

Network Working Group
Request for Comments: 2458
Category: Informational

H. Lu
Editor
M. Krishnaswamy
Lucent Technologies
L. Conroy
Roke Manor Research
S. Bellovin
F. Burg
A. DeSimone
K. Tewani
AT&T Labs
P. Davidson
Nortel
H. Schulzrinne
Columbia University
K. Vishwanathan
Isochrome
November 1998

Toward the PSTN/Internet Inter-Networking
--Pre-PINT Implementations

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1998). All Rights Reserved.

Abstract

This document contains the information relevant to the development of the inter-networking interfaces underway in the Public Switched Telephone Network (PSTN)/Internet Inter-Networking (PINT) Working Group. It addresses technologies, architectures, and several (but by no means all) existing pre-PINT implementations of the arrangements through which Internet applications can request and enrich PSTN telecommunications services. The common denominator of the enriched services (a.k.a. PINT services) is that they combine the Internet and PSTN services in such a way that the Internet is used for non-voice interactions, while the voice (and fax) are carried entirely over the PSTN. One key observation is that the pre-PINT implementations, being developed independently, do not inter-operate. It is a task of the PINT Working Group to define the inter-networking interfaces that

will support inter-operation of the future implementations of PINT services.

Table of Contents

1.	Introduction	3
2.	Terminology	3
3.	PINT Services	4
4.	Architectural Overview	5
4.1	Public Switched Telephone Network	5
4.2	Pre-PINT Systems	9
5.	IN-Based Solutions	20
5.1	The Lucent System	20
5.1.1	Roles of the Web Server, Service Node, and SMS	20
5.1.2	A Click-to-Dial-Back Service Scenario	21
5.1.3	Web Server-Service Node Interface	22
5.1.4	Web Server-SMS Interface and SNMP MIB	24
5.1.5	Security Considerations	26
5.2	Siemens Web Call Center	27
5.2.1	Service Description	27
5.2.2	Implementation	29
5.2.3	Derived Requirements/Lessons	35
6.	Alternative Solutions	37
6.1	The AT&T System	37
6.1.1	High Level Architecture	38
6.1.2	IP Client to CallBroker Interface	39
6.1.3	Protocol	40
6.1.4	APIs Exposed to the IP Client	41
6.1.5	Voice-Bridge Control API	41
6.2	Simple Computer Telephony Protocol	41
6.2.1	Overview	41
6.2.2	How SCTP Fits in with the Reference PINT Services ..	42
7.	Session Initiation Protocol--An Emerging Standard ..	43
7.1	Overview	43
7.2	SIP Protocol	44
7.3	SIP Entities	45
7.4	Providing Call Control Functionality	46
8.	Overall Security Considerations	47
9.	Conclusion	48
10.	Acknowledgments	48
11.	Appendix	49
11.1	PSTN/IN 101	49
11.1.1	Public Switched Telephone Network	49
11.1.2	Intelligent Network	51
11.2	Call Center Features	54

12. References 56
 Authors' Addresses 57
 Full Copyright Statement 60

1. Introduction

This document contains the information relevant to the development of the inter-networking interfaces underway in the Public Switched Telephone Network (PSTN)/Internet Inter-Networking (PINT) Working Group. It addresses technologies, architectures, and several (but by no means all) existing pre-PINT implementations of the arrangements through which Internet applications can request and enrich PSTN telecommunications services. The common denominator of the enriched services (a.k.a. PINT services) is that they combine the Internet and PSTN services in such a way that the Internet is used for non-voice interactions, while the voice (and fax) are carried entirely over the PSTN.

The organization of the document is as follows. First, the basic terminology and a short "intuitive" description of the PINT services are provided. The rest of the information deals, in one way or the other, with the pre-PINT support of these services where they are used as a benchmark. Thus, an architectural overview common to all present solutions is presented. The flow of the document then divides into two streams: one is dedicated to the Intelligent Network (IN)-based solutions; the other explores alternative means (i.e., CallBroker and Computer-Telephony Integration (CTI) approach). At this point, the emerging standards are explored, in particular, the Session Initiation Protocol (SIP), which promises an elegant solution to the PINT problem. Each of the above developments is addressed in a respective section. The final sections of the document contain the overall security considerations, conclusion, acknowledgments, appendix, and a set of references. The security section summarizes the PINT security requirements derived from the pre-PINT experiences and the appendix presents a tutorial on the PSTN, IN, and Call Center functions.

2. Terminology

This document uses the following terminology:

Authentication -- verification of the identity of a party.

Authorization -- determination of whether or not a party has the right to perform certain activities.

PINT Gateway -- the PSTN node that interacts with the Internet.

User or Customer -- the person who asks for a service request to be issued. In the context of PINT Services, this person will use an Internet host to make his or her request. The term "user" is also used to describe a host originating the PINT service request on behalf of this person.

3. PINT Services

This document addresses four services initially identified by the PINT Working Group and presently supported by pre-PINT implementations. These services are: click-to-dial-back, click-to-fax, click-to-fax-back and voice-access-to-content.

Note that the word "click" should not be taken literally. It is rather used to point out that initiation of the related services takes place on the Internet, where point and click are the most prevalent user actions. In other words, a service request could originate from any type of IP-based platforms. There is no implication that these services must be implemented by a device within the PSTN or the Internet running a Web server.

The common denominator of the PINT services is that they combine the Internet and PSTN services in such a way that the Internet is used for non-voice interactions, while the voice (and fax) are carried entirely over the PSTN. (An example of such a service is combination of a Web-based Yellow Pages service with the ability to initiate PSTN calls between customers and suppliers in a manner described in what follows.)

Some of the benefits of using the PSTN are high quality of the voice, an ability to route the call to different locations depending on pre-set criteria (for example, time of the day, day of the week, and geographic location), outstanding security and reliability, and access to flexible, low cost, and secure billing and charging systems. The benefits of using the Internet are the uniform, well-defined, and widely-used interfaces available anywhere, anytime.

Click-to-Dial-Back

With this service, a user requests (through an IP host) that the PSTN call be established between another party and himself or herself. An important pre-requisite for using this service is that the user has simultaneous access to both the PSTN and Internet.

One example of an application of this service is on-line shopping: a user browsing through an on-line catalogue, clicks a button thus inviting a call from a sales representative. Note that (as is the case with the all-PSTN Free-Phone, or "800", service) flexible

billing arrangements can be implemented here on behalf of the service provider. In addition (and also similarly to the Free-Phone/800), the PSTN could route the call depending on the time of day, day of week, availability of agents in different locations, and so on.

Click-to-Fax

With this service, a user at an IP host requests that a fax be sent to a particular fax number. In particular this service is especially meaningful when the fax is to be sent to someone who has only a fax machine (but no access to the Internet). Consider, as an example, a service scenario in which a Web user makes a reservation for a hotel room in Beijing from a travel service page containing hotel information of major cities around the world. Suppose a specific Beijing hotel chosen by the user does not have Internet connection but has a fax machine. The user fills out the hotel reservation form and then clicks a button sending out the form to the travel service provider, which in turn generates a fax request and sends it together with the hotel reservation form to the PSTN. Upon receiving the request and the associated data, the PSTN translates the data into the proper facsimile format and delivers it to the Beijing hotel as specified in the fax request.

Click-to-Fax-Back

With this service, a user at an IP host can request that a fax be sent to him or her. (Consider the user of the previous example, who now requests the confirmation from the Beijing Hotel. Another useful application of the service is when size of the information that a user intends to get is so large that downloading it to the user's PC over the Internet will require a long time and a lot of disk space.)

Voice-Access-to-Content

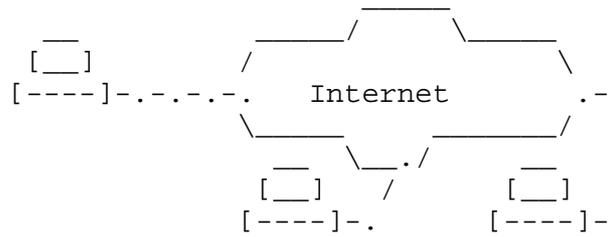
With this service, a user at an IP host requests that certain information on the Internet be accessed (and delivered) in an audio form over the PSTN, using the telephone as an informational appliance. One application of this service is to provide Web access to the blind. (This may require special resources--available in the PSTN--to convert the Web data into speech.)

4. Architectural Overview

4.1 Public Switched Telephone Network

From an application perspective, Internet nodes are interconnected directly, as shown in Figure 1. When two machines are to communicate, they will have the address of the destination end system, and will

send network level datagrams, assuming that the underlying infrastructure will deliver them as required.

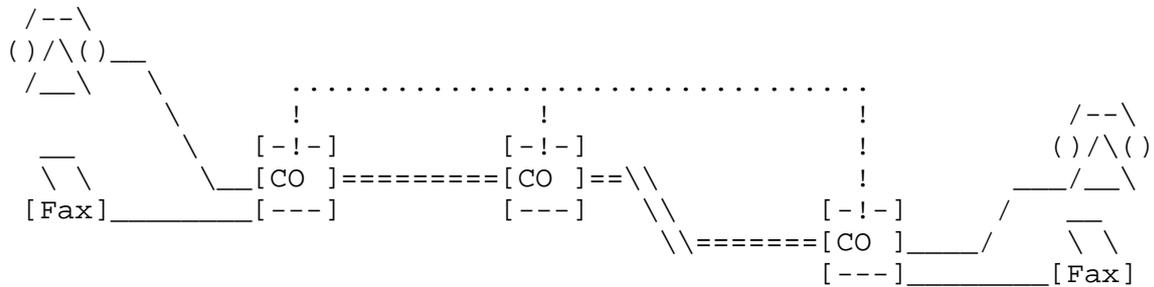


Key: .-.-. Internet Access Link

Figure 1

Where all nodes are on the same (broadcast) network, there is no need for intervening routers; they can send and deliver packets to one another directly. The Internet nodes are responsible for their own communications requests, and act as peers in the communication sessions that result.

This contrasts with the situation in the PSTN. There, the end systems are configured as shown in Figure 2. The end systems tend to be specific to a particular type of traffic, so that, for example, the majority of terminals are dedicated to carrying speech traffic (telephones) or to carrying facsimile data (fax machines). The terminals all connect to Central Offices (COs) via access lines, and these COs are interconnected into a network.



Key: ___ Access Lines
 === Trunk Links (inter-CO user data links)
 ... Inter-CO signaling network links
 CO Central Office (Telephone Exchange)

Figure 2

Communications between the terminals are all "circuit switched", so a dedicated synchronous data path (or circuit) needs to be placed between the end terminals for carrying all communications. Arranging for such a circuit to be made or removed (cleared) is the responsibility of the Central Offices in the network. A user makes a request via his or her terminal, and this request is passed on to the "local" Central Office. The relationship between the terminals and the local Central Offices to which they are connected is strictly Client/Server.

The COs are interconnected using two different types of connections. One of these is called a trunk connection (shown as a double line in the above figure) and is used to carry the data traffic generated by the terminals. The other connection acts as part of a separate network (and is shown as a dotted line in the above figure). This is the signaling network, and is used by the Central Offices to request a connection to be made between themselves and the destination of the required circuit. This will be carried across the trunk link to the "next" Central Office in the path. The path, once in place through the PSTN, always takes the same route. This contrasts with the Internet, where the underlying datagram nature of the infrastructure means that data packets are carried over different routes, depending on the combined traffic flows through the network at the time.

The call set up process can be viewed as having two parts: one in which a request for connection is made, and the other in which the circuit is made across the PSTN and call data flows between the communicating parties. This is shown in the next pair of figures (3a and 3b).

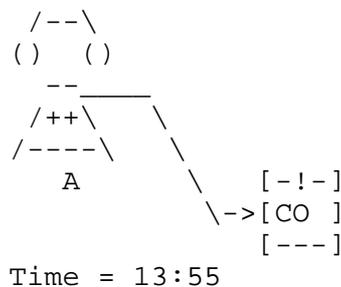


Figure 3a

- Key: ___ Access Lines
- === Trunk Links (inter-CO user data links)
- ... Inter-CO signaling network links
- CO Central Office (Telephone Exchange)

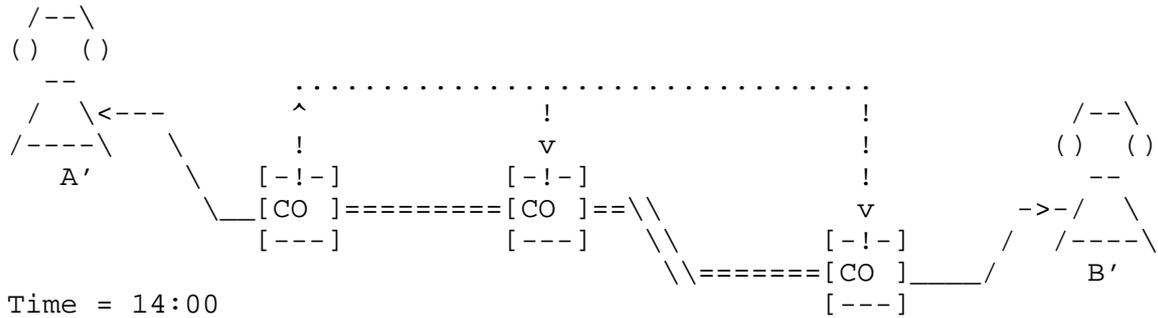


Figure 3b

Figure 3 shows a particular kind of service that can be provided; call booking. With this service, a request is sent for a connection to be made between the A and B telephones at a specified time. The telephone is then replaced (the request phase is terminated). At the specified time, the CO will make a connection across the network in the normal way, but will, first, ring the "local" or A' telephone to inform the user that his or her call is now about to be made.

For more complex services, the requesting telephone is often connected via its "local" CO to a Service Node (SN), where the user can be played prompts and can specify the parameters of his or her request in a more flexible manner. This is shown below, in Figures 4a and 4b. For more details of the operation of the Service Node (and other Intelligent Network units), see the Appendix.

When the SN is involved in the request and in the call setup process, it appears, to the CO, to be another PSTN terminal. As such, the initial request is routed to the Service Node, which, as an end system, then makes two independent calls "out" to A' and B'.

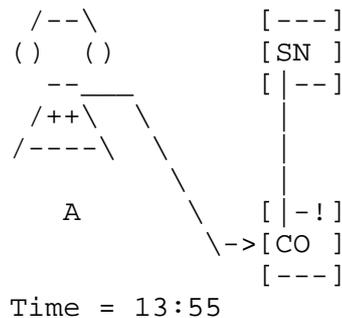


Figure 4a

Key: ___ Access Lines
 === Trunk Links (inter-CO user data links)
 ... Inter-CO signaling network links
 CO Central Office (Telephone Exchange)
 SN Service Node

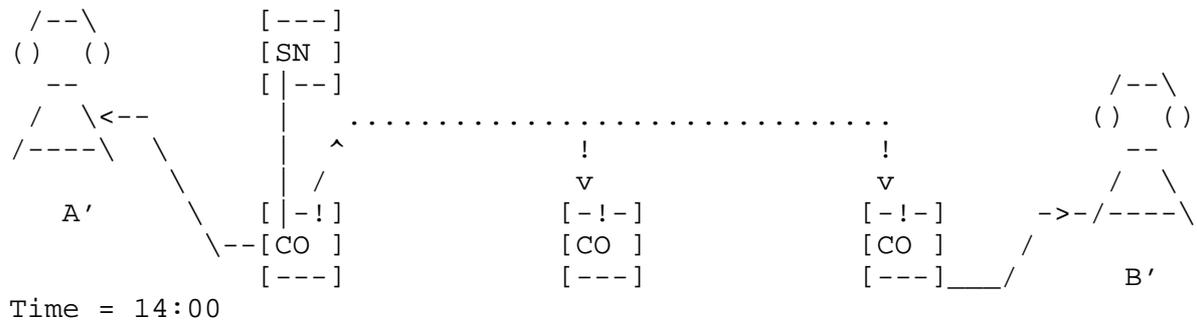


Figure 4b

Note that in both cases as shown in Figures 3 and 4 a similar service can be provided in which the B' telephone is replaced by an Intelligent Peripheral (or an Special Resource Functional entity within a Service Node), playing an announcement. This allows a "wake up" call to be requested, with the Intelligent Peripheral or Service Node Special Resource playing a suitable message to telephone A' at the specified time. Again, for more details of the operation of the Special Resources (and other Intelligent Network units), see the Appendix.

4.2 Pre-PINT Systems

Although the pre-PINT systems reported here (i.e., those developed by AT&T, Lucent, Siemens and Nortel) vary in the details of their operation, they exhibit similarities in the architecture. This section highlights the common features. Specific descriptions of these systems will follow.

All of the systems can be seen as being quite similar to that shown in the following diagram. In each case, the service is separated into two parts; one for the request and another for execution of the service. Figure 5 summarizes the process.

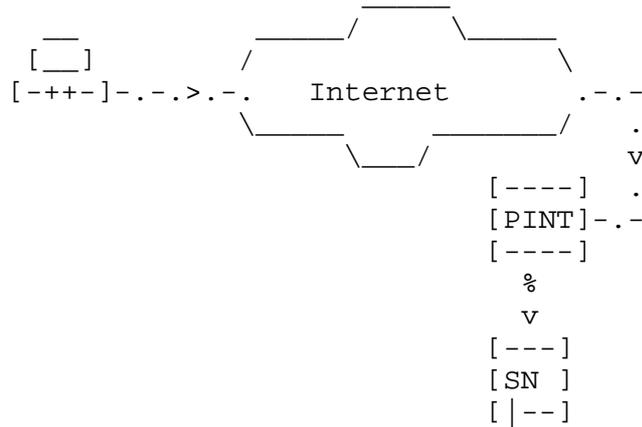


Figure 5a

- Key: CO Central Office (Telephone Exchange)
 SN Service Node
 PINT PSTN/Internet Gateway
 .-.-. Internet Access Link
 %%% Gateway/Service Node Link
 ___ PSTN Access Lines
 === PSTN Trunk Links (inter-CO user data links)
 ... Inter-CO signaling network links

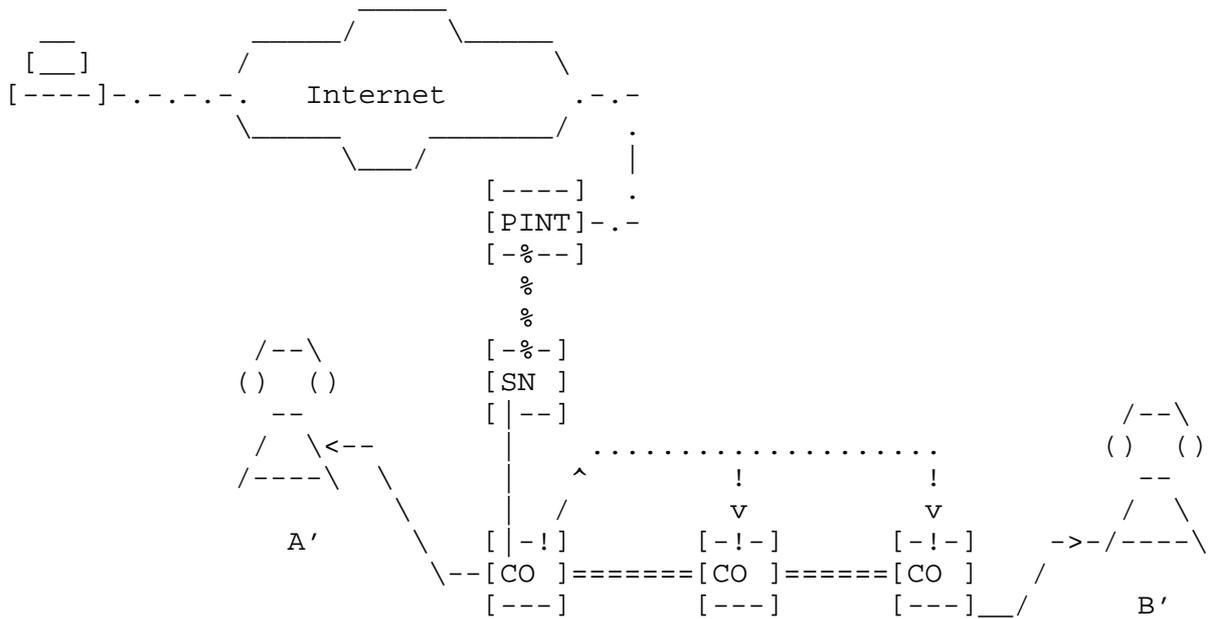


Figure 5b

Comparing Figure 4a with Figure 5a, the differences lie in the way that the information specifying the request is delivered to the Service Node. In the PSTN/IN method shown in the earlier diagram, the user connects to the SN from the telephone labeled A, with the connection being routed via the CO. In the latter case, the request is delivered from an Internet node, via the PINT gateway, and thence to the Service Node over a "private" link. The effect is identical, in that the request for service is specified (although the actual parameters used to specify the service required may differ somewhat).

The figures depicting the respective service execution phases (Figures 4b and 5b) show that the operation, from the IN/PSTN perspective, is again identical. The Service Node appears to initiate two independent calls "out" to telephones A' and B'.

The alternative systems developed by AT&T and by Nortel allow another option to be used in which the PINT Gateway does not have to connect to the PSTN via a Service Node (or other Intelligent Network component), but can instead connect directly to Central Offices that support the actions requested by the gateway. In these alternatives, the commands are couched at a "lower level", specifying the call states required for the intended service connection rather than the service identifier and the addresses involved (leaving the Intelligent Network components to coordinate the details of the service call on the gateway's behalf). In this way the vocabulary of the commands is closer to that used to control Central Offices. The difference really lies in the language used for the services specification, and all systems can use the overall architecture depicted in Figure 5; the only question remains whether the Intelligent Network components are actually needed in these other approaches.

The interfaces are:

- A The interface over which Internet requests for service are delivered to the Service Node
- B The interface over which Service Management requests are sent from the Internet to the Service Management System
- C The interface over which the Service Node sends call control requests to a connected Central Office
- D The interface over which the Service Management System manages the Service Node
- E The interface over which Internet requests for service are delivered to the Service Control Point
- F The interface over which the Service Control Point sends service call control requests to the Mobile Switching Center
- G The interface over which the Service Control Point sends service control requests to the Service Switching Point
- H The interface over which the Service Management System manages the Service Control Point
- I The interface over which the Service Node sends service call control requests to the Mobile Switching Center

In practice, a number of the interfaces have very similar purposes to one another. The means by which these purposes are achieved differ, in that some of the interfaces (C and I) reflect access arrangements, whilst others (F and G) imply a "core" signaling relationship. However, it is possible to categorize them in terms of the "intent" of messages sent across the interfaces.

For example, Interfaces A and E are similar; one of the main aims of PINT work is to ensure that they are the same. Similarly, Interfaces D and H imply similar actions and are likely to carry similar messages. Interfaces C, F, G, and I are all used to request that a call be initiated, albeit via access or core signaling relationships.

The interfaces can also be viewed in terms of the kind of components that are involved and the bodies by which they are codified. Interfaces A, B, and E are all going to be realized as Internet Protocols. All of the others use existing protocols in the PSTN/IN. Traditionally, these have been codified by different groups, and this is likely to be the case in the PINT work.

The general arrangements for the different systems are shown below (Figures 7, 8, 9, and 10). They differ in the details of their configurations, but the main tasks they perform are very similar, and so the overall operation is similar to the generic architecture shown in Figures 5 and 6.

Key for following diagrams:

Components:

W3C World Wide Web Client
 W3S World Wide Web Server
 WSA Web Server "Back End Program" Interface (CGI or Servlet interface)
 Srvlt Servlet "back end" program/objects
 FS Finger Server
 SCTPC Simple Computer Telephony Protocol Client
 SCTPS Simple Computer Telephony Protocol Server
 CBC CallBroker Client
 CBS CallBroker Server
 SSTPC Service Support Transport Protocol Client
 SSF Service Switching Function
 SCF Service Control Function
 SRF Special Resource Function
 CO Central Office/ Public Telephone Exchange
 SSP Service Switching Point
 SCP Service Control Point
 SR/I.IP Special Resource/ "Internet" Intelligent Peripheral
 SMS Service Management System
 INAPAd Intelligent Network Application Part Adaptor
 PktFlt Packet Filter (Firewall)
 SNMPAg Simple Network Management Protocol Agent

Protocols:

P0 HyperText Transfer Protocol
 P1 HTTP Server <-> "Back End Program" internal protocol
 P2 CallBroker Client <-> CallBroker Server protocol (AT&T system),
 or SCTP Client <-> Server protocol (Nortel system)
 P3 PINT User Agent <-> PINT Gateway protocol
 P4 Intra-Intelligent Network protocol (e.g., INAP)
 P5 Proprietary (INAP-based) Gateway-> I.IP protocol
 P6 Finger protocol
 P7 Digital Subscriber Signaling 1 protocol
 P8 Simple Network Management Protocol
 P9 SMS <-> Service Control Point/Service Node protocol

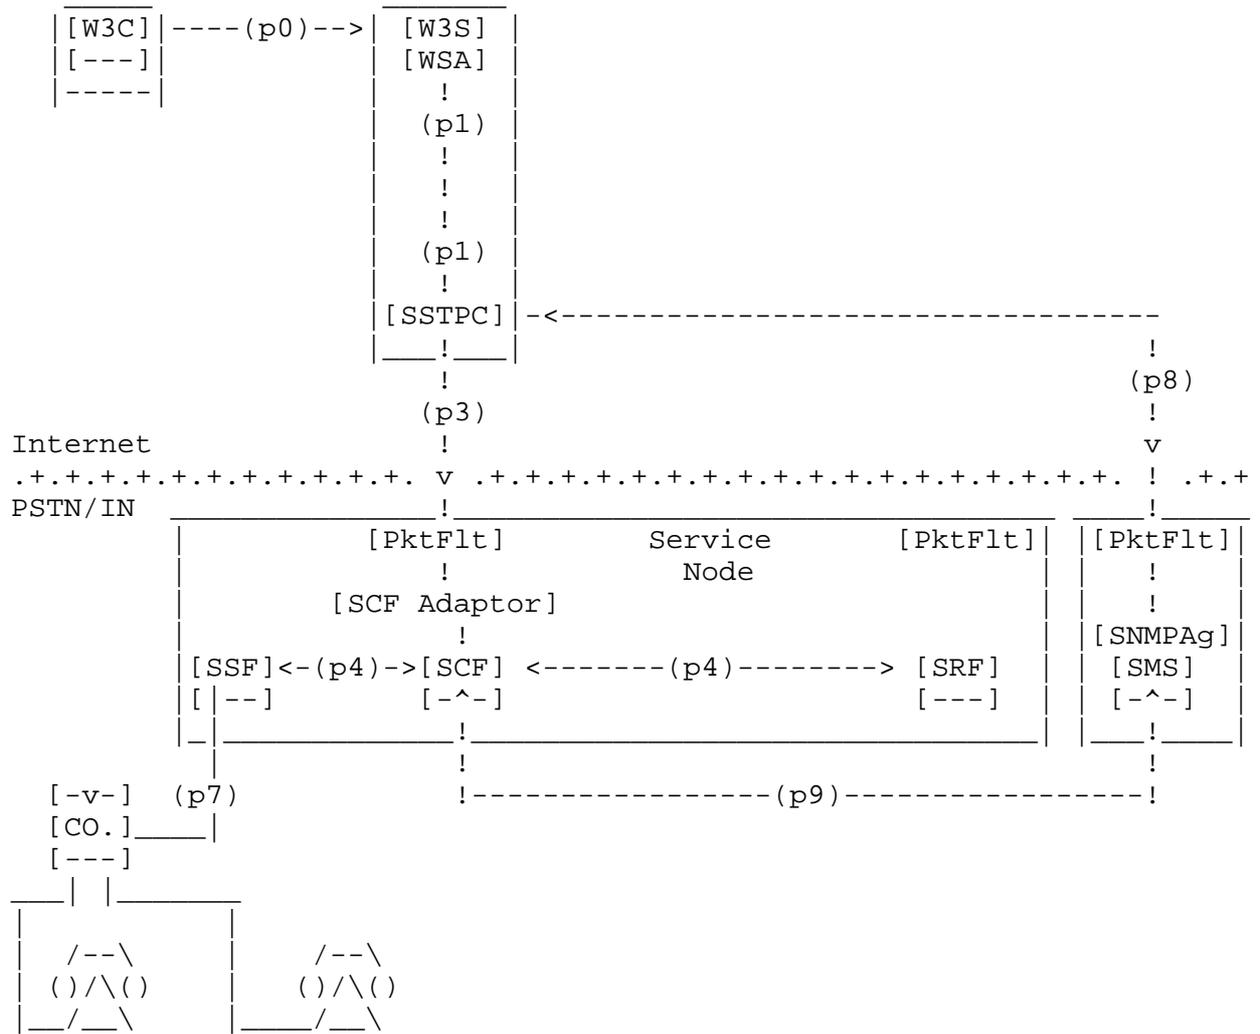


Figure 8: The Lucent System

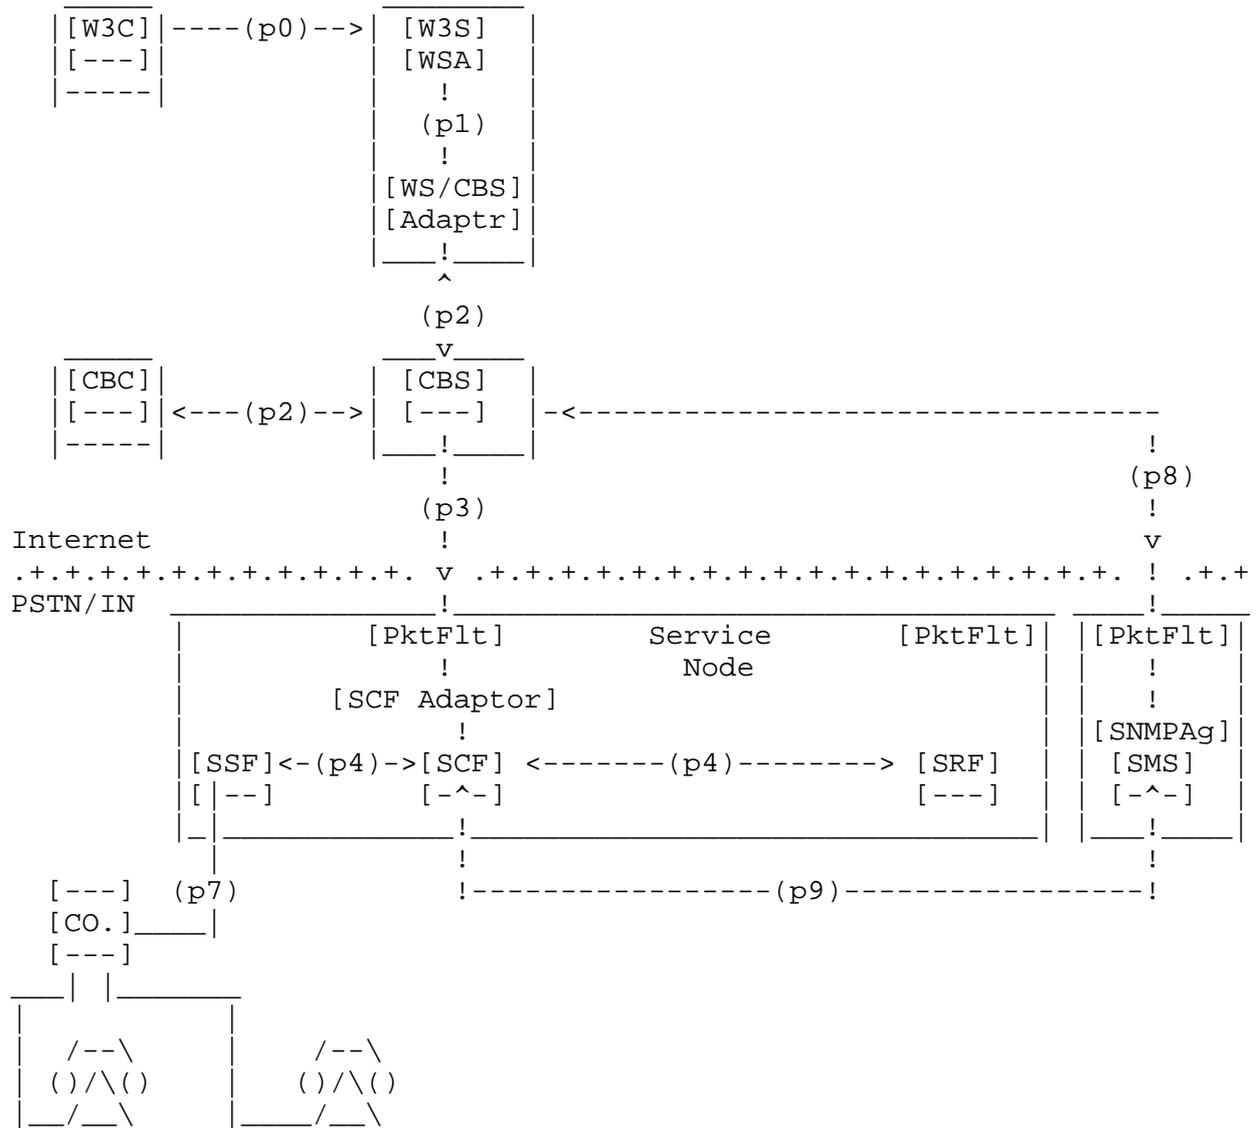


Figure 9: The AT&T System

Several Intelligent Network Functional Entities are combined into a Service Node in the Lucent, AT&T, and Nortel systems, while in the Siemens scheme they are separate units. However, this is not particularly important for the provision of the services they offer.

The main difference lies in whether or not the SCF is "aware" of the Internet interface and has been modified to be "complicit" in supporting these Internet requests. The Siemens approach was to re-use an existing SCP, providing a gateway function to translate as needed. The Lucent system used a "lighter weight" SCF adapter to terminate the Internet protocols, as the SCF was modified to support the Internet interface directly.

The AT&T CallBroker and Nortel SCTP Servers introduce an intermediate protocol (labeled p2) that allows an alternative to the Web based interface supported by the others. This protocol matches the "CallBroker Client API", or the "SCTP Client API". These options provide for a bi-directional protocol, with indications sent from the Call Broker or SCTP Server to the Client as needed. This is not easily possible using an HTTP-based scheme (and in the Siemens case, a dedicated Finger client/server pair was used to emulate such an interface)

The protocol between the Internet server and the Intelligent Network (labeled p3 in the above diagrams) differs in each of the systems. One of the main aims of future work will be to develop a common protocol that will support the services offered, so that the p3 interface will allow different implementations to inter-operate. In the Lucent, Siemens, and Nortel systems, this was an "internal" protocol, as it was carried between entities within the Service Node or Gateway.

Other contrasts between the systems lie in the support for Internet access to Service Management, and access to the Internet by Special Resources. Internet Management access was most developed in the Lucent system, in which a Simple Network Management Protocol (SNMP) agent was provided to allow inter-operation with the SMS controlling the Service Node. In the Siemens scheme, the SMS had no direct Internet access; any management actions were carried out within the normal PSTN management activities. As for Internet access to special resources, this was only required by the Siemens system as part of its support for Call Center agent notification. Equivalent functionality would be provided in the AT&T and Nortel systems as mentioned above, and this would in turn be associated with event notifications being sent as part of their (p3) Internet/IN protocol. These differences reflect the different emphases in the products as they were developed; again, future work will have to ensure that common protocols can be used to support the chosen services fully.

5. IN-Based Solutions

5.1 The Lucent System

Figure 11 depicts the overall interconnection architecture of the Lucent prototype in support of the four PINT services. The IN-based architecture utilizes the Service Node and Service Management System in addition to the Web server, which enables Web-based access to the PINT services. This section summarizes the roles of these elements (complemented by a click-to-dial-back service scenario), outlines the interfaces of Web Server-Service Node and Web Server-Service Management System (i.e., the interfaces A & B), and addresses the common security concerns.

5.1.1 Roles of the Web Server, Service Node, and Service Management System

Web Server

The Web Server stores the profiles of content providers as well as pre-registered users. The content provider profile contains information such as content provider ID, telephone number, and fax number. In addition, the profile may also include service logic that specifies, for example, the telephone (or fax) number to be reached based on time of the day, day of the week, or geographical location of the user, and the conditions to accept the charge of the calls.

Similar to the content provider profile, the pre-registered user profile contains information such as user name, password, telephone number, and fax number. The last two pieces of information can also be linked to time of the day and day of the week so the user can be reached at the appropriate telephone (or fax) number accordingly.

Service Node

Situated in the PSTN, the SN, like the SCP, performs the service control function [1, 2, 3]. It executes service logic and instructs switches on how to complete a call. The SN also performs certain switching functions (like bridging of calls) as well as a set of specialized functions (like playing announcements, voice recognition and text-to-speech conversion).

Service Management System

The SMS performs administration and management of service logic and customer-related data on the SN. It is responsible for the replication of content provider profiles and provision of these data on the SN. These functions are non-real time.

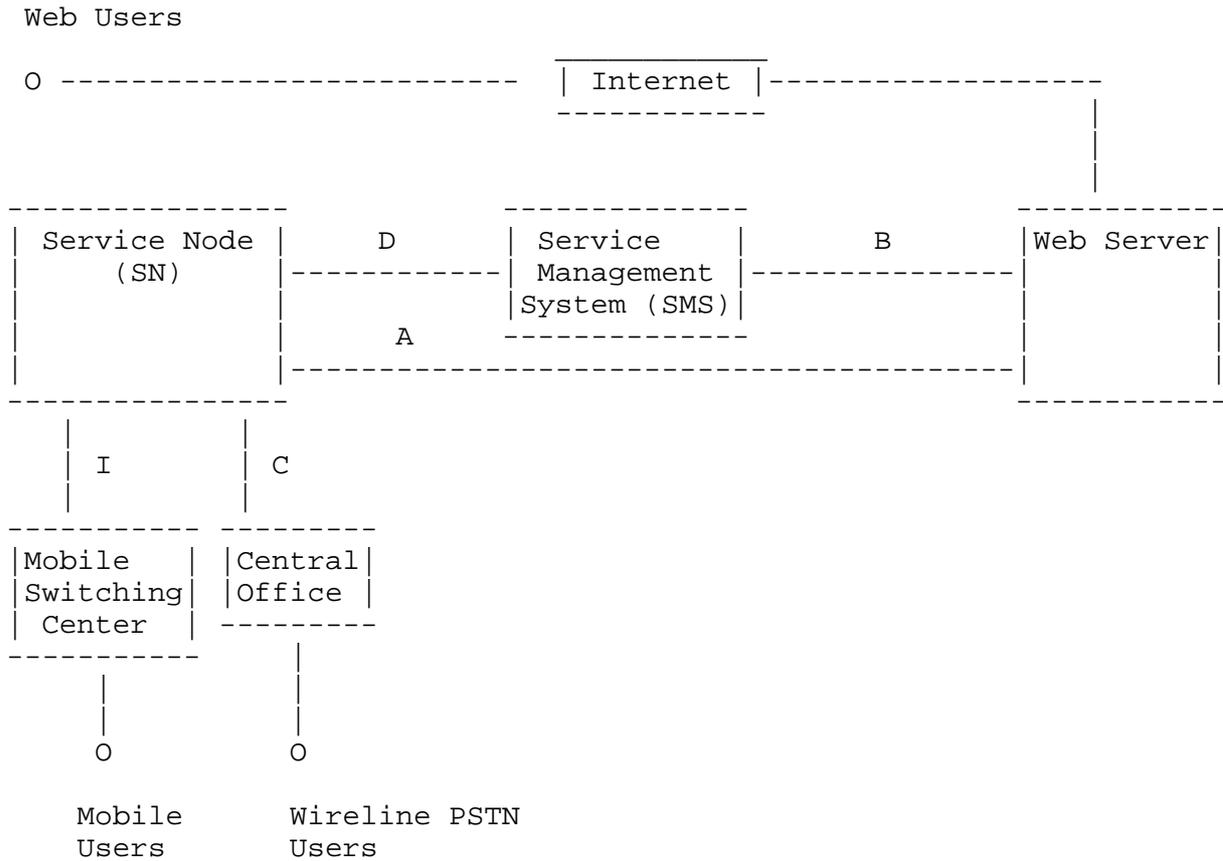


Figure 11: Overall Interconnection Architecture of the Lucent System

5.1.2 A Click-to-Dial-Back Service Scenario

A Web user, who has simultaneous access to the Web and telephone services (this can be achieved, for example, by having an ISDN connection), is browsing through a sales catalogue and deciding to speak to a sales representative.

When the Web user clicks a button inviting a telephone call from the sales office, the Web Server sends a message to the SN over the A interface, thus crossing the Internet-to-PSTN boundary. By matching the information received from the Web Server with the content provider profile that had been previously loaded and activated by the SMS over the D interface, the SN recognizes the signal.

At this point, the SN calls the Web user. The user answers the call, hears an announcement, e.g., "Please wait, while we are connecting you to the sale agent", and is waiting to be connected to the sale agent. Then the SN invokes service logic as indicated in the profile.

The execution of this logic selects an appropriate sales agent to call based on the time of the day. It is 8 P.M. in New York where the Web user is located, and the New York sales office has closed. The San Francisco office, however, is still open, and so the SN makes a call to an agent in that office. Finally, the SN bridges the two calls and establishes a two-party call between the sales agent and the Web user.

5.1.3 Web Server-Service Node Interface

Lucent developed the Service Support Transfer Protocol (SSTP) for communications between the SN and Web Server. SSTP is of a request/response type running on top of a reliable transport layer, such as TCP. The Web Server sends a request to the SN to invoke a service and the SN responds with a message indicating either success or failure. Note that SSTP engages only the service control function [1, 2, 3] of the SN.

5.1.3.1 Web Server to Service Node

In this direction, three kinds of messages may be sent: the Transaction Initiator message, the Data Message, and the End of Data message.

The latter two messages are needed if the service to be invoked involves data (such as the case in click-to-fax, click-to-fax-back and voice-access-to-content). This was so designed to handle the varying size of data and to ensure that the size of each stream is within the allowable size of the underlying transport packet data unit (imposed by some implementations of TCP/IP).

a. Transaction Initiator

This message provides all the necessary information but data for invoking a service. It includes the following information elements:

- + Transaction ID, which uniquely specifies a service request. The same transaction ID should be used for all the accompanying data-related messages, if the service request involves data. One way for generating unique transaction IDs is to concatenate the information: date, time, Web Server ID (uniquely assigned for each one connected to the SN), and transaction sequence number (a cyclic counter incremented for each service request).

- + Service ID, which specifies the service to be invoked. The service may be click-to-dial-back, click-to-fax, click-to-fax-back or voice-access-to-content.

- + Content Provider ID, which uniquely represents the content provider. This information is the key to accessing the content provider's service logic and data on the SN.
- + Content Provider Directory Number, which is the telephone or fax number of the content provider to be called through the PSTN.
- + User Directory Number, which is the telephone or fax number of the user requesting the service.
- + Billed Party, which specifies the party (either the user or content provider), to be billed.

In addition, optional parameters may be sent from the Web Server to the SN. For example, a retry parameter may be sent to specify the number of times the SN will attempt to complete a service request upon failure before the transport connection times out.

b. Data Message

This message provides the (encapsulated) user data part of a service request. For example, in the case of click-to-fax-back such data are the content to be faxed to the user. Each message is composed of the transaction ID and a data segment. The transaction ID must be the same as that of the transaction initiator part first invoking the service.

c. End of Data Message

This message contains the transaction ID and the end of data delimiter. The transaction ID is the same as that of the relevant transaction initiator message.

5.1.3.2 Service Node to Web Server

The SN must respond to a service request from the Web Server. The response message consists of the information elements:

transaction ID, service type, result, time, and error code.

- + Transaction ID, which is the same as that of the original service request.
- + Service Type, which is the same as that of the original service request.
- + Result, which is either success or failure.

- + Time, which indicates the time of the day completing the request.

- + Error Code, which gives the reason for failure. Possible reasons for failure are content provider telephone (or fax) busy, content provider telephone (or fax) no answer, user telephone busy, user refusal to complete, user no answer, nuisance control limit reached, and content provider telephone (or fax) not in the SN database.

5.1.3.3 Usage Scenarios: Click-to-Fax and Click-to-Fax-Back

For the click-to-fax and click-to-fax-back services, the Lucent system implemented only the case where the data to be sent as facsimile reside in the Web server. There are at least three messages that need to be sent from the Web server to the Service Node for these services.

The first message is the Transaction Initiator that identifies the service type as well as a unique Transaction ID. It also includes the sender/receiver fax number.

The next is one or more messages of the data to be faxed. Each message carries the same unique Transaction ID as the above.

Last comes the end of message. It consists of the Transaction ID (again, the same as that of the messages preceding it) and the end of data delimiter.

Upon receiving these messages, the Service Node, equipped with the special resource of a fax card, converts the data into the G3 format, calls the receiver fax, and sends back the result to the Web server immediately. Note that the receiver fax busy or no answer is interpreted as failure. Further, while the receiver fax answering the call is interpreted as success, it does not necessarily mean that the fax would go through successfully.

5.1.4 Web Server-SMS Interface and SNMP MIB

This interface is responsible for uploading the content provider profile from the Web Server to the SMS and for managing the information against any possible corruption. The SN verifies the Content Provider ID and the Content Provider Directory Number sent by the Web Server with the content provider profile pre-loaded from the SMS.

The content provider profile was based on ASN.1 [4] structure and SNMP [5] was used to set/get the object identifiers in the SMS database.

Following is an example of the simple MIB available on the SMS.

```

inwebContProviderTable OBJECT-TYPE
    SYNTAX          SEQUENCE OF InwebContProviderEntry
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION
        " A table containing Content Provider profiles "
    := { inweb 1 }

inwebContProviderEntry OBJECT-TYPE
    SYNTAX          InwebContProviderEntry
    MAX-ACCESS      not-accessible
    STATUS          current
    DESCRIPTION
        " A conceptual row of the inweb. Each row
          contains profile of one Content Provider"
    INDEX          { inwebSmsNumber }
    := { inwebContProviderTable 1 }

InwebContProviderEntry := SEQUENCE {
    inwebSmsNumber          Integer32,
    inwebContentProviderId Integer32,
    inwebContentProviderPhoneNumber Integer32,
    inwebContentProviderFaxNumber Integer32
}

inwebSmsNumber OBJECT-TYPE
    SYNTAX          Integer32
    MAX-ACCESS      read-only
    STATUS          current
    DESCRIPTION
        " Serial number of the SMS - used for SNMP indexing "
    := { inwebContProviderEntry 1 }

inwebContentProviderId OBJECT-TYPE
    SYNTAX          Integer32
    MAX-ACCESS      read-create
    STATUS          current
    DESCRIPTION
        " A number that uniquely identifies each Content
Provider "
    := { inwebContProviderEntry 2 }

inwebContentProviderPhoneNumber OBJECT-TYPE
    SYNTAX          Integer32
    MAX-ACCESS      read-create
    STATUS          current

```

DESCRIPTION

```

" Content Provider's Phone Number "
:= { inwebContProviderEntry 3 }

```

```

inwebContentProviderFaxNumber OBJECT-TYPE

```

```

SYNTAX          Integer32
MAX-ACCESS      read-create
STATUS          current
DESCRIPTION

```

```

" Content Provider's Fax Number "
:= { inwebContProviderEntry 4 }

```

5.1.5 Security Considerations

The Lucent prototype addressed the security issues concerning the interface between the Web Server and the SN. Those concerning the interface between the Web Server and SMS, which was based in SNMP, were handled by the built-in security features of SNMP.

+ Secure Communication Links

If the Network Operator (PSTN provider) is also the Web Service provider, the Web Server and SN/SMS will communicate over a corporate intranet. This network is almost always protected by the corporation's firewall and so can be deemed secure. This was the case handled by the Lucent prototype.

Nevertheless, if different corporations serve as the Network Operator and the Web Service Provider, then it is likely that there may not exist a dedicated secure communication link between the Web Server and SN/SMS. This raises serious security considerations. One possible solution is to use Virtual Private Networks (VPN). VPN features support authentication of the calling and called parties and encryption of the messages sent over insecure links (such as those on the Internet).

+ Non-Repudiation

All transactions were logged on both the Web Server and the Service Node to account for all operations in case of doubt or dispute. The log information on the SN may also be used to generate bills.

+ Malicious Requests of Users

A user may make repeated requests to a content provider directory number maliciously. This scenario was handled by setting a Nuisance Control Limit (NCL) on either the SN or the Web Server or both. The NCL has two parameters: one defining the number of requests from a

user and the other the period over which these requests takes place.

A user may also attempt to request a call from a directory number other than that of a content provider. This scenario was handled by verifying the directory number (and the content provider ID) against the database on the SN containing all the content provider information. If the directory number (or the content provider ID) was not in the database, the request would be rejected.

5.2 Siemens Web Call Center

5.2.1 Service Description

The Web Call Center is an Intelligent Network System that accepts requests from Internet nodes for services to be provided on the PSTN. As the name suggests, it was designed to support a cluster of services that, taken together, provide a subset of the features of a Call Center, with almost all user interactions provided via World Wide Web requests and responses. See the appendix for a background description of Call Center Features.

From an Intelligent Network perspective, there are a number of services that, when combined, provide the Call Center features. The Call Center features as implemented supported the scenario in which a customer makes a request to be called back by an agent at a time of the customer's choosing to discuss an item of interest to him or her. The agent will be selected based on his or her availability and expertise in this topic; the agent will be told whom he or she is calling and the topic of interest, and then the agent will be connected to the customer.

In addition, the individual services that were deployed to support this scenario provided support for management of the list of available agents as well. This involved allowing the agent to "log into" and "out of" the system and to indicate whether the agent was then ready to handle calls to the customer. The list of services, as seen from a user perspective, follows.

The services support:

i) Customer Request service - the customer explores a corporate Web site, selects a link that offers to request an agent to call the customer back and then is redirected to the Web Call Center server. This presents customer with a form asking for name, the telephone number at which he or she wishes to be called, and the time at which the call is to be made. Note will also be made of the page to which the customer was referred to when he or she was redirected. Once the form has been returned, the customer receives an acknowledgment page

listing the parameters he or she has entered.

ii) Agent Registration/Logon - An agent requests a "login" page on the Web Call Center server. The service checks whether it has a record of an agent present at the Internet node from which the call is made. If not, then the caller will be sent a form allowing him or her to enter the service identity, the company's agent identifier and password. On return, the service identity and company agent identifier will be checked against a list of known identities. If found, the password will be checked, and if this matches the record held by the service then a new session record is made of this identity and the Internet node from which the call has been made.

NB: This is very similar to the Universal Personal Telecommunications (UPT) service feature "register for incoming calls". It implies that the identified person has exclusive use of the Internet node from that point onwards, so messages for them can be directed there.

iii) Agent Ready - an agent who has already logged on can indicate that he or she is ready by requesting an appropriate "ready" page on the Web Call Center Server. The service will match the agent by the Internet node Identifier and Agent Identity passed along with the Web request against its list of "active" agents. It will mark them as being ready to handle calls in its list of available agents (with their pre-defined skill set).

iv) Agent Not Ready - an agent can request an appropriate "ready" page on the Web Call Center Server to indicate that he or she is temporarily not ready to handle calls.

v) Agent Logoff - an agent can request an appropriate "Logout" page on the Web Call Center Server to indicate that he or she is no longer associated with a particular Internet node. The service will match the agent by the Internet Node Identifier and Agent Identity passed along with the Web request against its list of "active" agents. Once found, the session record for that agent is removed and the caller is notified of this with an acknowledgment page.

NB: This is very similar to the UPT "unregister" service feature.

vi) Call Center Agent Selection and Notification - When the time that the customer selected has arrived and an available agent with the right skills has been selected from the appropriate list, this service will send a notification to the Internet node associated with that agent. A dedicated server is assumed to be running on the agent's machine that, on receiving the notification, triggers the agent's browser into requesting a "Agent Call In" page from the Web Call Center Server. Once the agent's machine has made this request,

he or she will be told that there is a customer to call.

NB: This is similar to a "Message Waiting" or "Wake Up Call" service.

Note: As implemented, the agent is led automatically into the following service (the returned Web page includes an automatic reload command).

vii) Agent Instruction - a selected agent makes a request of the "Customer Processing" page on the Web Call Center Server. The Internet node Identifier and Agent Identity the agent uses will be matched against a list of agents expected to handle calls, and the instructions for the calls will be returned to the agent.

NB: This is similar to a "Voice Mail Replay" message service, but in this case the message is automatically generated; there is no associated voice mail record feature accessible.

Note: As implemented, the instructions page will include a number of buttons, allowing the agent to view the page the customer was looking at when he or she made the request, and to trigger the customer callback (as described next).

ix) Agent/Customer Telephony Callback - the agent will make a request of a "dial-back" page on the Web Call Center Server. The Internet node Identifier and Agent Identity he or she uses will be matched against a list of agents expected to handle calls, and, when the appropriate records have been found, the service will make the telephone call through to the customer and then connect the agent to this telephone call (using the telephone number registered in the respective Call Center service record).

5.2.2 Implementation

5.2.2.1 Introduction

The Siemens Web Call Center used an existing IN system and service logic that supported Call Center features. The scenario it supports is very similar to the Siemens IN-based Call Center on which it was based; one of the goals was to minimize changes to the service offered. It is also virtually identical to the service "Internet Requested Telephony Dial-back" provided by the Lucent system.

As provided via the Internet, the services involved are mostly the same as those provided via the PSTN and IN alone. The main differences lie in the use of the World Wide Web as an interface to the services rather than a telephone, SSP, and Intelligent Peripheral. Also, the feature by which a telephone call is made

between the agent and the customer is implemented within the IN system in a different way; this is the only element in which the PSTN is involved.

5.2.2.2 Web Call Center Configuration

The general arrangement for the Web Call Center system is shown in Figure 7. The components that were added to an existing IN system to deal with the Internet interface are described next.

In addition to the SCP, SSP and SMS that were part of the original IN-based system, another unit was included to send notification messages to agents; in the IN system the agents were sent "wake up" telephone calls when they were required to handle their next customers' call back. This unit is called the "Internet Intelligent Peripheral", and its use is described later under "Non-World Wide Web Interactions".

As there was a need to re-use as many of the existing IN components unchanged, a Gateway unit to deal with the interface between the Internet and the SCP was provided. This injected INAP (Intelligent Network Application Protocol) messages into the SCP, making it think that it had received an Initial DP trigger from an SSP. It also intercepted the "Connect To Resource" and "Prompt and Collect" INAP messages sent from the SCP, acting on these to return the parameters generated by the Internet users when they filled in the forms that triggered the service transaction. It also translated the "Play Announcement" message sent to the Intelligent Peripheral into a form that it could use. Finally, it passed on the INAP message used by the SCP to trigger SSP into making the telephone call back.

5.2.2.3 User Interaction

In the IN/PSTN-based system, the services have contact with the customers and agents via their telephones, SSPs, and Intelligent Peripherals programmed to play announcements to them and to capture their responses. These responses are indicated by DTMF tones sent by pressing keys on the telephones.

In this case, almost all interactions are provided via World Wide Web requests and responses. The sequence of announcements and responses for each service are "collapsed" into individual form filling transactions, and the requests are not limited to digits (or "star" and "hash"). The implications of the use of forms on service operation are covered in more detail later (under HTTP/IN Service mapping).

5.2.2.4 Service/Caller Identifiers

When provided via the IN/PSTN-based system, the services are passed the Calling Line Identity (CLI) of the caller and the number the caller dials (the DN). The CLI value is used extensively to identify the caller and (in the case of the agent) to index into service data tables to decide what to do next. While an equivalent value to the DN is passed to the Web-based transactions as the requested Universal Resource Locator (URL), the CLI cannot be given reliably. The nearest equivalent caller identifier is the IP Address of the customer or agent's machine. However, the use of HTTP proxies means that this "original" Internet node Address may not be available; if a proxy is used then its IP Address will be associated with the request.

In providing these Call Center features the customer only has one Web-based transaction; that of providing the initial request for a PSTN telephone callback. To do so he or she will have to fill in a form so as to specify not only the time to be called back, but also the telephone number to be reached. These values can be used if needed to identify the customer, and so the problem of originating Internet Node ambiguity is not relevant.

With the agents, however, there are sequences of coupled transactions, and the particular sequence must be identified. There will be a number of such transactions being carried out at once, and there needs to be some identifier to show which agent is being handled in each case.

Such an identifier is not part of a sequence of basic Web transactions. In a Web transaction, the HTTP Client/Web Browser makes a request, and the HTTP Server will respond to this, normally including some content in its reply message that will be processed by the browser, after which it closes the TCP connection. That's the end of the transaction; the HTTP client and server cannot normally maintain state information beyond this point. Any sequence is reduced to a set of unrelated transactions.

A result of this simple pattern is that any state information reflecting longer or more complex interactions must be stored (at least partially) in the client system. One approach is the use of cookies [6]. These can be set by HTTP servers as part of their response to a request, and will be sent back with all subsequent requests for appropriate URLs as extra HTTP headers. These cookies allow the HTTP server to identify the client in the following requests, so that it can continue an extended session with the client.

Cookies are used in providing the Internet Call Center. Persistent cookies are installed into the Web Browser on machines that are to be used by call center agents as a service management (pre-service) task. The cookie value is unique to the machine and is used to index into a list of machine IP addresses that is stored as part of the service data.

Also, a session cookie is stored onto the agent's machine when the agent registers, and is cleared when he or she de-registers. This is used to identify the agent and so the IP address of the node with which the agent is associated (and from which the agent's subsequent requests should originate). The services that interact with Call Center agents use the agent session cookie value as an identifier; in principle this is unnecessary but it does simplify the session data lookup procedure. The rest of the services use the persistent machine identifier in place of the CLI, indexing into their service data using it. Both cookies are sent with each agent request; if they are not present, then the request is redirected to other services (for example to the agent Logon service).

5.2.2.5 Mapping from HTTP Transactions to IN-Based Service Features

All of the client-initiated services require user interaction. With the IN/PSTN-based system, the majority of the services are typified by the callers being connected to an announcement unit that plays them a list of choices and captures their selection. The caller can pre-dial the digits needed; in this case the prompts are not needed and are not made.

The pattern of operation is somewhat different in the Internet case, as the initial HTTP request returns a response, after which the Web transaction has ended. Where that initial response returns a form to be filled in by the caller, subsequently submitting the form initiates a new HTTP transaction. This is all part of one instance of service, however. The service consists of two request/response pairs in tandem.

Although it is possible to design a service to handle this pair of Web transactions as a single unit, it may be better to reconfigure it. The design of a service that deals with two Web exchanges as a single extended transaction is quite complex. It must maintain state across the pair of Web exchanges, and it has to handle a number of failure cases including dealing with time-outs and "out of time" submission of forms. The alternative is to split the service into two sub-features. The first of these reflects the initial request and delivery of the form by return, with the second one dealing with processing of the submitted form and returning any confirmation by reply.

The services offered don't all require form-filling, and so can be treated as a single IN feature. There are two cases where forms are required. The first of these is the Customer Request service, while the other one is the "Agent Registration" service. In both cases the initial Web transaction (by which the form is requested and returned to the client) need not involve specific service logic processing; the initial delivery of the form to a customer or agent can be handled by a "normal" Web Server. In both cases the service logic is only triggered when the form is submitted; this means that, again, each of the services can be treated as a single IN feature.

The IN service logic that deals with these requests has a general pattern of action. An HTTP request is received, and this triggers the IN service logic into action. The service logic "sees" this as an Initial DP message and starts its processing as if it had been sent from an SSF. The SCF uses what appears to it to be an Intelligent Peripheral to collect the parameters of the request, and then to send back final announcements to the requesting entity.

The main difference, from the perspective of the IN service logic running on the SCF, is that the service does not need to instruct the SSF to make a temporary connection to the Intelligent Peripheral. It is as if this connection had already been made. Similarly, there is no need to close the service transaction by sending an explicit "Continue Execution" message to the SSF.

The sequence of "prompt/collect" instructions used to collect service parameters from a caller in an IN service maps quite well to a sequence of requests to extract a data value from the HTTP request, based on a tag. This is a fairly standard feature of Web Server CGI or Servlet processing. Using this mapping minimizes the changes to the service design, in that the service logic "sees" an Intelligent Peripheral to which it sends normal "Request Report Prompt & Collect" messages, and from which it receives data values in response.

All services have to fit in with the underlying HTTP interaction pattern, and so will be expected to send a final "Announce" instruction to the Intelligent Peripheral at the end of the service; this is done in many IN services anyway and in all of the service features described here. These announcements form the content returned to the Web Client.

5.2.2.6 Non-World Wide Web Interactions

There are two exceptions to the sole use of the World Wide Web for interaction. The first one occurs in the "Message Waiting"/"Wake Up Call" service by which the selected agent is informed of a callback request. World Wide Web transactions are very simple; the client browser makes a request for content associated with a particular HTTP URL, and the server sends a response, marking the end of the transaction. The server cannot make a spontaneous association with a client; it must be initiated by the client request.

While it would be possible for the server to defer closing an earlier transaction (by not sending back all of the content specified and leaving the TCP connection open) it was decided that an alternative scheme would be more convenient. The "wake up call" was arranged by an "Internet Intelligent Peripheral" sending a request to a daemon process running on the selected agent's machine, using the Finger protocol [7]. The daemon sent back a standard response, but in addition the Web Browser on the agent's machine was triggered into making a further HTTP request of the server. In this way the "Agent Instruction" transaction is started automatically, while still allowing it to use a normal HTTP request/response pattern.

The second exception occurs in the final "Agent/Customer Telephony Callback" service. While this transaction is initiated by the agent selecting a link on the "call instructions page" returned to them, and includes a "confirmation" page being sent back to them in an HTTP response, the purpose of this service is to make a telephone connection via the PSTN between the agent's telephone and the customer's telephone. It is the only service element that involves the PSTN directly. From an IN/PSTN perspective, the resulting telephone connection is different from that provided in the scheme using the IN and PSTN alone. In this case, a PSTN call is made out to the agent's telephone, another call is made out to the customer's telephone, and these calls are bridged. This differs from the earlier scheme, in which the agent originated a call to the voice mail replay system, and this call was redirected to a new destination (the customer's telephone). As this feature differs in purpose from the other services, and it requires a different implementation within the IN and PSTN system, it was organized as a separate service in this case.

5.2.2.7 Security Considerations

In the case of this system, assumptions were made that the interface presented to requesting agents and customers was provided via a fire wall to deal with most attacks on the IN components. The interface appeared as a Web Server, and there was no direct access to the HTTP documents served, nor to the servlets providing the service logic.

The Callback service was deemed to have simpler security requirements than other IN services as it was akin to a free phone "1-800" service access number; the agents work for the service subscriber and are not charged directly. Similarly, the requesting customer is not charged for his or her request, nor for the resulting call back. Service subscribers would be willing to pay the costs of telephone calls generated as a result of this cluster of services, and the costs of running the agent services could be charged directly to them. As such the authorization for service is defined by the contract between the service subscriber and the service provider.

Authentication of agents was seen as a problem. As an interim measure, cookies were used, but this scheme delivers the cookie data as a plain text item (a header of the Web request). Secure Socket Layer connections were required for communication with the agent services, and this had an impact on the performance of the IN system.

5.2.3 Derived Requirements/Lessons

Security is seen as a major issue. A firewall was used to control access to the IN Components. Similarly, SSL was used for communication with the Agents, so as to protect the cookie values that they were sending with their requests.

For other services, it is likely that the entity from which requests appear to originate will be charged for the service to be rendered. This has implications in terms of authentication and authorization of service provision at the time of the request. It is necessary for the service to be authorized in such a way that non-repudiation is ensured; this is likely to mean that a certificate of identity be provided from the person making the request, and that this can be tied in with a financial account that that person has with the service provider. The certificate can then be stored as part of the billing record. While the process of electronic commerce is outside of the scope of this work, the mechanism by which a request for confirmation of identity is passed out to the requesting user and is delivered back to the service logic must be considered.

When changing from a "pure" IN/PSTN system to one supporting requests via the Internet, the differences in the way that clients interacted with the services meant that the service logic had to be redesigned. It was realized that maintaining the state of a service during its processing was going to be a problem; this problem was side-stepped by re-engineering the services as form processors, allowing them to deal with fully specified requests as a single (Web) transaction. In addition, a "normal" Web Server was used to deliver the forms to the users. This is a change from the IN system, where the equivalent of the form (the prompts) were sent in sequence as part of the same service process.

The Call Center features provided suited this change. However, this may not be the case for other IN services. It is quite common for services to be designed such that the user is prompted for a response, and the service continues dependent on this response. The Web form presents all of the options at once, so this kind of variant prompt/collect sequence is not possible. From this, it is difficult to see how an IN service could be reused without some degree of modification.

An intermediate "gateway" system was provided to "cocoon" the service logic as far as possible from the details of the components with which it was working. Where needed, this unit translated calls from the service logic into commands that operated with the Internet (and the Web Server that acted as the interface). Our experience was that an SCP could be "spoofed" into thinking that it was operating with other IN components in the normal way. Within the limits of the service used, this proved simpler than was originally expected.

Selecting this simple approach still allows a considerable range of services to be provided while maintaining any investment in existing IN systems. Modification of existing IN service logic was also easier than feared. All of the services examined provided announcements at the end of the service transaction, and this could be used to trigger a Web response to be sent back to the requesting Internet user. The changes to the Call Center service logic turned out to be minor; it took as long to analyze the service and see how it could be arranged as a sequence of "form processing" transactions as it did to make the changes to the service logic.

In the Siemens Web Call Center, the "Internet Intelligent Peripheral" with which the service logic communicated was running as a separate program on the same node. Where more complex behavior is required of it (such as conversion of text to speech data and interface with the PSTN) then it would almost certainly be on a separate node. If data is transferred from the Internet in such a scheme, any intermediate gateway would be involved in relaying the data to this node.

6. Alternative Solutions

6.1 The AT&T System

AT&T developed a framework for controlling voice and voice-band data (e.g., fax) and for providing PINT services. Key to the framework is CallBroker, a logical entity that acts on behalf of a user to set up sessions and make requests for PSTN resources. The sessions typically include initiation of calls between two or more end points specified by the user. In addition to its interactions with the PSTN for call setup, the CallBroker is responsible for other functions, when necessary, such as authentication and usage recording.

This section briefly discusses the protocol at the two interfaces that need to be defined and the corresponding APIs to provide the above services. The two interfaces are (1) the one between the CallBroker (or Web Server) and the Service Control Function in the Service Node in the PSTN and (2) the one between the IP client and the CallBroker. The latter interface, in particular, will enable service providers to extend the architecture defined here to serve as a platform for other advanced/value-added services (to be identified later). In addition, the view taken here is that the IP client is more general, and implements a protocol for communication with the CallBroker that allows full two-way communications. For example, this is required for the cases where a called party hangs up and an indication may be necessary to be given to the IP Client about this status/progress. This is also necessary when conferencing to give an indication/status of various parties joining the call.

6.1.1 High Level Architecture

A high level architecture depicting various logical entities and the Interfaces among these logical Entities and the IP Client is shown in Figure 12.

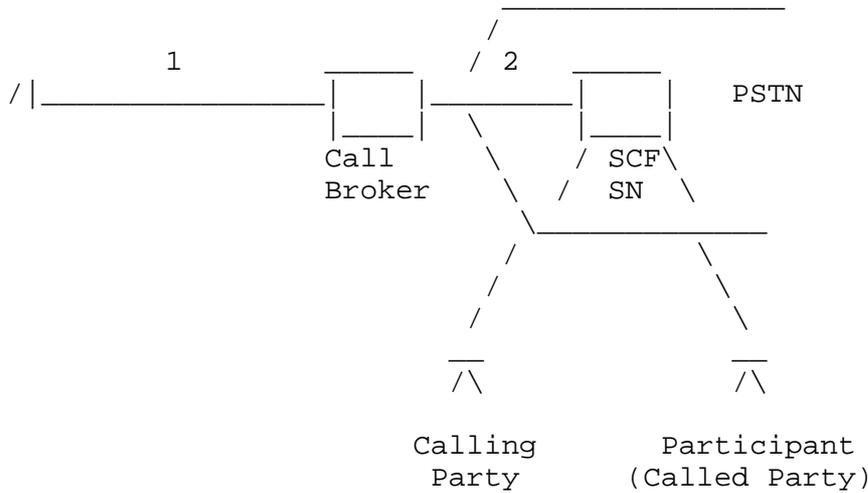


Figure 12: The CallBroker Architecture

The CallBroker, in addition to the initiation and control of calls on behalf of the user, performs additional functions. These functions include authenticating the IP Client, usage recording, and management of the session for the IP Client for the telephony call. The notion of the session requires that a client state machine be maintained in the CallBroker. This also helps in notifying the IP Client about the status/progress of the requests generated from the IP Client.

From the perspective of the IP Client, the logical entities needed for the above functions are within the CallBroker and are as shown in Figure 13 below. These correspond to the functions already discussed: Usage Recording Function, Session Management Function, Voice Bridge, and the Authentication Function. The fact that some of these functions may be physically separate from the CallBroker (such as the Voice Bridge being in the PSTN) is not inconsistent with the general view adopted here. Thus, the CallBroker Model mediates requests for network services and enables us to define various value added services in the future.

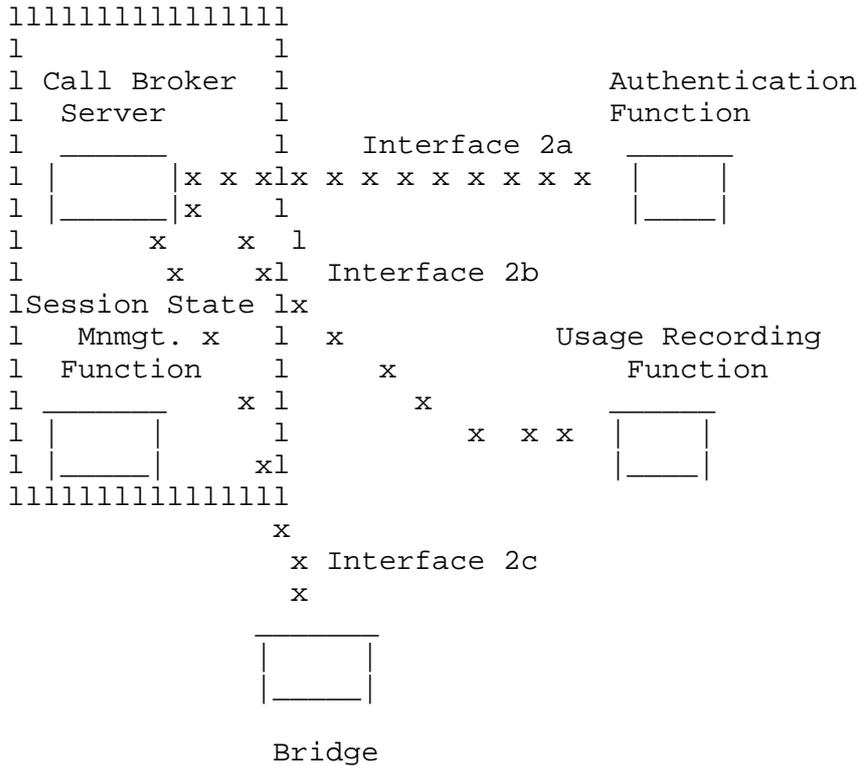


Figure 13: Functional Entities in the Call Broker

Various interfaces (i.e., 2a, 2b, 2c in Figure 13) between different functional entities in the CallBroker may also be standardized. The Session State Management Function may be physically realized as part of the CallBroker Server.

6.1.2 IP Client to CallBroker Interface

Communication on the IP Client to CallBroker Interface (Interface 1 in Figure 12) is a simple ASCII based protocol running directly on TCP. The messages on this interface are primarily requests from the client to the CallBroker, responses from the CallBroker to the IP client responding to the requests and unsolicited events from the CallBroker to the IP client. Since the communication is not strictly transaction oriented, traditional encapsulation protocols like HTTP cannot be used. There has been some ongoing work attempting to use multiple concurrent HTTP POST requests to support event delivery but, without too much difficulty, the ASCII protocol specified here can easily be mapped to the POST payload of the HTTP protocol.

6.1.3 Protocol

Basic Format

The basic format of the protocol is as follows:

```
[header]<<LF>
<<LF>
[body]<<LF>
<<LF>
<<LF>
```

The header and body of the protocol are separated by 2 line feed characters. The format of the header and the body is described below. Line feed characters in the header or body will be escaped using simple URL encoding.

Header

```
[session-id | 0]<<LF>
[message-id]<<LF>
[version-info]<<LF>
```

All CallBroker transactions are identified by sessions. A session does not necessarily correspond one-to-one to a TCP session. If the IP client is attempting to initiate a new session with the CallBroker the session-id field is populated with '0' to indicate session creation request. Every session request needs to be accompanied by sufficient information regarding authentication for the CallBroker to create the session.

Message-id represents the operation of the message.

Version-info contains optional version information of the protocol. This is to aid possible version mismatch detection and graceful error recovery.

Body

The body of the protocol messages consists of name value pairs. These name-value pairs are interpreted with reference to the message-id which signifies the operation to be performed by the CallBroker.

6.1.4 APIs Exposed to the IP Client

The APIs of the CallBroker exposed to the IP client are distinct and different from the APIs that the CallBroker uses from the different supporting subsystems including the authentication subsystem and the usage recording subsystem. The IP client APIs enable clients to effectively control voice conferencing.

6.1.5 Voice-Bridge Control API

The Voice Bridge Control API is used by CallBroker applications to access voice bridging functionality. The API distinguishes between sessions and calls. Calls represent actual voice calls placed from/to the voice bridge. These calls can be grouped together in sessions. All the calls that belong to a session are bridged. Calls have a significance outside the scope of sessions. Every call can be associated with multiple sessions with different weights at the same time. The advantage of this approach is the ability to support concepts like whispering in a conference call. Calls can also be dropped from a conference session and bridged together in a new session to give the notion of a sub-conference. These calls can later be re-added to the main conference session.

6.2 Simple Computer Telephony Protocol

6.2.1 Overview

The Simple Computer Telephony Protocol (SCTP) is a third party call control protocol and as such does not comply with the PINT charter. SCTP is described in this section to show how PINT services could be implemented using SCTP, and where SCTP fits into the PINT architecture.

In addition to third party call control, SCTP also provides subscriber (i.e., user) feature management (e.g., allows a user to set do not disturb, call forwarding parameters), and subscriber monitoring of terminal, line and address status. SCTP is strictly client/server-based. It has no provisions for peer to peer communications. SCTP runs as a TCP application protocol. It is ASCII-based and uses sockets. The SCTP Server is usually connected to a switch via a CTI (Computer-Telephony Integration) connection. Because of this, feature interactions are limited to those within the context of a single call, and not between PSTN services. The SCTP Server within a PINT Gateway could also be connected to an SN, or an SCP. See figures below. SCTP does NOT carry media.

6.2.2 How SCTP Fits in with the Reference PINT Services

SCTP Client as Part of a Web Server

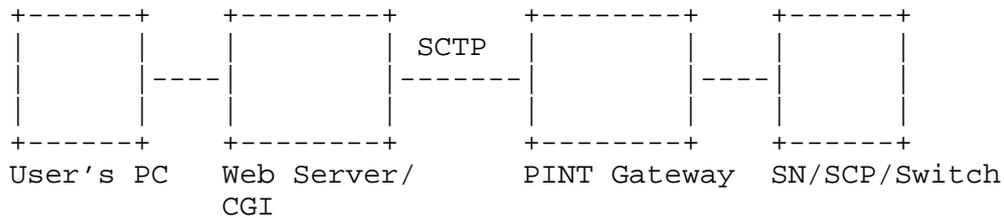


Figure 14: SCTP Client as Part of a Web Server

In this architecture, the SCTP Client is embedded in the Web Server. It is there for the specific purpose of initiating calls to the PSTN based on user requests. The SCTP Server is within the PINT Gateway. We go through the classic PINT examples:

Click-to-dial-back: The SCTP Client issues an SCTP MakeCall to the SCTP Server with the calling number supplied by Web page, and called number supplied by the user.

Click-to-fax-back: SCTP Client issues an SCTP MakeCall to the SCTP Server with called number set to user's fax machine, and calling number set to Web Server's fax machine, and treatment set to the URI for the file to be faxed. The SCTP Server takes the file and feeds it into the call just as a fax machine would.

Click-to-fax: SCTP Client issues an SCTP MakeCall with calling number set to user's fax machine, and called number set to Web Server's fax machine. How the file is supplied to the user's fax machine is outside the scope of SCTP.

Voice-access-to-content: SCTP Client issues an SCTP MakeCall with called number set to user's telephone number, and calling number set to Web Server and treatment set to a URI for the file of the particular Web page to be read to the called number. The SCTP Server takes care of the file to voice conversion and this is fed into the call as if it were voice.

In all of the above cases, the SCTP Client can generate a variety of different Web pages to send to the Web Server via CGI (Common Gateway Interface). The content of these pages is based on the call completion status of the CallMake SCTP action.

SCTP Client Running on the User's PC

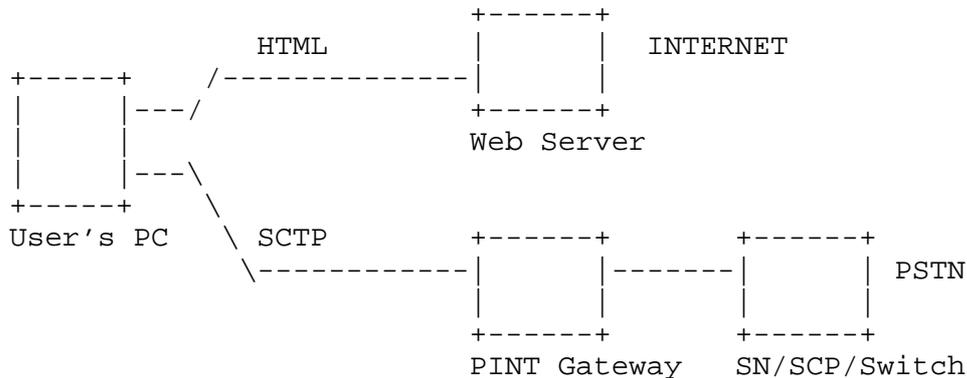


Figure 15: Sctp Client Running on the User's PC

In this architecture, the user has an Sctp Client co-located with it. If the user is using the telephone line for connection to a Web Server and there is an incoming call, then the Sctp Server in the PINT Gateway will post this event to the Sctp Client. A window will pop up on the user's screen with options available to the user for handling of the incoming call. The user can choose to take the call, send it to voice mail, or send it to another number.

For the Fax back service, for example, if the user had a separate fax machine from his or her PC, then the Sctp Server would tell the Sctp Client there is an incoming fax. The user would end or suspend his or her Internet connection, the fax would come in, and the user could then resume the Internet connection.

7. Session Initiation Protocol--An Emerging Standard

7.1 Overview

SIP, the Session Initiation Protocol, is a simple signaling protocol for Internet conferencing and telephony. It is currently under development within the IETF MMUSIC (Multiparty Multimedia Session Control) Working Group.

SIP provides the necessary mechanisms to support the following services:

- call forwarding, including the equivalent of 700-, 800- and 900-type calls;
- call-forwarding no answer;

- call-forwarding busy;
- call-forwarding unconditional;
- other address-translation services;
- callee and calling "numbers" delivery, where the numbers can be of any (preferably unique) naming scheme;
- personal mobility, i.e., the ability to reach a called party under a single, location-independent address, even when the user changes terminals;
- terminal-type negotiation and selection: a caller can be given a choice of how to reach a party, e.g., via Internet telephony, mobile, phone, and an answering service;
- caller and callee authentication;
- blind and supervised call transfer;
- user location; and
- invitation to multicast conferences.

Extensions of SIP to allow third-party signaling (e.g., for click-to-dial-back services, fully meshed conferences and connections to Multipoint Control Units (MCUs), as well as mixed modes and the transition between those) have been specified.

SIP addresses (URLs) can be embedded in Web pages. SIP is addressing-neutral, with addresses expressed as URLs of various types such as SIP, H.323 or telephone (E.164). A purely representational example of a SIP URL might be sip:+12125551212@foo.example.com, where foo.example.com is the host serving as a gateway into the PSTN.

SIP is independent of the packet layer and only requires an unreliable datagram service, as it provides its own reliability mechanism. While SIP typically is used over UDP or TCP, it could, without technical changes, be run over IPX, or carrier pigeons, ATM AAL5 or X.25, in rough order of desirability.

SIP can set up calls "out-of-band". For example, while the SIP protocol exchanges use IP, plus UDP or TCP, the actual data transport can take place via the PSTN. This feature makes it possible to use SIP to control a PBX or send requests to a Service Control Point. The PINT services make use of this flexibility.

7.2 SIP Protocol

SIP is a textual client-server protocol, similar in syntax to HTTP and RTSP. Requests consist of a method (INVITE, BYE, ACK, or REGISTER), a list of parameter-value pairs describing the request and an optional request body. Parameters include the origin and destination of the call and a unique call identifier. They may indicate the caller's organization as well as the call's subject and priority. The request body contains a description of the call to be

established or the conference to be joined. The description format is not prescribed by SIP; SDP is one possibility being standardized within the IETF. For the purposes of providing PINT services, an additional phone number address format is to be added to SDP.

Responses indicate whether a request is still being processed, was successful, can possibly be satisfied by another node or failed. When a call is redirected, the response indicates the name of the node to be tried. Unsuccessful calls may also return a better time to try again.

In a typical successful call, the caller sends an INVITE request to the callee. The callee accepts the call by returning a response code to the callee, which then confirms the receipt of that acceptance with an ACK request. Either side can terminate the call by sending a BYE request.

Requests can be authenticated using standard HTTP password and challenge-response mechanisms. Requests and responses may also be signed and encrypted.

7.3 SIP entities

SIP distinguishes three kinds of entities:

User agents receive and initiate calls and may forward the call.

A proxy server is an intermediary program that acts as both a server and a client for the purpose of making requests on behalf of other clients. Requests are serviced internally or by passing them on, possibly after translation, to other servers. A proxy must interpret, and, if necessary, rewrite a request message before forwarding it. A proxy server may, for example, locate a user and then attempt one or more possible network addresses.

Redirect server accepts a SIP request, maps the address into zero or more new addresses and returns these addresses to the client. Unlike a proxy server, it does not initiate its own SIP request. Unlike a user agent server, it does not accept calls.

Proxy and redirect servers may make use of location servers that determine the current likely location of the callee.

A PSTN gateway initiates phone calls between two parties. This may be a server that sends requests to an SCP in an IN environment or it may be a CTI-controlled PBX.

A SIP call may traverse one or more proxy servers.

The servers that control a PBX or an SCP act as user agents. A Web server may also act as a SIP user agent.

7.4 Providing Call Control Functionality

The SIP for PINT specification provides details on how to use SIP to initiate phone calls between two PSTN end points. (SIP can also initiate calls between Internet end points and between an Internet and PSTN end point, but this is beyond the scope of this document.)

It should be noted that the SIP client for initiating such phone calls can be either at the user's location (his/her workstation) or can be a Web server that calls up a SIP client via a CGI program. There is no difference in operation or functionality, except that the owner of the Web server may be legally responsible for the calls made.

A SIP client needs to convey two addresses to the PSTN gateway: the party making the call and the party to be called. (The party to be billed also needs to be identified; this can either be done by a SIP header or by having the server look up the appropriate party based on the two parties. This aspect is for further study.)

Described below are three ways these addresses can be conveyed in SIP. In the example, the address of party A is +1-212-555-1234 and that of party B is +1-415-555-1200. (The URL types in this and other examples are representational; they may but do not have to exist.)

(1) The two PSTN addresses are contained in the To header (and request-URI) and an Also header. For example:

```
INVITE sip:+1-212-555-1234@pbx.example.com SIP/2.0
To: phone:1-212-555-1234
From: sip:j.doe@example.com
Content-type: application/sdp
Call-ID: 19970721T135107.25.181@foo.bar.com
Also: phone:+1-415-555-1200
```

```
v=0
o=user1 53655765 2353687637 IN IP4 128.3.4.5
c=PSTN E.164 +1-415-555-1200
t=0 0
m=audio 0 RTP/AVP 0
```

In that case, the gateway first connects to party A and then party B, but without waiting for A to accept the call before calling B.

(2) Parties A and B are indicated by separate invitations. This allows the gateway to make sure that party A is indeed available before calling party B. After calling party A, the gateway could play an announcement indicating that the call is being connected using, for example, RTSP with appropriate Conference header indicating the call.

```
INVITE sip:+1-212-555-1234@pbx.example.com SIP/2.0
To: phone:1-212-555-1234
From: sip:j.doe@example.com
Content-type: application/sdp
Call-ID: 19970721T135107.25.181@foo.bar.com
...
INVITE sip:+1-415-555-1200@pbx.example.com SIP/2.0
To: phone:+1-415-555-1200
From: sip:j.doe@example.com
Content-type: application/sdp
Call-ID: 19970721T135107.25.181@foo.bar.com
...
```

(3) The two PSTN addresses are conveyed in the To header of the SIP request and the address in the SDP media description. Thus, a request may look as follows:

```
INVITE sip:+1-212-555-1234@pbx.example.com SIP/2.0
To: phone:1-212-555-1234
From: sip:j.doe@example.com
Content-type: application/sdp
Call-ID: 19970721T135107.25.181@foo.bar.com

v=0
o=user1 53655765 2353687637 IN IP4 128.3.4.5
c=PSTN E.164 +1-415-555-1200
t=0 0
m=audio 0 RTP/AVP 0
```

Here, pbx.example.com is the name of the PSTN gateway; the call will be established between 1-212-555-1234 and +1-415-555-1200.

Users can be added to an existing call by method (1) or (2).

8. Overall Security Considerations

Inter-networking of the Internet and PSTN necessitates the introduction of new interfaces (e.g., the A, B and E interfaces in Figure 6). To ensure that their use does not put the networks, in particular the PSTN, at additional security risk, these interfaces need to be designed with proper security considerations. Sections

5.1.5 and 5.2.2.7 describe how two of the pre-PINT implementations, the Lucent and Siemens systems, handle the security aspect, respectively.

Worth noting are the security requirements suggested by pre-PINT experiences. They are:

- +Peer entity authentication to allow a communicating entity to prove its identity to another in the network (e.g., the requesting IP-host to the PINT gateway, and the PINT gateway to the PSTN node providing the service control function).

- +Authorization and access control to verify if a network entity (e.g., the requesting IP-host) is allowed to use a network resource (e.g., requesting services from the PINT gateway).

- +Non-repudiation to account for all operations in case of doubt or dispute.

- +Confidentiality to avoid disclosure of information (e.g., the end user profile information and data) without the permission of its owner.

In the course of the PINT interface development, additional requirements are likely to arise. It is imperative that the resultant interfaces include specific means to meet all the security requirements.

9. Conclusion

This document has provided the information relevant to the development of inter-networking interfaces between the PSTN and Internet for supporting PINT services. Specifically, it addressed technologies, architectures, and several existing pre-PINT implementations of the arrangements through which Internet applications can request and enrich PSTN telecommunications services. One key observation is that the pre-PINT implementations, being developed independently, do not inter-operate. It is a task of the PINT Working Group to define the inter-networking interfaces that will support inter-operation of the future implementations of PINT services.

10. Acknowledgments

The authors would like to acknowledge Scott Bradner, Igor Faynberg, Dave Oran, Scott Petrack, Allyn Romanow for their insightful comments presented to the discussions in the PINT Working Group that lead to the creation of this document.

11. Appendix

11.1 PSTN/IN 101

11.1.1 Public Switched Telephone Network

What is normally considered as "the Telephone Network" consists of a set of interconnected networks. Potentially, each of these networks could be owned by a different Network Operator. The official name for such a network is Public Switched Telecommunications Network (PSTN). A simple PSTN consists of a set of Switches (called Central Offices or Telephone Exchanges) with links interconnecting them to make up the network, along with a set of access connections by which terminals are attached. The PSTN is used to deliver calls between terminals connected to itself or to other PSTNs with which it is interconnected. Calls on the PSTN are circuit switched; that is, a bi-directional connection is made between the calling and called terminals for the duration of the call. In PSTNs the connection is usually carried through the network in digital format occupying a fixed bandwidth; this is usually 56 or 64 Kbps. The overall configuration of the PSTN is shown in Figure 16.

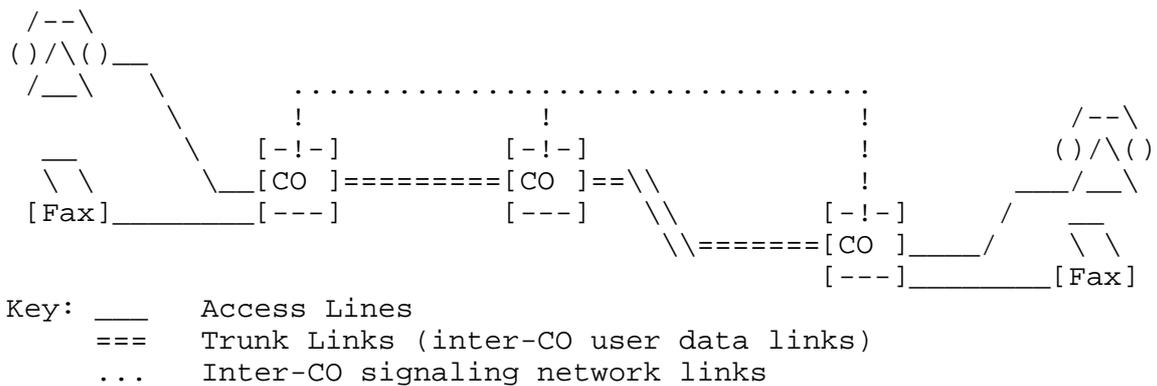


Figure 16

Messages are sent between the Switches to make and dissolve connections through the network on demand and to indicate the status of terminals involved in a call; these "signaling" messages are carried over a separate (resilient) data network dedicated to this purpose. This signaling network is also known as the Common Channel Signaling (CCS) or Signaling System Number 7 (or SS7) network after the names of the signaling protocol suite used.

As yet, the majority of access connections to a PSTN carry analogue signals, with simple (analogue) telephones or Facsimile machines as terminals. Call requests are indicated to the Central Office to which

a telephone is connected either by a sequence of pulses or tone pairs being sent. Notifications on the status of the request are sent back to the telephone in the form of tones. Indication from a Central Office that a call is being offered to a telephone is arranged by sending an alternating voltage down the access connection which in turn causes the ringer in the telephone to sound. These access lines have a unique address associated with them and can support a single call.

However, with analogue or digital multi-line connections, or Integrated Service Digital Network (ISDN) Basic or Primary Rate Interfaces (BRI or PRI), several concurrent calls are possible and a set of addresses are associated with them. The new ISDN access connections are designed so that data exchanged with the network is in multiplexed digital form, and there is an individual channel for each of the potential connections, together with a separate channel dedicated to sending and receiving call request and call alert data as well as carrying packet switched user data. These call request and call alert messages act as the equivalent of the pulses or tones that are sent when dialing, and the ringing signal that is sent to a telephone when a call is being made to it.

The operation of the call request is fairly simple in most cases and is shown in Figure 17.

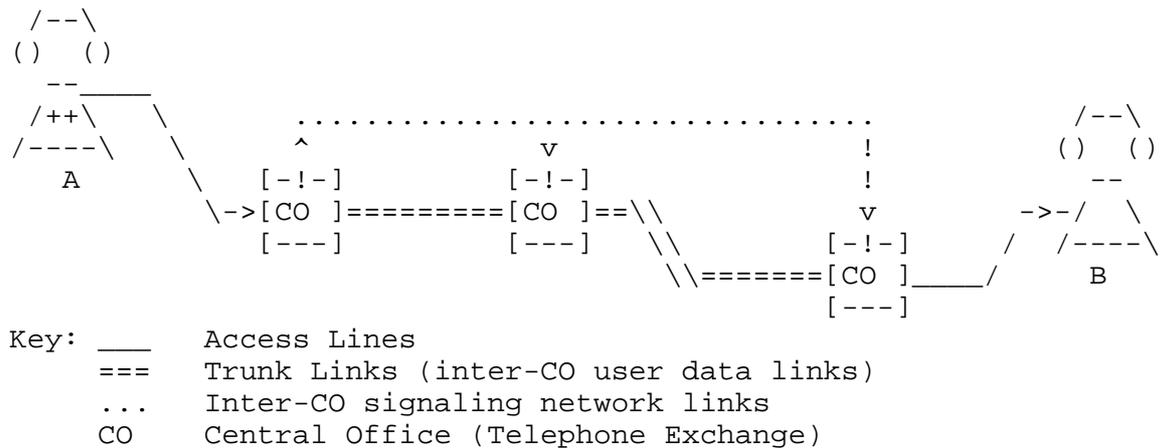


Figure 17

The user presses a sequence of numbers on a telephone handset (labeled A), and the telephone passes a sequence of digits (either as pulses or tone pairs) to the Central Office via the access line. The Central Office contains a processor that will be notified that the user has made a request and the digit string that is the sole parameter of the request. This digit string is taken to be the unique

address of an access line connected either to itself or to another Central Office. There is a hierarchical addressing scheme, so that the digit string can be parsed easily. A call request to a terminal (labeled B) connected to a remote Central Office can be routed by examining the digit string passed; the Central Office will extract the part of the passed address that corresponds to the remote Central Office in question, and can route the request onward, forming an inter-Switch call request and passing it via the signaling network. At the same time it will allocate one of its available transmission channels towards the remote Central Office.

11.1.2 Intelligent Network

This scheme has been used since the 1950s, and suffices for the majority of calls. However, there are a range of other services that can be (and have been) provided, enhancing this basic call processing. Freephone or Premium Rate services (1-800 or 1-900 services) are good examples of the supplementary services that have been introduced. Apart from the important feature that the cost of these calls is varied so that the caller does not pay for a free-phone call, or pays an extra charge for a premium rate call, they have the similarity that the number dialed must be translated to arrive at the "real" address of the destination terminal. They are known as number translation services, and make up the bulk of all supplementary services delivered today.

These were originally programmed into each Central Office, but the complexity of maintaining the data tables on each processor grew cumbersome, so a more general solution was sought. After a considerable gestation period, the eventual solution was the Intelligent Network. This takes the separation of Central Offices and the network links interconnecting them a stage further.

The Central Offices are considered to provide the Call Control Function (CCF). In addition, the Service Switching Function (SSF) is provided to "enhance" the operation of these Switches by detecting when a particular request has been made (such as by dialing 1-800). If this pattern is detected, the equipment implementing the SSF will send a specialized request message over the signaling network to a separate computer that implements the Service Control Function (SCF). This entity is responsible for querying service specific data (held in a unit providing the Service Data Function, or SDF), performing any digit translations necessary, and sending the details of how to proceed back to the SSF, where they are obeyed and the call is put through to the "real" destination. In many implementations, the SDF is closely coupled to the SCF. This configuration is shown in Figure 18.

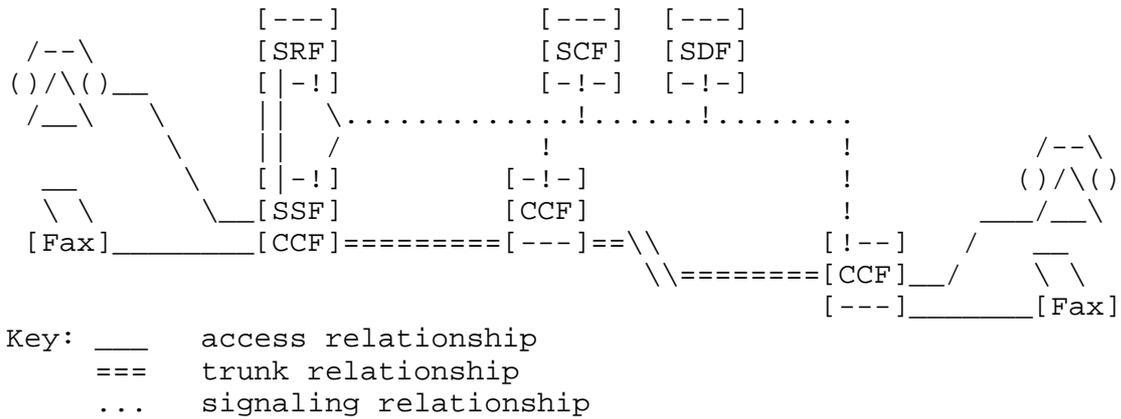


Figure 18

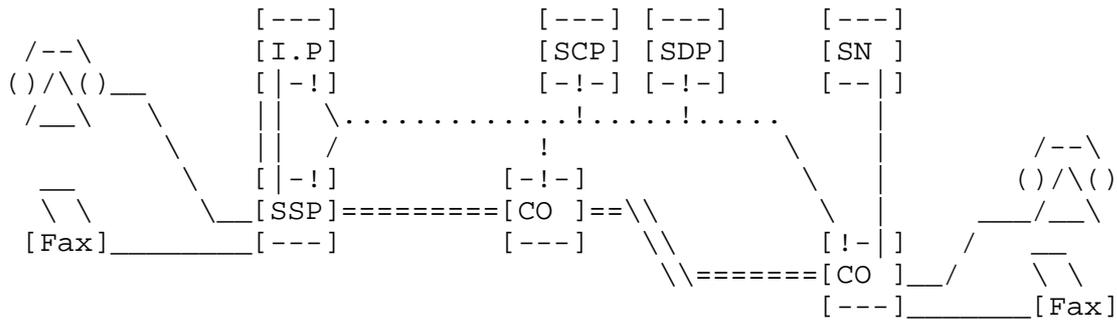
The advantage is that there can be a much smaller number of physical units dedicated to the SCF, and as they are connected to the signaling network they can be contacted by, and can send instructions back to, all of the units providing the SSF and thus the CCF.

In another enhancement, a separate entity called the Special Resource Function (SRF) was defined. Equipment implementing this function includes announcement units to play recorded messages (for example, prompts to enter digits) to callers. It will also include the tone decoders needed to capture any digits pressed by the caller in response to the prompts. It is connected to the rest of the PSTN usually via trunk data links. It will also include a signaling connection (directly or indirectly) back to the SCF, via the PSTN's core signaling network.

As an example of the way that these different functional entities interact, the SCF can ask an SSF handling a call to route the caller temporarily through to an SRF. In response to instructions sent to it from the SCF over the signaling network, the SRF can play announcements and can collect digits that the user presses on their terminal in response to prompts they are played. Once these digits have been collected they can be passed on to the SCF via a signaling message for further processing. In normal operation, the SCF would then ask the SSF to dissolve the temporary connection between the user's terminal and the SRF. This allows the collection of account numbers or passwords (or PINs) and forms the heart of many "Calling Card" services.

This pattern of user interaction is also used in a wide variety of other services where extra account information and PINs are needed. They are collected as just described and can be checked against the correct values stored in the service database prior to allowing the call to proceed.

The Intelligent Network functional entities can be realized as physical units in a number of different combinations. A common configuration is shown in Figure 19.



- Key:
- ___ Access Lines
 - === Trunk Links (inter-CO user data links)
 - ... Inter-CO signaling network links
 - SSP Service Switching Point - a unit that implements the Service Switching Function
 - CCP Call Control Point - a unit that performs call control functions.
This is normally a kind of Central Office (shown as CO above)
 - SCP Service Control Point - a unit implementing the Service Control Function. NOTE that this is connected to the SS7 Network and uses this connection for all of its communications.
 - I.P Intelligent Peripheral - a unit that contains specialized resources (like announcement units, tone decoders).
In effect, it implements Special Resource Functions.
 - SN Service Node

Figure 19

This diagram also shows a unit called a Service Node, or SN. This contains components that realize all of the operational Intelligent Network functions (SSF, SCF, SDF, and SRF). It is sometimes more convenient to have all of these elements in one node (for example, for operations and maintenance reasons), particularly within smaller PSTNs or where there is a relatively low level of requests for

particular services. Another difference is that, as they are all co-located, proprietary protocols can be used for internal communication, rather than the full Intelligent Network Application Part (INAP) protocol used over the core signaling network between discrete units. It also differs from the "unbundled" approach in that it is connected to the COs within a PSTN as a peripheral, having only an access connection to a Central Office; there is no connection to the core signaling network. Other than this, it operates in a similar way, and can provide the same kinds of services. Information on the specification of the Intelligent Network can be found in the ITU recommendations [1], while two books ([2] and [3]) describe the system, its history, operation, and the philosophy behind it.

11.2 Call Center Features

A Call Center is a system that allows a company to be organized with a group of similar individuals (agents), all of whom can either make calls to, or take calls from, customers. The system distributes incoming calls to the agents based on their availability and automates the placement of outgoing calls, selecting an agent to handle the call and routing the call to them only once the call request has been made of the PSTN.

The incoming call distribution feature ("automatic call distribution", or ACD) is usually coupled with a call queuing scheme. In this scheme, the callers are connected temporarily with an announcement unit that normally plays music. The calls are treated in sequence so that (once the caller is at the front of the queue) the ACD system selects the next available agent and routes the call through to them.

Another feature connects a customer making an incoming call to a unit that asks them for some information on the purpose of their call, selecting the agent to handle the call based on the particular area of expertise needed; to do this, the agents are further categorized by their knowledge (or "Skill Set"). If this skill set categorization is used then by implication there will be separate queues for each of the skill sets. This user selection scheme can be used independently of the others. For example these so-called "voice navigation systems" can be used to select a particular department extension number, based on the function required by the customer; as such, they can automate the job of company telephone receptionist in routing incoming calls.

Where possible, the information gleaned from the customer can be provided to the selected agent, usually via a separate networked computer connection. Similarly, if an outgoing call is being made to one of a list of customers, information on the customer and the purpose of the call can be provided to the agent selected to handle

the call. Such configurations are generally called "Computer Telephony Integration" or CTI systems. Strictly, a CTI system can be arranged to handle routing of incoming calls and automation of outgoing calls only (also known as computer integrated telephony features), without the agents having access to a network of computers. However, the business case for combining the telephony functions of the call center with provision to the agents of computers with customer information can be compelling.

This is often further combined with a company's order and service processing computer system. In this case, a call is treated as part of a business transaction, with the information to be exchanged captured as fields of a computer form. While such a computer system is not, strictly, part of a call center, integrating the company computer system with the call center is very common. This allows the details of the call to be stored on a centralized database, allowing further automated order processing, for example. It also allows the call to be transferred from one agent to another where needed, ensuring that the new agent has the information already captured. This might be useful if someone with a different area of expertise were to be needed to handle the customer's requirements.

Traditionally, Call Centers have been used to support teams of agents working at a single site (or a small number of sites, with private telephony trunks interconnecting them). The site Private Automatic Branch eXchange (PABX) was integrated with a computer system to provide these features to people at that site. There can be a business case for provision of such features to distributed teams of workers as well. In particular, the possibility of providing support for people working from home has been seen as important. Some of the Call Center features have been incorporated into public telephone exchanges or Central Offices (COs) from many manufacturers as part of their "Centrex" service offerings.

There are practical limitations in providing such features on COs. Apart from the procedures needed to configure these features for any telephone line that is to use them, the basic requirement that every agent must have a connection to the supporting CO can limit its usefulness. Another approach is to provide Call Center features via the Intelligent Network. The features might thus be provided over a Telephone Operator's entire network, and would mean that the Call Center could be configured centrally while still allowing agents to be located anywhere within the telephone network. It also means that the supported company can pay for the Call Center features "as they go" rather than having a high "up front" cost.

12. References

- [1] ITU-T Q.12xx Recommendation Series, Geneva, 1995.
- [2] I. Faynberg, L. R. Gabuzda, M. P. Kaplan, and N. J. Shah, "The Intelligent Network Standards, their Application to Services", McGraw-Hill, 1996.
- [3] T. Magedanz and R. Popesku-Zeletin, "Intelligent Networks: Basic Technology, Standards and Evolution", Intl. Thomson Computer Press, 1996.
- [4] Information processing systems - Open Systems Interconnection - Specification of Abstract Syntax Notation One (ASN.1), International Organization for Standardization, International Standard 8824, December, 1987.
- [5] McCloghrie, K., Editor, "Structure of Management Information for Version 2 of the Simple Network Management Protocol (SNMPv2)", RFC 1902, January 1996.
- [6] Kristol, D. and L. Montulli, "HTTP State Management Mechanism", RFC 2109, February 1997.
- [7] Zimmerman, D., "The Finger User Information Protocol", RFC 1288 December 1991.

Authors' Addresses

Steve Bellovin
AT&T Labs
Room E-215
180 Park Ave. Bldg. 103
Florham Park, NJ 07932-0000
USA

Phone: +1 973 360 8656
Fax: +1 973 360 8077
EMail: smb@research.att.com

Fred M. Burg
AT&T Labs
Room 1N-117
307 Middletown Lincroft Road
Lincroft, NJ 07738
USA

Phone: +1 732 576 4322
Fax: +1 732 576 4317
EMail: fburg@hogpb.att.com

Lawrence Conroy
Roke Manor Research Limited
IT&N-INIA Group
Roke Manor, Old Salisbury Lane,
Romsey, Hampshire SO51 0ZN
U.K.

Phone: +44 1794 833666
Fax: +44 1794 833434
EMail: lwc@roke.co.uk

Paul Davidson
Nortel
P.O.Box 3511 Station "C"
Mail Stop 242
Ottawa, Ontario, Canada K1Y 4H7

Phone: +1 613 763 4234
EMail: pauldav@nortel.ca

A. DeSimone
Lucent Technologies
Room 6H510
600-700 Mountain Avenue
Murray Hill, NJ 07974-0636
USA

Phone: +1 908 582 2382
Fax: +1 908 582 1086
E-Mail: tds@lucent.com

Murali Krishnaswamy
Bell Laboratories
Lucent Technologies
Room 2G-527a
101 Crawfords Corner Road
Holmdel, NJ 07733-3030
USA

Phone: +1 732 949 3611
Fax: +1 732 949 3210
EMail: murali@bell-labs.com

Hui-Lan Lu
Bell Laboratories
Lucent Technologies
Room 4K-309
101 Crawfords Corner Road
Holmdel, NJ 07733-3030
USA

Phone: +1 732 949 0321
Fax: +1 732 949 1196
EMail: hui-lan.lu@bell-labs.com

Henning Schulzrinne
Dept. of Computer Science
Columbia University
New York, NY 10027
USA

Phone: +1 212 939 7042 (@Bell Labs: 732 949 8344)
Fax: +1 212 666 0140
EMail: schulzrinne@cs.columbia.edu

Kamlesh T. Tewani
AT&T Labs
Room 1K-334
101, Crawfords Corner Rd.
Holmdel, NJ 07733
USA

Phone: +1 732 949 5369
Fax: +1 732 949 8569
EMail: tewani@att.com

Kumar Vishwanathan
Isochrone
EMail: kumar@isochrone.com

Full Copyright Statement

Copyright (C) The Internet Society (1998). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

