Global Knowledge®

Expert Reference Series of White Papers

IP Version 6 Address Types

IP Version 6 Address Types

Joe Rinehart, MBA, CCIE #14256, CCNP/DP/VP

Introduction

In 1998, the Internet Engineering Task Force (IETF) released RFC 2460, outlining the technical specifications of IPv6, which addressed the shortcomings of the aging IPv4 protocol. As with any evolution of technology, new elements exist in the protocol that may seem strange and unfamiliar. This certainly includes address representation, space, and so forth, but also includes a number of different **types** of addresses as well. A subset of these new addressing types has corresponding types in IPv4, but many will seem significantly different. The purpose of this white paper is to examine addressing classifications in detail and outline their functions within the context of the protocol.

To begin with, using the word types immediately implies the existence of numerous possible categories, and that impression is absolutely correct. Since it would be understandable to think of IPv4 addressing as somewhat monolithic, the existence of more than one addressing type can sound daunting to remember and understand. To put this concept into perspective, consider Figure 1 below, depicting various types of cellular telephones. The various "species" of devices displayed have differing colors, form factors, manufacturers, and capabilities, but in reality all are **phones**. In our discussion of IPv6 addressing types, just think of the different classifications have specific characteristics, but still having the same basic purpose: **communication**!

Figure 1: Cell Phone Types

commons.wikimedia.org

Three major addressing classification types exist in IPv6, with some subtypes as well. These consist of multicast, anycast, and unicast addresses, and each is technically worthy of separate consideration.

Multicast Addresses

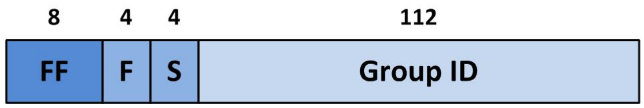

Figure 2: IPv6 Multicast Address Format

Purpose

In the classful addressing world of IPv4, multicast addressing was neatly confined into the Class D range, from 224.0.0.0 to 239.255.255.255. Within this group were specific allocations, such as the 239.0.0.0/8 **administratively scoped** range. Unicast represented one-to-one communication between hosts, and broadcast represented one-to-every communication. Multicast revolves around the concept not unlike being part of a club; communication and inclusion is a matter of whether or not you belong to that particular group or not. In IPv6, multicast takes on an even greater significance: first, because all broadcasts have been removed from protocol operations, and second, because a variety of protocol functions take place by means of multicast. Understandably, the address itself has several distinct features, as reflected in Figure 2.

Indicator or Prefix

Multicast addresses begin with eight bits, indicating the function involved, composed of all ones (1s), and yielding the hexadecimal characters of FF. The general designation for the entire multicast range is **FF00::/8**, which is readily recognizable with even casual observation.

Flags (F)

While the RFC specifies several values for the flag field, at present only one (the T bit) is in current use and identifies the lifetime of the address. 0 indicates a permanent address, while 1 denotes a temporary one.

Scope (S)

The 4-bit scope field describes to what degree the multicast address may be forwarded throughout the network, though not all values of this field have been defined. A partial list of these values is as follows:

- 0 Reserved
- 1 Interface-Local
- 2 Link-Local
- 3 Subnet-Local
- 4 Admin-Local
- E Global

Group-ID

The remaining 112 bits of the IPv6 multicast address identify the multicast group itself. In fact, some multicast addresses are already defined, and every network professional should be able to readily identify:

FF02::1	All Hosts
FF02::2	All Routers
FF02::5	OSPFv3 Routers
FF02::6	OSPFv3 Designated Routers
FF02::A	EIGRP Routers

Anycast Addresses

Purpose

The **anycast** addressing type deserves a special explanation, not only because it may appear unfamiliar, but also because it seems contrary to fundamental networking principles. To begin with, anycast employs the use of identical IP addressing to multiple devices within an internetwork, with nodes relying on routing protocols to determine which is physically closest. While not deployed extensively even in IPv4, this type of approach could be applied to shared resources such as DNS servers, for example. In IPv6, anycast is called out as a distinct addressing type, even though it is still not widely implemented.

No Distinct Format

As if the anycast concept itself is not confusing enough, the issue becomes more complicated by the fact that any IPv6 unicast address can be used as an anycast address. Unlike multicast addresses, there is no easily distinguishable format by which to recognize them. The most obvious distinguishing characteristic of an anycast address is whether it exists on more than one device within a routing domain.

GPS Analogy

As an engineer, I frankly have had a difficult time really grasping the concept of anycast addressing, though I have read many detailed explanations of the topic. The basic idea is to use identical addressing with shared resources, and have the natural process of IP routing select the closest resource, but that was as far as my understanding went. During the course of using a Global Positioning System program on my Android Phone, I inadvertently discovered a completely new method of understanding this previously ambiguous topic.

GPS devices exist seemingly everywhere in the modern technology landscape, and range from dedicated devices to running as a software application on a smart phone. When users have a particular need, they input a search string to locate the desired resource, such as a gas station for refueling their vehicle. Using GPS satellites, the device learns of its location and sends the search string to a database system that sifts through various data points to locate suitable establishments. The device then loads a list of the closest ones (see Figure 3), and may even recommend the facility physically closest to the user. If the search were the name of a particular establishment, all of the selected locations would have **identical names**, much like the duplicate addresses in anycast operations. Just as a GPS end-user would select the closest resource from that list, hosts perform the same operation. While the analogy is certainly not perfect, it does communicate the essentials in a much less confusing manner.

Figure 3: Anycast Addressing Functionality

Adapted from http://www.clker.com/

Unicast Addresses

Purpose

Unicast addresses have the same purpose in general, regardless of the protocol or mechanisms involved, namely, one-to-one communication between end devices or hosts. While this is certainly familiar with IPv4 addressing, it also takes place on the LAN at Layer 2 using unicast MAC addresses, as well as at Layer 3 in IPv6. Unlike IPv4, however, multiple sub-classifications exist, with specific formats that are easily recognizable.

Unique Local

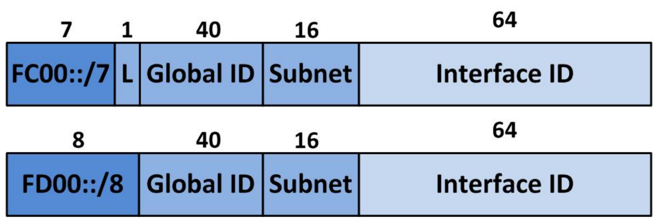

Figure 4: Unique Local RFC and Local Scope Address Formats

One of the first IPv6 unicast address subtypes you will probably encounter is named **Unique** Local, and analogous to the RFC 1918 Private Addressing space in IPv4. You may recall that these types of address exist solely **within** an enterprise, and are not valid for routing on the global Internet. When the IPv6 specifications were originally created, this function was assigned to another class of addressing called **Site Local**, with the prefix **FEC0::/10**, but soon after was deprecated.

RFC 4193 actually specifies a large address space for Unique Local addresses with the prefix **FC00::/7**, which actually contains two separate ranges. The **FC00::/8** range is designated for **centrally** assigned addressing (yet undefined) , while the **FD00::/8** range is for **locally** assigned addressing (see Figure 4). In the narrowest technical sense, Unique Local addressing is represented by the **FC00::/7** prefix, but, in actual practice, you will end up using the **FD00::/8** prefix/range instead.

Link Local

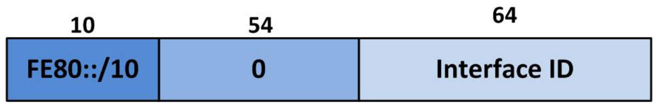

Figure 5: Link Local Address Format

Link local addresses do not have a direct correspondence to an addressing type in IPv4, as some of the other IPv6 addresses or subtypes do. This type of unicast address is dynamically created when IPv6 is enabled on an interface and is used only on the link on which it exists (hence the name link **local**). An entire set of IPv6-specific protocol mechanics (beyond the scope of this discussion) rely directly on these addresses, including routing protocols. One of the most important facts to recall about link local addresses is that they are never forwarded off of the local subnet in use.

Unspecified

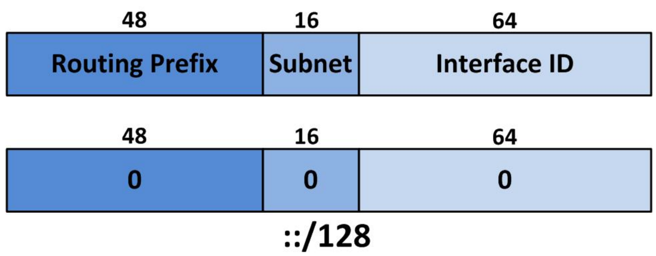

Figure 6: Unspecified Address Format

One of the many improvements built into the core protocol mechanics of IPv6 is autoconfiguration and dynamic host addressing. While Dynamic Host Configuration Protocol is alive and well in the newer version of the Internet Protocol, hosts have the capability of gaining connectivity to the network even if a DHCP server is not reachable or not present. This functionality is designed to create initial connectivity for the purpose of gaining more meaningful access. If you have ever connected to a wireless network and failed to receive an IP address, then Windows most likely assigned an autoconfiguration address from the **169.254.0.0/16** range. This is a similar function to the unspecified address of **::/128** (Figure 6).

Loopback

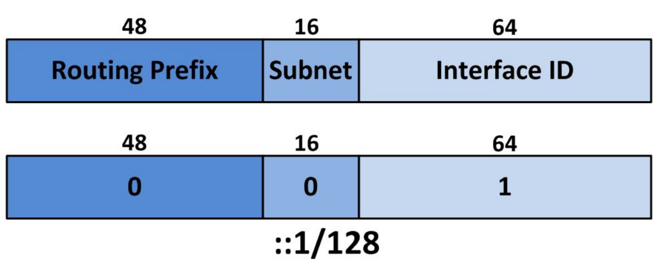

Figure 7: Loopback Address Format

Internal testing using loopback addressing is a familiar concept in IPv4 (using address 127.0.0.1), and this same functionality was carried over into IPv6. Testing TCP/IP functionality in IPv6 is also possible, but the addressing involved is simply a value of ::1/128 (Figure 7).

Global Unicast

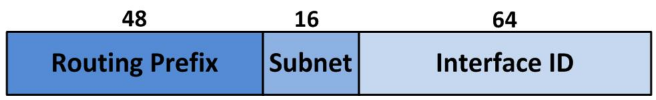

Figure 8: Global Unicast Address Format

Due to the vast addressing space built into the foundation of IP Version 6, the last and largest of the unicast addressing types is **Global Unicast**, sometimes referred to as **Aggregatable Global Unicast** (see Figure 8, above). As you might suspect from the name, these addresses are intended for routing on the global IPv6 Internet, and sections of the address space are assigned to regional registries in contiguous blocks to allow for simplified route advertisement. According to the Internet Assigned Numbers Authority (IANA), the defined prefix for global unicast addresses is **2000::/3**.

Conclusion

IP Version 6 is clearly a much more complex protocol than its still-popular predecessor, IPv4. Even aside from the unfamiliar address representations (i.e., hexadecimal format), many other intricate details exist that make mastering the new protocol somewhat challenging. One such area in IP Version 6 is the various classifications and subtypes of addressing, each with a specific function, format, and usage. As a network professional, you may not encounter each type on a regular basis, but knowing how to effectively recognize and interact with them is essential to personal and professional success.

Learn More

To learn more about how you can improve productivity, enhance efficiency, and sharpen your competitive edge, Global Knowledge suggests the following courses:

IPv6 Fundamentals Course

Introduction to IPv6: Protocols, Services, and Migration

IP6FD - IPv6 Fundamentals, Design, and Deployment v3.0

Visit **www.globalknowledge.com** or call **1-800-COURSES (1-800-268-7737)** to speak with a Global Knowledge training advisor.

About the Author

Joe Rinehart, MBA, CCIE #14256, CCNP/DP/VP is a professional trainer specializing in technology, business, and social media. He is also a successful speaker and published author, as well as a columnist for the Federal Way Mirror. He is active in the social media space, managing one of the largest groups on LinkedIn, as well as serving on the national steering committee of the Cisco Collaboration Users Group. Joe also serves as president of the Seattle Cisco Users Group, serving technology professionals throughout the Puget Sound region.

Joe Rinehart

MBA, CCIE #14256, CCNP/DP/VP

President and Chief Edutainment Officer

Gracestone Professionals, LLC

jrinehart@gracestonecompany.com

Twitter: jjrinehart

