Published on *Jisc community* (https://community-stg.jisc.ac.uk)

Videoconferencing Traffic: Network Requirements

This chapter aims to describe in some detail the demands that videoconferencing traffic places on the network, along with the metrics that can be used to predict – to a certain extent – the behaviour of a videoconference. Readers who may be less interested in the specifics of 'Why is this important?', and would like to move swiftly on to 'What should I look at and do?' are welcome to skip this section completely and move along to chapter 3.

Videoconferencing is about interaction, a two way exchange of information. Video streaming is a related technology to consider, but is generally used for one-way transmission rather than interaction. In order to interact effectively, participants in a videoconference need to be able to communicate in real-time, or as close to real-time as possible – much as we do every day using the telephone.

Significant processing is done within the videoconferencing endpoints, to reduce the latency of processes that compress the raw audio and video into a data stream which can be sent across the network. However, there is a limit as to how much the endpoints can deal with traffic/packet loss or delays across the intervening networks – hence the requirement to provide some protection to the traffic flowing between the endpoints.

The usual metrics, or measurements, used to determine how a network is performing are actually very few. They are:

- latency (end-to-end delay)
- packet loss
- IPDV (Inter Packet Delay Variation) / Jitter
- bandwidth

Latency

Usually, in the context of networking the term 'latency' is only used to define the network delay for IP packets between points A and B. However there are other latencies inherent in videoconferencing, which give rise to the need to be more careful with network latency than might otherwise be expected.
The end-to-end latency experienced by the video and audio through a videoconference, i.e. from source camera and microphone to destination screen and speaker, is the sum of a number of elements:

- the coding delay to compress the audio and video data
- the intervening delay across the network between the NICs (Network Interface Cards) at each of the communicating endpoints

- the decoding delay to decompress the video and audio.

With the CODECs (COders and DECoders) currently used in videoconferencing endpoints, a latency in the region of 100ms can be expected in both the coding and decoding processes. As a general rule of thumb, 300ms is a reasonable maximum target to aim at to get audio and video through the endpoints and network, and still allow relatively unhindered communication, without the satellite-type long delays encountered occasionally with TV interviews between the UK and USA. So allowing for the CODEC delays, there is still some 100ms to get the data across the network. The shorter the overall delay, the better the experience of the participants in a videoconference, so it is beneficial to reduce this figure to the minimum possible.

Packet Loss

As with all forms of compressed data, data loss becomes far more critical with respect to the original data. Losing a few packets from a compressed stream has a far higher impact on the integrity of that stream than if the raw data had been transmitted, as it makes the decompression far more likely to contain errors across a wider range of the received data.

The general network solution for reliable delivery of data lies in the use of TCP (Transmission Control Protocol) connections. TCP is a ‘reliable’ delivery protocol and traffic is ‘guaranteed’ to be delivered between the applications of communicating devices. The TCP protocol includes error checking, re-transmission and back-off under congestion to ensure reliable delivery.

TCP is an excellent delivery mechanism, and it guarantees that http requests and e-mail do get to their destination. However, reliability comes with a price – latency. In order for reliable delivery to take place, handshaking occurs between communicating devices to ensure that all data that has been sent is received. This process introduces peaks and troughs into the data stream as each end waits for the other to confirm what has, or has not, been delivered.

For videoconferencing, the delay and ‘burstiness’ of the reliable TCP transport mechanism is unacceptable. Delivery of real-time traffic across networks has therefore tended towards the use of UDP (User Datagram Protocol), which is an ‘unreliable’ delivery protocol. The network will then simply send and receive packets without considering what, if anything, has been delivered. It is then up to the application to decide whether any processing is required to counter the effects of missing or delayed packets. This means that latency through the network layers is reduced to a minimum, though obviously there is a trade-off in that packets may, or
may not, be delivered.

IPDV (Inter Packet Delay Variation) / Jitter

In a well-behaved, uncongested network, packets will be delivered fairly regularly with little difference in the gap between each (as long as the sending application is sending packets smoothly). However, in a network that is experiencing congestion, the delay between packets will vary, in some cases quite considerably. Usually, the IPDV of packets across a network will be measured in milliseconds; it should be noted though that when things go wrong, the IPDV can be measured in seconds.
Across the JANET network IPDV will normally be well below 20ms with only a few ms being

the norm.

Bandwidth

In many texts there is little discussion of bandwidth requirements. To a certain extent this is because a lack of bandwidth will simply be translated into the above metrics of packet loss, latency and jitter. It should, however, be noted that videoconferencing data streams, whilst having a nominal bandwidth of 768kbit/s or whatever the conference bandwidth is set to, can vary considerably around this if measured on a sub-second basis.

Data gathered some time ago by Cisco Systems, Inc. (Cisco, 2001) suggested that, whilst the average bandwidth used over a period in a 384kbit/s videoconference will be 384kbit/s or lower, the peaks of bandwidth, when measured at shorter time periods, are likely to exceed this significantly as key (whole picture) frames are transmitted. These peaks in the Cisco® data were shown to reach almost 600kbit/s for a 384kbit/s videoconference. This is far in excess of figures traditionally given by suppliers, who advise that conferences should be allowed 10% headroom to be successful. This 10% figure has recently been discounted by Cisco® as out of date and is no longer applicable, but it still appears in some suppliers' presentations.

As shown in Figure 1 (opposite), a Testbed was set up to gather data for analysis in order to examine more closely the traffic profile in current videoconferencing equipment. This has shown that while the data can certainly exceed the nominal bandwidth by more than 10%, it has not been seen to reach the bandwidth levels of the Cisco® data.

Table 1 (above) shows the summary data for a 20 second slice of a 768kbit/s videoconference.
The conference is hosted on the JVCS Leeds MCU, so is representative of standard conferences that take place across JANET. The data is gathered next to the endpoint at UWS (University of Wales Swansea), so data from the Testbed is measuring the output from the videoconferencing endpoint at the first network device it encounters, whereas data to the testbed is measuring the effect of the intervening network as well as the MCU on received traffic.

Figure 1: Testbed equipment and topology.
Measurement From Testbed To Testbed
No of packets in 20 seconds 2330 Packets 1843 Packets
Sum of packet size 1,818,977 Bytes 1,769,746 Bytes
Packet size average 781 Bytes 960 Bytes

Average bandwidth 727,591 bit/s 707,898 bit/s

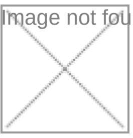

Figure 1: Testbed equipment and topology.

Measurement	**From testbed**	**T**
No of packets on 20 seconds	2330 packets	1
Sum of packet size	1,818,977 Bytes	1

Packet size average	781 Bytes	9
Avergage bandwidth	727, 951 bit/s	7

Testbed is measuring the effect of the intervening network as well as the MCU on received traffic.

The total data transferred in each direction is very close, 1.81 and 1.76 Mbytes, but as can be seen from the packet count and average of packet size, the method of sending the data is different. The question raised is, how different? Figures 2 and 3 (overleaf) show the data for one second of the same conference.The data in Figures 2 and 3 shows a different pattern of packet delivery onto the network, with data from Leeds MCU to the Testbed showing fewer deep, low bandwidth troughs. Figures 4 and 5 (below) show the same data, but summarized by bandwidth use in 0.1 second blocks.

Table 1: 768kbit/s videoconference summary data

Figure 2: One second data from Leeds MCU to Testbed.

In these graphs the bandwidth usage can be seen to vary considerably. The bandwidth from Leeds MCU to the Testbed peaks at 804kbit/s, well within a 10% allowance. However, traffic from the Testbed to Leeds MCU peaks at 877kbit/s which is some 15% higher than the 768kbit/s nominal bandwidth.

Care must therefore be taken when specifying nominal versus peak bandwidth requirements for videoconferencing traffic, and allowance should be made at those points in the network where bandwidth may be limited – for example by QoS policing – to ensure that the peak traffic level can be processed correctly.

Summary

In order to support reliable videoconferencing, the network needs to be provisioned with the above issues and metrics in mind. Latency should be kept to an absolute minimum, and as traffic is sent as UDP datagrams, the network should do its best to deliver those datagrams with the minimum packet loss possible.

In order of impact on the integrity of a videoconferencing stream, packet loss has the most significant effect with half a percent loss causing picture break-up, pixelation or blocking. High packet latency leads to more difficult interaction for conference participants, but will rarely cause significant problems. In some cases, of course, latency is naturally high, as occurs with a videoconference which involves sites spread across the US, Europe and Asia. The rest of this document will aim to describe some methods that may be employed to protect traffic and improve these metrics to the point that videoconferencing can reliably and consistently take place. These will include taking a close look at the existing network, using physical link separation to dedicated equipment, and implementing Layer 2 VLANs and Quality of Service.

Source URL: https://community-stg.jisc.ac.uk/library/videoconferencing-booking-service/videoconferencing-traffic-network-requirements