



WHITE PAPER

IPv6 ACCESS SERVICES

Last Updated: July 2005

Because of the exponential growth of Internet services and increasing number of end users, service providers are looking for new ways for their current network architecture to meet the needs of Internet-ready appliances, new applications, and services. IPv6 is designed to enable service providers to meet these challenges and provide new services to their customers. As the number of devices per broadband users exponentially increases worldwide, cable, DSL, Ethernet to the home, wireless, and other always-on access technologies can benefit from the huge address range of IPv6 (SP6).

In a convergent movement, the development of the IPv6 Internet has diversified the ways in which it can be accessed. By working from a campus or an enterprise network, IPv6 users initially stayed within the reach of hardwired IP networks. As IPv6 functionality and benefits have become generally accepted and understood, the demand for accessing IPv6 Internet, including various access technologies (XDSL, Ethernet, and IEEE 802.11), regardless of the use of Point-to-Point Protocol (PPP) on shared or dedicated media, from any location has grown.

At this point, IPv6-based services are seen as a differentiator that enables service providers to take advantage of the large IPv6 address space and allows them to better position themselves against the competition. The IPv6 deployments can be seen as an impetus to lower service support costs by eliminating Network Address Translation (NAT), with its negative consequences on applications and its complex behavior.

Cisco IOS® Software has been extended to meet this demand and thus scales for large IPv6 deployments. These extensions encompass IPv6 address assignment; authentication, authorization, and accounting (AAA); and RADIUS extensions for IPv6 and Dynamic Host Configuration Protocol (DHCP) Version 6. This document covers the deployment of those features and discusses IPv6 access operational design choices.

This document focuses on the Cisco IOS Software IPv6 Broadband Access feature set, which is available in Cisco IOS Software Release 12.3T and later. The same level of functionality will be provided in a future release of Cisco IOS Software Release 12.2S.

To determine the feature set supported in a given release or hardware, please refer to:

- Cisco® Feature Navigator: <http://www.cisco.com/go/fn>
- Cisco IOS Software documentation: "IPv6 Start Here" at http://www.cisco.com/en/US/products/sw/iosswrel/ps5187/products_configuration_guide_chapter09186a00801d65ed.html

ARCHITECTURE

Deployment of IPv6 highlighted operational requirements differs from deployment of IPv4. Because of this new protocol, the access architecture needs to be reexamined.

Address Allocation

The principal reason for creation of IPv6 was to provide more IP addresses to the Internet community. Because many more IPv6 addresses are available, IPv6 can offer long-lived address allocation. When using IPv6, there is no need to share addresses from a pool or to use NAT. Although with IPv4 end users could be allocated as little as a single address, the smallest address block allocation in IPv6 is for a single link (/64 prefix). The typical IPv6 allocation is a /48 prefix, which allows the end site to have 64,000 links in its network.

Static allocation is a requirement resulting from the foreseen proliferation of servers, typically home appliances and game stations. Those servers must be accessed using a naming service—for example, a Domain Name System (DNS) entry—which is practically preventing the use of temporary addresses. This option was also possible with IPv4 through the configuration of Dynamic DNS and NAT devices supporting port mapping, but it created operational problems.

Technically, IPv4 had no mechanism to automatically offer more than one address to a remote site. For deployments relying on PPP, IP Control Protocol (IPCP) suffered from the inability to deliver more than a single address to end customers. Because of a very conservative distribution policy for IPv4 addresses, DHCPv4 has never been required to offer more than one address.

With IPv6, a user receives an address prefix rather than a single address as with IPv4. Typically, the ISP will assign a /64 or a /48 address prefix. Note also that an IPv6 host can use multiple addresses within one address prefix (for example, temporary addresses) or multiple prefixes.

Stateless Address Autoconfiguration

In IPv4, PPP is the preeminent technique to connect single users or sites to the Internet. A number of functions that enable automatic address assignment occur at the PPP layer.

Similarly, large-scale IPv6 access deployments have to be able to automatically assign IPv6 address prefixes to end users, connecting using PPP. In IPv6, however, address assignment occurs exclusively at the network layer, using Internet Control Message Protocol (ICMP)-based mechanisms.

As IPCP configures IPv4 over PPP, IPv6CP also configures IPv6 over PPP, but it negotiates only a unique interface identifier, and all other functions are performed in Layer 3. With the PPP link established, the address prefix is advertised to the user in a router advertisement message. The address prefix advertised in the router advertisement can come from various sources (AAA server, manual configuration, or allocation from a prefix pool). The host receives a /64 address prefix and automatically configures its address based on the address prefix.

Numbering Sites

IPv6 allocates many more IP addresses to the end user than IPv4 does. This can lead to a proliferation of routers (customer premises equipment [CPE]) connected to always-on services (for example, DSL and cable).

In the IPv6 access model architecture, a primary requirement is to dynamically provide a CPE router with a variable-length address prefix for its parts to be assigned to several of its interfaces. This prevents mistakes in manual configuration and eases the propagation of addresses changes. The prefix must meet the following conditions:

- It must be shorter than 64 bits. It numbers several links.
- It must not be temporary. The customer is not willing to renumber frequently.
- It must be administratively assigned to a specific customer in the service provider database.

The CPE will typically receive a /48 address prefix and, out of it, will number its interfaces with /64 address prefixes. A standard-based solution is DHCPv6 (RFC 3315) prefix delegation: DHCPv6 provides variable-length prefixes to CPE routers. The advantage of DHCPv6 is that it can potentially be used for other services as well as DNS, domain names, and NTP (Network Time Protocol) server discovery. This technique is media independent, because it operates at Layer 3.

Delegating a prefix to an entire site is commonly a stateful operation, because the service provider routing scheme must always know the topological location of a site. A packet targeted to a site must be routed back to that site. This topological information is reflected in inserted routes in the routing table that make the sites reachable.

DHCPv6 PD (Prefix Delegation) fulfils the requirements listed above and offers a mechanism to delegate prefixes larger than /64 to remote sites. The Cisco IOS Software DHCPv6 function runs in routers. It is based on the DHCPv6 specification. Prefix delegation (RFC 3633) and DNS (RFC 3646) DHCPv6 options are supported and allow distribution of a prefix as well as a list of DNS servers and domain names.

IPv6 RADIUS

Current RADIUS attributes show IPv4 dependencies. They are encoded to support 32-bit addresses and run over an IPv4 transport. Although the IPv4 dialup model assigns only IPv4 host addresses, IPv6 RADIUS attributes carry IPv6 addresses and make room for concepts such as IPv6 address prefix assignments. To allow the RADIUS protocol to run over an IPv6 transport is also desirable.

The current Cisco IOS Software implementation of IPv6 on RADIUS supports both Cisco vendor-specific attributes (VSAs) and RFC 3162 attributes. RFC 3162 support on the ISP RADIUS server requires an upgrade of this software. Few widely deployed RADIUS solutions current supporting this functionality.

Storing Site Prefixes

In a typical IPv4 PPP deployment, stateful information related to users is stored in a RADIUS database. This is particularly convenient because RADIUS interacts well with PPP.

User authentication already occurs at the PPP level, and then addressing and routing information related to the user is logically stored alongside the authentication data.

This model can be replicated with IPv6: prefix information related to a site would be stored in the RADIUS server and provided to the provider edge router. Then the provider edge router acts accordingly to respond to DHCPv6 requests.

Alternatively, a DHCPv6 server can be used to store the prefixes assigned to each user. This can immediately pose an authentication problem. The user has been authenticated against the RADIUS server to get the connectivity, but authentication has to occur at the DHCPv6 level as well.

Routing

The IPv4 routing architecture relied heavily on pools of addresses either stored in routers or centralized in a RADIUS server. In both cases, address allocation was network driven and not user based. In the IPv6 case, where most if not all address allocations are static, a bit more care must be taken to avoid unmanageable routing table deaggregation within the ISP or telco networks.

In a wholesale architecture, the user is tunneled to its home gateway, so if there is a means to terminate a given user PPP session on the same L2TP network server (LNS) at all times, it does not pose specific problems. If, for redundancy purposes, the terminating LNS is not the same, that will lead to the insertion of a user route in the LNS. This does not scale well, because LNSs will have noncontiguous IPv6 prefixes for the users connecting to them. The same problem appears in a scenario where there is no L2TP and the user can terminate the PPP session on a different network access server (NAS) at a different time. To prevent this, the routing aggregation scheme must be carefully planned to aggregate users' prefixes as much as possible. If the ISP makes sure that all users terminate their PPP session on the same router at all times, the long-lived user IPv6 prefix can be based on the termination router. Each terminating router aggregates a user's prefixes in one larger routing announcement.

Access Architecture Models


Two architecture models that offer large scale Internet access are prevalent:

- A closed model, based on L2TP and adapted to wholesale-oriented service providers
- An open model without L2TP

They are both based on PPP/RADIUS and do not include a dedicated DHCPv6 server. If a DHCPv6 server is needed, its functions are fulfilled by the router terminating the PPP sessions.

Figure 1 depicts a wholesale architecture with L2TP access concentrator (LAC) and L2TP network server (LNS) elements. The ISP-customer link supports IPv6 in order to provide an IPv6 service to the end user. L2TP providing the tunneling mechanism between the LAC and the LNS is operated over IPv4. The RADIUS dialog between the LNS and the AAA server is done over an IPv4 transport as well.


Figure 1. Broadband Access Wholesale Architecture



In Figure 1, the RADIUS dialog between the NAS and the AAA server occurs over an IPv4 transport.

Figure 2 depicts ISP-operated broadband access architecture with a NAS element. The ISP-customer link supports IPv6 in the same manner as in Figure 1.


Figure 2. ISP-Operated Broadband Access Architecture



Both scenarios are PPP based. This is the most common way to offer an IPv4 Internet service, and many service providers initially chose it to offer IPv6 connectivity. Connecting users to the Internet with dedicated Ethernet (VLAN) or Wi-Fi links can lead to deployment models not based on PPP. In those scenarios, DHCPv6 (RFC 3315) could act at the central point of concentration of users' information. At this time, this solution is not as technically mature as the PPP-based scenario, but studies are under way to close this gap.

Figure 3 depicts possible ISP-operated broadband access architecture with a NAS element. There is no PPP on the CPE-NAS link. Each provider edge router acts as a DHCPv6 relay, and the DHCPv6 function is centralized on a server. The CPE prefix is acquired using DHCPv6.

Figure 3. Deployment with DHCPv6 Relay/Server Not Based on PPP



Customer Link Encapsulation

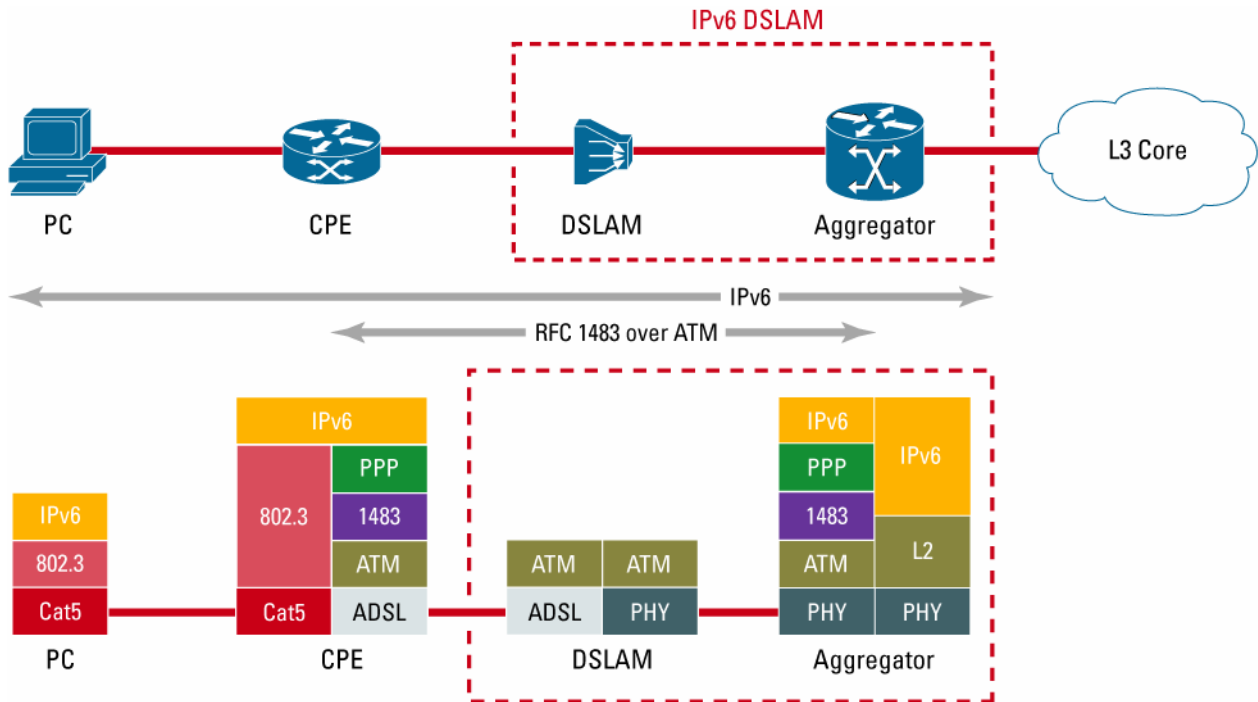
Point-to-Point Protocol over Ethernet (PPPoE), Point-to-Point Protocol over ATM (PPPoA), and routed bridge encapsulation (RBE) access methods are good encapsulation candidates to offer IPv6 connectivity in a variety of access service ISP designs. Those methods can leverage the IPv6 extensions to the AAA function. Offering IPv6 on top of Ethernet (or an Ethernet-like technology such as Wi-Fi) without PPP is certainly a model that will be explored by many ISPs in the near future.

This section focuses on DSL-based accesses: PPPoA, PPPoE, and RBE.

PPPoA Access

PPP over ATM adaptation Layer 5 (AAL5) (RFC 2364) is used between the CPE and the access concentrator. The user's PC IPv6 traffic, which flows over Ethernet to the CPE, is encapsulated over PPP to flow between the CPE and the access concentrator. Unlike the PPPoE approach, this one necessitates a Layer 3-aware (and thus IPv6-aware) CPE. After a PPP session is established, the CPE and the access concentrator must allocate the resources for a PPP virtual interface and configure IPv6 over it. Figure 4 depicts the global architecture for PPPoA.

Figure 4. Architecture for PPPoA Access

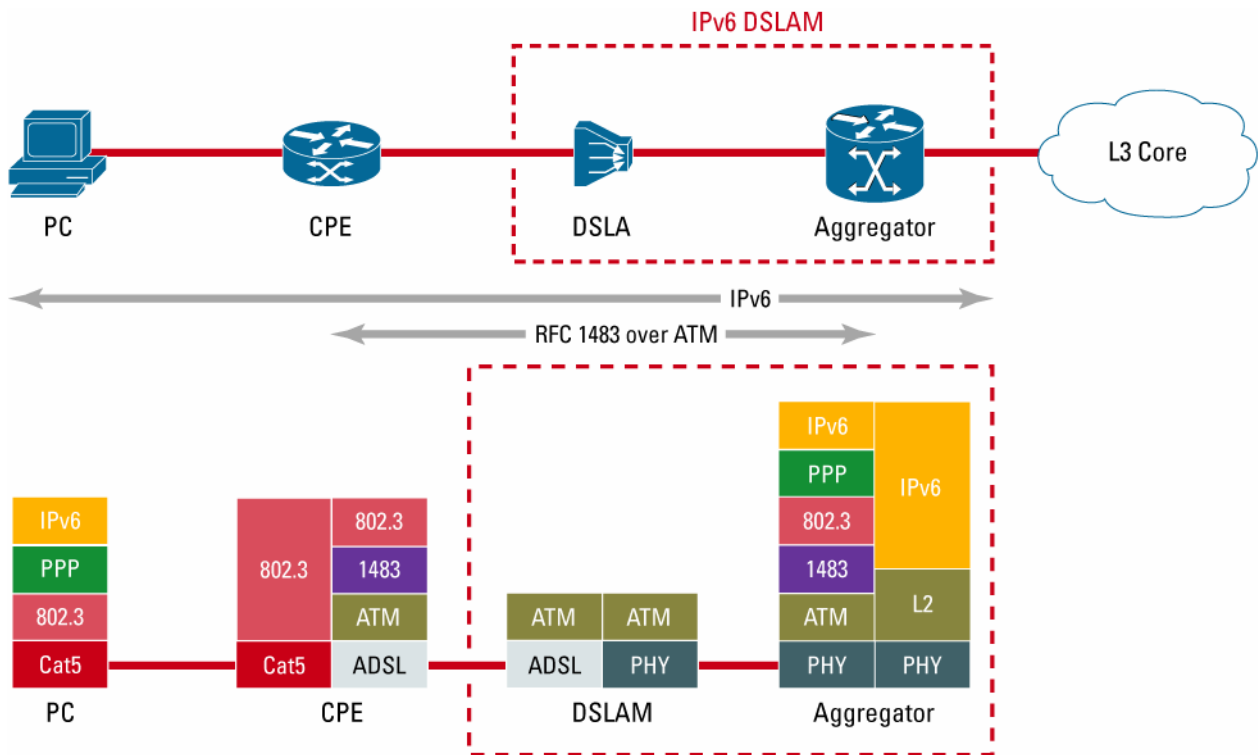


PPPoE Access

The PPPoE feature defined in RFC 2516 is used to encapsulate and transmit IPv6 traffic between the PC and the access concentrator through the CPE. The CPE is Layer 3 unaware and consequently IPv6 unaware. After a PPP session is established, both the host and the access concentrator must allocate resources for a PPP virtual interface and configure IPv6 over that interface. Figure 5 depicts this architecture.

Alternatively, the CPE can be the PPPoE session endpoint, which eliminates the need to install specific software on the customer's PC, but makes the CPE Layer 3 aware.

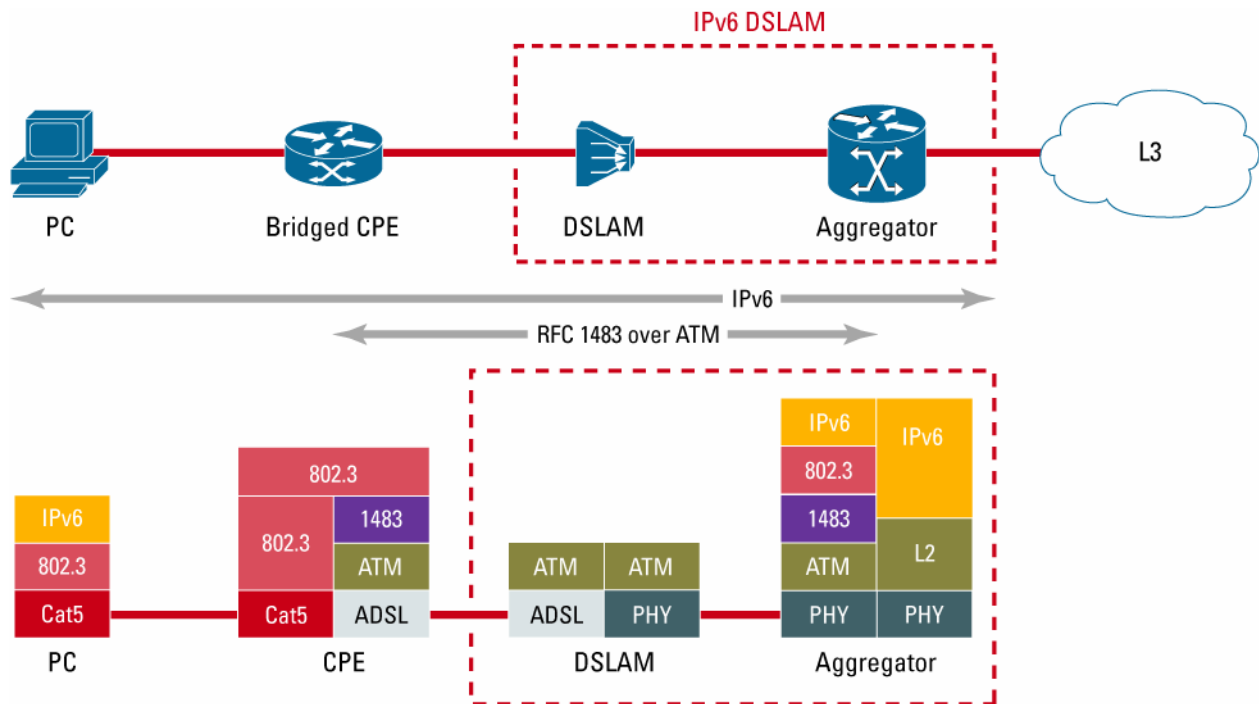
Figure 5. Architecture for PPPoE Access



RBE Access

The RBE feature routes IPv6 traffic received in ATM bridged (RFC 2684) encapsulation. An ATM interface that is configured in IPv6 route-bridged mode recognizes the IPv6 packets based on the type and routes them. Such interfaces use the characteristics of a stub LAN topology, which is commonly used for DSL access and offers increased performance and flexibility over integrated routing and bridging (IRB). Unlike in the IPv4 scheme, each user gets a different subnet. Figure 6 depicts the global architecture for RBE.

Figure 6. Architecture for RBE Access



DEPLOYMENT SCENARIOS

When creating an IPv6 access service, the ISP must make decisions in several areas, which sometimes depend on each other.

Most commonly, a /48 prefix will be delivered to every remote site with more than one subnet. A /64 prefix will be assigned to a customer with only one subnet or a host. As a last resort, a /128 prefix might be assigned to individual remote PCs.

The customer address allocation will be either static or dynamic:

- Static: when the customer network is always numbered with the same address prefix
- Dynamic: when the assigned address prefix changes with each connection

Connecting a Single Host


Permanent /64 Prefix

One addressing option is the assignment of a permanent /64 address prefix to a single PC: the provider edge router will send a router advertisement along the PE-CPE link. (See Figure 7.)

There are two options available for a single PC:

- Upon reception of the router advertisement, the PC completes the 64 least significant bits of the IPv6 address on its own.
- Before reception of the router advertisement, at the IPv6CP level, an interface identifier is given to the PC. The “Interface-Id” attribute in the user profile is used to provide a fixed interface identifier to the PC.

Figure 7. Permanent /64 Prefix Assignment



RADIUS Configuration

This is an example of the user configuration in the case where a /64 prefix is assigned permanently:

```
Auth-Type = Local, Password = "foo"  
User-Service-Type = Framed-User,  
Framed-Protocol = PPP,  
cisco-avpair = "ipv6:prefix#1=2001:db8:1:1::/64",
```

Interface identifier attributes can be specified:

```
Interface-Id = "0:0:0:1",
```

Short-Lived /64 Prefix

It is possible to assign a short-lived /64 prefix to a single PC. A different /64 prefix router advertisement is sent every time. This approach limits the capabilities of the user because it does not allow the PC to keep a long-lived address and provide content. If there is a requirement (for example, a tracking purpose) for a permanent interface ID, then the Interface-Id RADIUS attribute can be stored in the RADIUS profile.

RADIUS Configuration

The RADIUS profile references a shared pool on the NAS:

```
Auth-Type = Local, Password = "foo"  
User-Service-Type = Framed-User,  
Framed-Protocol = PPP,  
cisco-avpair = "addr-pool=foo-shared-64"
```

Provider Edge Configuration

The NAS is configured with a prefix pool, distributing /64 prefixes:

```
ipv6 prefix-pool foo-shared-64 2001:db8:e::/48 64
```

Short-Lived /128 Prefix

Assigning a /128 address prefix to a single PC is a possible addressing option. This solution can be deployed in environments where the connection is temporary; however, IPv6 does not bring anything on top of a classical IPv4 remote access scheme in this case. The same /64 prefix is advertised in router advertisements out of all interfaces. A route is installed only for the /128 prefix.

RADIUS Configuration

The user configuration in the case of a temporarily assigned /128 prefix is the following:

```
Auth-Type = Local, Password = "foo"  
User-Service-Type = Framed-User,  
Framed-Protocol = PPP,  
cisco-avpair = "addr-pool=foo-shared-128"
```

Provider Edge Configuration

The NAS is configured with a shared pool, distributing /128 prefixes:

```
ipv6 prefix-pool foo-shared-128 2001:db8:f::/64 128 shared
```

Connecting a Home or a Small Business


A service provider that offers connectivity to remote sites is likely to deploy a network in topology similar to the one depicted in Figure 8. The CPE is a Layer 3-capable device. In this case, where CPE is a Layer 2 device, the solution is similar to the assignment of a /64 prefix to a host (see the previous section). Figure 8 generically describes the access part of the service provider network and does not assume the use of a particular link-layer technology.

DHCPv6 Prefix Delegation

This scenario will need to deal with several deployment issues:

1. Numbering of the PE-CPE link with global addresses (link-local addresses can be used, but do not help network management)
2. Authentication of the connecting CPE
3. CPE prefixes database
4. Injection of the downstream link network and delegated customer prefixes in the ISP routing
5. Delegation of a shorter than /64 prefix to the CPE router
6. Autoconfiguration of hosts on links attached to the CPE router: IPv6 addresses, Internet parameters

Figure 8. Primary Deployment Issues and Their Location in a Service Provider Edge Network



DHCPv6 PD is the solution to aforementioned deployment issues 3, 4, and 5. DHCPv6 PD delegates prefixes from the provider edge router to the CPE and operates on the PE-CPE link.


Either the provider edge router or the RADIUS database can store the DHCPv6 PD delegated prefixes. If a local database is maintained on the provider edge router for all the connecting CPEs, the stored DHCP unique identifier (DUID) of the CPE indexes the binding between a given prefix and the corresponding CPE. In the other case, the /48 prefixes are stored as prefix attributes in the RADIUS server along with the other user's attributes.

The ISP-maintained RADIUS database performs CPE (user) authentication. This RADIUS database contains username/password pairs, corresponding /64 prefixes numbering the PE-CPE link, and delegated prefixes. This provides a solution to deployment issues 1, 2, and 3.

Standard autoconfiguration mechanisms as well as stateless DHCPv6 solve deployment issue 6.

In order to implement the proposed solution (Figure 9), the following must happen:

Figure 9. Proposed Solution for Prefix Delegation in a Service Provider Edge Network



1. A PPP session must be established over a Layer 2 link between the CPE and the provider edge routers. The CPE router authenticates itself with username “user1” in the PPP authentication phase of the negotiation. ¹
2. From the username contained in the PPP negotiation, a RADIUS request is sent to the service provider RADIUS server. If the username/password pair is validated, the result of this request returns a /64 prefix to the provider edge router. This prefix is then included in the router advertisement messages sent on the link connected to the CPE. The corresponding /64 prefix route is injected into the service provider routing system.
3. When the link between the CPE and the provider edge router comes up, the CPE issues a DHCPv6 SOLICIT message to discover DHCPv6 servers on the link.
4. The provider edge router, acting as a DHCPv6 server, sends a DHCPv6 ADVERTISEMENT message.
5. The CPE router uses this piece of information to issue a DHCPv6 REQUEST message to acquire a /48 prefix from the provider edge router. Note that a shorter sequence of message exchange is available in order to quicken the process.
6. The provider edge receives the DHCPv6 REQUEST message and issues a RADIUS request for the user (“user1-dhcpv6”).
7. The RADIUS server responds with the /48 prefix to assign it to the CPE.
8. The provider edge responds with a DHCPv6 REPLY message, including the /48 prefix assigned to this particular CPE. This response can include Internet configuration items (for example, DNS addresses, domain list). A /48 static route is automatically inserted in the provider edge routing table. An alternative is to store DHCPv6 bindings (between CPE identifiers and prefixes) in the provider edge router.
9. To assign /64 prefixes to connected interfaces, CPE derives (by configuration) them from the DHCPv6-assigned /48 prefix.
10. CPE interfaces configured with the /64 prefixes above start sending router advertisement messages on their links. Hosts on these links autoconfigure their respective IPv6 interface addresses based on the received router advertisements.

Note that this procedure is simplified if the PE-CPE link is not numbered with global addresses. In this case, the “user1-dhcpv6” profile does not exist, and the delegated prefix is stored in the “user1” profile.

¹ PPP is not mandatory (Ethernet could also be used), but it does offer client authentication.

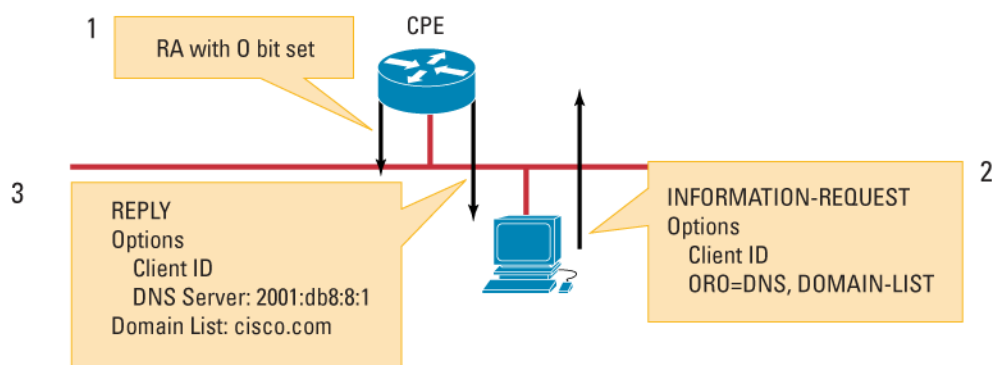
Stateless DHCPv6

From the address assignment perspective DHCPv6 PD is the primary element of an easy IPv6 deployment. However, in order to offer a fully functional IPv6 service, hosts behind the CPE need to be provided with DNS server addresses and possibly other parameters (for example, domain lists). This additional information should be provided automatically.

Following the Figure 9 scenario, after the CPE has received its /48 prefix and assigned /64 prefixes to its interfaces, it begins to send router advertisements with the corresponding prefixes. This way hosts on the link can autoconfigure their addresses. DHCPv6 can be used to provide other parameters in a stateless manner. This adds a few more steps to the process depicted in Figure 9:

11. The O bit can be set in the router advertisement messages sent by the CPE, so hosts on the link will know that other configuration parameters should be retrieved statelessly. They can then issue a DHCPv6 INFORMATION-REQUEST message to retrieve additional parameters (for example, DNS addresses, domain list).
12. The CPE or an external DHCPv6 server builds a DHCPv6 REPLY message with the responses to the parameters requested by the hosts.

Figure 10. Stateless DHCPv6 Message Exchange



PPPoA CPE Configuration

In the case of a PPPoA deployment, the CPE acts as a Layer 3 device. It receives a /48 prefix from the ISP and acquires Internet information (domain name, DNS addresses, and so on) from the DHCPv6 ISP server.

A PPP session is built over ATM. The DHCPv6 client is configured over the dialer interfaces. The CPE router is configured to send DHCPv6 SOLICIT messages on interface Dialer1. This interface is configured to accept router advertisements from the provider edge in order to configure its own IPv6 address.

Through the use of the “general-prefix” concept, prefix “PREFIX1” is used to number interfaces FastEthernet 0 and FastEthernet 1. The first 48 bits are coming from a prefix named “PREFIX1”; the last 16 bits are statically specified under the interface itself to build a complete /64 prefix. The prefix received using DHCPv6-PD is called “PREFIX1” by the CPE.

The FastEthernet interfaces are relaying the PC stateless DHCPv6 requests to a DHCPv6 server in the ISP network:

```
ipv6 unicast-routing
!
ip cef
ipv6 cef
!
interface Atm0
pvc <vpi/vci>
  encapsulation aal5mux ppp dialer
  dialer pool-member 1
!
interface Dialer1
  encapsulation ppp
  dialer pool 1
  ipv6 address autoconfig
  ipv6 dhcp client pd PREFIX1
  ppp authentication chap foo
  ppp chap hostname user@domain.net
  ppp chap password 7 1111111111
!
interface FastEthernet0
  ipv6 address PREFIX1 ::1:0:0:0:1/64
  ipv6 dhcp relay 2001:db8:4::1
!
interface FastEthernet1
  ipv6 address PREFIX1 ::2:0:0:0:1/64
  ipv6 dhcp relay 2001:db8:4::1
!
ipv6 route ::/0 Dialer1
!
```

RBE CPE Configuration

The following configuration (not IPv6 specific) applies to DSL access with RBE encapsulation. There is no PPP session to authenticate the user and to allow storing the /48 delegated prefix in the ISP RADIUS server. The CPE is a pure bridge. A router advertisement received on the ATM interface is bridged to the interface FastEthernet 0 to allow for PC autoconfiguration. Prefix delegation does not occur:

```
no ip routing
!
interface ATM0
  bridge-group 1
  pvc <vpi/vci>
  encapsulation aal5snap
!
interface FastEthernet0
  bridge-group 1
!
bridge 1 protocol ieee
```

RADIUS Configuration with DHCPv6 PD

When the PE-CPE link is numbered with global addresses, two user profiles are stored in the RADIUS server: "user1" and "user1-dhcpv6pd." The "user1" profile stores the PE-CPE link prefix, and "user1-dhcpv6pd" stores the /48 delegated prefix:

```
"user1"
Auth-Type = Local, Password = "foo"
User-Service-Type = Framed-User,
Framed-Protocol = PPP,
cisco-avpair = "ipv6:prefix#1=2001:db8:4567:1234::/64"
```

```
"user1-dhcpv6pd"
Auth-Type = Local, Password = "foo"
User-Service-Type = Framed-User,
Framed-Protocol = PPP,
cisco-avpair = "ipv6:prefix#1=2001:db8:aaaa::/48"
```

When the PE-CPE link is not numbered with global addresses, only one user profile is stored in the RADIUS server: the "user1" profile stores the /48 delegated prefix:

```
"user1"
Auth-Type = Local, Password = "foo"
User-Service-Type = Framed-User,
Framed-Protocol = PPP,
cisco-avpair = "ipv6:prefix#1=2001:db8:aaaa::/48"
```

PPPoA Infrastructure Configuration

As mentioned earlier, PPPoA can be deployed with or without L2TP. This section gives these two basic configurations.

NAS Configuration (Without L2TP)

This NAS numbers the link with the CPE. User authentication is provided by a RADIUS server. PE-CPE and delegated prefixes are stored in the RADIUS server.

When the PE-CPE link is not numbered with global addresses, "no ipv6 nd prefix framed-ipv6-prefix" is added under "interface Virtual-Templatel":

```
ipv6 unicast-routing
!
ip cef
ipv6 cef
!
aaa new-model
aaa authentication ppp default group radius
aaa authorization network default group radius
aaa authorization configuration PRELIST group radius
aaa accounting network default start-stop group radius
!
interface ATM0/0/0.1 point-to-point
 no ip directed-broadcast
 pvc <vpi/vci>
 encapsulation aal5mux ppp Virtual-Templatel
!
interface Virtual-Templatel
 ipv6 enable
 no ipv6 nd suppress-ra
 ipv6 dhcp server SRV-P1
 ppp authentication chap
!
ipv6 dhcp pool SRV-P1
 prefix-delegation aaa method-list PRELIST
!
radius-server host 192.168.2.20
radius-server key radius-password
```

LAC Configuration (with L2TP)

This LAC configuration is not IPv6 specific. It is a classical LAC configuration:

```
aaa new-model
aaa authentication ppp default group radius
aaa authorization network default group radius
aaa accounting network default start-start group radius
!
vpdn enable
!
vpdn-group 1
  request-dialin
  protocol l2tp
  domain domain.net
  initiate-to ip 192.168.1.27
  local name lac
!
interface ATM0/0/0.1 point-to-point
  pvc <vpi/vci>
  encapsulation aal5mux ppp Virtual-Templatel
!
interface Virtual-Templatel
  ppp authentication chap
!
radius-server host 192.168.2.20
radius-server key radius-password
```

LNS Configuration (with L2TP)

This LNS delegates /48 prefixes, numbers the link with the CPE, terminates the L2TP tunnels, and offers IPv6 connectivity. The IPv6 prefixes are stored in a RADIUS server.

When the PE-CPE link is not numbered with global addresses, "no ipv6 nd prefix framed-ipv6-prefix" is added under "interface Virtual-Templatel":

```
ipv6 unicast-routing
!
ip cef
ipv6 cef
!
aaa new-model
aaa authentication ppp default group radius
aaa authorization network default group radius
aaa authorization configuration PRELIST group radius
aaa accounting network default start-stop group radius
```

```

!
vpdn enable
!
vpdn-group 1
  accept-dialin
  protocol l2tp
  virtual-template 1
  terminate-from hostname lac
  local name lns
!
interface Virtual-Template1
  ipv6 enable
  no ipv6 nd suppress-ra
  ipv6 dhcp server SRV-P1
  ppp authentication chap callin
  !
  ipv6 dhcp pool SRV-P1
  prefix-delegation aaa method-list PRELIST
!
radius-server host 172.22.66.16
radius-server key radius-password

```

RBE Infrastructure Configuration

The following configuration applies to DSL access with RBE encapsulation. Unlike IPv4 RBE deployment, the ATM interface is numbered.

NAS Configuration

```

ipv6 unicast-routing
!
ip cef
ipv6 cef
!
interface ATM1/0.57 point-to-point
  atm route-bridged ipv6
  ipv6 address 2001:db8:6:6::1/64
  pvc <vpi/vci>
  encapsulation aal5snap
!

```

CONCLUSION

Cisco IOS Software currently enables the deployment of large-scale IPv6 access solutions. The IPv4 deployment models are able to support the addition of IPv6 in a dual-stack fashion. This is primarily achieved by DHCPv6, IPv6 RADIUS, and AAA extensions.

The introduction of these features is included in Cisco IOS Software releases and is a primary milestone in the broad adoption of IPv6. More tools are provided to service providers to simplify the deployment of value-added IPv6 features and to fulfill the basic needs of all Internet deployments.

Furthermore, Cisco Systems® is committed to broadening the support for Internet access methods that benefits of IPv6 support.

REFERENCES

- SP6: Salman Asadullah, Adeel Ahmed, Ciprian Popoviciu and Jordi Palet Martinez, “ISP IPv6 Deployment Scenarios in Broadband Access Networks,” draft-ietf-v6ops-bb-deployment-scenarios-02, work in progress, May 2005.
- [RFC 3315](#): R. Droms and others. “Dynamic Host Configuration Protocol for IPv6 (DHCPv6),” July 2003.
- [RFC 3633](#): O. Troan and R.Droms, “IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6,” December 2003.
- [RFC 3646](#): R.Droms, “DNS Configuration Options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6),” December 2003.
- [RFC 3162](#): B. Aboba, G. Zorn, and D. Mitton, “RADIUS and IPv6,” August 2001.
- [RFC 2364](#): G. Gross, M. Kaycee, A. Li, A. Malis, and J. Stephens, “PPP over AAL5,” July 1998.
- [RFC 2516](#): L. Mamakos, K. Lidl, J. Everts, D. Carrel, D. Simone, and R. Wheeler, “A Method for Transmitting PPP over Ethernet (PPPoE),” February 1999.
- [RFC 2684](#): D. Grossman and J. Heinanen, “Multiprotocol Encapsulation over ATM Adaptation Layer 5,” September 1999.

**Corporate Headquarters**

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters

Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters

Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on **the Cisco Website at www.cisco.com/go/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel
Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal
Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan
Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2005 Cisco Systems, Inc. All rights reserved. CCSP, CCVP, the Cisco Square Bridge logo, Follow Me Browsing, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Access Registrar, Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, Packet, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StrataView Plus, TeleRouter, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0502R) 205283.H_ETMG_SH_7.05

