IPv6 Deployment Best Practices

A Technical Paper prepared for the Society of Cable Telecommunications Engineers
By

Jeff Riddel
Solutions Architect
Cisco Systems, Inc.
7200-12 Kit Creek Road
Research Triangle Park, NC 27709
919-392-7911
jriddel@cisco.com

Overview

We've all heard the "song and dance" concerning the need to transition your network to IPv6 for some time now. American registry of Internet numbers (ARIN) and other IP space registries are at or near exhaustion of IPv4 address space (http://www.potaroo.net/tools/ipv4/index.html); yet many cable service providers have not begun or are still in the early stages of deploying IPv6. Techniques such as Carrier Grade NAT (CGN) have stretched the lifetime of IPv4, but the bottom line is with the explosion of devices requiring IP addresses IPv6 MUST be deployed and the sooner the better.

I was recently on vacation at some large amusement parks and experienced what most people experience; long lines for a majority of the rides and attractions and thus a degraded quality of experience (QoE). This brought forth the analogy that these rides and attractions are very much like the current state of IPv4 networks. Just like there aren't enough seats to go around; with the explosion in the amount of Internet enabled devices there's simply not enough IPv4 addresses go around. In the amusement park case the limited amount of resources requires complex planning on the number of seats allocated per attraction similar to the necessity to carefully scrutinize the size of each IPv4 subnet allocation. The disparity between demand and the amount of people each ride or attraction can service also often necessitates shorter ride times similar to the necessity to use short dynamic host configuration protocol (DHCP) lease times with IPv4 addressing. Lastly, these amusement parks often implement workarounds where vouchers for the more popular rides and attractions can be obtained giving the customer

access to shorter lines. These workarounds usually have restrictions such as they may be valid only for certain time periods and the vouchers themselves often come at an additional cost to the end user. Furthermore, there is still a strong possibility that the demand for these vouchers will exceed the supply. I equate these techniques to the use of CGN with IPv4, it helps mitigate the limited amount of public IPv4 address space but the private IPv4 address space can also become exhausted. Not to mention connectivity problems with running certain applications thru CGN as well as issues with billing and lawful intercept when IPv4 addresses are shared.

Now, what if each ride and attraction had a limitless availability of seats. People could enjoy them whenever they wanted without waiting in line. Ride and attraction times could be longer and no workarounds or tricks would be required. The user QoE would be at an all time high! This may be a pipe dream for an amusement park but the equivalent for a service provider is a reality with IPv6 and its 128-bit address space. Each end-host subnet can simply be assigned as a /64 enabling more hosts than would ever be present. The careful and complex allocation procedures required with IPv4 are no longer necessary. The abundance of address space enables end-users addresses to be allocated for much longer periods (DHCPv6 lease times) reducing churn and load on the provisioning systems. Finally, workarounds such as CGN are no longer necessary reducing complexity and restoring the natural end-to-end architecture of the Internet.

Designing IPv6 for an access network is different from designing IPv6 in the network core; for example a solid understanding of IPv6 neighbor discovery (ND) functionality and address acquisition is vital to ensuring success. A cable access network also has its own intricacies in regards to how multicast functions in DOCSIS. This paper covers these topics and best practices when introducing IPv6 into the network. Topics include the use of IPv6 for cable modem (CM) management, IPv6 usage for CPEs, home router IPv6 usage, addressing strategy, migration plans, and security considerations.

Contents

A high-level picture of the basic components of a cable network is shown below.

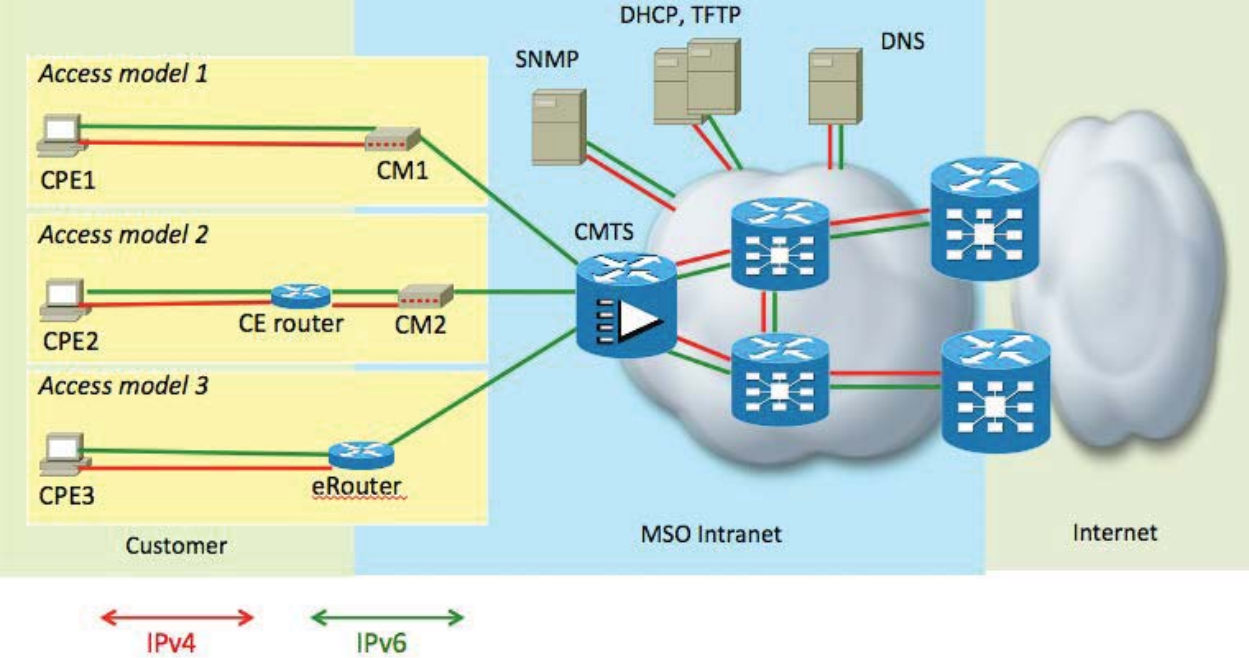

Figure 1 Moving to the IPv6 World

Moving to the IPv6 world won't happen over night and is best performed in an incremental approach.
In many cases hardware needs to be replaced along with hardware and software upgrades. IPv6 considerations should be incorporated into any network upgrades. In order to provision CPE, DHCPv6 support in the operational support systems (OSS) is required. DNS server records should be updated to support IPv6 both in terms of DNS over IPv6 transport access and the entries themselves updated to return AAAA IPv6 host records. In order to provision CMs for IPv6 the OSS needs to support DHCPv6 as well as the trivial file transfer protocol (TFTP) and time of day (TOD) protocols over IPv6. With IPv6 for CM management the network management systems (NMS) also need to be upgraded to support simple network management protocol (SNMP) transport over IPv6. The cable modem termination system (CMTS) needs to be of version DOCSIS 3.0 and the CMs need to be of DOCSIS version 3.x or DOCSIS 2.0+.

The recommended approach is to start by enabling the network for dual-stack IPv4 and IPv6 operation as indicated by the “red” IPv4 and “green” IPv6 lines between network components. Keep in mind that a dual-stack implementation enables dual communication paths to/from a device, thus both of these communication paths must be

protected from a security perspective. This means any kind of security implemented for IPv4 is also now required for IPv6; examples include protecting telnet/secure shell (ssh), SNMP, control plane policing (CoPP), routing protocols, etc. Tunneling can be used to workaround IPv4 only devices but not without a price. Also, keep in mind that the use of "dual-stack" does not result in the freeing of IPv4 address space. The ultimate goal is to remove IPv4 completely from the cable network and have only IPv6 deployed; but in all likelihood this isn't going to happen any time soon.

The typical migration approach is to begin by enabling IPv6 in the network core and move to the edge. Upgrading the OSS (and possibly the NMS) components may also be required before IPv6 can be enabled. In most deployments a redundant pair of DHCP servers are typically configured on a CMTS. Therefore the presence of DHCPv6 failover functionality in the DHCP servers is an important design factor.

Notice in the diagram the communication path to the CMs is shown as IPv6 only instead of dual-stack. The rationale here is that the cable multi-system operator (MSO) has complete control of the devices that need to communicate with the cable modems and as such can ensure all these devices are IPv6 capable limiting the need to consume IPv4 address space for cable modem management. Having no IPv4 access to the CMs also provides a measure of security against hackers only capable of IPv4. Granted if private IPv4 addressing is in use for CM management the motivation to move to IPv6 isn't as urgent but should still be part of the migration plan.

Typically there are multiple IP devices in the home enabled via the use of a home router. This router can either be a standalone device or integrated with the CM as shown by access models 2 and 3. IPv6 makes use of prefix delegation (PD) to allow global IPv6 addressing in the home rather than requiring the home router to perform NAT, as is the case with IPv4.

Global IPv6 address space is allocated from the regional Internet registries (RIR) (e.g. ARIN, RIPE, APNIC, LACNIC); generally a service provider should get a /29 to a /32 address block. If possible separate blocks should be obtained for infrastructure usage and end-user customer traffic. If a separate block isn't available a recommendation is to use either the first or last /48 block from the allocated space for infrastructure addressing. Infrastructure blocks are normally not announced to the Internet. Consequently one option is to use unique local addressing (ULA) for infrastructure addressing; as these addresses are typically blocked at peering points. If a device using ULA addressing needs to communicate to the Internet Network Prefix Translation (NPTv6) can be utilized. The use of ULA addressing allows for the internal network to remain operational during any global re-addressing event. Certain functions like Path MTU Discovery (PMTUD) may not work properly if ULA addressing is used as devices within the path are likely to filter packets sourced with ULAs. Of course IPv6 endpoints can have multiple addresses so using both globally unique addressing (GUA) and ULA is an option. Another option is to use only link-local addressing for infrastructure point-to-point links and using a loopback interface for management. Note that routing

protocols use link-local addresses for operation thus global addressing is not needed. The advantage here is that link-local addresses are only accessible from that link and hence unreachable from the Internet.

From the infrastructure block one /64 is recommended to be used for loopback addressing where each loopback would be defined with the full /128 mask. For infrastructure point-to-point links there's three options typically used for prefix length: 1) /64, 2) /126, 3) /127. The use of /64 prefixes allows a very simple design; it's also the recommended prefix length for any subnet with hosts and a point-to-point link can be thought of as a subnet with only 2 hosts. The /126 length is similar to the use of the /30 mask with IPv4 and the /127 length is similar to the /31 mask. If using /127 prefixes be careful that the two ends are in fact in the same subnet. It is recommended that even if configuring /126 or /127 prefix lengths that /64s be allocated for each link. This means using different /64s for each p2p link and only using multiple /126 or /127 prefixes from the same /64 once all /64s have been exhausted.

Any IPv6 subnet with end-hosts should be defined as a /64 regardless of the number of hosts on the segment. The practically unlimited number of hosts in a /64 network eliminates the need to scrutinize subnet size by anticipated device count as was the case with IPv4. The /64 recommendation is detailed in RFC 5375 "IPv6 Unicast Address Assignment Considerations"; also many aspects of IPv6 only work with /64 subnets.

When designing the subnetting strategy it is recommended to address on nibble boundaries; a nibble is 4-bits. Doing so makes addresses easily recognizable as IPv6 addresses are displayed as a series of hexadecimal digits (i.e. nibbles). The hierarchy is usually based on a combination of geographic factors (such as by region) and service factors (such as separating DOCSIS access from Metro Ethernet). A sample is in the following illustration:

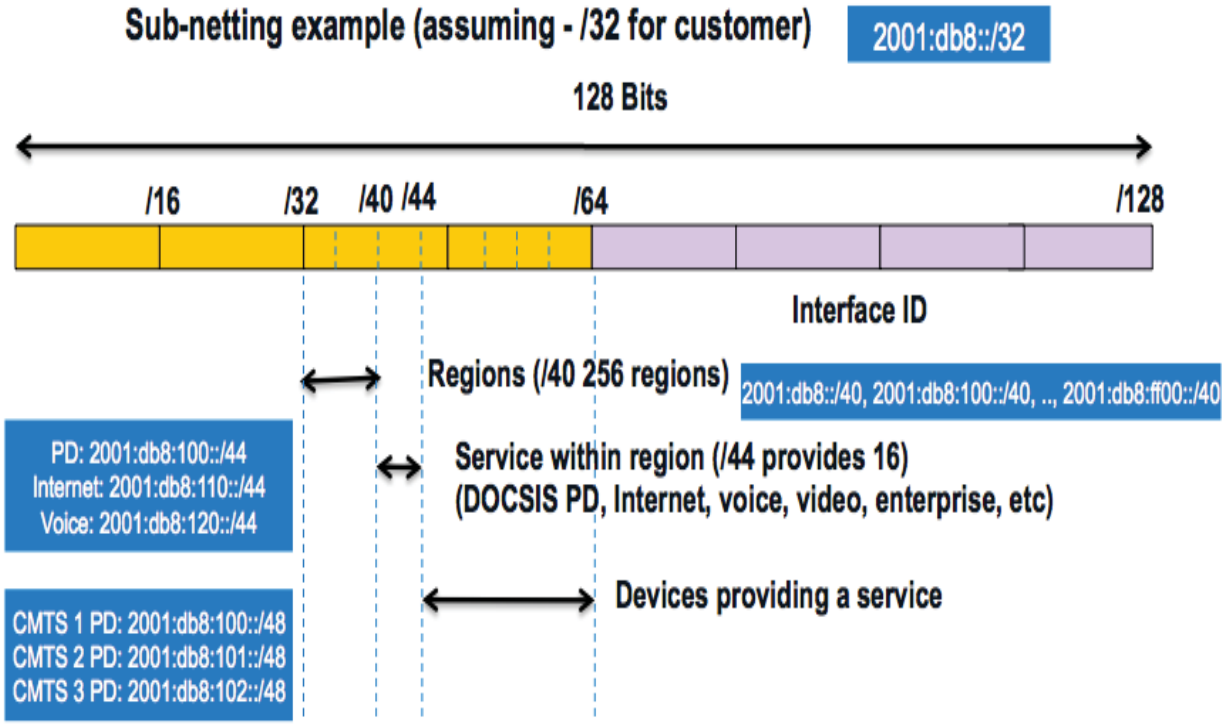

Figure 2 IPv6 Sub-netting strategies

First separating by regions allows for easy route aggregation. In this example the service provider (SP) receives a /32 allocation and uses the next 2 nibbles (8-bits) to identify up to 256 regions. It's recommended to err on the side of caution and allow for plenty of growing room; also when allocating addresses it is advisable to leave some room between allocations in case a region needs to expand. In the example the next nibble is then used to identify the service within the region such as PD, DOCSIS high-speed Internet (HSI), voice, etc. Remember the end-host networks should be defined as /64s; so in this case 20-bits (5 nibbles) are available for per service manipulation. The following URL has a great deal of information concerning best practices on IPv6 addressing - http://www.cisco.com/en/US/docs/solutions/SBA/February2013/Cisco_SBA_BN_IPv6AddressingGuide-Feb2013.pdf.

Next, let's review how neighbor discovery works in IPv6.

How devices discover each other and establish 1st hop connectivity is considerably different in IPv6 compared to IPv4 and a solid understanding is crucial especially when it comes to considerations for neighbor tuning. IPv6 ND is detailed in RFC 4861 and consists of four main areas:

1. Router discovery
2. Next-hop determination (via router or directly)
3. Address resolution
4. Neighbor Unreachability Detection (NUD)

IPv6 ND works over ICMPv6 (RFC 4443) using the Router Solicitation (RS - ICMPv6 message type 133), Router Advertisement (RA - message type 134), Neighbor Solicitation (NS - message type 135), Neighbor Advertisement (NA - message type 136), and Redirect (message type 137) messages. These messages always utilize multicast for delivery.

Cable subscriber hosts discover CMTSs via the RAs sent by a CMTS. RAs are periodically sent by a CMTS; but can also be triggered by the host sending a RS. The periodic RAs are sent to the link-local all-nodes multicast address (FF02::1) with the router link-local address as the source address. Hosts record this CMTS link-local address as a default route; note this is different than IPv4 where hosts typically learn of default routes during the DHCP process. For the cable access network the CMTS should be the only device generating RAs and thus the only default router for hosts. Consequently, aspects of IPv6 designed to provide fast detection and failover of an IPv6 default router aren't necessary in a cable environment.

The next-hop layer-2 address used by a host is based on whether the IPv6 destination prefix is "on-link" or "off-link". The "on-link" versus "off-link" concept exists in IPv4 as well as IPv6 but how it's determined is a bit a different. The following examples serves as a refresher of how it works in IPv4.

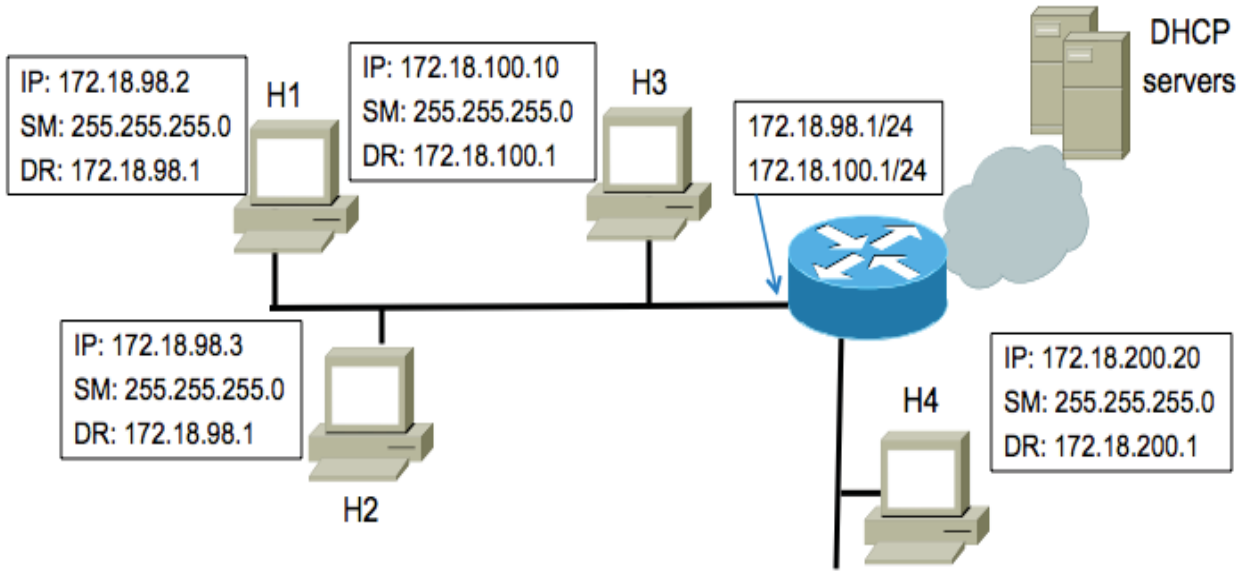

Figure 3 "On-link" versus "Off-link" in IPv4

In this example there are four hosts (H1 thru H4) all obtaining IPv4 addresses dynamically via DHCP. Hosts H1, H2, H3 are all on one network (N1) while H4 is on a different network (N2) on the router. The router has been configured with two IP subnets on N1 (172.18.98.0/24 and 172.18.100.0/24). During the DHCP process all hosts learn the /24 subnet mask and one of the router IP addresses to use as its default route.

When H1 needs to send a packet to H2 it determines the IP address of H2 is "on-link" because the IP address of H2 is on the same subnet as H1. Therefore, H1 broadcasts an ARP request on the LAN to resolve H2's IP address, which H2 (and hopefully only H2) responds to with an ARP reply containing it's layer-2 address. Note that because ARP is broadcast all devices on N1 will listen to the ARP request. When H1 needs to send a packet to H3 it determines the IP address of H3 is "off-link" because the IP address of H3 is on a different subnet than H1. Therefore, H1 sends packets destined to H3 to its default router link-layer address. The default router would then send these packets back out the same interface, as that's where H3 is located. Similarly, when H1 needs to send a packet to H4 it also determines H4 is "off-link" and sends packets to the default router.

Now lets examine how things change when the network is IPv6 instead of IPv4 as shown in the following figure:

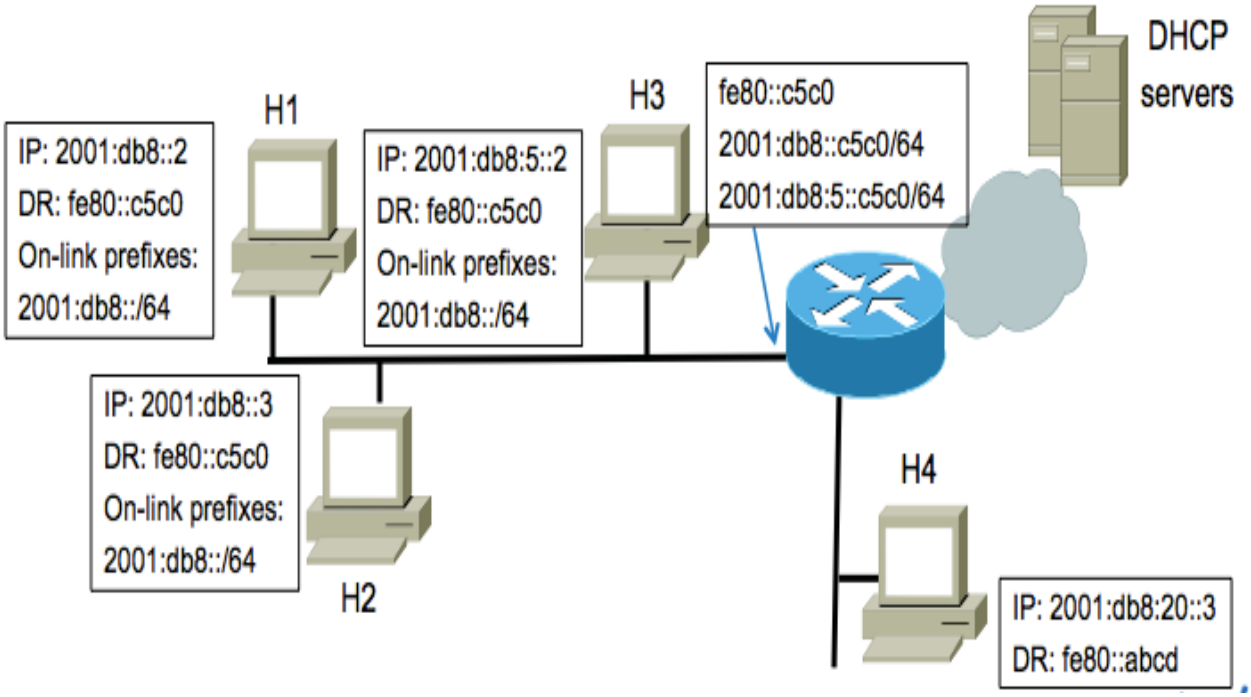

Figure 4 "On-link" versus "Off-link" in IPv6

Just as before all hosts obtain their IP address via DHCP; but now hosts learn of their default route by the receipt of RAs. The RAs also designate which prefixes are "on-link" by including the Prefix Information Option (PIO) with the L-bit set. The format of the RA message is shown in the following figure (the PIO option is in dark blue):

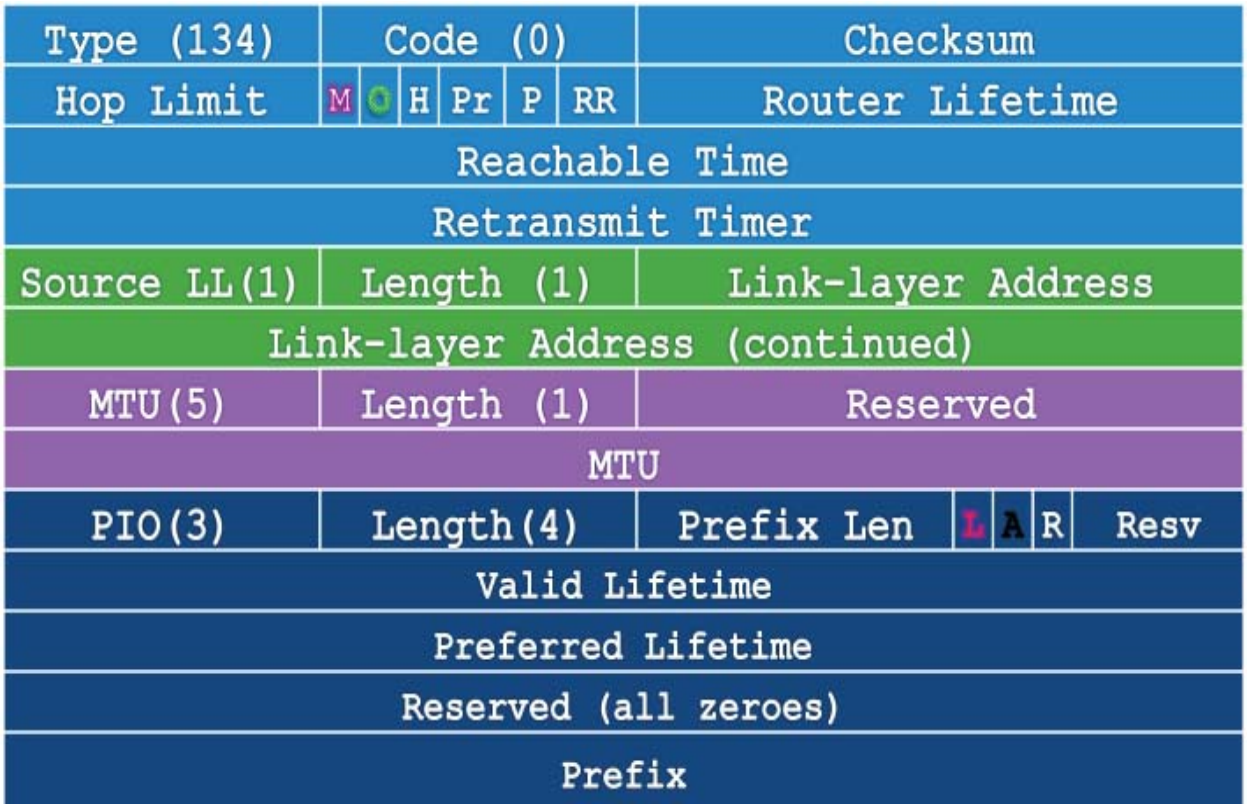

Figure 5 IPv6 Router Advertisement (RA) Format

As before hosts H1, H2, H3 are all on one network (N1) while H4 is on a different network (N2) on the router. The router has been configured with two global IPv6 addresses on N1; 2001:db8::c5c0 and 2001:db8:5::c5c0 each with a /64 length but only 2001:db8::/64 has been configured as "on-link". Thus in the router's RA message a PIO would be included with 2001:db8::/64 and the L-bit set. Logically, the router also has a link-local address fe80::c5c0 on N1.

Now when H1 needs to send a packet to H2 it sees that H2's address falls into its known list of on-link prefixes so it sends a NS to resolve H2's IPv6 address into the link-layer address. The NS is sent to the solicited-node multicast address of the target which is obtained by concatenating the prefix FF02:0:0:0:0:1:FF::/104 with the last 24 bits of the IPv6 address. H2 is the only device that should respond to the NS as it should be the only devices that has joined that solicited-node multicast address. An IPv6 device joins the solicited-node multicast address for each of its IPv6 addresses. H2 responds by sending a NA to H1 containing its link-layer address. Conversely, when H1 needs to send a packet to H3 it sees that H3's address does NOT fall into its list of known on-link prefixes so it sends the packet to its default route. At this point the router may send a Redirect message to H1 or it may simply just send the packet to H3. Note that when H3 needs to send a packet to H1 it will send a NS to resolve H1's IPv6 address as its prefix falls into the category of on-link even though H3's IPv6 address is not in this network.

In a DOCSIS network, all upstream communication is to the CMTS and all downstream communication is from the CMTS; there is no direct host-to-host communication. Consequently, it makes sense for all GUA and ULA prefixes to be configured "off-link".

For IPv6 not only is the IPv6 address to layer-2 mapping kept but the reachability state of a neighbor is tracked as well. The reachability state is tracked using the Neighbor Unreachability Detection (NUD) mechanism and has one of five values:

1. **INCOMPLETE** – initial NS sent but no NA received
2. **REACHABLE** – recently (tens of second ago) reachable by NS/NA or via confirmation from upper layer protocol e.g. TCP ACK
3. **STALE** - no longer known to be reachable but until traffic is sent to the neighbor, no attempt should be made to verify its reachability
4. **DELAY** – no longer known to be reachable and traffic has recently been sent to the neighbor, before sending NS wait for potential upper layer protocol confirmation
5. **PROBE** – no longer known to be reachable so (unicast) NS is sent

The amount of time a neighbor stays in the REACHABLE state is determined by the reachable time and is controlled by the router and signaled to hosts in the RA packets. If left unspecified the default value is 30 seconds. Confirmation from upper layer protocols that 2-way communication is successful can automatically refresh reachability. But in the case of the CMTS all CPE traffic is normally transient traffic thus CPE entries quickly transition to the STALE state. When traffic needs to be sent to entries in the STALE state it will most likely trigger probes (unicast NS) to determine reachability. Thus if there are thousands of IPv6 CPEs on a CMTS a great deal of neighbor churn can occur. Consequently, for a CMTS it is recommended to increase the reachable time from its default for a large-scale deployment.

NS packets are also used for Duplicate Address Detection (DAD). Whenever an address is assigned (no matter the method) a NS (DAD) packet is sent to confirm that address isn't already in use. This NS packet is sourced from the unspecified address (::) and sent to the solicited-node multicast address of the IPv6 address. No response is expected but if another device has this address it responds with a NA destined to the all hosts multicast address (FF02::1).

A thorough knowledge of device provisioning is also vital when it comes to deploying IPv6. As MSOs are providing a managed service they need full control of IP address assignments; thus devices should only be able to obtain their address using stateful DHCPv6. In IPv6 there is another method for hosts to dynamically obtain addresses known as StateLess Address AutoConfiguration or SLAAC. SLAAC works by hosts generating a 64-bit interface identifier and concatenating that with prefixes received in RA PIOs with the A-bit set. By not including any PIO options in the RA the cable MSO can ensure devices are unable to gain network access with SLAAC obtained addressing.

DHCPv6 is similar to DHCPv4 but it is not DHCPv4 running over IPv6. DHCPv6 is defined in RFC 3315. By default address acquisition is a four-phase process consisting of Solicit, Advertise, Request, and Reply but the Rapid Commit option can be used which eliminates the middle two phases. Hosts indicate their willingness to use Rapid Commit by including the Rapid Commit option in the Solicit message. When Rapid Commit is used the server commits the address assignment without any knowledge on if the client actually uses it or not. This can be a problem on the CMTS as it's gleaning information from DHCP and won't know which addresses a client will actually use. Hence Rapid Commit should not be used with multiple servers unless they are deployed such that only one server will respond to a Solicit message (i.e. DHCPv6 failover).

Hosts send DHCPv6 packets to the all DHCPv6 relay agents and servers multicast address (FF02::1:2). A DHCP relay agent can then send packets to a server either unicast or using the all DHCPv6 server multicast address (FF05::1:3). DHCPv6 message consist of a series of options and often times options within options. Refer to http://www.iana.org/assignments/dhcpv6-parameters/dhcpv6-parameters.xml for a list of currently defined options. Clients and services identify themselves via a construct known as a DHCP Unique Identifier (DUID). DUIDs must be unique and stable over time. There are three formats defined by RFC 3315.

1. Link-layer address plus time – DUID-LLT (1)
2. Vendor assigned unique ID based on Enterprise Number – DUID-EN (2)

3. Based on link-layer address – DUID-LL (3)

The Identity Association (IA) option is used for address assignment. There are two types: Identity Association Temporary Address (IA_TA) or Identity Association Non-Temporary Address (IA_NA). Temporary addresses have no lifetime and are assumed to be for short duration usage; non-temporary addresses are the standard type of address assignment and have an associated lifetime. These options have an Identity Association Identifier (IAID) parameter that is a four-byte value that uniquely identifies an interface of the client. The actual IPv6 address assigned is via the IADDR option as a sub-option to the IA_TA or IA_NA option. Remember in IPv6 it's perfectly valid for an interface to have multiple IPv6 addresses thus more than one IADDR option could be returned in the IA_NA or IA_TA option. From experience when multiple IADDR options are returned this is typically unintentional and a result of misconfiguration on the DHCPv6 server(s).

If the device requesting an IPv6 address is a router it should also request a prefix for delegation. This is made possible via the Identity Association Prefix Delegation (IA_PD) option, which has a similar format to the IA_NA option. An IAID parameter is included but in this case this value could identify multiple interfaces of the router. The actual IPv6 prefix assigned is via the IAPREFIX option as a sub-option to the IA_PD option and there could be more than one of these present. Refer to RFC 3633 for more information on the PD option.

Hosts are told to obtain addressing via DHCPv6 by the Managed bit or M-bit being set in the RA message. Note that the M-bit being set does NOT prohibit a device from using SLAAC; in fact it's valid for a device to obtain addressing using both DHCPv6 and SLAAC. When setting the M-bit it's common practice to also set the Other bit or O-bit indicating information other than addressing is available via DHCPv6. On the CMTS both the M-bit and O-bit should be set in the RAs.

The use and role of IPv6 in a DOCSIS network is documented in several CableLabs® specifications. Notable specifications are as follows:

- DOCSIS 3.0 MAC and Upper Layers Protocol Interface (CM-SP-MULPIv3.0)
- DOCSIS 3.0 Operations Support System Interface (CM-SP-OSSIv3.0)
- DOCSIS 2.0 + IPv6 Cable Modem (CM-SP-DOCSIS2.0-IPv6)
- eDOCSIS (CM-SP-eDOCSIS)
- IPv4 and IPv6 eRouter (CM-SP-eRouter)
- PacketCable Multimedia (PKT-SP-MM)
- PacketCable 2.0 E-UE Provisioning Framework (PKT-SP-EUE-PROV)
- PacketCable 2.0 Dual-Stack IPv6 Architecture (PKT-TR-DS-IP6**)**
- OpenCable Host Device 2.1 Core Functional Requirements (OC-SP-HOST2.1-CFR)
- CableLabs' DHCP Options Registry (CL-SP-CANN-DHCP-Reg**)**

Support for IPv6 in DOCSIS was introduced with DOCSIS 3.0. Equipment compatible with prior versions of DOCSIS does not support IPv6 functionality. Besides IPv6 other key advancements in DOCSIS 3.0 include the enablement of much larger service tier offerings via channel bonding and many enhancements relating to the handling of multicast traffic. These multicast enhancements are important to enabling support for IPv6. Recall IPv6 RA and NS messages are sent via multicast.

IPv6 functionality for the CMTS includes generating RAs towards the cable access network, participating in ND, acting as a DHCPv6 relay agent, and supporting IPv6 address assignment for CM management and hosts. IPv6 functionality for a CM includes obtaining an IPv6 address for management, participating in ND, and passing IPv6 multicast traffic to/from CPEs.

The typical home has multiple devices requiring Internet connectivity. The device enabling the complex home network is commonly referred to as a customer edge router or “CE router” which can be implemented as a standalone device or integrated into the CM (latter case referred to as CableLabs’ embedded Router (eRouter)). The eRouter is one type of embedded service/application functional entity (eSAFE) device; other types of eSAFE devices include the embedded digital voice adapter(eDVA) and embedded set-top box(eSTB). eSAFE devices undergo their own provisioning process that can be IPv4 only, IPv6 only, or dual-stack.

The basic requirements for an IPv6 CE router are defined in RFC 6204. In this document the basics of IPv6 provisioning for both the router itself as well as for connected hosts is detailed. The IPv6 CE router implements IPv6 routing between the home network(s) and the SP network. The SP facing or wide area network (WAN) side interface of the CE router acts as an IPv6 host and obtains an address using DHCPv6. The CE router also requests IPv6 prefixes for delegating on its end-user networks. The following diagram illustrates an eRouter in the home.

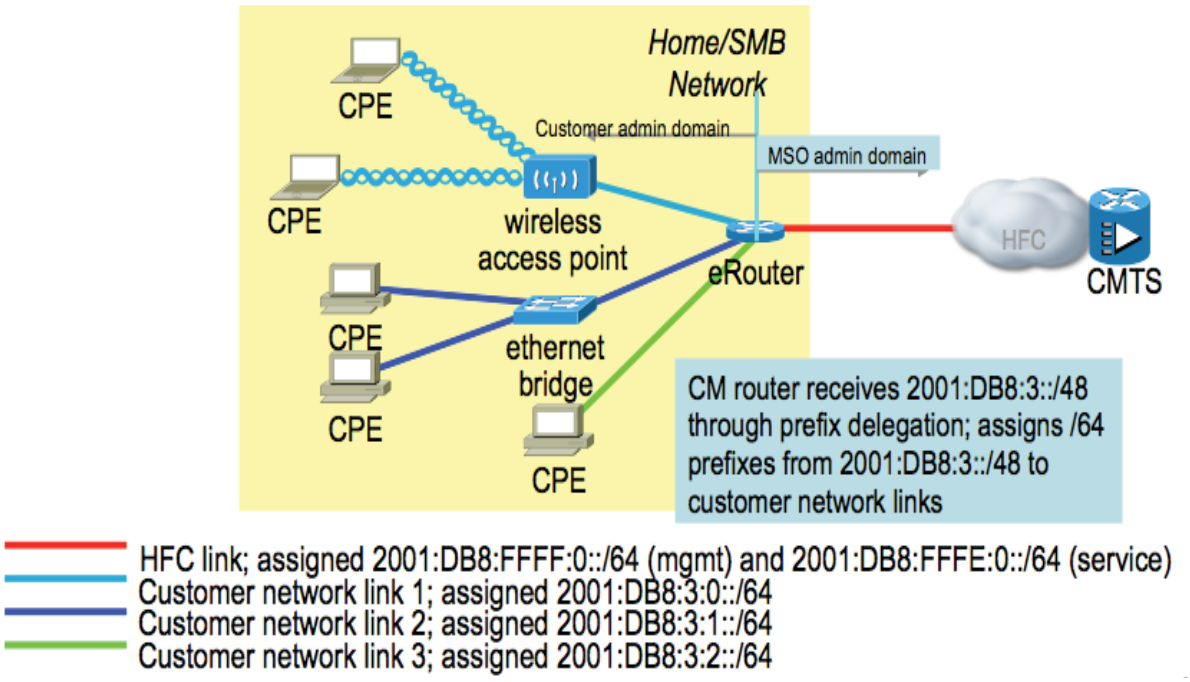

Figure 6 eRouter Reference Diagram

In this case the CM has been allocated an IPv6 address from the 2001:db8:ffff::/64 prefix for management and the eRouter component has been allocated an IPv6 address from the 2001:db8:fffe::/64 prefix for it's CMTS facing interface. These two prefixes would also be configured on the CMTS. The eRouter is also allocated the prefix 2001:db8:3::/48 for prefix delegation. This prefix is NOT configured on the CMTS but is gleaned by the CMTS during the DHCPv6 process and inserted into the IPv6 routing table. In this example the eRouter uses that /48 prefix and carves out three different /64 prefixes for use on its internal networks. The use of /64 prefixes permits the CPE within the home to use SLAAC for address acquisition.

The CPE within the home have globally routable IPv6 addresses and are typically also operating in dual-stack mode. As is normally the case with home routers today CPE are given private IPv4 address space that is NAT'ed by the home router to communicate to the outside world.

In DOCSIS 3.0, multicast operation is exclusively controlled by the CMTS. The CMTS assigns a Downstream Service ID (DSID) label used to identify a particular multicast stream. The multicast stream could be for user joined multicast traffic, multicast control traffic, or it could be for IPv6 ND traffic. All packets for a particular multicast stream in the same media access control (MAC) domain will have a common DSID label in the DOCSIS extended header. DSIDs are communicated to cable modems one of three ways:

1. Pre-registration DSID (MDD)
2. Static DSID(s) (REG-RSP)

3. Dynamic DSID(s) (DBC-REQ)

The pre-registration DSID is advertised in the MAC Domain Descriptor (MDD) messages periodically sent from the CMTS. Logically, the pre-registration DSID is used for traffic that the CM needs to process prior to reaching operational state on the CMTS (example usage is the IPv6 ND messaging needed for a CM to obtain an IPv6 address for management). Static DSIDs are communicated during the registration process and dynamic DSIDs are communicated via Dynamic Bonding Change (DBC) messages that are typically triggered off an event such as an Internet group management protocol (IGMP) or multicast listener discovery (MLD) join. The following example illustrates a DOCSIS trace of an IPv6 RA with the DSID label highlighted in the DOCSIS extended header. In this case the pre-registration DSID was used.

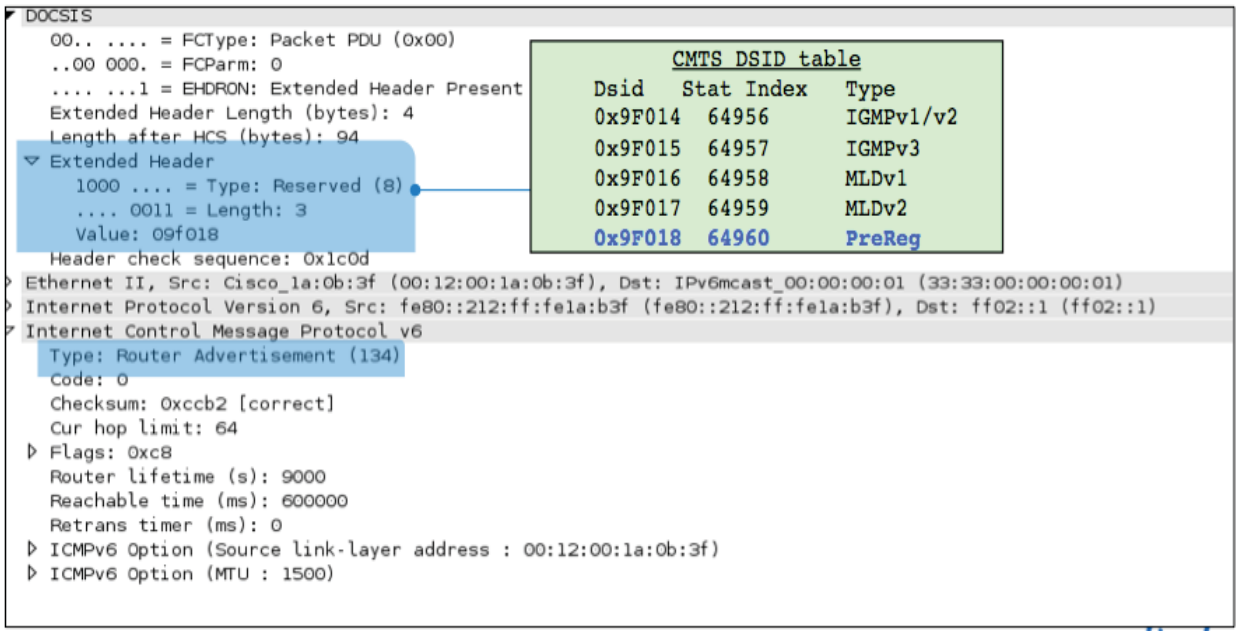

Figure 8 DSID Labeled RA

In today's deployments there are still a great deal of DOCSIS 2.0 devices. The DOCSIS 2.0 + IPv6 specification was introduced to provide aspects of DOCSIS 3.0 required to support IPv6 for DOCSIS 2.0 modems and embedded multimedia terminal adapters (eMTAs). There are two modes of multicast DSID forwarding (MDF) operation for DOCSIS 2.0+ IPv6 modems: MDF incapable (0) and MDF Group MAC (GMAC) explicit (1). Note: DOCSIS 3.0 modems operate in MDF GMAC promiscuous (2) mode.

MDF incapable modems do NOT make forwarding decisions based on DSIDs but instead are programmed to process well-known IPv6 multicast addresses (such as the all-nodes link-local MAC address). These modems also process the solicited node multicast MACs corresponding to the CM management IPv6 address and eSAFE IPv6 addresses. Since the identity of external CPE is unknown until they are connected to

the modem, solicited node multicast MACs of external CPEs must be learned via some other means such as snooping. The mechanism for multicast forwarding of external CPEs is vendor proprietary.

MDF GMAC explicit modems do make forwarding decisions based on DSID labels and thus learn about DSIDs the same way a DOCSIS 3.0 modem does. However, these modems are unable to filter unknown DSIDs in hardware and thus must be told what GMACs to forward. These modems process multicast traffic with a known DSID and a known GMAC.

MDF GMAC promiscuous (i.e. DOCSIS 3.0 modems) forward and filter all GMAC addresses with a known DSID.

If a MDF GMAC explicit or MDF GMAC promiscuous modem has MDF mode disabled by the CMTS it will most likely be unable to pass the IPv6 multicast traffic needed by external CPE.

The following diagram details the process by which a DOCSIS 3.0 or DOCSIS 2.0+ IPv6 modem obtains an IPv6 address for management.

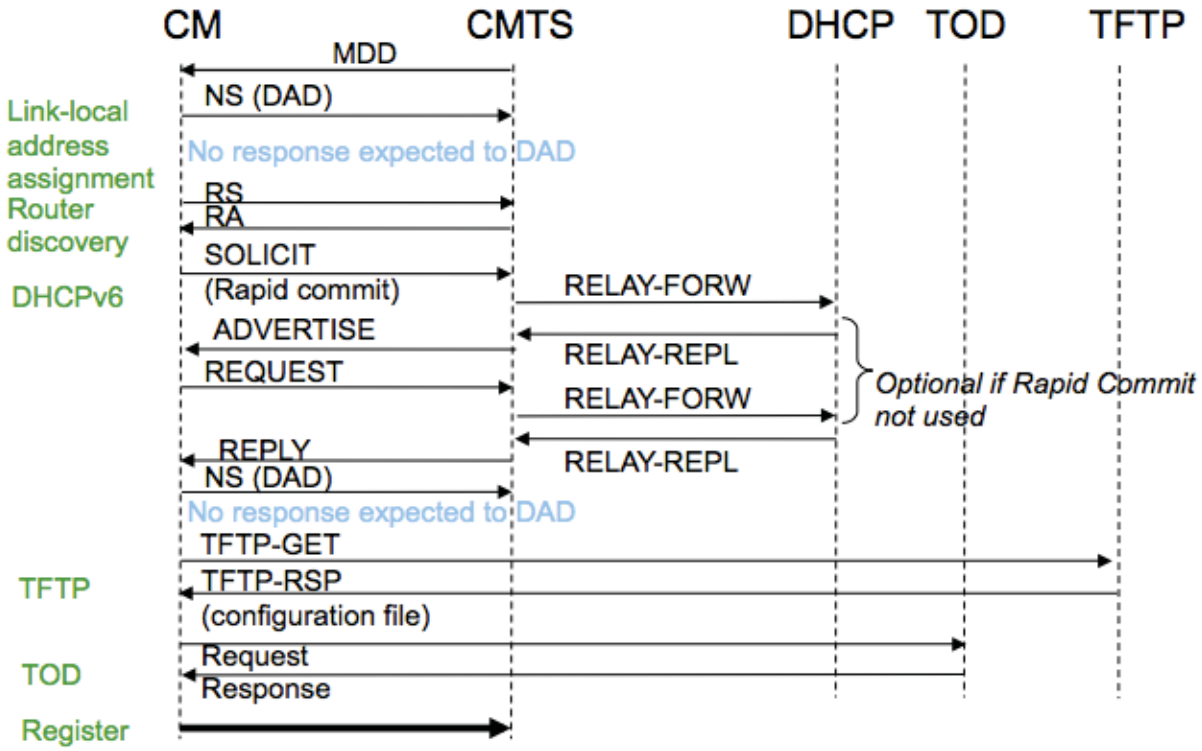

Figure 9 CM IPv6 Provisioning

If you're familiar with how channel bonding functions in DOCSIS 3.0 you know that the CMTS periodically generates a MAC message called a MAC Domain Descriptor (MDD) on its downstream channels. This MDD message lets DOCSIS devices know what downstream and upstream channels are available in the MAC domain; it also tells devices what type(s) of IP address to obtain for management. Devices can be told to obtain an IPv4 address only, an IPv6 address only, both an IPv4 and IPv6 address (dual-stack provisioning mode - DPM), or to first try to provision via IPv6 then only if that fails revert to provisioning via IPv4 (alternate provisioning mode – APM). DPM and APM provide protection that devices can still come online in the event of a problem with IPv6. These modes are not widely used, as the MSO is able to thoroughly test and ensure that IPv6 is functional; generally it either works or it doesn't. Additionally, often times the same provisioning entities are used for IPv4 and IPv6 so if the entity is down both IPv4 and IPv6 will be broken. Thus it is recommended to use IPv6 only mode for CM management.

As shown in the following illustration the MDD message will control provisioning for ALL DOCSIS 3.0 and DOCSIS 2.0+ modems in the MAC domain.

- CMTS periodically sends MDD messages which dictate CM IP establishment procedure for all DOCSIS 3.0 and DOCSIS 2.0+ CMs

IP Provisioning Mode set to IPv6
Pre-Registration DSID included

```
MDD MESSAGE
  <snip>
    dcid                          - 8
    MDD TLV, Total TLV size - 379
    MDD TLV
      Downstream Active Channel List
        Channel ID:               1
        Frequency:                687000000Hz
        Modulation Order/Annex:   256 QAM/Annex B
        Primary Capable:          Primary-Capable
  <snip>
      Downstream Active Channel List
        Channel ID:               8
        Frequency:                729000000Hz
        Modulation Order/Annex:   256 QAM/Annex B
        Primary Capable:          Primary-Capable
        CM-STATUS Event Bitmask:  0x36
                                  MDD Timeout
                                  QAM FEC failure
                                  MDD Recovery
                                  QAM FEC recovery
      MAC Domain Downstream Service Group
        MD-DS-SG ID:              1
        Channel IDs:              1
                                  2
                                  3
                                  4
                                  5
                                  6
                                  7
                                  8
  <snip>
      IP Initialization Parameters
        IP Provisioning Mode:     IPv6
        Pre-Registration DSID:    585767
  <snip>
```

Figure 10 Dictating CM IP Provisioning

As some particular modem models may have issues with provisioning IPv6 this behavior can be undesirable. The MSO can have tighter control of what modems can provision IPv6 by setting the CableLabs SNMP object CmMdCfg in either the CM configuration file or via a SNMP Set. The format of this object is shown below:

Attribute Name	Type	Access	Type Constraints	Units	Default
ifIndex	InterfaceIndex	key	InterfaceIndex of MAC Domain interface	N/A	N/A
IpProvMode	Enum	read-write	ipv4Only(0) ipv6Only(1) honorMdd(4)	N/A	honorMdd
IpProvModeResetOnChange	TruthValue	read-write	true(1) false(2)	N/A	false
IpProvModeResetOnChange HoldOffTimer	Unsigned32	read-write	0…300	seconds	0
IpProvModeStorageType	StorageType	read-write	volatile(2) nonVolatile(3)	N/A	nonVolatile

Figure 11 MDD Override SNMP object

The MSO could configure MAC domains for IPv4 provisioning and override the setting on known “certified” modem models to IPv6. Alternatively, they could configure the MAC domains for IPv6 provisioning and override the setting for all “uncertified” modem models to IPv4.

Once told to provision via IPv6 the CM sends a DHCPv6 Solicit message; an example is shown below:

```
CMTS DHCPV6: Incoming DHCPv6 SOLICIT from 38c8.5cb2.540a (sid 51 prim_sid 51) in Cable6/0/0.
IPv6 DHCP: detailed packet contents
  src FE80::3AC8:5CFF:FEB2:540A (Bundle99)
  dst FF02::1:2
  type SOLICIT(1), xid 5531996
  option RAPID-COMMIT(14), len 0
  option RECONF-ACCEPT(20), len 0
  option VENDOR-CLASS(16), len 15
  option ORO(6), len 2
    VENDOR-OPTS
  option VENDOR-OPTS(17), len 284
  option CLIENTID(1), len 10
    0003000138C85CB2540A
  option IA-NA(3), len 12
    IAID 0x5CB2540A, T1 0, T2 0
  option ELAPSED-TIME(8), len 2
    elapsed-time 0
CMTS DHCPV6: Adding CMTS relay agent options DOCSIS version 3.0 and CM MAC 38c8.5cb2.540a
```

Lists requested CableLabs options

Figure 12 Example DHCPv6 Solicit

If no DHCPv6 Solicit is seen and modems are getting stuck in the ranging complete state it most likely means IPv6 RA's with the M-bit set are not being generated by the CMTS. If modems are provisioning via IPv4 instead of IPv6 it most likely indicates a problem with CMTS configuration or MDD override being set to IPv4 for the modems.

Notice the solicit is sourced from the modem's link-local address and destined to the all DHCPv6 server and relay agent multicast address. The rapid commit option is included as well as a vendor class option that identifies the request as coming from a DOCSIS 3.0 CM. The vendor options includes CableLabs specific options such as the CableLabs option request option asking for time protocol server, time offset, TFTP server address, configuration file name, and Syslog server address. Lastly, notice the DUID identifying the client (in this case using the LLC format) and the request for a non-temporary IPv6 address.

When the CMTS receives this message it relays the message to the configured DHCPv6 servers. The relayed message contains several important parameters in addition to the message sent by the client. For example the link-address which is a CMTS IPv6 address from the link where the client is located; this is used by the provisioning system when determining what IPv6 prefixes are eligible for the client. The interface-ID references the physical interface where the DHCPv6 originated. Relay agent options are also included which specify the CMTS's DOCSIS capabilities and the

MAC address of the cable modem originating the DHCPv6 packet (this information is used by the provisioning system to associate CPE with the connected CM).

Assuming an acceptable address is available the provisioning systems will then respond with either a DHCPv6 advertise or DHCPv6 reply message depending on whether or not rapid commit is enabled. Remember if multiple DHCPv6 servers are deployed and no DHCPv6 failover mechanism is in place (such that only one DHCPv6 server responds) it is strongly recommended that rapid commit be disabled. The DHCPv6 preference option can be used to try to make addresses from one DHCPv6 server preferred over the other. This is just an 8-bit value where the client should prefer higher values; however note that some clients may not correctly honor the preference setting. One workaround to not having a DHCPv6 failover mechanism is to use a deterministic method for allocating IPv6 addresses such that it is guaranteed multiple servers will allocate the same address to a given client. A method to do this is to use the extended unique identifier 64 (EUI-64) format for the interface-identifier portion of the IPv6 address; this method constructs a 64-bit interface-identifier generated from the device MAC address. The hexadecimal digits “FF FE” are inserted into the middle of the 48-bit MAC address and the Universal/Local or U/L bit 7 is toggled. With this method any client should be guaranteed not to get the same IPv6 address regardless of which server they communicate with.

As mentioned before any IPv6 subnet with end-hosts should be configured as a /64 network; this is the case for CM management networks. It is strongly recommended that only /64 networks be defined. As a matter of habit many operators tend to setup DHCPv6 scopes similar to how DHCPv4 scopes are setup; remember that exhaustion of addresses from a DHCPv6 scope should never happen. Consequently, operators may choose to use longer lease times for DHCPv6 to reduce the amount of renew traffic hitting the provisioning servers.

If the CM never receives an advertise or reply message it more than like means there is a problem with IPv6 routing or a misconfiguration in the provisioning system. An exhaustion of IPv6 address space should never happen. In fact a single /64 prefix should be more than sufficient for any type of service. This is in contrast to IPv4 provisioning where often times multiple scopes had to be defined for a service. Be aware that if more than one IPv6 scope for a service is setup be sure the provisioning system has the intelligence not to assign an address from each scope.

The client device should collect advertise messages from all servers and respond to one with a request message (assuming the advertise is valid and contains all mandatory requested options). Again if more than one advertise message is received and all are valid the client **should** select the one with the highest preference. Note that if a client receives an advertise with the maximum preference value of “255” it can immediately respond with a request as its impossible for another advertise to have a better preference. If a request message is never sent this more than likely means all advertise messages are invalid such as if a mandatory option is lacking. If rapid commit is

enabled the symptom would be the provisioning process not progressing past the client receiving the reply message.

All servers should receive the request message, which indicates the server the client, chose. Thus servers whose advertises weren't selected are free to use the address they offered for other clients. The selected server would then commit the binding and respond with the Reply message.

The provisioning process from here on out is nearly identical to IPv4 provisioning. The TFTP and time of day (ToD) stages occur followed by the modem registering with the CMTS. A failure in the TFTP process most likely means a reachability issue to the TFTP server, a missing configuration file, or an invalid configuration file.

In order to manage (i.e. SNMP poll) CMs in IPv6 only mode either SNMPv3 must be used or SNMPv2c Coexistence (DOCSIS type length value or TLV 53) needs to be configured. The legacy DOCSIS TLV 11 for SNMPv2c does NOT support IPv6. The modems should also support HTTP, HTTPS, telnet, SSH, etc. over IPv6 transport. Just as with IPv4 these protocols need to be secured or disabled as appropriate.

Next, let's examine the use of IPv6 for CPEs. A CPE will typically have at least two IPv6 addresses, a link-local address and one or more global addresses learned via DHCPv6 and/or SLAAC. The link-local address may or may not use the EUI-64 format for the interface identifier. In fact many common operating systems use randomly generated interface identifiers by default. For security reasons the EUI-64 should not be used for CPE global IPv6 address either. When the EUI-64 format is used it is possible for other systems to track the location of a CPE. The last 64-bits of the CPE address would remain constant whereas only the first 64-bits would change as a device moves from home to work for example. Also, the EUI-64 format can also give clues to the manufacturer of the CPE such that a malicious observer could use that information to launch an attack against known vulnerabilities. Consequently, it is recommended to use randomly generated interface identifiers for global CPE addresses. Recall, in the service provider model all CPE global addresses should be assigned via DHCPv6 and CPEs should not be able to obtain network access using SLAAC addresses. Thus, the provisioning system should be setup to assign randomly generated interface-identifiers in DHCPv6.

The following diagram is a trace of a CPE provisioning:

Source	Destination	Protocol	Info
::	ff02::1:ff49:25c1	ICMPv6	Neighbor Solicitation for fe80::426c:8fff:fe49:25c1
fe80::426c:8fff:fe49:25c1	ff02::1	ICMPv6	Neighbor Advertisement fe80::426c:8fff:fe49:25c1 (ovr) is at 40:6c:8f:49:25:c1
fe80::426c:8fff:fe49:25c1	ff02::2	ICMPv6	Router Solicitation from 40:6c:8f:49:25:c1
fe80::212:ff:fe1a:b3f	ff02::1	ICMPv6	Router Advertisement from 00:12:00:1a:0b:3f
fe80::426c:8fff:fe49:25c1	ff02::1:2	DHCPv6	Solicit XID: 0xbc989d CID: 000100011811aac2406c8f4925c1
fe80::212:ff:fe1a:b3f	fe80::426c:8fff:fe49:25c1	DHCPv6	Advertise XID: 0xbc989d CID: 000100011811aac2406c8f4925c1 IAA: 2001:db8:fffe:0:514c:6
fe80::212:ff:fe1a:b3f	fe80::426c:8fff:fe49:25c1	DHCPv6	Advertise XID: 0xbc989d CID: 000100011811aac2406c8f4925c1 IAA: 2001:db8:fffe:0:f060:7
fe80::426c:8fff:fe49:25c1	ff02::1:2	DHCPv6	Request XID: 0x6ba2c7 CID: 000100011811aac2406c8f4925c1 IAA: 2001:db8:fffe:0:f060:794
fe80::212:ff:fe1a:b3f	fe80::426c:8fff:fe49:25c1	DHCPv6	Reply XID: 0x6ba2c7 CID: 000100011811aac2406c8f4925c1 IAA: 2001:db8:fffe:0:f060:7945:
::	ff02::1:ffb2:188	ICMPv6	Neighbor Solicitation for 2001:db8:fffe:0:f060:7945:20b2:188
fe80::426c:8fff:fe49:25c1	ff02::16	ICMPv6	Multicast Listener Report Message v2
fe80::426c:8fff:fe49:25c1	ff02::1	ICMPv6	Neighbor Advertisement 2001:db8:fffe:0:f060:7945:20b2:188 (ovr) is at 40:6c:8f:49:25:

Figure 13 CPE Provisioning trace

If the CM the CPE is behind is not MDF enabled then the CMTS RA will not be passed thru to the CPE (unless the modem is of the DOCSIS 2.0+ MDF incapable type) preventing the CPE from initiating the DHCPv6 address acquisition process. Consequently, make sure both the CMTS and CM are capable and configured for MDF operation.

Note in the figure multiple advertise messages are received due to no DHCPv6 failover mechanism being in place. Just as was the case with CM management a single /64 prefix is recommended for CPE addressing. With no DHCPv6 failover this prefix needs to be split amongst the servers – thus the first /65 would be defined as the available allocation range on one server and the second /65 would be defined as the available allocation range on the other server. Also just like the case with cable modem management rapid commit needs to be disabled, preference can be used to influence CPEs to prefer one server over the other (again no guarantee CPEs will correctly honor the preference setting), and longer lease times can be used compared to DHCPv4.

For security reasons it is recommended to only allow CPE with DHCPv6 obtained global addresses to gain network access and not populate neighbor tables and CPE tables based on ND messages alone. There may be circumstances where CPEs with static IPv6 addresses are required; for example a business customer hosting a server. In this case the DOCSIS source address verification (SAV) group functionality should be used for that particular modem. But as a general measure DHCPv6 should solely be used to populate neighbor tables and CPE tables. The NA messages from CPE are untrusted and should not be used to populate databases. Malicious clients could attempt to masquerade as other users by falsifying NA messages or they could attempt to bombard the CMTS with bogus NAs consuming valuable memory and CPU resources.

There may be circumstances where the CMTS loses data for CPE (for example a chassis reboot). In this case the use of DHCP lease queries should be available to repopulate these entries. This requires lease query support on both the CMTS and the provisioning system. Note there are two versions of DHCPv6 lease query: unicast lease-query (RFC 5007) and bulk lease-query (RFC 5460). When the CMTS receives data for an unknown CPE it issues a lease-query request to the provisioning system, which responds with the bindings for that CPE (including the corresponding cable modem identity) assuming that CPE is known. If the CPE is unknown those packets should be dropped by the CMTS.

Once the CMTS is configured for IPv6 and starts sending RAs any CPE behind an IPv6 capable CM is able to attempt to obtain IPv6 network connectivity. Most client devices are enabled for IPv6 by default and are just waiting for the RAs. The MSO may want to restrict the modems that can have IPv6 CPEs; perhaps to only "certified" modem types or they may just want to limit the amount of CPEs that can provision via IPv6. There are a couple ways to accomplish this task:

1. One is to set the DOCSIS DEVICE-MIB objects docsDevFilterLLCUnmatchedAction and docsDevFilterLLCTable in the CM configuration file to explicitly block IPv6 to/from the modem's Ethernet port.
2. Another option is to use Subscriber Management Filter Groups (DOCSIS TLV 37 in the cable modem configuration file or via CMTS configuration) to block IPv6 traffic on a per device basis.
3. A third option would be to use Upstream Drop Classifiers (DOCSIS TLVs 60 and 62) in the modem configuration file to drop the IPv6 traffic.

Recall in the case of an IPv6 CE router (either an eRouter or a standalone device) it requests an IPv6 address for it's CMTS facing interface and an IPv6 prefix for delegation on it's LAN interfaces. Whether or not an eRouter provisions IPv6 can be controlled by a TLV in the cable modem configuration file. The eDOCSIS specification defines TLVs for each type of eSAFE entity; for an eRouter TLV 202 is used and sub-TLV 1 dictates the eRouter initialization mode as shown in the following figure:

Mode	IPv4 Behavior	IPv6 Behavior
Disabled (0)	CM bridges all traffic	CM bridges all traffic
IPv4 Protocol Enabled (1)	IPv4 traffic forwarded via NAPT	IPv6 traffic not forwarded
IPv6 Protocol Enabled (2)	IPv4 traffic not forwarded	IPv6 traffic forwarded
Dual IP Protocol Enabled (3)	IPv4 traffic forwarded via NAPT	IPv6 traffic forwarded

Figure 14 eRouter Provisioning TLV

For example the following setting will trigger the eRouter to provision dual-stack:

GenericTLV TlvCode 202 TlvLength 3 TlvValue 0x010103; /* dual IP */

For a standalone CE router ideally all they would require is to see a RA message with the M-bit set to trigger DHCPv6. Unfortunately in practice this is not always the case and often times some initial configuration is required to turn on IPv6 operation.

The next figure shows a DHCPv6 solicit from an eRouter as its relayed from the CMTS to the DHCPv6 provisioning servers.

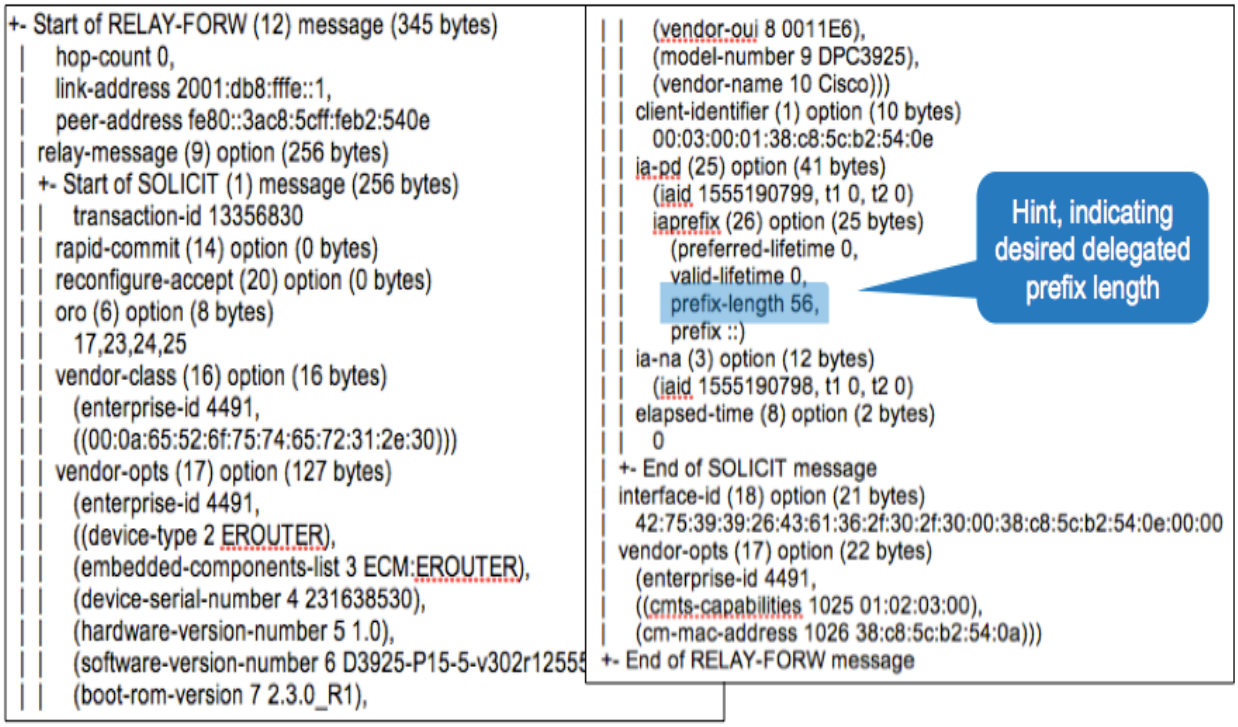

Figure 15 eRouter DHCPv6 Solicit

Notice both the IA_NA and IA_PD options are included and within the IA_PD option the eRouter includes a hint indicating the desired prefix length it would like to receive. There's no guarantee this request will be granted by the provisioning system and in fact as can be seen in the next figure a /60 prefix was returned.

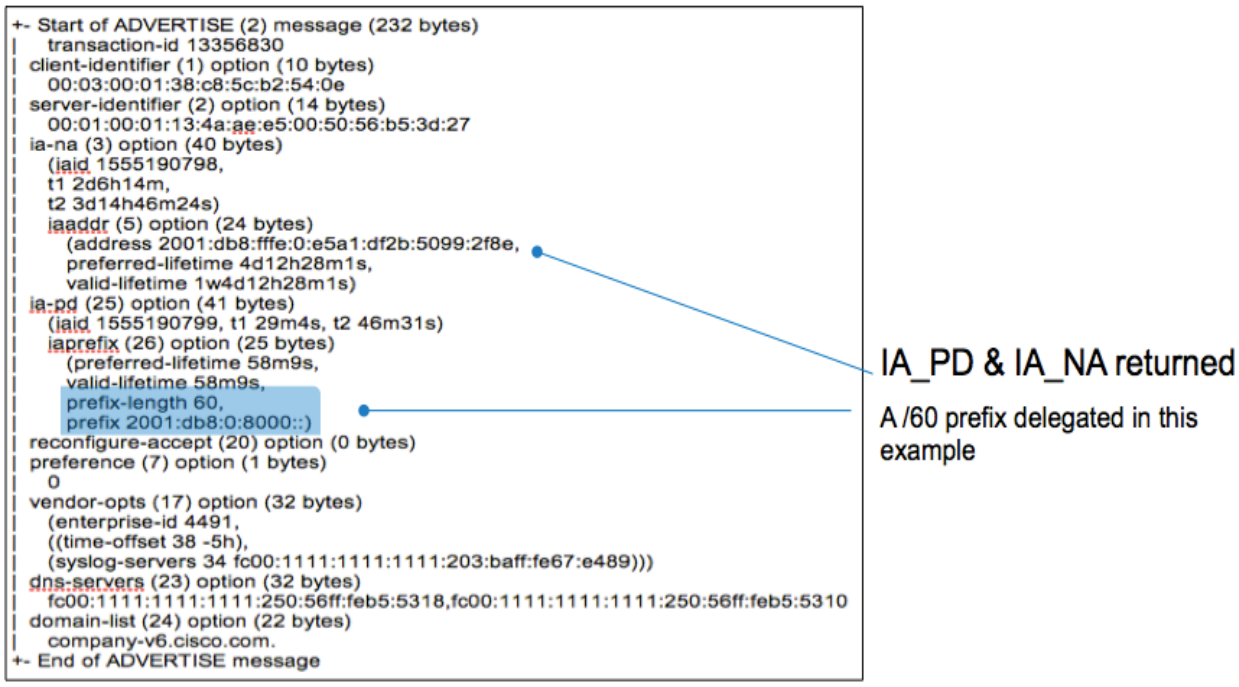

Figure 16 eRouter DHCPv6 Advertise

Just as was the case with ordinary CPE if no DHCPv6 failover mechanism is in place rapid commit needs to be disabled; preference can be used to attempt one server to be preferred over the other; and the prefix delegation range needs to be split between the provisioning servers. So the question may come up on how big to define the delegated prefix range in the provisioning system? The recommendation is to make it big enough to accommodate the worst case for the foreseeable future. If you're delegating /64 prefixes today you may be delegating /60 or /56 prefixes five years from now, so if possible set aside that address space from the beginning.

Delegated prefixes do not need to be defined on the CMTS, as these are one-hop away routes. However, CE routers do not communicate to the CMTS with a routing protocol to advertise these routes either. In its role as a relay agent the CMTS gleans the PD routes from the DHCPv6 messaging and inserts these routes into the IPv6 routing table and into the cable modem host database. These gleaned routes do need to be redistributed into the routing protocol so entities north bound of the CMTS learn these routes; typically these routes are aggregated before redistribution.

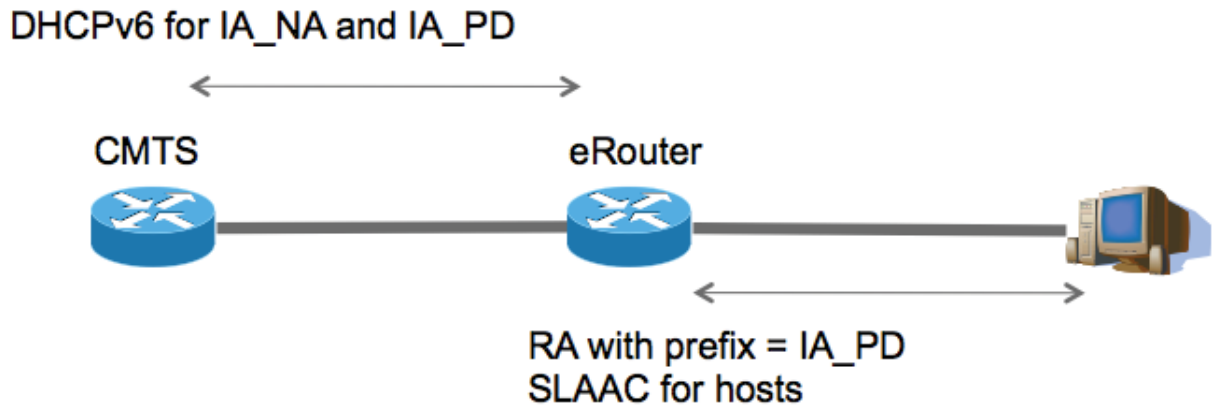

Figure 16 eRouter Functionality

Once a CE router receives a delegated prefix it divides that into /64 prefixes for advertisement on its LAN interfaces. Within the CE premises the use of SLAAC for host address assignment is expected. SLAAC provides ease of dynamic provisioning for hosts. Thus the CE router generates it's own RA messages containing the PIO with a sub-prefix from the delegated prefix with the A-bit set (the L-bit is set as well); the Managed bit is off telling hosts not to use DHCPv6 for address assignment but the Other bit is enabled telling hosts that DHCPv6 can be used to obtain non-address information such as DNS server addresses. The CE router acts as the DHCPv6 server in this case and proxies the DNS information it learned during its provisioning process to the hosts. Alternatively, the RA message itself can contain the DNS information (refer to RFC 5006 and 6106) however not all hosts support this functionality. The following figure is a trace of a host provisioning using SLAAC.

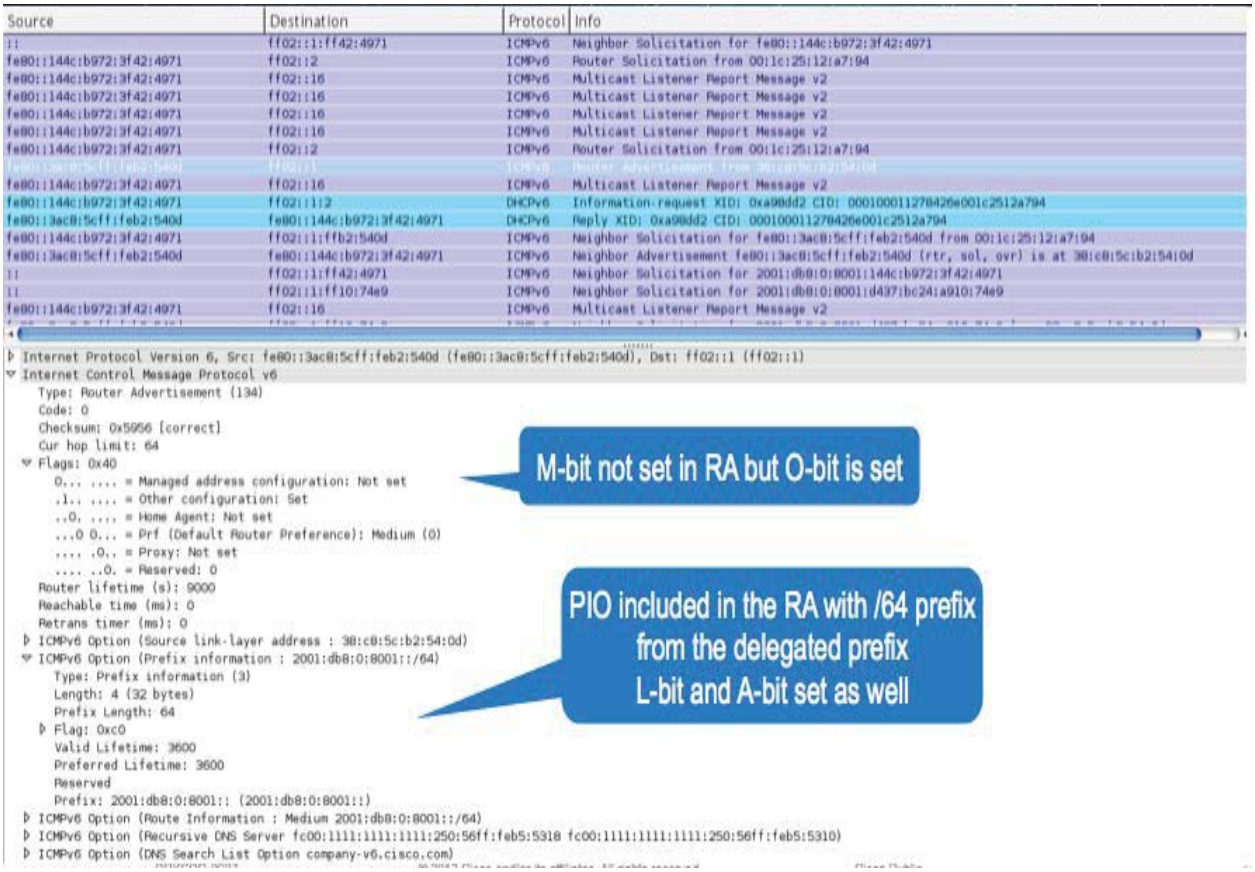

Figure 17 CPE SLAAC Provisioning

Typically a range of IPv6 prefixes available for delegation are allocated on a per CMTS basis. Thus if a CE router is moved form CMTS "A" to CMTS "B" due to a node-split by the MSO its delegated prefix changes. Node-splits are where a MSO alters the RF plant resulting in a portion of the customer base moving to a new MAC domain in order to alleviate congestion. This new MAC domain may or may not be on the same CMTS. If it's a different CMTS the delegated prefix changes, which could be problematic for some end customers; such as a business customer that requires it's addressing to remain constant.

Fortunately, the DOCSIS specifications provide a solution to this known as "prefix stability". With prefix stability the prefix delegation range is allocated to a group of CMTSs in a common physical location. A routing protocol runs between these CMTS and delegated prefixes are redistributed into the routing protocol such that all CMTSs in the site learn which CMTS the route is currently located. Each CMTS that receives a routing update containing one of these delegated prefixes checks to see if it has a local CM where this prefix was also learned (via the gleaning process). If the CM is currently online the routing update is ignored. If the CM is offline the routing entry for that prefix is updated and now points to the new CMTS instead of this CM's CPE. The solution as defined by DOCSIS assumes the CPE router will reset thus the new CMTS gleans the

PD route and sends a routing update to the old CMTS to purge the PD. Note, the ability to allocate prefix delegation ranges per location instead of per CMTS may be reason enough to use prefix stability. When prefix stability is used each CMTS in the location needs the exact routes of each delegated prefix in the location; thus route aggregation cannot be performed on the CMTSs but instead on the aggregation routers.

The discussion thus far has assumed a dual-stack implementation for CPEs. A dual-stack implementation does nothing to help IPv4 address exhaustion issues; it's simply an interim step for the migration to IPv6. The ideal solution is to have an IPv6 only network and thus IPv6 only CPEs. The reality is that not all the Internet is reachable via IPv6 and this is likely to continue to be the case for the foreseeable future. In order to enable IPv6 only CPEs and continue to be able to communicate with the IPv4 Internet the SP needs to implement an address family translation mechanism; IPv6 is not backward compatible with IPv4. Thus the translation mechanism will have to include some type of IPv4 address sharing mechanism. A number of different translation mechanisms exist today and new ones are evolving. Some differentiating criteria for these mechanisms include:

1. Is translation functionality required on the CPE itself?
2. Is a router/gateway with translation functionality need at the customer premises?
3. Is the address family translation stateful or stateless?
4. Is tunneling an aspect of the technology?

The following figure shows some of the translation mechanisms.

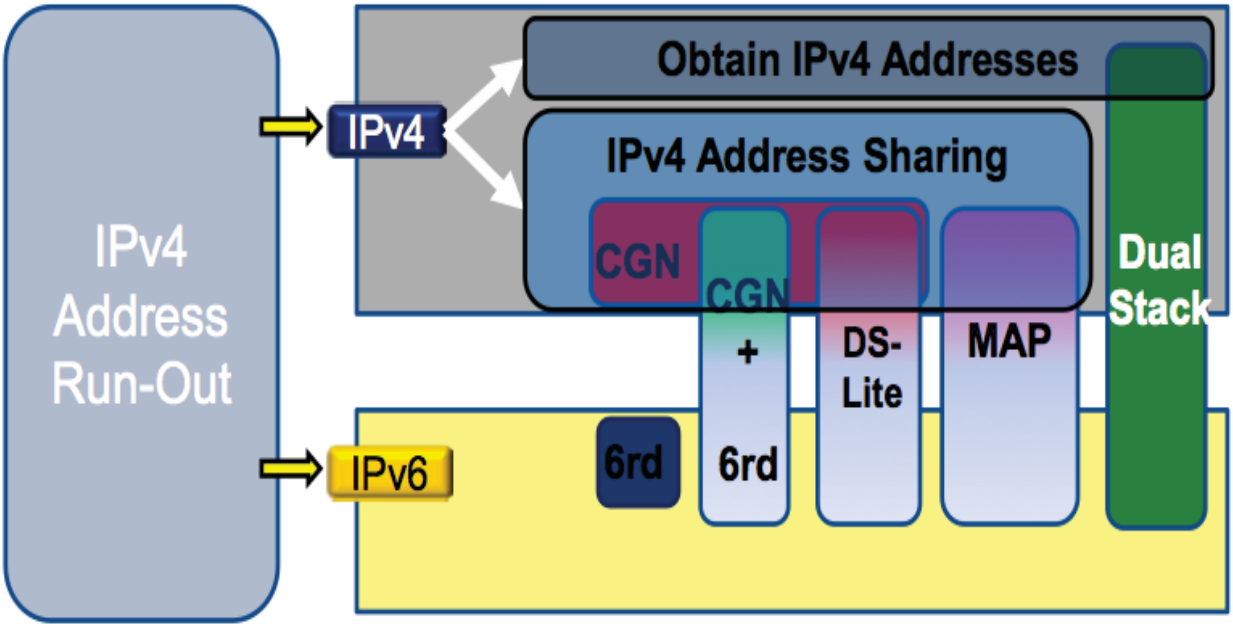

Figure 18 IPv6 Transition Mechanisms

The mechanisms of interest are the ones that allow the cable access network to be IPv6 only; recall 6rd relies on a home router tunneling IPv6 traffic over an IPv4 network. Private IPv4 addressing could be used in combination with CGN however this solution doesn't take advantage of an IPv6 enabled CMTS.

One option would be to use stateful NAT64 with DNS64 (also referred to as address family translation 64 or AFT64). With this option CPE can be made IPv6 only without needing additional translation functionality on the CPE or a router within the customer premises. If the customer attempts to access a domain name that resolves to only IPv4 addresses; the service provider's DNS server detects this and synthesizes an IPv6 address to return to the CPE. Consequently, the CPE thinks the domain name does have IPv6 accessibility. A NAT64 device exists in the SP network and translates this synthesized IPv6 address to the actual IPv4 address. The source IPv6 address is also translated to an IPv4 address.

Another option is to use Dual-Stack Lite or DS-Lite. DS-Lite requires some translation functionality either directly on the CPE or on a router within the customer premises. This functionality is known as the B4 element (Basic Bridging BroadBand element). The B4 element uses private IPv4 addressing locally – no IPv4 address is obtained from the SP. When the CPE needs to communicate to an IPv4 endpoint the B4 tunnels the IPv4 packet inside an IPv6 packet and sends it to a device known as the AFTR element located in the core infrastructure. The AFTR identity is learned during the CPE provisioning process. The AFTR takes the IPv4 packet and translations the CPE IPv4 address to a public IPv4 address.

Both of the aforementioned options use stateful translation. Stateful translation has many drawbacks: 1) CPU and memory resources are consumed on the translating devices, 2) Traffic must be sent and received thru the same device, and 3) Redundancy of the translating devices is difficult.

The Mapping of Address and Port (MAP) translation technologies provide stateless translation in the SP network. There is a tunneling/encapsulation version of MAP (MAP-E) and a pure translation version (MAP-T). MAP allows sharing of IPv4 addresses across an IPv6 network by allocating ranges of TCP/UDP ports via "rules". These rules allow the translation between IPv4 and IPv6 to be stateless thus there is no single point of failure and traffic flows can by asymmetric. The CPE has translation functionality to translate between IPv4 and IPv6 by these "rules". Traffic destined to IPv4 endpoints is sent to the MAP Border Relay (MAP BR) in the core.

In summary the transition to IPv6 is a challenging one. The information provided in this paper hopefully has supplied you with strategies and lessons learned to make the transition to IPv6 successful and easier. The first step is to have a solid understanding of how IPv6 works specifically regarding the access leg of the network. Having a solid grasp of IPv6 ND and address acquisition are vitally important. Intricacies of how IPv6

functions on a DOCSIS network are not to be overlooked as well. Using IPv6 for CM management saves IPv4 address space and allows a controlled environment for rolling out IPv6. For CPEs IPv6 can be deployed with PD enabling the support for advanced home networks. Remember, incremental migration to a dual-stack network is the recommended first step but the end goal should be to move to an all IPv6 access implementation. Translation mechanisms will in all likelihood need to be implemented, as it will be quite some time before IPv4 only parts of the Internet or your intranet are obsoleted.

Bibliography

World IPv6 Launch Information - http://www.worldipv6launch.org/

Comcast's IPv6 Information Center - http://www.comcast6.net/

Cisco's IPv6 Addressing Guide - http://www.cisco.com/en/US/docs/solutions/SBA/February2013/Cisco_SBA_BN_IPv6AddressingGuide-Feb2013.pdf

Cisco IPv6 Web Site - http://www.cisco.com/ipv6

Cisco IPv6 Lab Web Site - http://6lab.cisco.com/index.php

Abbreviations and Acronyms

AFTR	Address Family Transition Router
APM	Alternate Provisioning Mode
B4	Basic BroadBand Bridging
CGN	Carrier Grade NAT
CoPP	Control Plane Policing
DAD	Duplicate Address Detection
DPM	Dual-stack Provisioning Mode
DSID	Downstream Service ID
DS-Lite	Dual Stack Lite
DUID	DHCP Unique ID
eDVA	embedded Digital Voice Adapter
eSAFE	embedded Service/Application Functional Entity
eSTB	embedded Set-Top Box
GMAC	Group MAC
GUA	Global Unique Address
IA	Identity Association
IA_NA	Identity Association - Non-temporary Address
IA_PD	Identity Association - Prefix Delegation
IA_TA	Identity Association - Temporary Address
MAP	Mapping of Address and Port
MDD	MAC Domain Descriptor
MDF	Multicast DSID Forwarding
NA	Neighbor Advertisement
ND	Neighbor Discovery
NMS	Network Management System
NS	Neighbor Solicitation
NUD	Neighbor Unreachability Detection
OSS	Operational Support System
PIO	Prefix Information Option
PMTUD	Path MTU Discovery
RA	Router Advertisement
RS	Router Solicitation
SAV	Source Address Verification
SLAAC	StateLess Address AutoConfiguration
ULA	Unique Local Address