Campus Traffic and Common Network Issues

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. t 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).

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.

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.

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.

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.

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.

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.
