RTC
The magazine of record for the embedded computing industry
September 2011 www.rtcmagazine.com

IPv6 GETS READY TO TAKE THE STAGE

Inside: World is moving to IPv6
by Thomas Volz

An RTC Group Publication

The World Is Moving to IPv6: Are You and Your Product Ready?

The transition to IPv6 is more than just a move to an expanded address space. It involves a system of routing and address formats that will make the Internet more versatile. Compatibility with IPv4 is not automatically assured, and a rigorous program of certification is needed to ensure interoperability.

by Thomas Volz, EBSnet

On February 3, 2010, in a ceremony broadcast live over the Internet from a Miami Hotel, the last blocks of available IPv4 addresses were allocated to the Regional Internet Registries. Experts anticipated these blocks would be used up by the Regional Registries within a few months. IPv4 has finally hit its address space limitation. The future growth of the Internet will come through the transition to the next generation IPv6 protocol.

One of the main design goals of IPv6 was to increase the limited size of the IPv4 address pool available to network connected nodes. To this end the IPv6 address size was increased from 32 to 128 bits, and along with new features and enhancements to the IP protocol itself, it is a more efficient and interoperable network architecture than IPv4. The IPv6 address structure is much more than just an extension of IPv4's classless inter-domain routing (CIDR) structure. It uses the increased size in a very flexible way that preserves many of the CIDR concepts, can encapsulate existing IPv4 addresses, and can even accommodate unique worldwide global addressability for any network node interface. It is not however intended to enable worldwide global addressing for every interface. In fact network nodes bound to a local subnet for a dedicated purpose may derive no benefit at all from IPv6. But when considering the need or the possibilities from a greater network visibility and interaction, IPv6 offers many improvements over IPv4.

FIGURE 1

IPv6 Ready Logos—Phase-1 Silver logo, Phase-2 Gold logo.

FIGURE 2

IPv6 Ready Logos for the Phase-2 Extended categories.

Since the address size and structure play such an important role in the overall IPv6 functionality, it is definitely worth a closer look. IPv6 nodes can and often do use multiple addresses simultaneously for a variety of purposes. In general, the 128-bit address is comprised of two major fields, the subnet prefix and the interface identifier. Addresses are interpreted from left to right starting with the subnet prefix. This prefix may be any length, with the interface identifier represented by the remaining number of the least significant bits. Within this architecture there are many reserved prefix patterns that serve to define some commonly known address structure formats. However, as long as there is no conflict with one of the predefined formats, an address need not imply any structure at all or may be interpreted privately within a local network. The general 128-bit IPv6 address format is:

Prefix (n bits)	Interface Identifier (128-n bits)

Among the reserved prefix patterns mentioned above there are three types that are essential for interoperability within the general IPv6 framework: multicast, link local and global.

IPv6 Multicast Address

In IPv6 the multicast address type is used extensively and is required. There is no broadcast address that is often used in IPv4. The multicast prefix is identified by a leading hexadecimal 0xff with the next eight bits defining the address scope and flags. Combined, these eight bits specify several potential sub formats embedded within the remaining 112 bits, the multicast group identification field. The scope bits generally relate information to potential router nodes while the flags indicate the format interpretation. Multicasts are used by a node to establish and validate its primary local address, validate any intended additional addresses, locate other directly addressable IPv6 nodes, and to discover resources such as routers, servers and their capabilities.

IPv6 Multicast Address Format

0xFF	Flags (4 bits)	Scope (4 bits)	Group Identifier (112 bits)

IPv6 Link Local Address

A link local IPv6 address prefix is indicated with the leading hexadecimal 0xfe8/10 bit pattern. The next 54 bits must be zero and the remaining 64 bits are the interface identifier, which must be unique within a local network link. The interface identifier is typically derived from its network hardware's MAC address but this is not required. These addresses are never forwarded by routers; they are intended to be used between peers in a link and between a node and its locally connected routers and servers. Each IPv6 node is required to have at least one link local address but may have as many as desired for other on link special purposes, providing each one is unique within the link.

IPv6 Link Local Address Format

0xFE8 (10 bits)	All Zeroes (54 bits)	Interface Identifier (64 bits)

IPv6 Global Unicast Address

An IIPv6 global address uses a variable length prefix that includes a global routing prefix and a subnet identifier. These address types are routed beyond the local link. Global address identification may be easier thought of as those addresses that are not one of the multicast, local, or a variety of other specialized and well defined types. At present there are two general categories of global addresses, those beginning with a three bit binary 000 and those that do not. Again, these addresses may not overlap the other predefined types after considering any additional bits beyond the first three. Essentially the 000 bit prefixed addresses may be considered unstructured by conventional specifications. The others comply with an assignment and management structure similar to IPv4. These managed types are defined below. The interface identifier must be 64 bits, but the routing prefix and subnet identifier are variable given that they must total to 64 bits.

IPv6 Global Unicast Address Format

Global Routing Prefix	Subnet Identifier	Interface Identifier (64 bits)

RFC 2460 (Phase-1, Phase-2) *Internet Protocol,Version 6 (IPv6) Specification*	**RFC 4862 (Phase-1, Phase-2)** *IPv6 Stateless Address Autoconfiguration*
RFC 4291 (Phase-1, Phase-2) *IP Version 6 Addressing Architecture*	**RFC 4443 (Phase-1, Phase-2)** *Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification*
RFC 4861 (Phase-1, Phase-2) *Neighbor Discovery for IP version 6 (IPv6)*	**RFC 1981 (Phase-2)** *Path MTU Discovery for IP version 6*

TABLE 1

The IPv6 Core Protocols testing covers the following RFCs.

UNH-IOL (US) *www.iol.unh.edu*	**TTA (Korea)** *http://test.tta.or.kr*
IRISA (France) *www.irisa.fr*	**BII (China)** *www.biigroup.com*
JATE (Japan) *www.jate.or.jp*	**CHT-TL (Taiwan)** *www.chttl.com.tw*

TABLE 2

Approved labs in the IPv6 Ready Logo program.

Focus on Interoperability

While IPv6 has been in development for some time, there is anxiety among the engineering community about just how smoothly the transition will unfold, especially with regard to an existing IPv4 node's ability to co-exist with its new IPv6 neighbors and how well new IPv6 applications will perform. The consortium of Internet steering organizations has had a technical roadmap for the transition in place for some years. To address these issues, the IPv6 Forum launched a formal Certification program in 2002, called The IPv6 Ready Logo program. The goals for the program, as stated officially on the Forum website, are three-fold:

- To verify protocol implementation and validate interoperability of IPv6 products
- To provide access to testing tools
- To set up a network of IPv6 Ready Logo testing labs around the world to assist with testing

The scope of the IPv6 Ready Logo testing program has steadily evolved over the years: from the Phase-1 Silver logo program, which covered Core IPv6 Protocols with only 170 tests, to the current Phase-2 Gold logo program, featuring a much broader coverage of some 450 tests for Core Protocols, as well as additional, extended protocol categories (Table 1). Tests are designed to verify an IPv6 implementation's conformance against the Internet Engineering Task Force's (IETF) RFCs. The IETF is the internationally recognized organization responsible for maintaining the Internet Protocol specifications and standards. They use a system of Request for Comment (RFC) technical documents to organize and catalog the multitude of IP protocol specifications. Literally thousands of RFCs combine to specify the Internet Protocol as we know it. Upon successful completion of the test programs, an IPv6 network stack implementation may use and display the Ready Logo as a testament to its capability to properly interact in an open network environment (Figure 1).

To illustrate the depth and the amount of the testing involved, the RFC 2460 test section alone includes 54 individual tests against different sections of this RFC, and each test may include dozens of test packet scenarios multiple test/response sequences. All told, a complete test run takes many hours using a test target with three or more known testing hosts. On top of the Core Protocols testing, the Phase-2 Gold logo program covers the extended IPv6 test categories, which include IPsec, IKEv2, MIPv6, NEMO, DHCPv6, SIP and SNMP-MIBs. The steering committee of the IPv6 Ready Logo program recently introduced a set of individual logos to mark the successful completion of the respective protocol section of the testing program (Figure 2).

To obtain either phase of the IPv6 Ready Logo certification, applicants need to complete the tests with the product passing 100% each of the appropriate conformance and interoperability test requirements.

The self-testing tools and interoperability test scenarios, which need to be executed against the product being tested, are made available for download on the IPv6 Ready Logo website www.ipv6ready.org. Alternatively, you can submit the product for testing to one of the approved labs in the IPv6 Ready Logo program, which typically requires joining their membership at hefty annual fees well into five figures (Table 2).

As the pressures related to impending depletion of the IPv4 address space have been publicized in the media, it's no wonder that the pace of new applications for the IPv6 Ready Logo has dramatically picked up in the past two years. Figure 3 illustrates the trend well.

The sheer amount of testing required to meet the requirements of the IPv6 Ready Logo program makes achieving certification a daunting task, especially for smaller development teams. Fortunately, the committee administering the program reserves a provision called *OEM Licensee clause*. The clause stipulates that OEM manufacturers can attain the IPv6 Ready Logo Certification without doing actual testing, so long as there is a *"one-to-one IPv6 stack transfer certified by the OEM Licensor."* Applicants taking advantage of the clause must specify the original product name and Logo ID on their application.

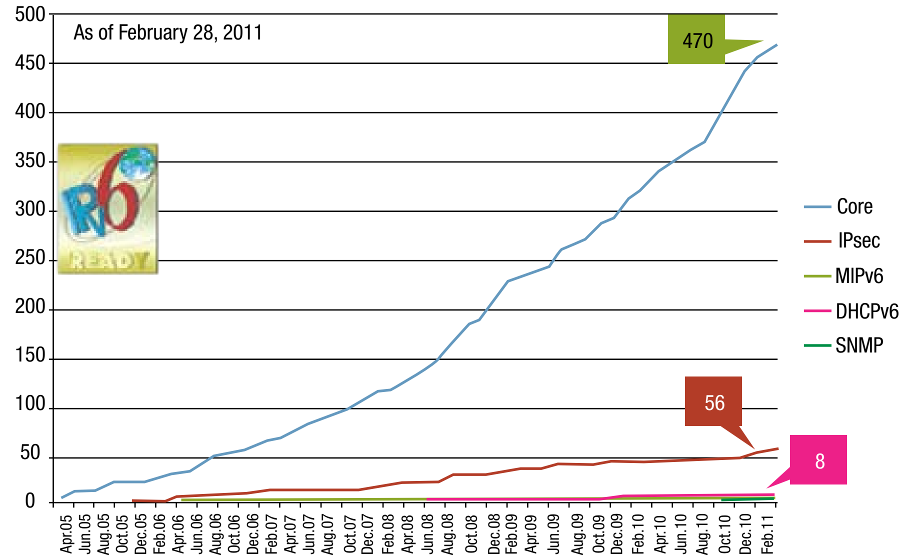

FIGURE 3

Total Applications approved for IPv6 Ready Logo certificate, including OEM licenses (source: IPv6 Forum).

Additional IPv6 Transition Challenges

Yet another challenge inherent in the transition to an entirely new IP protocol is that application developers and system integrators must devise plans on how newly deployed IPv6-enabled applications will continue supporting legacy IPv4 traffic. This is not a trivial task to address, as the IP protocol header got a dramatic makeover in the IPv6 version and is not directly compatible with IPv4. A number of techniques and approaches have been put forward by the industry to address the issue. While comprehensive coverage of various technology options providing a bridge between IPv4 and IPv6 network applications available or in the works would be beyond the scope of this article, one commonly used approach employs the use of a dual IPv4/IPv6 stack in the application. The approach, first described in the RFC 4213, is defined as a technique providing complete support for both Internet protocols—IPv4 and IPv6—in the hosts and routers. In this type of implementation network nodes will have the ability to send and receive both IPv4 and IPv6 packets. They can therefore directly interoperate with IPv4 nodes using IPv4 packets, and also directly interoperate with IPv6 nodes using IPv6 packets.

A variation of the concept is sometimes referred to as a dual-mode IPv4/IPv6 stack, which provides application developers with a facility to dynamically switch on and off either of the dual-stack layers, providing for operational efficiency. Yet another important transitional mechanism is a tunneling of IPv6 over IPv4. This technique provides for encapsulating IPv6 packets within IPv4, so that they can be carried across the existing IPv4 routing infrastructure.

EBSnet
Littleton, MA.
(978) 486-4000.
[www.ebsnetinc.com].