IPv6 Access Services

1 Introduction

The development of the IPv6 Internet has diversified the ways in which it is accessed. IPv6 users initially stayed within reach of hardwired IP networks by working from a Campus or an Enterprise network.

As IPv6 functionality and benefits have become generally accepted understood, the demand to access IPv6 Internet from any location has grown. This includes various access technologies with PPP and DSL. Cisco IOS® Software has been extended to meet this demand, and thus scale large IPv6 deployment. Typically that means automatic and per user IPv6 address assignment, which covers the Authentication, Authorization, and Accounting (AAA) extensions for IPv6. This includes address pool management and RADIUS extensions, as well as deployment of those features. IPv6 access operational design choices are addressed as well.

The content of this paper focus on the Cisco IOS IPv6 Broadband Access feature set, as introduced Cisco IOS Software Release 12.2(13)T and Release 12.2(13)B or above.

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

- Cisco Feature Navigator: http://www.cisco.com/go/fn/
- Cisco IOS documentation “IPv6 Start Here:” http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122t/122t2/ipv6/ftipv6s.htm

2 Architecture

The deployment of IPv6, and its specific nature have highlighted the operational requirements that differ between IPv6 and IPv4 access technologies operations. For example, IPv6 enables long-lived address allocation, as there are much more IPv6 addresses; this is the solution in most cases.

Static allocation is also a requirement resulting from the proliferation of servers, typically home appliances. This option was not possible with IPv4, because of the IP address shortage and the ability of IPCP to deliver a non-permanent, single address to end-customers.

An IPv6 user receives an address prefix, rather than a single address, in contrast to an IPv4 user. 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 (e.g. temporary addresses). It generally leverages IPv6 to assign long-lived address prefixes to end-users. In a typical architecture, the user is tunnelled to their home gateway, so this does not pose specific problems. However this will have impact the routing system during a scenario in which each Network Access Server (NAS) has a pool of addresses, and users receive address from this pool when they dial. If users dial another Point Of Presence (POP), or are connected to another NAS at the same POP, they will receive a different IPv6 address prefix. If there is a permanent IPv6 address prefix, a route for each user’s address prefix

must be advertised into the routing system. This is a problem, because it introduces potential routing table de-aggregation in the cases of broadband access and roaming customer.

While it is possible to deploy IPv6 on all access media, including low speed media, some access methods are clearly better positioned to leverage IPv6 specificity. DSL and other always-on access technologies will benefit from long-lived address assignments to host servers, which enable many inbound connections. PPPoE, PPPoA, and RBE[1] access methods are good candidate encapsulations to offer IPv6 connectivity in a variety of access service ISP designs.

Current RADIUS attributes show IPv4 dependencies, which are encoded to support 32 bit addresses and run over an IPv4 transport. IPv6 RADIUS attributes carry IPv6 addresses and make room for concepts such as IPv6 address prefix assignments, as the IPv4 dialup model assign only IPv4 host addresses. It will soon also be desirable to allow the RADIUS protocol to run over an IPv6 transport.

It is crucial that large-scale IPv6 access deployments be able to automatically assign IPv6 address prefixes to end users, connecting via PPP. In IPv6, address assignment occurs at the network layer; conversely, in IPv4, a number of functions occur in the PPP layer. As IPCP configures IPv4 over PPP, IPV6CP configures IPv6 over PPP. IPv6CP negotiates a unique interface identifier, while all of functions are performed in Layer 3. Host address auto-configuration is performed at reception of Router Advertisement messages (RA).

IPv6 allocates many more IP addresses to the end user than does IPv4. This will clearly lead to a proliferation of routers (Customer Premise Equipment (CPE)) connected to always-on services (i.e.: Digital Subscriber Line (DSL), cable).

No IPv6 mechanism offers native auto-configuration for the configuration of router interfaces, but this does exist for host address configuration. However, in the IPv6 access model architecture, it is key to dynamically provide a router with a variable length address prefix for the assignment of interfaces addresses. This prevents mistakes in manual configuration, eases the propagation of addresses changes. The CPE will typically receive a /48 address prefix and, out of it, will number its interfaces with /64 address prefixes.

The IETF is currently considering several propositions to fulfill this requirement:

- *DHCPv6 Prefix Delegation:* DHCPv6 is put to contribution (Request and Reply messages) to provide variable length prefixes to CPE routers. The advantage of DHCPv6 is that it can potentially be used for other services, as well as DNS, NIS, and NTP server discovery.
- *ICMP Prefix Delegation:* New ICMP messages (Prefix Request and Prefix Delegation) are created to query and send prefixes of any length.

Both of the aforementioned techniques are media-independent, as they operate at Layer 3.

2.1 Access global architecture

Follow are two possible access architectures.

Figure 1 depicts a wholesale architecture with L2TP Access Concentrator (LAC) and L2TP Network Server (LNS) functions. The ISP-customer link supports IPv6 in order to provide an IPv6 service to the end-user.[2]

1. IPv6 over RBE encapsulation will be supported in later version of Cisco IOS.
2. L2TP providing the tunnelling 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.

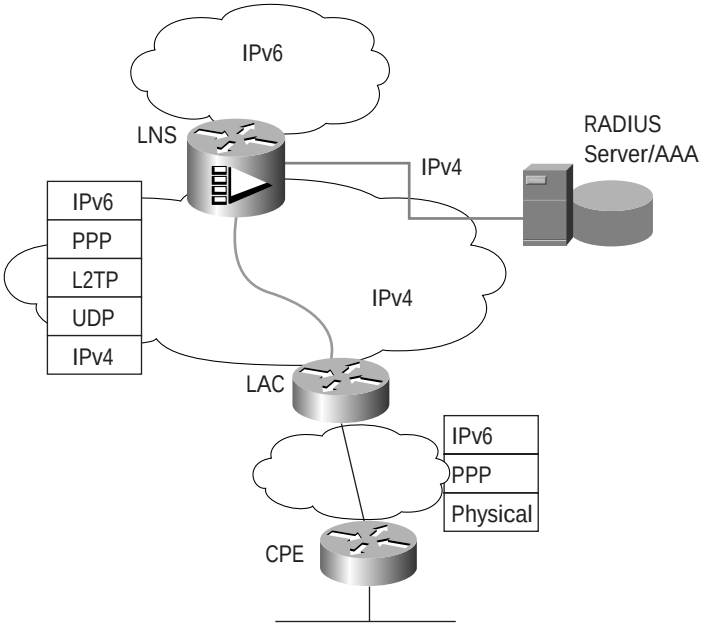

Figure 2 depicts ISP operated broadband access architecture with the Network Access Server (NAS) function. The ISP-customer link supports IPv6 in the same manner as in Figure 1. However the RADIUS dialog between the NAS and the AAA server occurs over an IPv4 transport.

Figure 2
ISP operated broadband access architecture.

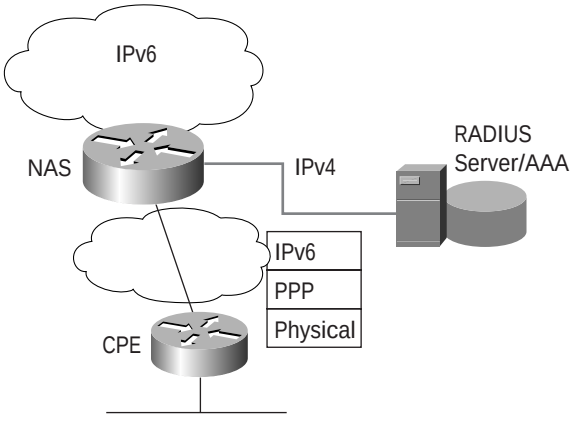

2.2 Customer link encapsulation

Encapsulations such as PPPoA, PPPoE, RBE, ISDN and Asynchronous access can leverage the IPv6 extensions to the AAA function. It is clear that always-on technologies, including DSL, will be predominant in the coming years, to the detriment of classical IPv4 access technologies (i.e.: ISDN or Asynchronous access). This section thus focuses on models working with DSL: PPPoA, PPPoE and RBE.

2.2.1 PPPoA access

PPP over ATM adaptation Layer 5 (AAL5) [RFC 2364] is created between the CPE and Access Concentrator. The IPv6 traffic originated by the user's PC flows over Ethernet to the CPE is encapsulated over PPP to flow between the CPE and the Access Concentrator. Unlike the PPPoE approach, this necessitates a Layer 3 aware (and thus consequently 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 3 depicts the global architecture for PPPoA.

Figure 3
Architecture for PPPoA access.

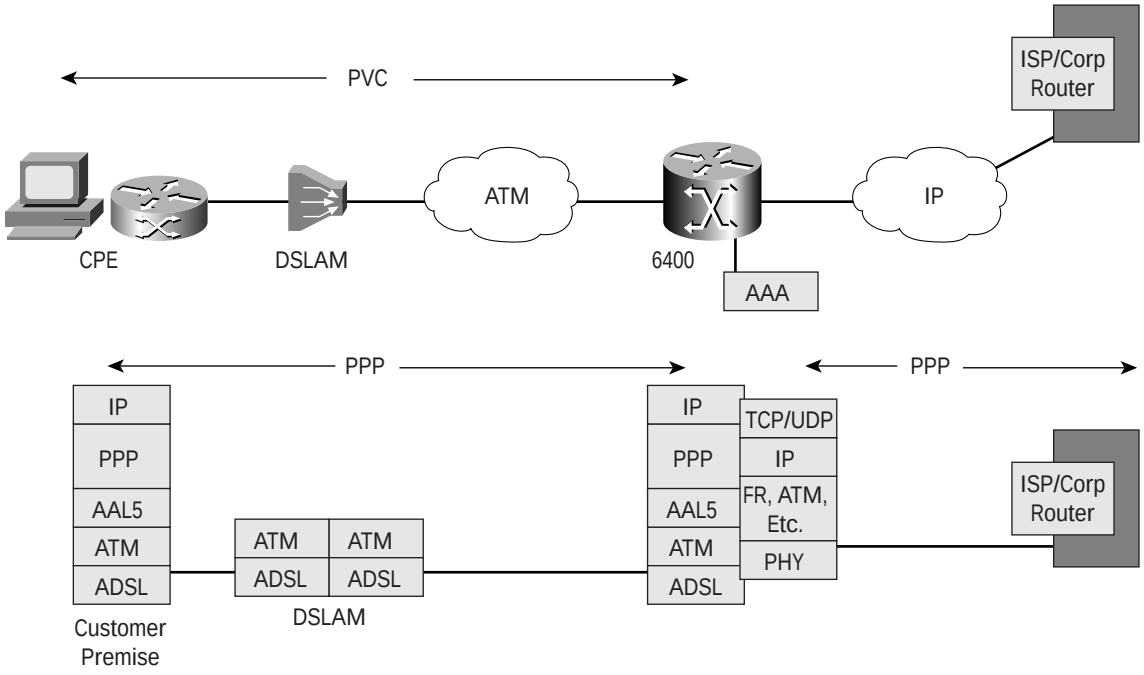

2.2.2 PPPoE access

The PPPoE feature defined in RFC2516 is put to contribution to transmit IPv6 traffic between the user's PC and the Access Concentrator through the CPE over PPP. The CPE is Layer 3 unaware, 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 4 depicts the global architecture for PPPoE. Alternatively, the CPE can be the PPPoE session end-point, which eliminates the need to install specific software on the customer's PC.

Figure 4
Architecture for PPPoE access.

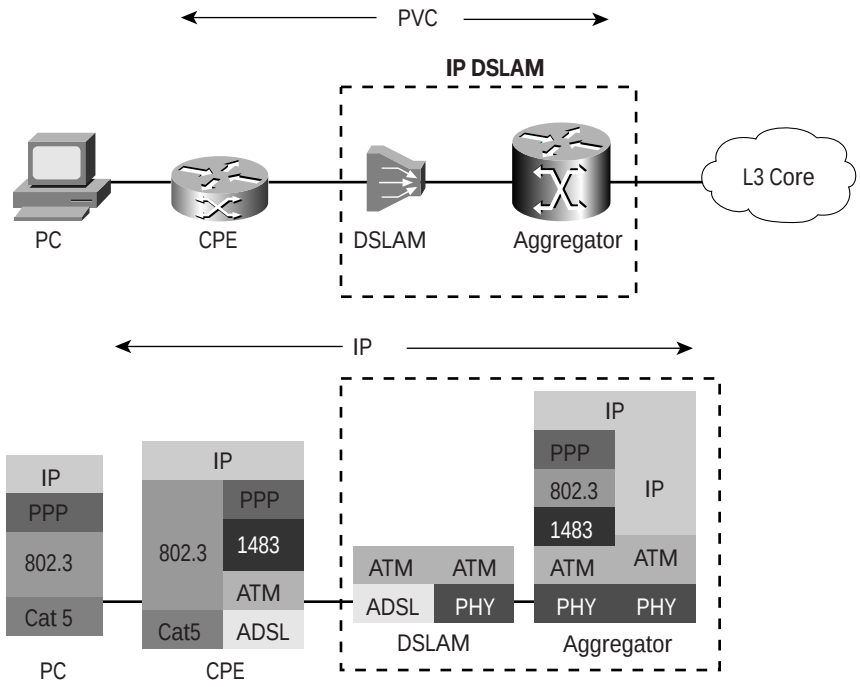

2.2.3 RBE access

The ATM Routed Bridge Encapsulation1 (RBE) feature routes IPv6 over bridged [RFC 2684] Ethernet traffic from a stub-bridged LAN. An ATM interface that is configured in route-bridged mode receives bridged IPv6 packets, which are routed via the IPv6 header. Such interfaces leverage 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). Figure 5 depicts the global architecture for RBE.

Figure 5
Architecture for RBE access.

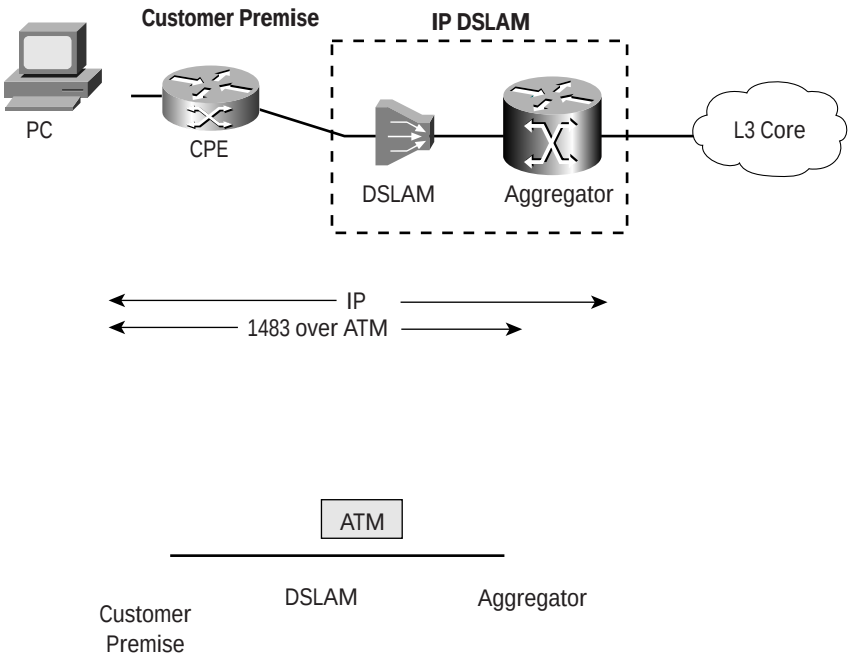

2.3 Stateless address auto-configuration

Stateless address auto-configuration works when the CPE is a single PC or a router limited to only one subnet.

When a single host connects via a PPP link, the address prefix is advertised to the user in a Router Advertisement (RA) message. The address prefix advertised in the RA can come from an AAA server (the prefix attribute), manual configuration, or allocation from a prefix pool. The host gets a /64 address prefix and automatically configures its address based on the address prefix. The Interface-ID portion of the address can come from Framed-Interface-Id AAA attribute, a derivation of the MAC address, or arbitrary assignment.

When the CPE is a router, auto-configuration of hosts attached to a segment of this router is achieved in bridging the RA messages from the PPP link to the local Ethernet segment.

3 AAA user attributes

The current Cisco IOS Software implementation of IPv6 on RADIUS uses Cisco Vendor Specific Attributes (VSA). The corresponding standard RADIUS attributes defined in [RFC 3162] exist. [RFC 3162] support on the ISP RADIUS server requires an upgrade of this software, and no widely deployed RADIUS solution currently supports this functionality. However, Cisco recognizes the importance of standards and will deliver an RFC3162 solution in future releases of Cisco IOS Software.

Table 1 depicts the correspondence, when it exists, between RFC3162 and Cisco VSA. As the current Cisco IOS Software implementation does not support RADIUS over an IPv6 Transport, some translation did not make sense. However, “Framed Interface ID” is already supported.

Table 1 RFC 3162 vs. Cisco VSA correspondence

RFC 3162	Cisco VSA
NAS-IPv6-Address	
Framed-Interface-Id	
Framed-IPv6-Prefix	IPv6 prefix
Login-IPv6-Host	
Framed-IPv6-Route	IPv6 route
Framed-IPv6-Pool	IPv6 pool
	IPv6 ACL {in, out}

Apart from the new “IPv6 prefix” attribute, the new Cisco VSAs are IPv4 attribute extended to support the IPv6 protocol. In addition to the aforementioned translated Cisco VSA, “IPv6 ACL {in, out}” is a Cisco VSA attribute with no correspondence in RFC3162.

3.1.1 IPv6 route attribute

The “route#” attribute allows the operator to specify a per-user static route. The format is the same as a Cisco IOS Software “ipv6 route” command. The following example specifies a route to “3ffe:ffff:1::/48” and a route to 3ffe:ffff:2::/48”.

```
cisco-avpair = "ipv6:route#1=3ffe:ffff:1::/48",
cisco-avpair = "ipv6:route#2=3ffe:ffff:2::/48",
```

3.1.2 IPv6 ACL attributes

The “inacl#<n>” and “outacl#<n>” attributes allow the user to specify a complete IPv6 access-list. The unique name of the access-list is generated automatically. When the user logs out, the access-list is removed, and the previous access-list on the interface (if any) is re-applied.

```
cisco-avpair = "ipv6:inacl#1=permit 3ffe:ffff:1::/48",
cisco-avpair = "ipv6:outacl#1=deny fec0::/10",
```

The “inacl” and “outacl” attributes let the user refer by name to an existing access-list configured on the router.

```
cisco-avpair = "ipv6:inacl=filter-n1",
```

The corresponding Router configuration would be the following:

```
ipv6 access-list remove-n1
 permit ipv6 3ffe:ffff:1::/48 any
 deny ipv6 fec0::/10 any
```

3.1.3 IPv6 prefix attribute

The “prefix#” attribute allows the user to specify which address prefixes to advertise in Neighbor Discovery Router Advertisement messages. The format used is as specified for the “ipv6 nd prefix” command. A corresponding route (marked as a per-user static route) will be installed in the RIB for the address prefix.

The following RADIUS configuration associates the address prefix “3ffe:ffff:5::/64” to a user.

```
cisco-avpair = "ipv6:prefix#1=3ffe:ffff:5::/64",
```

3.1.4 IPv6 pool attribute

This attribute extends the existing IPv4 “addr-pool” RADIUS attribute to support the IPv6 protocol. It specifies the name of a local pool from which to get the address prefix, and is used with “service=ppp” and “protocol=ipv6”. The specified pool must be pre-configured on the NAS. The following command creates the “foo” IPv6 pool.

```
cisco-avpair = "ipv6:addr-pool=foo",
```

3.1.5 Prefix pools

IPv6 prefix pools and IPv4 address pools have similar functions. Contrary to IPv4, entire address prefixes are allocated, typically /64s, and not single addresses. A pool is locally configured and cannot be changed after it is configured, as a change in configuration would remove and recreate the pool. All address prefixes already allocated will also be freed.

3.1.5.1 Shared prefix pools

A /64 address prefix is shared between all users of the pool. All interfaces send RAs for the same /64 address prefix. The user gets a /128, based on the address prefix and specific Interface-Identifier attribute. A per-user static route is installed in the RIB for the /128. Each user can receive only one address from the pool. Shared pools are not a recommended way to assign IPv6 addresses to end-users, as it would preclude auto-configuration, and unnecessarily restricts each user to a single address. The following command creates address prefix “3ffe:ffff:7::/64” as a shared pool.

```
ipv6 prefix-pool foo 3ffe:ffff:7::/64 128 shared
```

3.1.5.2 Normal prefix pools

Separate /64 are assigned to each user. The address prefix is advertised in RA and a route is installed in the RIB. The following command creates address prefix “3ffe:ffff:8::/48” as a prefix pool for /64 allocation.

```
ipv6 prefix-pool foo 3ffe:ffff:8::/48 64
```

3.1.6 RADIUS configuration

The RADIUS configuration corresponding to this feature requires that “users” files, which describe users’ configuration, be edited. The following example creates an account for user “user@isp.net”. Outgoing packets with site local address are blocked. A fixed /64 address prefix is assigned to this user.

```
user@isp.net  Password = secret
        Service-Type = Framed,
        Framed-Protocol = PPP,
        Cisco-AVpair = "ipv6:prefix#1=3ffe:ffff:5::/64",
  Cisco-AVPair = "ipv6:outacl#1=deny fec0::/10 any"
```

3.2 Address allocation priority

When there are several sources for obtaining an address prefix, follow this sequence:

1. AAA provided address prefix or prefixes from the pool named by AAA
2. An address prefix from a local prefix pool
3. Manually configured through the “ipv6 nd prefix” command.

4 Cisco IOS Command Line Interface

4.1 Configuration commands

4.1.1 ipv6 local pool

This global command configures an IPv6 prefix pool. The address prefix is included in the advertised RA on the interface, and it is possible to specify the length of this prefix that will be assigned to the customer. However, it is strongly recommended that the assigned address prefix length be 64, which allows for stateless address auto-configuration to work. If the “shared” keyword is specified, each user of the pool shares a single /64.

The prefix pool cache tracks the most recently used entries. The cache maintains information about users that disconnect, so they are re-assigned the same address prefix upon re-connection. The cache sized is adjustable, because of the cache size parameter.

The following command creates a /56 pool “foo” and from which each user is assigned a /64 address prefix. The cache size is 1000 entries.

```
ipv6 local pool foo 3ffe:ffff:a::/56 64 cache-size 1000
```

4.1.2 peer default ipv6 pool

This interface command configures which pool the client should assigned from. In the following example, the “foo” pool is assigned to Virtual-template1.

```
interface Virtual-Template1
 ipv6 enable
 no ipv6 nd suppress-ra
 peer default ipv6 pool foo
 ppp authentication chap
```

4.1.3 atm route-bridge ipv6

This interface command enables ATM Routed Bridge Encapsulation feature to route IPv6 over bridged [RFC 2684] Ethernet traffic from a stub-bridged LAN.

```
atm route-bridge ipv6
```

4.2 Show commands

4.2.1 show ipv6 local pool [<pool> [cache]]

The following command displays IPv6 Prefix Pool information, which includes cache information.

```
Router#show ipv6 local pool foo
Prefix is 3FFE:FFFF:A::/56 assign /64 prefix
2 entries in use, 254 available, 0 rejected
0 entries cached, 1000 maximum

User           Prefix               Interface
joe            3FFE:FFFF:A::/64     Vi1
john           3FFE:FFFF:A:1::/64     Vi2
```

4.2.2 show ipv6 interface [<interface>] [brief] [prefix]

This command displays IPv6 interface related parameters, addresses, and address prefixes.

```
Router#show ipv6 int vi1
Virtual-Access1 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::205:5FFF:FEAF:2C08
  No global unicast address is configured
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FFAF:2C08
  MTU is 1480 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 30 milliseconds
  ND advertised reachable time is 30 milliseconds
  ND advertised retransmit interval is 0 milliseconds
  ND router advertisements are sent every 200 seconds
  ND router advertisements live for 1800 seconds
  Hosts use stateless autoconfig for addresses.

Router#show ipv6 int vi1 prefix
IPv6 Prefix Advertisements Virtual-Access1
Codes: A - Address, P - Prefix-Advertisement, O - Pool
       X - Proxy RA, U - Per-user prefix, D - Default
       N - Not advertised, C - Calendar

O   3FFE:FFFF:A::/64 [LA] valid lifetime 2592000 preferred lifetime 604800
```

4.3 Debug commands

4.3.1 debug ipv6 pool

This debug command enables IPv6 Prefix Pool debugging.

```
2w4d: IPv6 Pool: Deleting route/prefix 3FFE:FFFF:A::/64 to Virtual-Access1 for cisco
2w4d: IPv6 Pool: Returning cached entry 3FFE:FFFF:A::/64 for cisco on Virtual-Access1 to
foo
2w4d: IPv6 Pool: Installed route/prefix 3FFE:FFFF:A::/64 to Virtual-Access1 for cisco
```

5 Deployment scenarios

When creating an IPv6 access service, the ISP must make decisions in several areas, sometimes showing dependencies to each other's.

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

5.1.1 Permanent /64 prefix

One possible addressing option is the assignment of a permanent /64 address prefix to a single PC or a to a router with a single segment. However, this can limit a router's functionality, because it does not leave room for extension behind the router. For a PC or a router, the NAS will send an RA along the link.

There are two options available for a single PC:

- Upon reception of the RA, the PC completes its own the 64 least significant bits of the IPv6 address on its own
- Before receipt of the RA at IPv6CP level, an Interface Identifier is given to the PC. It is possible to add the "Interface-Id" attribute in the user profile to provide a fixed interface identifier to the end PC.

When the router has a single link, the router LAN may be configured based on this permanent address prefix, because the /64 address prefix is permanent. An RA with this permanent address prefix can be sent to the router, increasing the flexibility of this deployment.

Figure 6
Permanent /64 prefix assignment.

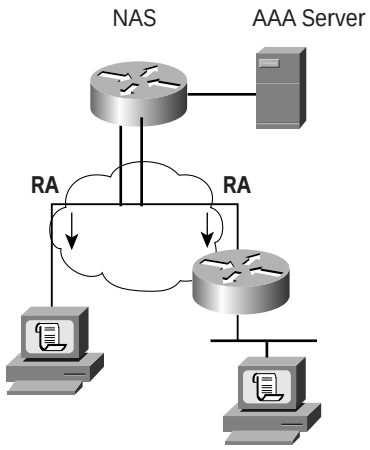

5.1.1.1 CPE config

```
ipv6 unicast-routing
!
interface Ethernet0
ipv6 address 3ffe:ffff:4567:1234::1/64
```

5.1.1.2 RADIUS config

This is the appropriate user configuration when there is a permanently assigned /64 address prefix:

```
Auth-Type = Local, Password = “foo”
User-Service-Type = Framed-User,
Framed-Protocol = PPP,
cisco-avpair = “ipv6:prefix=3ffe:ffff:4567:1234::/64

Interface Identifier attribute can be specified:

Interface-Id = “0:0:0:1”,
```

5.1.2 Permanent /48

The preferred addressing option is to assign a permanent /48 address prefix to a router. It is not limiting in the sense that it allows any network extension behind the router. However, as there is no current standard means of transmitting this address assignment to the remote router, manual address assignment occurs on the router. It would not make sense to assign temporary /48 address prefixes, which would generate several manual routers renumbering.

With the help of a future Prefix Delegation protocol, it will be possible to transmit the address prefix to the remote router for auto-configuration.

5.1.2.1 CPE config

The CPE interfaces will be configured manually from the assigned /48 address prefix.

5.1.2.2 RADIUS config

In the user profile a /48 route pointing to the customer link is added.

```
Auth-Type = Local, Password = “foo”
User-Service-Type = Framed-User,
Framed-Protocol = PPP,
cisco-avpair = “ipv6:route=3ffe:ffff:b::/48
```

5.1.3 Short lived and permanent /128

Assigning a /128 address prefix to a single PC is a possible addressing option. This solution may 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. If there is a requirement for a permanent address, then the Interface-Id RADIUS attribute may be used. An RA with a /128 address prefix is sent along the link with the address prefix.

Figure 7
/128 prefix assignment.

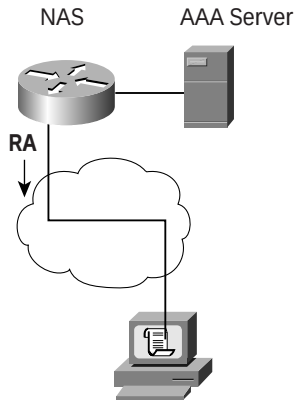

5.1.3.1 RADIUS config

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

```
Auth-Type = Local, Password = “foo”
User-Service-Type = Framed-User,
Framed-Protocol = PPP,
cisco-avpair = “addr-pool= 1”

cisco-avpair = "ipv6:pool#1=3ffe:ffff:c::/64",
```

5.1.4 Short lived /64

It is possible to assign a short live /64 address prefix to a single PC or a very simple network. This is limiting, as it forces renumbering at each connection attempt. A different /64 RA is sent every time if there is a single PC. Manual router renumbering is required for routers. Router configuration does not differ from permanent /64 case.

5.1.4.1 RADIUS config

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

5.1.4.2 NAS configuration

```
ipv6 prefix-pool foo-shared 3ffe:ffff:e::/48 64
```

6 Configuration Guide

6.1 PPPoA encapsulation

The following configuration applies to DSL access with PPPoA encapsulation as described in Figure 1, 2, and 3.

6.1.1 CPE configuration

```
interface FastEthernet0
 ipv6 address 3FFE:ffff:123:1999::1/64
!
interface atm0
 no ip address
 no ip directed-broadcast
 no ip mroute-cache
 pvc <vpi/vci>
  encapsulation aal5mux ppp dialer
  dialer pool-member 1
!
interface dialer1
 encapsulation ppp
 dialer pool 1
 dialer-group 1
 ipv6 address autoconfig
 ipv6 nd ra-interval 180
 ipv6 nd ra-lifetime 3600
 ppp authentication chap foo
 ppp chap hostname user@domain.net
 ppp chap password 7 1111111111
 ppp ipcp address accept
!
ipv6 route ::/0 Dialer1
```

6.1.2 NAS configuration (without L2TP)

```
aaa new-model
aaa authentication login default none
aaa authentication ppp default group radius
aaa authorization network default group radius
aaa accounting network default wait-start group radius
!
interface Loopback0
```

```
 ipv6 address 3FFE:ffff:123:1::1/64
!
interface ATM0/0/0.1 point-to-point
 no ip directed-broadcast
 pvc 1/32
  encapsulation aal5mux ppp Virtual-Template1
 !
interface Virtual-Template1
 ipv6 enable
 ipv6 mtu 1480
 ipv6 nd reachable-time 30
 no ipv6 nd suppress-ra
 ppp authentication pap
 peer default ipv6 pool foo
!
 ipv6 local pool foo 3ffe:ffff:a::/56 64 cache-size 1000
!
radius-server host 192.168.2.20 auth-port 1645 acct-port 1646
radius-server key radius-password
```

6.1.3 LAC configuration (with L2TP)

```
aaa new-model
aaa authentication login default none
aaa authentication ppp default  group radius
aaa authorization network default group radius
aaa accounting network default wait-start group radius
!
vpdn-group 1
 request-dialin
  protocol l2tp
  domain domain.net
 initiate-to ip 192.168.1.27
 local name sp_lac
!
interface Loopback0
 ip address 192.168.100.1 255.255.255.255

interface ATM0/0/0.1 point-to-point
 no ip directed-broadcast
 pvc 1/32
  encapsulation aal5mux ppp Virtual-Template1
 !
interface Virtual-Template1
 ip unnumbered FastEthernet0
 no ip directed-broadcast
 ppp authentication pap
!
!
radius-server host 192.168.2.20 auth-port 1645 acct-port 1646
radius-server key radius-password
```

6.1.4 LNS configuration (with L2TP)

```
ipv6 unicast-routing
!
aaa new-model
aaa authentication login default local
aaa authentication enable default local
aaa authentication ppp default group radius
aaa authorization network default group radius
aaa accounting network default start-stop group radius
!
vpdn enable
!
vpdn-group 1
accept dialin l2tp virtual-template 1 remote sp_lac
local name lns
!
interface Loopback0
ipv6  address 3ffe:ffff:bbbb::1/64
!
interface Virtual-Template1
 ipv6 enable
 ipv6 mtu 1480
 ipv6 nd reachable-time 30
 no ipv6 nd suppress-ra
 ppp authentication chap default
 !
radius-server host 172.22.66.16 auth-port 1645 acct-port 1646
radius-server key radius-password
```

6.2 PPPoE

The following configuration applies to DSL access with PPPoE encapsulation as described in Figure 1,2 and 4.

6.2.1 CPE configuration with PCs as PPPoE client

In this example the router connected to the DSL line does not originate the PPPoE sessions. PCs connected on the CPE LAN (FastEthernet0) originate the PPPoE sessions.

```
!
interface FastEthernet0
 no ip address
bridge-group1

!
interface ATM0
no ip address
no ip directed-broadcast
no ip mroute-cache
no atm ilmi-keepalive
pvc <vpi/vci>
encapsulation aal5snap
!
bundle-enable
bridge-group 1
```

```
hold-queue 224 in
!
bridge 1 protocol ieee
```

6.2.2 CPE with PPPoE client function configuration

In this example, the router connected to the DSL line originates the PPPoE sessions.

```
vpdn enable
!
vpdn-group pppoe
 request-dialin
  protocol pppoe
!
ipv6 unicast-routing
!
interface FastEthernet0
 ipv6 address 3FFE:ffff:123:1999::1/64
!
interface atm0
 no ip address
 bundle-enable
 dsl operating-mode auto
!
interface atm0.1 point-to-point
 pvc <vpi/vci>
 pppoe-client dial-pool-number 1

interface Dialer1
encapsulation ppp
 dialer pool 1
 ipv6 address autoconfig
 ipv6 nd ra-interval 180
 ipv6 nd ra-lifetime 3600
 ppp authentication chap callin
 ppp chap hostname john@domain.net
 ppp chap password 7 2104166219
 ppp ipcp address accept
!
ipv6 route ::/0 Dialer1
```

6.2.3 NAS configuration

```
aaa new-model
aaa authentication login default none
aaa authentication ppp default group radius
aaa authorization network default group radius
aaa accounting network default wait-start group radius
!
vpdn enable
!
vpdn-group pppoe
 accept-dialin
  protocol pppoe
  virtual-template 1
```

```
!
interface Loopback0
 ipv6 address 3FFE:ffff:123:1::1/64
!
interface ATM0/0/0
 no ip address
 no atm ilmi-keepalive
 hold-queue 500 in
!
interface ATM0/0/0.132 point-to-point
pvc <vpi/vci>
  encapsulation aal5snap
  protocol pppoe
 !
interface Virtual-Template1
 ipv6 enable
 ipv6 mtu 1480
 ipv6 nd reachable-time 30
 no ipv6 nd suppress-ra
 ppp authentication chap
!
radius-server host 192.168.2.20 auth-port 1645 acct-port 1646
radius-server key radius-password
```

6.2.4 LAC configuration with L2TP

See the PPPoE section above.

6.2.5 LNS configuration with L2TP

See the PPPoE section above.

6.3 RBE

The following configuration applies to DSL access with RBE1 encapsulation as described in Figure 1,2 and 5.

6.3.1 CPE configuration

```
no ip routing
!
interface FastEthernet0
 ipv6 address 3FFE:ffff:123:1999::1/64

interface ATM0
 mac-address  <FastEthernet0 MAC address>
 ipv6 address 3FFE:ffff:123:1999::1
 no atm ilmi-keepalive
 pvc <vpi/Vci>
  encapsulation aal5snap
 !
 bundle-enable
 bridge-group 1
!
!
bridge 1 protocol ieee
```

6.3.2 NAS configuration

```
interface Loopback0
 ipv6 address 3FFE:ffff:123:1::1/64
!
interface ATM0/0/0.132 point-to-point
 ipv6 unnumbered Loopback0
 atm route-bridged ipv6
 pvc <vpi/vci>
  encapsulation aal5snap
!
!
ip route 3FFE:ffff:123:1999::1 ATM0/0/0.132
```

6.3.3 LAC configuration with L2TP

See the PPPoE section above.

6.3.4 LNS configuration with L2TP

See the PPPoE section above.

7 Conclusion

Cisco IOS Software currently enables the deployment of large-scale IPv6 access solutions. This is primarily achieved by Cisco IPv6 RADIUS extensions and IPv6 AAA support; however, it will evolve towards standards as they are defined.

At this point of the ongoing standardization effort, IPv6 address prefix delegations mechanisms [MIAKAWA] are the most desired improvements. The IETF is currently considering several propositions for site address delegation (/48 address prefix), of which DHCP is the most promising solution [TROAN]. This was demonstrated at N+I Tokyo in July 2002.

8 References

[RFC 2364] G. Gross, M. Kaycee, A. Li, A. Malis, J. Stephens, PPP Over AAL5, July 1998.

[RFC 2684] D. Grossman, J. Heinanen, Multiprotocol Encapsulation over ATM Adaptation Layer 5, September 1999.

[RFC 2516] L. Mamakos, K. Lidl, J. Evarts, D. Carrel, D. Simone, R. Wheeler, A Method for Transmitting PPP Over Ethernet (PPPoE), February 1999.

[RFC 3162] B. Aboba, G. Zorn, D. Mitton, RADIUS and IPv6, August 2001.

[RFC 2461] T. Narten, E. Nordmark, W. Simpson, Neighbor Discovery for IP Version 6, December 1998.

[MIAKAWA] S. Miyakawa, draft-miyakawa-ipv6-prefix-delegation-requirement-00, work in progress, June 2002.

[TROAN] O. Troan, R. Droms, draft-ietf-dhc-dhcpv6-opt-prefix-delegation-00, work in progress, June 2002.

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.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 317 7777
Fax: +65 317 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 Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia
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

202823.L/ETMG 11/02