



VTAS Guide to Reliable Campus H.323 Networks

Steve Williams

Technical Guide

GD / VTAS / 010

V.1.1

This document was commissioned by VTAS as a result of a conversation with Steve Williams. It was borne out of a desire to ensure that expertise, developed by the WVN Support Centre in the deployment of IP (H.323) videoconferencing, is made available to the wider JANET community. The other driver has been Steve's desire to ensure that high quality documentation is widely disseminated. We hope you find this document, which was made possible by VTAS funding, useful.

UKERNA Technical Guides

UKERNA Technical Guides are one of a series of user guides available to JANET customers. Technical Guides contain detailed technical information and are intended for technical support staff at JANET sites, network specialists or those with a particular interest in the specialist area.

If you have any queries or comments about the Guide or would like to obtain copies, please contact:

JANET Customer Service

UKERNA

Atlas Centre, Chilton, Didcot

Oxfordshire, OX11 0QS

Tel: 0870 850 2212

Fax: 0870 850 2213

E-mail: service@janet.ac.uk

Further details of the documents in this series are available at:

<http://www.ja.net/documents/>

Contents

Table of Figures	4
Acknowledgements	5
1 Introduction and Scope of Document	6
1.1 A Word of Caution	6
1.2 Chapter Outlines	6
2 Videoconferencing Traffic: Network Requirements	8
2.1 Latency	8
2.2 Packet Loss	8
2.3 IPDV (Inter Packet Delay Variation) / Jitter	9
2.4 Bandwidth	9
2.5 Summary	13
3 Campus Traffic and Common Network Issues	14
3.1 Layer 1 - The Physical Layer	14
3.2 Layer 2 and 3 Issues	14
3.2.1 Speed and Duplex Settings	14
3.2.2 Spanning Tree Protocol Updates	15
3.2.3 Broadcast Traffic	15
3.2.4 H.323 Firewalls and Network Address Translation	16
3.2.5 General Traffic Profile	16
4 Network Engineering - Physical Separation of H.323 Traffic	18
5 Traffic Engineering 1: Logical Separation at Layer 2 - the VLAN	20
5.1 Inter-VLAN Routing	24
6 Traffic Engineering 2: Layer 2 Prioritisation - CoS (Class of Service)	26
6.1 Classification and Marking	26
6.2 A Note on Marking	29
6.3 Policing	29
6.4 Queuing	29
6.5 Configuration Samples	30
6.5.1 Applying QoS on a Cisco® 3524 Switch	30
6.5.2 Applying QoS on a Cisco® 2950 Switch	31
6.5.3 Applying QoS on a 3Com® 4400 Switch	32
6.6 Summary	33
7 Case Study: The QoS H.323 Network at University of Wales, Aberystwyth	34
7.1 Description	34
7.2 Videoconferencing Endpoints Connected over Trunked VLANs	36
8 List of References	39
9 Further and Recommended Reading	40

Table of Figures

Figure 1: Testbed equipment and topology	10
Figure 2: One second data from Leeds MCU to Testbed	11
Figure 3: One second data from Testbed to Leeds MCU	11
Figure 4: One second bandwidth from Leeds MCU to the Testbed	12
Figure 5: One second bandwidth from the Testbed to Leeds MCU	12
Figure 6: Physical separation	18
Figure 7: Dual-homed gatekeeper/proxy	19
Figure 8: Simple VLANs	21
Figure 9: Trunked VLANs	23
Figure 10: Directly connected videoconferencing endpoints at UWA	35
Figure 11: UWA H.323 QoS trust network	38

Acknowledgements

Geoff Constable, Welsh Video Network Support Officer at the UWA (University of Wales, Aberystwyth), contributed substantial input to the document, especially relating to Layer 2 QoS (Quality of Service) provisioning and the UWA Case Study. Philip Davison and Richard James at the UWS (University of Wales Swansea) were involved in building and operating the Testbed as well as being part of the numerous discussions about the pros and cons of various technologies.

Thanks go to Hefin James, Network Manager and Ian Jones, Welsh Video Network Support Officer at UWA and to Paul Matthews, Network Manager at UWS for providing advice and knowledge as well as data, testing facilities and production network tweaks.

The JANET Video Technology Advisory Service commissioned this document, so thanks are due to UKERNA for their commitment to providing advice and support to the academic community.

We would like to give our heartiest thanks to those who participated in writing this document – whether they had a small part, a large part or, more critically, provided the tea, coffee and biscuits.

Any errors remaining, or omissions from the document, are, of course, solely my responsibility.

Steve Williams

University of Wales Swansea

s.r.williams@swansea.ac.uk

1 Introduction and Scope of Document

This guide is aimed at network engineers and technicians, primarily in educational organisations, who need to provide a network capable of handling videoconferencing traffic.

It aims to inform the reader about the ‘special’ requirements relating to real-time voice and video traffic, as opposed to http, ftp or other traffic types. It then discusses techniques that can be applied to provide a network that can be made to carry such traffic reliably without the need for continual network changes or upgrades.

It will also, therefore, provide information that can help in determining the likely causes of problems experienced with real-time traffic.

This guide will consider the types of equipment and topologies that are frequently encountered in educational environments. It will primarily look at switched campus topologies, but with reference to the necessary routing that will take place, and will assume in the main 100Mbit/s or higher core and 10/100Mbit/s edge links. Issues related to running non-switched, i.e. repeated, networks will be briefly addressed, but as real-time traffic and non-switched networks rarely go well together, it will be only a plea to sites to implement completely switched systems – at least where videoconferencing traffic will flow.

The *List of References* and *Further Reading* sections provide links to more detailed information about many of the areas discussed in this document. By its nature, a document like this has to be highly selective in what is included – as an example of how selective, the Cisco® guide to implementing Campus QoS (Flannagan et al, 2003) alone runs to 400 pages.

Those looking for further information about WAN (Wide Area Network) QoS provision on JANET should take a look at the UKERNA QoS Project at:

<http://www.ja.net/development/qos/>

This document is not an introduction to H.323 and the reader is assumed to have a working knowledge of H.323 infrastructure components such as endpoints, gatekeepers and MCUs (multipoint control units). An excellent introduction to H.323 can be found in the VTAS Guide *An Introduction to H.323 Videoconferencing* which can be found at:

<http://www.video.ja.net/323intro.pdf>

1.1 A Word of Caution

Applying QoS in a network can allow certain applications that require real-time data throughput, such as videoconferencing and VoIP (Voice over IP), to work apparently effortlessly even in networks under significant load. It should be noted that if switches and routers have overloaded CPUs and no memory, then applying QoS will just add additional load to an overloaded system and that will not help, irrespective of link capacity.

1.2 Chapter Outlines

It is not necessary to read through this document in any particular order, and it is recommended that you pick and choose as best fits your needs.

- **Chapter 2** is a closer study of the network profile and requirements of videoconferencing traffic. It discusses the relative importance of metrics such as packet loss, latency and jitter as well as typical bandwidths required for IP videoconferencing.
- **Chapter 3** looks at some of the issues and types of traffic that can commonly be found on LANs (Local Area Networks). This section covers issues such as speed and duplex settings, network broadcasts and spanning tree protocol updates, and takes a look at common campus-edge issues with H.323 videoconferencing, such as the use of NAT (Network Address Translation) and firewalls.
- **Chapter 4** looks at network engineering – providing physically separate links for real-time traffic.

- **Chapter 5** looks at traffic engineering – providing logical traffic separation in the form of VLANs (Virtual LANs).
- **Chapter 6** looks at providing QoS by means of policing, classification and priority queuing on Layer 2 networks – some configurations for Cisco® and 3Com® are included.
- **Chapter 7** is a case study at the University of Wales, Aberystwyth, where some of the above methods have been used to provide reliable videoconferencing across campus.

2 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

2.1 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.

2.2 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.

2.3 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.

2.4 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.

Figure 1: Testbed equipment and topology.

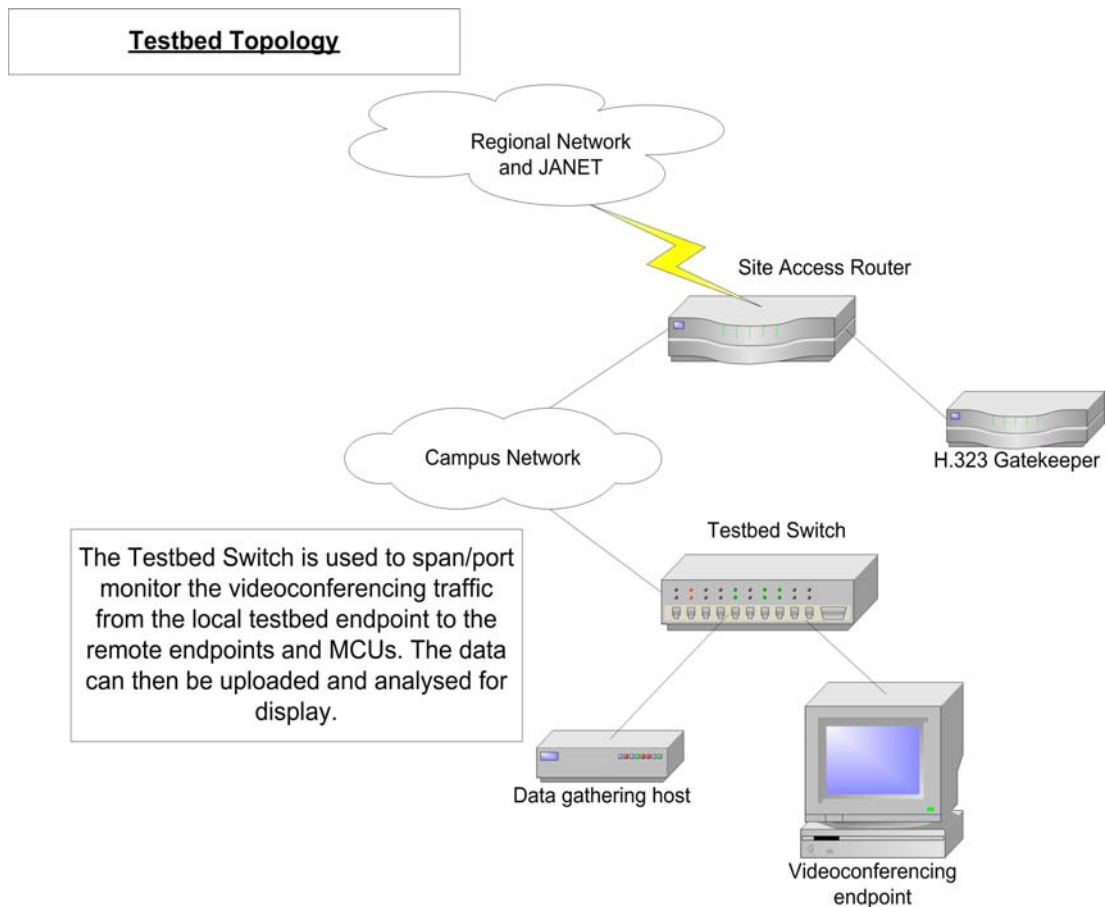


Table 1: 768kbit/s videoconference summary data

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

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.

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.

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





Figure 3: One second data from Testbed to Leeds MCU.



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.

Figure 4: One second bandwidth from Leeds MCU to the Testbed.

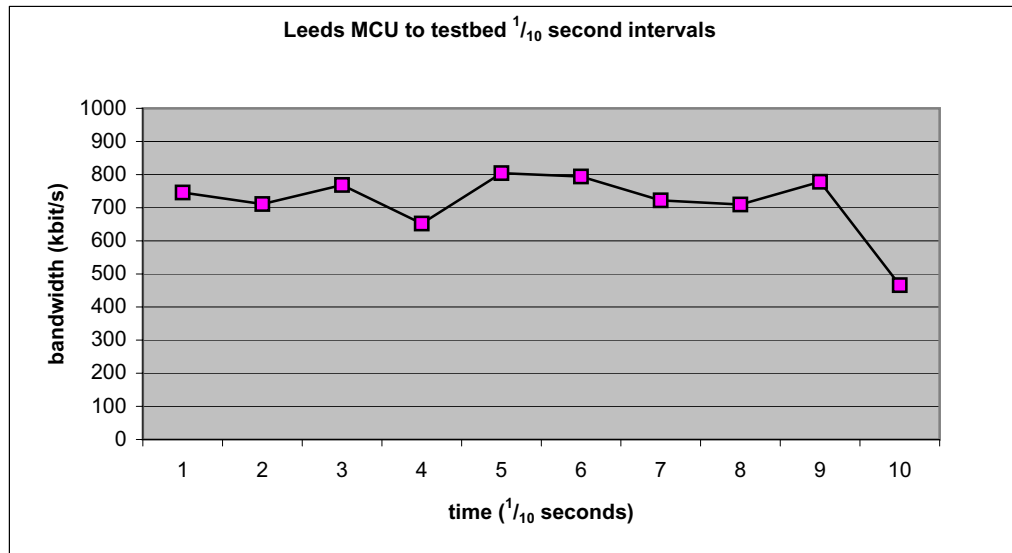
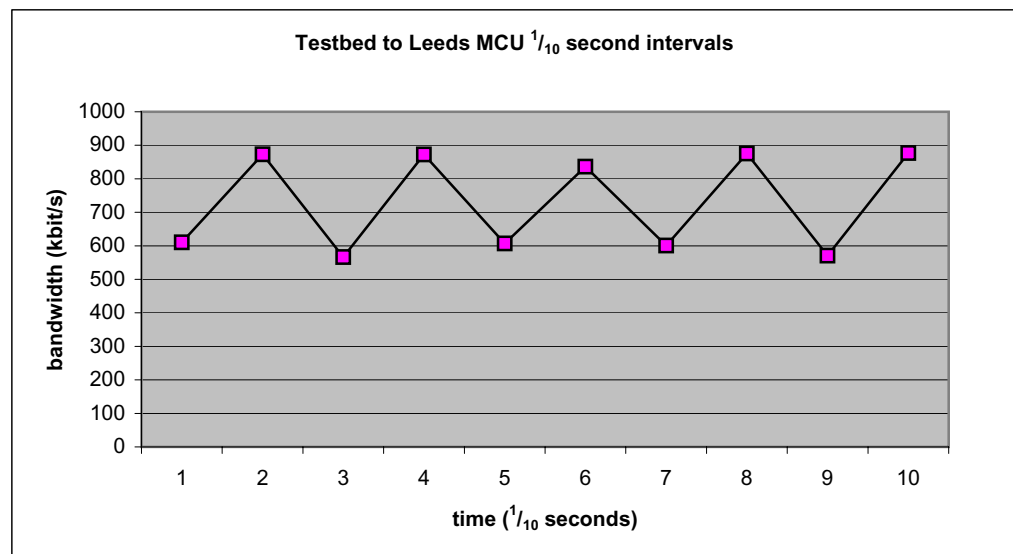


Figure 5: One second bandwidth from the Testbed to Leeds MCU.



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.

2.5 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.

3 Campus Traffic and Common Network Issues

3.1 Layer 1 - The Physical Layer

Most modern campuses have installed switched networks but there remain sections of some networks that have hubs or co-axial cabling with repeaters. Because these networks are built on protocols that accept collisions, and hence congestion, as a normal part of network life, their traffic forwarding algorithms will back off from sending frames in the face of congestion. This, added to the relatively low percentage load that a link will handle before retries and retransmission overload a segment, dictates that non-switched networks are not generally suitable for the demands that real-time videoconferencing traffic makes on the infrastructure.

It is highly recommended that H.323 systems, and indeed any real-time application, should be deployed in a purely switched network. Some introductory material on building switched networks can be found from (Cisco, 2003a) or (Long, 2002).

3.2 Layer 2 and 3 Issues

Whilst a switched environment has huge benefits over repeated networks, there are still limitations on size and scale and there are certain commonly found issues that may, or may not, be causing problems – but should be monitored. These issues are discussed below.

3.2.1 Speed and Duplex Settings

A major benefit of a switched environment – and UTP cabling – is that it allows the network to transport Ethernet frames highly efficiently, even when there is significant traffic – and even congestion – in a network.

Most NICs will now drive frames at the rate of 100Mbit/s (or 1Gbit/s) dependent on both the network media the NIC is connected to, and what is at the other end of that media (usually a port on a switch or router). On most network equipment, different ports can be set to transport frames at different speed settings, according to the capabilities of the host that the particular port is attached to. The main issue here is that on most ports there are two available options to select for the speed setting: 10Mbit/s or 100Mbit/s.

The advantage of the switched environment is that it allows hosts to communicate in full-duplex mode, i.e. with packets travelling ‘up’ and ‘down’ the link at the same time. In a shared Ethernet environment this is not supported: as the link only supports the transit of a single frame at a time, it has to be in either one direction or the other. So, as well as having a configurable network speed setting, individual ports also allow the duplex setting to be Full or Half (except in the case of Gigabit links which are always full-duplex).

For most manufacturers’ equipment, the default setting (and the setting that will be in place to start with) will be *auto-sense* (sometimes *auto-negotiate*, or just *auto*). For the majority of situations (and where equipment allows it), auto-sensing works fine. Both connected devices send a ‘handshake’ and establish the best speed and duplex setting that they can both support.

Unfortunately, there are times when auto-negotiation either does not work successfully, due to failed negotiation, or where a port on one of the devices has been configured explicitly to operate at a particular speed and duplex but the port at the other end of the ‘wire’ has not. Either of these scenarios can create a speed or duplex mismatch. If this does occur, the two machines will generate more frame or packet errors, leading to re-transmission of frames and packet loss.

Most network equipment that is ‘managed’ will have interface counts that can be monitored from a command line, web page or fully-fledged SNMP (Simple Network Management Protocol) monitoring system. The issue of speed and duplex mismatches can be traced by monitoring error counts on switch interfaces to ensure that they are not increasing.

In the case of H.323 terminals (endpoints) a speed/duplex mismatch can be disastrous to the quality of a conference as experienced by the user. Typically, if there is a speed/duplex mismatch in the path between two videoconferencing endpoints, conference participants are likely to notice the effects on quality of media play-out immediately. There will be stuttering and drop-outs in the audio playback, and blocky artefacts on the video, accompanied by jerky movement.

For this reason it is recommended that the NIC on the H.323 endpoint and the router/switchport to which the H.323 endpoint is connected are manually configured to be at the highest speed and duplex settings that the H.323 endpoint will support. In most cases this will mean setting both to 100Mbit/s, full-duplex. But there may be cases where equipment will not be capable of supporting this – for example the Polycom® Viewstation® 512 only supports 10Mbit/s half-duplex. If this is the case, ensure that the corresponding port or terminal is manually and explicitly set to be at the same speed/duplex settings as the equipment to which it is connected.

3.2.2 Spanning Tree Protocol Updates

STP (Spanning Tree Protocol) was designed to prevent packets being forever passed around an accidental, or deliberate, ‘loop’ in the network at Layer 2. Layer 3 routers have a time out count (TTL – ‘Time-to-Live’) which will drop an IP packet that is looping, but Layer 2 devices lack this safeguard. STP prevents loops by each switch allowing only one route to another destination and effectively putting the ‘other’ link or links into a standby mode, where they will not pass traffic.

STP will force switches to update their MAC (Medium Access Control) address tables in the event of a topology change being detected – and this is where some problems can arise. Topology changes are detected by a switch interface changing state from up to down, or from down to up. Every time that occurs, an STP update will be triggered – this is the default behaviour of most switches.

When an STP update is triggered, all switches affected flush their MAC address tables and commence re-learning of MAC/switchport mappings. In a large network this can take a considerable length of time – tens of seconds. During this period the switches will, for an unknown MAC address, act as they are designed and forward the frame to all ports on the switch except the one on which it was received. This creates a burst of traffic across the network and also can have other implications, such as traffic being able to be ‘sniffed’ at locations on the network where normally it would not be available. This can be a significant security risk if any usernames and passwords pass unencrypted over the network.

The surge in H.323 traffic being effectively broadcast following an STP update could cause problems, especially in situations where the network is working at a relatively high average load. The additional traffic caused could well overload switches causing increased latency and jitter and, potentially, lost frames.

There are some ways to limit the effect of STP updates:

- 1) Remove STP from core switches under central control. One of the main reasons for running STP is to prevent network loops affecting the network. In a centrally managed campus core this should not be an issue.
- 2) Tell each port that has a directly connected host that it should not generate update messages. Switching a machine’s power off or on will make the interface transition up/down or down/up. Setting the port with ‘Portfast’ or a similar command will prevent the transition from forcing a wider STP update.

3.2.3 Broadcast Traffic

Hosts will send broadcast frames for a number of reasons, the simple effect of this being that as you increase the number of hosts on a network, the broadcast traffic will increase in a relatively linear way.

It used to be fashionable to build large, flat networks, especially at sites with class B IP

addresses. Now, however, as the number of hosts at these organisations has increased — probably exponentially in most cases — these networks have become increasingly unmanageable and frequently have been split up into more manageable ‘chunks’ by utilizing Layer 3 subnets or Layer 2 VLANs.

The worst case situation is during a ‘broadcast storm’. These can be caused by many things, such as wrongly configured equipment or faulty hardware, but the effect can be catastrophic to traffic on the network. Effectively the system becomes overloaded with broadcast traffic and fails to deliver normal traffic effectively. Traffic monitoring will see broadcast rates increase sharply above the 5% or 10% baseline for the network.

3.2.4 H.323 Firewalls and Network Address Translation

H.323 is one of the few protocols that dynamically allocate destination ports to traffic. It is also one of the few protocols that embed IP address information within the packet’s data payload, rather than simply as source and destination addresses in headers. For these reasons, historically H.323 has not got on very well with firewalls. However, most recent firewalls are ‘H.323-Aware’ and have the ability to provide added security to endpoints and other H.323 equipment.

The availability of H.323-aware firewalls is matched by H.323-aware NAT in equipment – frequently in the same box as the firewall – allowing sites to deploy, or continue to use, private IP addressing on campuses and still have connectivity with the outside ‘public’ IP world. Whilst in general there seem to be fewer and fewer problems in the firewall and NAT areas, there have been instances where the introduction of these systems into an organisation has had an impact on H.323 traffic. These cases have been failure of traffic throughput, rather than occasional instances of loading or packet loss etc., so are fairly easy to trace.

In some instances the presented issues have been highly irregular in nature, such as all videoconferences from an organisation suddenly terminating after 30 minutes or so. In this case, the fault was endpoint manufacturer specific, in that certain manufacturers’ equipment functioned perfectly, whilst others failed specifically after the 30 minutes. This was eventually traced to the interoperation between the newly installed firewall and the videoconferencing endpoints, and was quickly fixed with a patch from the firewall manufacturer.

If a user has private IP networks deployed and wishes to run videoconferencing from their private space to the world, then there are basically two options available: to use an H.323-aware NAT device or to use a H.323 proxy, much in the same way as an http web proxy would be used.

As will be seen in the next chapter, the Welsh Video Network decided to deploy H.323 proxies at organisations rather than wrestle with firewalls and NAT, but both options are equally valid providing the NAT device will not impact too heavily on the metrics discussed in Chapter 2.

3.2.5 General Traffic Profile

It is usual to see fluctuations in bandwidth use through the day and week in campus networks. It is worth building up a picture of what traffic patterns are usual for your network, through habitual monitoring of the network.

One common use of this is to detect hacked machines on campus which may be scanning other machines, either internally or externally. If traffic to a department is habitually greater than traffic coming from a department, then a sudden change in the direction of greater traffic will usually be significant.

This is especially true of JANET access links. Most sites will expect to see higher bandwidth into their site than leaving it, as a small http ‘get’ results in a large http transfer to the requesting machine. The other way round, with more traffic leaving the site, can often be the sign of compromised hosts. This may not be true for your network, but having an understanding of how traffic is normally behaving will lead to faster resolution when something is not right.

It is true that larger organisations have tended more towards being net traffic exporters, but it is worth checking whether hosts sending a lot of data onto the network really should be doing so.

It can also be the case that networks that perform well most of the time have problems during peak load times (usually during lesson/lecture switchovers and during lunchtimes). Again, habitual monitoring will indicate potential issues at peak times.

4 Network Engineering - Physical Separation of H.323 Traffic

This section examines the role of physical separation in the provision of reliable and secure links for real-time traffic.

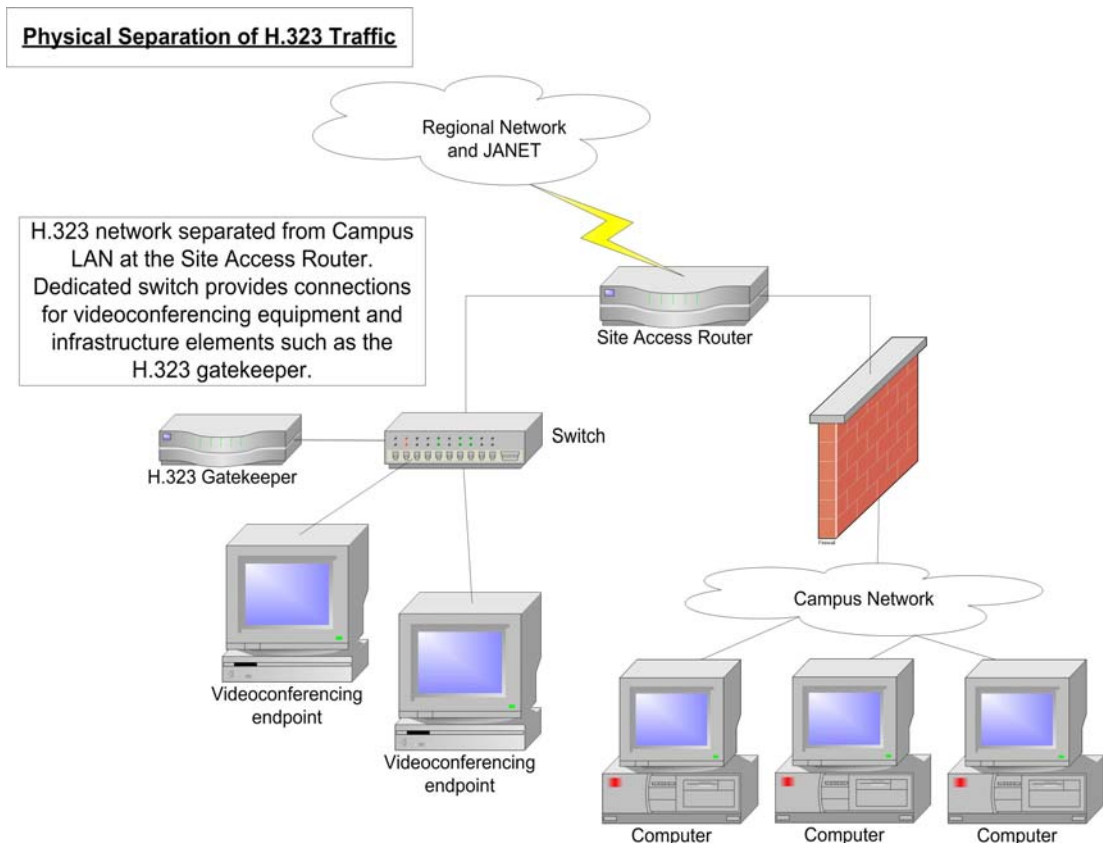
In many cases equipment can be directly connected together without being also plugged into the campus LAN directly. This can be especially beneficial when equipment is dedicated to a task – such as videoconferencing equipment – that operates stand-alone, i.e. is not part of a desktop PC used for normal network access/applications.

Direct physical connection of the equipment can be arranged to a suitable point, usually to a dedicated switch or even directly connected to the SAR (Site Access Router). The actual topology will depend on a number of factors such as the equipment available, the distance between equipment and the topology that will be most beneficial for the type of traffic involved.

Providing dedicated physical links to distant equipment would previously have necessitated the installation of dedicated cable runs from source to destination. With the advent of structured cabling installation, this process has become far simpler. It is still dependent, however, on the availability of cables between patch panels and the overall distance between equipment — Fast Ethernet limits cable length to about 100 metres.

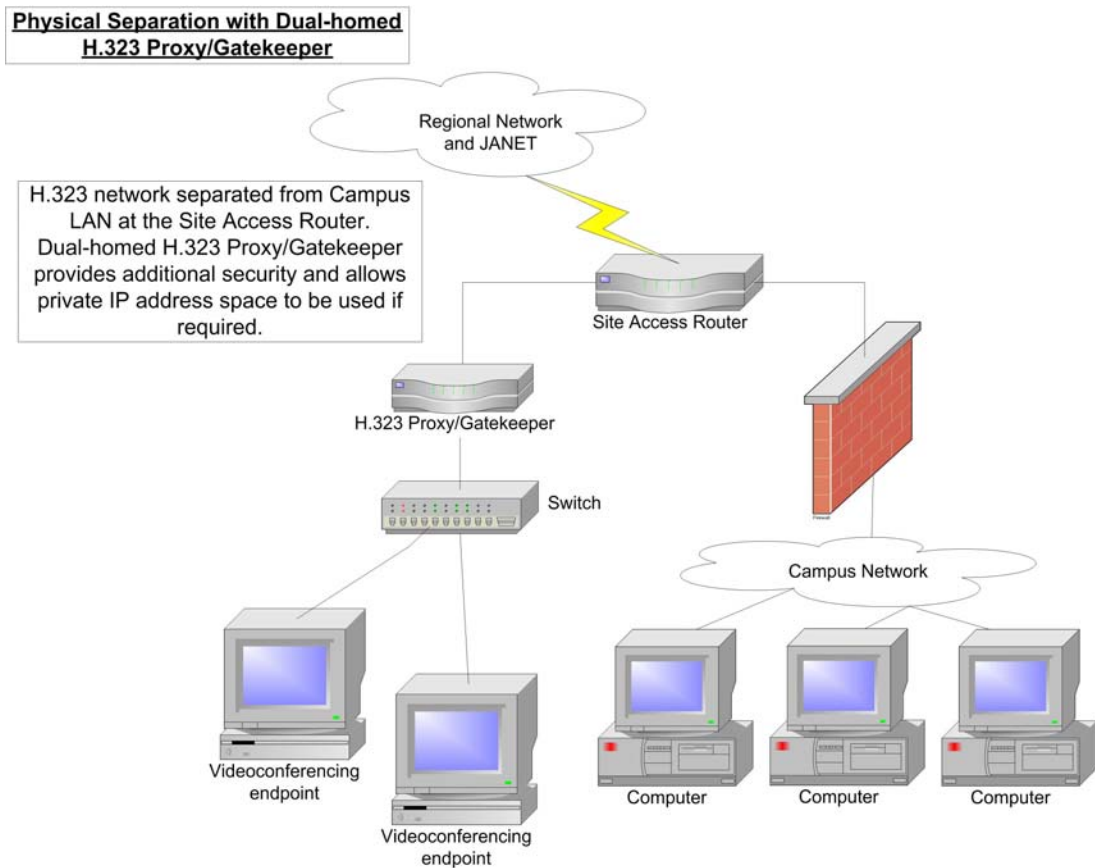
The Welsh Video Network used physical separation of equipment in their deployment of studios across Wales. Not all studios could be directly patched, but those that could were connected to a dedicated switch directly connected to the SAR (see Figure 6 (*below*) for details). This topology lets the H.323 traffic bypass the core of the network, so avoiding any potential campus network issues impacting on H.323 videoconferences. It also allows work to be done on the campus network without impacting on scheduled videoconferences.

Figure 6: Physical separation.



A number of refinements can be made to this topology if, as in the case with the Welsh Video Network, the gatekeeper is also capable of proxying H.323 traffic. To increase security, and in those cases where sites may wish to use private IP address space for videoconferencing endpoints, the gatekeeper/proxy can be moved to be in line between the switch and the site access router, as in Figure 7 (below).

Figure 7: Dual-homed gatekeeper/proxy.



It would also be possible to put a switch between the site access router and the gatekeeper/proxy. In the scenario above there would be little benefit, but this can be useful if you need to connect into the campus network.

5 Traffic Engineering 1: Logical Separation at Layer 2 - the VLAN

In many cases, it will not be possible to physically separate H.323 and campus network traffic, and it will be necessary for all traffic to share the same physical links. In this case there are some methods that can be used to provide some level of protection to H.323 traffic, above that provided to the campus traffic.

This section will look at the concept and practice of providing Virtual LANs or VLANs and then Chapter 6 will look at ‘persuading’ network equipment to queue and forward traffic differently, based on what that traffic is, i.e. providing preferential treatment to H.323 traffic.

A VLAN is, in effect, an overlay network that uses the underlying physical network to provide one or more virtual LANs on top of it. At installation, a physical LAN will map to the default VLAN (usually VLAN 1) and so any broadcast traffic sent by any host will be forwarded to all other hosts on the network.

The physical network can then be further split up into a number of VLANs, such that a broadcast will be restricted to be within the VLAN only, and will only be seen by other hosts in the same VLAN, rather than all the hosts on the same physical LAN.

A typical use of a VLAN is where a common group of users are topologically spread out, but for whatever reasons need to be in the same broadcast domain. That could be for server requirements or other application needs, but the network administrator does not want their traffic to be seen on the rest of the network, whether that is for load, security or other reasons.

Placing all the hosts in the same VLAN will allow hosts spread around the campus network to act as though they are in their own, dedicated LAN. They will receive no broadcast traffic from outside the VLAN and will send no broadcast traffic out of the VLAN.

VLANs can be used to:

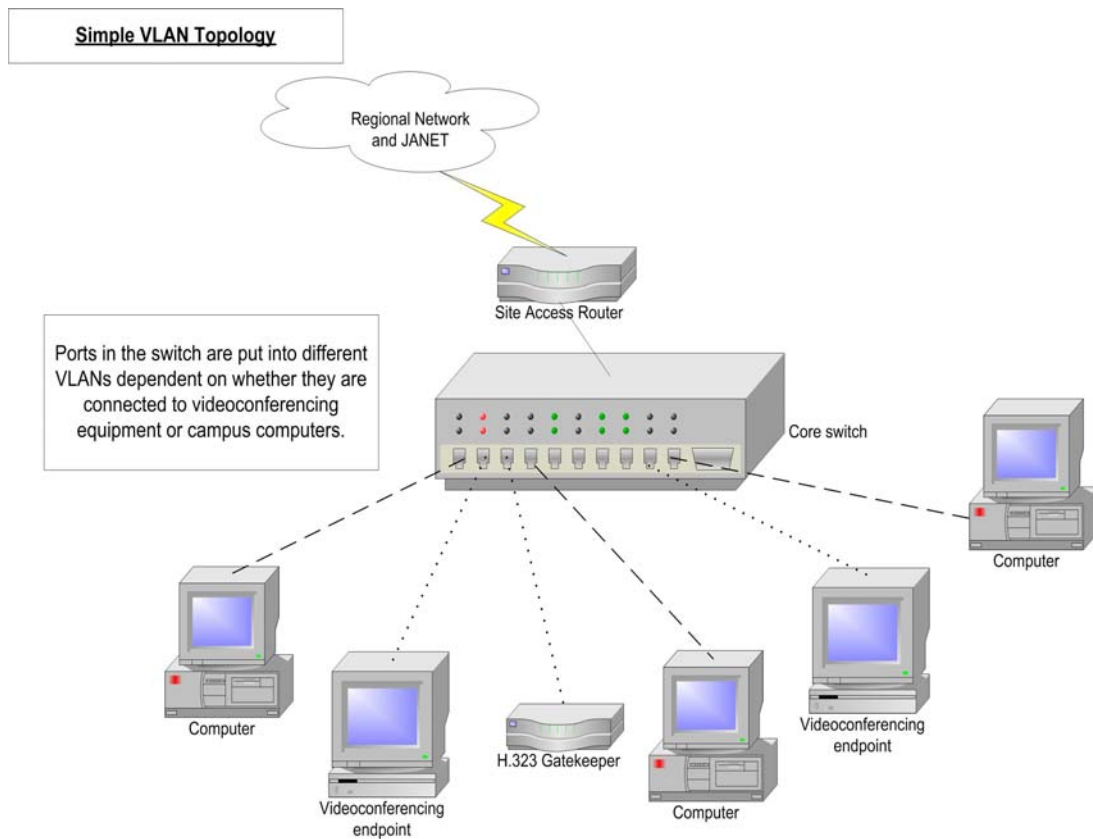
- provide extra security
- create logical groups that reflect organisational structure
- cut down on unwanted or unnecessary traffic
- reduce broadcasts
- ease network management.

It should be remembered that VLANs are Layer 2 constructions and any traffic that needs to move outside the VLAN will require a Layer 3 routing decision to be taken – in the same way as if it was a collection of separate physical LANs connected to a router. For this reason VLANs and IP subnets are generally identical in scope.

In the case of H.323 equipment, a VLAN can be used to protect the videoconferencing equipment, and the network links to that equipment, from receiving certain broadcast traffic, thus relieving the equipment’s network link and interface of handling that additional load.

Figure 8 (*overleaf*) shows the use of a VLAN to link H.323 videoconferencing equipment together.

Figure 8: Simple VLANs.



To configure a VLAN port on a Cisco® 2950 switch:

```
! in Interface Configuration mode
Switch(config)#int fa 0/3
Switch(config-if)#switchport access vlan 323
%Access VLAN does not exist. Creating VLAN 323
Switch(config-if)#int fa0/4
Switch(config-if)#switchport access vlan 323
Switch(config-if)#CTRL-Z
```

As can be seen, if the VLAN does not currently exist the switch will automatically create it as soon a port is put into it.

Showing the running configuration gives:

```
...
!
interface FastEthernet0/2
!
interface FastEthernet0/3
  switchport access vlan 323
!
interface FastEthernet0/4
  switchport access vlan 323
!
...
```

To configure a VLAN port on a 3Com® 4400 switch

Firstly, the VLAN must be explicitly created

Select Menu Option: bridge vlan create 323 VIDEO-VLAN

Here 323 is the VLAN ID and VIDEO-VLAN is the text name of the VLAN

Select Menu Option: Bridge vlan modify addport 323 1:2 untagged

This adds VLAN 323 to port 1:2 (unit 1: port 2). The untagged keyword tells the switch that a host will be connected to this port and it is not a VLAN trunk.

This will allow the equipment connected to VLAN 323 ports to communicate with each other, and all the other ports in the default VLAN 1 to talk to each other, but at the moment there is no method of sending traffic from one VLAN to the other which of course is essential. In most cases, the core network switch at a campus will be a device that is capable of Layer 3 operation as well as Layer 2 switching. In this case, the routing decision for inter-VLAN traffic can take place within the same switch. See section 5.1 on inter-VLAN routing.

The simple example above shows VLANs on a single switch, but in most situations the VLANs will need to span more than a single switch. One way to link two switches together that have VLANs configured is simply to provide two separate physical links between the switches, one for each VLAN, and put those link ports into the appropriate VLANs. However, in general, where there are a number of VLANs in use, providing physical links per VLAN will not be scaleable.

In order to allow VLANs to traverse inter-switch links, VLAN trunks can be established which consist of a single physical link that is capable of supporting as many VLANs as are configured on the switch. It is worth remembering that a single host connected to a port can only be in one VLAN, but inter-switch links can support as many VLANs as required.

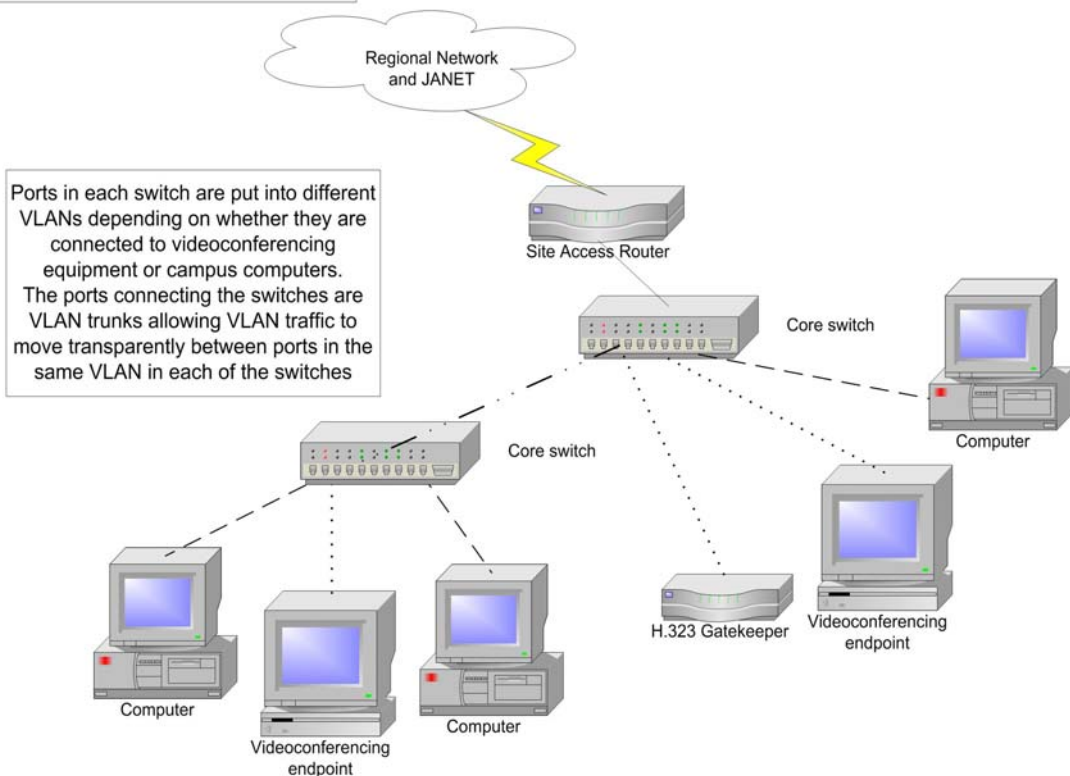
There is a standard for providing inter-switch VLAN trunks – however, that does not mean that there are no proprietary options available as well. The standard for VLAN trunking is based on 802.1q – also referred to as simply ‘dot1q’.

In theory, and mostly in practice, 802.1q allows different manufacturers’ switches to trunk VLANs successfully between them. Cisco® have an alternative which is known as ISL (Inter-Switch Link). This document will assume 802.1q VLANs will be used.

Figure 9 (*overleaf*) shows the topology with trunked VLANs between the switches. At this stage this will allow the dotted and dashed links to communicate between the switches – now, there are truly two virtual LANs, and at this stage it should be remembered that there is no provision for inter-VLAN traffic as Layer 3 routes between the VLANs have not been configured.

It is of course possible to place an entire switch, and all its connected hosts and downstream switches, into a VLAN by simply configuring a VLAN on the port on the upstream switch to which it connects. The alternative would be specifically to place all the ports in the switch, and the link ports, into the VLAN.

Figure 9: Trunked VLANs.

Trunked VLAN Topology

To configure a VLAN trunk for all VLANs between two Cisco® switches:

```
Switch 1 (Cisco® 3524)
! Port 0/24 is the link to the second switch
Int fa0/24
  switchport trunk encapsulation 802.1q
  switchport mode trunk

Switch 2 (Cisco® 2950)
!
Int fa0/24
  switchport mode trunk
!
```

You will see that for Switch 1, the Cisco® 3524, the VLAN trunk encapsulation has been explicitly set to 802.1q. By default the 3524 uses Cisco® ISL rather than 802.1q encapsulation for VLAN trunks.

However, Switch 2, the Cisco® 2950, has no support for ISL and so uses 802.1q by default – hence there is no encapsulation command for Switch 2.

To configure a VLAN trunk for all VLANs between two 3Com® 4400 switches:

On each switch, create the VLANs as described above and then use the command:

Select Menu Option: bridge vlan modify addport 323 1:24 tagged

This will add VLAN 323 to port 1:24 which is the link port to the other switch. The tagged keyword tells the switch that this is a VLAN trunk as opposed to a single host connection.

To configure a VLAN trunk for all VLANs between a 3Com® 4400 and a Cisco® 2950 switch:

The configuration is the same as the configuration detailed above. As the trunk is based on the 802.1q standard it will work between different manufacturers' equipment.

5.1 Inter-VLAN Routing

Provisioning VLANs can provide a solution which will allow a LAN to expand beyond what was possible before, and allow better management and control of the network. Most of the large, flat networks that were previously deployed at some campuses are gradually being replaced by VLANed and subnetted networks, to allow better segmentation and control of broadcast and other traffic.

Remember, though, that VLANs exist at Layer 2 only. To allow traffic to move between VLANs, a Layer 3 routing decision must be taken.

Our example network has a Cisco® 6500 switch at the core of the network, so VLANs are provisioned from the core 6500 out towards the edges, and the 6500 router blade provides the Layer 3 inter-VLAN routing capability.

Cisco® 6509 Layer 2 Switch VLAN commands

To set a VLAN's name

```
set vlan 323 name Video_Conference_Suite
```

To disable Spanning Tree on VLANs

```
set spantree disable 323
```

To place switchports into VLAN 323

```
#module 2 : 8-port 1000BaseX Ethernet
```

```
set module name 2
```

```
set vlan 323 2/5
```

To configure a port as a 802.1q VLAN trunk

```
set trunk 2/2 on 802.1q
```

Creating a VLAN Interface on a Cisco® 6509 with MSFC

```
interface vlan172
description Video Conference Network
ip address a.b.c.d 255.255.255.0
```

As inter-VLAN traffic is a Layer 3 decision, stopping traffic moving between VLANs needs to be done either by not routing traffic, or by controlling traffic by means of a firewall or access-list restrictions. This is especially true of IP directed-broadcast or Layer 3 broadcast traffic, e.g. to x.255.255.255, which you would normally not want to allow inter-VLAN.

It must be noted that whilst the provisioning of VLANs is relatively simple, provisioning very strict Layer 3 filters on inter-VLAN traffic in an environment where there may be Microsoft®, Novell® or other complex client-server systems in place can be complex and demanding. A full understanding of the protocols, ports and traffic between clients and servers is needed.

6 Traffic Engineering 2: Layer 2 Prioritisation - CoS (Class of Service)

Local area networks of any significant size, which almost certainly encompasses all those at educational organisations, are complex and unpredictable systems. The traffic flows produced within these networks, and the interactions between different flows within network components such as switches, are highly complex. Classifying, policing and priority queuing allow the network administrator some control over how these flows transit the network, and – crucially for voice and video traffic – allow time-critical traffic to have priority over other, less time-sensitive traffic.

In cases where physical or logical traffic segregation is not possible, or not adequate in the case of VLANs alone, the traffic prioritisation (or QoS) features available on most Layer 2 and Layer 3 network equipment can be used to provide reasonably robust traffic flows through a normally congested network. It should be noted, however, that this will help only where traffic is oversubscribing an output interface on a piece of equipment. If the traffic level is causing high CPU or memory usage, applying QoS may not help at all.

The most common bandwidth problems in the LAN arise where either high speed core links, frequently now at gigabit speeds, break out to lower speed links to departments or when multiple links into a switch feed into a single uplink towards the core. It is at these points, where the link speed falls or where many links are aggregated into one, that prioritisation will ensure that the time-critical traffic will receive priority treatment and not have to sit in queues.

The effect of traffic overload in these scenarios is either an increase in latency, or, if the latency becomes large enough, the packets risk being dropped off the tail of the queue, resulting in packet loss.

In order to configure QoS, there are a number of elements or processes that need to be considered – Classification and Marking, Policing and Queuing.

6.1 Classification and Marking

In order to treat some traffic in the switch differently to the rest of the traffic flowing through it, it is first necessary to divide the traffic up into different ‘Classes’. Depending on the equipment in your network, it may be possible to select eight or more different classes of traffic and treat them all differently. However, in practice, due to the limitations of most switches currently in circulation, it is normal to only consider two or four classes. This also makes configuration, monitoring and control far simpler.

Once the traffic has been classified it must be marked, via one of the standard methods, in order to indicate to later processes within the switch and to next-hop switches that this is the traffic that should receive some specific treatment; in this case prioritisation over non-H.323 traffic.

Traffic is normally classified by explicitly configuring the switch to select certain traffic flowing through it. It may be that all ingress traffic on a certain switchport is selected, or all traffic to and/or from a certain MAC address. Increasingly, Layer 2 switches are also adding Layer 3 functionality into their code, so in many cases with newer switches such as the 3Com® 4400s and Cisco® 3550s it is possible to select traffic by Layer 3 identifiers such as source and/or destination IP address.

Example.

To classify traffic on a Cisco® 2950 switch the following commands could be used:

```
Class-map match-all VIDEO
  Match access-group 190
  !
  !
access-list 190 permit ip host 10.1.1.1 any
!
```

In the above case a class-map called VIDEO is created into which is placed all traffic that matches the constraints of the access-list – i.e. all IP traffic from host 10.1.1.1.

Once the traffic has been put into classes, it is marked to enable treatment appropriate to its classification. Traffic is usually classified and marked on ingress to a port, but prioritisation, i.e. queuing, is only applied on egress from a port. Marking traffic at ingress allows the egress port to identify and prioritise traffic. In our Cisco® 2950, marking is achieved through the use of policy-map commands.

Example: Marking Traffic

```
Class-map match-all VIDEO
  Match access-group 190
  !
policy-map INCOMING
  class VIDEO
    set ip dscp 46
  !
access-list 190 permit ip host 10.1.1.1 any
!
```

In this case, adding to the above class-map example, a policy-map called INCOMING is created in which the class VIDEO has its DSCP value set to 46.

So in Cisco® terms – and this applies to Cisco® Layer 3 devices as well as switches – the normal course of action is to create one or more CLASS-MAPS which classify the traffic and then apply a POLICY-MAP to attach a defined policy to the classes created in the CLASS-MAPS.

There are a number of ways that traffic can be marked, as defined by standards that have now become well established. The two standards most frequently used are the 802.1p CoS (Class of Service) and DSCP (Differentiated Services Code Point). At Layer 2, strictly speaking, only the 802.1p values are used; however, as was mentioned above, many new Layer 2 switches are able, at least partly, to parse Layer 3 headers in order to extract or match certain information.

At the moment all that has been done is to specify a policy to mark the traffic with DSCP value 46. DSCP 46 is defined as being for EF (Expedited Forwarding) traffic flows and is the value usually assigned to all interactive voice and video traffic.

So, whereas previously it would have been necessary to mark all traffic in the LAN with 802.1p priority values from 0 to 7, and then remark traffic at the Layer 2 to Layer 3 boundary to DSCP values, it is now possible to use DSCP values across the LAN as well as the WAN, depending on the equipment in the network path.

At Layer 3 the IP Precedence value was previously used; DSCP is designed to replace it. However, some equipment at Layer 3 will only mark with IP Precedence values which, like 802.1p values, range from 0 through 7.

The 802.1p standard defines seven levels of CoS from 0 through to 7 (highest priority). 802.1p is a sub-set of the 802.1q standard which added additional fields into the header of a standard Ethernet frame allowing it to contain VLAN identifiers as well as the priority values.

DSCP defines 64 values (0-63) which can be used to treat traffic in a very granular manner, assuming the equipment has enough queues to support queuing at large numbers of different rates. However, most switches support only 2, 4 or 8 queues per output port so DSCP values in certain ranges, as well as multiple 802.1p priority values, will tend to map to a smaller range of queues, as shown in table 2 (*below*).

Table 2: 3Com® 4400 CoS to queue map

Priority	Queue Index	802.1p Name
0	1 (lowest)	Best Effort
1	1 (lowest)	Background
2	1 (lowest)	Reserved
3	2	Excellent Effort
4	2	Controlled Load
5	3	Video
6	4 (highest)	Voice
7	4 (highest)	Network

It is very important, having decided to run QoS at Layer 2 or Layer 3, that a policy is designed to specify which traffic flows will be marked, what value and scheme they will be marked with, and what sort of treatment the flow requires. The earlier in the process this 'design' is formulated, the simpler later configuration becomes.

It is also clear from the above that it is possible to configure all switches to classify and mark frames/packets on ingress. However, most large networks (at least in some areas) consist of edge switches which connect hosts, and core switches which only connect switches. It is critical that all the edges of the network are treated the same – which will consist of marking traffic in a consistent manner at all ingress points. If this is done at all the edges, it is no longer necessary to re-classify/police that traffic on ingress to next-hop switches as it moves closer to the core of the network. This can simplify configuration in core switches as they will only be required to queue traffic previously marked by other switches.

Again, as with designing a marking policy, designing a trust policy early on can simplify QoS management and administration later. It is possible to run QoS on one 'corner' of the network and gradually expand it outwards, but the ingress point to this 'corner' will need to have classification applied to prevent spuriously marked traffic from interfering with deliberately marked traffic.

It has been found that traffic in a number of networks is already being marked. Applications running on all current versions of desktop and server operating systems have the ability to mark traffic leaving their host, and some do. Most commonly these are voice- or video-based applications, but just as easily it could be game or denial-of-service applications that use it.

It is highly advisable, therefore, not only that organisations classify and mark traffic that is explicitly selected for better treatment, but also that they ensure that other traffic entering the switch has not already been marked by an external host or application. It is possible to re-mark all traffic, except selected traffic, to CoS/DSCP/IP Precedence 0; or at least to re-mark to 0 traffic that is trying to use other CoS values. So for example, either all traffic entering the

switch is re-marked to CoS 0 (except for H.323 traffic from certain IP/MAC addresses or from a certain switchport); or all traffic marked as CoS 5 is re-marked to CoS 0, except for relevant H.323 traffic. This way an organisation can be sure that the only traffic entering their Premium flow is traffic that they specifically want to allow.

6.2 A Note on Marking

At Layer 2 it is frequently the case that it is only necessary, or possible, to configure ingress marking to the switch and the priority treatment is pre-defined: in some cases behaviour can be changed, in others it is fixed. As an example the Cisco® 3524 switch allows no configuration at all; incoming packets with CoS value 4-7 will be prioritised at egress above those marked with CoS 0-3.

Other switches may have four queues and define the prioritisation treatment differently, e.g. CoS 5 in one switch may be the equivalent of CoS 6 in another. This can become more complex if DSCP and even IP Precedence values are all mixed in, and even different switch models from the same manufacturer may map 802.1p priority values into different output queues. Fortunately many switches also now allow these mappings to be changed in order to suit your specific needs.

As mentioned above, defining your classification scheme early on will simplify matters. At the simplest level this would just be to define which values should be applied to which types of traffic, as shown in table 3 (*below*).

Table 3: Sample classification scheme

Traffic Type	Description	CoS value	DSCP Value	IP Precedence
Best Effort	All web, e-mail etc	0	0	0
Premium IP	H.323 and VoIP	5	46	5

6.3 Policing

If you imagine the case where priority queuing is applied without any restriction, then a host, or hosts, which managed to inject the full link bandwidth, suitably marked as Premium traffic, could completely starve the Best Effort flow of any bandwidth. This is obviously not desirable, and policing allows the network administrator to limit the amount of traffic entering the Premium flow at a given point in the network. Policing could also re-mark traffic, either from unauthorised hosts or from a host that was trying to inject too much traffic into the Premium class. This must of course be applied carefully, as policing a flow and not allowing adequate bandwidth for all possible Premium flows that may be set up will severely impact all those flows – causing packet loss and hence degradation in video or voice applications. Policing, unfortunately, is not available on all switches.

In many circumstances, traffic shaping is often applied in conjunction with policing. For H.323, traffic-shaping is not recommended as, given the nature of UDP traffic, there is no mechanism, and no time, for a retransmission of traffic should it be dropped by a shaper. Shaping should only be applied to UDP voice and video streams where an adequate buffer size can be allocated to the shaper to prevent packet loss from occurring due to the shaping.

6.4 Queuing

Once the traffic has been selected, marked and policed, it must be treated differently in some manner in the switch. Usually this is effected by putting different classes of traffic into different egress queues at the output port. The queues are then emptied by the port onto the link in deference to their priority, so high priority queues will be served by the port more frequently than low priority ones. The frequency and nature of the difference between queues can be fixed in a switch, or on larger switches may be fully configurable.

Some options that may be encountered are:

- **Strict Priority Queuing.** Some traffic is selected and put into a strict priority queue. This queue is always serviced by the switch if there is any traffic in it, irrespective of any other traffic in any other queues. This can mean that if there is enough traffic in the strict priority queue, other traffic can be completely starved of bandwidth.
- **Priority Queuing (with bandwidth limit).** The use of strict priority queuing (*above*) is far more useful if it provides a mechanism to police the bandwidth allowed into the priority queue. Voice and video traffic will normally use a priority queue with a limit to prevent starvation of other flows, whilst protecting adequate bandwidth for the calls.

Note that while the 802.1p standard specifies eight different priority levels, many switches, including many high-end ones, simplify the choices available, sometimes giving only a high and low priority queue.

There are also often many options available for drop-precedence in queues, allowing traffic in certain classes to be dropped before others. In most cases, as the aim is to provide the best possible performance for the H.323 traffic, these will not apply. In general the default behaviour for the Best Effort class should be satisfactory.

On most edge switches, such as the Cisco® and 3Com® switches, there is a default setting for queues, and frames/packets marked with a certain CoS or DSCP value will be placed into the appropriate queue and shipped out onto the network.

As has been mentioned above, the hard part is not getting a particular piece of network equipment to prioritise traffic, but rather designing a scheme that will suit the network, and which will cover all the capabilities of the various switches in that network. Simplifying provisioning to four, or fewer, classes makes this process easier and is highly recommended in the first instance.

Table 4 (*below*) shows detail from a 3Com® 4400 switch which has four independent queues, and the default mapping between CoS value and queue:

*Table 4: Current CoS to queue mappings from switch command line
(from: trafficManagement|qos|trafficQueue|Summary)*

Priority CoS Value	Queue Index
0	1 (lowest)
1	1 (lowest)
2	1 (lowest)
3	2
4	2
5	3
6	4 (highest)
7	4 (highest)

6.5 Configuration Samples

The following sections look at specific configurations on Cisco® 3524, Cisco® 2950 and 3Com® 4400 switches. It is only possible here to cover a fraction of the options available. Please see the Further Reading section for links to more comprehensive documentation.

6.5.1 Applying QoS on a Cisco® 3524 Switch

The standard Cisco® 3524 switch has two available queues – high and low. By default, incoming traffic to the switch which is tagged with a CoS of 0-3 will be placed in the low-

priority queue. Traffic with CoS 4-7 will be placed in the high-priority queue. This behaviour is not configurable and is enabled by default. Traffic entering the switch which is untagged can be marked with a CoS value.

The issue with this, of course, is that any host attached to a Cisco® 3524 switch can put priority marked traffic into that switch and it will be honoured, above other traffic. This could be used as a denial-of-service mechanism.

Note that the Cisco® 3524-PWR-XL switch with Software version 12.0(5)XU or higher does support CoS remarking (Flannagan et al, 2003).

6.5.2 Applying QoS on a Cisco® 2950 Switch

QoS is not enabled by default on the Cisco® 2950 switch.

On many Cisco® switches there is the option of using basic commands (e.g. mls qos trust dscp) or using the Cisco® MQC (Modular QoS Command line). In general the MQC is preferable as it is useable not only across much of the Layer 2 range but also across the Layer 3 routers as well.

This code sample (*below*) uses Modular QoS Command-line commands to configure priority queuing for all traffic from host 10.1.1.1, policed at (limited to) 20Mbit/s. See the previous sections on Classification, Marking and Policing for an explanation of the commands.

```
!  
Class-map match-all VIDEO  
  Match access-group 190  
!  
policy-map INCOMING  
  class VIDEO  
    set ip dscp 46  
    police 20000000 4096 exceed-action drop  
!  
...  
!  
interface FastEthernet 0/23  
  service-policy input INCOMING  
!  
access-list 190 permit ip host 10.1.1.1 any  
!
```

6.5.3 Applying QoS on a 3Com® 4400 Switch

Firstly, create a Classifier for your chosen traffic. The code sample below will create a Classifier numbered 323 and named VIDEO-CLASSIFIER. It will class traffic from the stated IP address – here 10.1.1.1:

```
: Traffic|qos|classifier|create
Enter Classifier Number(): 323
Enter Classifier Name:VIDEO-CLASSIFIER
Enter Classifier Type: ipaddr
Enter IP address[0.0.0.0]: 10.1.1.1
Enter number of bits in subnet mask: 32
```

Next, a QoS Profile must be created. Drop down a menu level by typing q and then:

```
: profile | create
Enter profile number: 323
Enter profile name: VIDEO-PROFILE
```

At most points entering the command ‘summary’ will show the current status of the configuration of each element.

The profile just created will include the default classifier 1. This is not required as it is the default traffic running through the switch. Remove it as follows:

```
: Remove
Select profile number: 323
Select classifier number: 1
```

Add to this (323) profile the classifier (323) that has already been created (*see above*) and make it use service profile 5, as follows:

```
: AddClassifier
Select profile number: 323
Select classifier number: 323
Enter service level number: 5
```

(See table 5 (*overleaf*) for the service levels available)

Finally, this profile must be applied to all necessary ports on the switch:

```
: Assign
Select ports (unit:port...,?): 1:1-24
Enter profile number: 323
```

Table 5: Service levels available in the 3Com® 4400 switch
(from: trafficManagement|qos|serviceLevel|summary)

Num	Name	Conforming to		Used in QoS Profile
		Priority	DSCP	
1	Drop		-	None
2	Best Effort	0	0	None
3	Business Critical	3	16	None
4	Video Applications	5	24	None
5	Voice Applications	6	46 (EF)	323
6	Network Control	7	48	None

6.6 Summary

Priority queuing can be applied to H.323 traffic in order to increase the reliability of delivery across a campus network. Queuing and VLANs can be configured independently or together. Providing a dedicated H.323 VLAN would seem to be the best currently available method of running voice and video services reliably and consistently.

In testing the configurations in this document, the team has been pleasantly surprised by the fact that the equipment behaved exactly as expected when QoS was applied, even with severe overloading of links. Only when the edge switches were pushed to their packet switching limits, with the injection of very large numbers of small packets, did the switches have problems maintaining forwarding rates.

The complexity of configuring QoS is frequently more in the planning and administration rather than in the actual command-line configuration on the switches. In networks with homogenous equipment, it is far simpler and quicker to provision QoS than in the case where a number of different manufacturers' equipment interleaves. As QoS is strictly provisioned on a per-hop basis, there are no issues with inter-working between equipment. However, ensuring that different configurations maintain the appropriate behaviours for specific classes of service, and maintaining CoS ingress protection at all the edges of the network, can become somewhat challenging in a large network – which is often a dynamic and fluid environment.

Even with the added complexity, QoS is one of those tools that will add value in the reliability to video or voice services – and its deployment should be seriously considered.

7 Case Study: The QoS H.323 Network at University of Wales, Aberystwyth

The UWA (University of Wales, Aberystwyth) gatekeeper currently has seven H.323 videoconferencing endpoints in its zone. These consist of:

- four UWA WVN PictureTel 970 CODECs (each with a potential bandwidth of up to 2Mbit/s).
- two Leadtek BVP 8770 H.323 videophones. These have a maximum bandwidth of 640kbit/s.
- a Tandberg® 8000, with a maximum bandwidth of 768kbit/s.

These are distributed around the University as follows:

- ABER-PENGLAIS-HOL (HOL), a 970, is on the main campus.
- ABER-PENGLAIS-WVN (WVN), a 970, is also on the main campus.
- ABER-DILS-125 (DILS), a 970, is on a remote campus, located a mile away from the main campus. The two campuses are linked by a dedicated fibre. As well as University departments, this link also supports the Welsh Institute for Rural Studies (formerly the Welsh Agricultural College).
- ABER-OLDCOL-1LP (1LP), a 970, is situated in the Old College and is linked to the main campus over 100Mbit/s microwave link.
- ABER-PENGLAIS-IS (IS) is a Tandberg® 8000 which is also on the main campus.
- The Leadtek videophones are currently set up to be portable around the college, so are usually only connected temporarily. They can register with the gatekeeper with various IP addresses, depending on their physical location.

There are also two WVN studios on adjacent organisations' networks that are registered with the UWA gatekeeper and whose data is also routed through the UWA Cisco® MCM (Multimedia Conference Manager) H.323 proxy. These are the National Library of Wales (referred to here as NLW) and the second campus of Coleg Ceredigion, the local Further Education College (referred to here as CC).

This example examines how the different studios at UWA have been connected to the gatekeeper, and to the SAR, and what QoS measures have been taken to support the use of these CODECS.

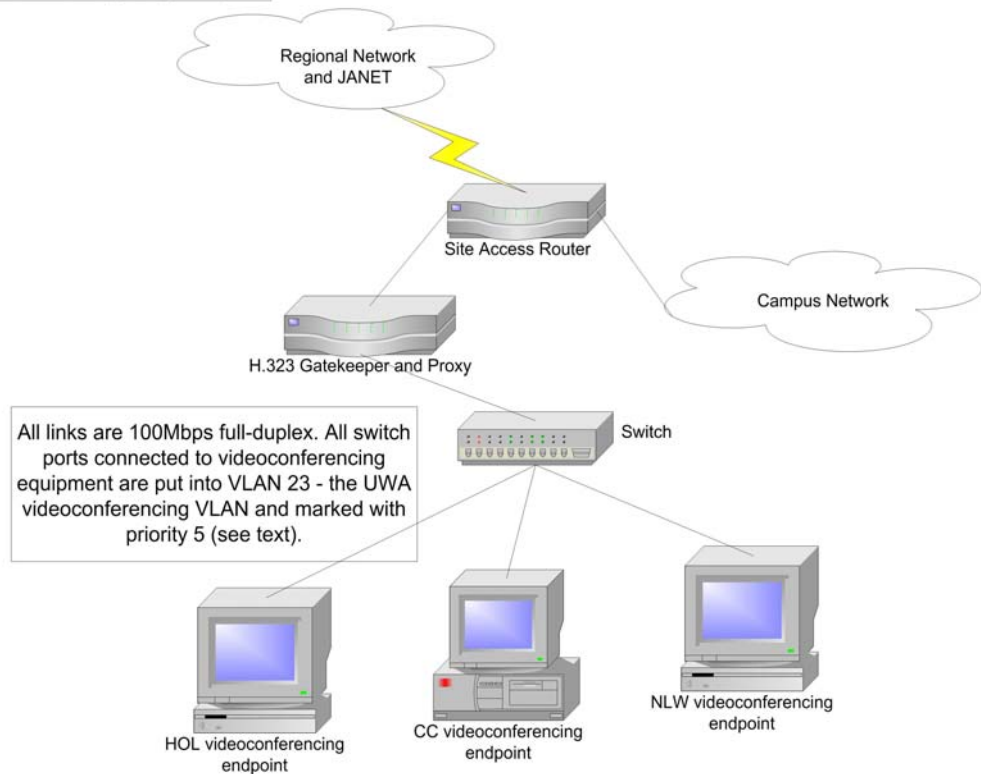
7.1 Description

All studios are on the same Class C IP subnet, which has been designated for H.323 endpoints. The HOL, CC and NLW studio CODECs are physically directly connected through patched fibre to ports on a dedicated switch. This is a Cisco® Catalyst® 3524, a 24-port switch. Another port on the switch connects to a Cisco® Catalyst® 3662 router (the MCM router) that is running the H.323 gatekeeper and H.323 proxy based on Cisco® MCM. This is a security measure rather than a QoS one.

Many of the techniques for offering QoS support to real-time applications described in this report are deployed at UWA. There are directly connected studios and also videoconferencing endpoints that are connected across the campus LAN with VLAN and CoS support, in most cases via trunked VLANs.

All links shown in Figure 10 (*overleaf*) are 100Mbit/s full-duplex – set on both the switch end of the link and on each of the videoconferencing endpoints. The switchports are also set to mark traffic entering them with Class of Service 5 – the standard CoS marking for video traffic.

Figure 10: Directly connected videoconferencing endpoints at UWA.

Aberystwyth - Directly Connected Videoconferencing Endpoints

The switch configuration uses the following commands to achieve this:

```

! Enter interface configuration mode for the relevant Ethernet
! physical port.
interface FastEthernet0/1
!
! The description gives a reminder of what is attached on this
! interface
description Codec aber-penglais-hol
! The speed and duplex settings are manually and explicitly set
! to match those of the codec (which is also manually and
! explicitly set to the same values).
duplex full
speed 100
! The switchport command causes all frames arriving on this
! physical port to be placed in VLAN 23, the VLAN for H.323
! equipment.
switchport access vlan 23
! By default the 802.1q tags that are attached by the switch to
! frames arriving on ports that have VLANs configured do not have
! a priority value set. This command, which can have any value
! between 0 and 7, attaches a Class of Service value to those
! tags, which will be retained as the frames are passed to the
! next switch. CoS value 5 is standard for video traffic.

```

```

switchport priority default 5
! This command causes a port to enter the spanning-tree
! forwarding state immediately, bypassing the listening and
! learning states. It can be used when the switchport is
! connected to a single device or workstation, but should not be
! used if there are multiple devices attached on that port.
! This should also decrease the boot time of the connected device.
spanning-tree portfast

```

All switches that connect videoconferencing endpoints do so by placing them into a VLAN (with the exception of the videophones, which move around to different locations at the University). All of the switches that are directly connected to videoconferencing endpoints place them in VLAN 23 and assign priority 5 to inbound frames on that port. So the same set of instructions would be used at those locations.

H.323 traffic that is received by the gatekeeper/proxy and forwarded to the videoconferencing endpoints also needs to receive QoS support as far as possible. The gatekeeper and H.323 proxy are also in the same subnet and VLAN as the other H.323 equipment. A similar set of commands is issued to place the gatekeeper in VLAN 23, and to tag frames received on the gatekeeper's port with CoS priority 5 for forwarding to the videoconferencing endpoints.

7.2 Videoconferencing Endpoints Connected over Trunked VLANs

The situation is a little different where a switch is receiving frames from videoconferencing endpoints and also from other equipment, such as workstations etc. Care then has to be taken to apply the correct VLANs and priorities to the different ports.

In the next command extract, a workstation is attached on an interface and is placed in VLAN 120 – the default VLAN for staff workstations.

```

interface FastEthernet0/3
description pciaj
duplex full
speed 100
switchport access vlan 120
spanning-tree portfast

```

The next extract configures an IP phone. Currently there are a few IP telephones deployed at UWA and these are placed in the same VLAN, as they are also H.323 equipment. It should be noted, however, that the QoS settings for a discrete VoIP deployment would differ from those deployed for a videoconferencing network. In these cases manufacturers recommend different parameters and – in some cases – have developed specific commands and architectures for supporting VoIP/IP telephony in the LAN, which are beyond the scope of this report.

```

interface FastEthernet0/4
description iaaj ip phone
duplex full
speed 100
switchport access vlan 23
switchport priority default 5
spanning-tree portfast

```

The switchport that is connected to the H.323 videoconferencing endpoints is configured in the same way as the directly connected example:

```
interface FastEthernet0/6
description Codec aber-penglais-wvn
duplex full
speed 100
switchport access vlan 23
switchport priority default 5
spanning-tree portfast
```

Many irrelevant commands that are in the Cisco® running configuration of this switch are omitted here. The next commands that are of interest are those that trunk the various VLANs:

```
interface GigabitEthernet0/1
switchport trunk encapsulation 802.1q
```

Trunking is a way of carrying traffic from several VLANs over a point-to-point link between two devices. The trunk combines a number of VLANs over a particular link. There have to be corresponding instructions at the other end of the link – in other words, the VLAN trunk link needs to be set up explicitly at both ends of each link. The command line at the other end of the trunk link will be identical. Because there are standard and proprietary methods of configuring VLAN trunks (i.e. there are proprietary alternatives to 802.1q), the command above is used to specify the mode of trunk signalling to be used. In this case the mode is 802.1q, and is the open standard method and not a proprietary one.

```
switchport mode trunk
```

There are various arguments to the switchport mode command – in this case the argument trunk has the effect of forcing the interface to become a VLAN trunk irrespective of the configuration of the device at the other end of the link.

```
switchport priority extend trust
```

This command tells the switch to trust, act on and pass on the priority or CoS levels found on the 802.1q tagged frames (i.e. those in VLANs) arriving on this physical port. This command can also be used when a H.323 terminal (phone or video videoconferencing endpoint) is the only device attached on a particular port, and the terminal has configurable QoS/CoS parameters.

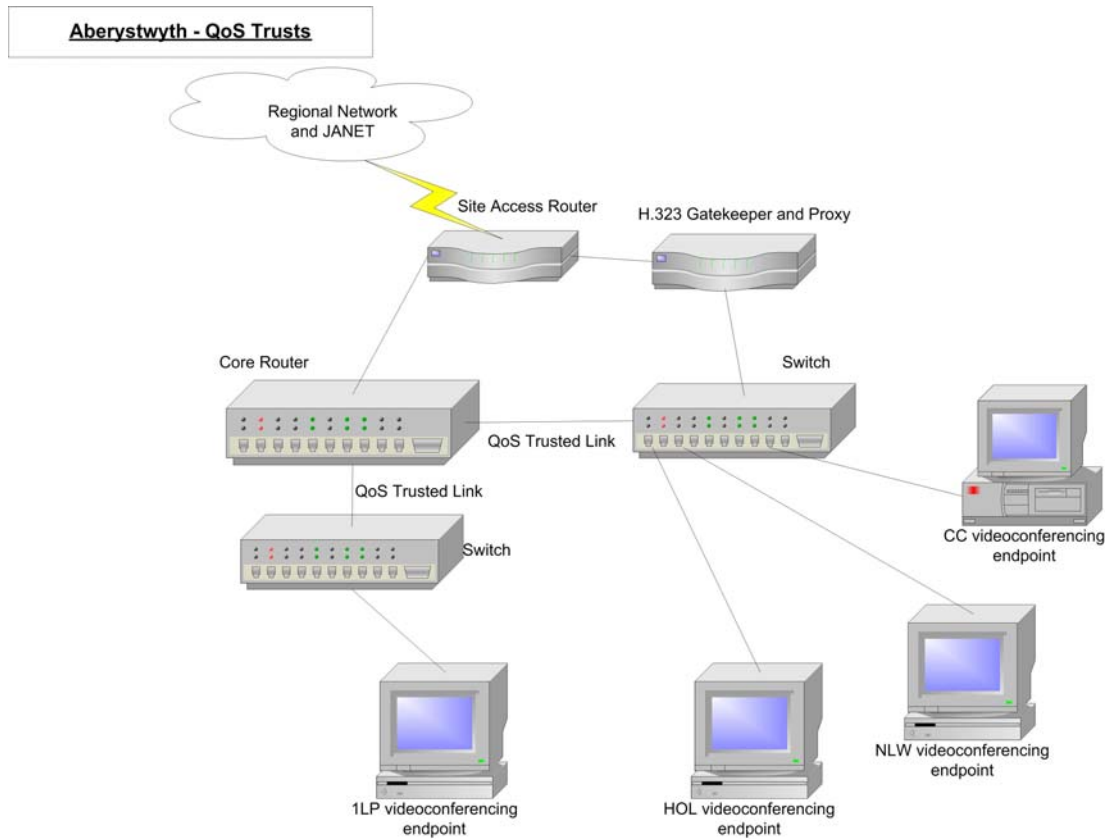
```
udld enable
```

This command enables the ULDP (Unidirectional Link Detection Protocol) which is necessary whenever STP is running. The ULDP needs to be enabled on both ends of a link to work, but uses information from Layers 1 and 2 to ensure bi-directional physical links are still capable of sending frames in either direction (even though they may be blocked from doing so). If it detects that a link has become unidirectional, it disables the port completely in order to stop possible loop and/or ‘black holes’. For more information please see (Cisco, 2003b).

The above commands are run on the switch (and port) that the 1LP CODEC is attached to. This means frames issuing from the 1LP CODEC are prioritised in the switch(es) between it and a central Layer 3 router. The Layer 3 router sends non-H.323 service packets towards the SAR,

but in-service H.323 packets are routed towards the gatekeeper. Because of the network of trust (see Figure 11 *below*) that has been built, the frame will continue to receive priority treatment as it traverses the switch(es) between it and the gatekeeper/proxy.

Figure 11: UWA H.323 QoS trust network.



It is worth noting that, depending on the way a videoconference call is dialled from a videoconferencing endpoint, the H.323 traffic can flow in two different ways. H.323 calls set up using E.164 Global Dialling Scheme numbers (Williams, 2002) will use the H.323 proxy and all the H.323 traffic in that call will go through the H.323 proxy.

Conversely, if a videoconference call is dialled using an IP address or hostname, the H.323 traffic that is part of that call will not use the H.323 proxy and will follow the normal, default traffic path into and out of the campus. This is due to the different way that E.164 GDS numbers are resolved compared to IP addresses that need no resolution or hostnames that use simple DNS name to IP address mapping. Other non H.323 traffic, such as web page accesses from a PC-based videoconferencing endpoint, will also be routed out to JANET in the normal way.

The inward facing port on the gatekeeper and H.323 proxy is part of the same subnet as the rest of the H.323 equipment on the campus. The outward facing port (i.e. the link to the SAR) has an IP address seen by external internet hosts. The proxy acts as an IP/IP gateway between these two nets, so all GDS dialled set-up, control and media packets are routed through this proxy, and receive (or will receive) QoS support in the LAN, the MAN and the JANET core, based on the parameters of these packets. Out-of-service calls, dialled by IP address, do not use the H.323 proxy and so are routed in a Best Efforts fashion, taking their chances with all of the other Internet packets taking this route.

8 List of References

- CISCO, 2001. *Deploying QoS for Voice and Video in IP Networks*. Cisco® Networkers 2001 Conference presentation VVT-213. Cisco®.
- CISCO, 2003a. *How LAN Switches Work*. [WWW 6 February 2004]
<http://www.cisco.com/warp/public/473/lan-switch-cisco.pdf>
- CISCO, 2003b. *Understanding and Configuring the Unidirection Link Detection Protocol Feature*. [WWW 6 February 2004]
<http://www.cisco.com/warp/public/473/77.pdf>
- FLANNAGAN, M., Froom, R., Turek, K. 2003. *Cisco® Catalyst® QoS: Quality of Service in Campus Networks*. USA. Cisco Press.
- LONG, C. 2002. *IP Network Design, Part 4: Campus LAN Design*. [WWW 6 February 2004]
http://searchnetworking.techtarget.com/originalContent/0,289142,sid7_gci803121,00.html
- WILLIAMS, S. 2002. *H.323 Global Dialing Scheme (GDS)*. [WWW]
<http://www.wvn.ac.uk/support/h323address.htm>

9 Further and Recommended Reading

Academic Networks

- UKERNA Development Programme – includes work on QoS, Multicast, IPv6 etc.
<http://www.ja.net/development/>
- The JANET VTAS (Video Technology Advisory Service) provides advice to the UK education community through publications and their web site:
<http://www.video.ja.net/>
- The TERENA (Trans-European Research and Education Network Association) TF-NGN (Task-Force on Next Generation Networking) looks at lower layer issues and, together with DANTE (Delivery of Advanced Network Technology to Europe), who manage GÉANT (the Gigabit European Academic Network), has done work on QoS provisioning and testing across GÉANT as well as issues related to inter-domain QoS provisioning, Service Level Agreements and Service Level Specifications.
<http://www.terena.nl/tech/task-forces/tf-ngn/>

General Networking

- Baccala, B. 1997. *Programmed Instruction Course*.
<http://freesoft.org/CIE/Course/>
(6 February 2004)

VLANs

- 3Com® 4400 Switch Implementation Guide.
Full configuration guide covering VLANs, QoS etc.
<http://support.3com.com/infodeli/tools/switches/ss/fast/dua1720-3baa03.pdf>
- Creating and Maintaining VLANs on Cisco® 3550 Switches.
<http://www.cisco.com/univercd/cc/td/doc/product/lan/c3550/1214ea1/3550scg/swvlan.htm>

QoS Provisioning

- 3Com® 4400 Switch Implementation Guide.
Full configuration guide covering VLANs, QoS etc.
<http://support.3com.com/infodeli/tools/switches/ss/fast/dua1720-3baa03.pdf>
- Cisco® QoS configuration links page – covers provisioning QoS on Cisco® 2950 through to Cisco® 6000 switches.
http://www.cisco.com/cgi-bin/Support/browse/psp_view.pl?p=Technologies:Quality_of_Service_LAN&s=Implementation_and_Configuration
(Requires Cisco.com username and password)
- Cisco® Catalyst® 6000 family Configuration Guide.
http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_7_3/config_gd/qos.pdf

Spanning Tree

- Understanding Spanning-Tree Protocols from Cisco®.
http://www.cisco.com/univercd/cc/td/doc/product/rtrmgmt/sw_ntman/cwsmain/cwsi2/cwsiug2/vlan2/stpapp.htm
- Spanning Tree Tutorials.
<http://tutorials.beginners.co.uk/read/category/recent/id/451>

Subnetting

<http://freesoft.org/CIE/Course/Subnet/>

Tell us what you think

Technical Guides are produced and published by UKERNA for use within the JANET Community. We welcome your comments on all aspects of this document and on any other UKERNA publication. Please direct feedback to JANET Customer Service, at the address below, or e-mail us directly at:

documentation@ukerna.ac.uk

UKERNA manages the networking programme on behalf of the higher and further education and research community in the United Kingdom. JANET, the United Kingdom's education and research network, is funded by the Joint Information Systems Committee (JISC).

For further information please contact:

JANET Customer Service
UKERNA
Atlas Centre, Chilton, Didcot
Oxfordshire, OX11 0QS

Tel: 0870 850 2212
Fax: 0870 850 2213
E-mail: service@janet.ac.uk

Copyright:

This document is copyright The JNT Association trading as UKERNA. Parts of it, as appropriate, may be freely copied and incorporated unaltered into another document unless produced for commercial gain, subject to the source being appropriately acknowledged and the copyright preserved. The reproduction of logos without permission is expressly forbidden. Permission should be sought from JANET Customer Service.

Trademarks:

JANET[®], SuperJANET[®] and UKERNA[®] are registered trademarks of the Higher Education Funding Councils for England, Scotland and Wales. The JNT Association is the registered user of these trademarks.

3Com[®] is a registered trademark of 3Com Corporation. All rights reserved. All other company and product names may be trademarks of their respective companies.

Cisco[®] and Cisco[®]Catalyst[®] are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

Microsoft[®] is a registered trademark of Microsoft Corporation in the United States and/or other countries.

Novell[®] is a registered trademark of Novell, Inc. in the United States and other countries.

Polycom[®] and Viewstation[®] are registered trademarks of Polycom, Inc. in the United States of America and various countries.

Tandberg[®] is a registered trademark in the US and certain other countries. All other trademarks are property of their respective owners.

Disclaimer:

The information contained herein is believed to be correct at the time of issue, but no liability can be accepted for any inaccuracies.

The reader is reminded that changes may have taken place since issue, particularly in rapidly changing areas such as internet addressing, and consequently URLs and e-mail addresses should be used with caution.

The JNT Association cannot accept any responsibility for any loss or damage resulting from the use of the material contained herein.

Availability:

Printed copies of this document may be obtained from JANET Customer Service at the above address.

Further PDF copies of this document are also available at: http://www.ja.net/documents/technical_guides.html



© The JNT Association 2004

JISC