

IPv6 CE Router Interoperability Whitepaper

November 2011



**InterOperability Laboratory
University of New Hampshire**

**121 Technology Drive, Suite 2
Durham NH, 03824
Phone: +1-603-862-3941
Web: www.iol.unh.edu**

Table of Contents

Table of Contents 2
Overview 3
Test Methodologies 4
Observable Results 5
Conclusion 9
About the UNH-IOL 10
Terminology 11
References 12
Author: 13

*University of New Hampshire
InterOperability Laboratory*

Overview

The University of New Hampshire InterOperability Laboratory (UNH-IOL) hosted its third IPv6 Customer Edge (CE) Router Interoperability Test Event the week of November 7-11, 2011. The event brought together users and suppliers of CE Router equipment in order to gain perspective on the current status of interoperability against the Internet Engineering Task Force's (IETF) [Basic Requirements for IPv6 Customer Edge Routers](#) (document draft-ietf-v6ops-6204bis-02). During the IPv6 CE Router event the UNH-IOL used publically routable IPv6 addresses, allowing participants to connect to the global IPv6 Internet. The eight participating vendor companies tested a total of 12 distinct CE Router implementations throughout the week. Participants included Actiontec, Broadcom, Cisco, D-Link, Lantiq, Motorola Mobility, and Time Warner.

An IPv6 CE Router is a customer edge router intended for use in a home or small office environment. The router connects the end-user network to a service provider network and forwards packets not explicitly addressed to it. Implementing IPv6 on CE Routers is necessary in order to sustain growth and usability of the Internet.

While IPv6 is the solution for keeping current customers connected and adding new customers to the network as the supply of remaining allocated IPv4 addresses reaches exhaustion, it is not widely deployed in broadband networks at this time. With that said, operators are deploying native IPv6 on operational networks when possible. One effect of the IPv6 transition is that some operators are choosing to implement native IPv6 without native IPv4 on new network deployments. In cases where an access network is not dual-stack (both IPv4 and IPv6), operators are looking to transition mechanisms, such as 6RD and DS-Lite, to help connect their subscribers to both IPv4 and IPv6 networks. As more networks become IPv6 operational, these mechanisms will continue to be refined and help ease end-users through the transition.


Going forward, the UNH-IOL will continue to host IPv6 CE Router Interoperability Test Events in order to help operators to find solutions to interoperability challenges that may be experienced in the transition, thereby cost-effectively speeding IPv6 broadband deployments.

Test Methodologies

The tests executed during this event were performed in order to verify that an IPv6 CE Router implements IPv6 routing; that is, that the IPv6 CE Router properly looks up IPv6 addresses in the routing table and sends them to the correct interface, as well as acts like a proper IPv6 node as defined by the IETF. Tests were also designed to verify the WAN side configuration, specifically that a node supports protocols necessary to enable IPv6 deployment on multiple network access architectures. LAN side configuration testing was limited due to time constraints.

The following common topology was used for all test cases.

Common Topology



- The WAN interface was a DOCSIS, DSL or an Ethernet network for all CE devices.
- Cisco Network Registrar (CNR) acted as both the DHCP-Server1 and DNS-Server.
- TAR-Host1, REF-Host1 and REF-Host2 consisted of a mix of popular Operating Systems including Microsoft Windows 7, Linux and Apple Mac OS X Lion.

Observable Results

IETF Standard: Basic Requirements for IPv6 Customer Edge Routers

Detailed test cases were developed from the requirements designated as “MUST” in the IETF [Basic Requirements for IPv6 Customer Edge Routers](#).

Test Event Evolution

This event was critical for providing measured progress of CE Router implementations over the past year. Several essential requirements were successfully demonstrated during this event including delegating prefixes from a prefix pool larger than /64, preventing forwarding loops, and duplicate address detection. Although implementations have shown vast improvements compared previous IPv6 CE events, issues remain and are documented in this paper as follows:

- Experience with the M and O Flags in the Router Advertisement
- Routing Delegated Prefixes
- No Route Available
- No Address Available
- Support for 6RD

Importantly, many of the issues identified were resolved during the week and implementations were then retested in order to pass the entire test plan by the end of the event.

Experience with the M and O Flags in Router Advertisements

A recent topic of IETF discussions has been the possibility of using the M and O flags in Router Advertisements to control the CE Router’s DHCP client. The table below documents how the M and O flags affected CE Routers starting DHCP Clients and DHCP Prefix Delegation.

	M=1/O=1	M=1/O=0	M=0/O=1	M=0/O=0
DHCP IA_NA	10	9	7	6
No DHCP IA_NA	0	1	3	4
DHCP IA_PD	10	8	0	9
No DHCP IA_PD	0	2	10	1

As the above table shows, implementations vary their behavior based on the M and O flags’ configuration. Some implementations treated the O flag as an indicator for turning on Prefix Delegation. However, this behavior is not conformant to RFC 6204, as it states Prefix Delegation should be done regardless of the M and O flags’ state. Other implementations ignored the flags altogether and always started DHCP IA_NA and IA_PD acquisition. This behavior is acceptable per the current RFC 6204. The varied and sometimes non-conformant behavior of CE Router devices observed indicates a need

*University of New Hampshire
InterOperability Laboratory*

for further clarification in the IETF about the relationship between the M and O flags and DHCP in CE Router devices. From these observations, it is clear that current requirements have led implementations to react differently depending on the M and O flags configured. This results in confusion and misconfiguration for administrators and operators as how DHCP clients are operating.

Routing Delegated Prefixes

DHCP Prefix Delegation is the mechanism intended for delegating a long-lived prefix from a delegated router to a requesting router. This allows Internet Server Providers to delegate prefixes to CE Routers to be distributed. During previous events, assigning prefixes to LANs from a delegated pool larger than /64 was deemed problematic for some implementations. However, at the most recent event, implementations demonstrated the ability to delegate a /64 prefix to the LAN networks from a /60 and /52 delegated pool. Another additional - but important - component is properly routing the assigned and unassigned prefixes. Incorrect routing of prefixes leads to forwarding loops, and excess traffic in the network.

In order to test the routing of delegated prefixes the following conditions were setup. A DHCP Server delegated a prefix of /60 to the CE Router through DHCP Prefix Delegation. The CE Router then assigns a /64 prefix from the /60 delegated prefix pool to the LAN interface, advertising the prefix in a Router Advertisement. An IPv6 Host, TAR-Host1 in the Common Topology, was attached to the LAN network and transmitted data destined to an IPv6 address in the /60 delegated prefix pool but unassigned. The node identifier (suffix) of the address is unimportant as the prefix is unassigned.

Several behaviors were witnessed. Some CE Router implementations routed the packets destined for the non-delegated prefixes to TAR-Router1, the default route rather than the null destination. TAR-Router1 had the CE Router as the next hop for the delegated prefix and routed the packet back to the CE Router. This behavior caused a forwarding loop, and excess network traffic.

A previously unobserved behavior seen under the described conditions was the transmission of ICMPv6 Redirect messages in response to the data packets destined for what was assumed to be a delegated but unassigned prefix. These types of messages indicate that the CE Router installed a route for the entire delegated prefix pool to the LAN interface, instead of only the /64 prefix assigned to it. This behavior is problematic because the IPv6 host will try to route the traffic on-link, again causing excess network traffic and inconsistent behavior for applications and ultimately the user. Most implementations addressed these behaviors and corrected them by the end of the event.

No Route Available

When no route is available to a CE Router on the WAN interface, the router must transmit a Router Advertisement with a router lifetime of zero to the LAN interfaces which indicates to hosts that this device may not be used as a default router. Three implementations at the event transmitted Router Advertisement with a default lifetime value greater than 0 instead of transmitting a router lifetime of zero. This behavior would improperly tell the IPv6 hosts on the LAN Network that there is a route available when in fact there is not, and could contribute to improper routing in the face of multi-homed networks.

One implementation transmitted a Router Advertisement on the LAN with a default router lifetime (a value greater than 0) prior to receiving a Router Advertisement on the WAN interface. Once a Router Advertisement with a default router lifetime of zero was received on the WAN interface, the CE Router transmitted a Router Advertisement with a router lifetime of zero on the LAN interface. While this is not prohibited in an IETF standard, this could cause delays to an IPv6 host if the second Router Advertisement, with a router lifetime of zero, is lost.

No Addresses Available

An important scenario was unable to be tested during the event due to DHCP servers and CE Routers support for handling a no address available situation. This scenario occurs when a CE Router has been delegated a valid prefix via DHCP, but no DHCP client unicast addresses are available on the server. This situation arises on networks wishing to use stateless global addresses, where stateful DHCP addresses are not required. DHCP prefix delegation however, will still be desired. The detailed explanation below documents the original RFCs and updated RFC describing the behavior each would exhibit in an implementation.

A CE Router will request both an IA_NA and IA_PD in the DHCP Solicit but will receive a DHCP Advertise with a status code option of NoAddrsAvail, indicating no addresses are available. A CE Router that received the NoAddrsAvail option in the DHCP Advertise will not process the remaining options.

According to RFC 3315 a DHCP server must do the following: “If the server will not assign any addresses to any IAs in a subsequent Request from the client, the server MUST send an Advertise message to the client that includes only a Status Code option with code NoAddrsAvail and a status message for the user, a Server Identifier option with the server's DUID, and a Client Identifier option with the client's DUID.”

Also, a client must react according to the following text from RFC 3315: “The client MUST ignore any Advertise message that includes a Status Code option containing the value NoAddrsAvail, with the exception that the client MAY display the associated status message to the user.”

*University of New Hampshire
InterOperability Laboratory*

The IETF issued errata in August 2010 that changes the behavior described in RFC 3315. The RFC has been updated to state the following for a DHCP server: “If the server will not assign any addresses to an IA in a subsequent Request from the client, the server MUST send an Advertise message to the client that includes the IA containing a Status Code option with status code NoAddrsAvail and a status message for the user, a Server Identifier option with the server's DUID, and a Client Identifier option with the client's DUID. The server SHOULD include other stateful IA options (like IA_PD) and other configuration options in the Advertise message.”

Also the DHCP client has been updated with the following behavior: “The client MUST ignore any Advertise message that includes a Status Code option containing the value NoAddrsAvail, with the exception that the client MAY display the associated status message to the user.”

Thus for future events, the UNH-IOL will use an updated DHCP Server that will support sending the status code option value NoAddrsAvail in the IA. One implementation worked around this problem. When a Status Code option with a value of NoAddrsAvail was sent outside the IA, the implementation sent a DHCP Solicit with only an IA_PD to get the Prefix.

Support for 6RD

During the event eight of twelve implementations showed support for IPv6 Rapid Deployment (6RD). 6RD is a mechanism that allows a network operator to roll IPv6 out on an existing IPv4 network. All of the implementations were capable of using the 6RD DHCPv4 Option for acquiring 6RD Parameters. All but one of the implementations had to be configured to use 6RD; the other implementation started to use 6RD when it received a 6RD DHCPv4 Option. It should be noted that one implementation preferred routing IPv6 traffic to an established 6RD tunnel instead of a native IPv6 interface if both were available.

Conclusion

IPv6 is paramount for Internet access as IPv4 addressing runs out. It is vital that IPv6 be implemented in customer edge equipment in order for operators to have the ability to use IPv6 in deployments. The participants in the IPv6 CE Interoperability Test Event demonstrated that IPv6 is implemented and is becoming deployable in IPv6 CE Routers.

Furthermore, participants are making great strides forward in quickly resolving issues in CE Routers once discovered. It's important that implementers thoroughly test to ensure that limited issues arise in the field and do not hinder the deployment of IPv6. It is also understood that additional standards work is needed detailing the interaction between transition mechanisms and native IPv6.

Moving forward, the UNH-IOL is working in conjunction with the IPv6 Forum to create an IPv6 Ready CPE Logo program that will help both operators and equipment manufacturers understand what features they need to support in order to aid in the worldwide deployment of IPv6. After the basic requirements are complete, new areas such as transition mechanisms and routing protocols will need to be covered in future testing.

About the UNH-IOL

Founded in 1988, the UNH-IOL provides independent, broad-based interoperability and standards conformance testing for data, telecommunications and storage networking products and technologies. Combining extensive staff experience, standards bodies' participation and a 32,000+ square foot facility, the UNH-IOL helps companies efficiently and cost effectively deliver products to the market. For more information, visit <http://www.iol.unh.edu/>.

The UNH-IOL hosts multi-vendor group tests (often called “plugfests”) as often as four times a month. These group test events compliment over 20 year-round standards-based testing programs that are managed and operated by the UNH-IOL. Each of the testing groups, called “consortiums,” represents a collaboration of industry forums, service providers, test equipment vendors and otherwise competing companies who benefit each other by:

- Distributing the cost of testing
- Lowering R&D and QA expenses
- Reducing product time to market
- Obtaining trusted vendor-neutral verification

The laboratory maintains a strong reputation for independent, vendor-neutral testing with a focus on quality assurance. The confidential test reports the UNH-IOL provides to its members are recognized throughout the data communications industry as evidence of interoperability and conformance to technical standards.



The UNH-IOL collaborative testing model distributes the cost of performing trusted, independent testing and validation through an annual membership.

Terminology

6RD: IPv6 Rapid Deployment

CE LAN: CE Router LAN Interface

CE WAN: CE Router WAN Interface

DAD: Duplicate Address Detection

DHCP: Dynamic Host Configuration Protocol

DS-LITE: Dual-Stack Lite

ICMP: Internet Control Message Protocol. ICMP Echo Requests and Replies facilitate troubleshooting at Layer 3 for both IPv4 and IPv6. IPv6 has built extra features into ICMP.

IETF: Internet Engineering Task Force

LAN: Local Area Network

NS: Neighbor Solicitation

NA: Neighbor Advertisement

NCE: Neighbor Cache Entry

NUT: Node Under Test

PD: Prefix Delegation

RA: Router Advertisement

RS: Router Solicitation

WAN: Wide Area Network

References

The following documents are referenced during the test event:

- [RFC 2460] R. Hinden, S. Deering, Internet Protocol, Version 6 (IPv6) Specification, RFC 2460, December 1998.
- [RFC 3315] R. Droms, J. Bound, B. Volz, T. Lemon, C. Perkins, M. Carney, Dynamic Host Configuration Protocol for IPv6 (DHCPv6), RFC 3315, July 2003.
- [RFC 3633] O. Troan, R. Droms, IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6, RFC 3633, December 2003.
- [RFC 3484] R. Draves, Default Address Selection for Internet Protocol version 6 (IPv6), RFC 3484, February 2003.
- [RFC 3646] R. Droms, DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6), RFC 3646, December 2003.
- [RFC 3736] R. Droms, Stateless Dynamic Host Configuration Protocol (DHCP) Server for IPv6, RFC 3736, April 2004.
- [RFC 4191] R. Draves, D. Thaler, Default Router Preferences and More-Specific Routes, RFC 4191, November 2005.
- [RFC 4291] R. Hinden, S. Deering, IP Version 6 Addressing Architecture, RFC 4291, February 2006.
- [RFC 4294] J. Loughney, Ed. IPv6 Node Requirements, RFC 4294, April 2006.
- [RFC 4443] A. Conta, S. Deering, M. Gupta, Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification, RFC 4443, March 2006.
- [RFC 4861] T. Narten, E. Nordmark, W. Simpson, H. Soliman, Neighbor Discovery for IP Version 6 (IPv6), RFC 4861, September 2007.
- [RFC 4862] S. Thomson, T. Narten, T. Jinmei, IPv6 Stateless Address Autoconfiguration, RFC 4862, September 2007.
- [RFC 5942] H. Singh, W. Beebee, E. Nordmark, RFC 5942, July 2010.

- [RFC 6204] H. Singh, W. Beebee, C. Donley, B. Stark, O. Troan. Basic Requirements for IPv6 Customer Edge Routers, RFC 6204 April 2011.
- [RFC 6204bis-02] H. Singh, W. Beebee, C. Donley, B. Stark, O. Troan. Basic Requirements for IPv6 Customer Edge Routers, draft-ietf-v6ops-6204bis-02, October 2011.

Author:

Timothy Winters, Senior Manager
UNH-IOL
121 Technology Drive, Suite 2
Durham, NH 03824
USA
Email: twinters@iol.unh.edu
URL: <http://www.iol.unh.edu>

Biography for Timothy Winters:



Timothy Winters, Senior Manager, University of New Hampshire InterOperability Laboratory (UNH-IOL)

Timothy Winters is a senior manager at the University of New Hampshire InterOperability Laboratory (UNH-IOL). He works with companies from all over the world to develop broad-based, flexible testing strategies to cost effectively meet network interoperability requirements for the Internet Protocol version 6 (IPv6), Session Initiation Protocol (SIP), routing, and IP Multimedia Subsystem (IMS) communities.

Timothy is the United States Government IPv6 (USGv6) and IPv6 Ready Logo technical lead for the UNH-IOL. In this role, he oversees various aspects in testing of IPv6 technology, deals with various multi-vendor IPv6 testing scenarios and acts as a liaison between students and vendors during device testing and development. He was instrumental in the UNH-IOL's designation as an ISO-17025 USGv6 accredited laboratory.

Timothy has more than nine years of experience as a UNH-IOL leader and his involvement with cutting edge technologies has led to outstanding achievement. He is an expert in routing protocols: Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), and Protocol Independent Multicast (PIM), as well as a connectivity specialist and deployment advisor for Moonv6, the world's largest permanently deployed multi-vendor IPv6 network. Passionate about innovation, Timothy has taken the lead on developing SIPv6 and IMS IPv6 Ready Logos for the IPv6 Forum's IPv6 Ready Logo Program. In addition, Timothy assisted with the IPv6 description for SIPConnect 1.1 and has co-authored multiple white papers and test suites at the UNH-IOL.

His ongoing collaboration with industry forums including the North American IPv6 Task Force, IPv6 Forum, IPv6 Ready Logo Committee, Multiservice Switching Forum, IMS Forum and SIP Forum demonstrates his dedication and persistence in developing new standards, as well as assisting commercial services providers, network equipment vendors and government agencies cost effectively speed go-to-market time for products.

Timothy holds a Bachelor of Science in Computer Science from the University of New Hampshire, Durham.