

Using the ClearSight Analyzer to troubleshoot the top five VoIP problems and troubleshoot Streaming Video

With the prevalence of Voice over IP and Streaming Video applications within the enterprise, it is important to understand several of the most common problems encountered during implementation. This paper examines each of these problems and addresses how the ClearSight Analyzer can be used to isolate the cause of the problem.

Top five VoIP problems

The top five problems typically encountered when deploying a Voice over IP application are:

- Phone Connectivity
- Call Won't Connect
- Echo
- Broken Speech/Dropped Calls
- Voice Quality

The severity of these problems range from degraded voice quality to a complete failure. Existing analog and digital phone systems have set a very high standard for voice quality and reliability. To meet these expectations when it comes to VoIP applications, it is important to understand the common problems and how to isolate and resolve them should they arise.

Phone Connectivity

Before any calls can be placed with a Voice over IP telephone, it must first successfully connect to the network and communicate with the call manager. Unlike the traditional analog phones, a VoIP phone needs more than just a 48v signal.


Getting a DHCP Address

The first thing a VoIP phone must do is acquire an IP address. While this can be statically configured, in most cases DHCP is used to acquire the address. If the DHCP process fails, the phone will not be able to continue the connection

process. A number of problems can arise during the DHCP process. The ClearSight Analyzer can be used to troubleshoot all of these.

If the phone does not receive a response from the DHCP server, it will not be able to connect to the call manager. This missing response can be caused by packet loss on the network, misconfiguration of the phone, or an improperly configured DHCP server. By placing the ClearSight Analyzer in various places throughout the network, the exact cause of the missing DHCP response can be determined in realtime.

Placing ClearSight Analyzer is placed between the phone and the switch can be used to determine if the phone is sending out the DHCP Discovery packets. If the phone is not sending out DHCP Discovery packets, there may be a configuration problem on the phone itself.



Example of successful DHCP request

If it is found that the phone is properly sending the DHCP requests, the ClearSight Analyzer can then be placed between the DHCP server and the switch to which it is connected. If the DHCP server is not receiving the DHCP Discovery packets, one of several issues could exist within the network. The first issue is that the IP Helper address is not properly configured on the router attached to the network containing the phone with connectivity problems. The second problem is that the DHCP Discovery packets are being dropped between the phone and the DHCP server. By following the path between the phone and the DHCP server, the ClearSight Analyzer can be used to determine the exact point at which the packets are being lost.



Troubleshooting VoIP and Streaming Video Problems

Getting the right options


As part of the DHCP Offer, a number of options are passed to the phone. These options typically include the VLAN that the phone should use, the IP address of the TFTP server for downloading the configuration, and the IP address of the call manager.

If any one of these options is miss-configured, the phone will not be able to connect successfully. Since each of these options are typically configured on a per DHCP scope basis, one subnet may work well, while another does not.

The ClearSight Analyzer can be used to capture the DHCP Offer that is received by the phone from the DHCP server. By using the build in decoder, the options and their values can be reviewed. This analysis technique can be employed to validate that the phone is receiving the correct configuration values.

Communicating with the TFTP server

Once the phone has obtained an IP address and the DHCP option values, it must then download its configuration from the configuration server. The process is typically accomplished through the use of TFTP. If the phone is not able to successfully establish a connection with the TFTP server and download the configuration, it will not be able to connect to the call manager.



Successful download of phone configuration file

Through the use of the ClearSight Analyzer the download process can be monitored. The screen above shows an example of how the download process should work. If the download were to fail, the conversation diagram could be used to determine which side of the conversation failed.


Communication with the Call Manager

Once an IP address has been obtained and the configuration file downloaded, the phone must register with the call manager. This may be a simple SIP Proxy or an enterprise class

device. However, if the phone is not able to communicate with the call manager, the person using the phone will not be able to make or receive calls.


As with the download server, the address of the call manager is typically passed within the DHCP Offer. If this value is entered incorrectly into DHCP the phone will not be able to communicate with the call manager. Additionally, if the network is not able to route packets from the phone to the call manager, the phone will not be able to connect.

The use of the conversation diagram in the ClearSight Analyzer provides a clear view of the communications between the phone and the call manager. The screen shot below shows an improperly configured SIP phone trying to register with the SIP proxy.



Here we can see that the phone is trying to register with 10.1.101.31. The correct SIP proxy is 10.1.101.30. Once we know that the phone is not properly configured, it becomes easy to fix the problem.

Once the phone has been reconfigured it successfully registers with the SIP proxy, as seen in the screen shot below.



Successful SIP registration


The key to quickly resolving connectivity problems is to be able to view the conversations in a clear and concise manner. The use of the conversation diagrams with the ClearSight Analyzer facilitates this troubleshooting and problem isolation process.

Call won't connect

The first major hurdle has been overcome; the phone has obtained an IP address and is able to communicate with the call manager. The next step is to call another phone within the network. Dial tone is heard, the phone number is dialed, but there is no audio.


This problem is the result of the dialing phone being able to communicate with the call manager, but not being able to communicate with the dialed phone. This communication failure is often due to problems with routing between the calling and called subnets. Firewalls and other filtering devices can also contribute to communication failures.

When the call is setup and working properly, we expect to see RTP traffic in both directions. However, if we only see RTP traffic in one direction, the person at the other end of the call may hear you, but you can't hear them. Below is an example of a successful call setup and RTP traffic displayed using the ClearSight Analyzer's multi-segment ladder view feature.



Successful call setup and RTP traffic

In the following example, 10.0.0.101 can hear 10.0.0.114, but 10.0.0.114 cannot hear 10.0.0.101. We can see that 10.0.0.114 is successfully sending RTP frames to the 10.0.0.101 device, but there is no traffic flowing in the other direction.



Audio in one direction only

Again, the use of the conversation diagram clearly points out the problem with this call. Once it is determined that the traffic is only flowing in one direction, the route between the two phones can be examined to determine the cause of the lost packets.

Echo


Echo is a problem experienced by most network engineers. It is a problem when a person's own voice is heard on the earpiece of the VoIP handset. Almost all VoIP systems experience echo at some point. What makes troubleshooting echo problems very difficult is that it is almost undetectable at a packet level. No matter where the analyzer is placed, echo rarely manifests itself in packet form.

There are two major causes of echo on a VoIP system

- Electrical Echo
- Acoustical Echo

Electrical Echo

This is caused when a VoIP system is bridged to an analog phone system. Data is carried over 2 pairs of twisted pair cabling. A standard analog phone only requires 1 pair of cable to carry the call. When a call is placed to an analog phone (or sometimes even to a digital phone system) at some point a hybrid cable rolling 4 pairs into 2 must be used. At this exact point is where Electrical Echo can become involved. When the signal from the transmit pair crosses over onto the receive pair of the cable and is reinserted into the stream.




Echo Induced by cable after H.323 Gateway

The signal is then received by the handset and reassembled as echo. In extreme cases, these signals can actually be interpreted by the H.323 Gateway as a new packet, and will appear on the network as an actual packet. However, this occurs very rarely, and only where there is extreme crosstalk on the hybrid cable.


Acoustical Echo

This type of echo usually occurs on calls between VoIP handsets. Phone manufacturers build cancelling into the mouthpiece or speaker phone so the received voice is not bounced back into the transmitter. In most cases, this cancelling mechanism is not tuned properly, and the mouthpiece will retransmit what is received on the earpiece. This is especially the case when softphones are used on Laptops. The microphone picks up the PC speakers and acoustically echos the voice back to the call partner. To minimize this effect, use a handset or headset when communicating with VoIP. Or, reconfigure the microphone and speaker settings for proper cancellation.



Acoustical Echo

Using ClearSight, echo can usually be detected using the playback feature. Capture a call where echo is experienced and play the RTP stream back.



Replaying the call using the ClearSight Analyzer

When a call is being replayed pay attention to the direction of echo. Is it only occurring in one direction? What types of phones are involved in this call? (VoIP handsets on both ends, analog on one end, etc...) Are handsets being used on each end, or is one end a softphone using builtin microphone and speakers?

When echo is experienced, this means the other end is inducing the echo. Start troubleshooting at the end where the echo is not being heard. Remember that in calls restricted to a VoIP system, acoustical echo is commonly the cause of any echo problems.

Broken Speech/Dropped Calls

Once a call has been initiated, speech can become broken mid-call and may or may not recover. At times, the call quality may become so degraded that the call drops altogether. There are two common causes of broken speech or dropped calls in VoIP:

- WAN Congestion
- Bad Cables connecting the phones
- Duplex Mismatch between phone and switch

WAN congestion

The Realtime Protocol (RTP) is typically used for sending the voice portion of a VoIP call. These RTP packets are sent on at a specific interval, whether there is any voice traffic to send or not. Take for example the G.711 codec. A RTP packet is sent every 20 milliseconds. When the phone at the opposite end of the connection receives the RTP packets, it must reassemble them back into an audio signal. If packets are missing, the audio signal will contain gaps.


This packet loss can be caused by a number of network problems. One common cause of packet loss is congested WAN links. The RTP protocol does not provide for a backoff mechanism when the link reaches capacity. Instead the phone will continue to send frames at the same rate, regardless of the capability of the link between the phones.

If the output queues on routers between the two phones become full, the router will begin discarding packets. These packet discards will cause protocols such as TCP to backoff and transmit at a lower rate. However, RTP at either phone will continue to transmit at the same rate.

Out of sequence	96
Drop frame	217
Duplicated frame	0
Current jitter (ms)	2.759
Max jitter (ms)	6.850
Mean jitter (ms)	2.904


Dropped RTP frames

In the example above, we can see that for this call 217 RTP packets were dropped. This may have been caused by congestion on the WAN. Using ClearSight, analyze the LAN connection on the edge router looking for unnecessary traffic, or bandwidth hogs. Check for any large file transfers, UDP streams (Streaming radio) or other unusual traffic traversing the WAN link. Determine what traffic can be terminated or firewalled in order to preserve precious WAN bandwidth, providing more resource for voice traffic.



Monitor Traffic just before the WAN Link

In the example above, FTP encompasses a significant amount of the measured traffic. Click the FTP protocol for more details on WHO is involved in this traffic.




Ladder diagram showing FTP transfer between end points

Source and destination stations involved in the FTP transfer are displayed on the right. Is this acceptable traffic for the WAN link? If it is, and voice packets are still being dropped, a more thorough look at the QoS policy on the routers should be performed. Make sure that time-critical UDP-based RTP traffic receives priority forwarding.

Bad cabling

Another cause of broken speech or dropped calls is bad cable. From time to time a bad cable will be used to connect a phone to the switch. A bad cable includes miswires, opens, shorts, and long untwisted sections of cable. In some cases, the phone will link to the network with no problems, receive a DHCP address, and even initiate a call. However, this cable fault will cause FCS errors on some frames that traverse it, requiring these frames to be dropped at the switch or the phone, depending on the direction the traffic is going. Often, these problems occur intermittently, and are not the first thing we think of looking for when troubleshooting.




Bad cable can cause intermittent problems

Duplex Mismatch

A duplex problem can cause huge problems for a VoIP phone. When one side of the Ethernet connection between the phone and switch is half duplex and the other is full, packet loss will be frequent. Again, because VoIP uses RTP, these packets will not be retransmitted between phones. Speech will become broken and the call may eventually drop.

Duplex mismatch errors are never intentional. No one statically configures the phone to operate at half duplex and the switch to operate at full. They are usually a result of statically configuring one side of the connection and leaving the

other side to auto-negotiate its speed and duplex. When a station is set to auto-negotiate, it will communicate its link capabilities to the link partner. For example, it may advertise 10/100Mbps speeds at either full or half duplex. If the other side is configured statically for full duplex operation, this will not be communicated to the link partner. The auto-negotiating side has no way of knowing how the link partner is configured. Ethernet by nature is a half-duplex technology. It is designed to default to half duplex when in doubt. Therefore, in a situation where one side is set to full duplex and the other is autonegotiating, a duplex mismatch will result.



The duplex setting must match between the phone and the switch

Some engineers make it a policy to hard-set all switch ports to full duplex. When a phone is connected, it will auto-negotiate to half duplex. When a frame arrives at the switch destined for the phone, the switch will begin forwarding the frame whether the phone is currently transmitting or not. After all, the switch can tolerate full duplex communication. The phone however, cannot. It will stop transmitting as soon as it receives a frame from the switch. It cannot send and receive at the same time. If this occurs within the first 64 bytes of the transmitted frame, it will be retransmitted. At any point after 64 bytes, the frame will be dropped and will not be retransmitted. This is called a late collision. The switch will consider this event to be an FCS error.

Duplex problems absolutely cripple VoIP communications. Look for dropped frames in ClearSight and check that the switch ports on either end are set for auto-negotiate, as well as the phones.

Voice quality

The quality of the audio received when on a Voice over IP call is based on a number of factors.

These factors include, but are not limited to:

- Packet Loss
- Jitter
- Out of Sequence Packets

Each one of these factors can a less than satisfactory experience for the person making the VoIP call. The ClearSight Analyzer provides a number of statistics that allow us to measure the impact of each of these quality thefts.

Packet loss

This packet loss is caused by a number of network problems. One common cause of packet loss is the over subscription of WAN links. This is described in the Broken Calls section above.

Other causes of packet loss in a call are:

- Routes flapping on routers
- Wireless LAN at the last hop to the phone
- Bad Cabling
- Duplex Mismatch (discussed above)

Jitter

Jitter is the result of a variation in the inter-packet delay for the RTP packets. The G.711 codec sends a RTP packet every 20 milliseconds. In a network with 0 milliseconds of jitter, the RTP packets would arrive at the destination with an interpacket gap of 20 milliseconds. This is not the case in the majority of VoIP implementations.

Queuing delays within routers are a common source of jitter. If the RTP packet must wait for another packet to be transmitted before it can be transmitted, jitter is introduced. The amount of jitter will have a significant impact on the quality of the VoIP call. Most VoIP phones are able handle up to 40ms of jitter before the quality of the voice becomes degraded.


By placing the ClearSight Analyzer between the distant phone and the switch, we can analyze the amount of jitter introduced by the network between the two phones.

Current jitter (ms)	2.759
Max jitter (ms)	6.850
Mean jitter (ms)	2.904

Jitter measurement in the ClearSight Analyzer

In this case we can see that the mean jitter for this call is 2.904 milliseconds and the maximum jitter is 6.850 milliseconds. Both of these values are well within an acceptable range for a good quality VoIP call.


The jitter values can also be viewed by running the VoIP report for the phone call. In addition to providing the mean jitter value for the call, the VoIP report provides a distribution of the jitter.



Jitter Distribution from VoIP report

The jitter distribution provides a means of determining the significance of the jitter. Merely looking at the maximum jitter may make the jitter appear worse than it really is. As well, looking at only the mean jitter may make it appear better than it really is. The jitter distribution gives us a means to determine where most of the jitter is occurring.

To troubleshoot the jitter problem, we must locate the source of the jitter. The first step is to place the ClearSight Analyzer between the distant phone and the switch and measure the far end jitter. The analyzer is then moved progressively closer to the near end phone. Once a significant drop is observed in the mean jitter, the source of the jitter has been located.



Locating the source of jitter using the ClearSight Analyzer


In the example above the phone on the right is experiencing significant jitter problems. The ClearSight Analyzer is first placed at point A. Here we observe the jitter that is being experienced by the phone.

The analyzer is then moved to point B. At point B we measure the same jitter values as at point A. The analyzer is moved to point C. The measurement at point C shows a marked decrease in the jitter. This would imply that router R3 is the source of the jitter. We would want to examine the bandwidth between router R3 and router R4. The implementation of Quality of Service may be appropriate between these two routers.

The jitter on this link may be due to congestion caused by other network traffic. As part of reducing the amount of jitter across the link, the ClearSight Analyzer can be used to monitor the traffic going to and from the router. This information can then be used to determine if there are any “network thefts” present. These “network thefts” are applications that use up available bandwidth, but are not mission critical applications. These include downloads from the Internet, peer to peer communications programs, uploading pictures, and large file transfers. By identifying these “network thefts”, we are able to free up bandwidth required for the VoIP application.

Out of sequence packets

Each RTP packet receives its own sequence number as it is transmitted by the sending phone. As the destination phone receives the RTP packets, it reassembles them based on this sequence number. If a RTP packet is missing, the receiving device must wait to reassemble until the out of order packet, or the timeout period expires. In either case, the voice quality is impacted.



In the example above, we have two paths to the same destination. However, one of these paths has a higher latency than the other. This can result in the packets arriving out of sequence at the destination.

src	10.30
Out of sequence	96
Drop frame	217
Duplicated frame	0

From the RTP statistics screen we can see that 96 out of sequence packets were received. As with troubleshooting jitter, the analyzer is placed in various locations between the sender and receiver to determine where the packets are getting out of sequence.

Troubleshooting Streaming Video

So far this paper has covered the top five problems found in a VoIP implementation. In this section we will address some of the problems encountered when deploying streaming video and the metrics used to assess the health of the video streams.

There are two typical applications for streaming video on today's networks. The first is video conferencing. In the past to setup a video conference required a room full of expensive equipment and an ISDN or T-1 connection. Today with applications such as Microsoft's NetMeeting and Skype a video conference requires little more than a laptop and a \$40 USB video camera. Even though the investment in equipment has decreased over the years, the importance of video quality has not. In a world with HDTV migrating into homes, the expectations of video quality are higher than ever.

The second type of deployment seen is the implementation of Video-on-Demand systems. In this case video can be stored centrally on a server either at the same location as the viewer, or across an international Internet connection. As with the video conferencing, quality is of up most importance to both the viewer and the provider. Video streams add an additional strain to the network links due to the greater bandwidth requirements than those of most VoIP implementations.

As with the VoIP traffic, packet loss and jitter are two of the biggest causes of poor video quality.

Other factors that must be taken into consideration when working with video traffic is the management of the receive buffer on the player end of the connection. This buffer must contain enough video data to constantly provide the player with a constant video stream. It is important not to send so much data to the player that the receive buffer is overflowed. At the same time, too little data will create an underflow situation, causing the video to pause with the buffer is filled up.

Troubleshooting metrics

Packet Loss – This value represents the number of frames lost between the sender and the receiver. Each Real Time Protocol (RTP) data packet is assigned a sequence number by the sending device.

```
Payload type=Dynamic (96), SSRC=23913, Seq=14744,
Payload type=Dynamic (96), SSRC=23913, Seq=14745,
Payload type=Dynamic (96), SSRC=23913, Seq=14746,
```

When the receiving device receives the RTP packets, it tracks the sequence numbers and monitors for missing packets. If the missing packet is not received before the player must display it on the screen, the data contained in that packet is useless. In the example below we can see one of the values passed during the Real Time Streaming Protocol Setup (RTSP) setup process. The late-tolerance value tells the sending device how late a packet can be and still be used. In this case a packet can arrive 400 milliseconds and still be usable.

```
late-tolerance=0.400000\r\n
```

The loss of packets can be caused by a number of problems on the network. One of the most common is a mismatch in the duplex settings between two devices. Another possible cause is the overflow of buffers on networking devices between the sender and receiver. Unlike TCP traffic, the UDP frames containing the video data are sent without waiting for acknowledgements. If the sender begins overflowing buffers along the way, it will not back off.

When troubleshooting the loss of packets, one of the first steps is to inspect the error counters on the networking devices located between the sender and receiver. On a network no packet is ever truly lost. The packet may be corrupted by cabling problems, or aborted by duplex mismatch, but in both of these cases, the error counter on the receiving port will be incremented. If the packet is discarded due to a buffer overflow condition, this too will be seen in the error counters.

If the error counters are not accessible, or do not yield a definitive answer, the next step is to place an analyzer at various locations between the sender and receiver to monitor for packet loss. This can be easily accomplished through the use of the Network Time Machine. The distributed analyzer can be accessed remotely, allowing the analyzers to be placed at key points along the path of the packets and all controlled through a central console.

Jitter

The Jitter value represents a variation in inter-packet arrival time at the receiver. The importance of this value depends on the type of video stream that is being sent from the server to the player. In the case of video on demand, the video data may be sent with a small interframe gap at the beginning of the conversation. This is intended to fill the receive buffer so the player can begin displaying the video as quickly as possible. After the buffer is full, the video packets will be sent at a rate equal to the display rate of the video programming.

In the case of video conferencing, the video cannot be buffered ahead of time. In this case the packets are transmitted with a constant inter-frame gap over the length of the conversation. In this case, we want to see a low and consistent jitter value.

Out of Sequence

Out of sequence frames can be the result of two different situations. In the first situation, a frame is lost and the analyzer detects that there is a missing sequence number between two RTP packets.

The other situation is where packets within the same stream can take more than one path between the sender and the receiver. In this case, it is possible for a frame to be delayed along one path and arrive out of sequence. As long as the out of sequence frame arrives within the late-tolerance value specified by the player, the out of sequence frames will not impact the video quality.

Transaction throughput

The transaction throughput value provides us with an indication of the amount of bandwidth consumed by this video stream.

Transaction Throughput (kbps)	228.619
-------------------------------	---------

This takes into account not only the rate at which the video was encoded, but the bandwidth consumed by the protocol overhead associated with each packet. If this value approaches or exceeds the bandwidth capabilities of the circuit between the server and the player, the resulting video will be degraded.

Interval

The interval represents the gap between each video packet as it is transmitted by the server. As mentioned before, this interval may be constant for video conferencing applications and vary for applications where the video is stored on the server and buffered by the receiving device.

Min interval (ms)	0.015
Max interval (ms)	109.836
Mean interval (ms)	20.380

In this example we can see that the average time between video packets sent by the server is 20 milliseconds. This combined with the average packet size can give us an idea of the anticipated bandwidth required by the stream.

VQFactor

The VQFactor is a ClearSight Networks developed measurement of the video quality. The VQFactor is calculated using many of the measurements listed above. This single number provides a indicator as to the overall quality of the video stream.


Mean VQFactor	3.87
Min VQFactor	3.04
Max VQFactor	4.63

The VQFactor is evaluated on a one to five scale, five being the best score. This value can be used to determine how changes to the network have impacted the overall quality of the video being received by the player. In the case of quality of server, a measurement can be collected prior to the implementation of the QoS rules. After the rules have been configured on the routers, another measurement can be collected. By comparing the before and after results, the impact of the QoS rules can be evaluated to determine whether the desired result was achieved.

☐	VQFactor	3.0	Informational
--------------------------	----------	-----	---------------

By enabling the VQFactor problem threshold under the RTP settings, a drop in VQFactor for a video stream can be used to trigger an alarm. In the example above, the minimum threshold has been set to 3.0. Should the analyzer observe a video stream with a VQFactor of less than 3.0, a problem diagnosis will be displayed.

This threshold can also be used to trigger an alarm based on the VQFactor value. This alarm can then be used to stop the capture at the time the threshold is exceeded.



After the capture has been stopped, the Jitter, Out of Sequence, and Lost Packet values can be reviewed to determine why the VQFactor dropped. Based on these values, the appropriate action can be taken to ensure the problem is resolved before further video streams are impacted.

MDI DF

The Media Deliver Index is comprised of two components, the first of which is the Delay Factor. The Delay Factor is used to compare the number of bytes received by player over a given interval and the number of bytes drained from the player's receive buffer over a given interval. This metric is used to evaluate the potential for a buffer overflow or buffer underflow condition.

This calculation used to arrive at the Delay Factor (DF) is as follows:

$$X = |\text{Bytes Received} - \text{Bytes Drained}|$$

$$DF = [\max(X) - \min(X)] / \text{media rate}$$

A high DF value can indicate that the data is being received by the player at rate greater than the player can drain the data from the receive buffer. Should the receive buffer be exceeded, the excess packets will be discarded, resulting in degraded video quality. This can be an indication that the network latency between the server and the player is too low.

The MDI DF can be used to determine the depth of the player's receive buffer. Below are two examples of DF values from two different video streams.

In the first example we captured a video stream that was encoded at 100 kilobits per second, sent over a Wide Area Network. We can see that the Mean DF is 938 milliseconds. This tells us that on an average for the stream a receive buffer depth of 1 second would be enough to buffer most of the traffic. However, to ensure all of the traffic is buffered without overflow, the buffer should be no less than 2.6 seconds.

Mean MDI DF(ms)	938
Max MDI DF(ms)	2599
Min MDI DF(ms)	271

In the next example, a video stream was captured that was encoded at 220 kilobits per second on a Local Area Network. In this case we see a larger buffer would be required on the player. The higher Delay Factor indicates that the data is being delivered to player at a rate greater than the player is able to drain the data from the receive buffer.

Mean MDI DF(ms)	2603
Max MDI DF(ms)	5665
Min MDI DF(ms)	1675

In this case, the receive buffer on the player should be able to hold up to 5.7 seconds of video to prevent buffer overflow.

A low DF would indicate that the player is draining the buffer at a rate equal to the rate at which the data is arriving. In a network with varying rates of Jitter, a underflow condition may occur where the player must pause while waiting for data to arrive.

MDI MLR

The Media Loss Rate (MLR) is a measurement of the number of packets lost or out of sequence over the measurement interval. It is important to include the out of sequence packets in this calculation, since not all devices will attempt to reorder the packets as they arrive. The DSL Form's WT-126 recommends the following limits on the MLR value:

Codec	Maximum Average MLR
SDTV	0.004
VOD	0.004
HDTV	0.0005

Conclusion:

There are many factors that will determine the quality of the video displayed by the player. Some of these factors such as Jitter can be compensated for by the proper sizing of the receive buffer. Others, such as packet loss will have a significant impact on the quality of the video.

By carefully monitoring the video streams both prior to deployment in a production environment and after deployment thresholds can be determined to ensure the video quality meets or exceeds to expectations of the viewer. By monitoring the video stream at various locations along the path between the server and the player, sources of impairment can be identified and resolved.

As with VoIP troubleshooting, the playback feature of the ClearSight analyzer can be used to determine the quality of the video at each of these points within the network.



205 Westwood Ave
Long Branch, NJ 07740
1-877-742-TEST (8378)
Fax: (732) 222-7088
salesteam@Tequipment.NET