Preventing Bad Audio Quality

ISP Connection Congestion Culprit

While it is possible and is on occasion caused by other places along the network path and even the audio servers themselves they are almost neve the cause of bad audio. See the article Troubleshooting Bad Audio for how to identify the cause of bad audio.

Problems with audio quality are almost always caused by congestion on the Internet connection to the Internet Service Provider (ISP). In this article we will discuss why it is such a problem and what can be done about it.
We have learned from experience, research and education how to identify the sources of bad audio.

Do 10 Calls Hog a lot of Bandwidth?

Using the calculator http://www.bandcalc.com/ you will find that using G.711 will use 64Kbs for each channel. So we can calculate that 10 calls will require 640Kbs for both inbound and outbound.

Each call has two channels. One for inbound and one for outbound. So you put 1 channel, not 2 for each call you want to calculate for. The number calculated will be the total bandwidth needed in each direc1ion. http://www.bandcalc.com/ lets you select which codec you want to calculate for.

Considering the bandwidth provided by today's high speed lnternet connections, how much bandwidth is used by the audio of VOIP calls is almost never something the needs to be worried about or a factor in poor audio quality.

Killer Surfing

The reason we need such high speed Internet connection is because of the nature of our use of it. Web surfing. Today's web pages are filled with visual content and every1hing from the browser to the sites providing the content is designed to get that page fully loaded as fast as possible. Even as it is only starting to load the base page html document, a single request that is usually very small, the browser is finding the definitions that tell it what else needs to be requested to fully display the page and it immediately sends the requests for that content.

But Router Graph Shows Not Saturated

As data is sent through the router, it keeps a cumulative count of how much data is passed in each direction. But the graph system polls that counter periodically, almost no router more often than every 5 seconds and most monitoring systems not less than once a minute. The

graph program then takes the difference between the current count and the previous count to show as the data point in the graph. That means that number, point on the graph, is an average of the data that went through between the two times polled. Bursty traffic from web pages is short lived b design. So how long the connection is saturated will likely not be long enough to use the full bandwidth of the connection the entire time between polled times and thus the graph will not show the connection was saturated.

One of my favorite tools to investigate where bad audio cause is PingPlotter.

PingPlotter resources:

Getting started video https://www.youtube.com/watch?v=vTWsRrfJnEY&feature=youtu.be

Understanding and using PingPlotter video
https://www.youtube.com/watch?v=xh0_eMlD1iU&feature=youtu.be

Analyzing the data video https://www.youtube.com/watch?v=3Raj7tX0TUQ&feature=youtu.be

PingPlotter tutorials playlist
https://www.youtube.com/watch?v=3Raj7tX0TUQ&index=8&list=UUVNcmN-8lgJli8kUJXtVPsQ

PingPlotter Manual https://www.pingplotter.com/files/pdf/pingplotter_v5manual.pdf

PingPlotter Download https://www.pingplotter.com/download

What about videos?

Videos generally are not bursty. Most video services do burst at the beginning to buffer ahead so you don't experience pauses in the video at times when there is less available bandwidth than the video needs. But once it is going the data it uses is quite steady. Netflix says for its HD video it use: no more than SMbs. Out of a common business high speed connection that is 25Mbs or more it is a large chunk but only a part of the story.

There are other things like file sharing downloads and uploads that can use steady streams of data that can quickly all put together fill up that connection and make it really easy for web surfing to burst and fill up the remaining bandwidth or even saturate the bandwidth itself.

Do I have to Buy a HUGE connection then?

While buying a HUGE bandwidth internet connection that the bursts will not saturate is the simplest solution the question for that solution is how big is big enough and can you afford it.

All solution will cost money and/or time (money).

The other primary solution is to protect the audio from congestion caused by other traffic.

This is done generally by identifying traffic to/from the VOIP phones then either route that traffic out a secondary ISP connection or use queues that can prioritize the audio traffic or reserve a certain amount of the bandwidth.

QOS

Most VOIP phones, softphones and VOIP services set high priority QOS values on the audio traffic by default. However, it doesn't seem to prevent the audio from getting squashed.

Queues & Throttling

Some routers have pretty good support and tools called Queues to *do* good prioritization and reserve the bandwidth needed for VOIP audio. But most it either doesn't work well or is just too complicated to get to work right. Further because the limiting is done on the router that is after ISP's download connection throttle at best the router can only send "back off request on traffic that needs to slow down to reserve bandwidth for the audio. But some traffic doesn't support back off requests. UDP, which is commonly used by most audio and video services, by its nature is supposed to just send data without expecting responses and is "ok" if some of its data is dropped or "late". That means that the router with its bandwidth "guarantee" for VOIP cannot send back a request for other non-priority UDP traffic to slow down like can be done for TCP traffic. It goes even deeper with ISP equipment queueing data in both directions so you have to set overall queue limits smaller than what the ISP is selling and the max of what you test without queues active.

The good news, with care and experience it is possible to tune a customer's router Queue system to effectively reserve enough bandwidth for VOIP audio to nearly eliminate audio quality issues caused by congestion on the ISP connection. This solution can have a high upfront time investment

Dedicated ISP Connection for VOIP

most trouble free solution is to have an ISP connection that is dedicated to VOIP. That way there is no other traffic for the audio to be congested with. It doesn't have to be a large connection but does have to be reliable.

Like for queues, you still have to identify the VOIP traffic and have a way to route that traffic through the dedicated connection while the rest of internet traffic goes through the other connection.

Identifying VOIP SIP and Audio

Identifying SIP can be simple, just make a rule to match the port that it is set to communicate on. By default SIP willl use port 5060 but we recommend using port 7000 to avoid equipment with bad SIP ALG detecting the SIP and breaking it. Our SRV records define

SIP to use port 7000. Many modems these days are shipped with bad SIP ALG enabled.

The trick though is audio since the port it uses is set up dynamically within a range of ports. There are several approaches that can be used to separate both SIP and audio (RTP) traffic.

1. Use VLANs.
    1. VOIP phones support VLAN tagging and with our Provisioner you can specify what VLAN id the phone should use as well what VLAN id should be used for its second network port that a PC can be plugged into.
    2. Some DHCP servers can be configured to specify what VLAN id should be used in the DHCP offer. Further in some cases they can be configured to group rules/policies dependent on information sent in the DHCP request by the device requesting an IP. For example, you might be able to match on the mac address range, the device name or other DHCP request headers. That way you don't have to add each device manually to the group when you bring a new device to the network.
    3. Disadvantages of VLAN approach
        1. It requires network switches and routers that are capable of managing and passing the VLAN tagged traffic. So all layer 3 AKA managed switches. Which can add up $$. But in the big picture of time/money are often worth the initial money.
        2. Softphones will not be on the VLAN since they are on a computer that you want the bulk of its traffic to be on the data VLAN and it non VOIP traffic VLAN.
2. Use specific IPs, ports and port ranges on devices in the office.
    1. You can easily identify the SIP port used and make a routing rule for it but audio not so simple. Most VOIP devices can be configured to narrow the range of ports that might be used for audio. Our Provisioner exposes this functionality but must be set for each phone.
    2. Disadvantages
        1. On a computer even if you tell the softphone to use a specific range of ports so the router can trigger it's routing and/or queue rules for traffic using those ports, there is nothing on the computer that will prevent other programs from using ports in that range and thus having undesired VOIP traffic down the dedicated VOIP ISP connection. But with a properly sized audio port range that traffic should be negligible. But it could cause problems for the service that is "wrongly" routed if some parts of the service are routed through one ISP connection and some a different ISP connection. A problem that might be very elusive to identify later.
        2. Web phones becoming more prevalent as they are integrated with things like a web based Operator Console pose a special problem since you can't control what ports they use. Further, the often use encryption so the router has no way to tell that the traffic is VOIP.
3. Use the destination IPs of the VOIP service provider.

1. After considering the disadvantages of the other options this may actually be the most simple and comprehensive solution. It allows identifying all VOIP traffic and only VOIP traffic as what should be VOIP traffic for its special handling. Even for a web phone assuming the web phone connection goes the VOIP servers and not the server serving the web phone page.
2. Disadvantage
    1. You have to get updates from the VOIP service provider when there are changes to the IPs used for VOIP and then make the appropriate changes to the rules that identify the VOIP traffic.

4. Use a combination of the above solutions to cover your bases
    1. Disadvantage of potentially Disadvantage of potentially overcomplicating things and creating more work for yourself if you're not really careful.

Choosing a Router

There is a good discussion in the forum Network, Firewalls and Routers .

Good Audio Guarantee (NOT)

Unfortunately, there is no guarantee of good audio with VOIP because you can't control the entire path of the audio from end to end. But effectively separating the VOIP traffic from other data will take care of pretty much all cases of bad audio caused by congestion on the ISP connection. And since that is almost always the source of bad audio you should be sitting pretty good.

Don't guarantee the customer no audio problems. Internet backbone outages and congestion happen and in spite of our best efforts even out servers and network resources have on occasion been overloaded. But we endeavor to learn from each experience to prevent it re-occurring. See Troubleshooting Bad Audio for further info on that topic.