IPv6 Operations	J.B. Brzozowski
Internet-Draft	C.G. Griffiths
Intended status: Informational	Comcast
Expires: April 04, 2012	October 02, 2011

Comcast IPv6 Trial/Deployment Experiences
draft-jjmb-v6ops-comcast-ipv6-experiences-02

Abstract

This document outlines the various technologies Comcast has trialed as part of the company's ongoing IPv6 initiatives. The focus here are the technologies and experiences specific to enabling IPv6 for subscriber services like high speed data or Internet. Comcast has learned a great deal about various technologies that we feel are important to share with the community.

Status of this Memo

Copyright Notice

Table of Contents

1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Introduction

Beginning in early 2010 Comcast announced plans to leverage the work the company has been doing related to IPv6 to conduct a number of IPv6 technology trials. These trials were specifically aimed at enabling IPv6 for subscriber services. The purpose of this document is to outline the technologies that have been trialed thus far along with experiences and observations that adopters of the same may find valuable in their own planning and deployment processes.
Further, there may be some additional feedback that the various groups within the IETF may wish to take into account as part of ongoing standards efforts.

3. 6to4

During production deployment planning the widespread use of 6to4 [RFC3068] to access content and services over IPv6 was assessed. In some scenarios 6to4 usage increased several hundred times. At the time Comcast had not deployed its own 6to4 relay infrastructure as such open relays being operated by independent third parties were by default used to facilitate 6to4-based communications. The deployment and default use of open 6to4 relays appears to be a key variable behind the sub-optimal performance associated with the use of 6to4. An important thing to note is that some home gateway vendors have turned on 6to4 by default, and in some of these implementations, they have not presented a user interface a user interface to disable it. For operators that have not deployed IPv6 or have IPv6 incapable infrastructures should note that the use of 6to4 is likely occurring today across their infrastructure. Many operating systems and home networking devices continue to support 6to4 and in some cases have 6to4 and other transition technologies enabled by default.
As a community there appears to be some consensus that long term the use of 6to4 is not desirable, however, in the near term it is clear that 6to4 will be used in specific scenarios. The expectation and goal is to see 6to4 usage diminish over time until use of the same is displaced by an alternate technique to access content and services over IPv6. While the debate continues over how and when to deprecate 6to4, it is clear that 6to4 should not be recommended as a primary mechanism to access content and services over IPv6.
The following documents outline the recommendations surrounding the use and status of 6to4 from a standards point of view:

1. [draft-ietf-v6ops-6to4-advisory]

2. [draft-ietf-v6ops-6to4-to-historic]

Comcast deployed a series of five (5) 6to4 relays in a geographically dispersed configuration across our network. The purpose of these relays was to reduce the latency typically associated with 6to4 usage. During our analysis, the use of off network, open 6to4 relays was determined to yield nearly unusable conditions depending on the geographic location of the end user relative to the open 6to4 relay. By deploying on-network 6to4 relays, latency in most cases was reduced by over 50%, which instantly yielded considerable improvements from an end user point of view. The simplistic design and deployment of these relays enabled us to rapidly put them in network, and in some cases create a better experience for some of our users who had 6to4 enabled.
Through the use of commodity x86 based servers that run a standard Linux Operating System, we reduced deployment and operating costs, while still maintaining a fault tolerant design. Each 6to4 relay was dual stacked, and with a simple kernel module, we enabled the 6to4 configuration. Some 6to4 specific configurations were required to

ensure compatibility across a wide range of end points. The logic to anycast the 6to4 records was handled by the network infrastructure providing connectivity to the 6to4 relays, and health checking enabled us to automatically remove the route for any relay from the routing table in case of failure.

```
                     Router Announces
<---------.            IPv4 Anycast            .--------->
Redundant |            192.88.99.1             | Redundant
Network   |           +------------+           | Network
Path      |           |  Network   |           | Path
          +-----------+   Router   +-----------+
                      +------------+
                      IPv4 |  | IPv6
                 Interface |  | Interface
                        +--+--+--+
                        | Linux  |
                        |  6to4  |
                        | Relay  |
                        +--------+
```

4. 6RD

6RD [draft-townsley-ipv6-6rd] is another transition technology similar to 6to4 that Comcast has deployed as part of technology trials. While 6RD yields some improvements over 6to4, 6RD is ultimately a tunneling technology. As such, it is subject to the challenges faced by other tunneling technologies.

As advertised, 6RD frees adopters from some restrictions typically associated with 6to4. The use of anycast addressing (IPv4 and IPv6) is no longer required and the infrastructure, like 6to4, is straightforward to deploy. However, at the time of deployment it was observed that a limited number of border relay (BR) implementations were available. This appears to be an evolving area with more implementations becoming available. Similarly it was observed that there we few if any customer edge (CE) implementations available to support a trial of the technology. As such engineering implementations were leveraged to evaluate 6RD. Further, there were no implementations available that supported the 6RD DHCPv4 options [draft-ietf-softwire-ipv6-6rd]. Because of this, every 6RD CE used for trial was manually configured with the necessary information required to enable 6RD. In order to support a wide scale production deployment leveraging 6RD an operator would have to ensure their DHCP infrastructure supports the required 6RD DHCPv4 options along with targeted 6RD CE devices.

Trial configurations included two (2) 6RD BRs, which were intentionally deployed in geographically dispersed configuration. An anycast design was used to enable 6RD with a well known IPv4 anycast address and FQDN for the 6RD BR. The use of anycast eased manual configuration and deployment. Additionally, an IPv6 /32 was used to support the 6RD trials permitting subscriber devices were able to yield a usable IPv6 /64 on the LAN side of the 6RD CE.

The quantity and location of the 6RD BRs is a key variable when planning the deployment of 6RD. Comcast specifically deployed a limited quantity of BRs resulting in some end users being "closer" to the BRs than others. Proximity to the 6RD BRs is an important factor that impacts the end user experience. While 6RD yields some improvements over 6to4, 6RD is ultimately a tunneling technology as such use of the same is subject to the challenges faced by other tunneling technologies.

Placement and quantity of 6RD BRs is also a significant variable to consider when assessing impacts to performance and IPv6 geo-location. A centralized approach to deploying 6RD BRs will yield undesirable impacts to IPv6 geo-location in that end users leveraging a particular 6RD BR that is geographically distant from their true location will not accurately represent the origin of the end user request. Conversely, deploying 6RD BRs that are near to end users may require a substantial quantity of 6RD BRs depending on the operator network.

The following provides an overview of the Comcast 6RD trial network design:

```
      +--------------------------------------------------------------+
      |                           Internet                           |
      +-------+-------------------------------------------+----------+
       IPv6   |                                           |    IPv6
      +-------+--------+                          +-------+--------+
      |                |                          |                |
      |    Router      |                          |    Router      |
      |                |                          |                |
      +---+--------+---+                          +---+---------+--+
          |        |                                  |         |
      IPv4|        | IPv6                       IPv4  |         | IPv6
          |        |                                  |         |
      +---+--------+---+                          +---+---------+--+
      |                |                          |                |
      |   6rd-br-01    |                          |   6rd-br-02    |
      |                |     6rd.comcast.net      |                |
      +----------------+     _..-'      -.._      +----------------+
                         _.-'               `--._
                    _.--'                        ``-.__
                _.-'  IPv6                   IPv6      `-.._
         ,-..--'   encapsulated           encapsulated      ``-..
        (   )        in IPv4                in IPv4          (   )
         `-'                                                  `-'
       6rd CE                                                6rd CE
          |                                                    |
          |  IPv6                                              | IPv6
          |                                                    |
```

5. Native Dual Stack

Native dual stack is central to Comcast's IPv6 program for trial and production deployment. Native dual stack is the model where IPv4 services remain as-is with native IPv6 support introduced in parallel or simultaneously. Many of the details surrounding how this is achieved are documented as part of the CableLabs Data Over Cable Service Interface Specification (DOCSIS) 3.0 [DOCSIS3.0]. However, relevant trial and deployment specific information that is of interest to the IETF community will be documented.

Native dual stack trials depend on the upgrade and enablement of Cable Modem Termination Systems [CMTS] to support IPv6. A CMTS is a device that end users in a cable network connect directly to using their cable modem [CM]. As with IPv4, native support for IPv6 is critical for the delivery of services to end users in a DOCSIS network. Anything less could yield an undesirable end user experience or instability in the operator network that could adversely impact larger populations of users.

Given the CMTS requirements, native dual stack trials have initially been limited to specific areas of the network. Further, where CMTS platforms have been upgraded and enabled to support IPv6 end users have been incrementally enabled with support for IPv6. Again this is to ensure a controlled introduction with a specific focus on maintaining stability. Initially, a limited combination of cable modem and IGD devices are being used to support trial activities. Over time diversity for both cable modem and IGDs are expected to grow. To date a number of cable modems support the ability to enable native dual stack connectivity to CPEs devices behind them. A subset of pre-DOCSIS 3.0 and all DOCSIS 3.0 devices support this capability. The population of DOCSIS devices that support these capabilities varies from operator to operator.

Trial enablement requires the stateful provisioning of an IGD using stateful DHCPv6 [RFC3315] for the IGD WAN interface and delegated prefixes [RFC3633] for LAN side connectivity. Similarly, trial supported direct attachment of IPv6 capable CPE devices to the CM. In this configuration the CPE is provisioned with one or more IPv6 addresses via stateful DHCPv6 [RFC3315] in similar fashion to the IGD WAN interface. The quantity of devices supporting a native dual stack mode of operation is growing. While some devices are upgradable to support native dual stack many devices deployed today are not upgradable to support this functionality. Early implementations of devices or devices that are upgradable to support native IPv6 were found to only require and/or support the use of an IPv6 /64 for LAN side connectivity. This has been an acceptable mode of operation, however, over time IGDs will be required to support more advanced functionality including the ability to support multiple, routed IPv6 LANs. While support for a single IPv6 /64 is in place today support for shorter IPv6 prefixes is also supported. It is important for operators to ensure they design and plan support across their infrastructures for delegated prefixes that are shorter than /64.

```
    +--------------------------------------------------------------+
    |                           Internet                          |
    +-------+---------------------+---------------------+---------+
                                  | Native Dual Stack
                                  |
                          +-------+--------+
                          |                |
                          |    Router      |
                          |                |
                          +---+---+----+---+
                                  |
                                  | Native Dual Stack
                                  |
                      +-----------+------------+
                      |   Cable Modem          |
                      |   Termination System   |
        '''''''''''''''   (CMTS)               '''''''''''''
        |             +------------------------+           |
        |                                                  |
        |                                                  |
     +--+---+                                           +--+---+
     |      | Cable                                     |      | Cable
     +--+---+ Modem                                     +--+---+ Modem
     +--+---+                                              |
     |      | IGD     <----Native IPv4 and IPv6---->       |
     +--+---+                                              |
       ,+.                                                ,+.
      (   )   Computer                                   (   ) Computer
       `-'    (CPE)                                       `-'  (CPE)
```

6. Dual Stack Lite

Part of Comcast's trial plans includes the trialing of Dual Stack Lite. At this time trial planning for the same is underway. While Comcast plans on trialing Dual Stack Lite there are no plans at this time to deploy Dual Stack Lite beyond a limited technology trial.

7. Content and Services

During early phases of our trials Comcast leveraged reverse proxies to expedite the availability of content natively over IPv6. Open source technology running on Linux based servers was used to enable the reverse proxies. To ensure that the origin content, which is IPv4 only, is available natively over IPv6 the proxy servers required native dual stack connectivity. This model allowed us to ensure that Internet facing access to Comcast content occurred natively over IPv6.

As third party CDNs introduce production quality support for IPv6 we plan to move away from the use of proxy servers and fully towards native dual stack for Comcast content and services. Native dual stack content is but the first step to ensure the same can be IPv6 only at some point in the future. Observations from Comcast's participation in World IPv6 day suggest it is premature to rely on IPv6-only content at this time

Further as part of our trials Comcast has also recently enabled IPv6 Message Transfer Agents (MTA), in a limited fashion, to allow a subset of Comcast trial users to send electronic mail using SMTP over IPv6.. Due to the limited availability of spam mitigation for IPv6 Comcast trials does not include the receipt of electronic mail over IPv6. In order to enable the receipt of electronic mail over IPv6 spam mitigation must be in place.

8. Backoffice

We made the decision early on in our design discussions to move all systems to a dual-stack design since we felt that this was the best way to transition to IPv6. The re-architect of many core systems like DNS, DHCP, OSS/BSS, and Billing systems took many years to plan and complete and this approach has paid off and allowed us to rapidly move towards support for dual-stack at the edge of our network, including support for our customers devices.

9. World IPv6 Day

During World IPv6 day, Comcast observed a significant increase in native IPv6 traffic once content providers enabled AAAA records for their websites. The resulting traffic has continued to increase even after World IPv6 when about 50% of the websites that participated in World IPv6 Day left their AAAA records enabled after the day. We view this as a positive sign for continuing to drive more IPv6 traffic.

10. Conclusion

To date Comcast trial activities have yielded important, useful information about the various technologies that are available to facilitate the transition to IPv6. Observations and experience to date confirms that native dual stack is the preferred approach to transition to IPv6, where possible. While the various tunneling technologies are indeed straightforward to deploy there are a number of variables that must be considered when planning to deploy the same.

Support for native dual stack continues to evolve across various broadband technologies and within consumer electronics. As evidenced by World IPv6 Day many of the world's largest content providers are also making progress with their IPv6 capabilities.

11. IANA Considerations

This document makes no request of IANA.
Note to RFC Editor: this section may be removed on publication as an RFC.

12. Security Considerations

There are no security considerations at this time.

13. Acknowledgements

Thanks to the Comcast team supporting the various trial and production deployment activities:

- Jonathan Boyer
- Chris Griffiths
- Tom Klieber
- Yiu Lee
- Jason Livingood
- Anthony Veiga
- Joel Warburton
- Richard Woundy

14. References

[RFC2119]	Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

Appendix A. Document Change Log

[RFC Editor: This section is to be removed before publication]
-02: Grammatical items and re-wording of some sections. We have also added a new World IPv6 Day section.
-01: Added C. Griffiths as co-author. Currently working on ascii art and several new sections.
-00: First version published.

Authors' Addresses

John Jason Brzozowski Brzozowski Comcast Cable Communications One Comcast Center 1701 John F. Kennedy Boulevard

Philadelphia, PA 19103 US EMail: john_brzozowski@cable.comcast.com
URI: http://www.comcast.com

Chris Griffiths Griffiths Comcast Cable Communications One Comcast
Center 1701 John F. Kennedy Boulevard Philadelphia, PA 19103 US
EMail: chris_griffiths@cable.comcast.com URI: http://www.comcast.com