Main > Communication > **Conferencing**

# How VoIP Works

by Robert Valdes

### Introduction to How VoIP Works

If you've never heard of VoIP, get ready to change the way you think about long-distance phone calls. VoIP, or **Voice over Internet Protocol**, is a method for taking analog audio signals, like the kind you hear when you talk on the phone, and turning them into digital data that can be transmitted over the Internet.

How is this useful? VoIP can turn a standard Internet connection into a way to place **free phone calls**. The practical upshot of this is that by using some of the free VoIP software that is available to make Internet phone calls, you are bypassing the phone company (and its charges) entirely.

VoIP is a revolutionary technology that has the potential to completely rework the world's phone systems. VoIP providers like Vonage have already been around for a little while and are growing steadily. Major carriers like AT&T are already setting up VoIP calling plans in several markets around the United States, and the FCC is looking seriously at the potential ramifications of VoIP service.

Above all else, VoIP is basically a clever "reinvention of the wheel." In this article, we'll explore the principles behind VoIP, its applications and the potential of this emerging technology, which will more than likely one day replace the traditional phone system entirely.

The interesting thing about VoIP is that there is not just one way to place a call. There are three different "flavors" of VoIP service in common use today:

- **ATA** - The simplest and most common way is through the use of a device called an ATA (analog telephone adaptor). The ATA allows you to connect a standard phone to your computer or your Internet connection for use with VoIP. The ATA is an analog-to-digital converter. It takes the analog signal from your traditional phone and converts it into digital data for transmission over the Internet. Providers like Vonage and AT&T CallVantage are bundling ATAs free with their service. You simply crack the ATA out of the box, plug the cable from your phone that would normally go in the wall socket into the ATA, and you're ready to make VoIP calls. Some ATAs may ship with additional software that is loaded onto the host computer to configure it; but in any case, it is a very straightforward setup.

- **IP Phones** - These specialized phones look just like normal phones with a handset, cradle and buttons. But instead of having the standard RJ-11 phone connectors, IP phones have an RJ-45 Ethernet connector. IP phones connect directly to your router and have all the hardware and software necessary right onboard to handle the IP call. Wi-Fi phones allow subscribing callers to make VoIP calls from any Wi-Fi hot spot.

- **Computer-to-computer** - This is certainly the easiest way to use VoIP. You don't even have to pay for long-distance calls. There are several companies offering free or very low-cost software that you can use for this type of VoIP. All you need is the software, a microphone, speakers, a sound card and an Internet connection, preferably a fast one like you would get through a cable or DSL modem. Except for your normal monthly ISP fee, there is usually no charge for computer-to-computer calls, no matter the distance.

If you're interested in trying VoIP, then you should check out some of the free VoIP software available on the Internet. You should be able to download and set it up in about three to five minutes. Get a friend to download the software, too, and you can start tinkering with VoIP to get a feel for how it works.

Next, we'll look at exactly how VoIP is used.

### Using VoIP

Chances are good you are already making VoIP calls any time you place a long-distance call. Phone companies use VoIP to streamline their networks. By routing thousands of phone calls through a circuit switch and into an IP gateway, they can seriously reduce the bandwidth they're using for the long haul. Once the call is received by a gateway on the other side of the call, it is decompressed, reassembled and routed to a local circuit switch.

Although it will take some time, you can be sure that eventually all of the current circuit-switched networks will be replaced with **packet-switching technology** (more on packet switching and circuit switching later). IP telephony just makes sense, in terms of both economics and infrastructure requirements. More and more businesses are installing VoIP systems, and the technology will continue to grow in popularity as it makes its way into our homes. Perhaps the biggest draws to VoIP for the home users that are making the switch are **price** and **flexibility**.

With VoIP, you can make a call from anywhere you have broadband connectivity. Since the IP phones or ATAs broadcast their info over the Internet, they can be administered by the provider anywhere there is a connection. So business travelers can take their phones or ATAs with them on trips and always have access to their home phone. Another alternative is the **softphone**. A softphone is client software that loads the VoIP service onto your desktop or laptop. The Vonage softphone has an interface on your screen that looks like a traditional telephone. As long as you have a headset/microphone, you can place calls from your laptop anywhere in the broadband-connected world.

Most VoIP companies are offering minute-rate plans structured like cell phone bills for as little as $30 per month. On the higher end, some offer unlimited plans for $79. With the elimination of unregulated charges and the suite of free features that are included with these plans, it can be quite a savings.

Most VoIP companies provide the features that normal phone companies charge extra for when they are added to your service plan. VoIP includes:

- Caller ID
- Call waiting
- Call transfer
- Repeat dial
- Return call
- Three-way calling

There are also advanced call-filtering options available from some carriers. These features use caller ID information to allow you make a choice about how calls from a particular number are handled. You can:

- Forward the call to a particular number
- Send the call directly to voicemail
- Give the caller a busy signal
- Play a "not-in-service" message
- Send the caller to a funny rejection hotline

With many VoIP services, you can also check voicemail via the Web or attach messages to an e-mail that is sent to your computer or handheld. Not all VoIP services offer all of the features above. Prices and services vary, so if you're interested, it's best to do a little shopping.

Now that we've looked at VoIP in a general sense, let's look more closely at the components that make the system work. In order to understand how VoIP really works and why it's an improvement over the traditional phone system, it helps to first understand how a traditional phone system works.

### VoIP: Circuit Switching and Packet Switching

Existing phone systems are driven by a very reliable but somewhat **inefficient** method for connecting calls called **circuit switching**.

Circuit switching is a very basic concept that has been used by telephone networks for more than 100 years. When a call is made between two parties, the connection is maintained for the duration of the call. Because you are connecting two points in both directions, the connection is called a **circuit**. This is the foundation of the **Public Switched Telephone Network** (PSTN).

**Click "Play" to see how circuit switching works.**

Here's how a typical telephone call works:

- You pick up the receiver and listen for a dial tone. This lets you know that you have a connection to the local office of your telephone carrier.
- You dial the number of the party you wish to talk to.
- The call is routed through the switch at your local carrier to the party you are calling.
- A connection is made between your telephone and the other party's line using several interconnected switches along the way.
- The phone at the other end rings, and someone answers the call.
- The connection opens the circuit.
- You talk for a period of time and then hang up the receiver.
- When you hang up, the circuit is closed, freeing your line and all the lines in between.

Let's say that you talk for 10 minutes. During this time, the circuit is continuously open between the two phones. In the early phone system, up until 1960 or so, every call had to have a dedicated wire stretching from one end of the call to the other for the duration of the call. So if you were in New York and you wanted to call Los Angeles, the switches between New York and Los Angeles would connect pieces of copper wire all the way across the United States. You would use all those pieces of wire just for your call for the full 10 minutes. You paid a lot for the call, because you actually owned a 3,000-mile-long copper wire for 10 minutes.

Telephone conversations over today's traditional phone network are somewhat more efficient and they cost a lot less. Your voice is **digitized**, and your voice along with thousands of others can be combined onto a single fiber optic cable for much of the journey (there's still a dedicated piece of copper wire going into your house, though). These calls are transmitted at a fixed rate of 64 kilobits per second (Kbps) in each direction, for a total transmission rate of 128 Kbps. Since there are 8 kilobits (Kb) in a kilobyte (KB), this translates to a transmission of 16 KB each second the circuit is open, and 960 KB every minute it's open. So in a 10-minute conversation, the total transmission is 9,600 KB, which is roughly equal to 10 megabytes (check out How Bits and Bytes Work to learn about these conversions). If you look at a typical phone conversation, much of this transmitted data is wasted.

While you are talking, the other party is listening, which means that only half of the connection is in use at any given time. Based on that, we can surmise that we could cut the file in half, down to about 4.7 MB, for efficiency. Plus, a significant amount of the time in most conversations is dead air -- for seconds at a time, neither party is talking. If we could remove these silent intervals, the file would be even smaller. Then, instead of sending a continuous stream of bytes (both silent and noisy), what if we sent just the packets of noisy bytes when you created them? That is the basis of a packet-switched phone network, the alternative to circuit switching.

Data networks do not use circuit switching. Your Internet connection would be a lot slower if it maintained a constant connection to the Web page you were viewing at any given time. Instead, data networks simply send and retrieve data as you need it. And, instead of routing the data over a dedicated line, the data packets flow through a chaotic network along thousands of possible paths. This is called **packet switching**.

While circuit switching keeps the connection open and constant, packet switching opens a brief connection -- just long enough to send a small chunk of data, called a packet, from one system to another. It works like this:

- The sending computer chops data into small packets, with an address on each one telling the network devices where to send them.
- Inside of each packet is a **payload**. The payload is a piece of the e-mail, a music file or whatever type of file is being transmitted inside the packet.

- The sending computer sends the packet to a nearby **router** and forgets about it. The nearby router send the packet to another router that is closer to the recipient computer. That router sends the packet along to another, even closer router, and so on.

- When the receiving computer finally gets the packets (which may have all taken completely different paths to get there), it uses instructions contained within the packets to reassemble the data into its original state.

Packet switching is very efficient. It lets the network route the packets along the least congested and cheapest lines. It also frees up the two computers communicating with each other so that they can accept information from other computers, as well.

Next, we'll look at the advantages of using VoIP.

**Advantages of Using VoIP**

VoIP technology uses the Internet's packet-switching capabilities to provide phone service. VoIP has several advantages over circuit switching. For example, packet switching allows several telephone calls to occupy the amount of space occupied by only one in a circuit-switched network. Using PSTN, that 10-minute phone call we talked about earlier consumed 10 full minutes of transmission time at a cost of 128 Kbps. With VoIP, that same call may have occupied only 3.5 minutes of transmission time at a cost of 64 Kbps, leaving another 64 Kbps free for that 3.5 minutes, plus an additional 128 Kbps for the remaining 6.5 minutes. Based on this simple estimate, another three or four calls could easily fit into the space used by a single call under the conventional system. And this example doesn't even factor in the use of data compression, which further reduces the size of each call.

Let's say that you and your friend both have service through a VoIP provider. You both have your analog phones hooked up to the service-provided ATAs. Let's take another look at that typical telephone call, but this time using VoIP over a packet-switched network:

**Click "Play" to see how packet switching works.**

- You pick up the receiver, which sends a signal to the ATA.

- The ATA receives the signal and sends a dial tone. This lets you know that you have a connection to the Internet.

- You dial the phone number of the party you wish to talk to. The tones are converted by the ATA into digital data and temporarily stored.

- The phone number data is sent in the form of a request to your VoIP company's **call processor**. The call processor checks it to ensure that it is in a valid format.

- The call processor determines to whom to map the phone number. In **mapping**, the phone number is translated to an IP address (more on this later). The **soft switch** connects the two

**VoIP Terms**

The central **call processor** is a piece of hardware running a specialized database/mapping program called a **soft switch**. See the "Soft Switches" section to learn more.

devices on either end of the call. On the other end, a signal is sent to your friend's ATA, telling it to ask the connected phone to ring.

- Once your friend picks up the phone, a session is established between your computer and your friend's computer. This means that each system knows to expect packets of data from the other system. In the middle, the normal Internet infrastructure handles the call as if it were e-mail or a Web page. Each system must use the same protocol to communicate. The systems implement two channels, one for each direction, as part of the session.

- You talk for a period of time. During the conversation, your system and your friend's system transmit packets back and forth when there is data to be sent. The ATAs at each end translate these packets as they are received and convert them to the analog audio signal that you hear. Your ATA also keeps the circuit open between itself and your analog phone while it forwards packets to and from the IP host at the other end.

- You finish talking and hang up the receiver.

- When you hang up, the circuit is closed between your phone and the ATA.

- The ATA sends a signal to the soft switch connecting the call, terminating the session. Probably one of the most compelling advantages of packet switching is that data networks already understand the technology. By migrating to this technology, telephone networks immediately gain the ability to communicate the way computers do.

It will still be at least a decade before communications companies can make the full switch over to VoIP. As with all emerging technologies, there are certain hurdles that have to be overcome. We'll look at those in the next section.

**Disadvantages of Using VoIP**

The current Public Switched Telephone Network is a robust and fairly bulletproof system for delivering phone calls. Phones just work, and we've all come to depend on that. On the other hand, computers, e-mail and other related devices are still kind of flaky. Let's face it -- few people really panic when their e-mail goes down for 30 minutes. It's expected from time to time. On the other hand, a half hour of no dial tone can easily send people into a panic. So what the PSTN may lack in efficiency it more than makes up for in reliability. But the network that makes up the Internet is far more complex and therefore functions within a far greater margin of error. What this all adds up to is one of the major flaws in VoIP: **reliability**.

- First of all, **VoIP is dependant on wall power**. Your current phone runs on phantom power that is provided over the line from the central office. Even if your power goes out, your phone (unless it is a cordless) still works. With VoIP, no power means no phone. A stable power source must be created for VoIP.
- Another consideration is that many **other systems in your home may be integrated into the phone line**. Digital video recorders, digital subscription TV services and home security systems all use a standard phone line to do their thing. There is currently no way to integrate these products with VoIP. The related industries are going to have to get together to make this work.
- **Emergency 911 calls** also become a challenge with VoIP. As stated before, VoIP uses IP-addressed phone numbers, not NANP phone numbers. There is no way to associate a geographic location with an IP address. So if the caller can't tell the 911 operator where he or she is located, then there is no way to know which call center to route the emergency call to and which EMS should respond. To fix this, perhaps geographical information could somehow be integrated into the packets.
- Because VoIP uses an Internet connection, it is susceptible to all the **hiccups normally associated with home broadband services**. All of these factors will affect call quality:

  - Latency
  - Jitter
  - Packet loss

  Phone conversations can become distorted, garbled or lost because of transmission errors. Some kind of stability in Internet data transfer needs to be guaranteed before VoIP could truly replace traditional phones.

> **Testing, Testing...**
> Wondering if your broadband connection could support VoIP service? Brix Network offers a way to test your Internet connection to see how well it works.

- VoIP is susceptible to worms, viruses and hacking, although this is very rare and VoIP developers are working on VoIP encryption to counter this.
- Another issue associated with VoIP is having a phone system dependant on individual PCs of varying specifications and power. A call can be affected by **processor drain**. Let's say you are chatting away on your softphone, and you decide to open a program that saps your processor. Quality loss will become immediately evident. In a worst case scenario, your system could crash in the middle of an important call. In VoIP, all phone calls are subject to the limitations of normal computer issues.

One of the hurdles that was overcome some time ago was the **conversion** of the analog audio signal your phone receives into packets of data. How it is that analog audio is turned into packets for VoIP transmission? The answer is **codecs**.

**VoIP: Codecs, Soft Switches and Protocols**

A codec, which stands for **coder-decoder**, converts an audio signal into a compressed digital form for transmission and then back into an uncompressed audio signal for replay. This is the essence of VoIP. Digital-to-analog conversion is seen in everything from CD players to cell phones to video game consoles.

Codecs accomplish the conversion by **sampling** the audio signal several thousand times per second. For instance, a G.711 codec samples the audio 64,000 times a second. It converts each tiny sample into digitized data and compresses it for transmission. When the 64,000 samples are reassembled, the pieces of audio missing between each sample are so small that to the human ear, it sounds like one continuous second of audio signal. There are different sampling rates in VoIP depending on the codec being used:

- 64,000 times per second
- 32,000 times per second
- 8,000 times per second

A G.729A codec has a sampling rate of 8,000 times per second and is the most commonly used codec in VoIP. It is a compromised balance between sound quality and efficiency of bandwidth.

Codecs operate by using advanced algorithms that help them sample, sort, compress and packetize audio data. The **CS-ACELP algorithm** (CS-ACELP = conjugate-structure algebraic-code-excited linear prediction) is one of the most prevalent algorithms in VoIP. CS-ACELP helps to organize and streamline the available bandwidth. **Annex B** is an aspect of CS-ACELP that creates the transmission rule, which basically states "if no one is talking, don't send any data." As we learned before, the efficiency created by this rule is one of the greatest ways in which packet switching is superior to circuit switching. It is Annex B in the CS-ACELP algorithm that is responsible for that aspect of the VoIP call.

So the codec works with the algorithm to convert and sort everything out, but none of that is any good without knowing where to send the data. In VoIP, that task is handled by **soft switches**.

**E.164** is the name given to the standard for the North American Numbering Plan (NANP). Simply stated, this is the numbering system that phone networks use to know where to route a call based on the numbers entered into the phone keypad. In that way, a phone number is like an address:

**(313) 555-1212**

**313** = State
**555** = City
**1212** = Street address

In our example, the switches know to use "313" to route the phone call to the region denoted by the area code. The "555" prefix sends the call to a central office, and the network routes the call using the last four digits, which are associated with a specific location. So based on that system, no matter where you are in the world, the number combination "(313) 555" will always put you in the same central office, which has a switch that knows which phone is associated with "1212."

The challenge with VoIP is that IP-based networks don't read phone numbers based on NANP. They look for IP addresses, which look more like this:

**192.158.10.7**

IP addresses correspond to a particular device on the network. It can be a computer, a router, a switch, a gateway or, in this case, a telephone. To make matters worse, IP addresses are not always static. They are assigned by a DHCP server on the network and generally change with each new connection. So the challenge with VoIP is figuring out a way to translate NANP phone numbers to IP addresses and then finding out the current IP address of the requested number. This is the mapping process referred to earlier and is handled by a central call processor running a soft switch.

The **central call processor** is a piece of hardware running a specialized database/mapping program called a **soft switch**. Think of the user and the phone or computer associated with that user as one package -- man and machine. That package is called the **endpoint**. The soft switch connects endpoints.

Soft switches know:

- Where the endpoint is on the network
- What phone number is associated with that endpoint
- The current IP address assigned to that endpoint

So when a call is placed using VoIP, a request is sent to the soft switch asking which endpoint is associated with the dialed phone number and what that endpoint's current IP address is. The soft switch contains a database of users and phone numbers. If it doesn't have the information it needs, it hands off the request downstream to other soft switches until it finds one that can answer the request. Once it finds the user, it locates the current IP address of the device associated with that user in a similar series of requests. It sends back all the relevant information to the softphone or IP phone, allowing the exchange of data between the two endpoints.

Soft switches work in tandem with the devices on the network to make VoIP possible. In order for all of these devices to work together, they must communicate in the same way. This communication is one of the most important aspects that will have to be refined in order for VoIP to really take off. Currently, there are three protocols used for this communication. In the next section, we will learn about them.

**Protocols**

As we've seen, on each end of a VoIP call we can have any combination of an analog, soft or IP phone as acting as a user interface, ATAs or client software working with a codec to handle the digital-to-analog conversion, and soft switches mapping the calls. So how do you get all of these completely different pieces of hardware and software to communicate efficiently to pull all of this off? The answer is **protocols**.

There are several protocols currently used for VoIP. These protocols define ways in which devices like codecs connect to each other and to the network using VoIP. They also include specifications for audio codecs. The most widely used protocol is **H.323**, a standard created by the International Telecommunication Union (ITU). H.323 is a comprehensive and very complex protocol that was originally designed for **video conferencing**. It provides specifications for real-time, interactive videoconferencing, data sharing and audio applications such as VoIP. Actually a suite of protocols, H.323 incorporates many individual protocols that have been developed for specific applications.

| H.323 Protocol Suite | | | |
|---|---|---|---|
| **Video** | **Audio** | **Data** | **Transport** |
| H.261<br>H.263 | G.711<br>G.722<br>G.723.1<br>G.728<br>G.729 | T.122<br>T.124<br>T.125<br>T.126<br>T.127 | H.225<br>H.235<br>H.245<br>H.450.1<br>H.450.2<br>H.450.3<br>RTP<br>X.224.0 |

As you can see, H.323 is quite a large collection of protocols and specifications. That's what allows it to be used for so many applications. The problem with H.323 is that it is not specifically tailored to VoIP.

An alternative to H.323 emerged with the development of Session Initiation Protocol (SIP). SIP is a much more streamlined protocol, developed specifically for VoIP applications. Smaller and more efficient than H.323, SIP takes advantage of existing protocols to handle certain parts of the process. Media Gateway Control Protocol (MGCP) is a third commonly used VoIP protocol that focuses on endpoint control. MGCP is geared toward features like call waiting. You can learn more about the architecture of these protocols at Protocols.com: Voice Over IP.

One of the challenges facing the worldwide use of VoIP is that these three protocols are not always compatible. VoIP calls going between several networks may run into a snag if they hit conflicting protocols. Since VoIP is a relatively new technology, this compatibility issue will continue to be a problem until a governing body creates a standard universal protocol for VoIP.

The overall hurdle facing VoIP is that there are currently no overriding standards. This includes hardware, protocols and virtually every aspect of the system. In the end, VoIP is a vast improvement over the current phone system in terms of efficiency, cost and flexibility. Like any emerging technology, VoIP has some challenges to overcome, but it is clear that developers will keep refining this technology until it eventually replaces the current phone system.

For more information on VoIP and related topics, check out the links on the next page.

**Lots More Information**

## Related HowStuffWorks Articles

- How Telephones Work
- How Cell Phones Work
- How Wi-Fi Phones Work

- How WiFi Works
- How Municipal WiFi Works
- Skype vs. Vonage
- How WAP Works
- How Web Servers Work
- How Home Networking Works
- How Internet Infrastructure Works
- How Routers Work
- How Speakers Work
- How Sound Cards Work
- How Modems Work
- How Cable Modems Work
- How DSL Works

## More Great Links

- iptel.org: On-line Reference for Internet Telephony
- Cisco.com: Configuring VoIP for the Cisco 3600 Series
- Protocols.com: H.323 Protocol
- VoIP Calculator
- Columbia.edu: Session Initiation Protocol (SIP)

### Sources


- An Interview with Andy Melton of NOS Communications
- PC World: Avoid VoIP Gotchas by James A. Martin - June 2004
- PC World: Net phones evolve - Sept. 2004
- VoIP watch by Andy Abramson
- PC Magazine