COMCAST

Implementing IPv6 in Comcast's Voice Network

Carl Klatsky
Comcast NE&TO Product Engineering

Background

- During 2013 & 2014, Comcast defined an updated voice architecture to utilize IPv6 in its IMS-based voice network
- Comcast's voice network currently uses Internet-routable IPv4 addresses on the voice interface of its Customer Premise Equipment (CPE)
- Driven by the exhaustion of Internet routable IPv4 addresses, the near term goal of this effort is to update the voice network to allow IPv6 address usage for voice traffic to & from the CPE
- This presentation will review the successes & challenges encountered to date, focusing primarily on the access portion of the voice network

Network Overview – Core

- SIP-based IMS network deployed and operational since late 2010
- Currently 5 IMS core sites serving Comcast’s national voice footprint
- IPv4 used for signaling and bearer traffic
- IMS CSCF (SIP Proxy) IPv4 / IPv6 Dual-Stack support not available until mid-to-late 2015
- Multiple open source based Application Layer Gateways (ALG) will be deployed to support IPv6 to IPv4 interworking for all signaling and bearer traffic destined to / coming from the IPv6 enabled CPE

Network Overview – Access

- CPE are comprised of digital voice adapters embedded within cable modems, residing at the customer premises
- CPE are SIP-based and are being updated to support the IPv6 protocol on their voice interface. IPv6 is already in use in Comcast’s network for managing the cable modem portion of the CPE
- Cable Modem Termination Systems (CMTS) serve as termination point for the link layer DOCSIS protocol and the entry point onto Comcast’s layer 3 backbone network

Network Overview – Access

- CMTSes, aggregation routers, and backbone routers already support IPv6 routing & transport, and the IPv6 enabled CPE will leverage this support with limited configuration changes
- Access side provisioning systems (DHCP, DNS, NTP, TFTP) also already support IPv6, and will be modified to support IPv6 enabled voice CPE
- Packet Cable Multimedia (PCMM) system, used for QoS policy enforcement, also supports IPv6 and will only require a policy update to support voice IPv6

IPv6 Voice Network – Single Stack vs. Dual Stack

- Initial design anticipated IPv4 / IPv6 dual-stack support on the IMS CSCFs for signaling interworking, with a media translation GW for bearer traffic interworking when needed
- CPE would run in dual stack mode, only invoking the media translation GW when needed
- Vendor support for those elements was delayed. Application Layer GW approach was adopted to meet project timelines
- With inclusion of an ALG to mask all IPv6 traffic from the IPv4 network, it became an easier decision to place the CPE in IPv6 single stack mode as all of the CPE traffic would need to traverse the ALG anyway
- Placing the CPE in IPv6 single stack mode leads to an easier CPE implementation allowing a more focused development to ensue

IPv6 ALG Diagram

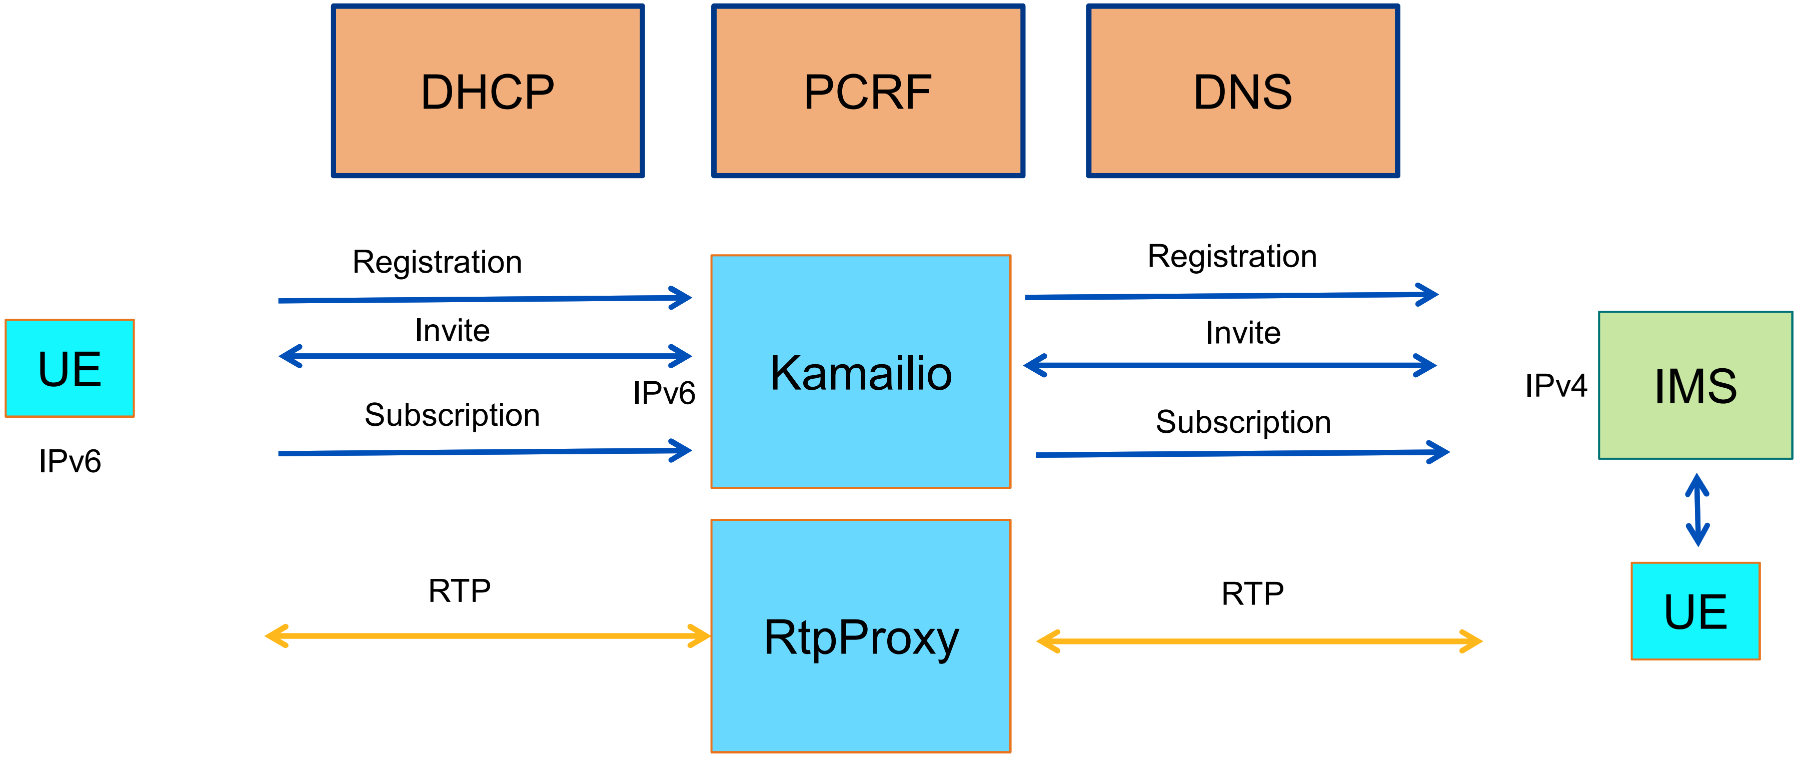

CPE Protocol Support

- Operating the CPE in single stack IPv6 mode required the IPv6 equivalent support for these protocols
  - DHCPv6 – Comcast’s address assignment plan in IPv6 uses DHCPv6, not SLAAC
  - DNS IPv6 support – CPE uses IPv6 transport to issue DNS queries and is requesting AAAA record resolution
  - TFTP IPv6 support – Existing configuration file transfer is through TFTP, re-using the current TFTP server supporting IPv6 transport
  - SIP IPv6 – Signaling protocol with IPv6 support
  - RTP & RTCP-XR – Media & Media Control protocols with IPv6 support
  - SNMPv6 – Comcast’s existing toolset has been converted for IPv6 operation, and will be re-used for managing Voice IPv6 devices

Success & Challenge Issues

- The first challenge to arise was related to Comcast’s provisioning systems
- To support a Production trial, we needed a means to enable IPv6 on the trial participant’s CPE, but not other CPE of that model type
- Comcast’s current provisioning system implements a “one size fits all” approach without customization on an individual subscriber basis
- An solution had to be designed to enable IPv6 support on a per subscriber basis
- To overcome this, a configuration parameter was defined, supported in non-volatile memory to override the default setting of IPv4, and instead support single stack IPv6 configuration
- This configuration parameter controls whether or not the CPE honors the value of the IP-PREF DHCP option returned from the DHCP server
- This option is used to signal which IP mode the DHCP server is directing the CPE to select

Success & Challenge Issues

- The next challenge to arise was related to CPE FQDN usage
- During provisioning, the CPE are assigned FQDN
- The CPE uses the domain portion of its FQDN to form a DNS SRV query to learn a prioritized list of SIP proxies (CSCFs) for registration
- In DHCP for IPv4, the host portion and domain portion of the FQDN are sent as separate DHCP options, with the domain being sent in option 15
- In DHCPv6, the host portion and domain portion are sent in one option, option 39
- The CPE code had to be updated to extract the domain portion only from option 39. The code change removed the first portion of the FQDN up to & including the first dot '.' and treat the remainder of the FQDN as the domain portion for the DNS SRV query

Success & Challenge Issues

- The next challenge was to populate the FQDNs in the response to accommodate both IPv4 and IPv6 queries
- The DNS server supports DNS SRV queries for both IPv4 and IPv6. If the same domain name is supplied in the DNS SRV query, the DNS SRV response will return the same list of SIP proxy FQDNs
- The CPE must then query the DNS server to resolve the returned FQDNs
- In order to retain the same list of CSCF FQDNs for IPv4 and IPv6 CPE, each CSCF FQDN in the list had to have a corresponding A and AAAA record in the DNS server, with each pointing to a different signaling element
- The A record for a given CSCF FQDN lists the IPv4 address of an existing signaling element, while the AAAA record for the same FQDN lists the IPv6 address of the new ALG element

Success & Challenge Issues

```
> set querytype=SRV
> _sip._udp.west.voice.cc38.ula.comcast.net
Server:  ns01.cable.comcast.com
Address:  69.252.80.80

Non-authoritative answer:
_sip._udp.west.voice.cc38.ula.comcast.net       SRV service location:
          priority       = 10
          weight         = 10
          port           = 5060
          svr hostname   = epcscf.voice.cc38.ula.comcast.net
_sip._udp.west.voice.cc38.ula.comcast.net       SRV service location:
          priority       = 20
          weight         = 10
          port           = 5060
          svr hostname   = hpcscf.voice.cc38.ula.comcast.net
```

Success & Challenge Issues

```
> set querytype=A
> hpcscf.voice.cc38.ula.comcast.net
Server:  ns01.cable.comcast.com
Address:  69.252.80.80

Non-authoritative answer:
Name:    hpcscf.voice.cc38.ula.comcast.net
Address:  10.253.156.134

> set querytype=AAAA
> hpcscf.voice.cc38.ula.comcast.net
Server:  ns01.cable.comcast.com
Address:  69.252.80.80

Non-authoritative answer:
Name:    hpcscf.voice.cc38.ula.comcast.net
Address:  2001:558:ff4f:14:10:253:178:4
```

Success & Challenge Issues

- The next challenge related to the routing of IPv6 of media packets
- Following IPv6 address acquisition, DAD, and Neighbor Discovery, the CPE listens for IPv6 Router Advertisements (RA)
- The RA includes a bit flag (O-bit) to indicate if the IPv6 prefix being advertised is “on-link” or “off-link, where “on-link” means directly accessible via layer 2 forwarding, and “off-link” means accessible via a router
- On Comcast’s network, all prefixes are advertised as “off-link” so that all IPv6 traffic leaving a CPE should be directed to the default router
- In this particular test, the next hop node for media interworking was on the same IPv6 prefix as the test CPE
- Vendor provided firmware was attempting to route signaling & media traffic on the same prefix directly via layer 2, which was not supported on Comcast’s network.
- Direct Neighbor Solicit messages from the CPE to learn the layer 2 address of the media interworking node were not receiving responses, hence the CPE was not transmitting the traffic out to the network
- Had to work with vendor to confirm that they follow the correct procedures for routing IPv6 packets (RFC 4861)

Summary

- Chose single stack IPv6 support if at all possible
- Be aware of the changes needed to provisioning endpoints (DHCP & DNS)
- Understand the difference in IPv6 packet forwarding and its implications on your network
- Understand any changes needed in QoS to support IPv6 signaling & media traffic