

Allworx VoIP White Paper

February 2005

Author: Jeffrey R. Szczepanski, Chief Technical Officer, InSciTek Microsystems, Inc.

Allworx
Division of InSciTek
Microsystems, Inc.
635 Crosskeys Office Park
Fairport, NY 14450

www.allworx.com
info@allworx.com
1.866.Allworx
585.421.3850 - Main
585.421.3853 - Fax

Copyright © 2005 All Rights Reserved - InSciTek Microsystems Inc. No part of this document may be used or reproduced in any manner whatsoever without written permission including quotations embodied in critical articles and reviews.

INTRODUCTION

The Allworx 10X Server and family of Digital VoIP phones are designed to meet the communications and networking needs of the typical small business, while also simplifying the setup and maintenance of the voice and IT infrastructure for the business owner. In fact, the primary mission of the Allworx product line is to deliver to small businesses the most recent round of IT technologies, including VoIP, in 'turn-key' solutions that previously were only practical for large enterprises with full time administration teams. While Allworx makes these technologies much easier to use and deploy, having a high level grasp of the various technologies involved goes a long way to understanding VoIP systems better and being able to diagnose issues, should they occur.

This paper is intended as a tutorial on several topics as they relate to deploying and maintaining digital phones on a VoIP network. In particular, this document aims to arm administrators, installers and network planners with information to help in successfully taking advantage of the new technologies associated with moving voice traffic over data networks, both on a LAN and across several sites via a WAN.

1 Phone system

2 Network server

3 Message center

WHAT DOES VOIP REALLY MEAN ANYWAY?

The term 'VoIP' is officially an acronym for **V**oice **o**ver **I**nternet **P**rotocol, but is also used to loosely refer to any application where packet based data networks are used to *packet switch* telephone calls in real time. This type of telephony contrasts to traditional hard wired analog telephony which is *circuit switched*. Going to VoIP based telephony technology has several advantages, both technical and economic, but also introduces some new complexities that must be a managed part of the data network.

Traditionally, data networks and the Internet in general were developed only as a best effort service. The network was designed to get the data there as fast as possible and when there are problems to get as much data there as possible, eventually. This is a good design characteristic for data, but it has problems where there are true real-time constraints to support *toll quality* telephone calls. For telephone audio, not only is bandwidth and throughput important, but packet loss, latency, and jitter performance are also critical factors to good sounding audio. Therefore real-time applications like VoIP gave rise to engineering and managing the **Q**uality **o**f **S**ervice (QoS) of data networks.

Designing networks for QoS factors and diagnosing QoS problems is an entirely new dimension in data networks for many people. In VoIP applications, not only is a valid data connection required to insure application success, but also a high quality and maintainable QoS. It is not uncommon for data networks to have throughput or packet loss problems that go completely unnoticed until VoIP systems are deployed. Therefore, when deploying VoIP systems it is important to inspect or validate the existing network to make sure it is going to be VoIP ready from a QoS perspective. This is especially important consideration when VoIP calls are going to be placed over data connections between physical locations. QoS topics are further explored in a later section of this document, but the main items of interest are network packet latency, jitter, and packet loss rates.

ALLWORX VOIP FEATURES

The Allworx 10x server is a sophisticated VoIP PBX and gateway designed specifically for use by small businesses. The Allworx 10x not only supports traditional analog telephony in a circuit switched fashion between its analog FXO loop interfaces and FXS analog extension ports, but also acts as a VoIP gateway to bridge the digital data packet based world together with the analog world. As a result, the Allworx product supports ordinary analog telephones, analog telephone lines and advanced VoIP telephones simultaneously in a seamless fashion.

The Allworx VoIP technology platform is intended to make the following key features as seamless as possible:

- **Auto Station Discovery:** Simply plug a new Allworx VoIP station into the network, and the Allworx server will automatically discover and configure the new station without manual intervention.
- **Simplify Move/Add/Change Administration:** With Allworx VoIP phones, station identity moves with the station, not with the physical cable drop used. Therefore, moving phones to new locations within the office is typically as simple as actually moving the phone to the new location and plugging it into another network drop.
 - **Remote Phone Capability with Remote plug and play type functionality:** Installing an Allworx VoIP phone at remote site, such as at home, and seamlessly operating it as if it was directly connected to the LAN at the office.
 - **Site-to-Site Toll Bypass Calling:** 7-Digit dialing from one Allworx system to another Allworx system at a different site, without placing a call on the regular phone network.
 - **Integration with Internet Telephony Service Providers (ITSP):** Allworx server acts as a proxy server placing and receiving low cost Internet telephone calls through ITSP providers without requiring use of regular phone lines.
 - **Third party SIP Gateway product integration:** Expand Allworx capabilities via LAN connected and SIP based FXO, FXS, and/or T1/PRI 3rd party gateway products.

SIP PROTOCOL AND VOIP

The Allworx VoIP platform is built around the industry standard VoIP protocol known as **S**ession **I**nitiation **P**rotocol (SIP). SIP is a packet-based protocol built on top of the standard IP stack using the User Datagram Protocol (UDP/IP). It is possible to do VoIP telephony using other standards besides SIP, such as H.323 or MGCP. However, SIP was specifically designed to be used with IP stacks, and was designed with the Internet protocols in mind. While historically much of the older installed base of packet based telephony used MGCP or H.323, it is generally accepted that SIP represents the future of VoIP and nearly all new installations using an industry standard protocol are deployed using SIP.

SIP is not the whole story when it comes to IP based VoIP and SIP itself only actually describes one out of the necessary three functional elements required. That is, when designing a VoIP protocol, three basic functions need to be provided:

1. **Call Control and Setup/termination** – A set of mechanisms to locate the intended dialed party, determine their availability and accept or deny their requests

1 Phone system

2 Network server

3 Message center

2. Session Negotiation – Once a new call is going to be accepted, determination of the format and network locations for transporting audio between the actual end points
3. Media Transport – Once accepted and then negotiated, providing the real-time audio transport between the end-points for the duration of the call.

SIP, itself, actually only provides the first item described above of call control and setup. Other protocols are actually required to perform the other two remaining primary functions. Although, when people talk about SIP, two protocols are also generally implied and come along for the ride:

- Session Description Protocol (SDP) to negotiate the session media types (G.711 or G.729) and the IP address and port number that each end-point should transmit towards.
- Real Time Transport Protocol (RTP) to actually move coded audio data during the live call. Therefore, when people talk about SIP VoIP telephony several things are implied here to be available and working properly for successful phone calls to carry on:
- Reliable IP routing data connectivity of UDP packets between associated phones and their associated gateways or proxy servers for basic network transactions (Network settings, DHCP, DNS, etc.)
 - SIP Protocol and Proxy configuration to locate intended parties and determine their availability (ringing or busy).
 - SDP Negotiation to determine the final coder type, IP addresses and port numbers that should communicate actual audio data
 - RTP to transport coded audio over a network with an acceptable QoS level from end to end

1 Phone system

2 Network server

3 Message center

The Allworx server acting as a VoIP gateway contains all the necessary facilities to make all of the above happen in as simple a manner as possible. However, when one of the above mechanisms is being interfered with on the network, certain types of interesting symptoms may result such as dropped calls, choppy audio, one-way audio, echo, etc. The focus for the installer and site administrator is to be certain that things are configured properly in the environment to ensure proper data connectivity and QoS between end points. Looking forward, the remainder of this document will be dedicated to helping administrators and installers understand the various potential pitfalls better, so that they may troubleshoot and resolve networking or configuration difficulties.

ECHO IN VOIP NETWORKS

From time to time, echo can be a problem in telephone calls. Certainly everyone has experienced echo in phone calls at some time or another, even for ordinary non-VoIP calls. Echo is most commonly experienced in international calls or long distance calls to rural areas, and of course in calls involving a cell phone. Unfortunately, the characteristics of VoIP telephony connections increase the opportunity for echo problems to occur. This happens due to the introduction of additional latency (delay) of the voice as it travels from source to destination over the packetized data network.

To the human observer, echo of his or her own voice is only noticeable when it is heard back with some amount of delay. Echo without delay simply sounds like side-tone. Side-tone is the sound of you talking simply fed back directly between microphone and speaker without any delay. Side-tone is

normally introduced purposely on phones so that the phone does not sound “dead” when you are talking. Therefore, when echo exists in the analog phone network, especially in local calls, any echo is completely covered up by the existing side-tone. However when a VoIP system is attached to that very same analog phone line, the additional latency of the data network now carrying the voice to/from the IP phone, now makes that echo noticeable to the user since the echo now arrives back well after the talker has finished each utterance.

Echo round trip delays of only a few milliseconds (ms) are not really noticeable, but as the delay accumulates into the area of 10ms, the system will start sounding hollow and eventually will start sounding like the reverb of a large stadium echo. As echo latencies run into the area of 50ms and beyond, the talker’s own speech will be followed by a very distinct echo of the same speech back in the talker’s ear. At this point, unless the echo is very quiet, it starts to become very annoying.

Generally speaking, the loudness and the latency impact how objectionable the echo sounds to the talker. That is, as the echo gets quieter and/or less delayed it becomes less objectionable. Alternatively, we can say that the more delayed the echo is, the quieter it needs to be, to be acceptable. Therefore, for echo to be acceptable in a VoIP system, it typically needs to be quieter to the VoIP caller than it was to start with in the analog only part of the network to still be acceptable. This is the role of the *echo canceller* in a VoIP system, but we will come back to that later.

Where does the echo come from?

Echo results in the phone network where 2-wire phone lines carrying voice in both directions on the same wire pair are converted into 4-wire circuits where a separate wire pair carries the audio in each direction. The analog device that does these conversions is called a *hybrid* and its job is to convert back and forth between the 2-wire analog loop world and the 4-wire Central Office (CO) switch world. If *hybrids* were perfectly matched devices to the particular phone and phone lines installed at every site, there would be no echo. However, in the real world, these hybrids are not installed with perfect impedance matches, and therefore echo results when the sounds “bounce off” these hybrid devices.

In a typical analog phone call, there will be at least two hybrids involved in the call. One at the CO for the calling party (near end echo), and one at the CO for the called party (far end echo). Beyond the hybrid *electrical* echo, there can be other sources of echo as well. The most common other type of echo is *acoustical* echo at the far end, when you call somebody that has a low quality telephone set, or is using a desktop type speakerphone. The near end echo is determined by the local phone line loop you are making the call on, while the far end echo depends on the party that is called. For this reason, near end echo is generally referred to as “Line Echo” and the other sources of echo are collectively referred to as “Network Echo”.

The delay times for “Line Echo” are not a concern in ordinary analog phone calls, because the delay path is short enough that the echo sounds like side-tone as described before. So, the phone company can ignore the effects of line echo. However, due to the delays inherent in “Network Echo” that typically is a problem, even for the regular phone network, especially with calls that cover long distances. Therefore, the phone company typically has to do something about that and have devices built into the long distance phone networks called **Network Echo Cancellers** (NEC) to remove such echo. Therefore, it is generally safe to assume for long distance calls, certainly, that “Network Echo” is not of a particular concern, even with VoIP systems attached to the PSTN.

1 Phone system

2 Network server

3 Message center

The need for Line Echo Cancellers in VoIP

As we stated above, the PSTN is not concerned with Line Echo since it will sound like side-tone. However, if we attach a LAN VoIP system to a PSTN gateway device, like Allworx, then Line Echo all of a sudden becomes a specific concern in the system because the hybrid echo coming back from the line is now delayed by 10's of milliseconds in the IP network and would no longer be acceptable to the VoIP phone station user. Therefore, VoIP gateway systems employ a device in their line interfaces called a **Line Echo Cancellor (LEC)** that can cancel up to say 16 or 32ms of echo resulting from the hybrid installed on the local phone line at the CO.

If VoIP systems have a LEC and the PSTN has NECs, why do we still hear echo at times?

This is a complicated question with several different answers:

- An echo canceller is a very sophisticated device that automatically attempts to dynamically detect, adapt to and remove all echo on the fly while still providing true full-duplex speech performance. Neither LECs nor NECs are perfect devices and depending on their design trade-offs of a given implementation, will exhibit certain strengths and weaknesses depending on their environment they are operated in.
 - The phone company NECs never perfectly converge down to zero residual echo. When a VoIP system is introduced at one end of the connection, the increased delay makes what ever residual echo is there, potentially more perceptible. As described previously, this added delay gives the perception that the echo is worse even though the magnitude of the echo signal is actually the same. In a given call, depending on the exact level of the residual echo, this may or may not end up being objectionable to the VoIP system user.
- Regional "intra LATA" calls can be very problematic relative to network echo. Because the latency in the network is not significant, the phone company doesn't usually bother to deploy the relatively expensive NEC's for intra-LATA calls. The delay is relatively short, so the echo is not typically objectionable when using an ordinary analog phone at each end. However, again, add a VoIP system to one end and the network echo can be a real problem. This most often occurs in markets where when placing short haul calls between competitive local or regional companies, and the called party has a particularly high level of far end echo coming back. The inherent latency of the echo falls into the hole between the LECs ability to combat the echo and the lack of deploying a NEC in the phone network. To be clear, that echo was always there, it just took having the VoIP system being installed to actually hear it.

Shouldn't a VoIP gateway then have both a LEC and a NEC?

Deploying both a Line Echo Cancellor and a Network Echo Cancellor in a VoIP gateway or PBX has some advantages. In particular, it can help with the regional calling area calls, that is intra LATA calls discussed above, which are typically the most problematic cases for a typical VoIP gateway solution. However, intra LATA calls are a small percentage of most user's calls, and having the NEC operating for all other types of calls too – the ones that already sounded really good - presents some problems. The fundamental concern is having two different NEC's operating on the same call, the NEC in the VoIP gateway and the one in the phone network. It can be difficult to get them to work reliably together and in many cases does more harm than good. You run into a situation where maybe 5% or 10% of calls are improved quite a bit, but the remaining 90% actually got a little worse. Which is a better trade-off?

1 Phone system

2 Network server

3 Message center

In the end, most VoIP systems installed as end user customer premise equipment, including systems like Allworx, employ only LEC's and do not employ NEC's. The reasons for doing this, as described before, are both technical and economic. NEC's require significantly more processing power than LEC's and therefore are historically very expensive to implement on several channels at a time. It seems that the market forces have shown that the added costs are not out weighed by the functional benefits.

Conclusions

Besides the new learned skills for both the installer and the administrator relative to data network Quality of Service (QoS) that is discussed in the next set of sections, echo is the biggest hurdle for VoIP systems to overcome, as they improve with each generation. Echo cancellers go a long way to maximizing the user's perceived audio quality, but still represent one of the areas for the relatively new VoIP technologies to improve. Looking forward, the quality and capabilities of the echo cancellers will continue to improve, but the only thing in the end that is going to completely eliminate echo sources is when VoIP systems no longer need to interact with the analog phone loops that date back to the designs of Alexander Graham Bell. That is, once calls between all end-points are completely digital, only then will these problems of occasional or persistent echo be a thing of the past.

1 Phone system

2 Network server

3 Message center

Bandwidth Calculations

This section attempts to provide details for a technical foundation of bandwidth calculations in Allworx supported VoIP applications across the LAN and WAN. Also included are deployment recommendations that may help resellers and end-users with their application rollout.

- In a VoIP telephone call the caller's spoken voice is converted to electrical signals which are then coded into data network packets or traffic. The coding/decoding (codec) scheme and the packet transmit interval collectively determine the amount of bandwidth consumed per call.
- G.711 calls send and receive a stream of Ethernet frames, each 214 bytes long, at a rate of 50 per second (20 ms interval). The bandwidth required is 85.6kbps in each direction of the call. Calls that traverse Frame Relay, ATM/DSL, PPPoE or VPN links will consume more bandwidth due to the additional encapsulation of the transport protocol(s).
- G.729A calls send and receive a stream of Ethernet frames, each 74 bytes long, at a rate of 50 per second (20 ms interval). The bandwidth required is 29.6kbps in each direction of the call. Calls that traverse Frame Relay, ATM or VPN links will consume more bandwidth due to the additional encapsulation of the transport protocol.
- Both G.711 and G.729A codecs may be used for VoIP calls through the Allworx. This typically involves a call from a VoIP telephone through the Allworx server to another VoIP telephone or soft-phone. Codec preference settings at each endpoint will determine which is used. For Allworx IP phones, this setting is configured on the handsets pages for each particular station.
- Only the G.711 codec is supported for calls in which the Allworx server is an endpoint. This typically involves a call from Allworx-attached analog telephone to another phone (VoIP or analog), a call from a VoIP telephone to the Allworx voicemail system or auto attendant, or a call from a VoIP telephone to an external user over analog public telephone network connections.

Codec Support Matrix

Calling/Called Endpoint	Codec Supported
Cisco 7905	G711
Cisco 7912	G711
Cisco 7940	G711 & G729A *
Cisco 7960	G711 & G729A *
Allworx 9102	G711 & G729A *
Allworx 9112	G711 & G729A *
Analog set Analog telephone line	G711
3 rd party gateways	G711 & G729A **
Auto Attendant	G711
On-Hold Music	G711
	Voicemail
	G711

* Depends on preference settings and the capabilities of the other endpoint.

**Subject to the capabilities of the 3rd party equipment.

Capacity Planning

Using the worst-case measurement of available bandwidth from above you can calculate the maximum number of simultaneous calls supported over the Internet connection. This of course assumes no other Internet activities are being performed by the local users such as web surfing, email, file transfers, music downloads, etc.). Moderate to heavy use of the Internet connection for other applications will degrade the quality of calls and may substantially limit the number of calls supported over the link.

Available Bandwidth	Simultaneous G711 Calls	Simultaneous G729A Calls
128K	1	4
256K	2	8
384K	4	12
512K	5	
768K	8	
1M	11	

SIP PROTOCOL AND NAT FIREWALLS

The SIP protocol was designed and first implemented before security issues and the necessity for NAT/Firewalls existed. This means that VoIP applications and the associated protocols of SIP and SDP were not designed with **N**etwork **A**ddress **T**ranslation (NAT) in mind. In fact, SIP/SDP negotiations are typically broken when a NAT device exists between the negotiating end-points

- 1 Phone system
- 2 Network server
- 3 Message center

whereby resulting audio not being available in one or both directions after a call is setup. While a full discussion of this topic is beyond the scope of this document, the glossary at the end of this document gives a brief description of NAT. Relative to the affect of NAT on VoIP protocols an understanding of the basic problem is useful, and is the topic of the remainder of this section.

NAT actually interferes with several different common protocols, with SIP/SDP being only one pair of them. NAT actually breaks almost all protocols that need to embed IP addresses and/or port numbers in their own protocol messages. This is best explained through an example. Lets assume two phones are trying to talk to each other over the Internet. Each phone is behind its own NAT firewall each at two different sites and the LAN network addresses of both sites is 192.168.1.0 with a subnet mask of 255.255.255.0. When a call is setup, each phone is going to report through its SDP information an address of 192.168.1.x (its local IP address and port number) to the remote party. However, it is clear that neither end is going to be able to send to the SDP reported address and get the intended recipient. In fact, this will be a problem not only if the LANs had used the same network address, but if they used any non-publicly routable network address. For this example, audio will not flow properly in either direction using normal SIP/SDP negotiations.

More typically, only one end is directly behind a NAT device, such as when contacting a remote VoIP gateway that connects to the PSTN. In these cases audio typically works in one direction but not the other. There are other problems beyond the basic logical NAT routing problem, even if we did get the IP addresses right somehow. The firewall function of the NAT firewall introduces other packet filtering problems since the whole point of the firewall is to prevent arbitrary packet data from entering the LAN network. To a simple firewall not tracking all the SIP/SDP sessions going back and forth, the audio data coming from the remote is simply blocked as illegal data, even if it did manage to route from end to end.

Since both the NAT/Firewall and VoIP services are desirable, what do we do about SIP and NAT? Well, typically, there are several pieces required to get this to work correctly, with the key element being a special device called an **Application Level Gateway (ALG)**. This is traditionally a separate type of special NAT firewall that is "SIP aware" and is specially configured to monitor the context of everything going back and forth between end-points, altering SIP/SDP/RTP packets and opening/closing holes through the firewall to allow the audio to negotiate and flow correctly.

Alternatively, you can use Allworx as your VoIP gateway with Allworx phones and they work together to solve all these problems for you automatically – even when using 3rd party firewalls! The details of this are explored in the next section.

ALLWORX SOLVES THE SIP NAT PROBLEM

Using VoIP protocols with NAT/Firewalls can be a big headache unless you are using Allworx equipment. Allworx products are designed to work together and automatically discover the networking topology between end-points and adjust all VoIP negotiations accordingly. Allworx products are actually able to do this even when 3rd party firewalls are involved in the path, but this requires the use of Allworx IP phones or servers at the associated end points. Non-Allworx end-points may not support all the necessary mechanisms to make this possible.

1 Phone system

2 Network server

3 Message center

Generally speaking, the primary requirement to make this possible is for each Allworx server to have its WAN port connected directly to the Internet at a publicly routable IP address. Note: The Allworx server does **not** have to be the primary data NAT/Firewall for the LAN, only that the Allworx server has a publicly routable WAN connection in parallel with an existing firewall.

Remote end-points, such as Allworx IP phones on a LAN can typically be behind any single NAT/Firewall whether it is an Allworx server acting as the firewall or any third party NAT/Firewall product. InSciTek specifically tests against Cisco/Linksys™ and Sonicwall™ product implementations as base verification reference points of 'typical' firewalls. Since InSciTek doesn't control the implementation of 3rd party products, it can't guarantee proper operation, but would typically expect things to work with most firewalls, including the ones InSciTek tests against.

QOS ON THE LAN

This section relates to configurations where both the Allworx server and the IP telephones are located at the same physical site. Typically, all devices on the site will be on the same LAN subnet. That is, when routers are not needed for internal LAN connectivity.

Strictly speaking, the textbook recommended configuration for such a site is to configure a pair of VLANs using managed Ethernet switches. One VLAN would be for voice traffic and the other one VLAN would be set at lower priority for data. Such a configuration basically guarantees that any amount of data traffic loading or problems cannot interfere with the voice traffic. Allworx phones and their built in switches have VLAN support to integrate with such a configuration, should it be desired. However, going to a fully managed VLAN configuration on a small business site is rarely actually done both for administration complexity and cost reasons. Generally speaking, following these tips will allow things to work well:

- Ensure the Local Area Network (LAN) is free of legacy hubs or repeaters and coaxial cable network segments. A fully modernized network with fully 10/100 switched Ethernet infrastructures is ideal.
- Minimize the number of Ethernet switches installed in the closet. Daisy chaining together small switches to add more ports, adds latency and increases traffic flow bottle necks. Installing one 48-port switch is much better than installing four 12-port switches.
- Group all VoIP devices onto the same Ethernet switch if possible
- Generally, there is no need for queuing capable switches or routers on the LAN. Both the Allworx and VoIP telephones employ sliding packet buffers that mask the modest packet loss and jitter (a.k.a. variable delay) associated with busy LAN networks. Managed routers and switches is typically only a concern in large enterprise networks.
- Don't bother with a VLAN setup on the network switches unless the customer's LAN is very large or the users are extremely heavy data users. If either of these is the case then Allworx should be configured as a LAN host and a voice VLAN be built on a separate switch to handle the telephony traffic.

1 Phone system

2 Network server

3 Message center

The simpler the site, the better the above set of tips and suggestions hold up. Where things start to break down is if network data traffic is regularly very heavy and the network is getting overloaded, or when the site has onsite routers to direct traffic between more than one local subnet. If voice traffic is going to be flowing through those routers along with data, significant attention to QoS topics will be required to insure proper operation 100% of the time. While a full discussion of this case is beyond the scope of this paper, the next section talking about QoS issues over a WAN hints to some of the issues involved.

QOS ACROSS A WAN

This section relates to configurations where the Allworx server and one or more IP telephones are located at different physical sites. Additionally, this section applies to cases where multiple Allworx sites are connected together in a site-to-site manner or when the Allworx server is configured to take advantage of an Internet Telephony Service Provider (ITSP) for calls to/from the Allworx server.

QoS topics across a WAN are of particular concern. The reasons are both physical and historical. The historical part of the problem is that IP protocols and the Internet in general were originally only really engineered to move data – all treated pretty much equally on best effort basis – it doesn't matter if the data is email or if the data is coded voice traffic. As it sits today, there is no standardized way for the public Internet to support prioritized traffic between arbitrary end-points. Those protocols are still evolving and the installed base of Internet infrastructure is not fully equipped to support the protocol standards, even where they do now exist.

The basic physical problem here is what do you do when bursts of data exceed the bandwidth of a rather limited size pipe. This is a complex topic that has several aspects to it. For example, there are priority and traffic shaping trade-offs that affect both the effective latency and available throughput of the different traffic classification types. The traffic patterns and needs of different sites are different and have to be administrated with site specific knowledge of policies and priorities desired.

With that said, in many circumstances, ordinary Internet connections will generally carry voice traffic pretty well most of the time. For most users the potential reliability disadvantages are greatly outweighed by the cost advantages of a simple ad-hoc WAN setup. While guaranteed operation can only be obtained through a carefully engineered and managed QoS plan, sticking to the following guidelines will pave the way for a cost effective solution using only ordinary Internet connections that may already be in place:

- Do not attempt to deploy VoIP service using a dial-up connection. Dial-up connections are too easily overloaded by even modest data traffic.
- For remote telephony applications to work through Allworx the Allworx server must have its WAN interface directly connected to the public Internet. This is discussed in more detail in previous sections. In particular, calls to/from an ITSP service will not typically work if the Allworx server is behind a firewall.
- Prior to deploying VoIP between two sites it is highly recommended that you first test your Internet connection to determine the speed of your link. Test the speed several times per day over the course of a week and base your planning on the worst rate measured.

1 Phone system

2 Network server

3 Message center

- A VoIP call consumes symmetrical data on the network. Be sure your speed test results account for uplink and downlink performance. The lesser of the two values should be used for bandwidth planning.
- Determine what percentage of the available bandwidth will be used for voice. Generally it is better to not use more than about 50% of the available bandwidth for voice, leaving the remaining 50% for data applications.
- Compute the maximum number of calls your voice allocated bandwidth will support and configure Allworx VoIP server settings to limit the maximum number of calls accordingly so the desired limits are not exceeded.
- Use your local ISP's speed test if the remote application will traverse the same provider's network. Check your ISP's home page for a speed test link or visit BroadbandReports.com for a comprehensive list of 216 global sites to test your bandwidth <http://speedtest.broadbandreports.com>.
- Test the speed to the remote user's ISP if they are on different providers. This will expose the performance of the peering connection between ISP's and provide a better perspective of the bandwidth available for the application.
 - Use the free Brix Networks test utility for VoIP for Internet connection assessments. This will test your Internet connection's ability to handle VoIP calls. It also gauges the quality of the call in comparison to traditional and cell phone call qualities. Check them out at <http://www.testyourvoip.com/>.
 - Allworx does support QoS tagging of voice traffic. However, it will make little difference in the caller's experience over the Internet with normal ISP based services. Bandwidth availability should be the main concern today because ITSPs do not manage call quality to the customer using QoS features. The ITSPs that manage call quality only do so within their core networks; quality to the customer through ISPs is considered a best effort and will not be managed towards your network unless you subscribe to a dedicated private service with a specific Service Level Agreement (SLA) in place.

1 Phone system

2 Network server

3 Message center

KEY SYSTEM FEATURE – ALLWORX BLF PROTOCOL

The Allworx server and IP phones, whether local or remote, have several differentiating features that enable emulation of the classical key system type capabilities such as line appearance, busy lamp field monitoring, and direct station selection. This is all done on an industry standards compatible SIP VoIP platform via the addition of some Allworx advanced mechanisms built on top of SIP.

Specifically, InSciTek has added some SIP specification compliant private headers to the Allworx implementations of SIP to activate the advanced features. Such features include support for NAT enabled remote phones and to operate things like automatic off-hook for direction station selection and such. However, the live system status monitoring required for things like busy lamp field indicators and line appearances, goes beyond what SIP was designed for. As a result, InSciTek developed a companion protocol for SIP to offer some of these advanced features. The protocol is referred to as Allworx Busy Lamp Field (ABLF) Protocol.

The internal syntax of the ABLF is not important here, since Allworx takes care of the details. But what is potentially helpful to the system administrator is some understanding of how ABLF operates at the IP level. This will assist in troubleshooting specific problems where ABLF or Line Appearance lights do not seem to always operate properly on a particular phone, or some set of phones.

ABLF Protocol Background

The Allworx BLF protocol is an event driven peer-to-peer protocol and is implemented by both the Allworx server and by all Allworx IP phones. Because ABLF is a peer-to-peer protocol that uses subnet broadcasts to reach all devices simultaneously, all ABLF devices associated with a particular Allworx server site must transmit and receive on the same UDP port number. By default this UDP port number is port 2088, but is configurable on a site-by-site basis using the “Servers -> VOIP Server” page of the Allworx administrative web site to change the BLF port option. Note when this setting is changed, all Allworx devices (IP phones and server) must be restarted to acquire the new configuration information and to keep all devices monitoring and sending with the same correct port number.

ABLF Protocol Operation

Each time an ABLF peer has a change in status, it broadcasts this information to all peers on its subnet, and if not located on the same LAN subnet as the Allworx server, also sends a directed packet with the same content to the Allworx server so the server can forward that packet to all other subnets with phones attached. If the phone is an Allworx remote phone, this directed packet goes to Allworx’s WAN address, otherwise this is typically directed at Allworx’s LAN port since the device is on the private side of Allworx’s firewall. The ABLF device automatically determines the interface to use when it is configured during startup.

The Allworx server keeps track of a list of all local subnets that contain ABLF devices and also contains a list of each remote device participating in the ABLF protocol. When the server receives a packet from any device, it automatically forwards the packet to a single device on every other subnet other than the originating subnet and requests that one device to broadcast the packet on its own local subnet. This mechanism allows learning of the topology of all ABLF devices and to managing the traffic so that every device gets each ABLF update notification in as efficient a manner as possible.

Troubleshooting ABLF

Generally speaking, the ABLF protocol and Allworx devices take care of themselves even when NAT/Firewalls exist at the remote sites sitting in front of the remote phones. This results, because the Allworx server and the phones work together to safely traverse NAT devices within the same restrictions explored previously when talking about SIP. Note that it is not even normally necessary to open a specific NAT firewall device port at the remote site to support ABLF as the phone will keep an automatic port open with the server. In the event that one or more devices or subnets are not correctly getting ABLF packet updates, usually something is specifically administratively blocking traffic on the configured ABLF UDP port, (port 2088 by default). Make sure that nothing is specifically blocking port 2088 to the Allworx WAN port and that no intermediary NAT/Firewalls specifically limits port 2088 traffic from local LAN to WAN. If port 2088 is limited or blocked you can change the Allworx servers receiving UDP port as specified previously.



REMOTE OFFICE PHONES

The Allworx server and Allworx VoIP phones are designed to work together as seamlessly as possible. In particular, InSciTek has tried to make it as straight forward as possible to install and maintain remote office phones – nearly as easy as it is to install and maintain local main office phones. In this section, remote phones refers to a standalone phone that operates in conjunction with an Allworx server located at distant site without having a local Allworx server installed. The most typical example would be a phone installed at an employee home, working off the Allworx installed in the main office.

In all circumstances, it is important to have the main Allworx server installed with its WAN port directly connected to a publicly routable IP address on the Internet. The Allworx server must not be behind a separate NAT/Firewall device, as the Allworx server will be dynamically creating inbound and outbound sessions to interact with the remote phone. The remote phone itself must have Internet access to the Allworx server, but the remote phone is allowed to be installed behind a single NAT/Firewall. Typically, this is the NAT/Firewall protecting the LAN that the remote phone resides on.

Configuration of the remote phone is mostly automatic and works similarly to installing phones on the main LAN of the Allworx server. The remote phone requires two specific settings to be set by the administrator:

Boot Server IP	The public IP address of the Allworx Server WAN port. This IP address tells the phone where to find the main office server on the Internet.
Remote plug and play key	<p>This numeric code number is the authentication key that the phone uses to authorize itself with the remote Allworx server. If the correct key is not entered, the Allworx server will not allow the remote phone to place calls through it. This is a security measure to help prevent unauthorized users from using services on the Allworx server.</p> <p>The current correct setting to enter is located in the configuration setting on the Servers->VoIP page of the server's administration screens. The value is common to all remote phones associated with a particular Allworx Server.</p>

1 Phone system

2 Network server

3 Message center

Beyond this, generally speaking, the remote phone operates the same as any other phone on the LAN of the main site. The station can make and receive calls and be configured just like any other station. There are only a few limitations in capabilities:

- Intercom works fine, but paging functions do not extend beyond the local LAN. That is, neither zoned nor overhead pages will typically play out at the remote site.
- If more than one remote phone is behind the same physical 3rd party NAT/Firewall, the remote phones will typically not be able to call each successfully due to limitations in how those firewalls operate. If this is a requirement, it is suggested that an additional Allworx server is installed and they are configured to work with each other in a true site to site fashion.
- Remote phones off a single Allworx server at different sites should be able to call each other, but if both those connected phones are behind NAT/Firewalls, getting audio between them can be problematic. It is sort of a chicken and egg type problem where neither phone can determine its public IP address nor RTP port numbers until after the opposite end receives the first audio packet. Since both ends are in the same situation, neither end ever receives the first packet. What is required is for at least one of the two phones to have a static mapping through the firewalls for its RTP ports. See the troubleshooting section for more details.
 - ITSP services configured on the Allworx server may not be accessible from remote phones that are behind NAT devices. This is a limitation of the service providers and not under control of the Allworx devices.

1 Phone system

2 Network server

3 Message center

Caveats and details about remote phones operating over the Internet, especially when using 3rd party firewalls:

- Normal ISP Internet access for regular residential and business customers is a best effort service. Specific QoS metrics are not guaranteed and poor quality audio can result at times, depending on traffic flow between providers and through Internet peering points.
- Most low cost NAT/Firewall routers do not prioritize traffic, and even if they do, the Internet service that they are being used with typically does not. Therefore, normal user data traffic activity can affect audio quality. For example, downloading email simultaneously while taking on the phone may cause audio interruptions.
- Depending on the service provider and the quality and bandwidth of your ISP service, the above issues may be rare or they may be regularly occurring events. See later sections for more details.
- If none of the above are acceptable from time to time, then leased line and/or virtual private circuit type service with a specific Service Level Agreement (SLA) is required – along with the proper router equipment – to guarantee that the desired QoS metrics are always met.

ZONED PAGING

Allworx system paging features use a special form of IP traffic routing called *multicasting* to direct zoned pages from the Allworx server to all Allworx phones configured for a particular zone. Multicasting is used because potentially many phones need to receive the page simultaneously and this is exactly what Ethernet and IP multicasting was designed for.

Phones subscribe to a multicast group to become a member of a particular paging zone since each zone has its own multicast address. Generally, this is all automatic and transparent to the end user, the administrator, and even Ethernet switches. However, when a particular site is having trouble with pages reaching one or more phones, understanding the configuration options and how Allworx transmits these pages should be helpful in diagnosing site configuration problems.

Allworx zoned paging is configured on the “Servers -> VOIP Server” page of the Allworx administrative web pages. In particular there are three parameters used to determine where Allworx transmits zoned pages:

Configuration Item	Description
Paging Base Multicast IP Address	This is the multicast base IP address (zone 0 - the overhead zone) of the system. Follow on zones are numbered sequentially from this base by adding the zone number to the base IP address. For example, zone 5 is located at base + 5 multicast address.
Paging Port Number	This setting is the UDP port number that the server transmits to at the multicast IP address for the current zone. All zones use the same port number, each with their own multicast address as described in the above field. The actual UDP payload is RTP packets of 20ms G.711 frames.
Paging Max Hop Count	This value controls the time-to-live count value in the IP header of all paging UDP/RTP frames. Typically this value is set to one (1) for a single LAN subnet, but if you have multiple LAN subnets with phones on them, this value may need to be increased.

Generally speaking, pages are only typically implemented on a single site and then typically only on a single LAN subnet. However, by manipulation of the above parameters, and through configuration of intermediate routers, it is possible for Allworx pages to span multiple subnets and even to tunnel across a VPN between sites. In short, what needs to be configured on the site is a way for Allworx server sourced multicast packets to be routed between subnets using the normal router

- 1 Phone system
- 2 Network server
- 3 Message center

specific configuration mechanisms. A full discussion of routing multicast IP packets and how this is configured into different brand routers is beyond the scope of this document.

Common Problems and Tips

This section explores some commonly observed problems and possible causes not specifically addressed by the previous sections. This section also provides some general troubleshooting tips.

One Way Audio

Having audio flow in one direction only after placing or receiving calls, or even perhaps having no audio flowing at all is a common problem with SIP and RTP, especially when NAT/Firewalls exist the path. As described in previous sections, Allworx products implement several mechanisms to automatically deal with the majority of configuration complexities required to get SIP and RTP to work with NAT firewalls. Still, it is possible for specific routing asymmetries or specific packet filtering rules to interfere with one or more protocol mechanisms.

The most common configuration problem in this regard exists in complex setups where the Allworx server is connected to the WAN and another firewall on the LAN is used for normal data traffic. The user level symptom is that a LAN phone will not get good audio either inbound or outbound to a remote phone located at another site. While one might typically conclude this is a remote site configuration issue since the LAN phone works normally when talking to other LAN phones, typically this is not the case. Generally, the LAN gateway is set in the phones so that Internet traffic traverses through the non-Allworx server firewall that is not SIP/NAT aware. In these cases, it is important to configure the LAN IP phones such that they use the Allworx server LAN IP address as their gateway to the Internet so that Allworx can correctly orchestrate firewall filtering to properly pass the remote phone audio traffic to the LAN properly.

Intermittent Connectivity or Devices dropping off the Network

In cases where connectivity between devices on the network or between IP phones and the Allworx server is spotty or random, usually some sort of configuration or network topology problem exists. The primary things to look for are duplicate IP addresses on the network, multiple DHCP servers enabled on the LAN or for more fundamental QoS issues including bad cables or improperly configured VLANs. Generally speaking it is best to start the investigation at the point of recent changes to the network topology or on the newest configuration settings – that is, whatever has changed recently.

Connectivity problems, especially when intermittent, can be very subtle and sometimes very difficult to track down. A network packet sniffer or protocol analyzer is a great companion to analyze packet flow and to look for problems. If all else fails, generally the best thing to do is isolate as many things as possible from the network and start adding back things one-by-one (over time) to discover what added device is interacting with the already existing devices on the LAN. Note, that just because a specific device starts the troubles occurring doesn't necessarily mean that it is at fault, only that it is a necessary accomplice. A careful review of packet data on the analyzer and/or configuration settings is always justified.

Common MAC Address on LAN and WAN ports of a router or Firewall



Most dual Ethernet interface firewall products, including Allworx use the same MAC address for both or all interfaces of the device. This is perfectly legal to do, however there is one specific instance where this can cause traffic routing difficulties. In particular, when two such dual interface devices are hooked together in parallel, both with a common routable path between networks, difficulties can arise. The most common case of this is when Allworx and some other router/firewall are hooked up with both the LAN and WAN interfaces hooked between the same pair of networks. I.e.: Logically the two WAN interfaces and the two LAN interfaces are connected together. This configuration causes a problem because of how some routing decision optimizations are done between layers two and three of the network protocol stacks. In effect, because the same MAC address is reachable from two different interfaces and both devices have visibility to both paths to those MAC addresses, neither device correctly can optimize the correct routable interface to use to reach the intended destination.

To prevent this from becoming a problem, one must configure the network such that both devices do not have visibility to both interfaces of the other device. Generally, the easiest solution is to avoid the above configuration, or if it must be used, configure the WAN sides of these devices to be on different VLANs so that they cannot see each others traffic. That is, assuming a hypothetical 3-port switch, put one device on port 1, VlanId=1, the other device on port 2 VlanId=2, and then the internet router on port 3, VlanId=1&2.

1 Phone system

2 Network server

3 Message center

DTMF Digits not Passed during live calls

Negotiation of how DTMF number key digits pass during live calls is the least standardized and most problematic aspect of doing VoIP telephony with SIP. Between Allworx branded equipment this will not be a concern, but it does often result as a problem when operating with ITSPs. The most typically exhibited symptom is that PSTN phone calls coming in from an ITSP will not be able to operate the Allworx auto-attendant or voice mail applications. The reverse may also be a problem where Allworx system phones will not be able to operate phone applications out on the PSTN. These problem(s) happen because the service provider may not have indicated correctly during call setup negotiations how it plans to send and receive DTMF digits.

Allworx expect all SIP parties to negotiate DTMF out of band so that DTMF can work correctly with all coder types (including G.729). If you are having such problems, the default negotiation settings can be controlled in the advanced settings tab of the Allworx server SIP Proxy and Gateway configuration pages of the administrative web site. Currently, the default RTP type and whether padding of RTP payloads should be applied is controllable.

Music On-Hold Sound Quality is Poor

Some of the voice coders are specifically designed to compress human voice signals and not arbitrary audio content. This is case of G.729 and its variants that take full advantage of the properties of human vocal tract to get such levels of high quality compression. However, when transmitting audio such as music through these compressed streams (which is not speech) the resulting audio is very distorted. Configure the system to use only G.711 coder if this is a concern.

Calling Auto Attendant

When a user is reporting problems with calls to or from a particular station, or when calling a particular party, it is usually best to let the Allworx auto-attendant assist in troubleshooting. If a station has correct connectivity to the Allworx server, dialing extension 400 should always return your own familiar auto-attendant outgoing greeting in short order. If there are delays in connecting

to the auto-attendant or a fast busy congestion signal is returned, then there are connectivity or connection problems that must first be addressed. Additionally, once the auto-attendant is successfully reached, dialing '#7' will have the auto-attendant report back audibly what station it thinks you are calling from so that your station identity can be confirmed. Additionally, once you hang-up, per the auto-attendant instructions, the system will automatically place a call back to your station to confirm return connectivity to you location. If the system is properly configured and good connectivity exists, this auto-attendant mechanism should always work and can be applied individually to each station to isolate which devices are having configuration or connectivity concerns.

Mapping Ports through Remote Phone Firewalls

When a particular Allworx server has remote phones at multiple sites, with each remote site having a NAT/Firewall device between the phone and the Internet – it is necessary for the remote phone's RTP ports to be statically mapped through the firewall if you wish the remote phones to be able to call each other successfully. It is not necessary to do this if the remote phones only wish to call the main office and not each other.

The steps to perform this operation vary by the type of firewall that the remote phone is being used with, but the basic goal is to map a range of UDP ports on the firewall's IP address to the LAN address ports of the IP phone. The Allworx server's admin page for each handset allows you to control the range of ports the phone will use for RTP sessions under UDP. What you need to do in the firewall product's configuration page is to map those same ports (one to one) through the firewall thereby allowing access from the WAN to LAN through those ports. Some products call this "DMZ'ing" those ports or perhaps "port mapping". Generally speaking the LAN IP address of the phone must also be entered into this firewall configuration page, therefore is it wise to make sure the phones LAN IP address will be fixed.

Note: That remote phones on the same site typically can't call each other directly. If you wish to be able to do this, generally you need to have an Allworx server at each site with multiple phones and then configure the servers in a full site to site topology.

1 Phone system

2 Network server

3 Message center

GLOSSARY

AEC – Acoustic Echo Canceller. An echo canceller that cancels the effects of audible room echo such as in a desktop speakerphone. See echo canceller.

Echo Canceller – An algorithmic device, usually implemented in software in a Digital Signal Processor chip to remove echo of the talker's voice from a full-duplex telephone connection. Echo cancellers are more sophisticated and differ from echo suppressors in that they eliminate echo without having to impose half-duplex communications on the link like echo suppressors do.

Full Duplex – Allowing communications including the movement of data or voice in both directions of the circuit at the same time. That is, simultaneous 2-way communications. A half-duplex connection allows communications both directions, but only one direction at a time.

Jitter – The amount of variability in latency as a function of time. In VoIP systems, jitter describes the irregularity of packet arrivals over the course of time. As jitter increases in the network, deeper receive buffers are required to smooth the effects of jitter, so as to not lose packets. Therefore, increased jitter typically results in increased overall end-to-end latency.

Latency – A specific type of delay in time of transmission or response to events. More formally, is the amount of time elapsed between two fixed reference points in a system that transmits information or takes action based on certain events. Latency is sometimes thought of as the time delay between cause and effect. In VoIP there are several sources of latency including network latency, coder latency, packetization latency, buffering latency, etc.

LEC – Line Echo Canceller. An echo canceller that cancels the effects of electrical circuit echo in phone lines. See echo canceller.

NAT – Network Address Translation. Generically NAT refers to devices that translate the IP addresses of packets as they transit from one subnet to another. NAT is normally associated with firewall devices and is used to hide private non-routable IP addresses of a LAN from the public Internet. Using NAT has security advantage and also works to minimize the number of public IP addresses required at a single site. Minimizing the use of public IP addresses is important because the pool of available addresses is a scarce resource.

Side-Tone – Feedback of talker speech from microphone to the speaker at their own ear piece to make speech sound as natural as possible. Without side-tone, phone handsets sound like they are broken even though speech is still being transmitted to the remote end.

For more information:

ALLWORX
635 CROSSKEYS OFFICE PARK
FAIRPORT, NEW YORK 14450
TOLL-FREE: 866-ALLWORX
TEL: 585-421-3850
WWW.ALLWORX.COM

© 2005 INSCITEK MICROSYSTEMS, INC. ALL RIGHTS RESERVED. ALLWORX IS A REGISTERED TRADEMARK OF INSCITEK MICROSYSTEMS.
ALL OTHER NAMES MAY BE TRADEMARKS OR REGISTERED TRADEMARKS OF THEIR RESPECTIVE OWNERS.

1 Phone system

2 Network server

3 Message center