P2P VoIP for Today's Premium Voice Service[1]

Ayaskant Rath†, Stevan Leiden*, Yong Liu‡, Shivendra S. Panwar‡, Keith W. Ross‡

†ARath01@students.poly.edu, ‡{YongLiu, Panwar, Ross}@poly.edu, *Steve.Leiden@verizon.com

†‡Polytechnic Institute of New York University, Brooklyn, NY, 11201, USA

*Verizon Network and Technology, Waltham, MA, 02451, USA

Abstract

With the integration of voice and data networks, service providers are gradually switching to Voice-over-IP (VoIP) for their voice services. Typically, today's centralized VoIP networks pay a high price for scalability, particularly, when they provide a modern premium voice service. However, as the Internet Service Providers are offering their customers high bandwidth access, scalability issues with VoIP systems can be addressed by using Peer-to-Peer (P2P) schemes for VoIP services. In this article, we propose a scalable and robust P2P VoIP service system that can provide all the features and services a premium voice service provides today. We discuss the network architecture on which the system runs, as well as the implementation of various features.

I. Introduction

A telephone service that includes just a voice service has now become a service of the past. Today's premium telephone services provide much more than just voice. Features such as voice mail, address books, contact specific ringtones and customized diversion of incoming calls are common features of a telephone service today. With the advent of Voice-over-IP (VoIP) the voice network is getting more and more integrated with the data network. Although VoIP started off as a by-product of the cheap and high trunk bandwidth provided by Internet Service Providers (ISPs), it has now become a standard application that service providers offer. However, most of the VoIP usage by telephony service providers is done by using centralized servers that provide the voice services and host subscriber features. The problem with such a system is that it scales with increasing number of subscribers and/or service requests only at great cost. There is an upper limit to the number of subscribers or service requests a centralized system can handle. As this maximum number is approached, particularly over 80% utilization, performance generally

[1]The system described herein is for discussion only and does not represent a system Verizon has fielded or planned

degrades in a non-linear way. Therefore, a second similar service host must be added to the system, making it an expensive scaling process. Apart from the scalability servers, the core network infrastructure must also be upgraded to handle the excessive traffic due to the increase in service requests. Often there are delay and bandwidth constraints involved in meeting the service requirements.

Today, ISPs are deploying Fiber-to-the-Premises (FTTP) networks [1] that will support voice, data and video over a single fiber to the home. The Optical Network in such systems terminates right at the subscriber premise where a device called the Optical Network Termination (ONT) is installed. The ONT takes care of demultiplexing data to various devices such as the Telephone, the Set Top Box (STB), an in-home router, etc. Thus, the ONT, located at the subscriber premise, potentially has substantial computational power.

With such systems fast evolving, it seems beneficial to address the scaling issue of a large scale VoIP Telephony service by utilizing Peer-to-Peer (P2P) technology. If the services provided by the centralized servers are pushed to the peers, making them work in a distributive manner, the system will not only scale well, but also will be far more resilient to failures. In this article we discuss how a P2P system can be used to replace the existing centralized system in order to cut costs, attain scalability, and enhance resilience while not compromising on the services provided to subscribers.

We describe below the most common features offered by today's voice services. We then discuss the network architecture on which the proposed system works, followed by the details of the system. At the end, implementation of each of the important features of the service is discussed.

II. Features in Today's Telephone Service

A subscriber being able to make a telephone call to another is merely one of the services provided in a Premium Voice Service today. Listed below are the most important features normally included in a Premium Voice Service bundle.

1) *Service Account* The subscriber is provided with an online Service Account, which they can manage from any location with an Internet connection. The account includes the subscribed and customized service features that can be modified by the user on the go.
2) *Voice Mail* A caller can leave a voice message when the subscriber does not respond to a call. This voice message should be accessible from the Telephone device, the Television or even from a computer connected to the Internet in or outside the user premise (using the service account).

3) *Address Book* The subscriber is provided with an address book, which stores the information about other subscribers including their telephone numbers, pictures and addresses for example. The address book can also be edited by the subscriber using their service account.
4) *Caller ID* When the subscriber receives a call, the caller telephone number is displayed on the telephone device. Also, depending on the subscriber preferences, the caller's picture (and possibly other information) is also displayed if the device makes this possible.
5) *Ringtones* The subscriber can set specific ringtones for specific entries in their address book. Ringtone preferences can also be edited using the service account.
6) *Ringback Tones* Ringback tone is what the caller hears when the called phone rings. The subscriber can choose specific ringback tones for specific callers calling them. The ringback tones can be one of several provided by the service provider, or even a customized tone made available by the subscriber.
7) *Click-to-Dial* The subscriber can log into their service account, open their address book and click on a contact to call them. When clicked, the subscriber's telephone device rings, indicating that the call is ready to be set up, and when the subscriber picks the phone, the call is set up.
8) *Three-Way Call* The subscriber can call two distinct parties at the same time and join the two calls to form a three-way call where all three parties can hear each other.
9) *Call Forwarding and Scheduling* The subscriber can choose when and what calls they want to receive on their telephone device. They can choose calls to ignore or to forward to other telephone numbers. The subscriber can also specify which calls are directed to which handset(s) in the premise. These preferences are set through the service account.
10) *Emergency Calls* The subscriber can make emergency calls from their telephone device. When the subscriber makes an emergency call, the service provider provides subscriber's location information to the emergency service.
11) *CALEA* The service provider has built-in surveillance capabilities, allowing federal agencies to monitor all voice traffic in real-time to enhance the ability of law enforcement and intelligence agencies to conduct electronic surveillance.
12) *Billing* The subscriber can get real time account billing information.
13) *Softphone* Softphone is Skype [2] like software that can be used by the subscriber as a telephone device. The subscriber can install the software on a computer at any location connected to the internet

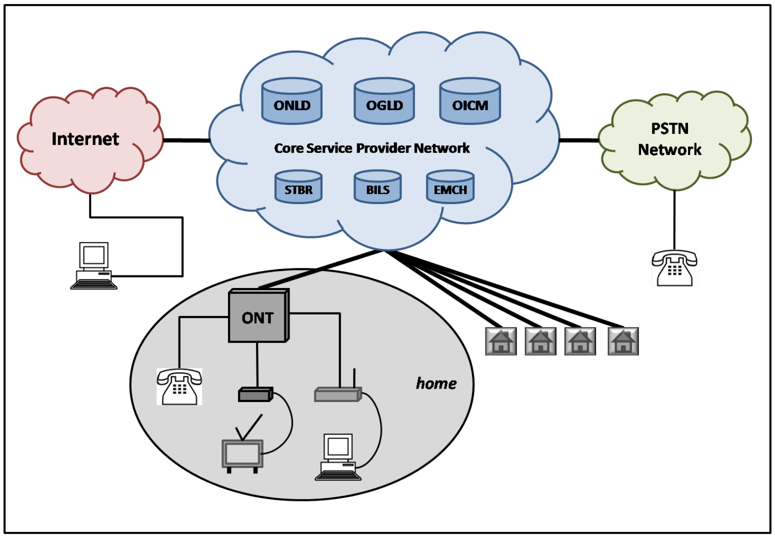

Fig. 1. Simplified schema of the Network Architecture

and then log in to make calls, just as if making calls from their telephone device on premise. (Such features, usually, do not provide emergency call services)

III. Network Architecture

Our proposed P2P System for VoIP runs on a Fiber Optic to Home network, where the fiber optic line terminates at the user premise. The network is assumed to be used to provide other services such as TV, Video on Demand, Internet, etc. apart from the voice (telephone) service discussed here. The ONT located at the user premise demultiplexes data for various services and is assumed to have sufficient computational power for our purposes, but no permanent storage. It also has battery backup to sustain power failures. Though located at the user premise, it is part of the service provider's network infrastructure and is fully controlled by the service provider. The ONT is assumed to be secure enough so that a subscriber or any other entity cannot manipulate it. Every ONT in the network is associated with a *Network Location* and a *Geographic Location*. The network location of an ONT is its identity in the network, for example, its IP Address. The Geographic Location of the ONT is its set of geographic coordinates. The network also has a set of servers, each serving a specific purpose. Figure 1 shows the network architecure described.

The ONT Network Location Database (ONLD) is a set of servers storing the network location information of the ONTs. The ONLD is implemented using a Distributed Hash Table (DHT) with the telephone number as the key and network location of the ONT as the value. Each server in the ONLD stores information about a set of ONTs (possibly covering a particular area). These ONTs are said to be associated with this particular ONLD server. A similar set of servers called the ONT Geographic Location

Database (OGLD) is used to store the Geographic Locations of the ONTs, also in a DHT. Use of a set of servers instead of a single one reduces individual load and adds a layer of redundancy. Any of the open source implementations available for DHTs [3] can be used to build the servers.

The ONT Initiation and Configuration Manager (OICM) is responsible for the configuration of the ONTs and the input of data into the ONLD and OGLD. OICM is the only entity that has write permission on the DHTs. The nodes in the ONLD and OGLD can only read (or fetch) data from the DHTs they store. For subscribers who subscribe to TV services, their STB is responsible for permanent data storage. For subscribers who don't subscribe to TV services and hence do not have a STB on their premise, and for subscribers who subscribe to TV services, but do not have a STB, a STB Replacer (STBR) is responsible for permanent data storage.

The Billing Server (BILS), maintains the billing information for all users in the network. The ONTs are responsible for updating and sending information to this server. The Emergency Call Handler (EMCH) manages emergency calls. It can retrieve geographic location from the OGLD for a given ONT.

IV. Installation of an ONT

Whenever a new ONT subscribing to the voice service is installed, the OICM takes care of adding the information about the new ONT into the ONLD, the OGLD and the BILS. The OICM also configures the new ONT. Once this is done, the new ONT becomes part of the network and calls can be made from and to it. Configuration of an ONT includes loading and running the application software, loading connectivity information about the associated ONLD, EMCH, BILS, STBR (when there is no STB) and the OICM itself, loading the default user preferences and settings, the Emergency Call Filters and the Ringtones and Universal Ringback Tones. Connectivity information loaded is used by the ONT to connect to the servers when necessary. Emergency Call Filters are used by the ONT to distinguish emergency calls from regular ones. Universal Ringback Tones are the standard set of ringback tones made available to the subscriber, by the service provider. The ONT forwards the default user preferences and settings, the Ringtones and the Universal Ringback Tones downloaded as part of initiation process to the STB or the associated STBR.

The above installation process also takes place in the rare occasion of an ONT getting restarted. In the event of an associated server going down, the ONT connects to the OICM to load connectivity information of its new associated server.

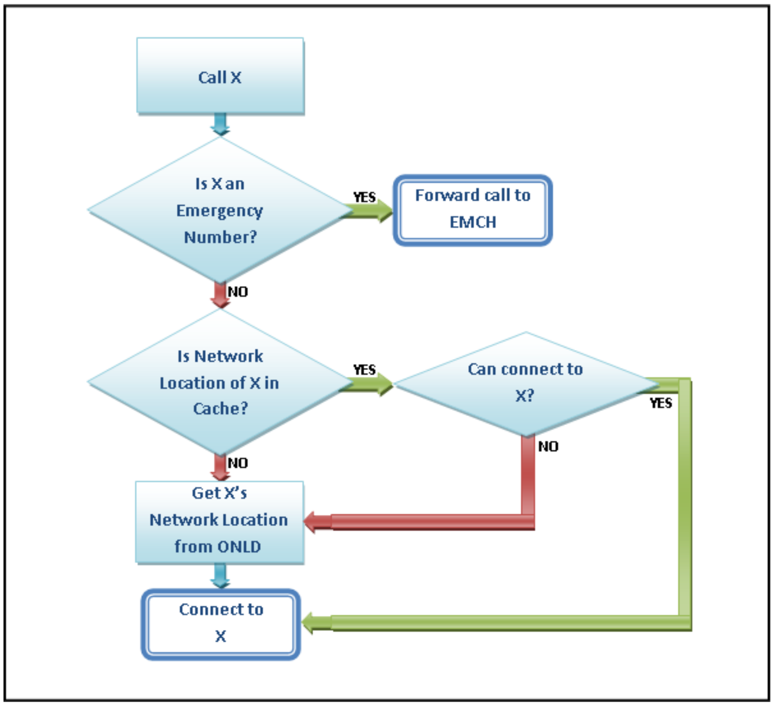

Fig. 2. Normal Call Set Up

V. Normal Call Set Up

When one subscriber makes a call to another subscriber, the caller ONT first checks if the number being called is an emergency number using its Emergency Call Filters. If the number is not an Emergency Number, the ONT then checks its cache for the network location of the called ONT. If found, the caller ONT connects to the called ONT. If connection fails because of stale cache, or if the called ONT network location is not found in the caller ONT cache, the caller ONT connects to the ONLD to retrieve network location information about the called ONT. Note that in a centralized server system, the network location of the called ONT is never revealed to the caller ONT as all communication happens through the centralized server. Our model satisfies this requirement. The Network Location of the called ONT is invisible to the user. Thus an ONT being given access to the network location information of another ONT does not raise any security issue. The normal call set up process described here is shown in Figure 2

Once the caller ONT gets the Network Location of the called ONT, it makes a connection to the called ONT and proceeds with the necessary exchange of control signals thus initiating the call. Then the voice data between the two ONTs is exchanged through the direct connection between them. During the call, both ONTs record necessary information about the call, for example, duration of the call, for accounting purposes including billing.

When an ONT receives a call, it first checks if it is really the right receiver of the call (the called number is actually its number). If not, it notifies the caller ONT resulting in a connection failure. If the number is correct, it checks for an entry in its user's address book and also checks its user's preferences. If the preferences allow for the call to be accepted, the ONT accepts the call, sends the right ringback tone and forwards the necessary ID and Picture information to the Phone/TV as per user preference. The call receiving options are discussed in detail later.

At the end of the call, both ONTs add the network location of each other to their cache and report the call duration and other necessary information about the call to BILS. The caching feature avoids connection to ONLD and DHT lookup for every call.

Compared to the centralized VoIP service architecture, this scheme reduces the entire DNS load on the servers by using DHTs. The scheme also reduces the load on the bottleneck servers which are otherwise responsible to forward all voice traffic by making ONTs directly connect to each other. This also helps reduce traffic on the backbone network to some extent. The central hardware infrastructure is also reduced as a result.

VI. Implementation of other Features

All the other features commonly used in a premium voice service can be implemented within the proposed architecture. The details of the implementations are given below.

1) *Service Account* The Service Account manager, which normally sits on a centralized server, instead runs on the ONT. The data it uses, such as the user preferences are stored in the STB and retrieved by the ONT when necessary. For subscribers who do not have a STB, the data is stored in the STBR. Initial Service Account parameters and default user preferences are downloaded as part of the ONT initialization process. Users can access all the information from their PC by logging into their service account. When the user logs in using a PC outside home, they connect to their ONT and work just as if working from home.

   Since all user information, address book and preferences for example, are stored in the STB or STBR, other features, like call notification via e-Mail or sending SMS to a mobile device, can be handled by the ONT very easily.
2) *Voice Mail* If the called user is busy or if the called user's preferences require it, the incoming call is forwarded to the called user's Voice-Mail. A greeting is passed on to the caller ONT. The voice message of the caller is then recorded by the called ONT. The called ONT then transfers the recorded voice mail to the STB instead of a central Voice Mail server. This is unlike recording

voice messages in a wireless handset because the ONT is not under the user's control and the data cannot be lost. Also, voice mail can now be accessed by the user when logged into their service account from any computer. In order to make the system more robust, there may be an option for backing up Voice Mail data onto other ONT/STBs.

Once a Voice Mail is received, depending on user preference, the ONT sends out notifications to the user's handset(s), TV (through STB) and/or their online account portal.

3) *Address Book* The address book of the user is cached in the ONT and stored in the STB/STBR. The ONT refers to the address book every time there is an incoming call for caller ID and picture settings. Also, the user can access the address book from the handset which is directly connected to the ONT in order to make a call. The user can also access and edit their address book from a PC by logging into their service account.
4) *Caller ID* When an ONT receives a call, it checks for the caller's entry in the user's Address Book. If found, depending on the preferences, the ONT transfers the name and picture of the caller to the user's handset(s) and/or TV (through STB).
5) *Ringtones* When an ONT receives a call, it chooses the ringtone associated with the caller from the address book or user preferences, retrieves the tone from the STB or STBR and transmits it to the handset(s). Alternatively, a set of standard ringtones can be added to the handsets, in which case the ONT simply needs to signal the handset which Ringtone it needs to play when there is an incoming call.
6) *Ringback Tones* As described before, Ringback tones can be chosen from a specific set of Universal Tones or can be custom made. Every time a called ONT accepts a call, it checks for a Ringback Tone entry in the Address Book. If the Ringback tone corresponding to the caller is one of the universal tones, the called ONT sends a signal to the caller ONT for the particular tone. The caller ONT then retrieves the corresponding tone from its own STB and plays it to the caller. If, on the other hand, the tone is a custom tone, the called ONT retrieves it from its STB and streams it to the caller ONT which then plays it.
7) *Click-to-Dial* When the user logs into their account using a PC and clicks on a number to make a call, the ONT rings the user's handset. When the user picks it up, the call set up process is automatically initiated by the ONT like a normal call.
8) *Three-Way Call* When an ONT needs to make a Three-Way call, it needs to simultaneously connect

to two ONTs. The caller ONT thus makes two connections by just doing everything it does for a normal call, twice. Once the two connections are established, the caller ONT signals the two called ONTs informing them about the Three-Way call and forwarding each of them information about the other's identity. The caller ONT is then responsible to forward voice traffic from each of the two parties to the other ONTs. Thus, if not for the signaling, the called ONTs cannot distinguish a Three-Way call from a normal call. The received voice data appears just like a normal call to the called ONTs.

Note that the above scheme may not scale to accommodate conference calls with a higher number of members as the ONT may not have enough resources to handle many simultaneous calls. In order to handle conference calls with more members, it is best to use a separate media server infrastructure, that handles these features including multimedia.

9) *Call Forwarding and Scheduling* The Call Forwarding and Scheduling preferences of the user are part of their profile data and would be cached in the ONT and stored in the STB (or the STBR) by the ONT. Every time an ONT receives a call, it checks for this information. If the call is supposed to be forwarded, then the ONT sets up another connection to the right destination.
10) *Emergency Calls* When a user makes an emergency call, the number is identified by the ONT using its Emergency Call Filters. Once identified as an Emergency Call, the ONT doesn't go through the normal call set up process. Instead, it connects to its associated EMCH. The EMCH then connects to the OGLD to retrieve necessary information about the caller's geographic location from the DHT. Finally, the EMCH forwards the call to the emergency service with the necessary emergency information attached.
11) *CALEA* The OICM has access to all the ONTs in the network. Whenever calls from and/or to an ONT need to be intercepted, the OICM signals the ONT informing it of the CALEA request. When the ONT receives this request, it forwards all necessary information to the OICM for every call involved, similar to a Three-Way call. The information is then packaged and forwarded to the appropriate body. The interception remains transparent to the user.
12) *Billing* The duration of every call is recorded at the caller's ONT as well as at the called ONT. As soon as a call is terminated the ONTs report the necessary information about the call to the BILS. The BILS then sends back billing information to the ONTs and the ONTs transfer this data to the STBs. The user can see this information via their Service Account. Whenever the user logs into

their service account to check the current billing status, for example, minutes used, free minutes left, the information is provided by the ONT (after being retrieved from the STB). However the BILS always maintains all billing information and updates the information at the ONTs from time to time.

Unlike other features, Billing features are not totally left on the ONTs because this information is very important for the service provider, and the service provider needs to be able to access/modify information frequently. This approach also provides an extra layer of security.

13) *Softphone* When a user starts the softphone application from any PC connected to the Internet, the client, after being authenticated, would automatically be connected to its respective ONT. A direct connection between the softphone and its ONT remains established as long as the user stays logged in. All the features in the softphone can then be provided by the ONT.

    Signaling would be backhauled to the ONT, but a voice path would be established client to client depending on whether the interface between the service provider's network and the outside network allows it. Whether other approaches are possible, or even needed, requires further research.

VII. Conclusion

With the increase in network bandwidth made available to users by ISPs, the feasibility of pushing the load on the centralized infrastructure to the edge of the networks must be considered. We propose here a network architecture built on a modern Fiber-to-the-premise network that can be used to provide premium Peer-to-Peer Voice-over-IP services to users. The load on the centralized voice servers is reduced tremendously as most of the job is now done by the ONTs which sit on the user premise at the network edge. The few servers that remain are simple databases running DHTs and are mostly used for data retrieval tasks. ONTs talk to each other directly for voice calls. In cases where ONTs involved in a call are located close to each other in the network, the voice data exchanged stays at the edge of the network. Such calls also save a lot of network bandwidth in the network core. Secondly, there is almost negligible effect of infrastructure failure in the network. For example, failure of a server is handled by updating the associated ONTs and attaching them to another existing server that is part of the DHT. Because of the use of DHTs, data redundancy ensures no data loss. Also, failure of an ONT only affects the user(s) directly associated with that ONT.

References

[1] M. Abrams, P. C. Becker, Y. Fujimoto, V. OByrne, and D. Piehler, “FTTP Deployments in the United States and Japan - Equipment Choices and Service Provider Imperatives,” *IEEE Journal of Lighwave Technology*, vol. 23, no. 1, pp. 236–246, January 2005.

[2] “Skype,” Website, http://www.skype.com.

[3] S. Rhea, B. Godfrey, B. Karp, J. Kubiatowicz, S. Ratnasamy, S. Shenker, I. Stoica, and H. Yu, “OpenDHT: A Public DHT Service and Its Uses,” in *SIGCOMM*, August 2005.