Tech-Report Number CUCS-002-99

Implementing Intelligent Network Services with the Session Initiation Protocol

Jonathan Lennox
Department of Computer Science
Columbia University
lennox@cs.columbia.edu
Phone: (212) 939 7018

Henning Schulzrinne
Department of Computer Science
Columbia University
hgs@cs.columbia.edu
Phone: (212) 939 7042

Thomas F. La Porta
Bell Laboratories
Lucent Technologies
tlp@bell-labs.com
Phone: (732) 949-2281

Abstract

Internet telephony is receiving increasing interest as an alternative to traditional telephone networks. This article shows how the IETF's Session Initiation Protocol (SIP) can be used to perform the services of traditional Intelligent Network protocols, as well as additional services.

Contents

1 Introduction

In the development of Internet telephony, we want to ensure that all the features supported by modern advanced telephony systems can be supported. This article describes many of the features as they are implemented in traditional telephone networks, and then describes how they can be implemented in Internet Telephony with the IETF's Session Initiation Protocol and its extensions.

The initial task of enumerating a large number of advanced telephony services is the same one that Study Group 11 of the International Telecommunications Union Telecommunications Standards Sector (ITU-T) addressed in the process of developing their standards for Intelligent Networks. The study group published its accumulated descriptions of services and service features in Annex B of ITU-T recommendation *Q.1211: Introduction to Intelligent Network Capability Set 1* [1]. Since these service descriptions were compiled from a number of disparate sources, the document acknowledges that they may be self- and mutually-inconsistent.

This paper will describe the implementation of Internet Telephony services by following Q.1211's descriptions of each service and service feature, noting the specific description of each one so as to clarify the ambiguities and inconsistencies in the descriptions.

Study Group 11 has written a follow-up document, *Q.1221: Introduction to Intelligent Network Capability Set 2*. This document has not yet been formally ratified or released by the ITU; I address it, in less exhaustive detail, in section 6.

2 Overview

The architectural model of Internet telephony is rather different than that of the traditional telephone network. The base assumption is that all signaling and media flow over an IP-based network, either the public Internet or various intranets. This is a dramatic change in the ability of nodes in the network to communicate: in the traditional telephone architecture, nodes can generally only communicate with those other nodes to which they are directly connected.[1] IP-based networks, on the other hand, present the appearance at the network level that any machine can communicate directly with any other, unless the network specifically restricts them from doing so, through such means as firewalls.

This architectural change necessitates a dramatic transformation in the architectural assumptions of traditional telephone networks. In particular, whereas in a traditional network a large amount of administrative control, such as call-volume limitation, implicitly resides at every switch, and thus additional controls can easily be added there without much architectural change, in an Internet environment an administrative point of control must be explicitly engineered into a network, as in a firewall; otherwise end systems can simply bypass any device which attempts to restrict their behavior.

In addition, the Internet model transforms the locations at which many services are performed. In general, end systems are assumed to be much more intelligent than in the traditional telephone model; thus, many services which traditionally had to reside within the network can be moved out to the edges, without requiring any explicit support for them within the network. Other services can be performed by widely separated specialized servers which result in call setup information traversing paths which might be extremely indirect when compared with the physical network's actual topology.

[1] This is somewhat of an oversimplification for the modern SS7 backbone signalling network. Signals do not necessarily follow the same physical path as the trunks; however, except for some specialized functions such as lookups in the 800-number database, call-setup signals must proceed hop-by-hop from one switch to another, indirectly following the trunks' path.

Most of the services and service features of ITU-T Q.1211 can be provided by the IETF's draft signaling standards for Internet Telephony, SIP (the Session Initiation Protocol) [2] and, for some more specialized features, its Call Control extensions [3].

2.1 Billing

The one broad class of features which cannot be directly provided by SIP and SIP-CC is those which involve payment responsibility. In Internet telephony, it is still somewhat unclear what services can actually be charged for; clearly, for those services which can be performed by end systems, an external entity cannot expect to exact any fee for them, other than the one-time sales price the vendor of the end-system or its software receives. Those services which *do* reside in the network can be generally divided up into three categories: those residing at a single point, such as user-location or premium-rate end systems; those which involve better-than-best-effort packet delivery; and those which leave the Internet for some other network, such as PSTN gateways.

A number of billing models are possible for each of these types of services. Three of them seem to be most likely: the subscription model, the "New Jersey Turnpike" model, and the "Garden State Parkway" model. The subscription model involves a user paying for an unlimited amount of service in advance; this is most likely applicable for the single-point services discussed above. The "New Jersey Turnpike" model is payment by settlement — when a service is initiated, the user commits to paying after completion for however much service he or she has used. This is the model of current residential or charge-card service in traditional telephony. The third model, "Garden State Parkway", is pay as you go; a user commits some token before the service is initiated which allows a certain amount of service, and commits further tokens along the way. This is the model of traditional coin-operated telephone calls, or pre-paid calling cards.

Regardless of the billing model used, the format of billing information for Internet telephony is essentially orthogonal to questions of its signalling. Work is currently ongoing in the IETF's Internet Open Trading Protocol working group to define standards for exchange of such information.

3 Architecture of IPtel signaling

The flow and control model of SIP signaling is generally based on the existing models of Internet e-mail (operating at signaling rather than background-batch speeds) and (to a somewhat lesser extent) the World Wide Web.

Every SIP address specifies (through the usual DNS-style addressing) a server domain in which the address resides. Addressing is deliberately reminiscent of e-mail: sip:⟨user⟩@⟨domain⟩ or sip:⟨user⟩@⟨host⟩. An end system, when placing a call, will either directly look up the remote destination specified in the address, to send it the SIP invitation directly, or it will forward the call invitation to a local proxy, which will perform this lookup function for it (and can perform other functions as well). The resolved destination may be the address of the actual destination end system; it may be a recipient-side proxy, which handles such services as user location and firewall punch-through; or it could be a redirection server, which informs the originating station of a different server at which it should search for the user.

It is important to note that unless specifically architected (as in firewalls) no proxy server can guarantee that it will be on the server path for calls between any pair of end systems. Also note that the majority of the Internet path over which both media and signals will flow (the backbone) will never see any signaling information as anything other than IP packets to be routed to their destinations.

Some more specific details of the Internet telephony architecture are discussed in the sections describing their corresponding IN services.

4 Capability Set 1: Service Features

Q.1211 divides the services it describes into two broad categories: "services," which are what an Intelligent Network vendor would actually wish to provide to customers; and "service features," which are lower-level building blocks used to construct the services.

This section describes all the service features listed in Q.1211, Annex B, section 2; Section 5 describes the services built out of these service features (from Annex B section 1), and any unique aspects of them in the Internet telephony environment.

For each service, Q.1211 lists a number of service features which are considered either core to it, those without which it would not be useful, or optional, which provide added value to the feature. For each service feature, this section lists the services which Q.1211 specifies use that feature. Similarly, in section 5, each service lists its component service features.

See table 1 for a summary of the characteristics of each service feature.

4.1 Abbreviated Dialing (ABD)

Used for: ABD (core); ACC (core); AAB (optional); CCC (optional); VPN (optional)

> Abbreviated dialing allows the definition of short (e.g., two digit) digit sequences to represent the actual dialing digit sequence for a public or private numbering scheme.

In Internet telephony, an end system would typically do this work; either by storing an internal table of locally-defined shortcut addresses for the actual addresses it would send, or (for setups more analogous to VPNs) by having end systems configured to consult a local database server (running, e.g., LDAP) for address-translation queries.

This translation could also be performed by a local proxy or redirection server through which the end system always sends its outgoing call requests.

4.2 Attendant (ATT)

Used for: VPN (optional)

> This allows VPN users to access an attendant (operator) position within the VPN for providing VPN service information (e.g, VPN numbers) by dialing a special access code.

An Internet telephony end system needs only to be configured with an address of an appropriate local operator to translate the special access code to the actual local address of an attendant, or some address which will resolve to that address.

4.3 Authentication (AUTC)

Used for: FPH (optional); SEC (core); VPN (optional)

> This allows verification that a user is allowed to access certain options in the telephone network.

SIP allows for transactions to be authenticated, using either the standard HTTP "basic" or "digest" authentication, or an extended authentication using more sophisticated methods such as PGP or S/MIME. Unlike in the traditional telephone system, calls can be authenticated end-to-end to remote parties, in addition to providing authentication to the signaling infrastructure; these are the Authenticate and Proxy-Authenticate headers respectively.

4.4 Authorization code (AUTZ)

Used for: ACC (core); AAB (core); CCC (core); UPT (core); VPN (optional)

> This allows a user (typically in a VPN) to override the restrictions placed on the system from which calls are made.

This can be accomplished by using Proxy-Authenticate as described in the previous service feature, or, in simpler situations, by specifying a password in a SIP URL.

Service Feature			**Characteristics**	
Code	Name	Section	Location	Call Time
Authentication				
AUTC	Authentication	4.3	End / Proxy	Setup
AUTZ	Authorization code	4.4	Proxy	Setup
Billing				
PRMC	Premium charging	4.33	-	-
REVC	Reverse charging	4.35	-	-
SPLC	Split charging	4.36	-	-
Filtering				
OCS	Originating call screening	4.30	End / Proxy	Setup
TCS	Terminating call screening	4.37	End / Proxy	Setup
Forwarding				
CD	Call distribution	4.6	Proxy / Redirect	Setup
CF	Call forwarding	4.7	Proxy / Redirect	Setup
CFC	Call forwarding on busy/don't answer	4.8	Proxy / Redirect	Setup
ONE	One number	4.28	Proxy / Redirect	Setup
ODR	Origin dependent routing	4.29	Proxy / Redirect	Setup
PN	Personal numbering	4.32	Proxy / Redirect	Setup
TDR	Time dependent routing	4.38	Proxy / Redirect	Setup
Translation				
ABD	Abbreviated dialling	4.1	End / Proxy / Redirect	Setup
ATT	Attendant	4.2	End / Proxy / Redirect	Setup
PNP	Private numbering plan	4.34	End / Proxy / Redirect	Setup
User interface				
CW	Call waiting	4.15	End	Setup
CRG	Customized ringing	4.20	End	Setup
Other				
ACB	Automatic call back	4.5	End	Setup
GAP	Call gapping	4.9	Proxy	Setup
CHA	Call hold with announcement	4.10	End	In call
LIM	Call limiter	4.11	End	Setup
LOG	Call logging	4.12	All	All
QUE	Call queueing	4.13	End / Proxy	Setup
TRA	Call transfer	4.14	End	In call
CUG	Closed user group	4.16	End / Proxy	Setup
COC	Consultation calling	4.17	End	In call
CPM	Customer profile management	4.18	End / Proxy / Redirect	Indep. of call
CRA	Customer recorded announcement	4.19	End	In call
DUP	Destinating user prompter	4.21	End	In call
FMD	Follow-me diversion	4.22	End	Indep. of call
MAS	Mass calling	4.23	Proxy	Setup
MMC	Meet-me conference	4.24	Other	Setup
MWC	Multi-way calling	4.25	End	Setup
OFA	Off-net access	4.26	All	All
ONC	Off-net calling	4.27	Proxy	Setup
OUP	Originating user prompter	4.31	All	Setup

Table 1: This table indicates the characteristics of each service feature, and sorts them into rough categories. For the "location" column, "End" means the service can be provided at an end system, "Proxy" means at a proxy server, and "Redirect" at a redirection server.

4.5 Automatic call back (ACB)

Used for: CCBS (core)

> This feature allows the called party to automatically call back the calling party of the last call directed to the called party.

This can be handled entirely by end systems, which need only store and make available to the user the previous invitations they have received. Note that this is not restricted to only the last call; the number of old invitations a SIP end system remembers is bounded only by its local storage.

4.6 Call distribution (CD)

Used for: CD (core); DCR (core); FPH (optional); MAS (optional); PRM (optional); SPL (optional); VOT (optional); UAN (optional); VPN (optional)

> This service feature allows the served user to specify the percentage of calls to be distributed among two or more destinations. Other criteria may also apply to the distribution of calls to each destination.

This is one of the core purposes of a proxy (or, potentially, redirect) server, which could make decisions for the destination addresses for a request based on any data or state it sees fit, including a random choice or previous call volumes.

4.7 Call forwarding (CF)

Used for: CF (core)

> This service feature allows the user to have his incoming calls addressed to another number, no matter what the called party line status may be.

This is the simplest case of proxying or redirection: unconditionaly directing the calls to the forwarding address.

4.8 Call forwarding on busy/don't answer (CFC)

Used for: CRF (optional); FPH (optional); PRM (optional); SCF (core); SPL (optional); UAN (optional)

> This service feature allows the called user to forward particular calls if the called user is busy or does not answer within a specified number of rings.

A proxy server can implement this by awaiting the response to a call invitation. If the target end system does responds that it is busy, or does not respond within a certain period of time, the proxy server can initiate a new proxy request, or return a redirection. Alternatively, an end system could be configured to return a redirection if it is busy or not picked up.

4.9 Call gapping (GAP)

Used for: FPH (optional); PRM (optional); SCF (core); SPL (optional); UAN (optional)

> This feature allows the service provider to restrict the number of calls to a served user to prevent congestion of the network.

The intended scenario of this service feature is that vast numbers of people simultaneously call the same destination address, for instance because it was announced on television, and the network needs to ensure that its servers and signalling network are not overloaded.

The simplest case of this is when the overloaded server does not have the necessary resources to completely fulfill the request, but it can still process it and send a basic response. In this case, a SIP server can send a 503 error response, "Server Unavailable."

Somewhat more complicated is the case when the volume of traffic expected is such that a single server will not be able to handle even this class of error responses. In this case, one solution would be to have an "outer ring" of proxy servers which can either statelessly forward the request to the actual server, or return a 503 error; the central server would communicate with the periphery through non-SIP means to tell them what fraction of requests to let through.

If network signalling bandwidth is an issue, the outer ring of servers can be dispersed over a large range of network locations; DNS load-balancing can be used to distribute signalling information among them.

4.10 Call hold with announcement (CHA)

Used for: VPN (optional)

> The call hold with announcement service feature allows a subscriber to place a call on hold with options to play music or customized announcements to the held party.

This can be handled simply by switching the media which is being sent to the remote party. It can either originate from the same end system, or from a media server, perhaps triggered by a media server control protocol such as RTSP [4]. (RTSP servers are normally restricted from sending media to a third party, but if the RTSP server and SIP server trust each other this could be overridden.)

4.11 Call limiter (LIM)

Used for: CRD (optional); FPH (optional); MAS (optional); PRM (optional); SPL (optional); VOT (optional); UAN (optional)

> This service feature allows a served user to specify the maximum number of simultaneous calls to a served user's destination. If the destination is busy, the call may be routed to an alternative destination.

This can easily be done by an end system, which may return redirection messages to an alternate location.

4.12 Call logging (LOG)

Used for: ABD (optional); ACC (optional); AAB (optional); CD (optional); CF (optional); CRD (optional); CCBS (optional); CON (optional); CCC (optional); DCR (optional); FMD (optional); FPH (optional); MCI (core); MAS (optional); OCS (optional); PRM (optional); SEC (optional); SCF (optional); SPL (optional); VOT (optional); TCS (optional); UAN (optional); UPT (optional); UDR (optional); VPN (optional)

> This service feature allows for a record to be prepared each time that a call is received to a specified telephone number.

Obviously, any element of a SIP system may log anything it wishes, if it has someplace to store the log.

4.13 Call queueing (QUE)

Used for: CRD (optional); FPH (optional); MAS (optional); PRM (optional); SPL (optional); VOT (optional); UAN (optional); VPN (optional)

> This service feature allows calls which would otherwise be declared busy to be placed in a queue and connected as soon as the free condition is detected. Upon entering the queue, the caller hears an initial announcement informing the caller the call will be answered when a line is available.

There are two approaches that can be taken to this: SIP-level signaling, or end-system-level queueing. SIP-level allows the recipient (whether an end system or proxy) to signal to the caller with a "182 Queued" message, informing it of its status. This message cannot specify a media connection per se, but it could include as a payload or a URL reference a static media file (or other format, such as HTML) describing the situation or current queue status.

Alternately, a system could do queueing entirely at an end system, accepting the call from a signaling level and directing the call's media to a server which plays appropriate announcements about the queued call's state. When the call is de-queued to be picked up, the queueing server either transfers the call (if the calling system supports SIP Call Control [3]) or proxies the signaling to the appropriate end system while directing the media transmission there. This solution is more similar to how call queueing is handled in the traditional telephone network.

4.14 Call transfer (TRA)

Used for: VPN (optional)

> The call transfer service feature allows a subscriber to place a call on hold and transfer the call to another location.

The SIP Call Control draft allows a wide and flexible variety of decentralized call transfer and multi-party call operations. (See section 7.2 for further discussion of this.)

4.15 Call waiting (CW)

Used for: CCBS (optional)

> This service feature allows a subscriber to receive a notification that another party is trying to reach his number while he is busy talking to another calling party.

Due to the separation of signaling and media in Internet telephony, this feature is entirely an end-system issue. An end system which receives a call invitation while in a call may alert the user however it wishes, if it so chooses.

4.16 Closed user group (CUG)

Used for: VPN (optional)

> This service feature allows the user to be a member of a set of VPN users who are normally authorized to make and receive calls only within the group.

Both making calls and receiving calls within a group are restrictions which in SIP require administrative control of end systems. The end systems can either directly enforce these calling restrictions, or they can insist that all their signalling go through a fixed local proxy, which enforces these rules.

Alternately, if the desired closed group corresponds to the end systems on some particular part of the underlying network topology, firewalls could keep calls restricted to that sub-network.

4.17 Consultation calling (COC)

Used for: CON (optional); VPN (optional)

> The consultation calling service feature allows a subscriber to place a call on hold, in order to initiate a new call for consultation.

Initiating new calls is possible at any time in Internet Telephony. Placing a call on hold is a matter of either ignoring its media, if bandwidth is not an issue; or, more efficiently, sending it a re-invitation with media turned off. In either case, the local end system would likely either stop sending media or transmit a recorded message to the remote party.

4.18 Customer profile management (CPM)

Used for: ABD (optional); CD (optional); CF (optional); CRD (optional); CON (optional); DCR (optional); FMD (optional); FPH (optional); MAS (optional); OCS (optional); PRM (optional); SEC (optional); SCF (optional); SPL (optional); VOT (optional); TCS (optional); UAN (optional); UPT (optional); UDR (optional); VPN (optional)

> This service feature allows the subscriber to real-time manage his service profile, i.e. terminating destinations, announcements to be played, call distribution, and so on.

Features that reside in end systems can obviously be configured at these end systems transparently. Features that reside in the network, in proxies or redirect servers, can be configured through any number of means; one under current development is a Call Processing Language.

4.19 Customer recorded announcement (CRA)

Used for: CDR (optional); FPH (optional); MAS (optional); PRM (optional); SPL (optional); VOT (optional); UAN (optional); UPT (optional); VPN (optional)

> This service allows a call to be completed to a (customized) terminating announcement instead of a subscriber line. The served user may define different announcements for unsuccessful call completions due to different reasons (e.g. caller outside business hours, all lines are busy).

A server receiving an incoming call can direct the media component of a call it accepts to any device which can send appropriate media, including, for instance, an RTSP server to play back media. (See the note about RTSP in section 4.10.)

4.20 Customized ringing (CRG)

Used for: FPH (optional); PRM (optional); SPL (optional); UAN (optional); VPN (optional)

> This service feature allows the subscriber to allocate a distinctive ringing to a list of calling parties.

SIP end systems can easily do this based on the incoming request. Note that the distinctive ring does not need to be based only on the identity of the calling party; rings could be assigned based on call priority, the caller's organization, or any other aspect of the invitation. Sophisticated "rings" such as text-to-speech announcement of the calling party's name are also possible.

4.21 Destinating user prompter (DUP)

Used for: ABD (optional); FPH (optional); PRM (optional); SPL (optional); UPT (optional)

> This service feature enables to prompt the called party with a specific announcement. Such an announcement may ask the called party enter an extra numbering, e.g. through DTMF, or a voice instruction that can be used by the service logic to continue to process the call.

It is not clear that this is necessarily the best solution to this problem in an Internet environment; for example, if an HTML viewer is available, browsing via HTML/HTTP would probably make for a much better user experience than listening to a voicemail tree. However; for voice-only environments, an end system can implement the same functionality as in the traditional telephone network.

The DTMF codes could either be sent literally as audio data, or encoded using the specific DTMF encoding described in [5].

4.22 Follow-me diversion (FMD)

Used for: FMD (core); UPT (core); VPN (optional)

> With this service feature, a user may register for incoming calls to any terminal access.

This is exactly the purpose of the SIP REGISTER message; it allows a user to register a destination SIP URL with a proxy server as being associated with an end station. Note that this is more general than IN's FMD, since multiple end stations can be registered simultaneously; the proxy server will perform a parallel search to all of them.

4.23 Mass calling (MAS)

Used for: FPH (optional); MAS (core); VOT (core)

> This service feature allows processing of huge numbers of incoming calls, generated by broadcasted advertisings or games.

This feature is similar to call gapping (section 4.9), except that rather than prevent the overloading calls from going through, they are processed in "early" in a distributed manner. Typically this is intended for services such as televoting.

This can be handled through similar techniques as were described for call gapping. In particular, since server addresses are resolved through DNS, load can be distributed over multiple servers with DNS load-balancing.

4.24 Meet-me conference (MMC)

Used for: CON (optional)

> This service feature allows the user to reserve a conference resource for making a multi-party call. At a specified date and time, each participant in the conference has to dial a designated number in order to have access to the conference.

In SIP, this can be implemented either with a conference bridge which performs the audio mixing, or (if third-party call control is available) sets up the various parties to talk to each other with appropriate Also messages.

4.25 Multi-way calling (MWC)

Used for: CON (core)

> This service feature allows the user to establish multiple, simultaneous telephone calls with other parties.

With SIP, an end system may initiate as many simultaneous calls as it wishes, subject to its available network bandwidth.

4.26 Off-net access (OFA)

Used for: VPN (optional)

> This service feature allows a VPN user to access his or her VPN from any non-VPN station in the PSTN by using a personal identification number (PIN).

This is really two issues: security authorization, and access to a VPN's services (numbering plan, etc). The former issue is best handled at a lower network layer, using appropriate data access tunneling protocols, to traverse the firewall; the latter can then be accomplished by simply treating the node like any other node on the intranet.

(I assume here that we substitute "Public Internet" for "PSTN" in the service feature description to achieve an equivalent Internet telephony feature. Connecting to an Internet VPN from the PSTN is a modem dialup running PPP or a similar protocol.)

4.27 Off-net calling (ONC)

Used for: VPN (optional)

> This service feature allows the user to call outside the VPN network.

This is exactly the firewall penetration problem; calls destined outside the intranet use the firewall's proxy server to reach the outside destination.

4.28 One number (ONE)

Used for: CD (core); CRD (core); FPH (core); PRM (core); SPL (core); UAN (core)

> This feature allows a subscriber with two or more terminating lines in any number of locations to have a single telephone number. This allows businesses to advertise just one telephone number throughout their market area and to maintain their operations in different locations to maximize efficiency. The subscriber can specify which calls are to be terminated on which terminating lines based on the area the calls originate.

This service can be achieved through use of a proxy or redirect server; alternately, if the desire is to distribute that as well, the DNS records for the domain could be pointed at multiple, geographically-separated sites.

The problem of assigning requests to "nearby" (in a network sense) servers is the wide-area service location problem, and is currently being researched in its general case [6].

4.29 Origin dependent routing (ODR)

Used for: CD (optional); DCR (optional); FPH (optional); MAS (optional); PRM (optional); SPL (optional); VOT (optional); UAN (optional); UDR (optional)

> This service feature enables the subscriber to accept or reject a call, and in case of acceptance, to route this call, according to the calling party geographical location. This service feature allows the served user to specify the destination installations according to the geographical area from which the call was originated.

Note that in this context, "routing" refers to the selection of an endpoint, not the network path a call takes to that endpoint. (Q.1211 is inconsistent in this usage.) In an Internet context, this is essentially the same service as One Number in section 4.28.

In the Internet environment, geographic location does not tend to be terribly relevant for network purposes, except for inter-continental connections. Thus, presumably some other criterion, such as shortest network path or cheapest acceptable QoS, will be used instead; again, this is the wide-area service location problem.

4.30 Originating call screening (OCS)

Used for: FPH (optional); MCI (core); MAS (optional); OCS (core); PRM (optional); SPL (optional); VOT (optional); UAN (optional)

> This service feature allows the served user to bar calls from certain areas based on the District code of the area from which the call is originated.

In an Internet context, the corresponding address attribute to the District code would be either the DNS domain or the IP network of the originating address (the address in the From field). Either a proxy or an end system could filter requests on these attributes, or the entire address. Alternately, the filtering could be performed on the media's destination network address.

However, it is important to realize that an unauthenticated SIP request provides no guarantee that it actually came from the party associated with the address claimed in the From field. The only way this can be verified is with cryptographic signing of the request, along with an infrastructure for distribution of public keys. Media addresses are more likely to be meaningful in the absence of any authentication (at least if the remote party does indeed appear to be receiving the media), though media could in principle be passing through a forwarding gateway.

Note that Q.1211's choice to name this service feature "Originating call screening" contradicts every other usage of the term. Normally, it is called "Terminating call screening", as indeed it is in the service cited in section 5.24. The distinction is whether the name is chosen based on the fact that the screening is done at the termination point, as is normally done, or on the fact that the filters act upon the originating address, as is apparently the case here.

4.31 Originating user prompter (OUP)

Used for: ACC (core); AAB (core); CCC (core); FPH (optional); MAS (optional); PRM (optional); VOT (optional); UAN (optional); UPT (optional); VPN (optional)

> This service feature allows a served user to provide an announcement which will request the caller to enter a digit or series of digits via a DTMF phone or generator. The collected digits will provide additional information that can be used for direct routing or as a security check during call processing.

It is not clear that this service per se is useful in the Internet context; the proper approach would be for the initial request to return a "Multiple choices," "Ambiguous," or "Authorization required" error condition, allowing the end system to resubmit its request with the appropriate additional information. The error code could, if desired, be accompanied by a web page or a sound file containing the announcement.

4.32 Personal numbering (PN)

Used for: UPT (core)

> This service feature supports a UPT number that uniquely identifies each UPT user and is used by the caller to reach that UPT user.

Reaching a user through a personal address is accomplished simply by a proxy or redirection server which locates the user.

There are three possible levels of scoping for personalized SIP addresses. The one that will be the most common is specialized usernames. Analogous to the current e-mail forwarding servers, services can be set up which permanent SIP addresses to their subscribers, redirecting to actual end-system addresses. The domain is that of the service provider. An intermediate sort of permanent personal addressing would be vanity DNS domains; permanent SIP addresses in the user's chosen name could be set up much as personal web pages can be set up today.

Finally, a provider could choose to set up a permanent addressing scheme of host-independent naming. End systems could access this information when encountering addressing schemes which do not conform to the standard SIP addresses; this is analogous to Netscape's Netcenter name lookup, for instance.

4.33 Premium charging (PRMC)

Used for: PRM (core)

> This service feature allows for the pay back of part of the cost of a call to the called party.

This is (in the U.S.) 900-number service; see section 2.1 for discussion of billing in the Internet context. This is perhaps the simplest case of billing, as in this case the telephone company is simply acting as a settlement authority so the called party can bill the caller for their services. In an Internet environment this service could either be negotiated directly between the two parties, or some external settlement authority trusted by both parties could be used.

4.34 Private numbering plan (PNP)

Used for: VPN (core)

> This service feature allows the subscriber to maintain a numbering plan within his private network, which is separate from the public numbering plan.

Since in the Internet environment, addressing is always controlled by independently administered systems, this largely becomes trivial. If a numbering plan should be hidden or partially hidden from the public, a proxy can pretend to know nothing about the private addresses when they come from the outside; or proxies or end systems can re-write addresses to maintain a mapping from local to global names. Call routing for internal calls can be based either on the internal addresses or performed transparently by internal routers.

For instance, many systems re-write internal e-mail addresses today so that usernames specified without corresponding hosts are delivered as though they were addressed to the user at the local domain. A similar re-writing system could easily be set up for SIP.

4.35 Reverse charging (REVC)

Used for: FPH (core)

> This service feature allows the service subscriber (e.g. freephone) to accept to receive calls at its expense and be charged for the entire cost of the call.

This is (U.S.) 800-number service; see section 2.1 for a discussion of billing in the Internet context.

4.36 Split charging (SPLC)

Used for: SPL (core); UPT (core)

> This service feature allows for the separation of charges for a specific call, the calling and called party each being charged for one part of the call.

This is a generalization of the previous billing problems, with multiple parties responsible for payment. Again, see section 2.1.

4.37 Terminating call screening (TCS)

Used for: TCS (core)

> This service feature allows the user to screen calls based on the terminating telephone number dialed.

This can easily be done by a proxy which can force itself to be on the outgoing signalling path of an end system, or by administrative control over an originating end system (see section 3).

As with OCS (section 4.30), the use of the term “Terminating call screening” for this service feature in Q.1211 contradicts every other usage of the term. It is much more commonly known as “Originating call screening,” since it occurs at or near the origination point of the call.

4.38 Time dependent routing (TDR)

Used for: CD (optional); DCR (optional); FPH (optional); MAS (optional); PRM (optional); VOT (optional); UAN(optional); UPT (optional); UDR (optional); VPN (optional)

> This services feature allows the served user to apply different call treatments based on the time of day, day of week, day of year, holiday, etc.

Any proxy or redirect server with a clock could be programmed to make decisions of this sort.

5 Capability Set 1: Services

This section describes all the services listed in Q.1211, Annex B, section 1. These services are built out of the service features listed in section 4. For each service, we list which service features are core (essential) to it, or optional, according to Q.1211.

See table 2 for a summary of the characteristics of each service.

5.1 Abbreviated dialling (ABD)

Uses: ABD (core); LOG (optional); CPM (optional); DUP (optional)

> This service is an originating line feature that allows business subscribers to dial others in their company using, e.g., only four digits even if the calling user's line and the called user's line are served by different switches.

This is entirely covered by the descriptions of the component service features; see particularly the discussion of the ABD service feature in section 4.1.

5.2 Account card calling (ACC)

Uses: ABD (core); AUTZ (core); LOG (optional); OUP (core)

> The account card calling service allows subscribers to place calls from any normal access interface to any destination number and have the cost of those calls charged to the account specified by the ACC number.

This is another generalization of the previous billing problems, namely the separation of the responsibility for payment with the recipient of a service. See section 2.1.

5.3 Automatic alternative billing (AAB)

Uses: ABD (optional); AUTZ (core); LOG (optional); OUP (core)

> This service allows a user to call another user and ask him to receive the call at his expenses.

This is another generalization of billing; again, see section 2.1.

Service		Characteristics		
Code	Name	Section	Location	Call Time
Authentication				
SEC	Security screening	5.17	All	All
Billing				
ACC	Account card calling	5.2	-	-
AAB	Automatic alternative billing	5.3	-	-
CCC	Credit card calling	5.9	-	-
FPH	Freephone	5.12	-	-
PRM	Premium rate	5.16	-	-
SPL	Split charging	5.22	-	-
UPT	Universal personal telecommunications	5.26	Proxy / Redirect	Setup
Filtering				
OCS	Originating call screening	5.15	End / Proxy	Setup
TCS	Terminating call screening	5.24	End / Proxy	Setup
Forwarding				
CD	Call distribution	5.4	Proxy / Redirect	Setup
CF	Call forwarding	5.5	End / Proxy / Redirect	Setup
CRD	Call rerouting distribution	5.6	End / Proxy	Setup
DCR	Destination call routing	5.10	Proxy / Redirect	Setup
FMD	Follow-me diversion	5.11	Proxy / Redirect	Indep. of call
SCF	Selective call forwarding on busy/don't answer	5.18	End / Proxy / Redirect	Setup
UAN	Universal access number	5.25	Proxy / Redirect	Setup
Multi-party				
CON	Conference calling	5.8	End / Other	In call
Other services				
VOT	Televoting	5.23	End	In call
UDR	User-defined routing	5.27	-	Setup
VPN	Virtual private network	5.28	-	-
Translation				
ABD	Abbreviated dialling	5.1	End / Proxy / Redirect	Setup
User interface				
CF	Call forwarding	5.5	End / Proxy / Redirect	Setup
Other				
CCBS	Completion of calls to busy subscriber	5.7	End	Setup
MCI	Malicious call identification	5.13	End / Proxy	Setup
MAS	Mass calling	5.14	Proxy	Setup

Table 2: This table indicates the characteristics of each service, and sorts them into rough categories. For the "location" column, "End" means the service can be provided at an end system, "Proxy" means at a proxy server, and "Redirect" at a redirection server.

5.4 Call distribution (CD)

Uses: CD (core); LOG (optional); CPM (optional); ONE (core); ODR (optional); TDR (optional)

This service allows a subscriber to have incoming calls routed to different destinations, according to an allocation law which may be real-time managed by the subscriber.

Three types of law may exist:

- circular distribution, where the calls are routed to the different locations with a uniform load;
- percentage distribution, where the calls are routed to the different locations according to a percentage;

- hierarchical distribution, where the first location to be chosen is the first met in the priority list.

  In addition, congestion at one location may cause overflow calls to be rerouted to an alternate location.

A call proxy can handle all of these points. A redirect server can handle all but the last. The real-time management can be accomplished by changing one's Call Processing Language scripts (see the discussion of CPLs in section 4.18).

5.5 Call forwarding (CF)

Uses: CF (core); LOG (optional); CPM (optional)

> Call forwarding allows the called user to forward calls to another telephone number when this service is activated. With this service, all calls destined to the subscriber's number are redirected to the new telephone number.
>
> This service is under control of the subscriber and can be activated/deactivated by the subscriber.
>
> When this service is activated, the subscriber's line will receive an alerting ring, "reminder ring," to indicate that the service is activated.

Call forwarding itself is described in section 4.7 in the discussion of the service feature of the same name.

Reminder ring is somewhat ill-defined in Internet telephony, due to the separation of end systems and addresses; if a proxy is doing the forwarding, it is not clear what end system should be alerted of the forward, since call forwarding is just one specific case of the general user-location problem. A proxy could be configured with a specific end system to alert, however; though no signals are currently defined in SIP to request that an end system perform a reminder ring, it would be a relatively simple extension to the protocol, and several have been informally proposed.

If an end system is doing forwarding on its own, of course, it may perform a reminder ring if it so wishes.

5.6 Call rerouting distribution (CRD)

Uses: LIM (optional); LOG (optional); QUE (optional); CPM (optional); ONE (core)

> This service permits the subscriber to have his incoming calls encountering a triggering condition (busy, specified number of rings, queue overload or call limiter) rerouted according to a predefined choice: the calls may be rerouted to another destination number (including pager or vocal box), rerouted on a standard or customized announcement, or queued.

This is all easily possible in a proxy or end system; proxies can also handle the additional case when a device is not responding, whereas end systems keep enough call state to do call limiting more reliably. For control, see the description under CPM in section 4.18.

5.7 Completion of calls to busy subscriber (CCBS)

Uses: ACB (core); LOG (optional); CW (optional)

> This service allows a calling user encountering a busy destination to be informed when the busy destination becomes free, without having to make a new call attempt.

(This is noted in Q.1211 as not actually being possible to implement using only "Type A" (single-ended, single-point-of-control) service features, the type of features that Capability Set 1 specifies.)

The full semantics of this feature are somewhat elaborate: the calling user is alerted to the fact that the destination is willing to accept the call, and can re-initiate the call (automatically) at his or her leisure, before the called user is alerted to the call attempt. This cannot be implemented with the basic SIP standard, since the protocol does not provide any way to convey that information.

There are two proposed ways to create a feature which simulates this behavior to some extent. One is simply polling the destination periodically with new call attempts, and then alerting the calling user once it is available; another is specifying “Call-disposition: queue” when the call is placed (as specified in the SIP Call Control draft), to request that the call be queued rather than rejected if the destination is not available. Both these strategies have the disadvantage that the calling party might forget about the call, or get involved in another one, in which case a call would be placed but the calling party would not be present when the called party picked up.

It has also been proposed that the current standards work on Presence Information Protocols [7] could also be used to implement this feature, since the semantics of “the individual is available for communication” are similar. One of the proposals for this functionality involves using some simple extensions to SIP.

Most existing traditional telephone systems also cannot implement this service in all cases, especially between systems under separate administrative control, as their signalling protocols do not support it.

5.8 Conference calling (CON)

Uses: LOG (optional); COC (optional); CPM (optional); MMC (optional); MWC (core)

> Conference calling allows the connection of multiple parties in a single connection. The number of parties connected simultaneously will vary based on bridging requirements.
>
> #### 5.8.1 Conference calling add-on
>
> This service allows the user to reserve a conference resource for making a multi-party call, indicating the date, time, and conference duration. Once the conference is active, the user controls the conference, and may add, drop, isolate, reattach or split parties.
>
> #### 5.8.2 Conference calling meet-me
>
> This service allows the user to reserve a conference resource for making a multi-party call, indicating the date, time, and conference duration. In due time, each participant in the conference has to dial a special number whcih has been attached to the booked conference, in order to access the conference bridge.

There are a number of ways that conference calling can be handled in SIP. The one requiring the least effort from end systems is an architecture similar to conference calling meet-me; the multiple parties all connect to a specified address, which controls a conference bridge. This conference bridge will perform mixing of the media from each station, and send it out to all the other stations involved in the call.

Another possibility, also part of the basic SIP standard, is to use a multicast conference. For this, an end system specifies a multicast address rather than a unicast address as its media destination. Any end system which is in the scope of the multicast group can send and receive media to that group; any number of additional parties can be invited to it by anyone who knows the multicast group.

Thirdly, if the SIP Call Control extension is used, a multipoint-unicast conference may be used, in which many parties send each other data over unicast connections. Any party in the call may add additional parties to the call, may request that others drop parties, etc.

Note that in these latter two cases, control is decentralized; any party may add other parties to the call, and so forth. (For multipoint unicast, one party may refuse to add another to its part of the mesh, however.) Only in the first case, where the conference’s media distribution is centralized, can one party have complete control over who is and is not admitted to the conference. Normally this control resides in the MCU itself; it could, of course, be controlled remotely, and the mechanisms of SIP Call Control could be used between the controlling party and the MCU to accomplish this.

5.9 Credit card calling (CCC)

Uses: ABD (optional); AUTZ (core); LOG (optional); OUP (core)

> The credit card calling service allows subscribers to place calls from any normal access interface to any destination number and have the cost of those calls charged to the account specified by the CCC number.

This is yet another case of the billing problem. See section 2.1.

5.10 Destination call routing (DCR)

Uses: CD (core); LOG (optional); CPM (optional); ODR (optional); TDR (optional)

> This service allows customers to specify the routing of their calls to destinations according to
>
> - time of day, day of week, etc.;
> - area of call origination;
> - calling line identity of customer;
> - services attributes held against the customer;
> - priority (e.g. from input of a PIN);
> - charge rates applicable for the destinations;
> - proportional routing of traffic

See the service feature descriptions, particularly CD, ODR, and TDR in sections 4.6, 4.29, and 4.38; all of this can be done in a proxy or redirect server, possibly under the control of a call processing language. The PIN mentioned for priority could be a password in the SIP URI or an explicit header; the network server could of course also choose priorities based on its own criteria.

5.11 Follow-me diversion (FMD)

Uses: LOG (optional); CPM (optional); FMD (core)

> This service allows the subscriber to remotely control his call forwarding capabilities, basically the number to which the calls are forwarded, from any point in the network.

This is the purpose of REGISTER messages; see the FMD service feature in section 4.22. Alternately, for longer-term control, the user could simply remotely change the call processing language they have installed on the proxy or redirect server.

5.12 Freephone (FPH)

Uses: AUTC (optional); CD (optional); CFC (optional); GAP (optional); LIM (optional); LOG (optional); QUE (optional); CPM (optional); CRA (optional); CRG (optional); DUP (optional); MAS (optional); ONE (core); ODR (optional); OCS (optional); OUP (optional); REVC (core); TDR (optional)

> Freephone allows the served user having one or several installations to be reached from all or part of the country, or internationally as appropriate, with a freephone number and to be charged for this kind of call.

See all the relevant service feature descriptions, particularly REVC (section 4.35), and the general discussion of billing in section 2.1.

5.13 Malicious call identification (MCI)

Uses: LOG (core); OCS (core)

> Malicious call identification allows the service subscriber to control the logging (making a record) of calls that are received that are of a malicious nature.

Any entity through which SIP messages pass (and which has some sort of storage) can log information about these messages. If legal non-repudiability is needed, a customer can have a trusted third party perform proxying and logging for them.

See under OCS (section 4.30) about the complexities of guaranteeing the authenticity of SIP messages, however.

5.14 Mass calling (MAS)

Uses: CD (optional); LIM (optional); LOG (optional); QUE (optional); CPM (optional); CRA (optional); MAS (core); ODR (optional); OCS (optional); OUP (optional); TDR (optional)

> Mass calling involves instantaneous, high-volume traffic which is routed to one or multiple destinations. Calls can be routed to these destination numbers based on various conditions, such as the geographical location or the time of day.

See the discussion of the MAS service feature in section 4.23.

5.15 Originating call screening (OCS)

Uses: LOG (optional); CPM (optional); OCS (core)

> This services allows the subscriber to authorize outgoing calls, through the use of a screening list. This list may be managed by the subscriber. The user may override the restriction by giving a PIN.

Note that unlike the service feature with the same name, the Q.1211 summary for this service describes it correctly; thus, the proper core service feature for this service is described under TCS, in section 4.37. See that section for a full description of the implementation of this service.

5.16 Premium rate (PRM)

Uses: CD (optional); CFC (optional); GAP (optional); LIM (optional); LOG (optional); QUE (optional); CPM (optional); CRA (optional); CRG (optional); DUP (optional); ONE (core); ODR (optional); OCS (optional); OUP (optional); PRMC (core); TDR (optional)

> This service allows to pay back part of the call cost to the called party, considered as an added value service provider.

See the discussion of the PRMC service feature in section 4.33, and of billing in general in 2.1.

5.17 Security screening (SEC)

Uses: AUTC (core); LOG (optional); CPM (optional)

> This capability allows security screening to be performed in the network before an end-user gains access to the subscriber's network, systems, or applications. Access code abuse detection is a capability which will generate a report on the invalid access attempts: how many, over what time period, by whom, and from where.

Facilities which implement authorization codes should have this kind of auditing in place; other Internet services which require authorization or other security approval should have a similar security infrastructure.

5.18 Selective call forwarding on busy/don't answer (SCF)

5.19 Selective call forwarding

5.20 Call forwarding on busy

5.21 Call forwarding on don't answer (no reply)

Use: CFC (core); GAP (core); LOG (optional); CPM (optional)

> This service allows the called user to forward particular pre-selected calls if the called user is busy or does not answer within Y seconds or X rings.

These services can easily be handled by either a proxy, which could also handle "selective call forwarding on device not responding", or an end system. They could be under control of a call processing language.

The "Selective call forwarding" case can also be handled by a redirect server.

5.22 Split charging (SPL)

Uses: CD (optional); CFC (optional); GAP (optional); LIM (optional); LOG (optional); QUE (optional); CPM (optional); CRA (optional); CRG (optional); DUP (optional); ONE (core); ODR (optional); OCS (optional); SPLC (core)

> This service allows a split charging, the calling and the called party being each charged for one part of the call.

See the discussion of SPLC in section 4.36, and of billing in general in 2.1.

5.23 Televoting (VOT)

Uses: CD (optional); LIM (optional); LOG (optional); QUE (optional); CPM (optional); CRA (optional); MAS (core); ODR (optional); OCS (optional); OUP (optional); TDR (optional)

> Televoting enables subscribers to survey public opinion using the telephone network. Persons wishing to respond to an opinion poll can call advertised televoting numbers to register their votes. The charging is to the discretion of the service subscriber.

It is not clear in an Internet context how useful this actually is; a web page would be a much more natural way to present this service, and would be more efficient as it would not require setting up a voice path. However, an end system could implement this service without any intrinsic signalling support being necessary. See the OUP service feature in section 4.31.

5.24 Terminating call screening (TCS)

Uses: LOG (optional); CPM (optional); TCS (core)

> Terminating calls may be controlled by the terminating call screening capability. This allows the subscriber to specify that incoming calls be either restricted or allowed, according to a screening list and optionally, by time of day control.

Note that unlike the service feature with the same name, the Q.1211 summary for this service describes it correctly; thus, the proper core service feature for this service is described under OCS, in section 4.30. See that section for a full description of the implementation of this service.

5.25 Universal access number (UAN)

Uses: CD (optional); CFC (optional); GAP (optional); LIM (optional); LOG (optional); QUE (optional); CPM (optional); CRA (optional); CRG (optional); ONE (core); ODR (optional); OCS (optional); OUP (optional)

> This service allows a subscriber with several terminating lines in any number of locations or zones to be reached with a unique directory number. The subscriber may specify which incoming calls are to be routed to which terminating lines, based upon the area the call originated.

This is the same as the ONE service feature (section 4.28), with other features optionally "behind" it.

5.26 Universal personal telecommunications (UPT)

Uses: AUTZ (core); LOG (optional); CPM (optional); CRA (optional); DUP (optional); FMD (core); OUP (optional); PN (core); SPLC (core)

> This service provides personal mobility by enabling a user to initiate any type of service and receive any type of call on the basis of a unique and personal network-independent number, across multiple networks, at any user-network access (fixed, movable or mobile), irrespective of geographic location, limited only by terminal and network capabilities.

For Internet Telephony features, this is essentially the same as the FMD service feature (section 4.22) with authentication; once billing is resolved, the Split Charging issue needs also of course to be resolved.

This service can easily be performed by a proxy or redirection server which tracks REGISTER messages.

5.27 User-defined routing (UDR)

Uses: LOG (optional); CPM (optional); ODR (optional); TDR (optional)

> This capability allows the subscriber to specify how outgoing calls, from the subscriber's location, shall be routed, either through private, public, or virtual facilities or a mix of facilities, according to the subscriber's routing preferences list. These lists will typically apply to individual lines or to several lines at the subscriber's location.

Note that despite the fact that both this and other services refer to the ODR and TDR features, they are referring to quite different things: what network is used to place a call, vs. the destination of that call. See section 4.29.

For most non-dysfunctional networks, the choice of network for connectivity is only of concern for media flows, not signalling flows. As such, this is not properly the domain of SIP, but rather that of the mechanisms which control the routing of media packets. Normally this would most likely be under the control of administrators who set up a network's routing table, but users could dynamically control it through loose source routing or a QoS routing protocol.

In somewhat unusual circumstances, however, a user may wish to be able to specify what networks signalling flows cross. For instance, the network's standard routing might cross networks that *are* dysfunctional; or security might be so important that encryption is not sufficient, for instance if the number of call attempts made is sensitive information. In these cases, the same techniques mentioned above for media routing could also be used for signal routing, since the underlying network does not, of course, distinguish between them.

5.28 Virtual private network (VPN)

Uses: ABD (optional); ATT (optional); AUTC (optional); AUTZ (optional); CD (optional); CHA (optional); LOG (optional); QUE (optional); TRA (optional); CUG (optional); COC (optional); CPM (optional); CRA (optional); CRG (optional); FMD (optional); OFA (optional); ONC (optional); OUP (optional); PNP (core); TDR (optional)

> This service permits the subscriber to build a private network by using the public network resources. The subscriber's lines, connected on different network switches, constitute a virtual PABX, including a number of PABX capabilities, such as private numbering plans (PNP), call transfer, call hold, and so on.

In an Internet telephony environment, this would consist of two parts: an IP-level Virtual Private Network, and an internal Internet telephony system running over it. The exact definition of an IP VPN is still a matter of some debate, but it would seem to include such features as quality of service support, network-level encryption, and on-demand dynamic bandwidth allocation, so as to provide the illusion of a private intranet between geographically separated locations. See Paul Ferguson's white paper *What is a VPN* [8] for a thorough discussion of this.

On this intranet, various private numbering plans can be maintained (see PNP, section 4.34), and within it any of the service features described in this document can be implemented much as they would be on the public Internet or a physically contiguous intranet.

6 Capability Set 2

The draft of recommendation Q.1221 introduces a large number of new services and service features beyond those listed in Q.1211. Many of these do not need to be addressed here in any great detail: a large number of them concern "Service Management Services," which are used to allow direct user control of the network's services. Since the architecture of Internet Telephony puts the network's services under the user's control intrinsically, addressing these services is not directly necessary.

Capability Set 2 also lists some services which in Capability Set 1 were merely service features, and it repeats some service features from CS1 with clearer or more cleanly-separated descriptions. The descriptions of those service features are not repeated here.

The other significant services and service features which are new in CS2 are addressed below.

6.1 Wireless services

A large number of the CS2 services are designed to support wireless (mobile) networks.

SIP should work transparently over Mobile IP; with SIP re-invitations, the media path can be optimized. Research into other approaches to this problem is also underway.

6.2 Inter-network services

There are a number of services of Capability Set 2, such as Internetwork freephone (IFPH), Internetwork televoting (IVOT), and so forth, which are essentially services of Capability Set 1 extended so that they are guaranteed to work between multiple networks.

As its name implies, the Internet is intrinsically a collection of multiple networks. SIP signalling is inherently multi-domain, and all the services described in sections 4 and 5 which operate between multiple machines will work regardless of whether those machines are on the same subnetwork or opposite sides of the planet.

6.3 Multimedia

> This service allows a subscriber to receive or send an integrated communication consisting of mixtures of voice, data, image, and video information. A key capability will be the ability to synchronize and control devicery of information from disparate sources (e.g. voice and data). This will include controlling delivery from multiple sources to a single recipient and from a single source to multiple recipients.
>
> The subscriber will also desire the ability to tailor a particular service to the type of terminating device or subscriber preference (i.e. turning off the video feed). Another key aspect of this service is that additional capabilities may be requested during the call (i.e. adding data capabilities to an existing voice connection).

The Internet's media delivery protocol, RTP [9], can deliver many simultaneous synchronized media formats; it can transparently perform either point-to-point or multicast delivery; it can disambiguate multiple media senders without their previous coordination; and many other features besides.

SIP requests normally use the Session Description Protocol [10] to describe their media capabilities; through it, they can specify any set of media capabilities and formats they desire. Through SIP's re-invitation mechanism, these capabilities and formats can be altered (either added or removed) at any time during a call.

6.4 Call pick-up

> This service feature enables a user to associate a call request to an already alerting call. The alerting call awaits answer while the user originating call pick-up signals to the network a desire to connect to the alerting call. The network then connects the call parties.

The SIP standard specifies that a REGISTER request for a currently alerting call should cause that call to be transfered (either proxied or redirected) to the location specified in the registration.

6.5 Calling name delivery

> This service feature gives to the network operator the capability to display/announce the name of the calling party to the calling name delivery user (the called party) prior to answer, thus allowing this user to screen or distinctively answer the call.

Every SIP request has a From field listing the address and optionally a display name of the calling party.

See section 4.30 for issues in guaranteeing the authenticity of the specified address.

7 SIP Services not in IN

SIP also supports a number of features not supported in traditional IN networks.

7.1 OPTIONS

The SIP OPTIONS request allows one end system to query another about what formats it supports. Thus, a user can passively determine, for instance, whether video is available, without having to explicitly initiate a call to the remote party.

7.2 Third-party Call Control

The SIP Call Control draft [3] allows a wide and flexible generalization of call transfer and conference calling in a fully distributed manner. Conference calls can transition smoothly between de-centralized multiple point-to-point links, IP-level multicasting, and having media flow through a multi-point control unit.

Third-party call control consists of two header primitives: Also, with which an end system requests another one to add a third party to the call; and Replaces, with which an end system requests that the other remove the third party from a call. Any number of either of these requests can piggy-back on top of any SIP request.

Thus, adding a party to a conference is sending it an INVITE with Also referring to every current party in the call. Blind call transfer is BYE with Also to the transferee. An MCU takes over from a multiple-point-to-point call by inviting existing call members with Replaces referring to the called parties who are already using the MCU. See [11] for more details and many more examples of the flexibility of this approach to multi-point call control.

7.3 Additional invitation parameters

SIP invitation requests have a number of optional parameters which traditional telephone networks lack, generally inspired by electronic mail.

Addresses can have display names, of the form
From: Jonathan Lennox ¡sip:lennox@cs.columbia.edu¿ which present the remote system with a sender-defined "true name."

Messages can have Subject fields specified, giving a texual intended subject for a call. Users can have Organization fields, similarly giving the organization to which they belong.

Messages can have priorities set, ranging from "non-urgent" through "emergency."

If desired, other information (such as a picture of the caller) can be added.

Any of these fields can be presented to the recipient, and can also be used by proxies or redirect servers to make forwarding decisions.

7.4 Forwarding short-circuiting

When an end system responds to an initial request, it puts a `Location:` header into the response giving the most direct address at which it should be reached. Thus, all subsequent signalling requests for that call can be sent directly to that location; there is no need for them to traverse the full path of forwarding and redirection servers that the initial invitation traversed. (This location will usually be an end system, but doesn't have to be; it might, for example, be a firewall.)

References

[1] International Telecommunication Union, "General recommendations on telephone switching and signaling – intelligent network: Introduction to intelligent network capability set 1," Recommendation Q.1211, Telecommunication Standardization Sector of ITU, Geneva, Switzerland, Mar. 1993.

[2] M. Handley, H. Schulzrinne, E. Schooler, and J. Rosenberg, "SIP: session initiation protocol," Internet Draft, Internet Engineering Task Force, Nov. 1998. Work in progress.

[3] H. Schulzrinne and J. Rosenberg, "SIP call control services," Internet Draft, Internet Engineering Task Force, Feb. 1998. Work in progress.

[4] H. Schulzrinne, R. Lanphier, and A. Rao, "Real time streaming protocol (RTSP)," Request for Comments (Proposed Standard) 2326, Internet Engineering Task Force, Apr. 1998.

[5] H. Schulzrinne, "RTP payload for DTMF digits," Internet Draft, Internet Engineering Task Force, Nov. 1998. Work in progress.

[6] J. Rosenberg, H. Schulzrinne, E. Guttman, and R. Moats, "WASRV architectural principles," Internet Draft, Internet Engineering Task Force, Feb. 1998. Work in progress.

[7] L. Dusseault and M. Calsyn, "Presence information protocol requirements," Internet Draft, Internet Engineering Task Force, Feb. 1998. Work in progress.

[8] P. Ferguson and G. Huston, "What is a VPN?," white paper, Cisco Systems; Telstra Internet, Mar. 1998.

[9] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, "RTP: a transport protocol for real-time applications," Request for Comments (Proposed Standard) 1889, Internet Engineering Task Force, Jan. 1996.

[10] M. Handley and V. Jacobson, "SDP: session description protocol," Request for Comments (Proposed Standard) 2327, Internet Engineering Task Force, Apr. 1998.

[11] H. Schulzrinne and J. Rosenberg, "Internet telephony: Architecture and protocols," *Computer Networks and ISDN Systems*, 1998.