nrri

The Transition to an All-IP Network: A Primer on the Architectural Components of IP Interconnection

Joseph Gillan and David Malfara

May 2012
NRRI 12–05

About the Authors

Joseph Gillan is a consulting economist who primarily advises non-incumbent competitors in the telecommunications industry. David Malfara is president of the ETC Group, an engineering and business consulting firm.

Disclaimer

This paper solely represents the opinions of Mr. Gillan and Mr. Malfara and not the opinions of NRRI.

Online Access

This paper can be found online at the following URL:
http://communities.nrri.org/documents/317330/7821a20b-b136-44ee-bee0-8cd5331c7c0b

Executive Summary

There is no question that the "public switched telephone network" (PSTN) is transitioning to a network architecture based on packet technology and the use of the Internet Protocol (IP) suite of protocols. A critical step in this transition is the establishment of IP-to-IP interconnection arrangements permitting the exchange of voice-over-IP (VoIP) traffic in a manner that preserves the quality of service that consumers and businesses expect. This paper provides a basic primer on how voice service is provided by IP-networks and discusses the central elements of IP-to-IP interconnection.

The ascendancy of IP networks is the architectural response to the growth of data communications, in particular the Internet. Although some voice services rely on the public Internet for transmission (a method commonly known as "over-the-top" VoIP), the Internet does not prioritize one packet over another; as a result, quality cannot be guaranteed for these services. To correct for this shortcoming, carriers are deploying managed IP networks that emulate PSTN-levels of quality. One of the challenges to IP interconnection is establishing traffic-exchange arrangements for these managed VoIP services that sustain quality of service on an end-to-end basis, even where multiple carriers are involved.

The paper explains VoIP and IP interconnection from a micro and macro perspective. First, it explains the general architectural arrangements that exist, especially during the transition to an all-IP network, when a variety of interconnection scenarios will occur (such as between IP networks and the traditional circuit-switched architecture of the PSTN). In addition, the paper provides a micro-look at packet-based voice service by discussing the different roles played by the various layered protocols that control VoIP communications.

Finally, the paper concludes with the caution that the evolution of telecommunications technology (and the regulatory issues that follow) will not end here. Newer architectural approaches (in particular, an approach known as the IP Multimedia Subsystem) are emerging that present their own challenges. This paper addresses the beginning stages of the transition, with a focus on education and providing a foundation for further research and understanding.

Table of Contents

Table of Figures

All figures were prepared by the authors.

I. Introduction

The familiar "public switched telephone network" (PSTN) is undergoing a fundamental change, becoming a network platform based on packet technology and the use of the Internet Protocol (IP) suite of protocols to support voice on a shared data network.[1] Although the conversion of the PSTN to an "all-IP" network is inevitable, a critical prerequisite to that outcome—establishing IP-to-IP interconnection arrangements for the exchange of voice traffic—is clouded in uncertainty. The purpose of this paper is to provide a basic primer on what is meant by IP-to-IP interconnection, with a focus on the exchange of traffic between the managed IP networks that we expect will become the PSTN of the future.[2]

The paper is structured as follows. In Section II, the paper provides a general overview, discussing the emergence of IP networks and, in particular, *managed* IP networks. The introduction of this technology is not an isolated event in North America; rather, it is a part of the worldwide transformation of the existing circuit-switched/TDM networks.[3] Section II also explains the difference between a managed IP network and the public Internet, a distinction that gives rise to the two types of voice (VoIP) services: those relying on managed networks, and those "over-the-top" applications that ride the public Internet. Crucial to understanding the migration of the services provided by the circuit-switched PSTN to a new IP-based network, as

---

[1] Circuit-switched networks (such as the PSTN) reserve a dedicated channel for the duration of a call, while packet networks constantly (and rapidly) break the call into individual "packets" that are transmitted individually on a shared channel and then reassembled at the far end. Readers may wish to consult additional texts to develop background before reading this paper. See, for instance, *Fundamentals of Telecommunications: Markets, Jurisdictions, and Challenges*, Peter Bluhm (as updated by Dr. Sherry Lichtenberg), National Regulatory Research Institute, January 2011, available at: http://communities.nrri.org/documents/317330/7bb0d474-4d21-479d-bf84-58eac5ef87a7

[2] As used in this paper, the term *PSTN* refers to the collection of *services* that end users have grown to expect, recognizable by common features such as dial tone, North American Numbering Plan addressing (phone numbers), and consistent service quality. In some contexts, others have used the term *PSTN* to describe the prevailing *technology* (circuit switching) that underlies the network today. This nomenclature—i.e., using the term *PSTN* to refer to a network defined by a specific technology—leads to confusion because it implies that the PSTN is disappearing (or "sun-setting") merely because its underlying technology is changing. Such a technology-specific definition of the term "PSTN" is inappropriate to this paper, however, because the central point of the paper is to discuss the interconnection arrangements needed to *preserve* end-to-end service quality as the PSTN moves between different technologies.

[3] The existing network architecture is frequently characterized by two defining characteristics: (1) its reliance on *time-division multiplexing* (TDM) to derive individual channels on a common transmission facility (be it copper or fiber), and (2) the use of *circuit-switching* to establish temporary connections between TDM channels on different facilities. We use both terms (TDM and circuit-switched) in this paper as shorthand for the architecture that underlies the PSTN today.

well as the important role that IP interconnection will play in that transition, is understanding the different ways in which VoIP is provided.[4]

Because a variety of interconnection scenarios exist for voice traffic, we devote Section III to explaining each of these scenarios. The most common "first-generation" interconnection scenario involves a call that is VoIP only on one end; that is, the call originates or terminates in VoIP, but the customer on the other end is served by a traditional circuit-switched carrier. This VoIP-to-PSTN interconnection scenario has historically been the focus of FCC orders. We also explain a variant of this scenario, commonly referred to as *IP in the middle*.[5] Increasingly, however, IP-to-IP interconnection will involve calls in which VoIP applies at both ends of the call, so we address that scenario as well.

Finally, we take a deeper look inside an interconnection architecture to explain the different layers (i.e., header and trailer components) of a typical VoIP packet call. This level of detail is essential to understanding the type of intelligent interconnection that will be needed to exploit fully the innovations made possible by IP technology. Thus, the paper provides a macro-view of the basic architectural components of an IP interconnection arrangement, as well as a micro-view of an individual call as it becomes part of a packet stream.

II. Placing IP Interconnection in Context

The emerging dominance of IP networks is a global development with public-policy implications. As a result, a discussion concerning changes brought about by IP technology is occurring in Canada and Europe, as well as the United States.[6] For instance, in the United States, the FCC has issued an FNRPM to consider IP interconnection,[7] while the Canadian

---

[4] Understanding the telecommunications industry is frequently complicated by its use of multiple terms that mean the same thing. For instance, *local-loop*, *twisted-pair*, and *last-mile* each refer to the final connection to the customer; similarly, *central office*, *class 5 switch*, and *end office* refer to the final point of switching. In this instance, however, confusion is created by using a single term, *VoIP*, to mean two different things: an over-the-top application such as Skype or Vonage, as well as a managed service, such as the VoIP offering of the typical cable company, facilities-based CLEC, or incumbent LEC.

[5] *IP in the middle* is the term used to describe a network configuration in which both ends of the call are legacy PSTN subscribers, but at least a portion of the transport facilities between them are packet-based and use Internet Protocol.

[6] Although beyond the scope of this paper, we include a short listing of key regulatory decisions/documents addressing IP interconnection in Canada and Europe in Appendix B.

[7] Report and Order and Further Notice of Proposed Rulemaking, *Connect America Fund et al*, FCC 11-161, WC Docket Nos. 10-90, 07-135, 05-337, 03-109, GN Docket No. 09-51, CC Docket Nos. 01-92, 96045, and WT Docket No. 10-208 (rel. Nov. 18, 2011) ("IP Interconnection FNRPM").

Radio-television and Telecommunications Commission has issued a decision addressing how IP interconnection will be introduced in that country.[8] Other countries are also addressing this issue. A list of key documents discussing the global development of IP networks appears in Appendix B.

A. Transitioning the PSTN to a data-oriented network

The ascendancy of IP networks is an architectural response to the ascendancy of data communications. The traditional telephone network was designed to support voice service, with data services provided as an overlay.[9] With IP, these roles are reversed; the underlying network is designed for data services; VoIP is only one of a number of data streams, but it is one requiring specialized treatment to ensure high quality.[10]

To be sure, the technology of the PSTN has changed before. It has moved from its roots as an analog network with in-band signaling to its current architecture characterized by high-speed fiber transport, digital switching, and call control managed by a parallel signaling network, Signaling System Seven (SS7). These evolutions, however, were largely silent—that is, the changes were known primarily to those who managed (or regulated) the network, but were generally invisible to consumers (and the political process).[11]

---

[8] Telecom Regulatory Policy CRTC 2012-24, Network Interconnection for Voice Services, Canadian Radio-television and Telecommunications Commission, January 19, 2012 ("CRTC IP Interconnection Order").

[9] A core benefit of a packet-based architecture is its flexibility. In a circuit-switched network, switching and transmission resources are configured into channels of a *fixed* size, designed to accommodate single conversations. Though many channels can be grouped onto a single physical facility, there is no dynamic ability to share the bandwidth of that physical facility; its configuration is fixed and dedicated to serve each assigned channel. This means that when a channel is not in use, its bandwidth is wasted because it cannot be made available to a different application. Packet networks, however, can *dynamically* adjust the bandwidth of a physical facility among a variety of applications, allocating more or less of the total bandwidth of the facility to each session as needed.

[10] The traditional circuit-switched architecture addressed quality by reserving a circuit for the duration of a call. By its nature, this design eliminates the possibility of congestion in the information flow. Because packet networks share resources to a much greater extent, other techniques must be used to prevent congestion from affecting quality. We address this issue below.

[11] This is not to say that prior technological changes did not affect consumers through improvements in quality (such as the significant post-dial delays eliminated by SS7), reductions in cost, or changes in competitive conditions. The point is rather that the technology itself did not become the source of different policies or consumer perceptions.

The emergence of IP technology is far more visible, but not well understood. A major explanation for this visibility is that IP technology supports one of the most dynamic innovations of our time, the global Internet. Yet Internet Protocol is used in networks other than the Internet, and the importance of IP technology extends beyond this one (albeit remarkable) network.

A key feature of today's PSTN is the quality expected of voice service. As the PSTN transitions from a circuit-switched to a packet-switched architecture based on IP, a key question is how we will maintain the quality of voice and data communications provided by today's circuit network. The quality dimension that is most important to understanding the transition to IP networks is dealing with congestion, and understanding how congestion will be addressed in these new networks.[12]

In a circuit-switched network, SS7 addresses the potential for congestion directly at its source, when a call is first dialed. SS7 immediately determines whether a circuit is available. If so, the call proceeds; if not, the caller is sent a busy signal. In other words, SS7 addresses congestion through avoidance, by blocking the congestion-producing call (and *only* the congestion-producing call).

The public Internet addresses congestion with a very different philosophy. The public Internet is referred to as a *best-efforts* network because it is designed to treat all packets the same way. In a best-efforts architecture, queuing is used to address congestion. The network delays some packets (i.e., it “buffers the traffic”), either at its source or at a router, while other packets are transmitted.[13] The result is that all packets are created equal—and, therefore, are equally likely to be delayed.

Because the process of buffering (or discarding) packets can be arbitrary, packets from *all* information flows are likely to be affected by packet loss and degradation during a period of congestion. On the Internet, therefore, every user can experience the consequences of congestion. In contrast, a circuit-switched architecture addresses congestion by blocking some calls entirely, but otherwise maintains the quality of the calls that do go through.

---

[12] Many factors affect the perceived (by an end user) quality of a phone call, including noise and the audio frequency range being transmitted. Because both traditional and packet networks are digital, noise issues are generally resolved. As to the available audio frequencies, this parameter is limited in circuit-switched networks by the network's design and standardization. Because packet networks can flexibly assign bandwidth to different applications, IP networks are capable of “high-definition” voice service (HD voice) that increase the fidelity of the audio stream and thus improve sound quality. The focus of this paper, however, is ensuring that voice packets arrive with a speed and reliability necessary to maintain, at a minimum, the audio expectations of the PSTN.

[13] If congestion becomes too much of a problem, the network discards some of the "overflow" packets.

B. Preserving voice quality through managed IP networks

A *managed-packet network* is able to combine the quality standards of the PSTN with the flexibility of a packet-based network. A managed-packet network is an IP network that is capable of defining class-of-service categories and then adapting network performance to maintain a quality-of-service level appropriate to each class. Internationally, these networks are known as next-generation networks (NGNs);[14] their importance stems from their ability to resolve congestion and thereby ensure service quality. Managed networks are particularly useful for a real-time service such as VoIP, which requires that packets be delivered reliably and without significant delay (latency).

Managed-packet networks are a key improvement over a best-efforts IP network such as the Internet. A managed-packet network can provide a higher priority to services (such as voice) that need greater reliability or reduced latency. The technique of *traffic shaping*[15] improves network performance for all information flows and provides the means to guarantee network performance for VoIP. Although packet networks are designed to support a number of media streams and applications—Internet access and video entertainment are two prominent examples—this paper focuses on the emergence of such networks as the replacement technology for the circuit-switched PSTN in the provision of voice service.

The difference between managed networks and the public Internet leads to a parallel distinction between two different categories of VoIP services: *over-the-top* and *managed-VoIP* services.[16] Over-the-top VoIP offerings, such as Vonage or Skype, ride the public Internet for some part of the service; these packets are treated no differently (during this part of their transmission) from other Internet traffic. In circumstances of low congestion, over-the-top VoIP services can provide quality equivalent to that of managed networks, although that quality cannot be guaranteed because the service relies on the best-efforts design of the Internet.

In contrast, most facilities-based providers of VoIP separate the voice packets from other traffic on the Internet in order to preserve quality. This is true, for instance, for incumbent local

---

[14] http://www.itu.int/ITUT/studygroups/com13/ngn2004/working_definition.html

[15] *Traffic shaping* is the process of dynamically managing the flow of packets through a network to conform to a defined *traffic profile*. The traffic profile identifies which data flows will be assigned what priority so that network resources can be allocated. Traffic shaping is a generic term for a variety of techniques to achieve a specified outcome (for example, low latency, low jitter, or low packet loss) for a particular data flow.

[16] The generally accepted nomenclature for some of these services is not particularly helpful in drawing the distinctions we feel are important. For instance, over-the-top services are sometimes called *nomadic*, as though the fact that they can be moved is their most important feature. Equally confusing, VoIP offerings that rely on managed networks are sometimes referred to as *fixed* or *facilities-based* VoIP services. We deliberately adopt the terms *over-the-top* and *managed* to distinguish what we believe to be the most critical difference between the offerings: the ability to guarantee service quality in one arrangement and not the other.

exchange carriers such as AT&T (U-verse)[17] and Verizon (FiOS),[18] as well as cable-based VoIP providers[19] and CLECs.[20] These examples are not exceptions; rather, this illustrative list validates the claim that PSTN is moving to a new packet architecture but is not simply becoming "just another application on the Internet."

The distinction between over-the-top VoIP and managed-VoIP services is critically important for interconnection. As explained above, the PSTN is migrating to managed-packet networks because of reliability, quality, and security. A central challenge for IP-to-IP interconnection arises from the need to provide quality-of-service assurances on an end-to-end basis across managed IP networks, just as the PSTN has historically provided quality of service end to end. Achieving this goal will require a degree of interconnection and interoperability beyond that which applies to the public Internet.

III. Basic IP Interconnection Architectures

In this section, we explore the basic architectural components of IP interconnection. We explore these questions from a variety of perspectives and for a variety of basic architectures. We describe the dominant (yet transitional) architecture for calls that traverse both circuit-switched and IP architectures. We also describe the more straightforward (albeit less prevalent in the near term) case in which both end points are IP.

To begin, VoIP-to-TDM[21] interconnection is likely to be a part of the landscape for years. As shown in Table 1, as of the end of 2010, there were still over 116 million traditional

---

[17] *See* http://www.att.com/u-verse/explore/home-alarm.jsp (emphasis added): "AT&T U-verse Voice service is provided over AT&T's world-class managed network and *not the public Internet*."

[18] See http://newscenter.verizon.com/press-releases/verizon/2010/fios-digital-voice-heres.html (emphasis added): "To understand the features and quality of FiOS Digital Voice, you first need to know that the service is not the same as the services you get with a little Internet adapter for your modem and phone, and *it does not ever touch the public Internet*."

[19] *See* http://www.comcastoffers.com/faq_phone/#q11:

Q: Are Comcast Digital Voice calls routed across the Internet?
A: No! Comcast keeps the calls on their private broadband network until our switch delivers the call on the PSTN (public switched telephone network). Calls are just as secure with Comcast as with any other telephone provider.

[20] *See* http://ir.cbeyond.net/index.cfm (emphasis added): "Cbeyond is a Voice over Internet Protocol-based *managed* services provider to small businesses."

[21] As noted earlier (fn. 4 *supra*), the existing network is typically referenced either by how calls are routed (i.e., using circuit-switching) or by how they are commonly organized to

switched-access lines in service. It will take years before the traditional architecture is modernized all the way to the customer premises itself (i.e., the end point is a customer subscribing to an IP service). Circuit-switched (TDM) networks will not vanish overnight, and a mix of IP and TDM networks will occur for the foreseeable future.

Table 1: Comparing VoIP and TDM (Year End 2010 - millions of lines)[22]

	ILEC	CLEC	Total
VoIP	2.9	28.8	31.7
TDM	94.7	22.2	116.9
	97.5	51.1	148.6

Table 1 reports the number of subscribers served by VoIP. This focus, however, understates the amount of IP technology in the network because it fails to appreciate that significant portions of the interoffice transport network are being replaced by IP, even if the subscribers themselves continue to be served by circuit switches. For instance, it has been estimated that 90% of the interLATA interoffice network has already been replaced by IP technology (along with 60% of the intraLATA interoffice network).[23]

IP interconnection issues arise in a variety of different scenarios. The first involves interconnection in which VoIP is converted to TDM at the network edge to interconnect with the TDM facilities of a traditional network. The second is similar, but assumes that the ILEC has deployed IP transport (which can be used to reach its TDM-based end offices). A similar scenario occurs when both carriers have deployed IP-transport networks (even if they also have TDM local switches), which gives rise to a configuration in which IP interconnection occurs "in the middle," with subscribers at both ends of the call connected to TDM switches. Finally, over time, more and more customers will have subscribed to VoIP services and IP interconnection will provide end-to-end service. We describe each of these configurations in more detail below.

A. VoIP-to-TDM interconnection

Today, most VoIP interconnection involves a VoIP call being converted to a TDM call at the boundary between the IP network and the traditional circuit-switched networks of incumbents (and other carriers that use circuit-switched technology). Traditional interconnection is often used even where both carriers have deployed IP transport networks on either side of the interconnection, because traditional interconnection is well-understood and fully defined by interconnection agreements, tariffs, technical standards, pricing standards (though these are

---

share a transmission facility (i.e., using time-division multiplexing, or TDM). We prefer the term *TDM* here because the more relevant facility for interconnection is the *transmission* link between the networks. In other words, while the calls within the existing network are circuit-switched, the traffic is formatted for TDM transmission at the network edge, and the conversion that must occur is between TDM and VoIP (although, as we explain in detail here, interconnection requires interoperability between signaling and other functions as well).

[22] *Local Telephone Competition, Status as of December 31, 2010*, Industry Analysis and Technology Division, Wireline Competition Bureau, October 2011, Table 8.

[23] Presentation of Carl Ford, Vice President, Crossfire Media, to National Association of Regulatory Utility Commissioners, Staff Telecommunications Subcommittee, February 14, 2009.

frequently in dispute), and operational systems. Figure 1 portrays a VoIP-to-TDM Interconnection, with the TDM portion shown in the upper half.

VoIP-to-TDM interconnection requires that both the media stream (i.e., the conversation itself) and the signaling information (for call setup) be interconnected and interoperate. These two interconnections are illustrated in Figure 1. An *SS7 STP* is a Signaling System 7 Signal Transfer Point. A *Class 5 switch* is a TDM switch that connects local loops and handles voice service within a local exchange.

Figure 1 introduces a number of terms common to IP networks, specifically *media gateway*, *signaling gateway*, and *call agent*. The media and signaling gateways provide the protocol conversions that must exist at the boundary of an IP-to-TDM network. Within both IP and TDM networks, digitized voice signals are coded and decoded using a *codec*.[24]

---

[24] Virtually all calls, whether transported in IP or TDM, are in digital form. They therefore must first be transformed from an analog form to a digital form. This conversion is accomplished by a *codec*, which is a device or computer program capable of encoding or decoding a digital data stream or signal. The word *codec* is formed by the combination of "**co**mpressor-**dec**ompressor" or, more commonly, "**co**der-**dec**oder." Packet networks can support many different codecs, including codecs that encode high-fidelity data (frequently called HD Voice) that is of a substantially higher quality than the fidelity provided by the circuit-switched network.

Figure 1: VoIP-to-TDM Interconnection

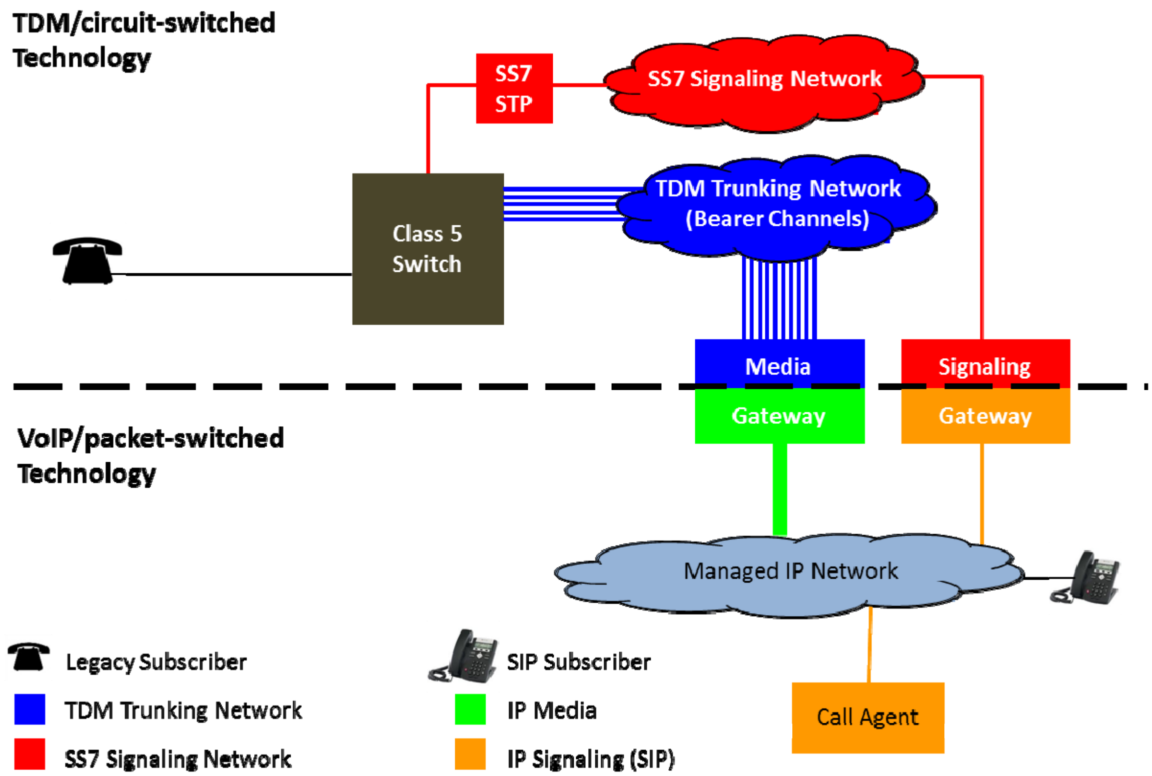

As shown in Figure 1, a media gateway converts the voice audio signal from the packet network's codec to the encoding used by the standard codecs in a conventional switched network, and from IP to TDM format. With respect to call control, Session Initiation Protocol (SIP)[25] is the most common signaling protocol for VoIP services, while SS7 signaling is generally used within the traditional network. Therefore, VoIP to TDM interconnection requires that SIP signaling from the VoIP network be translated and connected to the SS7 signaling that provides circuit control in the conventional circuit-switched network.

In addition, the packet network relies on a call agent that plays the role of controller. The call agent is responsible for registering the end points of an IP call session, such as the phone and media gateways. In addition, the call agent is responsible for establishing the calling session (and instructing the end points to end a session when the call is complete). In lay terms, the call agent functions as the "brains" of the connection by controling the call session. The call agent can be virtually anywhere on a packet network because it remotely performs its role through signaling (the sending/receiving of SIP packets that contain network instructions).

---

[25] *Session Initiation Protocol* is a signaling protocol used to set up, maintain, and tear down multimedia communication sessions such as voice and video calls over packet networks using Internet Protocol.

As Figure 1 indicates, media and signaling gateways provide conversion to (and from) TDM. To maximize the efficiencies of IP, it makes sense for these conversions to occur as close to a TDM end-point as possible. This arrangement is discussed below.

B. IP-to-IP interconnection (TDM end point)

Figure 2 (below) illustrates an IP interconnection that ultimately involves a call to or from a customer served by a traditional, circuit-switched end office. This interconnection would arise where an ILEC has deployed an IP transport network (and, perhaps, has decommissioned its traditional access tandems) but still serves a large number of end users from a traditional TDM switch.

Figure 2: IP-to-IP Interconnection with VoIP/TDM End Point

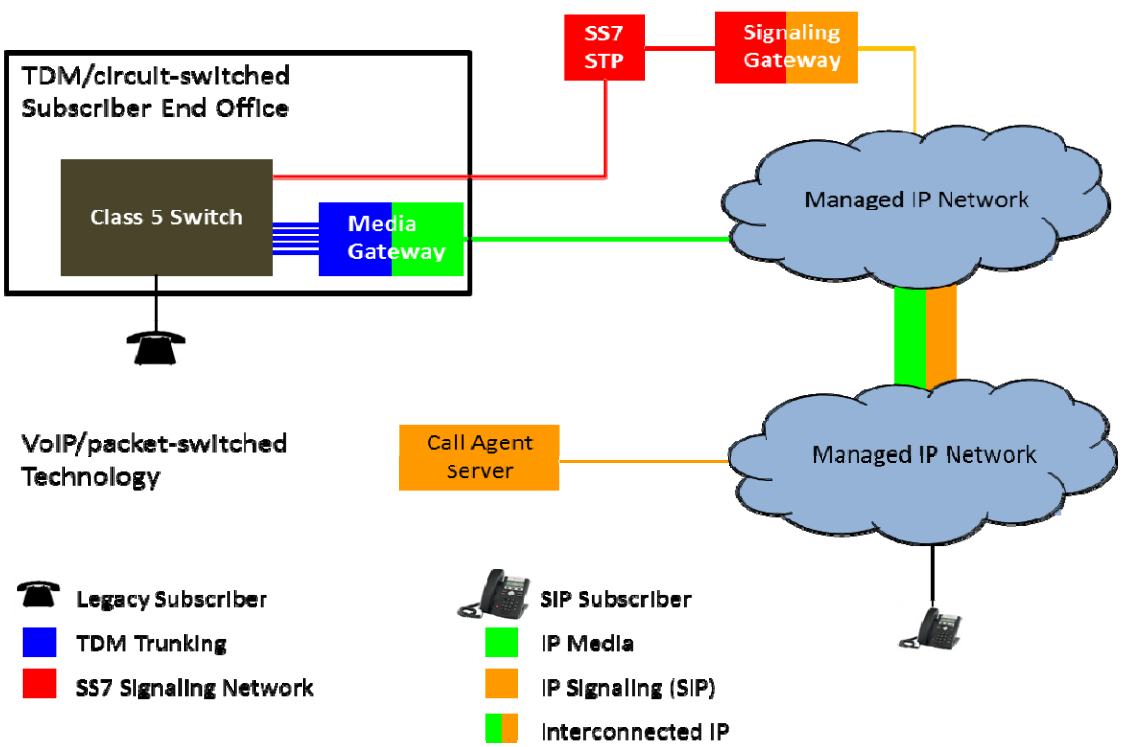

As shown in Figure 2, the Class 5 switch is connected to a media gateway in much the same way as it would connect to a tandem switch in a circuit-switched network (i.e., using Feature Group D trunk facilities). The principal difference between a tandem connection and the connection to the media gateway is that in this case the media gateway resides at the same geographic location as the Class 5 switch (i.e., in the same wire center).

Figure 2 shows the requirement for two gateways: one for the voice conversation itself (the media gateway), and a second gateway to provide the signaling interface (signaling gateway). Very little difference exists between Figure 1 (VoIP-PSTN) and Figure 2 above. The incremental change is the use of IP transport to reach the end-office (otherwise the call flow is equivalent to the architecture depicted by Figure 1).

C. IP-to-IP interconnection (both end points TDM)

Similar to the configuration above, it is also possible to have a transitional configuration in which both end points are served by TDM, but the transport networks of each carrier have been converted to IP. In this case, as illustrated in Figure 3 below, the traffic exchange can occur in IP, even though the calls would still need to be converted to TDM at the end office for delivery to the called party.

Figure 3: IP in the Middle (with End Points TDM)

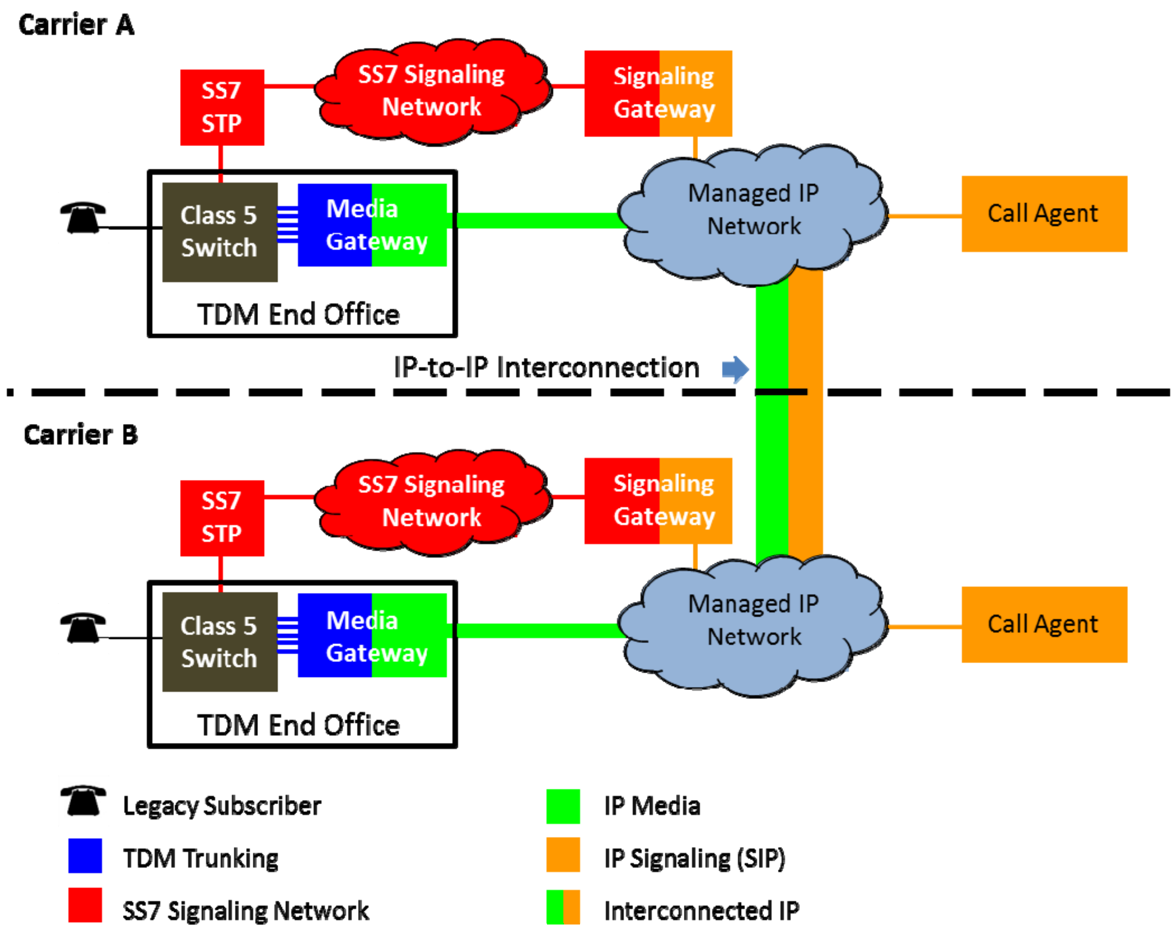

D. IP-to-IP interconnection (IP end points)

Of course, over time, more and more customers will be served by IP networks. Figure 4 illustrates the final scenario: interconnection between two IP networks when both end points are subscribing to a VoIP service.[26] Instead of a signaling "network" connection for call setup and control, the functions are provided via external server(s) accessed using SIP. These servers provide call setup, termination, maintenance, and redirection, as well as other capabilities.[27]

Figure 4: IP-to-IP Interconnection with VoIP/VoIP End Points

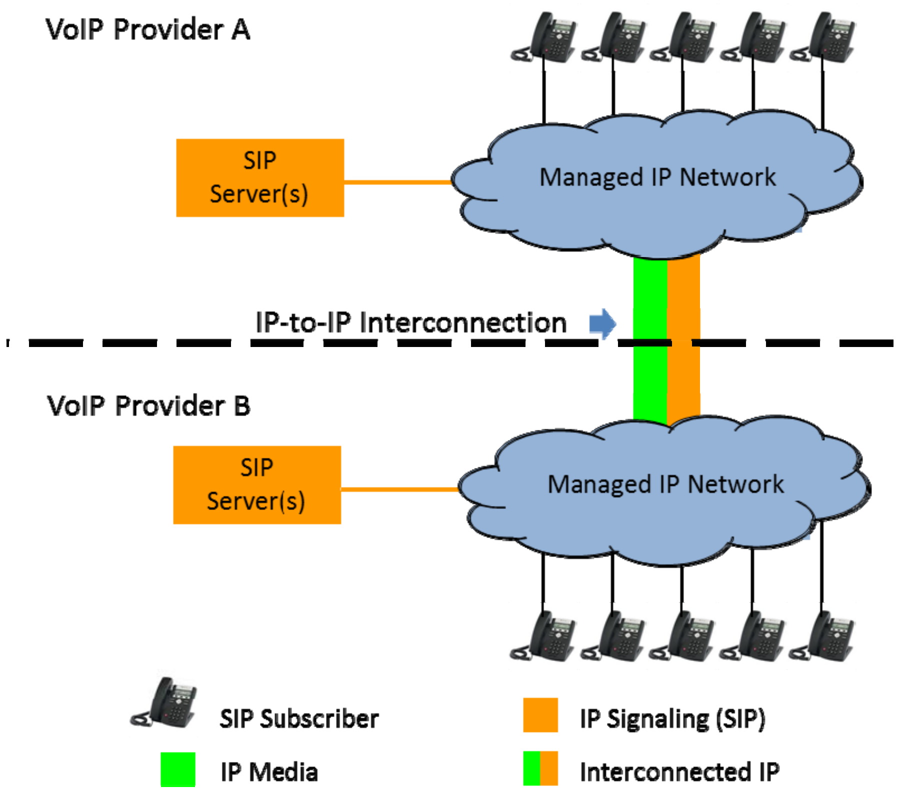

[26] Our discussion focuses on intelligent interconnection between managed-packet networks that we expect will be the architecture that replaces the circuit-switched architecture of the PSTN. This is not to say that interconnection and packet exchange will not occur between a managed-packet network and a best-efforts network. Indeed, we expect that some consumers will continue to subscribe to over-the-top VoIP services and will call (and be called by) consumers served by managed networks. Architecturally, such interconnections would be similar to the IP-to-IP interconnection scenario described here, although the benefits of network management will occur only within the managed network.

[27] As an example, *SIP forking* creates the ability to ring multiple end points simultaneously to locate the called party on the device (or at the location) of its choice.

It is important to appreciate that when both sides of the interconnection have VoIP end points, subscribers can receive a variety of additional service features. In order for this enhanced functionality to be retained as calls traverse different networks, however, the interconnection between the providers must allow for intelligent interaction between the various *layers*—i.e., the header and trailer instructions assigned to each packet—to ensure end-to-end service control.

IV. Inside IP Interconnection – Understanding the OSI Layers

Section III focused on the architectural components of an IP interconnection arrangement. In this section, we move from this macro view to a more micro look at exactly how VoIP packets are structured and what functions are provided by each layer. This discussion is important because IP interconnection must enable each layer to interoperate as a call moves between the IP network of one carrier and the IP network of another.

The common approach used to describe packet networks (at this level of detail) is the Open Systems Interconnection (OSI) Reference Model.[28] The OSI Reference Model uses seven layers of functionality to describe a complete system of communication. The highest layer is the data being transmitted (in the case of VoIP, the digital information that comprises the phone call).

Figure 5 illustrates each layer of the OSI reference model and uses the example of a traditional TDM call to illustrate what occurs at each layer. By mapping a traditional phone call to the OSI reference model, we can use the model to compare the information and media flows of TDM interconnection with those of IP Interconnection. A brief description of each layer follows:

> Layer 7 (Application Layer) represents the actual service (e.g., telephone calls) as perceived by the end user. It also contains the end-user interface, which conveys the transmitted information to and from the end user.
>
> Layer 6 (Presentation Layer) arranges the data from/to the end user into a format that will be understood by the receiving application (for example, audio encoding and decoding, encryption/decryption, and so on).

---

[28] The *OSI Reference Model* was developed by the International Organization for Standardization to depict Open Systems Interconnection. The OSI Reference Model is not the only abstraction model in common usage. The *Internet Model* is also an abstraction model for communication networks, but one which depicts a five-layer structure, collapsing the upper three layers of the OSI Model into a single application layer. We discuss the OSI Reference Model, with its seven-layer architecture, because it allows for a more granular discussion of the variances between circuit-switched and packet-switched networks than the Internet Model.

Layer 5 (Session Layer) is responsible for the initiation, identification, maintenance, and termination of a particular information exchange or dialog between two or more end users (for instance, to begin and end a specific telephone call).

Layer 4 (Transport Layer) provides the protocols that establish a reliable communications path between the sending and receiving end-user devices. Error correction and data flow control of information are two primary functions of this layer.

Layer 3 (Network Layer) provides the means for identifying originating and destination end users (for example, telephone numbers), as well as the communications path between networks to be used for connection (for example, the routing and translation functions).[29]

Layer 2 (Datalink Layer) is responsible for the intelligent conveyance of information between two devices on the same physical network. It includes a framing format that provides error correction, flow control, and acknowledgement processes to ensure the delivery of the frames carrying the information.[30] It is also responsible for access to the physical transmission medium and supports configurations where that medium is shared by multiple devices (for example, SONET) or dedicated devices (point-to-point T-1).

Layer 1 (Physical Layer) consists of the physical transmission medium used to convey electrical, optical, or radio-frequency signals at the bit level (i.e., digital ones and zeros) from the originating party to the terminating party. Layer 1 can be a twisted-pair, a copper facility, an optical fiber cable, a coaxial cable, a radio or free-space optical frequency, or another transmission medium.[31]

---

[29] The terms *transport layer* and *network layer* sometimes create confusion because the terms are frequently used to describe transmission *facilities* in discussions addressing traditional networks. In the OSI Reference Model, however, the physical facilities themselves comprise Layer 1; the "transport" and "network" layers referenced by the OSI model describe the functional *instructions* that are *communicated over* the physical elements of Layer 1.

[30] Layer 2 error detection/correction and flow control provide reliable transmission between two physical network elements (such as the optronics that exist at either end of a fiber span), while Layer 4 error detection/correction and flow control provide reliable transmission at a higher level in the hierarchy between hosts (such as two SIP phones). The lower layers (1 through 3) are sometimes referred to as the *chained layers,* because they are primarily focused on establishing connectivity between each individual link in a transport network; while the higher layers (4 through 7) are referred to as the *end-to-end layers*, because they are focused on achieving interoperability between the end points.

[31] Much of the discussion concerning broadband deployment has focused on fiber as the favored transmission medium at the physical layer. However, with new bonding and vectoring

Figure 5: Describing a Traditional Call Using the OSI Reference Model

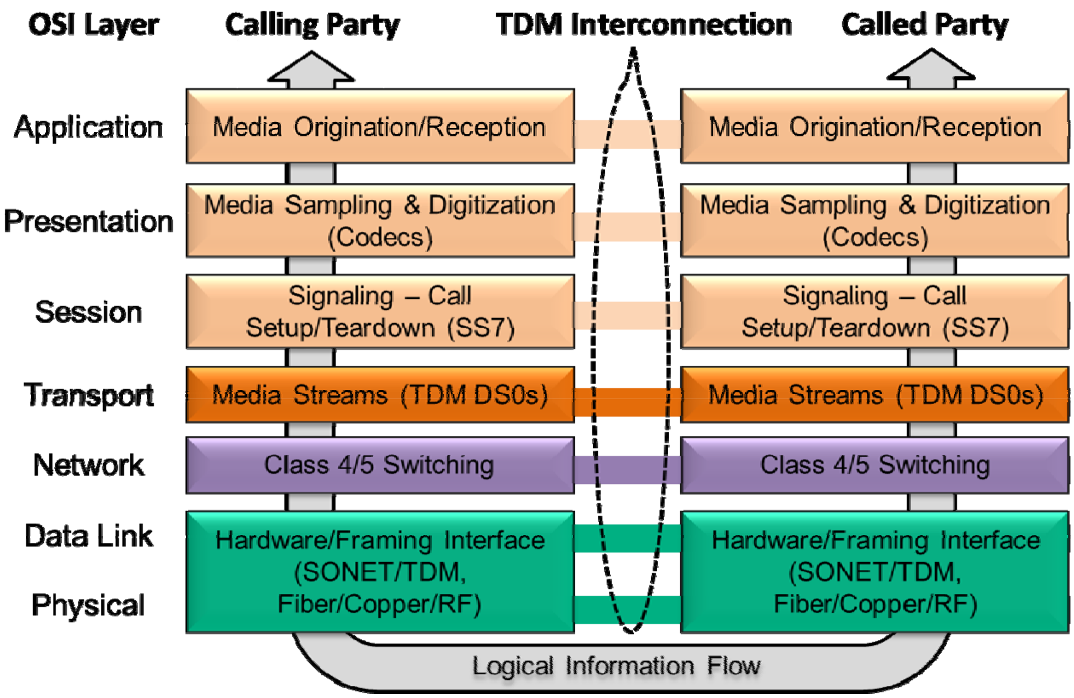

To begin a PSTN phone call, the calling party picks up a telephone, receives a dial tone, and presses the keypad to cause tones to travel down a subscriber loop to an end-office switch. These actions signal the application layer that the subscriber is making a telephone-call attempt. At this point, the end-office switch assigns a *dual-tone multi-frequency* (DTMF) *receiver* to the subscriber loop. The presentation layer interprets and collects the DTMF dialed digits, and then sends them to the switch's translation function. If the call attempt is being made to a destination that is not local to the switch, the session layer establishes the call using the signaling system to identify the destination switch and claim the resources needed to complete the call.

At this point, the signaling system notifies the called party's serving end-office switch of an incoming call. Additional signaling ensues between the transport and network layers of each switch to identify and reserve a *bearer channel*[32] between the two networks to carry the call. Once the called party answers, the signaling system notifies both parties to open the reserved

---

techniques, broadband is also being deployed on copper, and broadband wireless (4G) is actively being deployed. A critical feature of the IP architecture and NGN is to allow applications to ride seamlessly over a variety of physical layers without disruption.

[32] A *bearer channel* is a basic communication channel with no functionality other than transmission capability.

bearer channel; the datalink layer[33] and physical layer[34] then connect an audio path between the calling and called parties across the trunking network.

Figure 6: The Encapsulation of Functions in a Packet Architecture

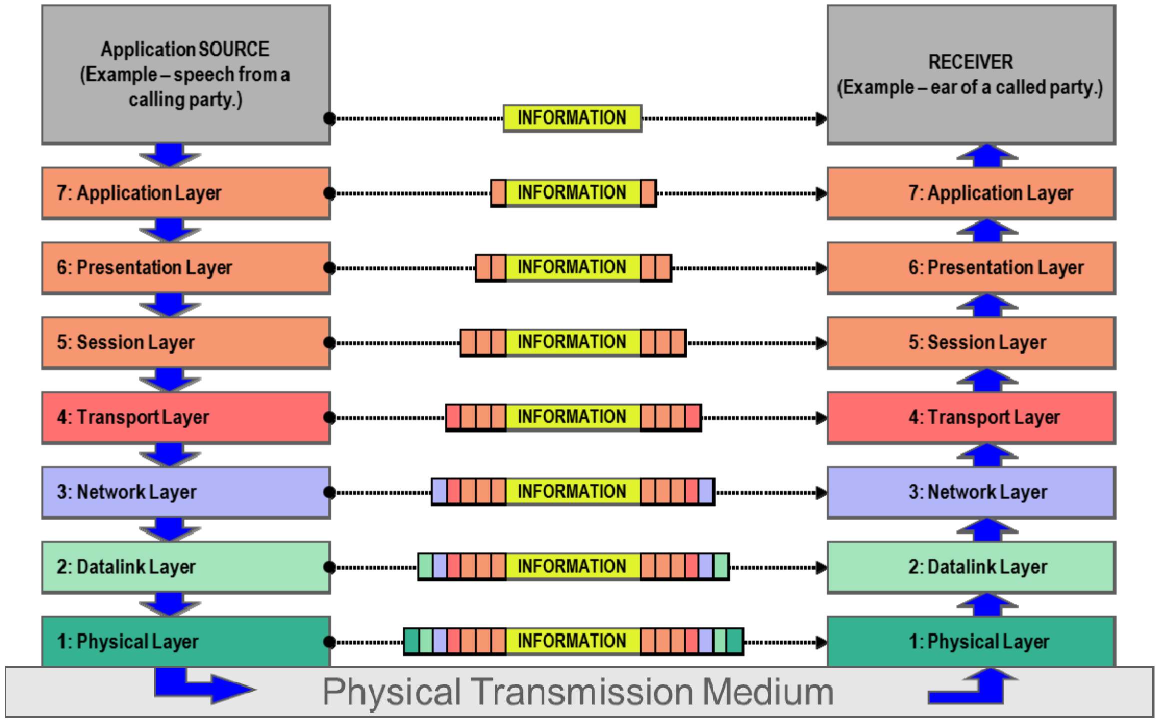

The routing of a telephone call using VoIP technology can be similarly mapped to the functions described by the OSI reference model. A key difference with VoIP is that much of the work performed by discrete systems in the circuit-switched world is performed in the packet-switched world by shared systems operating over a common network. In addition, as illustrated above in Figure 6, the different layers present themselves in a packet-switched VoIP environment as information or instructions encapsulated in defined headers, trailers, and extensions that wrap each individual packet of the information (or specially purposed packet[35]) being transmitted, rather than as a special-purpose system or the discrete elements typical in a TDM call.

---

[33] The datalink layer has the responsibility for reliably transmitting a frame containing the IP packet over the next layer down—the physical layer.

[34] The physical layer includes the actual medium used by the datalink layer, such as twisted-pair copper cable, coaxial cable, fiber-optic cable, or radio spectrum.

[35] An example of a specially purposed packet would be a packet that contains signaling information. Although most of our discussion focuses on packets that carry the phone conversation itself, the network will also be carrying packets that establish the session (such as the *SIP invite message* discussed below).

In a common VoIP-to-VoIP telephone call, a calling party picks up the receiver of a traditional telephone connected to an analog-to-IP adapter[36] that performs as a *user agent*.[37] Cable-based providers often provide the analog telephone adapters for their customers' use at home. In more sophisticated offerings, the user-agent function may be performed by an IP PBX system.[38]

Figure 7: Inter-carrier VoIP Call Interconnection

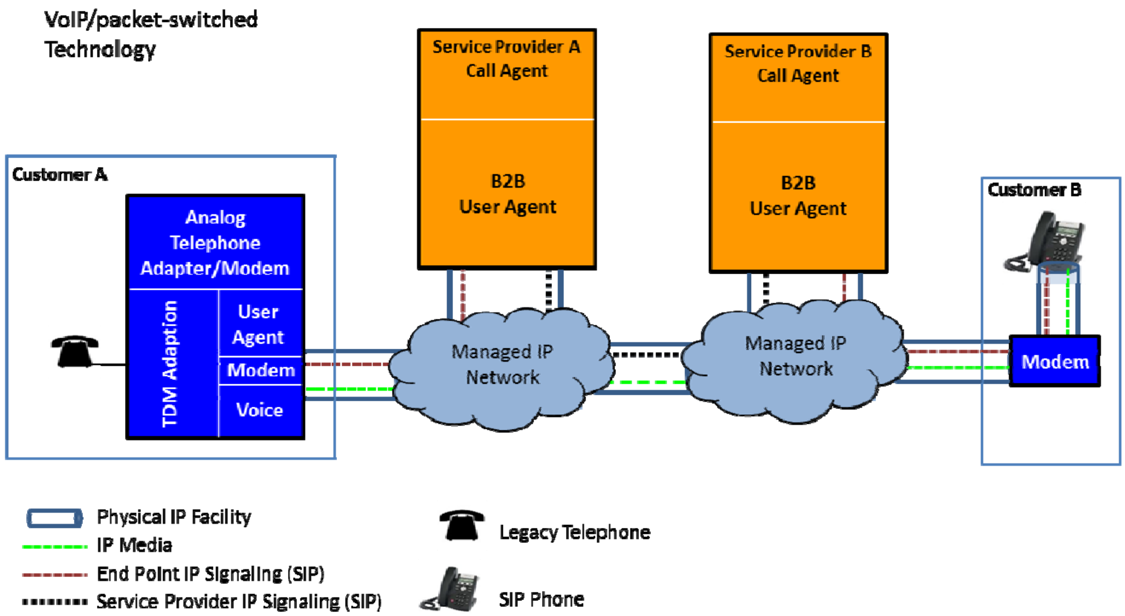

Figure 7 is a pictorial diagram of the network components and data "flows" comprising a call in which a customer of one VoIP service provider is calling a customer of a different VoIP service provider. Figure 7 illustrates the location and role of the various functions engaged in setting up a VoIP call. Figure 8 (below) illustrates the same process (i.e., establishing a single call) from the perspective of the role played by the each layer of the OSI model.[39] Together,

---

[36] Although the most common residential applications allow existing analog phones to connect to IP networks through an adapter at the customer's premise, many business offerings involve phones that are "VoIP phones" and act as the end point of a VoIP call.

[37] The term *user agent* refers to the collective SIP signaling activities associated with a VoIP endpoint in establishing, maintaining, or releasing a SIP session. It is also a term used to describe those same functions when they are housed in a SIP server.

[38] In other words, *user agent* refers to the *functionality*, while the ATA is the *device* that houses the functionality.

[39] Part of the challenge in understanding VoIP is the new vocabulary that must be learned (particularly in comparison to the TDM architecture with which most regulators are familiar).

these figures trace the progress of a VoIP call between two IP networks and illustrate the range of *physical* (best seen in Figure 7) and *logical* (Figure 8) interactions needed for IP interconnection.

Figure 8: Describing a VoIP Call using the OSI Reference Model

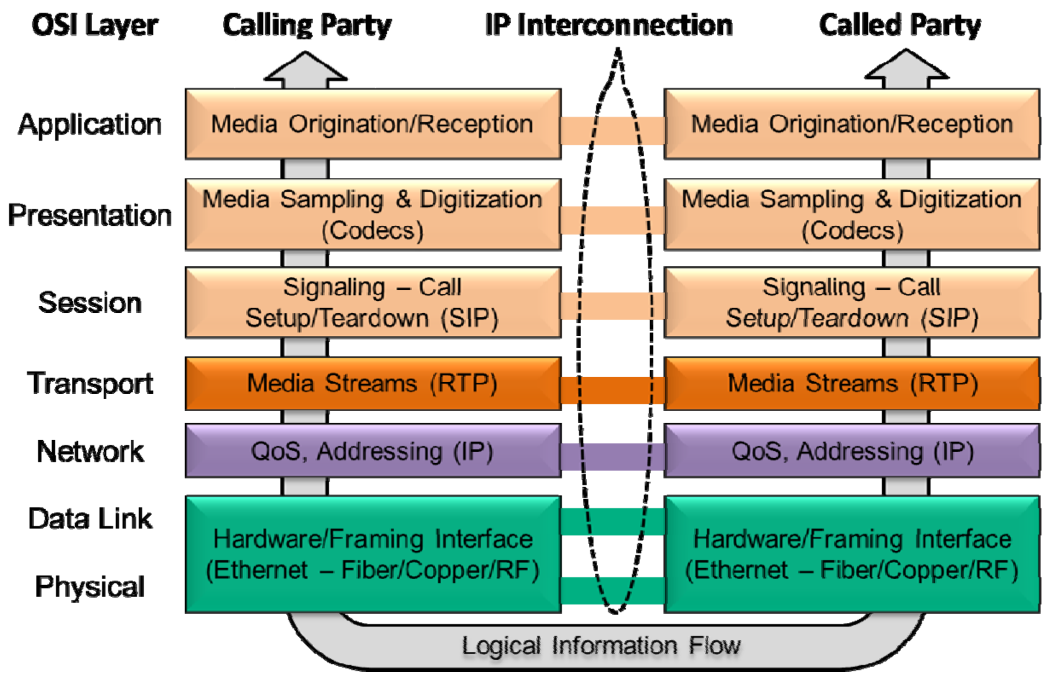

When the telephone goes off-hook, the adapter provides a dial tone and prepares to accept digits. In effect, these actions constitute the application layer—in other words, a telephone call attempt is in process. At this point, the adapter interprets the signals and, after collecting the dialed digits, formats a SIP message to invite the called party to a session.[40] The interpretation of the dialed digits and the conversion of the media to a digital format is the presentation layer. A SIP message typically specifies the session parameters, such as the preferred and supported codecs, data-sampling rate, media protocol, and port assignment.[41] The adapter then sends the SIP invite message across the IP network to the calling party's VoIP provider's SIP call agent.

The end points of a session occur at user agents, which are located in a variety of devices. The simplest to visualize (as discussed before) is the user agent that resides in the ATA at the

---

[40] What we would refer to as a "telephone call."

[41] These parameters are defined using the *Session Description Protocol* (SDP). SDP is a protocol used to describe the technical characteristics of the session. The SDP is encapsulated within a SIP message (i.e., it is a sub-protocol of SIP). Every *SIP invite* (which initiates the phone call) contains an "offer SDP" that requests how the Session should be configured. The response to the invite will either accept these parameters or initiate a negotiation to achieve agreement.

typical residential home (or within the SIP phone itself). The called party is also associated with a user agent that defines that end point. Directing how these user agents interact are *call agents*. Call agents are the "control centers" (lay term) that receive signaling messages from one SIP end point and convey those messages downstream to the destination network or end point. Because call agents both receive and send a given message, they "talk" to other user agents through a configuration referred to as a *back-to-back user agent configuration* (B2BUA).[42] As shown in Figure 6, when a VoIP call travels between two networks, two call agents (one for each network) are used to establish the session.[43]

Although many commenters have likened the call-agent function to the function of the traditional end-office switch, we intentionally avoid using the term "soft switch" to describe the call-agent function, because the term implies that the call agent is a direct replacement for the Class 5 end-office switch in a TDM/circuit-switched network. This is misleading, because the interactions between the user agents and call agent involve *signaling* functions that are communicated using the specially purposed packets described earlier and also called "SIP packets," but do not involve *bearer* channels to carry voice conversations, as a Class 5 switch would provide. Therefore, these interactions are more closely analogous to an SS7 signaling network component than to the Class 5 end-office switch of a TDM/circuit-switched network.

After the adapter sends its invite message to its call agent, the session layer is engaged to accomplish a variety of tasks. The call agent performs a directory lookup and determines the destination network of the called party. The call agent (again through the B2BUA) then signals the destination network that there is an incoming call attempt. The destination network signals the called party with the appropriate session-description parameters, and, once all session parameters are negotiated, the called-party user-agent client will initiate ringing.

When the called party answers, the destination network signaling system notifies the originating network that the called party has answered. For voice conversations, the packets that transport the actual voice audio (RTP packets) are formatted at the transport layer using *real-time transport protocol* (RTP).[44] Upon detection of RTP from the transport layer (speech payload),

---

[42] To be precise, when a user agent *initiates* a call, it is acting as a "user-agent *client*," and when *receiving* a call, it is acting as a "user-agent *server*." Because of this distinction, when a call goes between two networks, a *back-to-back user agent* configuration applies (as shown in Figure 6). That is, when the call agent receives a SIP invite from its customer, it is acting as a user-agent server. Then, when it recognizes that the call is destined for another network, it must initiate an invite (and thus act as a user-agent client). Because it is performing this dual role (acting as a user-agent server receiving an invite and then acting as a user-agent client to issue an invite), it is referred to as a *back-to-back user agent* (B2BUA).

[43] Depending upon the architecture of the VoIP provider network, the SIP call agent may be a session border-control function, call agent, call-session control function, or other device or function.

[44] Note that RTP packets, which carry the voice conversation, are formatted differently from the SIP packets, which provide the signaling necessary to establish, maintain, and release

the IP network layer provides IP packet assembly and routing, encapsulates the RTP payload in an IP packet, determines and inserts the destination IP address of the recipient, and sends the assembled IP packet to the datalink layer. At the datalink layer, the IP packet is further encapsulated in a frame specific to the requirements of that layer (example: Ethernet) and then sent for transmission across the network facilities. These network facilities (as with the TDM example) define the physical layer (of wired or wireless facilities). At the receiving end, the entire process is repeated in reverse.[45]

As the above discussion makes clear, each of the sending-network functions involved in a VoIP call must be able to communicate with its corresponding function on the receiving network in order for the service to work. Physical interconnection is not enough; intelligent interconnection must be in place that permits interaction between all functions comprising the entire seven-layer architecture of the OSI reference model.

For instance, the SIP signaling *between* the provider networks conveys call setup, progress, and disconnection information between the networks for all call traffic—just as SS7 performs that same role in a legacy network. These SIP packets need to be identified and treated as high-priority network traffic by both providers, then given the appropriate network resources to maintain availability and performance. Likewise, the RTP media flows (that is, the packets that contain the actual conversation) require network performance to ensure that the call will meet quality levels equal to those provided by the TDM technology used in legacy networks.

Although extensive interaction must occur for IP interconnection to succeed, each network must protect itself as well. For instance, network security can be compromised if the actual routable IP addresses of the service provider devices are disclosed. To protect each network, carriers rely on *session border controllers* (SBCs) to govern their interconnection.[46] The SBC can secure both the signaling and media traffic between a trusted network and an untrusted network. It is normally deployed between the customer and the service-provider core network, as well as between individual service-provider networks.

Figure 9 depicts a configuration of two service-provider networks using SBC functionality for IP interconnection at the border of each provider's network.

---

the call. The RTP packets will even travel a different route to the destination than will the SIP packets, which involve the call agent to establish the call (but which do not process the media).

[45] In a typical VoIP call, this cycle is repeated fifty times per second at each end of the call.

[46] For simplicity, we refer to a session border controller (SBC) as a *device* that is deployed at the boundary of two networks. Although standalone SBCs do exist, it is more valid to refer to an SBC *functionality* that can be located elsewhere in the network while performing the same role.

Figure 9: Inter-carrier VoIP Call Interconnection Using SBC Functionality

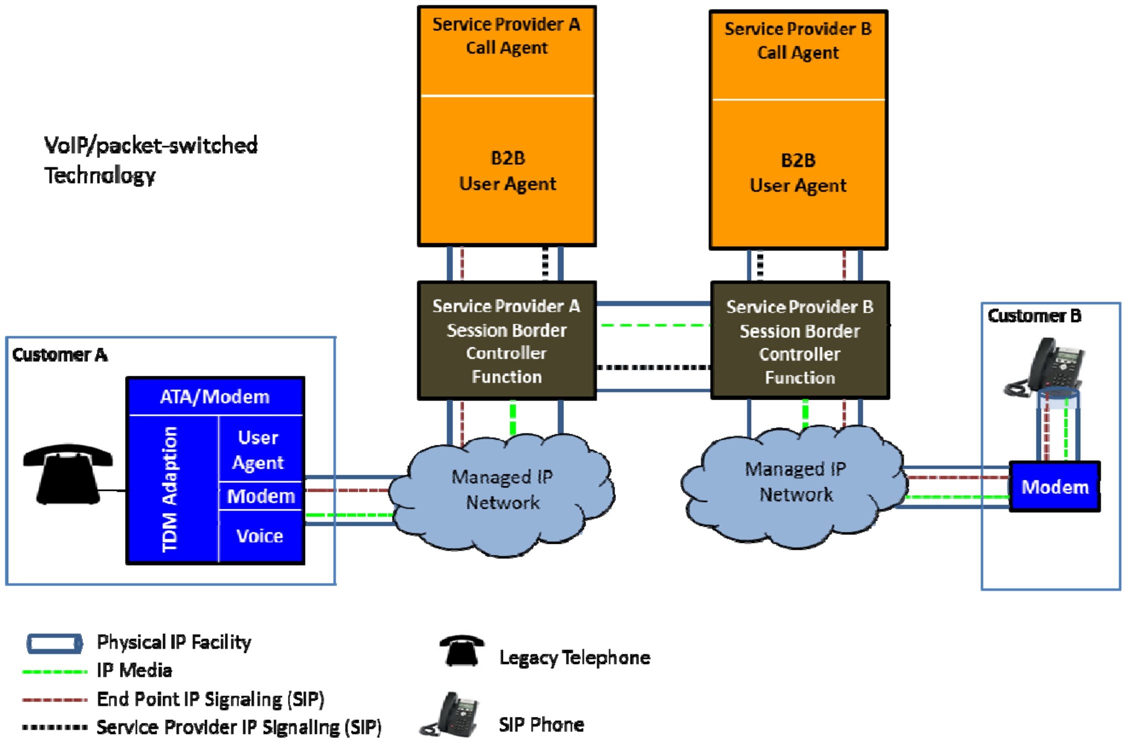

The SBC addresses a variety of issues to enable networks to interoperate on a secure basis. The SBC hides the actual IP addresses of network devices, provides codec transcoding between disparate codecs, and acts as a firewall by protecting the network against cyber-attacks. The SBC also acts as a load-balancing device for multiple service-provider call agents and off-loads the call agent of error handling for signaling messages and unauthorized call attempts, as along with other functions.

V. Conclusion

The future of the PSTN has never been brighter. Its migration to IP technology will create a flexibility that the rigid requirements of the traditional TDM architecture foreclose. For the full benefit of these networks to be realized, however, they must interconnect with one another on an intelligent basis that assures that quality, security, and new features can be preserved end to end, no matter which provider serves the end point.

In the discussion above, we have provided a basic primer that focuses on the architectural components of IP interconnection. But the issues that need to be addressed do not end here.

The central questions for IP-to-IP interconnection in the 21st century will be little different (conceptually) from the questions surrounding traditional interconnection in the late 20th century (or late 19th century, for that matter): *where* networks should interconnect, *how* *many* interconnection points are needed, *how* they should interconnect, and *what compensation* (if any) should apply.

When competitive interconnection was first introduced by the Telecommunications Act of 1996,[47] new entrants were deploying networks against the backdrop of the ubiquitous preexisting networks of incumbent carriers. Although the competitive providers' interconnection rights were defined to include "any technically feasible point within the [incumbent] carrier's network," the FCC's implementing regulations understandably focused on interconnection points defined by the existing architecture (such as the line side of a local switch), coupled with the right to collocate in the incumbents' buildings to establish physical connections.

IP interconnection not only involves a different technology but also arises at a different point in the industry's maturation. Its technological differences—as well the differences in the market today—will create different interconnection opportunities as well as needs.[48] There are no tandems—and, as such, no tandem serving areas or LATAs—in IP networks. As a result, we expect that fewer points of interconnection will exist between IP networks that will cover much larger areas.[49] Moreover, IP interconnection could take place at the carrier hotels and other venues where conventional Internet traffic is exchanged today (i.e., at recognized points beyond the wire centers of the incumbents).

Addressing these questions is beyond the scope of this paper (which is focused on creating a basic technical understanding of how managed IP networks interconnect). As noted earlier, the FCC has begun a process to address the rights and obligations (if any) that will apply to IP interconnection, and the initial filings in that proceeding have identified a large division of opinion.[50] We do not attempt here to describe, much less bridge, that division. We do, however, encourage the reader to develop a broader understanding of related issues by reviewing the status of this issue in other countries.[51]

---

[47] Telecommunications Act of 1996, Pub. LA. No. 104-104, 110 Stat. 56 (1996), available at http://transition.fcc.gov/Reports/tcom1996.pdf

[48] Nothing in this paper should be construed as suggesting that the authors believe that incumbent local exchange carriers no longer enjoy market power. The principal source of their market position—a preexisting customer base and infrastructure—provides benefits to their IP networks that are comparable to the benefits enjoyed by the existing TDM networks. This issue, however, is beyond the scope of this paper.

[49] In Canada, for instance, it is proposed that two POIs per province (presumably, with the second for redundancy) may be adequate. *CRTC IP Interconnection Order* at ¶53.

[50] *See* Comments filed in Response to *IP Interconnection FNPRM*, February 24, 2012.

[51] We have attached a short list of useful decisions and documents in Appendix B.

Finally, a comment (perhaps a caveat) is in order concerning the ongoing evolution of technology. Although the paper has focused on interconnection between *wireline* networks, *wireless* services and networks are becoming increasingly important as the primary vehicle for voice communications. As these networks move to 4G-based technologies, an increasing percentage of their traffic will be IP, because 4G standards abandon circuit switching in favor of packet switching (IP) for *all* traffic, including voice telephony. Consequently, interconnection with (and among) wireless networks will be a critical part of the network of the future.

In addition, just as the PSTN is evolving from a circuit-switched architecture into a packet-based architecture, an evolutionary change in packet technology is also underway. A full discussion of this architectural change is beyond the scope of this paper, but it is important at least to reference the emergence of networks that incorporate a network control layer known as the IP Multimedia Subsystem (IMS). IMS was originally designed by the wireless-standards body 3rd Generation Partnership Project (3GPP).[52] The goal of IMS is to establish standardized interfaces to create a horizontal control layer in the network that would enable any number of multimedia applications (for example, voice, video, and so on) to run without regard to the underlying transport network. In this way, the IMS architecture is intended to facilitate wireline/wireless convergence.

It is clear that the future will require IP-to-IP interconnection, and that technological advances that improve the management capabilities of modern packet networks will require increasingly sophisticated levels of interaction to ensure interoperability and end-to-end service quality. The transition from TDM technology to IP will have far-reaching consequences and requires that regulators become familiar with a very different network topology, involving a very different (and more extensive) vernacular than exists in the TDM environment. This paper is only the beginning of that educational process.

[52] The 3rd Generation Partnership Project (3GPP) is a collaborative effort between groups of telecommunications associations, known as the Organizational Partners. The organizational partner representing North America is the Alliance for Telecommunications Industry Solutions (ATIS).

Appendix A – Glossary

Analog Telephone Adapter	A device that turns a standard telephone into a VoIP device by providing the necessary translation and signaling to be a recognized endpoint of a VoIP call.
Back-to-Back User Agent (B2BUA)	Refers to the combined functionality of a SIP user-agent client (UAC) and a SIP user-agent server (UAS) such that it can not only *pass* SIP messages but also regenerate them, modify them, or create new ones as part of the call session. B2BUAs are found in intermediary devices such as a call agent or an IP-PBX.
Bearer Channel	A bearer channel is a basic communication channel with no functionality other than transmission capability.
Call Agent	The central signaling function (or device) that controls the establishment, maintenance, and release of telephone calls in a VoIP service. It also controls registration of SIP devices such as customer devices (for example, a SIP phone), media gateways, other interconnected VoIP networks, and so on. In lay terms, the call agent provides signaling control for the complete VoIP infrastructure within its domain.
Codec	A device or computer program capable of encoding or decoding a digital data stream or signal. The word *codec* is formed by the combination of "**co**mpressor-**dec**ompressor" or, more commonly, "**co**der-**dec**oder."
IMS	IP Multimedia Subsystem. IMS is a framework originally designed by the wireless-standards body 3rd Generation Partnership Project (3GPP) to document functions and establish standardized interfaces to create a horizontal control layer in the network that would enable any number of multimedia applications (for example, voice, video, and so on) to run without regard to the underlying transport network.
ITU	International Telecommunications Union. The ITU is the specialized agency of the United Nations responsible for information and communication technologies and the establishment of worldwide standards.

Appendix A – Glossary

Managed VoIP	A Voice over Internet Protocol service provided over a network that applies a segmentation or prioritization scheme to the VoIP service in order to ensure quality of service standards can be met.
Media Gateway	A translation device that converts a digital media stream (such as a voice conversation) between disparate networks, such as an IP network and the traditional TDM network.
Modem	A modem (**mo**dulator-**dem**odulator) is a device that modulates an analog carrier signal to encode digital information, and also demodulates such a carrier signal to decode the transmitted information.
NGN	Next-generation network. A next-generation network (as defined by the ITU) is a packet-based network that can provide services, including telecommunication services; and is able to make use of multiple broadband, quality-of-service-enabled transport technologies; and in which service-related functions are independent from underlying transport-related technologies.
OSI Model	The Open Systems Interconnection Model was created by the International Organization for Standardization (ISO) and is the common reference model used to describe the organization and role of different layers in a packet network.
Over-the-Top VoIP	Over-the-Top VoIP refers to a Voice over Internet Protocol service that relies on the public Internet for some of its transmission and, as such, is vulnerable to the Internet's undifferentiated packet congestion policies.
PSTN	Public Switched Telephone Network.
RTP	Real-time Transport Protocol (over UDP) is the protocol used at the Transport Layer in a packet network to provide end-to-end delivery services including payload type identification, sequence numbering, time stamping, and delivery monitoring for real-time applications such as VoIP. Generally, the public Internet uses TCP (Transmission Control Protocol) as the transport-layer protocol for Internet traffic. The Real-time Transport Protocol (over UDP) is used by applications that require timely delivery, but can tolerate some packet loss to achieve the goal. TCP is not

	normally used for a real-time service because TCP favors reliability over timeliness.
SBC	A Session Border Controller is a device (or functionality) that provides security and protocol conversions at the border between two networks. SBC functionality can hide a service provider's network topology by providing the necessary translations without disclosing proprietary information.
SDP	Session Description Protocol ("SDP"). SDP is a protocol used to describe the media characteristics of the session. The SDP is encapsulated within a SIP message (*i.e*., it is a sub-protocol of SIP). Every SIP invite (which initiates the phone call) initiates an SDP "offer/answer" sequence that determines how the Session should be configured.
Signal Transfer Point (STP)	A device that acts as a routing hub for messages between signaling points on a SS7 signaling network.
Signaling Gateway	A translation device that interconnects the signaling information between two different networks. For instance, a signaling gateway would translate SIP messages addressing call control function into System Signaling 7 messages used to control the PSTN.
SIP Endpoint	A function or device that can create and respond to SIP messages (i.e., it contains a UAC and a UAS).
SIP Server	A device that provides functionality to assist in the establishment, maintenance, release, or enhancement of the attributes of a SIP session.
TDM	Time Division Multiplexing. TDM is a technique that enables multiple date streams to share the same facility by assigning each a unique time slot. In effect, TDM enables different digital signals to "take turns" being transmitted.
User Agent (UA)	The SIP signaling functions associated with a VoIP endpoint in establishing, maintaining, or releasing a SIP session. It is the term used to describe those functions whether they are housed in a SIP server or a customer device.
User-agent Client (UAC)	When a user agent *initiates* a SIP request (such as a SIP Invite message to establish a phone call), it is acting as a

	user-agent client. The role of the UAC lasts only for the duration of that transaction. In other words, if a piece of software initiates a request, it acts as a UAC for the duration of that transaction. If it receives a request later, it assumes the role of a user-agent server for the processing of that transaction.
User-agent Server (UAS)	When a user agent *receives* a SIP request (such as a SIP Invite message to establish a phone call), it is referred to as a user-agent server. This role lasts only for the duration of that transaction. In other words, if a piece of software responds to a request, it acts as a UAS for the duration of that transaction. If it generates a request later, it assumes the role of a user-agent client for the processing of that transaction.
User Datagram Protocol (UDP)	User Datagram Protocol is a communications protocol used with RTP at the transport layer that provides port numbers to help distinguish between different user requests and a checksum capability to verify that the data arrived intact when messages are exchanged between end points.
VoIP	Voice over Internet Protocol. The term indicates that the Internet Protocol suite is being used to create and route a voice call using packet technology. The term does not require (or preclude) that the packets be routed over the public Internet.

Appendix B – Recommended International Resources/Decisions

1. Final Report on IP interconnection, European Regulatory Group,[53] Project Team on IP-Interconnection and NGN ERG (07) 09

This document describes the significant evolution taking place in (IP) interconnection and in the networks of most European operators, particularly PSTN incumbent networks, reflecting the development of next-generation networks (NGN) and the massive adoption of IP-based services such as VoIP.

http://erg.eu.int/doc/publications/erg_07_09_rept_on_ip_interconn.pdf

2. ERG Common Statement on Regulatory Principles of IP-IC/NGN Core – A Work Program towards a Common Position

This document presents the (then) current situation with regard to IP interconnection in Europe and outlines how the development of NGNs may affect market regulation. This document analyzes the effects this evolution might have on interconnection regimes and develops some general principles with regard to regulatory treatment of IP interconnection and interoperability.

http://erg.eu.int/doc/publications/erg_08_26_final_ngn_ip_ic_cs_081016.pdf

3. Canadian Radio-television and Telecommunications Commission, Telecom Regulatory Policy CRTC 2012-24, Network Interconnection for Voice Services

In this decision, the CRTC decides that it is in the public interest to establish a set of principles to facilitate IP voice-network interconnections between network operators while allowing market forces to shape the details of the arrangements with a structure of oversight and expected time frames.

http://www.crtc.gc.ca/eng/archive/2012/2012-24.htm

4. Telecom Decision CRTC 2009-139 CRTC Interconnection Steering Committee – Consensus Items – IP to IP Interconnection Guidelines

This decision approves a consensus report of the Network Working Group (NTWG) of the CRTC Interconnection Steering Committee (CISC) establishing the technical guidelines for Internet Protocol (IP)-to-IP interconnection.

http://www.crtc.gc.ca/eng/archive/2009/2009-139.htm

---

[53] The European Regulatory Group (ERG) was set up by the European Commission to encourage cooperation and coordination between national regulatory authorities and the European Commission. The ERG has been formally replaced by the Body of European Regulators for Electronic Communications (BEREC). http://erg.eu.int/