From POTS to PANS – A Commentary on the Evolution to Internet Telephony

Christos A. Polyzois, K. Hal Purdy and Ping-Fai Yang (AT&T Labs Research)
David Shrader (Telcordia Technologies, formerly Bellcore)
Henry Sinnreich (MCI)
François Ménard (Mediatrix)
Henning Schulzrinne (Columbia University)

March 26, 1999

Abstract

We describe some of the challenges and problems as we move from the world of plain old telephony services (POTS) to the promise of Internet-based pretty amazing new services (PANS). We present a set of ideas that relate the architectural evolution of networks and services to the underlying changes in technology. After examining the motivations behind the drive to Internet multimedia communications, we discuss the distribution of service intelligence in the future communications architecture. As with any period of evolution, thorny problems arise regarding interworking between the two network architectures. A brief exploration of the Intelligent Network (IN) is presented as a solution to address interworking. Finally, we examine how the transition to an Internet-based system might affect how transport and services are billed for.

1 Introduction

The transition from circuit-switched technology to a packet-based infrastructure affords us the opportunity to revisit many of the fundamental design choices that have governed the architecture of the world's most widely used telecommunications service, telephony. First, we look at some of the motivations for Internet-based telephony services. We believe that the incentive for a new infrastructure has to be more than just cheap phone calls. Rather, we hope that Internet telephony can deliver pretty amazing new services (PANS), as described in Section 2. We then take a closer look at the central architectural assumption of the current plain old telephone service (POTS), namely that end systems are dumb and that almost all services are provided by computing power inside the network. In Section 3 we discuss the trade-offs involved in locating services either at the core or the perimeter of the network.

Given the reluctance to just discard billions of dollars sunk into the existing circuit-switched network and the relative immaturity of Internet telephony products, it will take many years to transition to a purely packet-based environment. Thus, there is keen interest in methods that allow the existing phone system to connect to "islands" of Internet telephony, while maintaining the functionality of the existing network. Section 4 discusses some of the issues. Finally, Section 5 describes options for billing for Internet telephony. In this brief survey, we do not touch upon the crucial issue of providing quality-of-service appropriate for telephony services, as these are discussed at length in the literature [1].

It should be noted that we follow convention in using the term Internet telephony, which by its root indicates voice communications. However, unless we are discussing the interoperation with the existing telephone network, this encompasses all media types, as discussed below in Section 2.

2 Motivations for Internet Telephony

In the near term, Internet telephony is motivated by lower marginal costs. Note that the total costs depend on the Internet telephony mode used, whether it is "hop-on", i.e., using a regular telephone with a local or 800 number or true end-to-end packet telephony. In the latter case, only the marginal costs are lower, i.e., assuming that the caller owns a computer and subscribes to a flat-rate ISP for reasons other than Internet telephony. International settlement rates and domestic U.S. local access charges do not apply to IP telephony.

In the longer term, integration with other Internet services and more rapid service creation motivate Internet telephony. In particular, it may become feasible to create services that appeal only to small groups of subscribers. Also, the the transport of bits can be separated from the provision of services, similar to how today there are specialized companies providing "portal" services, web-based email or web hosting, without necessarily owning the Internet network infrastructure. Initially, services will probably be the familiar ones, such as call-forwarding, call waiting or caller id, although each of these has somewhat different properties in an Internet environment. However, other, Internet-specific features may emerge. An example may be time-of-day routing interfaced with a web calendar. Calls could then automatically be transferred or refused during meetings, for example. Integration with web services may be as simple as that on calling a company, one gets back a web page, listing the different destinations, rather than an interminable voice menu. In general, the integration allows a back-and-forth between the asynchronous, "canned" interaction of a web page and live human interaction.

Integration with email and media-on-demand also creates new opportunities. For example, unanswered calls could generate an email to the callee, containing a SIP [2] or H.323 [3] URL for one-click call returns. Instead of a separate voice mail system, calls can be recorded by standard Internet multimedia tools, with a URL pointing to the recording sent by email or added to the callee's web page.

Moving beyond carrying voice, Internet-based "telephony" applications can readily include text chat, video, shared applications, or games, as inviting someone for a phone conversation is not all that different, technically, from inviting that person to a networked game of chess.

3 Intelligent Endpoints vs. Network-Provided Services

A PSTN phone can only generate a small set of signaling events and tones (off/on-hook, "touch tones" or DTMF); it cannot receive and process signaling of any sophistication. Signaling is received in band, i.e., in the same voice channel as the phone call, and processed by the human using the phone (e.g., flash the switch hook as a response to a call waiting tone). The results of this limitation are that PSTN phones are considered "dumb" devices and that switches (transport elements) act as proxies for the implementation of several services such as automatic call return services and other custom calling services. In contrast, packet phones such as IP-based phones are able to receive and process signaling messages automatically. IP phones send signaling out-of-band, as a separable set of IP packets. Consequently, the endpoint can execute programs that reside on it or on a remote server; there is no need to involve a transport element. The ability of an endpoint to receive and act on signaling is the fundamental property that makes the endpoint intelligent and enables the complete separation of the service functionality from the task of transporting voice bits.

The separation of service from transport is very desirable. The attempt of the Intelligent Network to isolate PSTN programmability within computer nodes attests to this importance. ISDN has similar goals for the end system; it made the choice that user-to-network interfaces (UNI) and protocols and network-to-network interfaces (NNI) are different.

However, the lack of programmability in PSTN endpoints has prevented complete separation of transport and service in the PSTN. Even for ISDN, circuit-based bearer (B) and packet-based data (D) channels are closely tied together. There is direct packet-based signaling between the end systems.

We examine the ramifications of such separation in packet networks. The design of networks and network elements becomes simpler, since packet transport and service can evolve independently. For example, most packet routers are optimized hardware devices with fixed functionality, while services are often implemented in general-purpose computers. Services can be deployed faster, since they do not involve changes to transport elements. Furthermore, services can be deployed by third parties, not just by facilities-based carriers. The reliability of networks may also benefit from the separation, since faulty service logic will not affect the transport or other services (assuming proper design). Services can use general purpose computing facilities (instead of switches) and realize savings by riding the technology curve of those facilities and reducing the cost of their operations support aspects. With the separation of packet transport and services, services can be handled almost anywhere in the network, without forcing the packets to take detours from the shortest path between the caller and callee. (This separation is impossible in an ISDN environment, for example.) For example, the SIP call setup only takes three or four packets, so that the service provider can be far away from the subscriber without adding undue burden to the network.

There has been a lot of debate on the placement of services. The ability of packet phones to store programs has led to the view that all of the intelligence should be placed in the packet phones, so that the network's role is reduced to carrying packets. A more likely scenario is for some intelligence and state to reside on servers that are endpoints with respect to the transport network such as the Internet but are viewed as part of the "network" with regard to the service offering. In other words, these servers are hosts rather than routers in the Internet architecture, but are owned and operated by entities other than the parties making a phone call.

There are several reasons for this path. First, there is the need to manage and apportion transport resources, to record their usage, to require payment (by controlling resource allocation) and to prevent fraud; this task is better left to a trusted network server. Second, there is need to keep state in the network even while an endpoint is not connected or accessible; for example, such state maybe necessary for address translations or call dispatching (e.g., call forwarding to voice mail). The provision of emergency services is a third reason for having servers handle calls, as they may be needed to insert verifiable subscriber street address information and direct E911 (enhanced emergency) calls to the nearest PSAP (Public Safety Answering Point) such as the local fire department in small communities and a centralized emergency dispatch center in larger cities. [1] Emergency use also argues for low-complexity, low-power end systems, as either battery-powered or line-powered devices are limited to between 200 mW and 2 W of power consumption. (This value is based on the current of about 20 mA supplied by the phone company, and the capacity of about 600 to 2000 mAh of typical cordless phone or laptop batteries. The current phone system is designed to run on batteries for eight hours. Larger batteries located at the network termination point could be used to increase the power budget or the available time.)

Fourth, version control and new service deployment may be facilitated operationally by storing the programs implementing services in relatively few servers rather than millions of endpoints, although Java applets and JavaScript have shown that small programs can be distributed automatically and on demand to end systems without explicit upgrades. Finally, during the transition phase, interoperability with the PSTN may require access to the PSTN intelligent nodes, which may be better left to secure servers that are viewed as part of the telephone network.

Beyond the location of where services are implemented, we can distinguish whether end nodes are forced to use network servers or whether they "voluntarily" make use of these services. The various media gateway control protocols such as the Media Gateway Control Protocol (MGCP) [5] require the end system to be controlled by a network server. Peer-to-peer protocols like H.323 and SIP *allow* the use of servers – gatekeepers [6] and proxy/redirect servers [7], respectively – but also enables two end systems to connect

[1] Enhanced 911 services routes the call to the appropriate emergency service, based on the caller's address, and allows the emergency operator to see the caller's address in real time.

directly, without help of a third party. There have been proposals to couple the use of network servers to permission for voice packets to pass a firewall, for example, enforcing the use of servers even with the peer-to-peer protocols. Since telephone companies derive significant revenue from services such as caller-ID [2] or call waiting that come "for free" with peer-to-peer protocols, they also have a monetary incentive to retain control over these services.

Requirements by U.S. federal law enforcement agencies under the CALEA (Communications Assistance for Law Enforcement Act) provisions may lead carriers to enforce that calls traverse a carrier-provided server. Without this, "pen register" recordings indicating the call patterns of a suspect may be difficult to obtain. On the other hand, the ease of use of proxy servers will greatly simplify the making of anonymous phone calls. (Pen registers record the numbers dialed by a suspect, but not the actual phone conversation.)

The considerations above show that it is hard to make blanket statements as to where telephony-specific functionality should be placed in the network. Often, factors other than technology guide these decisions. That suggests that network design should offer flexibility in the placement of services to avoid having to start over at some point in the future.

4 Interworking with Existing PSTN Services

As Internet Telephony is introduced, there will be a need to interwork with the existing PSTN for some time. Some telephone companies (such as Qwest or 1WorldCONNECT)[3] are now using Internet Telephony to transport calls between separate traditional telephone networks. These companies are simply using the Internet as an alternative to the historical network of trunks and switching systems. The users on either end of a voice call established by these carriers are unaware that the call has traversed a packet IP network. The carriers are finding that although standards are not necessary for this type of operation, it is useful to develop some Internet standards in order to obtain equipment from different vendors that will interoperate. The next step in Internet telephony is to consider the interworking of calls and services between users of the PSTN and users of the Internet. In this scenario, for a basic two-person call, a call is placed from a user on a regular telephone on the telephone network to a user on a computer on the Internet. Below, we propose how the part of the PSTN infrastructure known as the Intelligent Network can be used to solve the problem of handling these calls. Additional perspectives on interoperation can be found in [5].

4.1 Intelligent Network Architecture

The Intelligent Network is a service infrastructure in the PSTN that provides a service provider with capabilities to provide custom services. The ITU Q.1200 series of recommendations [9] defines Intelligent Network architecture and protocols. In the United States, the Advanced Intelligent Network defined by Telcordia Technologies, formerly Bellcore, in its AIN Generic Requirements [10] provides the same network infrastructure and tailors the interface protocols to meet the needs of US service providers.

The basic IN architecture is shown in Figure 1. The IN components shown are the Service Switching Point (SSP), the Service Control Point (SCP), and the Intelligent Peripheral. A packet network consisting of Signaling Transfer Points (STPs) connects them. The packet network uses another international standard known as Signaling System No. 7 [11].

The SSP is integrated with a conventional telephone switching system. Using an abstraction of the basic call setup process, the SSP defines discrete points at which a service can interact with the call and modify its basic handling. The logic for an IN service is a program located on a remote computer platform known as

[2] Caller ID alone represented $166 million in revenue for SBC in 1995.

[3] A listing of next-generation telephone companies can be found at http://www.pulver.com/nextgen/. A taxonomy of Internet telephony modalities is provided in [8, 4].

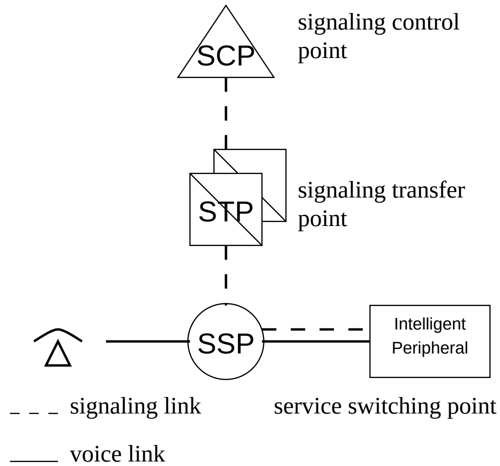

Figure 1: Intelligent Network Architecture

an SCP. An application protocol is defined between the SSP and the SCP to allow the service to instruct the SSP on desired processing of a call. For more information on how services can use the IN protocols, consult [12]. IN is the first step in an evolution towards having service intelligence resident in a network server.

4.2 Addressing and Routing Calls to Internet Users

Numbers are used in the telephone network to represent parties on a call. The routing structure of the network makes use of this digit representation of the end points of a call. Over the last few decades the telephone network has become highly optimized in its use of digits to represent users. ITU recommendation E.164 [13] and others determine the format of telephone numbers.

The Internet on the other hand, can use more meaningful names to address users. For example, to address the current president of the United States, one might use `president@whitehouse.gov` to direct the call. Unlike a telephone number, this email-like address identifies users by name and organization or network provider. To allow software to make a call to such an email-like address [4], we need to add a protocol identifier to make it a complete uniform resource identifier (URI). as in `sip:president@whitehouse.gov`.

Through the use of the Internet Domain Name System (DNS), the host name in the URI is translated to an IP address. Such addresses have the advantage that they do not depend on the location of the owner. For organizationally-based names such as the President's address, the address offers "local number portability" in that it does not change if the White House switches network providers. However, they do require their owners to change addresses if they move to a different organization. For most consumers, changing network providers means changing addresses, making this scheme less attractive. We can speculate that in the future, domain names will become much cheaper than the current $50/year (in the U.S.) so that families may own their name as a domain name – or at least those families blessed with a less common surname. In addition, professional organizations such as the IEEE or alumni organizations may offer permanent identifiers that are independent of the network provider.

Note that we can use a single identifier across all communications modalities [4], from email, to fax, pager and telephone – or as many identifiers as needed to distinguish a person's different roles in life (rather than the hardware used to communicate).

To make calls between a PSTN phone and an Internet phone, the telephone network has to translate between the telephone number and the URI.

The Intelligent Network provides an ideal mechanism for this capability. Currently IN is used to translate a given set of digits to another set of digits, a function known as address translation. The most common example is the translation of "800" numbers to a physical telephone network destination, based on factors such as time-of-day and geographical location. The Wireless Intelligent Network translates a Personal Communication Service (PCS) number to a routing number for the current location of the subscriber. The translation capabilities necessary to implement these services reside in an IN SCP. When the SSP detects a call to a number requiring translation, e.g., 800 number, it initiates a query to the SCP. The SCP responds with routing instructions to the actual physical destination. The SCP can be extended to identify a destination URI based on translation of the dialed number.

Once the network determines that the call is destined for the Internet, the call must be routed to an appropriate gateway between PSTN and Internet. In the IN architecture, this gateway is classified as a specialized resource function and can be housed in an Intelligent Peripheral. The Intelligent Peripheral receives a call from a PSTN facility, packetizes it and delivers it across the Internet to the destination user, alerting the Internet user via one of the Internet telephony signaling protocols that have been defined such as the Session Initiation Protocol (SIP) [2] or ITU Recommendation H.323 [3].

Figure 2 depicts a typical call flow for this architecture. In Step 1, the caller requests a call to a special number (e.g., 800 number). Upon determination that this number requires IN processing, the SSP sends a message to the IN SCP requesting instructions to process the call (Step 2). The SCP determines that the destination of the call resides on the Internet, requests the connection to an Intelligent Peripheral, and includes the Internet URI in the response (Step 3) to the SSP. The call is routed to the Intelligent Peripheral (Step 4) and the Internet URI is passed transparently (i.e., without interpretation by the SSP) to the Intelligent Peripheral. The Intelligent Peripheral initiates a call to the appropriate destination based on the URI received from the SCP (Step 5).

Alternatively, the Intelligent Peripheral can serve solely as a media gateway. The SCP, having full knowledge of the call on the telephone network side can serve as a call agent, or media gateway controller [5] on the Internet side to handle the invitation of the Internet user to the call. In this way, the URI information does not need to be delivered to the Intelligent Peripheral, only the voice data. Although this architecture is suitable from a technological standpoint, many network operators are concerned about the security of SCPs connected to the Internet, as the IN protocols have weak or non-existent provisions for authentication and encryption.

The Internet Telephony Gateway Architecture, based on an IN SCP, can serve as a generic building block that can be incorporated into many existing services. In the first three examples of Table 1 the caller is not aware that the called party is actually on the Internet. In the final example, the caller has explicitly requested that the call be delivered to an Internet computer.

The general problem of mapping telephone numbers to IP addresses is also considered by Lee and Orsic [14]. It has, for example, been suggested that IT subscribers obtain a different country code instead of numbers drawn from regular "domestic" area codes, but a country code makes it easy to route the calls to gateways that can connect the call to the Internet, rather than to the telephone exchange identified by the area code. This exchange may not have access to a gateway and, given that IP addresses don't have a direct correspondence to geographic location, may not even be close to the current location of the Internet terminal. Beyond the IN approach above, the Internet telephony gateway could use DNS, LDAP [15] or SIP for mapping phone numbers to (dynamic) IP addresses.

5 Billing for Internet Telephony

Internet telephony affords us the opportunity not only to re-engineer both network and devices to provide telephone service, but also the information systems required for such functions as buying and paying for communications services. A large part of the cost of present telephone service can be attributed to producing the telephone bill.

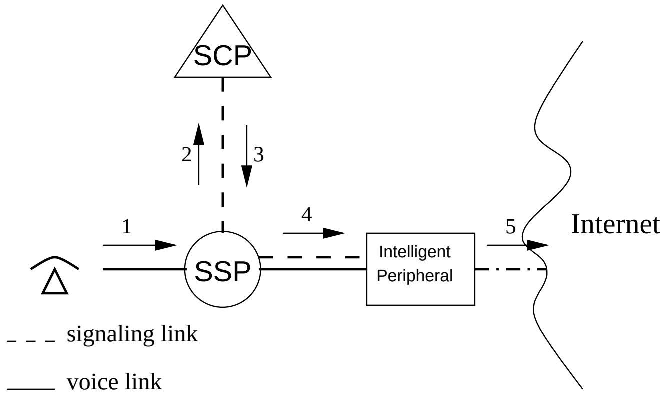

1. Caller dials special telephone number;
2. IN SSP triggers on dialed number and requests instructions from an IN SCP;
3. IN SCP redirects call to an Intelligent Peripheral;
4. IN SSP routes call to the Intelligent Peripheral;
5. The Intelligent Peripheral converts the switched voice call to a packet voice call over the Internet.

Figure 2: Internet Telephony Gateway Call Flow

Service Example	Description
Toll-free Service	A small business may provide a toll-free number for ordering products. This toll-free number could be routed to an Internet call center.
Multi-location Ringing	Multi-Location Ringing enables subscribers who have multiple work locations to program their incoming calls to sequentially ring at all locations. Locations can be those served by regular telephony, cellular, or Internet telephony.
Call Forwarding on Busy to Internet	When a person is connected to the Internet via an ordinary phone line, "call forwarding on busy to Internet" would allow that person to have calls that are destined for his/her home phone number to be re-routed to a personal computer on the Internet.
Voice Activated Dialing	Voice activated dialing (VAD) allows subscribers to pick up the telephone and speak the name of the person they wish to call. The VAD system will recognize the spoken name and dials either the appropriate telephone number or Internet telephony destination. This can be extended to allow the caller to speak or spell an Internet URI for the destination call.

Table 1: Example Applications for Internet Telephony Service Building Block

5.1 Billing for Transport and Services

As in the PSTN, it seems likely that transport and services will be billed separately.

For IT, packet transport will likely be billed either flat-rate (for best-effort service) or volume-based, if resource reservations or differentiated service is supplied. In either case, there is no particular reason why

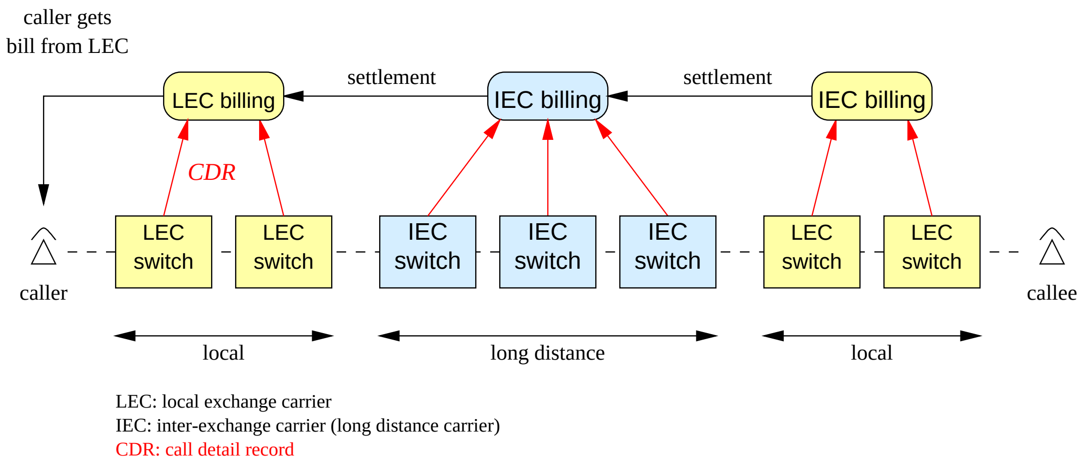

Figure 3: PSTN billing

IT packets should be billed any differently than, say, packets for video-on-demand. (On the other hand, IT may have specific charging requirements if traditional models of caller-pays or callee-pays, rather than data-sender-pays, are to be maintained.) However, other services such interconnection with the PSTN, call filtering, security and authentication, mobility or media translation and recording may well be charged on a per-use or subscription basis.

5.2 How to Bill

In the PSTN, each network provider, i.e., the local exchange carrier (LEC), inter-exchange carrier (IEC) and possibly international carrier, produces a call detail record (CDR) when the call is completed (Figure 3). The CDR contains information about the type, starting time and duration of the call, the originating and terminating number and other details [16, p. 445ff]. The CDR is delivered from the originating switch to billing center. For IT, we can model this behavior by having routers or their policy decision points submit a CDR-like record for resource reservations or calls [17]. AAA (authentication, authorization, and accounting) protocols like COPS (Common Open Policy Protocol) [18] or DIAMETER [19] may be used between the router and the policy decision point.

The Internet AAA architecture generally distinguishes between policy enforcement points (PEPs), for example, a gateway, router or server that has the power to drop packets, connection attempts or calls, and a policy decision point (PDP), that uses information supplied by the PEP to guide these decisions. Generally, the PEPs are hardware devices without a general-purpose operating systems (such as a router), while the PDPs are assumed to be regular computers with access to customer databases, rules and other information contained in access and billing policies. A single PDP may well serve several enforcement points, giving it a broader view of overall resource consumption or fraud attempts.

SIP proxy and redirect servers or H.323 gatekeepers (Fig. 4) may use AAA protocols, may generate log files that may be used to create per-use billing records or they may use RPC mechanisms, including HTTP, to convey charging information to billing centers, possibly using remote access accounting proposals like the Accounting Data Interchange Format (ADIF). “ADIF is designed to compactly represent accounting data in a protocol-independent manner. As a result, ADIF may be used to represent accounting data from any pro- tocol using attribute value pairs or variable bindings.” [20].

5.3 From Telephone Bill to Electronic Commerce

We can view communications services from two perspectives: as a subscription service where charges are aggregated monthly and as individual purchases of calls and services. In the latter model, services are purchased "on the spot", similar to how we might currently think of a coin phone or buying software on-line. The latter model is particularly attractive if users want to maintain business relationships with a large number of service and transport providers, so that the model of a monthly bill is not appropriate. Once we consider communications services as a good to be priced and purchased individually, electronic commerce models seem more applicable than billing models. For example, we can apply web technologies such as the Internet Open Trading Protocol [21] to identify customers and negotiate prices and conditions or simply convey account information using XML-based account records. However, one problem is that unlike purchasing online content, a typical phone call would generate several transactions, and possibly credit card records. One approach is to have signed copies of a call record proceed "laterally" from one proxy server or gatekeeper to another, with the aggregate record presented for billing or collection at call completion by one of the servers, e.g., the caller's proxy server. As an alternative, electronic payment protocols suited for low-value transactions, such as the Millicent protocol [22], may be more appropriate for telephone services. Information for these protocols can be carried as part of a multipart/sip-id MIME type in SIP [23].

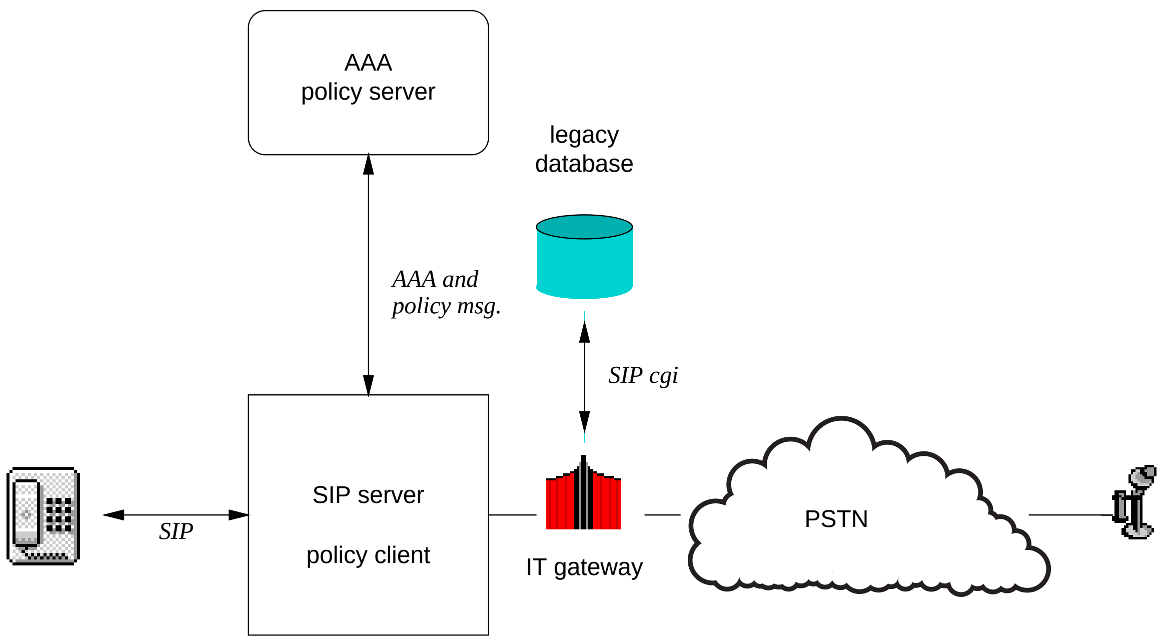

Figure 4: Billing in a SIP-Based Internet Telephony System

Full electronic commerce protocols may well find their place in inter-carrier commerce [24], allowing carriers to settle charges for communications services even when they have not established a settlement relationship. (Currently, the services of clearinghouses are used to avoid the problem of having to maintain N^2 trust relationships.)

6 Outlook

The integration of data and voice/multimedia services has been promised several times – first, as ISDN, then as B-ISDN or ATM. Maybe the third attempt, as Internet multimedia and telephony, will yield more than just hype. Unlike for the earlier attempts, no massive upgrades of end systems will be required, although many of the benefits of IT, such as quality beyond telephone quality or multimedia integration, can be reaped only once IP reaches the end device rather than stop at a gateway. Also, this time, there may well be an incentive for both users and providers, in the promise of, say, unified messaging for users, and lower costs for switching and service development for providers.

We have tried to identify some of the issues for the transition to a packet-dominated infrastructure. We use the term "packet-dominated" as it is likely that "remnants" of the legacy telephone network will be with us for many decades, so that, more so than with any other transition, interworking with the existing infrastructure will be a constant concern.

There may well be two converging evolutions: from the inside out and from the outside in. Packet transmission will grow from the inside of the network, while leaving the circuit-based interfaces to subscribers intact. Individuals and organizations, using packet modems and existing LANs, will deploy local packet-based telephone infrastructures that get converted to existing POTS soon after leaving the premises.

It is important now to develop sufficiently open technical frameworks that can accomodate the evolution of integrated telephone and other communications services for the second century of telecommunications.

References

[1] R. Guérin and H. Schulzrinne, "Network quality of service," in *Computational Grids: The Future of High-Performance Distributed Computing* (I. Foster and C. Kesselman, eds.), San Francisco, California: Morgan Kaufmann Publishers, 1998.

[2] M. Handley, H. Schulzrinne, E. Schooler, and J. Rosenberg, "SIP: session initiation protocol," Request for Comments (Proposed Standard) 2543, Internet Engineering Task Force, Mar. 1999.

[3] International Telecommunication Union, "Visual telephone systems and equipment for local area networks which provide a non-guaranteed quality of service," Recommendation H.323, Telecommunication Standardization Sector of ITU, Geneva, Switzerland, May 1996.

[4] H. Schulzrinne, "Re-engineering the telephone system," in *Proc. of IEEE Singapore International Conference on Networks (SICON)*, (Singapore), Apr. 1997.

[5] C. Huitema, J. Cameron, P. Mouchtaris, and D. Smyk, "An architecture for internet telephony service for residential customers," *IEEE Network*, vol. 13, May/June 1999.

[6] D. Rizzetto and C. Catania, "A voice over ip service architecture for integrated communications," *IEEE Network*, vol. 13, May/June 1999.

[7] J. Rosenberg, J. Lennox, and H. Schulzrinne, "Programming internet telephony services," *IEEE Network*, vol. 13, May/June 1999.

[8] J. Kuthan, "Internet telephony - an overview," tech. rep., GMD Fokus, Berlin, Germany, 1998. Available at http://www.fokus.gmd.de/research/cc/glone/projects/ipt/.

[9] International Telecommunication Union, "Q-series intelligent network recommendation structure," Recommendation Q.1200, Telecommunication Standardization Sector of ITU, Geneva, Switzerland, Mar. 1993.

[10] Telcordia Technologies (formerly Bellcore), "Advanced intelligent network generic requirements (fr-15)," 1998.

[11] J. G. van Bosse, *Signaling in Telecommunications Networks*. Telecommunications and Signal Processing, New York, New York: Wiley, 1998.

[12] I. Faynberg, L. R. Gabuzda, M. P. Kaplan, and N. J. Shah, *The Intelligent Network Standards*. New York: McGraw-Hill, 1997.

[13] International Telecommunication Union, "The international public telecommunication numbering plan," Recommendation E.164, Telecommunication Standardization Sector of ITU, Geneva, Switzerland, May 1997.

[14] C. Lee and M. Orsic, "A framework for E.164 number to IP address mapping," Internet Draft, Internet Engineering Task Force, Nov. 1998. Work in progress.

[15] W. Yeong, T. Howes, , and S. Kille, "X.500 lightweight directory access protocol," Request for Comments (Proposed Standard) 1487, Internet Engineering Task Force, July 1993.

[16] R. F. Rey, ed., *Engineering and Operations in the Bell System*. Murray Hill, New Jersey: AT&T Bell Laboratories, 1984.

[17] IP Telephony MoU, "Suggested cdr fields for internet telephony." http://pulver.com/mou/cdr.htm, Feb. 1998.

[18] J. Boyle, R. Cohen, D. Durham, S. Herzog, R. Rajan, and A. Sastry, "The COPS (common open policy service) protocol," Internet Draft, Internet Engineering Task Force, Feb. 1999. Work in progress.

[19] P. Calhoun and A. Rubens, "DIAMETER base protocol," Internet Draft, Internet Engineering Task Force, Nov. 1998. Work in progress.

[20] B. Aboba and D. Lidyard, "The accounting data interchange format (ADIF)," Internet Draft, Internet Engineering Task Force, Nov. 1998. Work in progress.

[21] D. Burdett, "Internet open trading protocol - IOTP version 1.0," Internet Draft, Internet Engineering Task Force, Mar. 1999. Work in progress.

[22] S. Glassman, M. Manasse, M. Abadi, P. Gauthier, and P. Sobalvarro, "The millicent protocol for inexpensive electronic commerce," *World Wide Web Journal*, pp. 603–618, Dec. 1995. Fourth International World Wide Web Conference Proceedings.

[23] C. Huitema, "The multipart/sip-id media type," Internet Draft, Internet Engineering Task Force, Feb. 1999. Work in progress.

[24] ETSI/TIPHON, "Inter-domain pricing, authorization and usage exchange," Technical Specification TS 101 321, ETSI/TIPHON, Dec. 1998.