RFC2784] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 2784, March 2000.
- [RFC3041] Narten, T. and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 3041, January 2001.
- [RFC3053] Durand, A., Fasano, P., Guardini, I., and D. Lento, "IPv6 Tunnel Broker", RFC 3053, January 2001.
- [RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, February 2001.
- [RFC3177] IAB and IESG, "IAB/IESG Recommendations on IPv6 Address Allocations to Sites", RFC 3177, September 2001.
- [RFC3180] Meyer, D. and P. Lothberg, "GLOP Addressing in 233/8", BCP 53, RFC 3180, September 2001.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [RFC3618] Fenner, B. and D. Meyer, "Multicast Source Discovery Protocol (MSDP)", RFC 3618, October 2003.
- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.

- [RFC3736] Droms, R., "Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6", RFC 3736, April 2004.
- [RFC3904] Huitema, C., Austein, R., Satapati, S., and R. van der Pol, "Evaluation of IPv6 Transition Mechanisms for Unmanaged Networks", RFC 3904, September 2004.
- [RFC3931] Lau, J., Townsley, M., and I. Goyret, "Layer Two Tunneling Protocol - Version 3 (L2TPv3)", RFC 3931, March 2005.
- [RFC4001] Daniele, M., Haberman, B., Routhier, S., and J. Schoenwaelder, "Textual Conventions for Internet Network Addresses", RFC 4001, February 2005.
- [RFC4029] Lind, M., Ksinant, V., Park, S., Baudot, A., and P. Savola, "Scenarios and Analysis for Introducing IPv6 into ISP Networks", RFC 4029, March 2005.
- [RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and More-Specific Routes", RFC 4191, November 2005.
- [RFC4213] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, October 2005.
- [RFC4214] Templin, F., Gleeson, T., Talwar, M., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", RFC 4214, October 2005.
- [RFC4380] Huitema, C., "Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)", RFC 4380, February 2006.

13.2. Informative References

- [6PE] De Clercq, J., Ooms, D., Prevost, S., and F. Le Faucheur, "Connecting IPv6 Islands across IPv4 Clouds with BGP", Work in Progress, December 2006.
- [AUTO-CONFIG] Wen, H., Zhu, X., Jiang, Y., and R. Yan, "The deployment of IPv6 stateless auto-configuration in access network", 8th International Conference on Telecommunications, ConTEL 2005, June 2005.

- [BSR] Bhaskar, N., Gall, A., Lingard, J., and S. Venaas, "Bootstrap Router (BSR) Mechanism for PIM", Work in Progress, June 2006.
- [DOCSIS3.0-OSSI] CableLabs, CL., "DOCSIS 3.0 OSSI Specification(CM-SP-OSSiv3.0-D02-060504)", May 2006.
- [DOCSIS3.0-Reqs] Droms, R., Durand, A., Kharbanda, D., and J-F. Mule, "DOCSIS 3.0 Requirements for IPv6 Support", Work in Progress, March 2006.
- [DynamicTunnel] Palet, J., Diaz, M., and P. Savola, "Analysis of IPv6 Tunnel End-point Discovery Mechanisms", Work in Progress, January 2005.
- [IEEE80211i] IEEE, "IEEE Standards for Information Technology: Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications, Amendment 6: Medium Access Control (MAC) Security Enhancements", July 2004.
- [IEEE8021X] IEEE, "IEEE Standards for Local and Metropolitan Area Networks: Port based Network Access Control, IEEE Std 802.1X-2001", June 2001.
- [IPv6-Multicast] Savola, P., "IPv6 Multicast Deployment Issues", Work in Progress, April 2004.
- [IPv6-Security] Convery, S. and D. Miller, "IPv6 and IPv4 Threat Comparison and Best-Practice Evaluation", March 2004.
- [ISISv6] Hopps, C., "Routing IPv6 with IS-IS", Work in Progress, October 2005.
- [ISP-CASES] Mickles, C., "Transition Scenarios for ISP Networks", Work in Progress, September 2002.
- [Protocol41] Palet, J., Olvera, C., and D. Fernandez, "Forwarding Protocol 41 in NAT Boxes", Work in Progress, October 2003.
- [RF-Interface] CableLabs, CL., "DOCSIS 2.0(CM-SP-RFiv2.0-I10-051209)", December 2005.
- [RFC4562] Melsen, T. and S. Blake, "MAC-Forced Forwarding: A Method for Subscriber Separation on an Ethernet Access Network", RFC 4562, June 2006.

- [Softwire] Dawkins, S., Ed., "Softwire Problem Statement", Work in Progress, May 2006.
- [v6tc] Palet, J., Nielsent, K., Parent, F., Durand, A., Suryanarayanan, R., and P. Savola, "Goals for Tunneling Configuration", Work in Progress, August 2005.

Authors' Addresses

Salman Asadullah
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134
USA

Phone: 408 526 8982
EMail: sasad@cisco.com

Adeel Ahmed
Cisco Systems
2200 East President George Bush Turnpike
Richardson, TX 75082
USA

Phone: 469 255 4122
EMail: adahmed@cisco.com

Ciprian Popoviciu
Cisco Systems
7025-6 Kit Creek Road
Research Triangle Park, NC 27709
USA

Phone: 919 392 3723
EMail: cpopovic@cisco.com

Pekka Savola
CSC - Scientific Computing Ltd.
Espoo
Finland

EMail: psavola@funet.fi

Jordi Palet Martinez
Consulintel
San Jose Artesano, 1
Alcobendas, Madrid E-28108
Spain

Phone: +34 91 151 81 99
EMail: jordi.palet@consulintel.es

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

