

The VoIP Peering Puzzle Part 5: Addressing and Numbering Concepts

By Mark A. Miller | Dec 7, 2006 | [Print this Page](#)

http://www.enterprisenetworkingplanet.com/unified_communications/The-VoIP-Peering-Puzzle151Part-5-Addressing-and-Numbering-Concepts-3647851.htm

In our last few tutorials, we have examined the work of the Internet Engineering Task Force's SPEERMINT (Session PEERing for Multimedia INterconnect) working group, which has developed an architecture and a set of procedures to facilitate peering between two or more SIP-based providers.

The world of VoIP communications that goes beyond SIP-to-SIP connections, however, involves additional challenges. Still quite typical of VoIP usage is the following case:

Your workstation is attached to an IP-based network (such as the Internet), and you want to place a voice call to a friend that is connected to the Public Switched Telephone Network (PSTN). Your workstation (the source) is identified with an IP address, and your friend is identified with a telephone number. Each of these addresses resides in a database, and is somehow associated with your workstation, name, location, or some other logical link to you. But here is the key question: Since the source workstation is on an IP network, and the destination telephone is on the PSTN, how do these two different addressing schemes come to some type of understanding so that my call via the Internet will be successful?

This, as we used to hear in graduate school, is a non-trivial question. The short answer is that a process called ENUM (for Electronic Numbering) handles those details. However, as with many protocols and systems, there are a number of underlying technologies that must collaborate in order for the end result to be successful. In this tutorial, we will define some of these underlying concepts and lay the groundwork for a more in-depth discussion regarding ENUM in the tutorials that follow.

Here are some of the subsystems that will enter into those future discussions:

E.164: The International Public Telecommunication Numbering Plan, a telecommunications standard developed by the International Telecommunications Union—Telecommunications Standardization Sector, or ITU-T (see www.itu.int). This standard defines formats for four distinct categories of numbers that are used in international public telecommunication services: geographic areas, global services, networks, and groups of countries. Each of these categories has a specific numbering format to enable calls to be routed from source to destination, with a maximum of 15 digits, and beginning with a Country Code of 1–3 digits in length. The balance of the address format is defined by the specific category of numbers (from the list of four, above) being defined and the intended application. For example, the E.164 address for geographic areas includes a Country Code (1-3 digits), a National Destination Code, and a Subscriber Number.

IP Address: an addressing format defined in the original standard for the Internet Protocol, RFC 791 (see <ftp://ftp.rfc-editor.org/in-notes/rfc791.txt>), published in 1981. This format defines an address of 32 bits in length, is now referred to as an IPv4 (IP version 4) address, and is comprised of a part that defines a Network ID, and another part that defines a Host ID. These 32 octets are typically represented in a more human-friendly form, known as "dotted decimal," such as 10.53.10.1, where each of the four sections represents the decimal equivalent of eight of the 32 bits within that address. With the rapid growth of the Internet, IPv4 addresses are beginning to run out, so another addressing scheme, defined as part of the IP version 6 (IPv6) development project, has been defined. The IPv6 address is 128 bits in length, which adds a significantly larger number of addresses to the available pool (see <ftp://ftp.rfc-editor.org/in-notes/pdf/rfc4291.txt.pdf>). Of note, one of the reasons that have been given in the past for the IPv4 address shortage is the expectation of IP-connected wireless devices (such as cell phones and PDAs), which is now becoming reality. A good resource for further information on IPv6 is the IPv6 Forum at www.ipv6forum.org.

URI: a Uniform Resource Identifier, described in RFC 3986, developed by the IETF (see <ftp://ftp.rfc-editor.org/in-notes/rfc3986.txt>). The concept of the URI dates back to the early days of the World Wide Web, and was developed to provide a simple means of identifying a resource, where that resource might be a document, an image, an information source, or a service. The URI can be further classified as a Uniform Resource Name (URN), a Uniform Resource Locator (URL), or both. The syntax for a URI is defined in RFC 3986 (and can get rather involved), so perhaps some examples from that document, would be helpful:

`ftp://ftp.is.co.za/rfc/rfc1808.txt`

`http://www.ietf.org/rfc/rfc2396.txt`

`mailto:John.Doe@example.com`

`news:comp.infosystems.www.servers.unix`

`tel:+1-816-555-1212`

`telnet://192.0.2.16:80/`

(Note that both telephone numbers and IP addresses are considered URIs.)

Domain Name Service or DNS: operates similarly to a white pages telephone directory, in that it translates host names and URLs into IP addresses. DNS was originally developed in 1983, and has been

implemented in a number of operating systems since that time. The benchmark standards for DNS are RFC 1034 (see <ftp://ftp.rfc-editor.org/in-notes/rfc1034.txt>), and RFC 1035 (see <ftp://ftp.rfc-editor.org/in-notes/rfc1035.txt>). Another important aspect of the DNS is the naming convention applied to domain names, which are specified in two parts: a host portion (e.g. jupitermedia) followed by a suffix identifying the type of organization (.com, .org, .mil, and so on).

Unknown domain names are resolved by sending a query to a DNS server, which performs a database lookup and responds with either the corresponding IP address or an error message. In most cases, the DNS server will reside somewhere on either the local or enterprise network. That local DNS server has access to other DNS servers (for example, a DNS server at the Internet service provider) in the event that your DNS server does not contain the answer to your query.

The data within the DNS is stored in one of several record types. These include: the A (address) record, which stores an IPv4 address, the AAAA record, which stores an IPv6 address, an MX (mail exchange) record, which maps a domain name to a mail exchange server for that domain, and a NAPTR (Naming Authority Pointer) record, which allows the information to be rewritten as the name analysis proceeds.

With that insight into some of the underlying subsystems, our next tutorial, which we'll post tomorrow, will dig deeper into the inner workings of ENUM, which employs the subsystems described above.

Copyright Acknowledgement: © 2006 DigiNet ® Corporation, All Rights Reserved

Author's Biography

Mark A. Miller, P.E. is President of DigiNet ® Corporation, a Denver-based consulting engineering firm. He is the author of many books on networking technologies, including Voice over IP Technologies, and Internet Technologies Handbook, both published by John Wiley & Sons.