

introducing voice over Internet protocol technology

Overly optimistic marketing during the so-called communications revolution has given voice over Internet protocol (VoIP) technology the stigma of being the next big thing that never materializes. I'll risk the same mistake by suggesting that VoIP is emerging as a viable Internet application.

Key VoIP Drivers

Broadly speaking, three factors will motivate the adoption of VoIP:

- Reduced ownership and operational costs
- Simplification
- A roadmap for building next-generation services

Operational cost is paramount. Long distance charges are the first and most obvious expense. Skeptics suggest that VoIP has missed its window because of fierce competition in the long distance market, but the market continues to grow, so the window remains open. Furthermore, even a few cents a minute is far more expensive than utilizing an otherwise unused – and already paid for – resource.

Personnel costs also add up. Most companies large enough to have their own private branch exchange (PBX) are staffed with a telecom group. Since VoIP is premised on open, interoperable standards, telephone services become an application more akin to running an HTTP server than a traditional phone system. One is able, then, to leverage the competencies of an IT networking group and invest resources there, which is attractive given the versatility of those staff.

Security is a feature that one gets “for free” with VoIP. VoIP is secured in the same way as other Internet services: by minimizing attack vectors, using strong authentication, and protecting important servers with a firewall. On the management front, IP phones can often be configured using DHCP and TFTP, for instance, which exploits services that usually exist already.

Simplification might save money directly but is a benefit in its own right. VoIP converges voice and networking, reducing the number of service providers a company must deal with. (Of course, I'm not predicting the death of the public switched telephone network or PSTN in the short or medium term.) On one hand, ISPs offer straightforward billing plans for bit pipes. On the other, VoIP empowers people to take control of the server infrastructure that telcos use to extract complex fees for every move/add/change operation.

The most important simplification, though, concerns the phones themselves. If VoIP is really just another Internet application, then are phones even required? So-called hard phones are available to recreate the experience of using a regular phone, but soft clients, ordinary PC applications that run on desktop computers, are making inroads. In time soft clients will dominate.

VoIP will offer new features over legacy phone systems, namely rich media and convergent integration. Rich media rejects the assumption that a communications channel spans only one medium (e.g., audio). Instead, voice, video, and text can be shared simultaneously. VoIP protocol engineers, particularly those with the Internet Engi-

by Ryan Kereliuk

Ryan Kereliuk is a contract developer focusing on systems and network software ranging from device drivers to protocol stacks. He holds a master's degree in CS from the University of Alberta.



ryker@ryker.org

Estimates suggest that VoIP can save 30% over conventional telecom rollouts after accounting for long distance charges, personnel, servers, and phones.

neering Task Force (IETF), have shown tremendous commitment to generality, so today's protocols are equipped to help people communicate in whichever ways serve them best.

Convergent integration – unified messaging – suggests that voice mail and faxes should, like email, be accessible using any device connected anywhere on the Internet. This unification might even extend to integrating VoIP software into customer relationship management suites or e-commerce applications.

The case for VoIP is closely tied to cost as well as potential new services. Estimates suggest that VoIP can save 30% over conventional telecom rollouts after accounting for long distance charges, personnel, servers, and phones.

Thinking About Deployment

A VoIP project, not surprisingly, begins with a requirements analysis that informs the trade-offs faced at every stage of the decision-making process. Key areas for consideration include cost, network requirements, protocol selection, client and server hardware and software, impact on existing infrastructure, and migration paths.

Regardless of the underlying technology, certain questions govern subsequent choices and, in the end, serve as evaluation criteria for any deployment. Three that deserve consideration are:

- System utilization
- Interoperability
- Quality

Utilization refers to the capacity of a given system and is related to the numbers of active and potential users: How many VoIP endpoints will be connected, and what percentage of those will be active at any one time? Interoperability is concerned with which users can connect with one another: Is the system meant for internal use only or for use with external parties? In the latter case, are external parties reached by using some (possibly different) VoIP protocol or by using a gateway to the PSTN? Finally, unlike legacy phone systems, VoIP offers the opportunity to trade quality for other benefits. What sort of quality do users expect in terms of system availability and media fidelity?

Call quality is governed by the codec in use, assuming the network transport is able to keep up. The term *codec* is familiarly expanded to coder/decoder but, these days, compressor/decompressor is an equally valid meaning. Different media codecs require different network resources but also provide different levels of quality. The G.711 codec, for example, provides quality equal to conventional telephone systems at a rate of 64Kbps. G.729A, by contrast, needs 8Kbps of bandwidth but sacrifices quality. There is no substitute for making test calls with different codecs, but my subjective impression is that G.729A offers a quality markedly better than cellular telephone service. Choosing a codec or building a model of codec usage is of great importance to resource planning. Simplistically, the codec bandwidth can be multiplied by the number of active users or voice paths to compute bandwidth needs.

A bandwidth number in hand, one can begin shopping for network service. Often forgotten, though, is that bandwidth is just one measure of a network. Latency, jitter, quality of service (QoS), and availability are other important considerations. VoIP systems are particularly sensitive to latency, which distorts the flow of conversation, because changes in speaking order are politely signaled by silence. High latency has the

effect of making multiple parties think the channel is open to new speakers – which of course results in a collision. A rough rule of thumb is that end-to-end latency should not exceed 250 milliseconds, but this number should be minimized. Jitter is a measure of the difference in inter-packet latency between packets leaving the sender and arriving at the receiver. Introduced by network elements, jitter can cause increased packet loss or perceptible delays, so its minimization is very desirable.

Making guarantees about bandwidth, latency, and jitter – QoS in networking parlance – is a hard problem that is not adequately solved in contemporary networks. These days, poor man's quality of service is achieved by overprovisioning, a very basic but effective technique. In summary, if one is outfitting a new location with network service or is changing service providers, it is a good time to ask harder questions than one might have in the past. Learn about a provider's reliability by interviewing existing customers, get comfortable with their problem-tracking and resolution procedures, and sign a service level agreement (SLA) that captures the important requirements.

Media codecs have direct implications for bandwidth. By contrast, calls are set up and torn down with a signaling protocol that is comparatively lightweight. When it comes to signaling, one is able to choose from a wide selection: Megaco, H.323, SKINNY, MGCP, SIP, etc. (to name a few); a comparative examination of these different protocols is beyond the scope of this article. The only protocol I work with today is the IETF's Session Initiation Protocol (SIP), which I strongly recommend to anyone implementing VoIP.

The signaling protocol one selects will obviously affect choices to do with client and server software and, to a lesser extent, client and server hardware. For SIP, many server choices are available, including high-quality open source and commercial packages. In terms of hardware, a good rule of thumb is that the machine used for an organization's corporate HTTP server is comparable to the machine needed to run SIP services.

On the client side, choosing between soft clients and hard clients is a key decision. Soft clients tend to be less expensive (or free), offer more features, and integrate more organically with existing desktop software. The downside, if there is one, is a learning curve similar to deploying any new application. Hard phones, predictably, offer a user experience that in most cases is identical or slightly improved over regular telephone handsets.

Regardless of the particular protocol choice, the infrastructure associated with deploying VoIP will be impacted. Hard phones, for example, don't reuse the resources allocated for existing PC desktops. When deploying hard phones, additional switch ports are required for each endpoint, more cabling may be needed, and routers are a factor if additional subnets are required.

Nodes (whether hard phones or desktops) configured with RFC 1918 private IP addresses will have issues communicating with people beyond the nearest network address translation (NAT) boundary. SIP, for example, includes routing information inside IP packet payloads, which means that vanilla translation systems do not work. Like FTP, it will be some time before NAT-enabled firewall devices are smart enough to rewrite these packets with appropriate translated address information. In the meantime, options include using public IP addresses for SIP endpoints or using a SIP-specific application-level gateway that is able to reconcile addresses inside SIP packets.

Lastly, any VoIP implementation needs to mesh with existing security policies and infrastructure. This might mean adjustments to deployed security systems such as fire-

Any VoIP implementation needs to mesh with existing security policies and infrastructure.

walls and RADIUS servers. VoIP security is particularly important, especially considering that authorized users can often access expensive resources such as PSTN lines through a gateway.

Finally, finding a migration path will be of key importance if one already has a significant investment in legacy phone technology or cares about reaching PSTN users. PBX users are best advised to contact their PBX vendor, as some vendors have VoIP options in the form of line cards for existing equipment. If PSTN connectivity is needed for a VoIP installation, a variety of choices are available in a variety of sizes. Running a T1 line to a VoIP-enabled router, for example, allows 24 simultaneous calls to be gated to the PSTN. On a smaller scale, some routers can take a module implementing a foreign exchange office (FXO) interface, which connects one or two lines to a telco. In short, telephony companies are sensitive to migration issues, and solutions are generally available.

An Example Implementation

Building out a VoIP network is not as complex as it might seem. In this section I will give a high-level description of one setup I've used that addresses some different requirements. This scenario describes the VoIP solution for a multi-office distributed company.

In this deployment, remote offices are connected to the Internet using business-grade DSL lines. When using this class of network connection, one generally doesn't have the influence necessary to negotiate a favorable service level agreement, but the good news is that these links are more than adequate for serving small satellite offices. The central corporate phone system is served by a VoIP-dedicated one-megabit link with an SLA guaranteeing latencies of 70 milliseconds or less to other predefined points on the Internet. This is a connection capable of serving 12 calls with G.711 or 42 calls with G.729A, though both codecs are used in practice.

This example uses the SIP signaling protocol exclusively, with a hodgepodge of different servers and clients. In terms of servers, for example, remote offices utilize open source SIP proxies, including VOCAL and SIP Express Router (SER), while the headquarters uses a commercial SIP server that supports unified messaging functions. The clients deployed vary widely, because the organization is a software development firm in which developers are permitted to use the client of their choice. Typically, however, desks are equipped with hard phones such as Cisco 7960 handsets, while roaming users use soft clients such as Microsoft's Messenger product.

The reality is that VoIP is not currently sufficient to reach all potential business partners and customers. In this deployment, then, PSTN connectivity is supplied using Cisco routers. One satellite office uses a Cisco 2600 router equipped with an FXO module to connect two lines to the PSTN for local dialing. Headquarters, however, uses a Cisco 3600 router with a T1 interface card to provide PSTN direct dial numbers into 24 different VoIP endpoints. By terminating these latter PSTN lines at SIP phones, some of them remote, the organization achieves the effect of virtualizing its geographically distributed operations at the head office. Both PSTN compatibility systems work very well.

On Timing

There is no question that VoIP represents a major paradigm shift. The legacy telephony model is very strong and thus will not be unseated easily – and change will not happen all at once even in an environment of enthusiastic adoption. So when, then?

Before I say when, let me describe how. The value of a communications medium correlates with the number of users reachable using that medium, so the adoption rate will accelerate due to what economists call a network effect. This is developing in a couple of ways.

First, trends in the service provider sector augur well for the emergence of VoIP. We're presently seeing major telecommunication providers size up the next generation of services. Looking to the competitive long distance market, we already see more agile providers using IP links to traffic international PSTN calls, thus gaining a competitive edge.

There is a new breed of company, those offering local and long distance telephone service using VoIP technology. Vonage, for instance, offers flat-rate local and long distance calling in the United States, with the option of choosing a phone number from several American area codes. Early entrants like Vonage are aiming to capture business and consumer customers who aren't in a position to manage their own server-side infrastructure or PSTN interconnectivity. This parallels managed Web services, which coincided with the growth of that technology.

Second, consumers, too, are seeing new reasons for VoIP, including the sort of managed services just described. Other reasons include network access and software availability. As anecdotal evidence, I've been broadband connected for over five years, and among my friends, family, and coworkers this is universal (although Canada is a bit ahead of the curve on this technology). And, now, inexpensive VoIP software is widely available to leverage such connections. By the time this article is printed, Microsoft, the desktop juggernaut, will have released the latest version of their Windows XP Messenger program, which includes a complete SIP client. This enhancement, to be released in version five at Microsoft's Windows Update site, means that over 17 million PCs are potential VoIP nodes using the SIP protocol.

Remember the network effect? Within 36 months it will be possible to satisfy a healthy percentage of one's calling needs, both personal and professional, using VoIP.

Conclusion

It has been the intent of this article to raise consciousness about VoIP, a technology that has been rightly considered imaginary. I use VoIP systems every day for my telephone needs, so, to me, they hold great promise. I'm convinced that VoIP is, at once, both a solution for next-generation communications and a challenge for IT departments, so it would be wise to keep VoIP in mind. The good news is that people comfortable with networking and Internet services will catch on to VoIP very quickly.

In a future article, I will focus on the Session Initiation Protocol for VoIP. In that article, I'll look at some of the lower-level nuts-and-bolts related to implementing SIP services. I'll describe the different actors in a SIP system, with examples drawing on open source SIP applications. SIP is just an application protocol that runs on the Internet, so, as developers, we can contemplate writing new and interesting phone services without investing huge amounts of time learning proprietary protocols. To that end, we'll take a tour of one open source application built on an open source SIP stack.

SELECTED POINTERS

<http://www.voip-calculator.com/> – An online script to calculate network requirements based on anticipated utilization.

<http://www.vovida.org/> – An open source community site that offers a variety of VoIP software including the VOCAL SIP proxy.

<http://www.iptel.org/> – Makers of the SIP Express Router proxy server and providers of personal SIP hosting services.

<http://pulver.com/fwd/> – A free SIP service provider that currently connects over 40,000 users.

<http://www.microsoft.com/windowsxp/windowsmessenger/> – Watch here for the new Windows XP Messenger or download the old Messenger 4.7, which is also SIP compatible.

<http://www.vonage.com/> – A service provider offering SIP services with interconnection to the PSTN.

<http://www.resiprocate.org/> – An open source SIP stack optimized for speed.

<http://www.asteriskpbx.org/> – An open source, multi-protocol PBX suite for Linux.

<http://www.xten.com/> – Xten offers the lite version of their SIP soft phone as a free download.

<http://www.sipphone.com/> – A SIP provider that sells hard phones for use with its free service.