

Ohio K-12 Network – Qos recommendations for IP video conferencing**June 30, 2007****Overview**

The Ohio K-12 network (the network) supports one of the largest interactive video networks in the country, and currently consists of approximately 1,000 endpoints statewide. The current standard for interactive video on the network is H.321, or video over ATM. However, starting with the 2007-2008 school year, the network will begin supporting H.323, or video over IP on a statewide basis.

The advantages to moving to an IP network are numerous, and include, but are not limited to the following:

- Generally less expensive endpoints
- Greater mobility
- Increased connectivity without the reliance of physical paths
- More widely available technical support resources

Along with the many advantages that IP video can provide, there are also some inherent disadvantages to the technology, which, if not addressed, can result in an unsatisfactory video experience for everyone involved. Most notable among these disadvantages is the issue of quality of service (or Qos) associated with IP networks. For the purpose of this document, Qos is defined as the ability of the network to “guarantee” a predetermined level of service necessary to support real-time audio and video signals without noticeable degradation.

Traditional IP networks are intended primarily for data applications that are much more tolerant of impairments such as packet loss and latency, both of which can adversely effect the quality of a video call. Unlike data applications, where, if a packet is dropped or it takes a couple of extra minutes to arrive at its destination, it has little or no effect on the users, video conferencing is a real-time application that needs certain levels of guaranteed quality in order to be of use to the end users.

There are roughly five factors involving the network which can affect the quality of IP video conferencing.

1. Bandwidth - typical video calls today operate at 384k – 512k per call. When you add in the overhead necessary to complete and to carry the call (roughly 20%), you must account for approximately 460k – 620k per call. If your network consists of T1 lines, you can see that an interactive video call can consume a large portion of your bandwidth. In situations like this, the presence of Qos is critical.

2. Packet loss – packet loss results when packets arrive late, incorrectly, or not at all. This is often the result of congestion, or other errors on the network. Packet loss will usually be noticed as “tiling” on the end users video screen.
3. Latency – latency (also referred to as delay) is the amount of time it takes a packet to traverse the network. Latency levels above approximately 150 ms will normally result in unfavorable video, and will often cause people to “talk over each other” during conferences.
4. Jitter – is described as the variance in the expected interval of packets, with some packets arriving faster or slower than other packets. This can cause packets to arrive out of order, and if the jitter values are too high, the codecs will not be able to buffer enough traffic to reassemble them properly.
5. Policies - policy issues are usually introduced by firewalls, network caching, and/or network address translation devices

Qos Implementation

In most cases, it's safe to assume that the bandwidth in the local area network (LAN) exceeds that of the interface leading to the wide area network (WAN). In these cases, it is necessary to implement some level of Qos in order to insure that higher priority traffic such as video and voice packets, are processed first, while lower priority traffic such as email, and other office applications are handled accordingly.

There are two primary mechanisms for insuring that IP video traffic will receive the service level guarantees within your network that it needs to conduct successful video conferences. The first is to identify packet classifications for different traffic types, and the second is to assign queuing priority to these classes.

Identifying packet classes

Within the IP header, there is an eight bit field known as the “type of service” (TOS) identifier. As its name implies, this field can be used to identify traffic types. Earlier on, IP precedence was the standard mechanism for classifying types of traffic. It used the first three bits of this field to establish eight different class types (0-7 respectively), with zero being the least critical and seven being the most highly critical. Today, differentiated services (or diffserv) has become the most widely used mechanism for identifying class types within the IP header.

Like IP precedence, diffserv uses the type of service field in the header, however it uses the first six bits of this field instead of the three bits used by IP precedence, thus providing 64 different classes. Diffserv introduces the concept of differentiated services

code points (or DSCP). Diffserv was designed to recognize, and be backwards compatible with earlier implementations of IP precedence.

Hardware devices in the network, such as routers and switches, can choose to ignore, use, or change the contents of the TOS field. Technical personnel at the site where the video endpoints operate, as well as any intermediate sites along the path (hub locations, or ITC's) must insure that their routers and switches are configured to "use" the TOS field for their video applications.

If your devices will only support IP precedence, then your video traffic should be classified as: IP precedence level (5), critical, hex value 101.

If you are implementing Diffserv, there are several code points that could be used for your video traffic. Although the "EF" (or expedited forwarding) class provides the highest priority, it may not be the best choice for your video traffic. This is because these particular queues are normally the smallest, and if too much traffic is sent to these queues, it may result in lost packets. Therefore, you should classify the video traffic as follows: set your DSCP value to (AF41), the hex value will be 100010, and the decimal equivalent will be (34).

If you are experiencing problems with getting calls established, due to signaling failures, it may be necessary to classify the signaling traffic (H.225 and H245) with a higher than "best effort" class such as CS3, hex value 011000, decimal value (24).

In order to achieve the highest level of traffic classification, it is recommended that network administrators configure their hardware devices to recognize and use the diffserv code points. On devices that are not capable of this, then IP precedence should be used for this purpose.

Establishing and implementing queues

Once the traffic has been classified into its appropriate type, the varying traffic types can be assigned to different queues based on the requirements of each class. Queuing allows you to assign different "weights" based on their requirements, to differing classes of traffic. Users can also go so far, as to assign certain amounts of bandwidth to different traffic classes. This is especially helpful on links with lower bandwidth capacity, as this will prevent higher priority traffic from consuming all of the WAN link bandwidth.

There are multiple methods and levels of queuing, however, for video conferencing, there are two preferred schemas known as weighted fair queuing (WFQ), and class based weighted fair queuing (CBWFQ).

Weighted fair queuing uses the IP precedence value as well as the number of active queues to assign weights to differing types of traffic. Once this is accomplished, it places the traffic with higher priority at the front of the line for processing. Based on the weighting system, the highest priority traffic will always be processed first. The one

drawback to the WFQ method is that it tries to insert smaller packets between the larger packets, and so if there a lot of different queues all active at the same time, WFQ may assign a higher value to the smaller packet size, even though it is lower priority traffic. However, in most cases WFQ will work for video applications.

Class based weighted fair queuing operates very similar to WFQ, however it allows for greater flexibility in assigning the appropriate weights. CBWFQ enables network engineers to assign bandwidth limits, and to establish drop policies for different traffic types. CBWFQ offers the most stringent guarantee of traffic delivery. However in some cases, especially where bandwidth is not overly congested, this may prove to be an unnecessary burden.

In most cases, weighted fair queuing will be satisfactory for successfully assigning video traffic to the proper queue. However, in locations where bandwidth is limited, such as those using a single T1, it may become necessary to limit the amount of bandwidth allowed for video. In these cases, class based weighted fair queuing can be implemented.

It's important to note, that in order to successfully implement these recommendations for classifying and establishing the proper queues for video conferencing, all of the hardware elements along the path (switches and routers) must be configured to be aware of these priorities. This includes hardware at the end user site, the next hop (be it a hub location or ITC), and the core of the network. Therefore it is critical that network administrators are involved in implementing any changes to the existing network.

Firewalls and NAT traversal

Another major concern when implementing IP video across the wide area network is how to deal with firewall and network address translation (NAT) devices. While both devices are needed, and are routinely implemented in networks today, they both can cause difficulties when trying to establish video connections with endpoints outside of your internal network. In most cases, firewall devices will have the NAT capabilities as a part of their platform, and as such, these two functions will be discussed jointly here.

Firewalls

The purpose of a firewall is to prevent unwanted, or untrusted access from a device on an external network, into your internal or trusted environment. By their nature, firewalls are not very "trusting" devices, and therefore are designed to keep any traffic, and/or certain traffic types from entering your network.

When establishing a video call, there are a number of signaling and capabilities exchange messages that go back and forth between endpoints prior to the call being established. These messages use a wide range of ports to pass these messages back and forth. Typically these ports are blocked by firewalls, and if a network administrator were to

open all of the possible ports necessary for these messages, the firewall would quickly become useless. In a typical scenario, video endpoint "A", which is behind a firewall, would be able to originate and establish a call to endpoint "B" on an external network (if it were not behind a firewall), however endpoint "B" would not be able to originate a call back to endpoint "A" because the firewall on endpoint "A's" network would reject the incoming signaling messages.

There are several methods for addressing this issue, and those will be discussed briefly here. These methods are outlined below:

- Placing video endpoints outside the network firewall
- Protocol aware firewalls
- Firewall transversal devices

Placing video endpoints outside the network firewall

The simplest, and obviously the least expensive method for avoiding this issue, is to place your video endpoints outside of your firewall/NAT device, and to provide them with a public IP addresses. While this approach seems fairly straight forward, it is not without its disadvantages. Primary among these, is network security. By placing the endpoint outside of your firewall, you are creating a potential security risk to your network. While this might not present a major problem for hardware devices that are designed solely as video endpoints, more and more video endpoints are being incorporated into desktop, or even laptop PC's, as software based codecs, and therefore, this is not a recommended long-term solution.

[In the event that you choose to implement this solution, it may be beneficial to create a separate VLAN for your video traffic. You then have the option of assigning a higher priority level to this particular VLAN, thus providing this traffic with some level of priority]

Protocol aware firewalls

Another method for addressing this issue, is to implement protocol aware firewalls into the network. These devices are sometimes referred to as application level gateways. As there name implies, these devices actually inspect the packets, and are able to determine information about the IP addresses, port assignment, and the actual protocol of these packets, and can do so for incoming and outgoing messages.

One concern to be aware of when choosing to implement the protocol aware firewall scenario, is that the various manufacturers may not be totally interoperable with each other, and may not support all the functionality of your particular device, especially as gatekeepers are introduced into the network. Another potential drawback to this solution, is that by inspecting all of the incoming and outgoing packets to this level, could result in undesirable delay depending on the amount of memory and the level of CPU in

the hardware device. This could adversely affect not only your video traffic, but also your data traffic. For these reasons, it is recommended that careful consideration be given when implementing this solution.

Firewall / NAT transversal

The most current solution for dealing with the firewall/NAT transversal issue is addressed by implementing and using the recent ITU H.460 standard. There are two elements of the H.460 standard that deal with the signaling (H.460.17 & H.460.18), and one for the actual media channel (H.460.19). Therefore, when implementing this solution, users will use a combination of either H.460.17 or H.460.18, along with H.460.19.

The primary difference between the H.460.17 and H.460.18 signaling standards deals with how the call is set up initially and the number of “pinholes” opened in the firewall/NAT device. The H.460.17 signaling option uses tunneling and multiplexing of the call set up elements (RAS, Q.931, etc..) to create a single TCP connection that is always open during the length of the call. The H.460.18 signaling standard does not use tunneling during call set up, and instead tries to mimic the H.323 standard while constantly looking for pinholes during the call set up procedure. The most recent reports regarding the use of the H.460 protocol indicate that the H.460.18 signaling method is becoming the de-facto standard. The H.460.19 media standard establishes solutions for opening the necessary real time transport protocol (RTP) and the real time transport control protocol (RTCP) pinholes, and for maintaining these pinholes through the use of keep alive messages.

Using the H.460 standards is the preferred, and safest way of conducting video calls in the current environment, as this does not require network administrators to modify their security measures introduced by the firewall and NAT devices. The one drawback to implementing this solution, is that the device behind the firewall/NAT device (normally the video endpoints) must be capable of initiating the H.460 elements. While this is not a problem for many of the newer, IP only video endpoints, the vast majority of the video endpoints deployed in the current network are older, ATM/IP codecs and may not be capable of initiating the necessary H.460 elements. In this case, an external device would be needed to be placed in the network, behind the firewall to accommodate this solution.

Summary

Video conferencing is an integral part of the Ohio K-12 network, and the education community in general. And as such, wholesale changes to the network must be made cautiously, and must involve a great deal of communication amongst those who are responsible for the operation and security of the network. Implementing and maintaining

IP video across a statewide network is a complex and fluid process, and as such, continual changes will be necessary over time as new technologies evolve.

It is imperative that some form of Qos be implemented across the entire statewide network in order to achieve successful IP video connections. This first set of recommendations is designed to provide network administrators with a basic blueprint that will help to insure successful video conferencing across the Ohio K-12 network.

From a Qos standpoint, it is recommended that network administrators configure their hardware devices to recognize and use diffserv wherever possible to identify traffic types, and to assign their video traffic with a higher priority than normal data traffic. In addition to this, it is recommended that all video traffic be sent to a higher queue within the overall traffic flow through the use of weighted fair queuing, or where bandwidth is extremely tight, using class based weighted fair queuing.

Due to the large number of legacy codecs in the network, the issue of firewalls and NAT transversal is not quite as straight forward. The recommended solution for dealing with this issue on a statewide basis is to implement the H.460 standards in order to maintain the security and integrity of the entire network. However, many of the legacy codecs are not capable of initiating the proper H.460 signaling elements. Therefore a secondary device behind the firewall would be required to enable this functionality. Hardware devices from Polycom, Radvision, and Tandberg to name a few are available for this purpose. Most of the newer, IP only codecs in the network are capable of being configured to use IP precedence and/or diffserv, and are capable of initiating the proper firewall/NAT transversal signaling elements.

In the current Ohio K-12 network environment, multipoint conferences are initiated from the core video bridges to each of the participating end points. Therefore, it is important to note that the video bridges will NOT be able to place calls to any video endpoint that is behind a NAT device. In this case, the video endpoints will be required to initiate the call into the video bridge. For this reason, we do not recommend that video endpoints use addresses that are part of a NAT pool, and it is recommended that all of the video endpoints be assigned static, public IP addresses.

Lastly, it is important that the network administrators work closely with their end users, and with the eTech technical staff to successfully implement these recommendations.

Disclaimer: Any mention of a particular manufacturer in this document was intended strictly as a reference. eTech is not making any recommendations as to a particular manufacturer's equipment.