The VoIP Peering Puzzle&#65533;Part 1: Concepts and Challenges

By Mark A. Miller | Nov 15, 2006 | Print this Page
http://www.enterprisenetworkingplanet.com/unified_communications/The-VoIP-Peering-Puzzle151Part-1-Concepts-and-Challenges-3644066.htm
If you've been reading Enterprise VoIPplanet.com, you have doubtless come across the term VoIP peering; it is certainly among the top buzzwords in today's networking culture.

Webster's defines peer as "one that is of equal standing with another." But to fully understand the concepts of VoIP peering, we must roll back the clock a century or so, and examine the architecture of the first telephone systems.

Alexander Graham Bell's invention of the telephone in 1876 was initially available to a limited few—those who could afford not only the telephone instrument, but physical connections to others that possessed similar technical foresight and economic means. Switching systems had not yet been developed, so if Dr. Smith needed a telephone line to Pharmacist Jones, a physical cable was installed between their respective locations to make that connection.

With that physical connection established between Smith and Jones, they became peers, or individuals with equal standing with another—at least as far as communicating by telephone was concerned.

Unfortunately, if either Smith or Jones wanted to consult with Dentist Brown, additional lines would have to be constructed, and before long, large cities such as New York, were draped with telephone poles and cables connecting various locations of the rich and famous.

Switching systems changed the telephone network from a point-to-point to a point-to-multipoint topology, and over time gave rise to the interconnection of central offices and switching systems that comprise the Public Switched Telephone Network, or PSTN, that we know today.

Island time
Fast forward a century, and we see that the packet-based Internet Protocol (IP), has emerged as the dominant communications medium. Many organizations, from small businesses to large, multi-location enterprises, have embraced IP as their data transport protocol of choice, creating both local and wide area networks (LANs and WANs) all based upon an IP infrastructure.

However, since they are not connected by a ubiquitous IP grid, these networks are frequently referred to as IP Islands, and communication between islands—not to mention with much of the outside world—requires additional connectivity.

For example, an enterprise, with, say, one location in New York and another in Chicago, could procure a high speed leased line, such as a T1or T3 circuit, from an inter-exchange carrier to connect its two "islands". However, in doing so, it would be doing little more than recreating the point-to-point network of Alexander Graham Bell's era. As long as you only want/need to talk with people within your privately interconnected island network, the enterprise VoIP system can stand alone.

However, as soon as you need to speak with someone in the outside world (often referred to as an "off-net call"), your inter-island connections no longer get the job done. To get off your archipelago, you need a ubiquitous bridging connection. Historically, this has been provided by the PSTN, and that PSTN connection can become the limiting factor—dictating the cost, quality and application support for the end-to-end connection.

Making VoIP the best it can be
Thus, if the PSTN connection between VoIP islands is replaced with an IP network, several immediate benefits emerge.

- First, the "per minute" charges associated with the PSTN go away, since IP networks transmit information in packets, independent of the connection time.
- Second, the hardware costs associated with IP-to-PSTN gateways are eliminated.
- Third, the related signal degradation, owing to the multiple signal format conversions that occur when the IP network connects to/from the PSTN is eliminated, which should improve the voice quality.
- Fourth, multimedia applications, such as video conferencing, which might not be able to traverse the PSTN because of bandwidth constraints, can now flow on an end-to-end basis.

There are two different types of peering arrangements.

The first is called bilateral peering, in which two locations are connected in a point-to-point topology with an IP network. In this case, the two peers may be two locations from the same enterprise, or two locations from two distinct enterprises (such as trading partners—for example, a manufacturer and supplier) that have a significant amount of inter-network traffic. Some agreement is reached between the various parties to share the costs of the connection.

The second type of peering arrangement is a called multilateral peering, also known as federation peering. This arrangement is similar to a star topology network, in which the islands all connect to a central location, typically a provider of peering services.

Peering—good, but not simple

As might be expected, the peering relationship is not as simple as just plugging one network into another. Much like other computer and communications architectures, there are a number of issues that must be addressed:

- The physical media for interconnection, such as a fast Gigabit Ethernet backbone.
- Signaling, or call setup (establishment) and teardown (disconnect) messages that are transmitted from the sender to the receiver. Different networks use varying signaling protocols, which requires a meeting of the minds (protocol conversion) before an end-to-end connection can be successful.
- Registry services, with databases to cross-reference telephone numbers and IP addresses.
- Business issues, defining pricing, billing, traffic reporting, and other contractual terms between the parties.
- Location services, identifying where the desired application is located.
- Network security, preventing network topology information, or other proprietary information, from inappropriate disclosure.
- Government compliance, assuring that access to emergency services (E911) and law enforcement call monitoring functions can be maintained.
- End-to-end quality of service, or QoS, as the original voice signal undergoes various format conversions from the original analog to digital to the recipient analog, perhaps traversing multiple IP-TDM-IP network connections in route.
- Identity notification, such as Caller ID.
- Preventing unwelcome calls, known as Spam over Internet Telephony (SPIT)
- Adherence to emerging standards, thus minimizing multi-vendor interoperability issues.

These, and other issues must be addressed in order to implement a successful peering operation, and will be examined in future tutorials. Our next tutorial will look at the developing standards for VoIP peering, beginning with the network peering architecture developed by the Internet Engineering Task Force (IETF).

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.