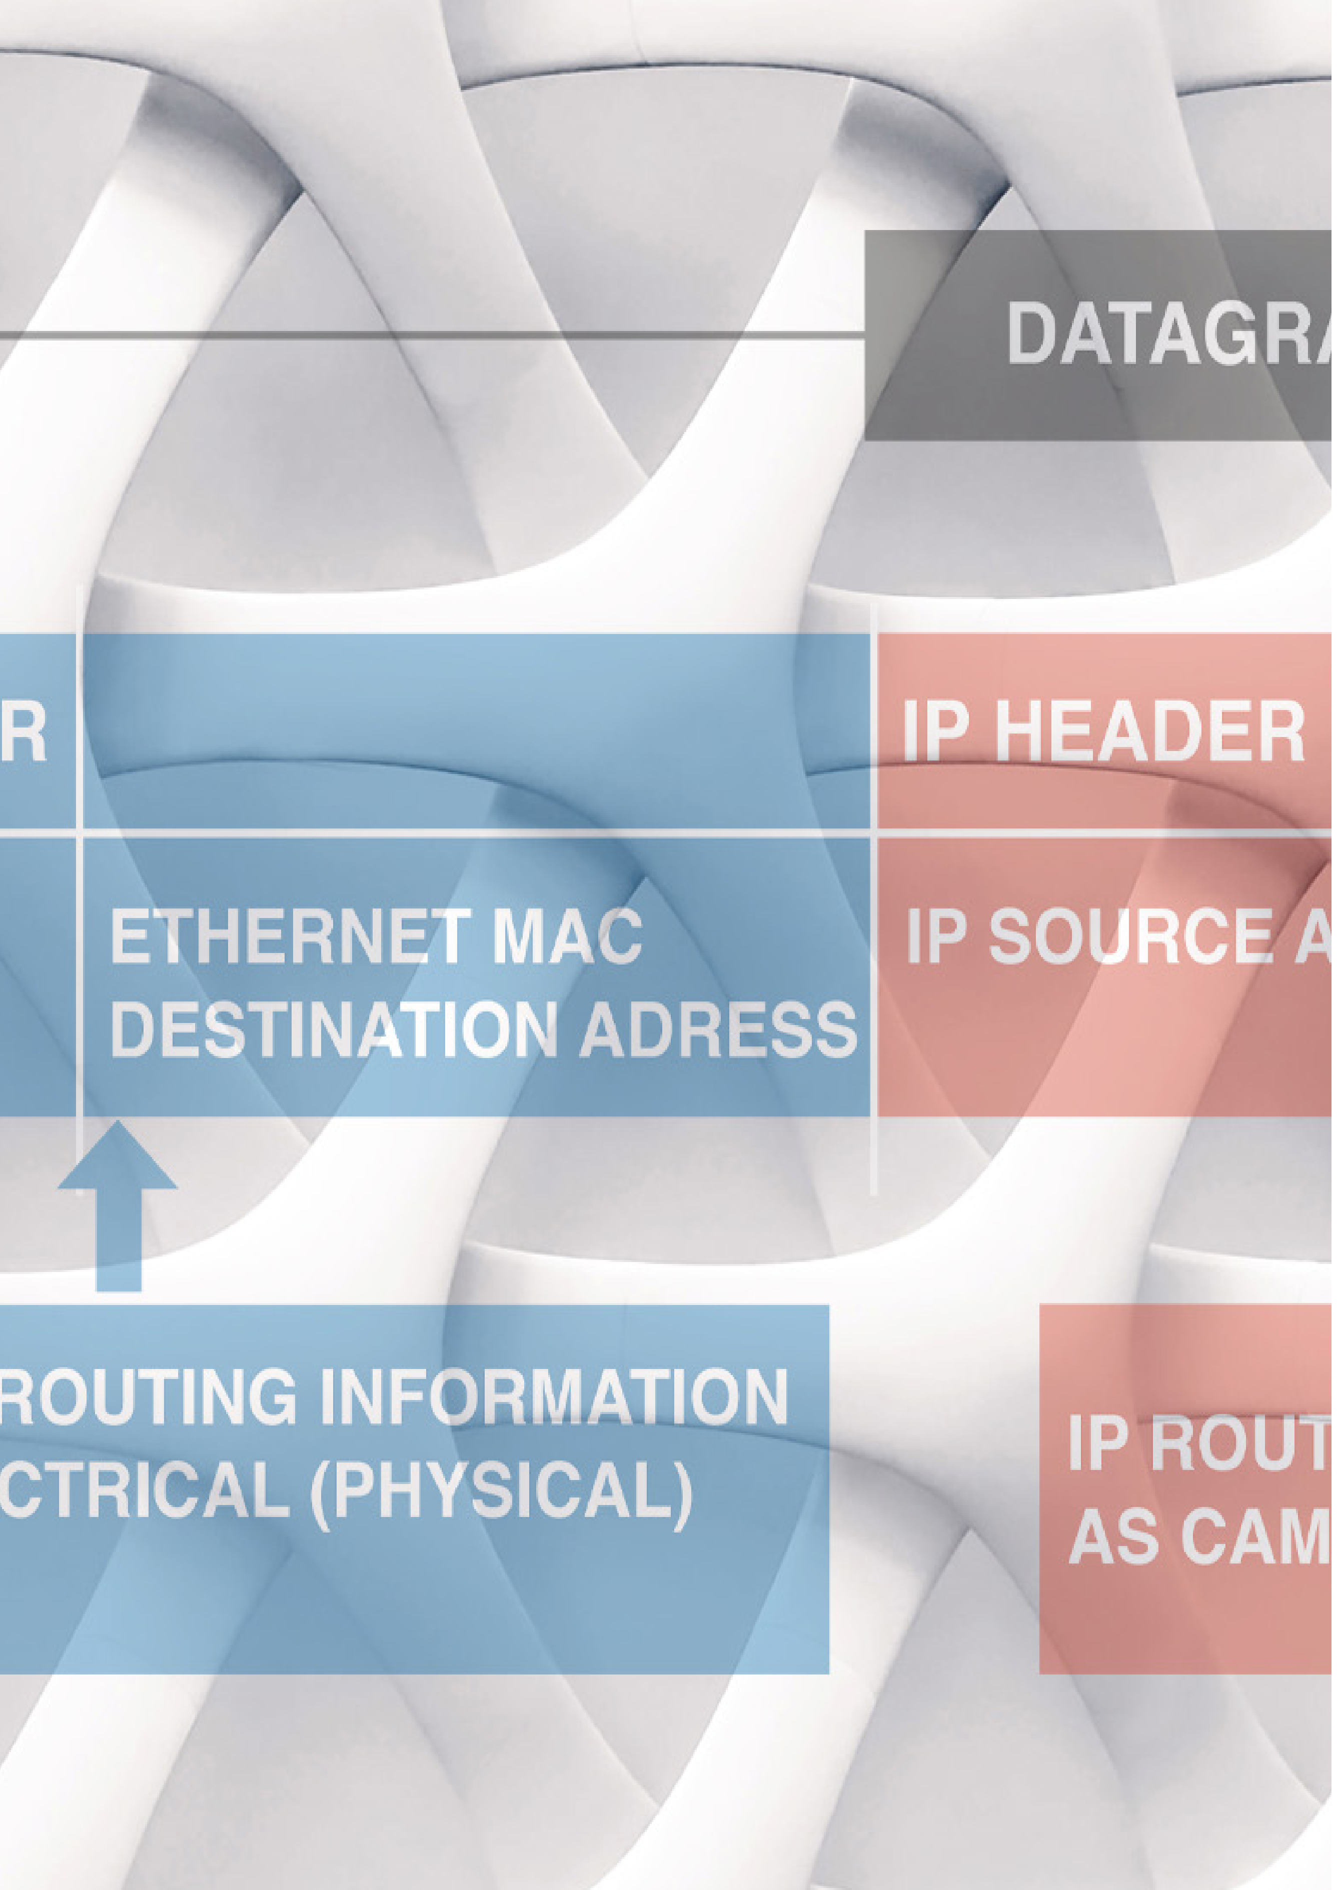

THE BROADCAST BRIDGE
Connecting IT to Broadcast

Understanding IP Broadcast Production Networks

By Tony Orme

A Themed Content Collection from The Broadcast Bridge

Themed Content Collection

Introduction

By Tony Orme. The Broadcast Bridge.

Understanding IP Broadcast Production Networks is a Themed Content Collection which provides the basic building blocks of knowledge required to understand how IP networks actually work in the context of broadcast production systems.

Much of the content in this book was originally published by The Broadcast Bridge way back in 2017. Back then IP for broadcast production was new and everybody was evaluating it's value and its practicality.

A mere six years on and IP is everywhere. The versatility and scalability it brings has moved it into the mainstream and established it as the central nervous system of broadcast.

Most new facilities are IP native, many broadcasters are running hybrid SDI-IP systems, and there are very few engineers for whom IP is not a key part of daily life.

The Broadcast Bridge has published hundreds of articles on IP and continues to push the boundaries, helping broadcasters evaluate the next generation of infrastructure - where IP is enabling cloud and microservices to create the next step change in our technology and workflow.

Throughout the evolution of IP for broadcast this series of articles has continued to draw significant traffic from search engines and our own site search. This series has become a reference work for engineers and operators who need to understand the fundamental principles and technology of IP.

The content has been updated and edited to reflect the changing times. We hope it continues to serve the industry well.

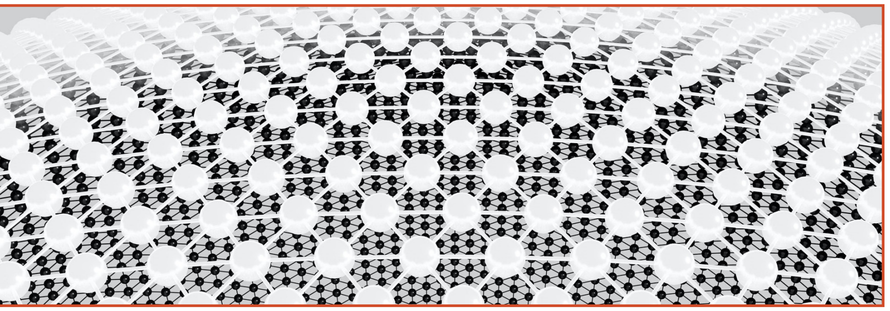

Contents

1

Basic Principles

The fundamental principles and terminology of IP based broadcast production systems.

Network timing requirements, differing working practices and protocols, and data integrity all help to deteriorate communication between broadcast and IT engineers.

Timing in broadcast is tightly defined and a thorough understanding of legacy television systems is required. IT engineers use asynchronous full duplex systems, expect there to be network failure, and use protocols that effectively slow transmission to make sure data has been accurately delivered. Broadcasters use synchronous one direction connectivity and assume the network is as robust and reliable as SDI.

In this book, we look at networks from a broadcast engineer's point of view, giving a better understanding of core IT concepts and enabling them to communicate with colleagues in the "IT department".

To fully understand IT networks, we must understand the problem we are trying to solve; a network is needed to allow users to exchange data predictably, reliably and securely, and provide control of one computer over another. This is true of PC's, servers, IP-camera's, production switchers and control panels, and the more secure and reliable a system, the more complex it becomes.

A network must be resilient, fast, and reliable to give the best user experience. To explain the roles of routers and switches we start with a basic network of four PC's and two servers connected in a simple IP over Ethernet network using CAT5.

Ethernet has three forms of physical interface: coaxial, twisted pair and fiber optic. They all send the same type of packets of data but differ in their duplex as twisted pair can send and receive data at the same time, but coaxial and fiber optic cannot. Transmission speeds are faster on fiber optic and coaxial.

Few computers use coaxial connectivity as twisted pair is cheaper and more robust. Fiber optic tends to be reserved for high bandwidth switch and router connection due to its higher cost and fragility.

A hub with twisted pair infrastructure (CAT5) could be used in a simple network.

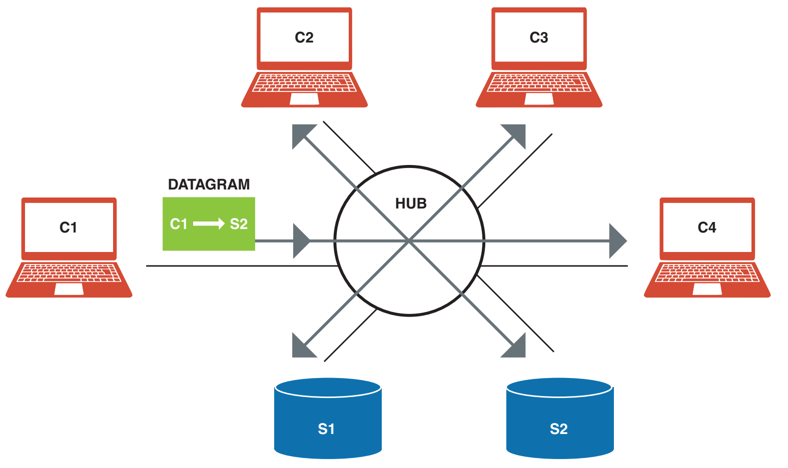

Figure 1 - SIMPLE HUB NETWORK – A datagram sent from C1 to S2 will be re-sent to all computers and servers on the network potentially causing security and congestion issues.

The hub is like a distribution amplifier allowing mapping of one-to-many transmit and receive pairs. The hub has no intelligence and will route a packet received on one port directly to all its other ports.

In a hub network, security becomes a problem as all users would be able to see data being exchanged between each other's computers and servers. For example, all users would receive transactions associated with the finance server.

Computer network cards receive all datagrams on the connected network and will usually discard those not intended for them. With the right software, it's easy to decode the datagram and view restricted and sensitive financial transactions. This is true of all the systems running on any of the servers.

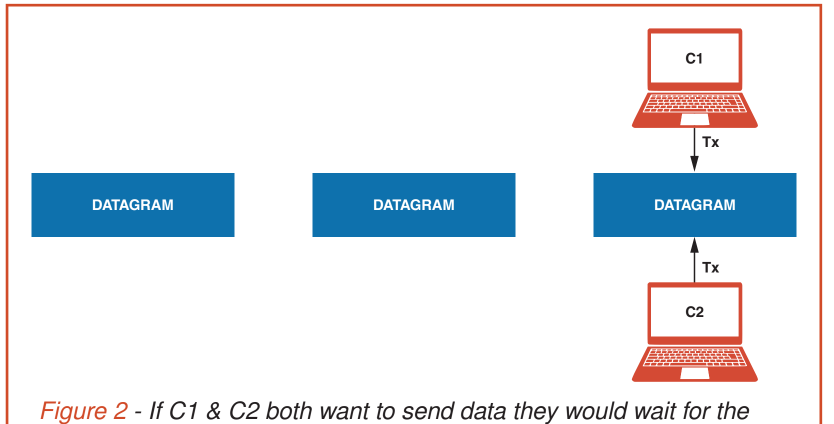

Figure 2 - If C1 & C2 both want to send data they would wait for the first available space on the transmission line, potentially sending at the same time and corrupting their data.

Lost packets of data occur as network traffic increases and the physical switch and router links become quickly saturated, and this is further exasperated by micro-bursts of data that can overflow egress buffers. Protocols such as TCP (Transmission Control Protocol) can remedy lost packets but do so at the expense of increased and variable latency. This is one of the reasons standards such as SMPTEs ST2110 uses UDP (User Datagram Protocol) as it operates a fire-and-forget transmission system, resulting in predictable and low latency. However, when using UDP, lost packets cannot be easily recovered.

Ethernet is a packet switched system, each PC will monitor the transmit line and wait for a gap so it can send its own packet. Although the packets are of a fixed size, the frequency with which they are sent is random across all the connected computers on the network. Another computer may be listening at the same time waiting for the same space, and two or more computers could try and simultaneously access the transmit pair resulting in a collision and packet loss, and slow response for the user.

A network router or switch will protect against collisions and congestion and is one of the reasons routers and switches are used, other reasons are to provide resilience and security. Switches send packets at the Ethernet packet level (layer 2) and routers route packets at the Internet Protocol level (layer 3).

In the ISO seven-layer model IP packets are encapsulated by the layer 2 Ethernet frames. This might seem like an unnecessary overhead; however, the IP protocol is independent of the transmission network and abstracts the data away from the hardware limitations of Ethernet. It's entirely possible, during the lifetime of an IP packet, that it will be routed over non-Ethernet networks such as ATM (asynchronous transfer mode) or WiFi. With IP, we need not be too concerned with the medium the data is travelling on.

Each Ethernet card in a PC or IP-camera has its own unique hard coded address called the media access control (MAC address). Each camera can be configured to have a unique IP address, so a faulty camera can be replaced with the same IP address. The MAC address will have changed but the address resolution protocol (ARP) in the routers would detect this and reconfigure themselves.

Managed Ethernet switches provide a better solution but have limited capability. The switch is configured with the MAC address of each computer connected to its ports and will send traffic only intended for the associated computer thus reducing network traffic on each connection. For these reasons Ethernet switches tend to be used in fixed high-speed applications such as core network switches and head of rack topologies. They are faster as there is less information to process in the Ethernet datagram header compared to an IP header. For example, there is no "time to live" value to be updated.

IP addressing schemes offer greater flexibility and allow administrators to specify their own IP number schemes. Security is improved as routers can be configured to make sure finance transactions only go to authorized computers and IP-camera's only send their pictures to monitors and production switchers in the studio. Programs such as Ping can be blocked to stop hackers detecting computers and attacking them.

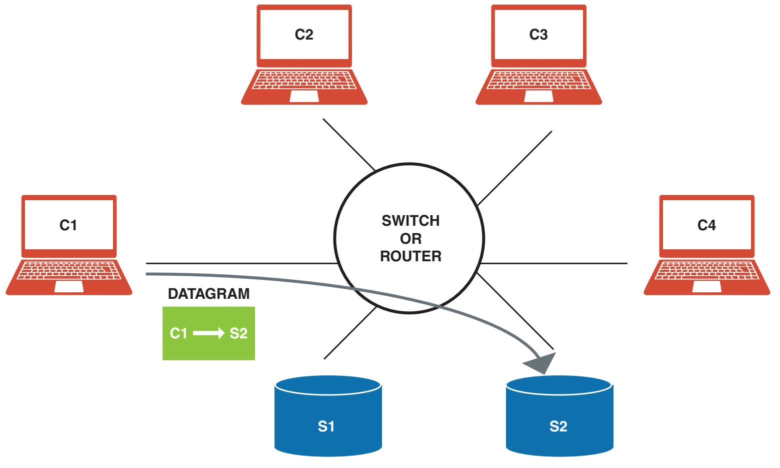

Figure 3 - SWITCH OR ROUTER NETWORK – a datagram sent from C1 to S2 will only be sent to S2, improving security and data speeds.

Automatic routing algorithms provide resilience by detecting a broken link and sending the data via a different route. Multi-path links can be used between studios and outside broadcast consisting of different types of media such as fiber optic and satellite. Users are unaware that routers have switched to a different path when a link breaks.

Even in a simple network routers and switches improve network speeds and security, and routers become essential when resilience is needed.

2

Routers & Switches

How Routers & Switches reduce traffic congestion and improve security.

Routers and switches reduce traffic congestion by sending packets and frames only to the devices that require the data and not spraying the whole network with them. As well as reducing congestion this will help improve security as administrators can limit which hosts can exchange data; IP-cameras, IP-production switchers, and computers etc.

An IP packet consists of two parts, the header, and the payload. In the case of video over IP, the camera will break its video stream into smaller chunks until they fit into the payload of an IP packet. The header consists of the source IP address, that is the address of the device sending the packet such as a camera, and the destination address of where the camera wants to send its' video stream to, for example a production switcher.

If we assume that the network we are using in the studio is CAT6/Ethernet, the camera will have its' own media access control (MAC) address and the production switcher will also have its own MAC address. These are unique Ethernet addresses that are hard coded into the devices during manufacture and are different for every Ethernet enabled unit that leaves the production line.

A network often consists of routers and switches where switches distribute traffic within the locality of the connected devices, and routers route packets between localized networks. For example, studio-1 consisting of four cameras and a production switcher will only distribute data within studio-1, however, if camera-1 from studio-1 needs to be sent to studio-2, then a router is required.

Ethernet networks are defined as a network which share the same Ethernet broadcast address. It is possible to have multiple switches cascaded across many studios but whenever an Ethernet broadcast message is sent, for example, when using the ARP protocol, then every device in every studio will receive this message, thus creating unnecessary traffic and potentially increasing latency. To stop this from happening, routers are often used to separate the networks, so the broadcast traffic is only sent to the devices within the locality of a studio.

VLANS also solve this problem and are discussed in a later chapter.

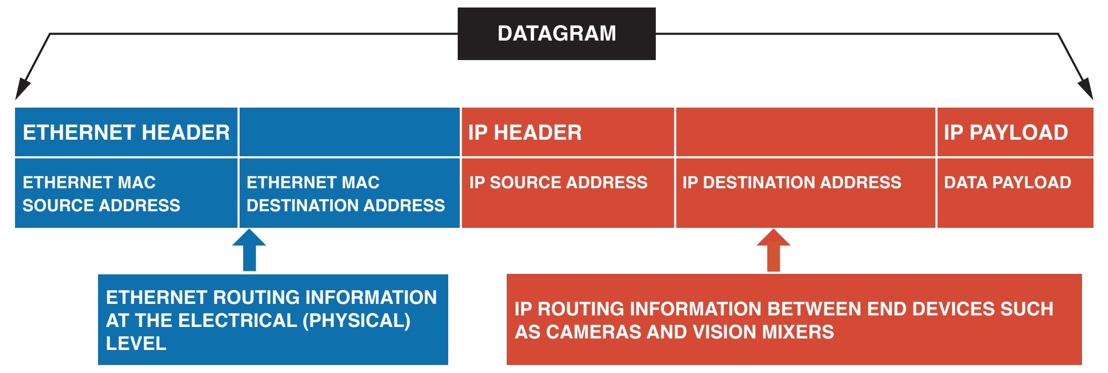
DATAGRAM
ETHERNET HEADER
IP HEADER
IP PAYLOAD
ETHERNET MAC SOURCE ADDRESS
ETHERNET MAC DESTINATION ADDRESS
IP SOURCE ADDRESS
IP DESTINATION ADDRESS
DATA PAYLOAD
ETHERNET ROUTING INFORMATION AT THE ELECTRICAL (PHYSICAL) LEVEL
IP ROUTING INFORMATION BETWEEN END DEVICES SUCH AS CAMERAS AND VISION MIXERS

Figure 1 - How an Ethernet datagram encapsulates the IP datagram.

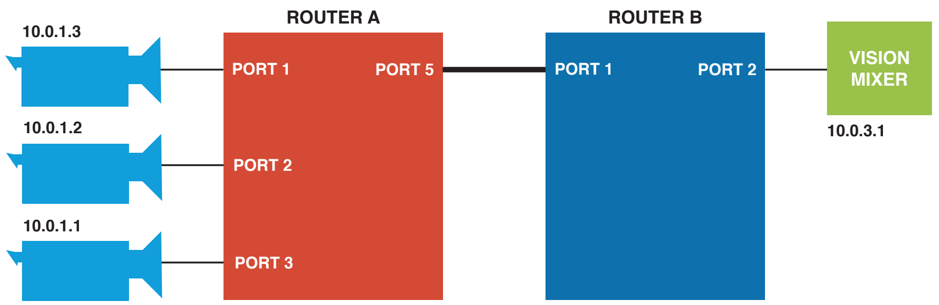

Figure 2 - A simple network using layer 2 switches

In diagram two we have a simple studio network where a production switcher is connected to three cameras via two different layer-2 switches. It is possible to connect all these devices together on one switch but that would leave us with a single point of failure and probably over-load the switch due to the high bandwidths video-over-IP demands.

The router uses a combination of the IP address and netmask, called the Network-ID to route packets to other routers. A typical network ID in the class-less system would be 10.1.1.0/24. The "24" refers to the subnet mask and the 10.1.1.0 is the network. In this instance the router would route any packet in the address range 10.1.1.000 to 10.1.1.255, as it only looks at the first 24 bits of the IP address, and each number within the dot is 8 bits. As another example, a network ID such as 10.1.1.0/8 will contain all hosts with an IP address range from 10.000.000.000 to 10.255.255.255. Subnets are often aligned with Ethernet networks to aid administration and security.

Static routers work on a next-hop system, that is each router only knows of the existence of subnets connected to its ports, this simplifies the network design and keeps the router database to a manageable size.

In diagram two, Ethernet is being used as the physical interface and Camera-1 will have to set the source and destination MAC addresses of the Ethernet frame leaving camera and entering the switch port. A MAC address is a unique number issued by the IEEE for every piece of equipment that is manufactured. When the IP packet leaves camera-1 it must also set the Ethernet MAC source and destination addresses for the encapsulating Ethernet frame, as well as those for the IP datagram.

A system called address resolution protocol (ARP) is instigated by camera-1 which broadcasts an Ethernet message saying "who has IP address 10.0.3.1 (the production switcher)? And what is your Ethernet address?". The broadcasting of ARP messages is restricted to local subnets and routers to stop the entire network being flooded with ARP query messages.

At this point switch A answers the ARP query from camera-1 with the address of the MAC address for port-1. Camera-1 sets its own MAC address as the source MAC address, and the destination address is the MAC address of port-1 on switch A. Port 5 on switch A will have a different MAC address than ports 1 to 3 as it's connected to a different network segment.

Each switch is able to build a database of the connected devices so will substitute the destination MAC addresses as required. The switches in diagram one provide an Ethernet network so, if the studio technical director wants to send camera-1 from studio-1 to studio-2, then a router with the appropriate forwarding tables will need to be provided. This stops unnecessary traffic from studio-1 flooding studio-2, and vice versa.

HEADER ADDRESSES	ROUTER A, PORT 1	Router A, Por t 5 to Router B, Port 1	Router B, Port 1 to Vision Mixer
MAC Source	MAC for Camera 1	MAC for Router A, Port 5	MAC for Router B, Port 2
MAC Destination	MAC for Router A, Port 1	MAC for Router B, Port 1	MAC for Vision Mixer
IP Source	Camera 1	Camera 1	Camera 1
IP Destination	Vision Mixer	Vision Mixer	Vision Mixer

Figure 3 - showing the changing Ethernet Addresses and fixed IP addresses as they datagram moves from the camera to the vision mixer.

Throughout the whole process of sending a packet from camera-1 to the production switcher via switch A and switch B, the source and destination IP addresses did not change at all. However, the source and destination Ethernet MAC addresses changed at each network segment.

It's entirely possible that the cameras on switch A were at the Superbowl stadium, with switch B and the production switcher at the studio in New York, with an IP satellite link connecting them. The source and destination IP addresses did not change, but the MAC and physical connectivity mapping automatically changed as the frames moved between devices. From an IP point of view, we do not know or care how the packet travelled between switch A and B, the underlying complexity of the physical, electrical and optical routing was abstracted away from us. And this is one of the greatest strengths of IP.

3

Resilience

How distance vector routing simplifies networks and improves resilience.

Routing protocols allow distributed routers to communicate with each other so that IP packets can be sent to different subnets in the most efficient way possible. One of the fundamental concepts of distance vector routing is that each router only has knowledge of the neighboring routers it's attached to. This greatly simplifies the network administration as each router does not need an understanding of the whole network topology.

Figure one shows a simple distributed studio network connecting cameras to a production switcher and monitors from different geographical dispersed layer-2 networks. Three routers are used to demonstrate the distributive approach of networks and resilience. For example, the production switcher has two distinct paths back to the cameras; link A-B, and link B-C to A-C. Deciding which path to take and then what happens in the event of a link failure is the subject of this chapter.

Distance vector routing (DVR) is a protocol which allows routers to decide which is the most optimal route to take to reach the destination host. Intuitively, we might think that if a datagram was sent form camera-1 to the rack monitors we should use link A-C, that is router A then C. But from the metrics we can see that route A-B and B-C has measure 4 (3 for A-B, and 1 for B-C), and route A-C has measure 5. In this example, the most optimum route is A-B then B-C.

The measures are calculated using several different protocols; routing information protocol (RIP), and interior gateway routing protocol (IGRP) being another. RIP is one of the oldest DVR protocols and uses hop counting as its measure. In diagram one, if RIP was used instead of IGRP, then a measure of one exists between router A and B, a measure of one between routers B and C, and a measure of one between routers C and A. So, to travel from router A to C, we

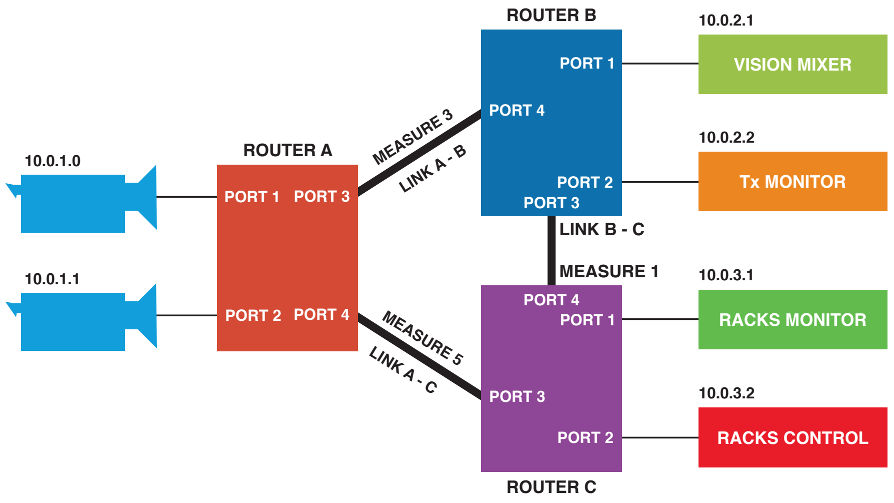

Figure 1 - A simple distributed studio network with cameras, monitors and vision switcher, all in geo-graphically different areas.

have a measure of two for A-B then B-C, and one for A-C. This implies A-C is the fastest route but as RIP only counts hops and does not take into consideration any propagation delays between the hops, the measure may not be accurate. Consequently, IGRP improves on RIP by adding node and propagation delays to the measure. Link B-C might be fiber and A-C might be satellite, so B-C is clearly faster and has a lower metric.

When the network is first switched on, each routers' DVR is zero and has no knowledge of any of its neighboring routers. As each router initializes it looks at its routing table and advertises the subnets it can see to its neighbors. Each router continues to periodically advertise the subnets it can see to its neighboring routers. As the routers build up knowledge of each other's routes, and build up knowledge of the measures, they are able to select the best route to send a packet.

If camera-1 sends video to the racks monitors, it will first break the video into small IP datagrams. The IP source address will be 10.0.1.0 and the IP destination address will be 10.0.3.1. Router A will have determined the most efficient route to send the datagram will be along link A-B. The IP source and destination addresses stay the same, but the Ethernet headers will have a source address of port 3 on router A and a designation address of port 4 on router B.

When router B receives the datagram it knows, from the metrics in the routing table, populated by the distance routing vector algorithm, that the most efficient route to 10.0.3.1 is via router C, therefore it will keep the IP source and destination addresses the same, but change the Ethernet source address to the MAC of router B port 3, and the destination MAC to port 4 on router C.

If link A-B was to fail, then router A would eventually realize that there was no traffic or control messages coming from router B, and also any packets sent by router C to either router A or the production swithcer would time out. Router C would remove the route on port 4 from its routing table and send error messages to its neighbors advising network timeouts on link A-B. From router A's point of view the next best metric to the racks monitor is 5, the link A-C. It would automatically start sending its camera-1 packets along link A-C. The network administrator would be alerted to the lost error messages reported by router C and take action to fix the link.

As networks increase in complexity the measure process can take time to reach its optimum routing, this is referred to as convergence.

There are some problems with distance vector routing especially with infinite loops. The count-to-infinity problem is one of these and suffers from the fact that new routes are advertised quickly, but problems are detected slowly, and the routers become unstable because a route that is being advertised as valid, is in fact failing.

Although DVR's use minimal resource and require basic administrator knowledge, they do have some other disadvantages namely the entire routing table is sent every 30 to 90 seconds taking up valuable bandwidth.

Link state protocol is the newest edition to routing protocols and overcomes some of these limitations by only sending changes in routing tables, and updates are triggered by events such as a subnet failing, instead of being sent periodically resulting in reduction of bandwidth used. They also allow routers to gain a greater knowledge of the whole topology of the network, increasing network efficiency and resilience. However, they are extremely complicated and need advanced administrator knowledge to make them work and keep them maintained

Distance Vector	Link State
Protocols - RIP, RIPv2, IGRP, EIGRP	Protocols - OSPF, ISIS
Routers advertise route and metric information directly to their neighbors.	Routers communicate with all other routers on the network to build knowledge of the topology of the whole network.

Distance = metric
Vector = direction (interface of subnet)
	Link state = interface connections (links) to other routers and networks.

Best for simple flat designs with minimal administrator knowledge and where fast convergence times are not required.	Best for large complex hierarchical designs requiring very fast convergence times. Requires advanced administrator knowledge.
Send copies of the routing table every 30 to 90 seconds.	Much more efficient as sends triggered updates when an event occurs such as adding a router or a link failing.

Figure 2 - The main differences between distance vector and link state routing protocols.

4

Host Configuration

All devices on a network are referred to as ‘hosts’ and they all need to be configured correctly.

IT convention refers to user equipment connected to routers and switches as hosts. This could be a desktop computer, laptop, camera, production switcher or sound console for example. All these devices share similar configuration in the way they are connected to IP/Ethernet networks.

There are two ways to set up a computer on a network, either manually, or using automatic configurations. The two key parameters that need to be configured are the IP address details, and the details of domain name system. DNS is a server service to associate web names with IP addresses and is not entirely relevant to cameras or production switchers.

IP configuration can be automated by using Dynamic Host Configuration rotocol. This is a service running on a network server which automatically provides an IP address for a host when it boots up. This is generally not used for broadcast kit as we want to be able to identify cameras and production switchers through their IP addresses. It’s possible for a host to change its IP address when using DHCP without warning. For these reasons, we tend to use static manually configured IP addresses.

The three key parameters that will need configuration are the IP address, subnet mask, and default gateway. When setting the IP address care must be taken not to use an address which has already been used, doing so will result in IP ghosting and provides some very interesting results as other hosts may try to send and receive packets to a camera without realizing they’re addressing the wrong or same device.

Two conventions of subnet are in use: classful and classless. Classful is an older rarely used system often referring to class A, B, C and D subnets using discontinuous address ranges. This was considered too complicated in normal use and has largely been dropped in favor of the classless system. Classless uses a sequence of bits to define an address range. For example, 10.0.1.1/24 refers to the range 10.0.1.0 to 10.0.1.255. Host configuration of the subnet value

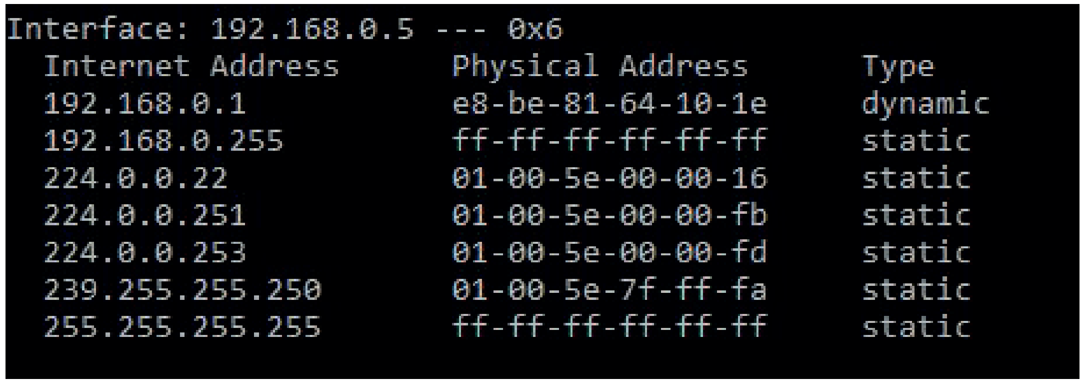

```
Interface: 192.168.0.5 --- 0x6
  Internet Address      Physical Address      Type
  192.168.0.1           e8-be-81-64-10-1e     dynamic
  192.168.0.255         ff-ff-ff-ff-ff-ff     static
  224.0.0.22            01-00-5e-00-00-16     static
  224.0.0.251           01-00-5e-00-00-fb     static
  224.0.0.253           01-00-5e-00-00-fd     static
  239.255.255.250       01-00-5e-7f-ff-fa     static
  255.255.255.255       ff-ff-ff-ff-ff-ff     static
```

Figure 1 - Screen grab of PRINT ROUTE showing the routing table of a host PC and the default gateway.

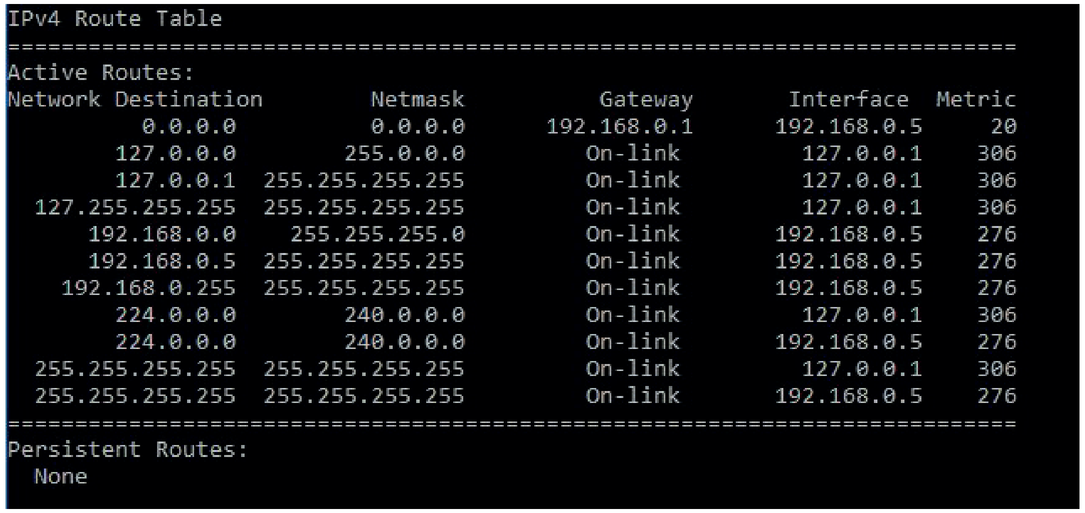

```
IPv4 Route Table
===========================================================================
Active Routes:
Network Destination        Netmask          Gateway       Interface  Metric
          0.0.0.0          0.0.0.0      192.168.0.1      192.168.0.5     20
        127.0.0.0        255.0.0.0         On-link         127.0.0.1    306
        127.0.0.1  255.255.255.255         On-link         127.0.0.1    306
  127.255.255.255  255.255.255.255         On-link         127.0.0.1    306
      192.168.0.0    255.255.255.0         On-link       192.168.0.5    276
      192.168.0.5  255.255.255.255         On-link       192.168.0.5    276
    192.168.0.255  255.255.255.255         On-link       192.168.0.5    276
        224.0.0.0        240.0.0.0         On-link         127.0.0.1    306
        224.0.0.0        240.0.0.0         On-link       192.168.0.5    276
  255.255.255.255  255.255.255.255         On-link         127.0.0.1    306
  255.255.255.255  255.255.255.255         On-link       192.168.0.5    276
===========================================================================
Persistent Routes:
  None
```

Figure 2 - Screen grab of ARP -a showing the Ethernet MAC address of some of the devices on the hosts network, specifically the default gateway.

tends to still use the dot notation. As an example, 10.0.1.1/24 would give a subnet of 255.255.255.0, and 10.0.1.1/20 would give a subnet of 255.255.240.0, that is the first 20 bits of the IP address.

The combination of the IP address and the subnet is often referred to as the network ID. In address 10.0.1.0/24, the part 10.0.1.nnn is the value of the subnet the camera is connected to. The camera will be able to send and receive packets to all devices in the range 10.0.1.0 to 10.0.1.255.

Default gateways are used by the host to work out what to do if they need to send data to a host outside of its subnet. For example, if we have a camera with address 10.0.1.0/24 and production switcher with address 10.0.2.0/24, the camera would not be able to send its packets directly to the production switcher as they are on different networks. The camera is on network 10.0.1.nnn and the production switcher is on network 10.0.2.nnn. In effect, they are physically separated.

Each host, whether it's a camera, sound console or desktop computer, will have its own routing table. This consists of a series of networks the host can see, and how it routes to the ones it can't.

When the camera needs to send its packets to the production switcher, it will look up the production switchers network ID in its own routing table and realize it doesn't have a listing for network 10.0.2.0/24. The default gateway is then used to resolve this.

When routing an IP packet, the source and destination IP addresses are kept intact throughout its whole journey from source to destination (unless Network Address Translation is used). However, the source and destination Ethernet MAC addresses do change at each node.

When the camera realizes it cannot send its packets directly to the production switcher, it will send it to the default gateway instead. It does this by finding the Ethernet MAC address of the default gateway, then setting the destination MAC address of the packets to be sent to be the MAC address of the default gateway. It's important to note, that the source and destination IP addresses do not change.

The default gateway address must be accessible by the host so will always be within the subnet of the hosts IP address.

Configuring cameras and sound desks can be frustrating for broadcast engineer's due to the lack of available tools within the equipment. For example, a PC will have command line programs such as IPCONFIG to show the connectivity of the network interface card, ARP to show the resolution of IP addresses and MAC addresses, and ROUTE PRINT to show the hosts routing table configuration. IP interfacing in broadcast kit sometimes appears to be a bit of an add-on, and the tools taken for granted in the IT world are generally not available.

Industry initiatives such as AMWA/NMOS Registration and Discovery system will help overcome this. But in the meantime, the broadcast engineer must collaborate closely with their IT colleagues to ensure broadcast IP configuration is configured correctly.

5

Audio Basics

How AVB, AES67 and proprietary solutions like Dante address the challenges of using asynchronous internet networks to distribute audio.

In television, audio is notoriously difficult to get right, the human audio- visual system can cope much better with disturbances in pictures than sound. If a television picture flashes or bangs for a few seconds, then we don't seem too concerned with it. If the audio in the same television transmission breaks up, stutters or distorts, then we are aware of the problems much more quickly, and we become easily frustrated or stressed if the problem continues.

Equipment manufacturers have gone to great lengths to make sure the audio provided is distortion free, doesn't break up and is delivered in a known timeframe. Analog audio used point to point connections with dedicated cable. Digital built on this to give us time division multiplexed systems such as AES3 and MADI, albeit still over point-to-point connections.

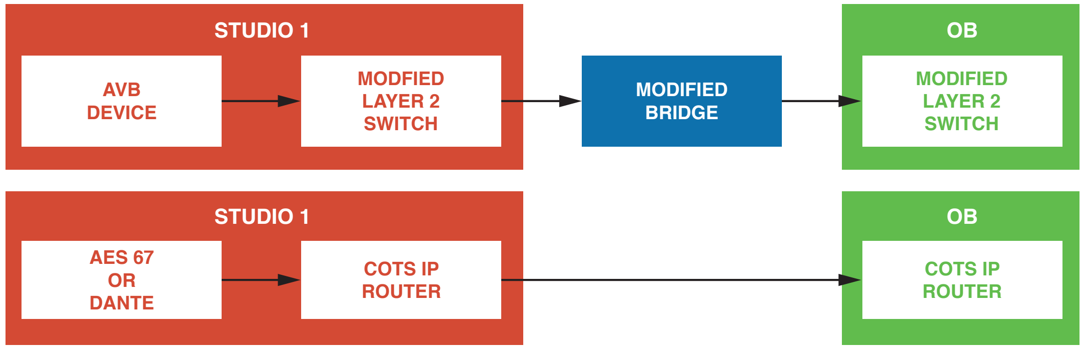

Figure 1 - comparison of AVB and AES67/Dante systems.

As we move to the IP world, we have to look at ways we can deliver audio over networks without delay, or breakup and distortion. IP networks were originally designed to transport non-real-time data such as web browser traffic. One of the benefits of point-to-point connections is that we can guarantee the audio will reach its destination within a predictable timeframe and without errors. We must now apply this philosophy to computer networks.

The Audio Engineering Society (AES) has provided two point-to-point standards that have been used extensively in the broadcast industry; AES3 and AES10. AES3 provides pulse code modulated (PCM) audio samples in real time over a point-to-point connection. MADI was adopted by the AES and became formally AES10, this builds on AES3 to provide up to 56 channels of audio.

Point-to-point connections work well within studios and have stood the test of time. The problems start to occur when we move outside of the studio to distribute audio to the wider TV station and sending to other facilities, such as OB units, as this form of connectivity becomes very inefficient.

A 100Mbit/s MADI connection needs a 125Mbit/s baud rate connection, even if you only send one channel consisting of 2Mbit/sec. When using packet switched systems such as Ethernet or IP, we only need to use the data bandwidth required, in this instance 2Mbit/sec, with some minor overhead for packet framing.

Two technologies have emerged as the leaders in audio over IP; IEEE 802.1BA and AES67. IEEE 802.1 is otherwise known as Audio Video Bridging (AVB) and works at layer 2 of the ISO seven-layer model. In previous chapters we demonstrated that layer 2 switches Ethernet packets and is limited to the domain of the broadcasters' network. Although its faster than IP switching it's difficult to move packets outside to other facilities and OB trucks.

To achieve synchronization data integrity and low latency, IEEE 802.1 uses new switcher protocols to provide rate shaping, a method to guarantee bandwidth, and clock synchronization. All of this requires the layer-2 switch to have IEEE 802.1BA protocols installed and configured. The IT department will need to understand this configuration and as it's not standard IT, the system is much more complicated.

For many commercial broadcasters, the benefits of moving to consumer of the shelf (COTS) products could easily be outweighed by the customization of switches required by AVB. Instead of using hot spares available that can be deployed anywhere in the network we now have to keep specific routers and switches for the broadcast part of the network.

The main problem with packet technology is jitter and delay throughout the network. Solutions are required to synchronize the codec frame and bit clocks at the send and receive ends of the chain.

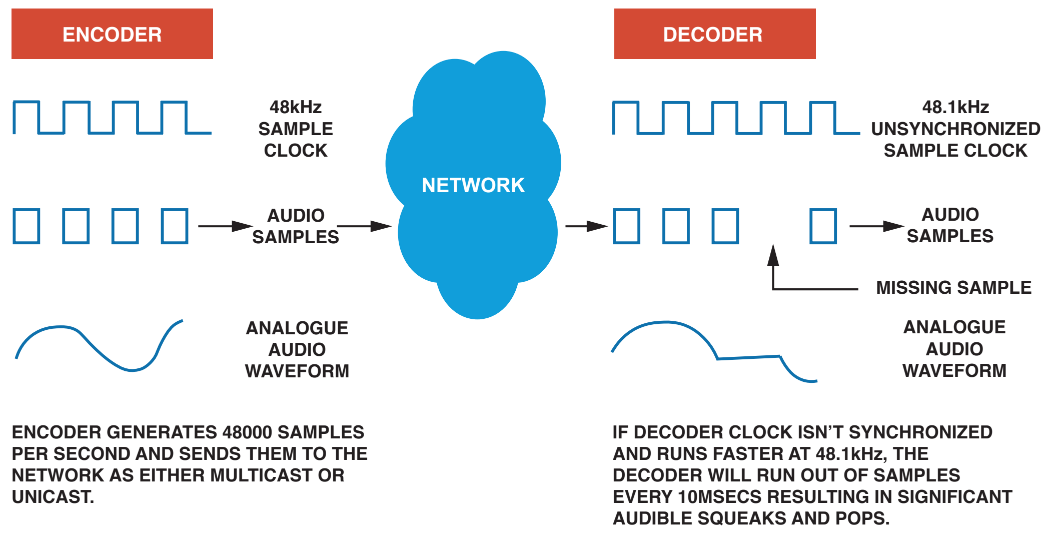
ENCODER
48kHz SAMPLE CLOCK
AUDIO SAMPLES
ANALOGUE AUDIO WAVEFORM
NETWORK
DECODER
48.1kHz UNSYNCHRONIZED SAMPLE CLOCK
AUDIO SAMPLES
MISSING SAMPLE
ANALOGUE AUDIO WAVEFORM
ENCODER GENERATES 48000 SAMPLES PER SECOND AND SENDS THEM TO THE NETWORK AS EITHER MULTICAST OR UNICAST.
IF DECODER CLOCK ISN'T SYNCHRONIZED AND RUNS FASTER AT 48.1kHz, THE DECODER WILL RUN OUT OF SAMPLES EVERY 10MSECS RESULTING IN SIGNIFICANT AUDIBLE SQUEAKS AND POPS.

Figure 2 - The importance of clock synchronization between encoder and decoder.

Without synchronization, the audio would degenerate to squeaks and pops and be completely inaudible.

AES67 is a packet technology working at the IP layer-3 level and provides a specification for the three main areas of audio over IP; synchronization and transport, encoding, and connection management. This is much easier to distribute outside of a facility as it uses IP routing, which can easily deliver to an OB truck. AES67 provides word and frame clock alignment to guarantee the delivery of high-quality audio over IT networks.

Dante by Audinate goes one step further as it abstracts away the IT network to deliver a user management system enabling simple discovery of connected equipment and interfacing to computers. In effect, providing a system that is easy to manage without having to have an in-depth understanding of the underlying IT infrastructure.

AES67 and Dante both have the major advantage that they do not need any modifications to industry standard IP routers and can work alongside existing IT networks.

Using Precision Time Protocol, timing signals are sent from a master computer to each connected encoder or decoder. This IEEE protocol allows sub-microsecond clock synchronization between devices guaranteeing synchronization of clocks and high-quality audio.

Network engineers must configure their routers to allow PTP timing packets to have the fastest access through the network, in affect providing the best quality of service (QoS). Broadcast engineers must discuss this with the IT department so that the network engineers can apply the correct QoS parameters to the PTP packets.

Other benefits of Dante are that it provides a plug and play facility. We saw in the previous chapter on host configuration, how easy it is to create ghosting IP addresses, or mis-configure a camera or sound console. Over three hundred partner companies build interface solutions to work with Dante making it a complete solution.

6

Video Streaming

How efficient one-to-many video distribution is achieved over IP networks using multicasting.

In traditional broadcast SDI systems, we use a point-to-point distribution system with one-to-one mapping. That is, to connect the output of a camera to a production switcher we take the SDI output of the camera and connect it directly to one of the production switchers inputs. If we then want to monitor the camera feed going to the production switcher we must either disconnect the input or introduce a Distribution Amplifier which has multiple outputs. Each one of these then forms a one-to-one mapping with monitors, routing matrices and other processing equipment.

By introducing a DA we have effectively provided a one-to-many mapping system.

In the IP world, to provide one-to-many mappings, we substitute the physical DA with an abstracted logical model using IP routers. Instead of having individual SDI cables connecting the equipment together, we take advantage of the IP-router's ability to be able to duplicate packets and create a multicast system.

IP routing generally works on the unicast model; when a computer requests data from a server, there will be just one source and one destination. However, if a playout server wants to send an online film to many users, a mapping of one-to-many is used and this is called multicast.

Multicasting is much more flexible than SDI routing using matrices and DA's as we are not limited to working in the SDI domain. If a user wants to watch the studio output from their office, or the green room, then we don't have to worry about running new SDI cables to them, or adding the studio to RF distribution, we just set up another multicast user in the IP-router.

The IP-router provides the one-to-many mapping by distributing the video at the IP packet level. Each IP datagram is copied and then sent out on the appropriate port of the router to the user.

The source device, such as camera, production switcher or sound console, has no knowledge of which devices are receiving their output in an analogous way to traditional broadcast equipment. The sound console has no knowledge of the destination of its outputs, only that they are seeing a high impedance load to a DA at the end of a twisted pair. The Ethernet output of a camera has no knowledge of where its packets are going and can only see the network interface card of the IP-router or switcher it is connected to.

An alternative method to multicasting would be to set the camera's destination address of each device requiring the stream. This would cause two problems, firstly, the overhead of having to set the IP destination addresses of six cameras in a studio, each connected to potentially ten different destinations, is an administrative nightmare and would be unworkable. Secondly, we would increase the network load by the number of destinations set in the camera's routing table, as each packet would be individually sent to each device requiring the stream.

Multicasting solves both problems as it relies on the receiving equipment to opt-in to the feed, and IP packets are only sent once to each group destination, with the routers providing duplication only on the network branches that need the stream. There are three key concepts in a multicasting system: the group address, reverse path forwarding and IGMP.

The multicast group address is the destination IP address of the equipment creating the audio or video stream, i.e.,

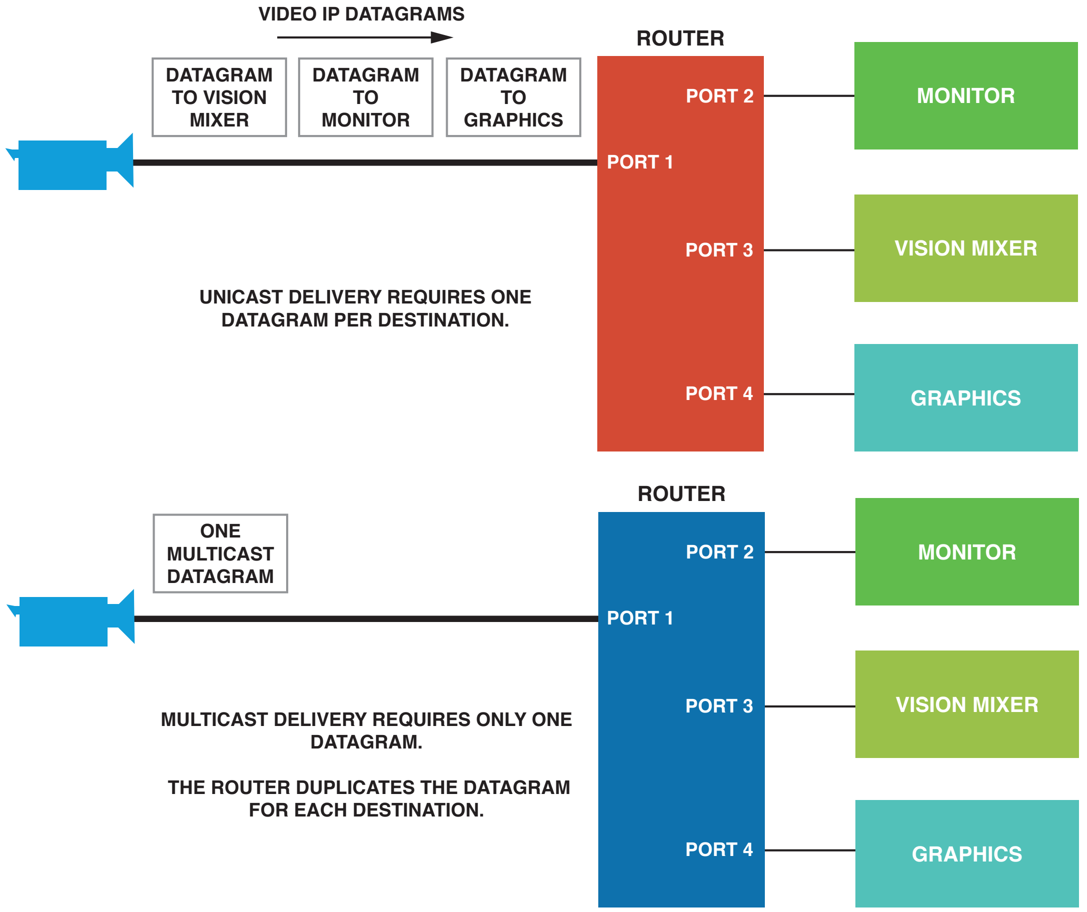

Figure 1 - Multicast distribution is more efficient than the unicast distribution for streaming to multiple destinations.

a camera or microphone. Each camera, microphone, production switcher output or sound console aux send will require its own group address. The multicast group addresses are special IP addresses constrained in the range 224.0.0.0 to 239.255.255.255, approximately 248 million groups. Some of these are reserved and a full list can be found at the Internet Assigned Numbers Authority (IANA).

Internet Group Membership Protocol (IGMP) provides the mechanism for equipment receiving the stream (monitors etc) to tell its router that it wants to receive a particular group. Each IP-router passes this information along the line to the source equipment so that only the networks that require the stream relay it, thus reducing network congestion.

In figure 2 the router will periodically send an IGMP Host Membership Query message to all devices connected to its networks, in this case the sound console responds with an IGMP Host Report Group 1 and Group 2 message, advising the router that it wants to receive Mic-1 and Mic-2 audio streams. Mic-3 (Group 3) is not sent to the sound console as the sound console did not opt-in to receive that group.

Unicast routing works by looking at the destination IP address of the packet and comparing it to its look-up table so the router knows where to forward the packet to. Multicast routing uses the source address in the IP-datagram to route back to find out where the source of the packet is, a system called Reverse Path Forwarding (RPF).

The beauty of IP networks in general, and multicasting specifically, is that the data can move in both directions simultaneously. For example, a sound console could have all its aux sends and returns on one Ethernet cable, potentially saving tons of twisted pair cable. The sound console becomes both a multicast receiver, and multicast broadcaster. If sufficient capacity IP routers are used the production switcher can do the same with its video inputs and outputs, with hundreds of feet of coax cable being substituted by CAT6 Ethernet and fiber optic cable.

If a producer wants to watch the studio output in their office, their desktop computer or tablet can be easily configured to receive the group address of the studio output. No configuration will be required by the IT department, and no additional cables will need to be run to their office. However, this introduces the topic of network security. It might be that the producers do not want the whole station to be able to see the studio output, so the routers would have to be configured to stop them sending certain streams to certain networks.

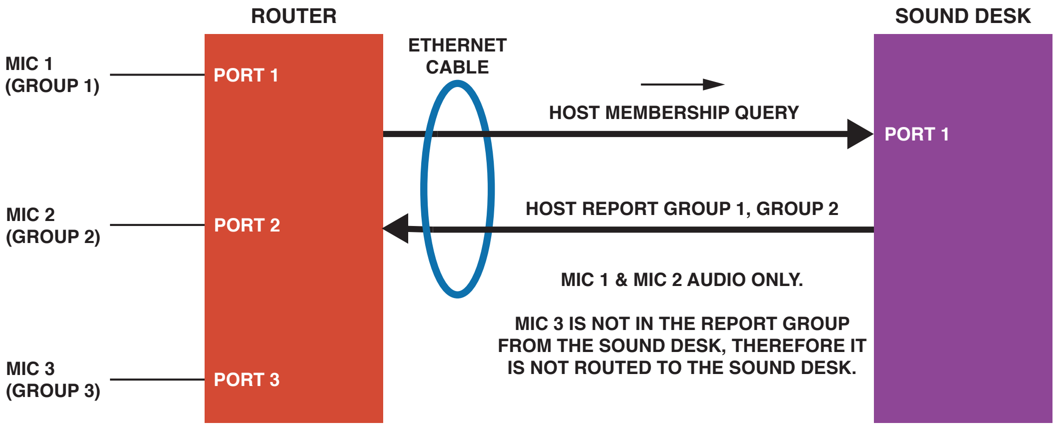
ROUTER
ETHERNET CABLE
SOUND DESK
MIC 1 (GROUP 1)
MIC 2 (GROUP 2)
MIC 3 (GROUP 3)
PORT 1
PORT 2
PORT 3
HOST MEMBERSHIP QUERY
HOST REPORT GROUP 1, GROUP 2
PORT 1
MIC 1 & MIC 2 AUDIO ONLY.
MIC 3 IS NOT IN THE REPORT GROUP FROM THE SOUND DESK, THEREFORE IT IS NOT ROUTED TO THE SOUND DESK.

Figure 2 - IGMP Host Membership Query messaging.

7

Timing

How the introduction of PTP addresses the critical challenges of timing in IP networks and brings additional flexibility to broadcast infrastructure.

Broadcast has timing intrinsically built into the signal paths. For example, analog PAL and NTSC have field and line sync pulses to synchronize the scanning process in cathode ray tubes. Color sub-carrier bursts synchronize the flywheel oscillator to lock the color demodulation frequency. And SDI and AES have bi-phase modulated clocks built into their signals so that the receiver clock can lock to the sample clock.

Clock synchronization is extremely important in both synchronous and asynchronous digital television systems. The problem we are trying to solve is to keep the encoder and decoder sample clocks at the same frequency and in phase. If we do not do this, then one clock will run faster than the other resulting in either too many or too few samples reaching the decoder.

Lost samples of data in uncompressed signals will cause an instantaneous audio splat or loss of a video pixel. In compressed systems, the effect could be much worse as forward and reverse compression can result in a prolonged error.

Broadcasters have gone to great lengths to provide master clock referencing for both audio and video in the form of master sync pulse generators.

Although Ethernet uses bi-phase modulation to encode its data and clock signal, the clocks are not synchronized between network interface cards (NIC's) so we cannot use this as a form of global synchronization.

GPS has been used in the past to lock encoders and decoders; however it's proved impractical when the signal path moves away from line of sight of a satellite.

Precision Time Protocol IEEE-1588 has been developed by the IEEE to address the issue of network timing. PTP was designed as a standard for many different industries and as it can provide sub-

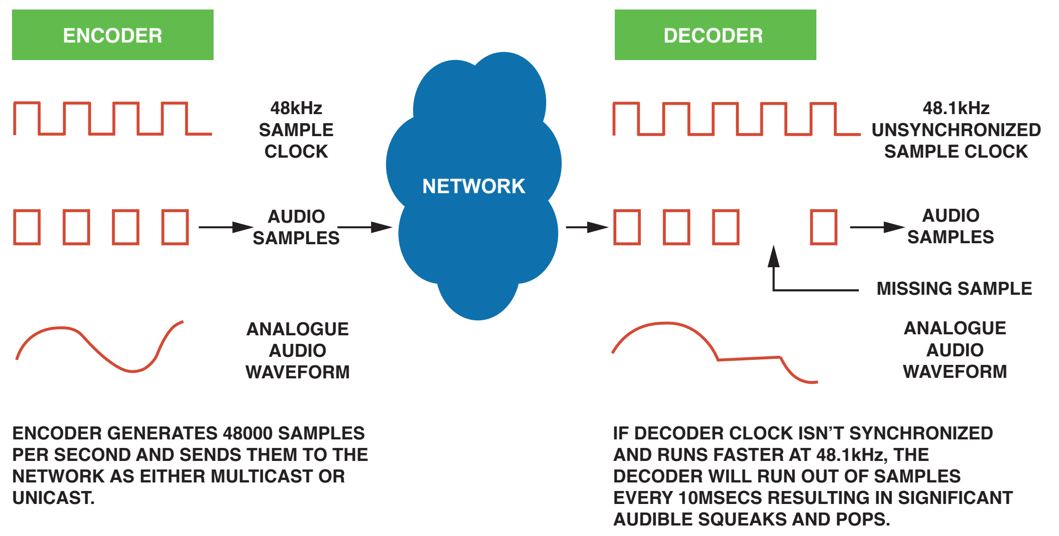

ENCODER GENERATES 48000 SAMPLES PER SECOND AND SENDS THEM TO THE NETWORK AS EITHER MULTICAST OR UNICAST.

IF DECODER CLOCK ISN'T SYNCHRONIZED AND RUNS FASTER AT 48.1kHz, THE DECODER WILL RUN OUT OF SAMPLES EVERY 10MSECS RESULTING IN SIGNIFICANT AUDIBLE SQUEAKS AND POPS.

Figure 1 - The effects of non-synchronized encoder and decoder clocks.

microsecond accuracy, it lends itself well to broadcast television.

PTP works in a master slave topology. One server or customized device is nominated as the master clock, and all other devices within the subnet synchronize to it forming a network of synchronized servers.

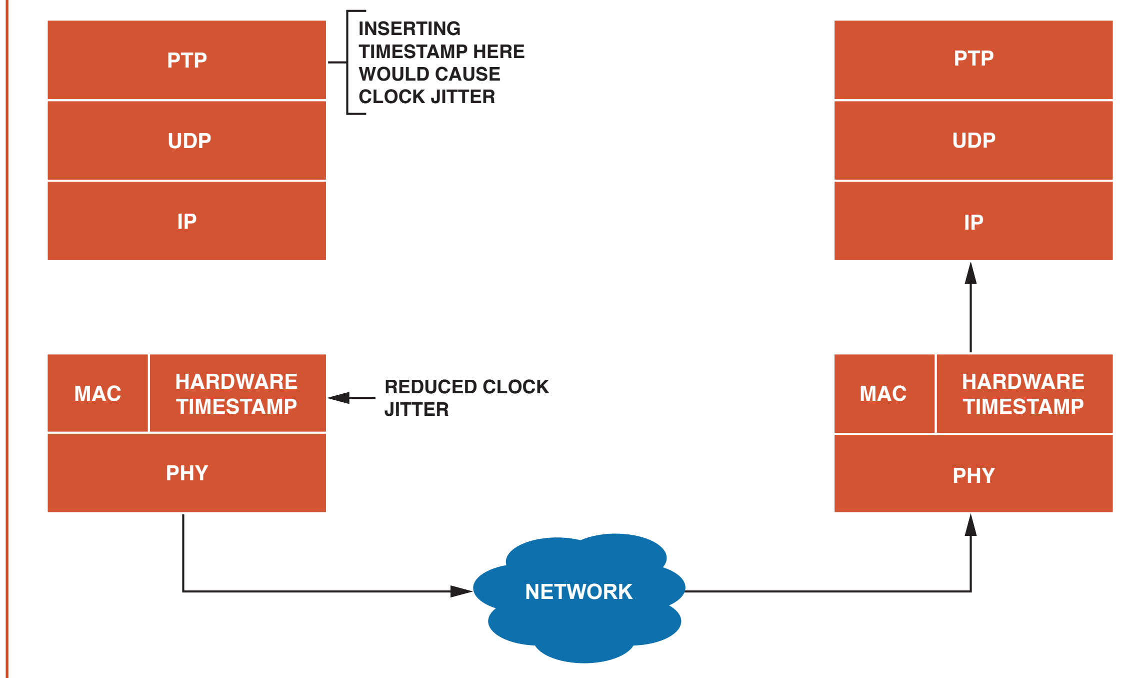

Figure 2 - Inserting time stamps into the NIC will reduce clock jitter.

Although the protocol can run on any router without modification, some configuration work must be done to provide the timing packets with the fastest and shortest delay path in the network. Network engineers achieve this by setting the quality of service (QoS) in the routers for specific types of packets by using a form of rate shaping that gives priority switching to the timing signals.

The time difference between the master and slave clocks consists of two components; the clock offset and the message transmission delay. To correct the master clock, synchronization is achieved in two parts, offset correction, and delay correction.

The master clock should be a very accurate generator capable of providing 1GHz clock samples, either locked to GPS or deriving its clock from an oven-controlled oscillator in a similar way to the broadcast sync pulse generator. Established manufacturers of SPG's are now including PTP clock outputs on their products.

In a similar way to Unix time systems, PTP uses the concept of an Epoch clock. This is an absolute time value when the clock was set to zero, and the number of 1GHz clock pulses that have occurred since provides the current time, these are converted into human readable time with software to provide year, month, day, hours, minutes, and seconds. The Epoch (or zero time) for PTP was set at midnight on the 1st January 1970.

As PTP uses a 1GHz master clock, the granularity of the slave clock can be accurate to 1nS. The clock should be thought of as an event clock or presentation time clock rather than an absolute pixel count.

Software timing is notoriously unpredictable hence the reason manufacturers have kept to hardware solutions for time critical processing such as video playout. When PTP masters create timing packets, and slaves receive them, the timestamp should be inserted within specially designed network interface cards at the Ethernet layer. If it was inserted by the software stack, then jitter would occur due to the unpredictable interactions of the operating system and software stacks.

When sending a video frame, some frames will arrive ahead of their display time and some behind. Buffers smooth this out and the internal presentation software will make sure the frame is constructed before the next field pulse comes along. In effect, the frame pulses are synchronized by the PTP, so the frame rate of the receiver is locked to the encoder.

The benefits of this method of synchronization go beyond video and audio playout. PTP now provides us with a predictable event clock so we can trigger events in the future instead of relying on centralized cues. If a regional opt-out of Ads was to occur in a schedule at 19:26:00hrs, the remote playout servers would be able to switch the program at 19:26:00hrs to play the regional Ads within a timeframe of 1nS. If the schedule database is correctly replicated to all the regional playout servers, we no longer have to rely on cue tones and in-vision prompts to provide opt outs.

PTP protocol allows for master and slave devices to be daisy chained together so a slave device can become master for another subnet. In this way, we can have entire LAN's and WAN's synchronized together to allow broadcast devices to accurately switch and mix between sources.

In traditional analog and SDI studio's there tended to be just one timing plane for the video; the production switcher. If multiple production switchers were to be used, then video synchronizers would have to be employed to provide another timing reference. PTP removes this need as the timing plane is essentially the same throughout the entire network as all slaves and masters become synchronous.

A new timing dimension has been brought to broadcast television.

8

VLAN’s

How VLAN’s enable the division of an Ethernet network into multiple smaller logical networks.

VLANs work at the layer 2 level, that is Ethernet. They are similar to subnets but not the same, they provide network security and improved performance.

IP has been successful within the internet and media domain as it is transport stream independent. That is, it can work with Ethernet, ISDN, ATM, serial and a whole plethora of different underlying hardware distribution networks. Video and audio streams provide a comparative analogy as they can both exist independently of SDI or computer networks.

A single Ethernet network can have thousands of devices connected to it using hubs, switches, and bridges.

Hubs are rarely used as they replicate all the traffic on one port to all the ports on the rest of the hub, causing an increased likelihood of congestion and collisions, especially in high bandwidth video and audio applications.

Switches are available in two varieties, managed and un-managed; An un-managed switch learns which devices are connected to each of its physical ports. When an IP camera wants to send video streams to a production switcher with IP address 10.2.1.9, it first sends an address resolution protocol (ARP) query, which says “who has IP address 10.2.1.9 and send me your Ethernet address?” The ARP query is sent to all devices connected to the layer 2 switch using an Ethernet broadcast message.

The production switcher responds with its Ethernet address; the camera then sets its destination Ethernet address to be that of the production switcher answering the ARP query. The un-managed layer 2 switch monitors this interaction and learns which port 10.2.1.9 is connected to, and from then on will only send traffic for the device to the port its connected to, in effect reducing congestion on the rest of the network, thus stopping collisions and improving efficiency.

Un-managed layer 2 switches do not require configuration and cannot be used as VLAN devices. Managed layer-2 switches allow more control over the network such as data rate shaping, quality of service configuration and VLAN ports.

A group of switches defines a network, bridges will link several networks together of the same protocol type, in the case of Ethernet this is a layer 2 bridge. So, if Studio-1 is in one building and Studio-2 in another, a bridge can be used to link the two different networks together at the layer 2 level.

At layer 3, if the IP address of a packet cannot be resolved in the network it is sent to a gateway router, the router has look up tables with destination addresses so that it can forward the packet to another network, which may also be a different protocol such as ATM or DSL.

Ethernet bridges differ from routers as they can only route layer 2 traffic between networks of the same type, for example Ethernet. But if the user wants to send a packet to a network of a different type, for example ADSL, a router must be used.

The fundamental problem with this approach is that within a single Ethernet network, all devices can be seen by all other devices. Camera-1 in studio-1 could send data to the sound console in studio-3, even with a managed switch. This may increase flexibility, but the network becomes congested extremely quickly and security is an obvious issue. Equipment may stop working properly as

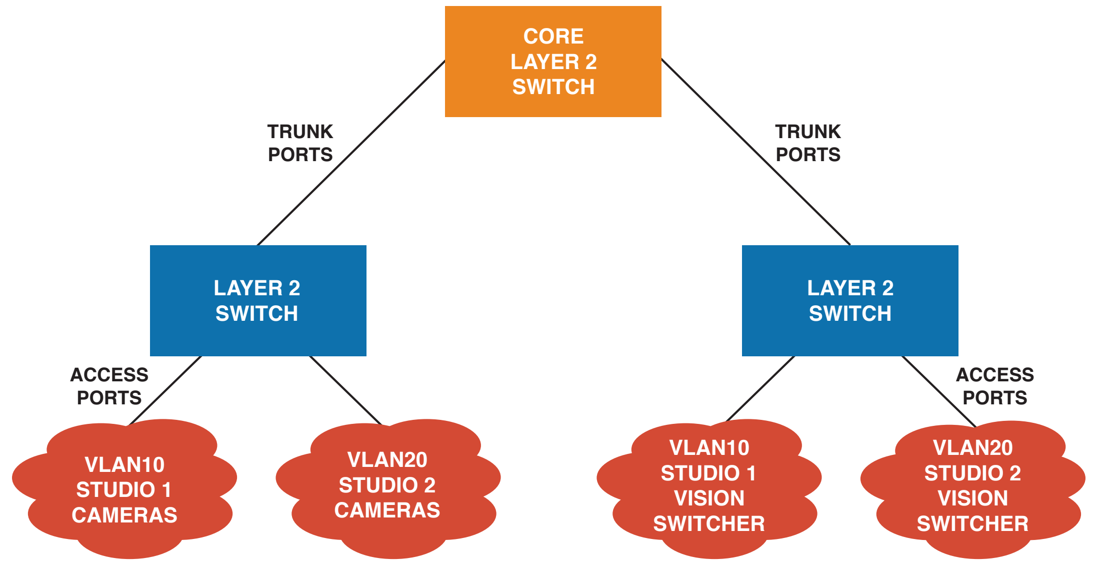

Figure 1 - Cameras can send video IP packets to Studio 1 Vision Switcher on VLAN10, but cannot send to Studio 2 as they are on a different VLAN.

a sound console would not respond well to having data from many cameras sent to it.

Virtual LAN's (VLANs) are a solution to both these problems.

Layer 2 switches, when configured to operate in VLAN mode can logically separate an entire Ethernet network into many different logical networks. The key here is "logical", the devices can still be connected to the same physical switcher, but the ports can be labelled with different VLAN identifiers (VLAN ID's), the algorithms within the switches stop Ethernet frames being sent to ports with different VLAN ID's, thus greatly improving security and congestion handling.

The layer 2 switch inserts the VLAN ID into the header of an Ethernet frame as it enters the switch, sends it to the appropriate port, and just as the frame leaves the switch at the destination end the VLAN ID data is removed. From the point of view of the user the VLAN ID is never seen.

Each port on the switch can be configured as either an access or trunk type. Access ports can have only one VLAN configured in the interface and carry traffic for only one VLAN. Trunk ports can have two or more VLANs configured on the interface and can carry several VLANs simultaneously. Trunk ports are generally used to route VLANs to different switches.

Although each access interface can have only one VLAN, they can all be different VLAN ID's on the same switch. This is where the logical separation, security and reduction of congestion takes place. If port 1 has VLAN1, port 2 has VLAN2 and port 3 has VLAN3 configured in the switch, then none of the devices connected on each of these ports can see the other devices. So, if camera 1 is connected on VLAN1, microphone 1 is connected to VLAN2, then camera 1 media streams cannot be sent to the microphone on VLAN2.

Sometimes a device may need access to a different VLAN. If the producers' computer in Studio-1 was attached to VLAN11 and they needed email access which was on VLAN90, then a router would be needed to connect the two networks together. This doesn't compromise security as the network administrator will be able to configure the router to allow only email traffic to the computer.

Generally speaking, each IP subnet is aligned to a VLAN ID, this makes administration easier and routing between different VLAN's more intuitive. Some layer 2 switches have layer 3 routers built into them allowing routing between VLANs. If the layer 2 switch does not have a router built into it, then an external layer 3 router must be used.

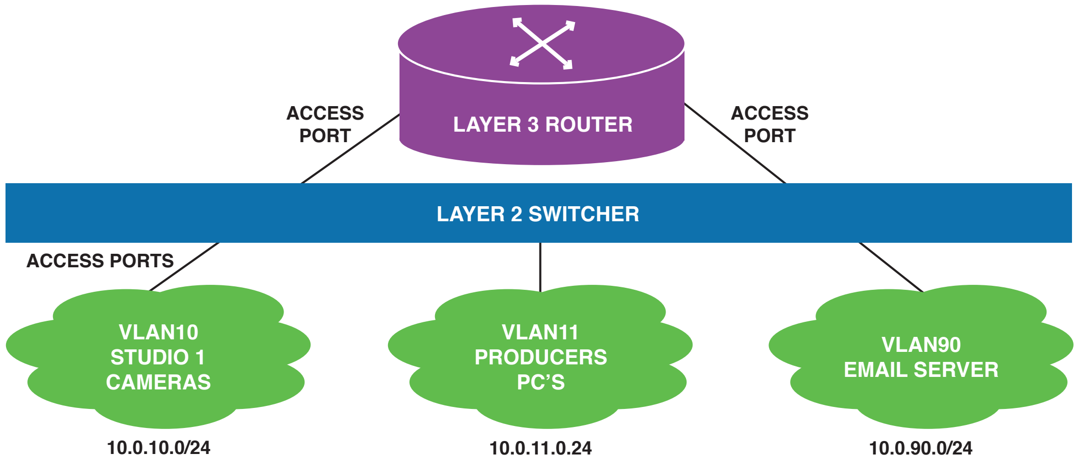

Figure 2 - For the producer to access emails, the network administrator must provide a route from 10.0.11.0/24 to 10.0.90.0/24 using a layer 3 router.

9

Ethernet

How Ethernet has evolved to combat congestion and how speeds have increased through the decades.

In the 1970's Ethernet was in its infancy and competing with two other proprietary networking systems, token ring and token bus. These systems were backed by heavy industry and the concept of each employee having a computer on their desk wasn't considered.

Token ring allowed computers to be connected to each other with coaxial cable in a ring topology. An empty packet of data held the electronic token and was passed between each computer, and only when the computer received the token could it send data. Various timeout protocols existed to stop a computer hogging the token.

Token bus uses the token ring protocol, and they only differ in that token bus can have an open-ended network connection, but token ring must have the LAN connected in a continuous loop. With token ring, each computer on the network must know the address of the previous and next computers within the LAN to allow them to pass the token.

Token systems could be unreliable, especially if a computer failed, or it was switched off. Adding extra computers to the network was difficult as it stopped all computers on the LAN from communicating. Furthermore, they were highly inefficient as they would have to receive, process, and transmit the empty token packet even if they had no data to send.

Ethernet was adopted by the IEEE in 1980 and given the project number 802. Although Ethernet originally used coax as its LAN cable, it soon moved to twisted pair providing a cheaper more flexible alternative with full duplex operation. Coax was limited to either send or receive data but could not do both at the same time. Hardware buffers within the Network Interface Cards (NIC) allowed the computer to preload memory at CPU speeds and then the NIC would send the data at Ethernet line speeds, thus making the whole operation significantly faster.

Ethernet was originally developed with a shared bus system in mind. All computers on the LAN would be either listening, sending, or receiving data. To send data, a computer would detect other transmissions, and if it found another communication was taking place it would wait.

During transmission, the network interface of each computer would continue to listen for other traffic on the coax, and if one or more computers started sending data, they would all detect this and stop their transmissions. An algorithm in each NIC randomly held back the transmission so statistically one NIC would transmit ahead of the others and stop an infinite race-off condition occurring, keeping the system stable.

This collision detection is called "Carriers Sense Multiple Access with Collision Detect", or CSMA/CD and was adopted by the IEEE as 802.3 and became a full published standard by 1985. CSMA/CD is used on twisted pair networks as multiple computers can be connected through a hub. These essentially connect all the transmits together (through buffer circuits) and all the receives together so CSMA/CD was still required.

With one cable or transmission system, packet collisions would increase as the number of computers increased and communicated more. More collisions result in reduced efficiency and hence lower data rates. Layer 2 switches solved this problem as they greatly reduced the number of devices on each segment and

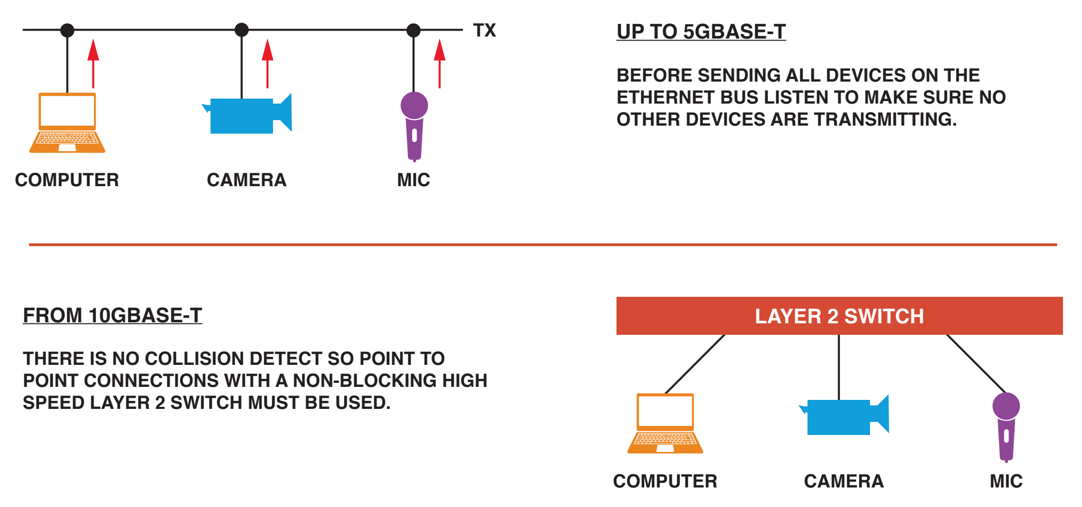
TX
COMPUTER
CAMERA
MIC
UP TO 5GBASE-T
BEFORE SENDING ALL DEVICES ON THE ETHERNET BUS LISTEN TO MAKE SURE NO OTHER DEVICES ARE TRANSMITTING.
FROM 10GBASE-T
THERE IS NO COLLISION DETECT SO POINT TO POINT CONNECTIONS WITH A NON-BLOCKING HIGH SPEED LAYER 2 SWITCH MUST BE USED.
LAYER 2 SWITCH
COMPUTER
CAMERA
MIC

Figure 1 - Point-to-point requirements for non CSMA/CD CAT8.

could allow a point-to-point connection between the computer and port of the switch.

Intelligent switches took advantage of input buffering and would decide when to schedule the sending of a packet and avoid collision altogether.

Layer 2 switches can be connected through bridges to keep maintenance easy and reduce complex cable runs.

Each Ethernet network interface card has its own unique Media Access Control address programmed during manufacture. Ethernet assumes that the manufacturers of the NIC have been assigned a MAC address from the IEEE's registration authority and that it has been programmed properly. A NIC will only respond to two types of received messages: its own MAC address in the destination address header, and the broadcast address in the same part of the header.

The broadcast address is always "ff-ff-ff-ff-ff-ff" and is used by protocols such as address resolution protocol (ARP) to resolve an IP address to a MAC address. The downside of this is that when one computer sends out a broadcast message every single device on the network will receive it and must process it. This is inefficient and a waste of valuable bandwidth. The solution to this is to split networks physically using bridges, or VLANs, which coincide with subnets on IP networks.

Ethernet was designed to work with many different protocols, IP is only one of them, and it is possible to send IP packets and ARCNET packets at different speeds on the same Ethernet network.

As Ethernet evolved a series of standards were published (5GBASE-T, 10GBASE-T etc) that defined the bandwidth supported by new and improved cable types (CAT5, CAT6, CAT8) etc.

In recent years, speeds of Ethernet networks have increased significantly.

In the 1980's network speeds were running at 100Kbps on CAT1 cabling, CAT5 provided 250Mbps and CAT8 now gives us 40Gbps by splitting the data over multiple pairs within a single cable or using fiber. However, only distances of 30m are achievable over CAT 8 cable.

The most important aspect of 40GBASE-T is that CSMA/CD has been dropped from the IEEE 802 specification and there is no backwards compatibility with CAT6 and earlier. This is fine for cameras as the output of the camera will connect directly to the high-speed non-blocking layer 2 switch. However, we can't mix other CAT6 devices within the CAT8 segment as they will require CSMA/CD to work correctly.

As speeds increase to 40Gbps and 100Gbps the RJ45 Ethernet connector has been succeeded by SFP and CXP connectors. The small form-factor pluggable (SFP) is a hot pluggable sub-module allowing connection for fiber or copper. CXP combines multiple copper pairs to provide 100Gbps with twelve 10Gbps pairs, or three 40Gbps links.

To network multiple 2160P120 cameras would need a seriously fast layer 2 switch to process all the 24Gbps network streams being sent to it. This would be a switch worthy of a Tier-4 datacenter and certainly wouldn't be Consumer-Off-The-Shelf.

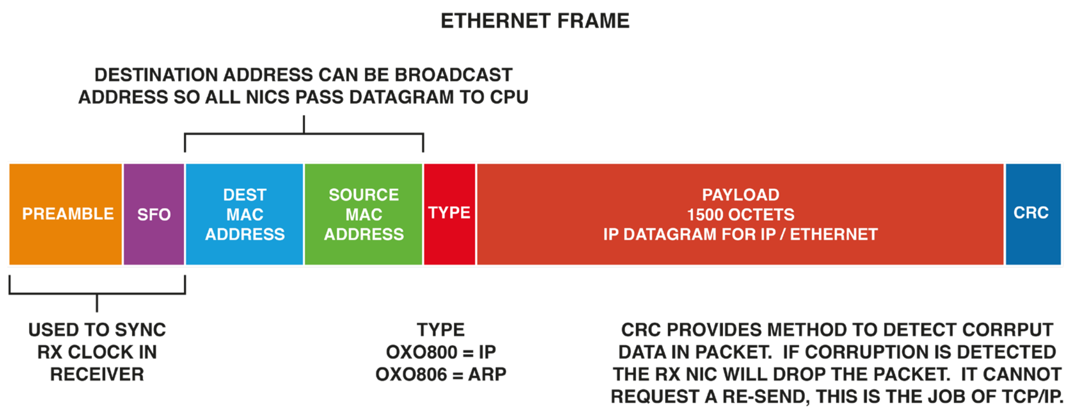
ETHERNET FRAME
DESTINATION ADDRESS CAN BE BROADCAST ADDRESS SO ALL NICS PASS DATAGRAM TO CPU
PREAMBLE
SFO
DEST MAC ADDRESS
SOURCE MAC ADDRESS
TYPE
PAYLOAD 1500 OCTETS IP DATAGRAM FOR IP / ETHERNET
CRC
USED TO SYNC RX CLOCK IN RECEIVER
TYPE OXO800 = IP OXO806 = ARP
CRC PROVIDES METHOD TO DETECT CORRPUT DATA IN PACKET. IF CORRUPTION IS DETECTED THE RX NIC WILL DROP THE PACKET. IT CANNOT REQUEST A RE-SEND, THIS IS THE JOB OF TCP/IP.

Figure 2 - Ethernet frame layout.

10

Security

The flexibility of IP and COTS brings with it all of the security dangers of the internet and the need for robust processes.

New questions need to be asked of broadcast equipment manufacturers, especially around the subject of internet security.

One of the great advantages of broadcast IP Networks is that we can take advantage of consumer off the self (COTS) routers and IT equipment, and we can reduce costs and scale designs easily. One of the disadvantages of using IT COTS is that we are potentially susceptible to the same security issues as the IT world and broadcast engineers must plan for this.

Broadcast IP equipment such as production switchers and cameras will have an Ethernet port running protocols to support at least UDP (User Datagram Protocol) to transmit a packet using the fire and forget policy, once the packet has left the camera or sound console, there is no guarantee that it will get to its destination.

Transmission Control Protocol (TCP) expands on UDP by adding congestion and flow control, and error checking. Many UDP packets are sent within a window and the receiving kit will send an acknowledge packet to tell the sender to send the next packets. If no acknowledge is received the sender will resend the packets, unfortunately this adds delay and is of little use for real time streaming.

On top of UDP and TCP we have protocols such as File Transfer Protocol (FTP) and Hypertext Transfer Protocol (HTTP) which send files and webpages to and from servers. These are all potential vulnerabilities for viruses and malicious hacks, and we must plan for IT style attacks.

Most computer operating systems, whether commercial, proprietary or open source will have a TCP/IP stack. If a camera outputs video over IP, then it at least supports UDP. If you can configure the camera using a web page, then it will almost certainly have TCP. At this point, as far as a hacker is concerned, there is no difference between a PC and a camera, or PC and a sound console.

Configuring a camera, production switcher or sound console over IP is an engineer's dream. No longer do we have to crawl around the floor trying to find the "config" port or having to push a button whilst powering the device to put it into maintenance mode. Being able to configure a sound console using a web page is incredibly powerful for today's broadcast engineer. However, if a studio engineer can gain access to the config screen of a production switcher, then there is potential for a malicious hacker to be able to do the same.

IT engineers go to great lengths to protect their passwords with well documented procedures in place to allow only authorized users administrator access to servers and routers. Broadcast camera operators, sound engineers and VT operators have traditionally had the luxury of being able to gain access to any menu within their equipment. Great care must now be taken when configuring these devices as a camera operator could easily cause problems for other parts of the network if they inadvertently incorrectly set one of the IP parameters. If the network is not designed properly, they could easily take the office telephone system down for example.

With no processes in place, malicious or disgruntled employees could gain the IP addresses of broadcast kit and perform external denial of service (DoS) attacks on the studio at a later date. A DoS attack occurs when an external computer bombards a device through its IP address with requests for data, this will render the camera unusable as the software that is responsible for sending and receiving IP packets will be spending all its time dealing with the data requests from the DoS computer.

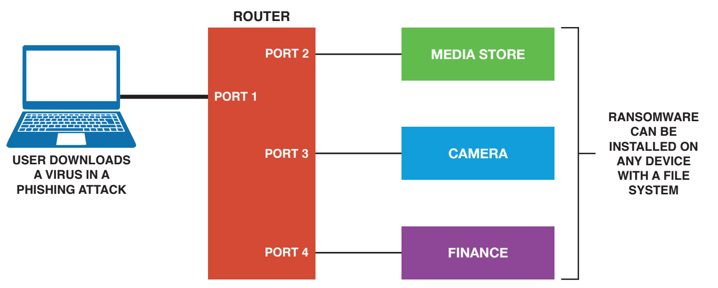
ROUTER
PORT 2
PORT 1
PORT 3
PORT 4
MEDIA STORE
CAMERA
FINANCE
USER DOWNLOADS A VIRUS IN A PHISHING ATTACK
RANSOMWARE CAN BE INSTALLED ON ANY DEVICE WITH A FILE SYSTEM

Figure 1 - Ransomware can easily affect media asset libraries.

One of the most common vulnerabilities are phishing attacks. When a user opens what appears to be a legitimate email only to find when they click on a link, a virus is installed on their computer that easily moves through the rest of the network, including the broadcast kit and any device with a file system on it. Ransomware viruses are installed in this way, replicating to the media store to encrypt media files. A ransom must be paid to the attacker before the files are decrypted so they can be played again.

For these reasons, office and broadcast networks must be separated within the network design. Without adequate security measures, a multi-million dollar media asset library could be rendered worthless in a matter of minutes.

Luckily, the IT department will have thought about and put into place procedures and systems to reduce the risk of attacks and virus downloads to broadcast equipment and media assets. This assumes the broadcast engineers haven't built a side-chain of IT kit as they didn't want to go through the change control processes that ITIL (IT Infrastructure Library) demands. IT engineers are process driven for good reason, that is they need to guarantee uptime and maintain the system without affecting the rest of the business.

Versions of ITIL have found their way into broadcast systems in recent years, especially in playout where service companies provide transmission to many different broadcasters with many channels. The potential to take one broadcaster off air due to the actions on another broadcasters' system must be understood and avoided.

With its change control processes ITIL will be quite new to many broadcast engineers. However, we must respect the IT procedures as we don't want to be responsible for taking payroll down on pay day by changing the IP address of a camera. With flexibility comes responsibilities.

Problems in networks are not always created intentionally or maliciously. Quite often a simple mistake can result in catastrophic failure. For example, if an engineer configures an IP address of camera 1 to be the same as the sound console, then IP ghosting occurs. The routers don't know whether to send the return packets to the camera or sound console, address resolution protocol (ARP) will be thrashing to ascertain the MAC address and either the camera or sound console will randomly respond causing unnecessary network congestion.

Network design consideration must be applied to the security of media being transferred. If a film is being played out to transmission, then there are potential copyright infringements that must be considered.

Somebody may be able to download the film and take it home on a portable disk drive, or they may be able to gain access to the edit storage and copy films, potentially gaining access to blockbusters that have yet to be released.

Big film companies will expect to audit a broadcasters' networks and be certain that nobody can gain unauthorized access to their material. And audit systems will need to be in place so broadcasters can see who is accessing the media and why.

The traditional approach to ring-fencing a network infrastructure has served facilities well, but the advances in cybersecurity has seen the development of a new type of security measure, that is, Zero Trust security. Whereas perimeter-type security systems assumed a user was friendly once they had successfully passed the user credential checks, Zero Trust does not make this assumption and instead relies of verification of every user access throughout the whole network infrastructure workflow.

Zero Trust may seem onerous as it implies that users must log into every device they need to operate, but this is not necessarily the case. By using verifiable context and user control, media assets and processes can be reliably protected. However, Zero Trust is not just an add-on to a network but instead embraces a whole security ecosystem that must be implemented from the start of the network infrastructure design.

11

Network Analyzers - Wireshark

Wireshark is an invaluable tool that enables engineers to examine network traffic in detail.

Wireshark is an open-source packet analyzer running on Linux, UNIX type systems and Windows. Originally called Ethereal and first released in 1998, the name was changed to Wireshark as Ethereal was already a registered trademark.

Wireshark allows engineers to see what is going on under the hood of a network by monitoring an Ethernet port in promiscuous mode and then decoding and displaying the packets. With the intuitive graphical interface, it's very easy to drill down into an IP packet, and then Ethernet frame to see the actual data.

Promiscuous mode is required as Ethernet interface cards generally only pass two types of packets to the CPU; when the destination Media Access Control (MAC) address is the same as that of the card, or the destination MAC is the broadcast address. This would be a serious limitation for any network monitoring device as it would not be able to see packets on the rest of the network.

To bypass this limitation promiscuous mode enables the NIC to pass all Ethernet frames to the CPU regardless of source and destination MAC addresses. Clearly this could be a major security issue as anybody operating in promiscuous mode with a packet analyzer would be able to view and decode all packets within a network.

To reduce security risks, the network administrator will only allow your desk computer to receive frames and packets associated with its VLAN or destination MAC address. In this case, promiscuous mode would have no effect as your computer would not be receiving frames from the rest of the network.

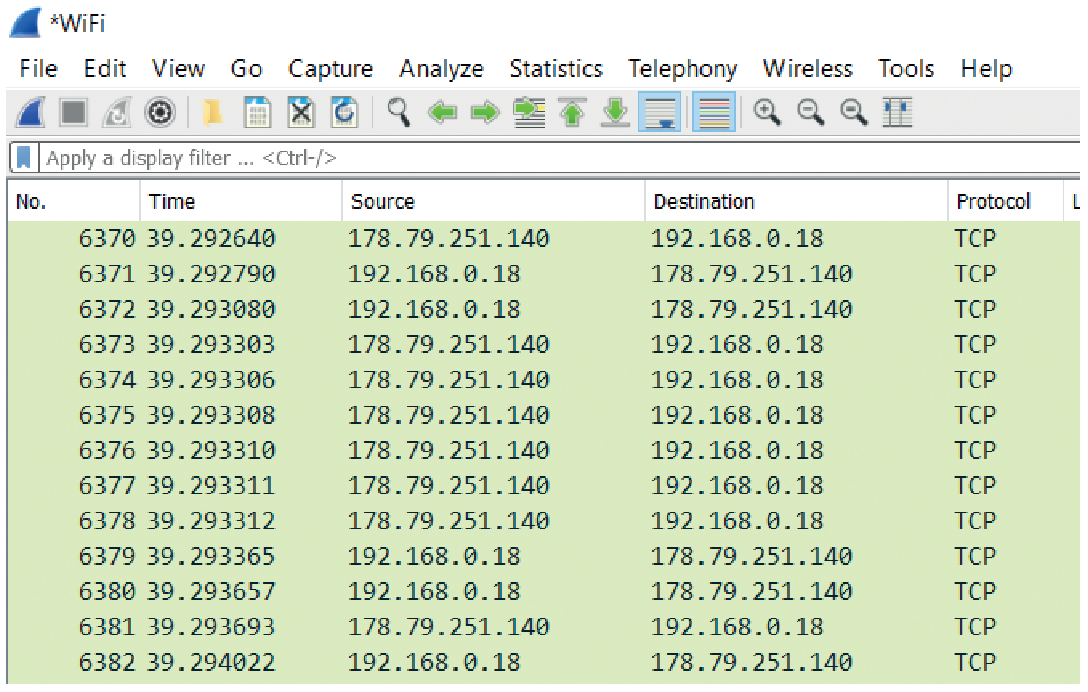

No.	Time	Source	Destination	Protocol	L
6370	39.292640	178.79.251.140	192.168.0.18	TCP	
6371	39.292790	192.168.0.18	178.79.251.140	TCP	
6372	39.293080	192.168.0.18	178.79.251.140	TCP	
6373	39.293303	178.79.251.140	192.168.0.18	TCP	
6374	39.293306	178.79.251.140	192.168.0.18	TCP	
6375	39.293308	178.79.251.140	192.168.0.18	TCP	
6376	39.293310	178.79.251.140	192.168.0.18	TCP	
6377	39.293311	178.79.251.140	192.168.0.18	TCP	
6378	39.293312	178.79.251.140	192.168.0.18	TCP	
6379	39.293365	192.168.0.18	178.79.251.140	TCP	
6380	39.293657	192.168.0.18	178.79.251.140	TCP	
6381	39.293693	178.79.251.140	192.168.0.18	TCP	
6382	39.294022	192.168.0.18	178.79.251.140	TCP	

Figure 1 - Wireshark Showing TCP Packets from Multiple Sources and Destinations.

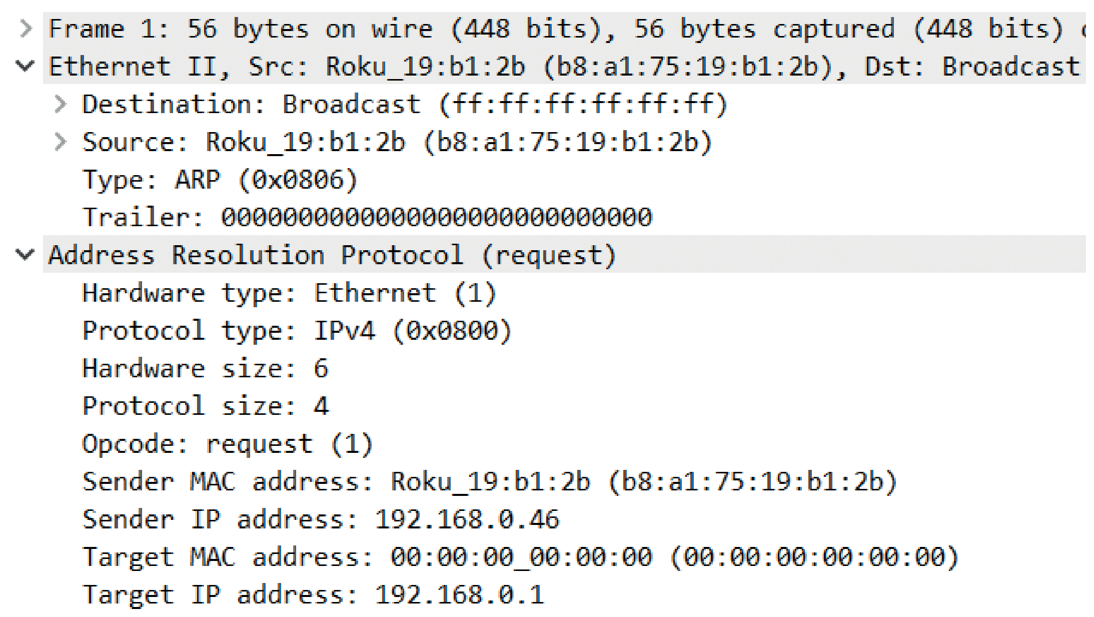

```
> Frame 1: 56 bytes on wire (448 bits), 56 bytes captured (448 bits)
v Ethernet II, Src: Roku_19:b1:2b (b8:a1:75:19:b1:2b), Dst: Broadcast
  > Destination: Broadcast (ff:ff:ff:ff:ff:ff)
  > Source: Roku_19:b1:2b (b8:a1:75:19:b1:2b)
    Type: ARP (0x0806)
    Trailer: 000000000000000000000000000
v Address Resolution Protocol (request)
    Hardware type: Ethernet (1)
    Protocol type: IPv4 (0x0800)
    Hardware size: 6
    Protocol size: 4
    Opcode: request (1)
    Sender MAC address: Roku_19:b1:2b (b8:a1:75:19:b1:2b)
    Sender IP address: 192.168.0.46
    Target MAC address: 00:00:00_00:00:00 (00:00:00:00:00:00)
    Target IP address: 192.168.0.1
```

Figure 2 - Wireshark Showing ARP Broadcast Packets.

Wifi, by its very nature will receive data from many areas of the network. Laptops vary in their ability to operate in Wifi promiscuous mode, but even if they don't then a cheap Wifi dongle can be purchased to allow it. Wireshark can be used to monitor applications used and type of traffic flowing on smart phones, tablets and other devices using WiFi, providing the Wireshark host system is equipped with a WiFi card that supports monitoring mode.

Voice Over Internet Protocol (VOIP) is becoming an industry standard enabling telephony over IP networks instead of having to run the traditional two-wire with ringers to each desk, and VOIP apps are readily available for smart phones. However, to allow their use the network administrator will have to enable VOIP traffic over Wifi.

A Wifi packet analyzer working in promiscuous mode will be able to receive VOIP traffic, and if it's not encrypted we will be able to listen to the conversation. As the packet analyzer is passive there is no way of detecting if somebody is listening to your conversation. VOIP must be encrypted to stop unauthorized snooping.

Broadcast engineers generally work with point-to-point connections, and monitoring consists of either inserting a jack on the listen socket of an audio jack-field or pulling the U-link on a video patch bay. In computer networks, this concept is not available to us, especially if the network is resilient. Pulling a patch-cord out of a server or switch could result in the network believing a link has failed and re-routing all the traffic through another switch or router, in effect removing the traffic we want to monitor.

Ports on managed switches and routers can be configured to work in monitor mode providing all the network traffic for that segment on one port. This can be connected to a computer with Wireshark installed so monitoring can take place. Close collaboration with the network administrator is required to make this possible and the request will probably raise a few eyebrows.

The limiting factor within Wireshark is the hardware it is running on and the speed of the monitoring port on the switch. It would be impossible to monitor a UHD camera feed running at 12Gbps on a 100 Megabit Ethernet server NIC, and a much faster NIC will be required. At these speeds servers start to become very expensive and disc drives fill up very quickly.

Installing Wireshark is very easy and consists of downloading the pre-compiled binaries and installing them onto the target server. For slower networks, a laptop could be used. As this is open-source software the source code is available, empowering engineers to really get under the hood and find out how the code works, or even develop it further.

Many a time engineers have been frustrated at the lack of documentation of a product when trying to configure and install it or suspect bugs that won't be acknowledged by the vendor. With open-source software that frustration goes away as we can literally look at the code to find out how it works and improve the product and increase our knowledge.

Operating Wireshark is very easy. Once installed the program is executed and the network interface selected from the intuitive GUI. Depending on the configurations chosen during installation there may be multiple network interfaces available, including USB connections. Start the capture by selecting the Ethernet interface and clicking on the record button, the screen will soon fill up with all the network traffic being presented to the Wireshark server.

Once a sufficient amount of data has been captured click on the "stop capture" button on the tool bar and analysis of the data packets can begin. By clicking on the arrows at the beginning of each packet, we can drill further and further into the data to find out what is happening in the network and how it is working.

The hex-viewer window even allows us to look at actual data within the packets. On a busy office network, it soon becomes apparent that some software is not as secure as it might first appear, especially when the passwords are sent in unencrypted open text mode.

Filters can be selected during capture mode to look for specific packets, either by protocol, type, or address. This makes Wireshark a fantastic tool and captures can be streamlined to find a problem and saved for off-line analysis using Wireshark's edit tools.

Broadcast engineers must understand what is going on inside an IP network to make sure the strict timing constraints we work to are respected, even more so than the network administrator. Wireshark is the ultimate network training and diagnosis tool and should be understood by every broadcast engineer wanting to excel in the IP domain. But expect a lot of resistance from the network administrator when you turn up with Wireshark.

12

Measuring Network Line Speeds

Broadcast and IT engineers take very different approaches to network speed and capacity - it is essential to reach a shared understanding.

In the analog days, broadcast engineers used frequency response to measure line capability. An audio twisted pair had a bandwidth of at least 20KHz and video had a bandwidth of approximately 8MHz. As we moved to digital audio AES3 and SDI HD these requirements increased to 3Mb/s and 1.485Gb/s respectively.

Digital audio and video systems send high and low voltages to represent the one's and zero's. However, at the higher frequencies we need to take into consideration the analog qualities of the transmission cable and equalizing circuits. Return loss, reflections and phase distortion are all significant factors needing consideration.

When speaking to the IT department, we might think that we should no longer worry about these qualities as signal paths are defined in bits per second. An internet line might be defined as 200Mb/s and an SFP link might be defined as 1Gb/s.

IT engineers tend to be product specialists in their own fields of Microsoft, Linux, and Cisco. Very few of them will study transmission theory in the way broadcast engineers have in the past, especially those who worked in transmitters. This can lead to some very frustrating and confusing conversations between IT and broadcast engineers.

An IT engineer might tell you that the bandwidth of a circuit is 200Mb/s, or the delay is negligible. At this point a broadcast engineers' blood would boil as they have flashbacks to their Telegrapher's equations and two port networks. The simple answer is that IT engineers think in terms of service level agreements, if a Telco has provided a circuit with 200Mb/s capacity then they assume and expect it to be true.

IT engineers think of network capacity in terms of bits/second, as opposed to broadcast engineers who think in terms of bandwidth, return loss, phase distortion and reflections. Further problems arise when we start discussing actual data throughput in a system.

Broadcast engineers will assume a 10MHz point to point network will allow them to send signals with 10MHz bandwidths. As IT networks are based on packets of data, there is an inherent loss of data in the system due to the headers and spaced distribution. A 10Mb/s circuit might only have a useful data-rate of 9.5Mbps.

Ethernet frames are generally split into two parts, the header and payload. The header includes information such as send and receive addresses, payload types, packet counts and flags. The payload will be a protocol type such as IP, which will also contain a header and payload.

Drilling down into the Ethernet packet we have the four octets of Cyclic Redundancy Check (CRC), and twelve octets of inter-packet gap appended to the end of the frame.

When discussing networks, engineers use octets instead of bytes to represent eight bits. This is to remove any ambiguity as computer science tells us a byte is "a unit of information", which is based on the hardware design of the computer, and could be as easily eight bits, ten bits or one hundred bits.

Our Ethernet frame generally consists of 1,542 octets, or 12,336 bits. Only the payload is available to us when sending audio and video which is 1,530 octets,

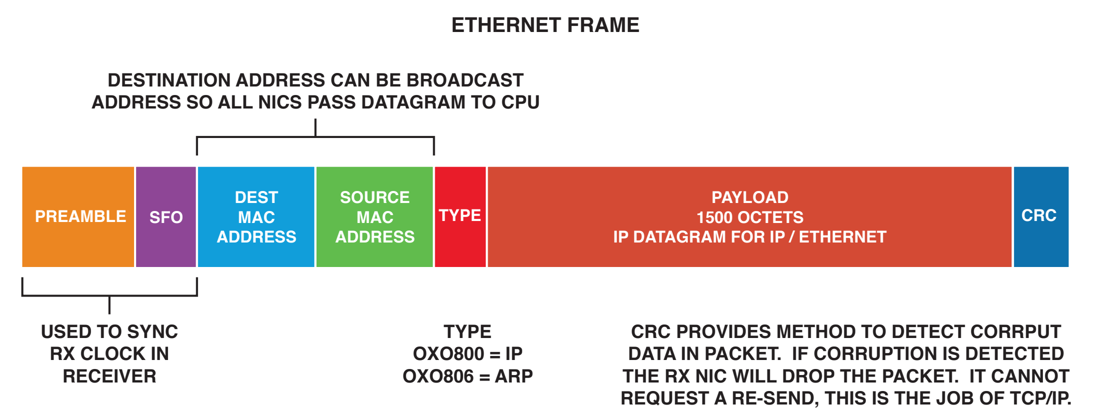
ETHERNET FRAME
DESTINATION ADDRESS CAN BE BROADCAST ADDRESS SO ALL NICS PASS DATAGRAM TO CPU
PREAMBLE
SFO
DEST MAC ADDRESS
SOURCE MAC ADDRESS
TYPE
PAYLOAD 1500 OCTETS IP DATAGRAM FOR IP / ETHERNET
CRC
USED TO SYNC RX CLOCK IN RECEIVER
TYPE OXO800 = IP OXO806 = ARP
CRC PROVIDES METHOD TO DETECT CORRPUT DATA IN PACKET. IF CORRUPTION IS DETECTED THE RX NIC WILL DROP THE PACKET. IT CANNOT REQUEST A RE-SEND, THIS IS THE JOB OF TCP/IP.

Figure 1 - Ethernet packets have an overhead from the header and CRC reducing data throughput.

or 12,240 bits. If a camera uses UDP/IP to send its data over Ethernet a further 20 octets are lost in the Ethernet payload to the UDP header and CRC information leaving 1,510 octets, or 12,080 bits, all resulting in approximately 98% data throughput.

The theoretical maximum of 98% assumes no congestion. The IT engineer is expecting 200Mb/s, but the Telco is providing 196Mb/s of usable data. Clearly the Telco will dispute this as they will show they are providing a 200Mb/s circuit, and the fact that you are using 2% of it on packet header information is your problem not theirs. To the letter of the contract, they are probably correct.

This may not sound a lot, but if you suddenly find your network has a 2% reduction in capacity, you could find yourself with some tough questions to answer when approaching the Finance Director for more cash.

When analyzing network throughput, a thorough understanding of protocols must be achieved.

Transmission Control Protocol (TCP) sits on top of IP/Ethernet and provides a reliable connection-based system to allow two devices to exchange data. Low latency video and audio distribution doesn't use TCP/IP for standards such as ST2110; its data integrity is very high, but the throughput can be low and highly variable. However, some video and audio distribution systems will use TCP/IP and understanding how it works is critical when considering data throughput and latency.

TCP works by sending a group of packets and then waits for an acknowledge packet from the receiver, if the "Ack" isn't received then the same packets are resent, eventually timing-out and sending an error to the user if too many of these errors happen in succession. If the "Ack" is received by the sender, then the next group of packets are sent. Data throughput is reduced and is the price we pay for this data guarantee.

UDP/IP protocols, which ST-2022 and ST2110 uses to transport video and audio, is a "fire and forget" system. The camera outputs the packets and doesn't use any form of testing to see if it was received at the destination, such as a production switcher. ST-2022-5 has forward error correction (FEC) built into it to provide some resilience over the network.

"IPerf" is used to actively measure the maximum achievable bandwidth in an IP network, it's the closest tool we have for measuring a networks capacity and is released under the BSD license. It runs on two computers, a transmit at one end of the network and receive at the other. It works by flooding the link with IP datagrams and using the receiver to check their validity, consequently the measurement is potentially destructive to other users of the network and must be used in isolation.

Operating as a command line tool IPerf can make many different measurements, from UDP/IP line-speed to TCP throughput. TCP measurements will always be slower as IPerf will be measuring the speed of the protocol, not the network line-speed.

When working with the IT department great care must be taken in understanding what exactly you are measuring.

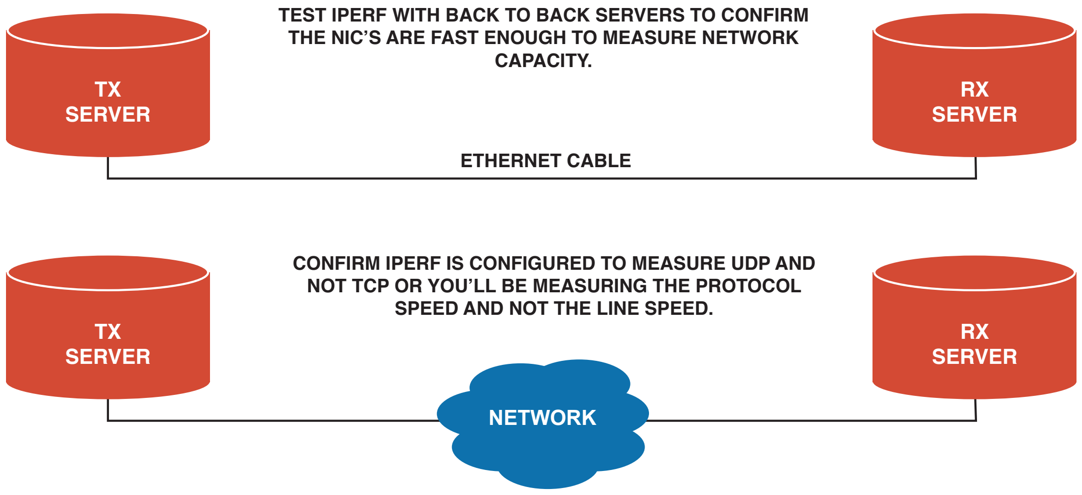
TEST IPERF WITH BACK TO BACK SERVERS TO CONFIRM THE NIC'S ARE FAST ENOUGH TO MEASURE NETWORK CAPACITY.
TX SERVER
RX SERVER
ETHERNET CABLE
CONFIRM IPERF IS CONFIGURED TO MEASURE UDP AND NOT TCP OR YOU'LL BE MEASURING THE PROTOCOL SPEED AND NOT THE LINE SPEED.
TX SERVER
RX SERVER
NETWORK

Figure 2 -IPerf line speed testing.

13

Quality of Service

How QoS introduces a degree of control over packet prioritization to improve streaming over asynchronous networks.

Broadcast engineers are accustomed to point-to-point, reliable, guaranteed bandwidth circuits. An SDI cable will guarantee 270Mbps for SD and 1.485Gbps for HD. Digital audio circuits will guarantee 2.5Mbps for uncompressed 48Khz sampled 24bit word stereo.

Computer networks do not offer this level of guarantee as they are shared packet switched systems and use best effort delivery mechanisms. A stream of video or audio is divided up into packets of 1,500-octets to fit in an IP packet, which is generally inserted into an Ethernet packet. Each of these packets is streamed in sequence at the video or audio data-rate and sent through the network to its destination using a best effort mechanism.

Routers and switchers vary enormously in their sophistication and data handling speeds. The routing protocols they support and their ability to provide services such as multicasting also vary. Consequently, if a packet is sent from a camera to a production switcher, through a network, we cannot guarantee, or even predict how long it will take to get there, whether it will get there, or whether it will be in sequence.

Each packet has a sequence number within its header enabling the receiver to re-order the packets within its input buffer. The packets can be received out of sequence if a route is temporarily interrupted and the packets sent by a different path, resulting in some packets taking longer than others. The receiver can re-order them, but if some packets take too long to arrive then they may be dropped as they will fall outside of the input buffer window.

Input buffers are usually of a fixed size and are used to re-time and re-order packets causing delays for video and audio streaming. If the buffer is too big then there will be an unacceptable delay, possibly even seconds. If the buffer is too small, then a disproportionate number of packets will be dropped as their arrival will fall outside of the buffer window.

To meet the demands of video and audio streaming, IT have adopted the term QoS, adding an element of control to packet arrival. Another term used is rate shaping. Essentially, QoS is helping us distribute a synchronous stream over an asynchronous network.

When IP and TCP were originally designed, there was no consideration or provision for synchronous services such as streamed audio and video.

Two strategies are available to help us reliably stream audio and video over an IT network; extra provision and packet prioritization.

Extra provision is providing more bandwidth than we need. If a streamed audio service requires 2.5Mbps, extra provisioning would demand 5Mbps or even 10Mbps. Clearly this is wasteful of bandwidth and doesn’t scale efficiently. The bandwidth requirements and switching speeds become absurd when we start looking at HD and UHD video.

Packet prioritization takes advantage of information inside the IP header. The Type of Service (ToS) field, recently renamed as Differential Services Code Point (DSCP), belongs to packet prioritization models called differentiated services. This information is used by the router to help it determine its routing order.

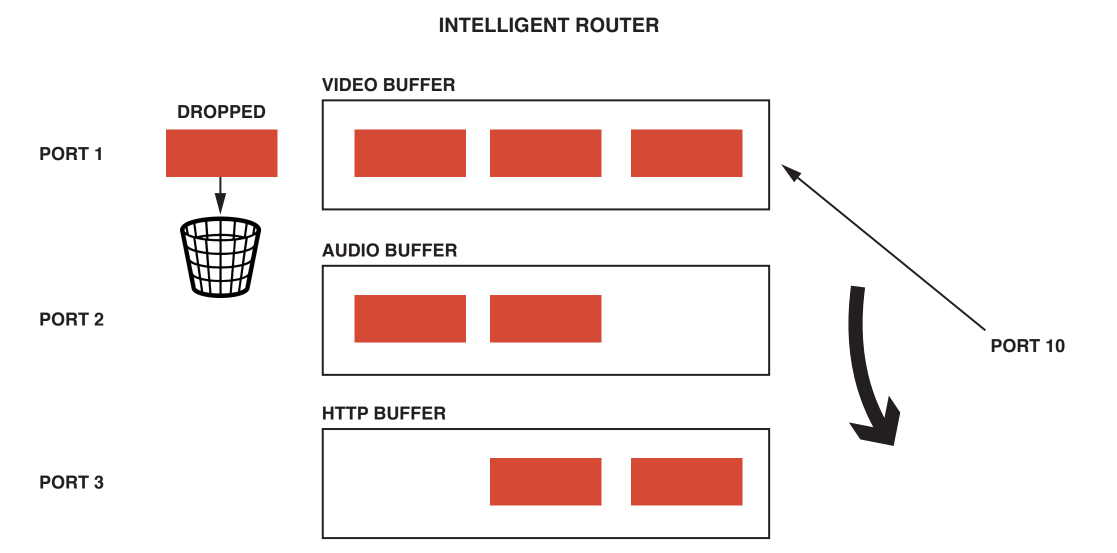

Figure 1 - When packets are queued, the routers algorithm must decide which packet to output next. The video buffer is full resulting in new packets being dropped.

If many streams are being switched to one port, then the router is left to decide which packets take priority over others. This might be a round-robin type strategy, or first come first served. Higher end routers use buffers to temporarily store packets as they enter the device. Algorithms within the router interrogate the packets, extract information such as the DSCP value, and decide its priority when sending to the next hop in the network.

The prioritization causes delay and packet loss within a network. If a buffer becomes overloaded, then it drops packets. Variations in delay give rise to jitter, both long and short term. Solving this problem is the essence of QoS.

Packet prioritization relies on routers within a network all agreeing on a prioritization strategy. From a broadcasting point of view, video and audio should take priority over HTTP and other traffic. Within closed private networks this may be possible, however, once we move into public networks the prioritization becomes less predictable. Some network providers might not even provide QoS or bit rate shaping, potentially resulting in a complete mess.

Packet prioritization relies on switchers and routers performing deep analysis of the packet to extract the necessary DSCP values, and even looking at the stream itself to determine whether it is audio or video, resulting in delays and bottlenecks. This is further compounded when we look at encrypted packets as the router will not be able to decode any of the payload and will not determine if it's switching video or audio.

Multi-Protocol Label Switching (MPLS) is designed to overcome these problems. MPLS is provided by a network supplier and is largely transparent to the end user, as a packet enters the MPLS network the ingress router adds a label to the packet, this is used by all subsequent routers within the providers' network to prioritize and route that packet.

Routing efficiency is improved as switchers use the label for prioritization and routing, but don't interrogate the packet further, thus improving throughput and reducing complexity. As QoS is an intrinsic part of MPLS it forms a fundamental part of the routing method instead of being an unwelcomed add on.

Interoperability between network providers is maintained, to be part of an MPLS system all providers must agree on using the same QoS strategies to guarantee streaming of video and audio.

MPLS can adopt diverse types of layer 2 connections including Ethernet, ATM and DSL. Combined with multi-vendor interoperability, MPLS is extremely flexible and lends itself well to broadcasting, especially when backhauling cameras and microphones from remote outside broadcasts.

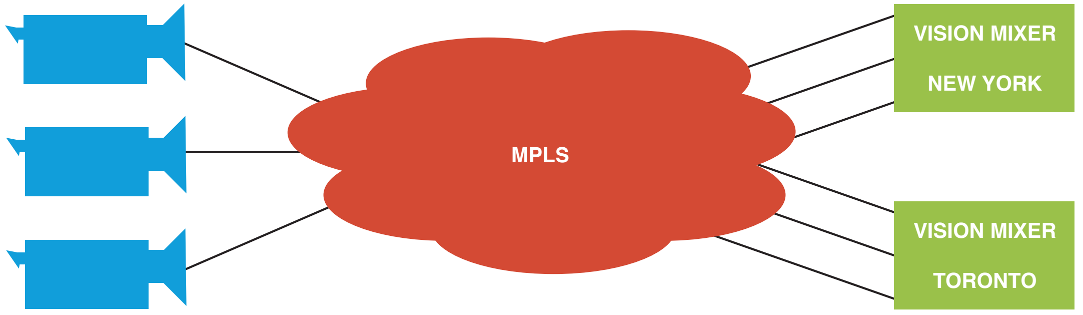

Figure 2 - MPLS service providers abstract away the network and guarantee a Quality of Service for video and audio streaming.

As IP grows within broadcast facilities QoS will become a fundamental consideration, and decisions on whether to use protocols such as MPLS will need to be taken at an early stage, especially when choosing network providers who may or may not be able to provide MPLS.

14

Delay Monitoring

We use buffers to reassemble asynchronous streams so we must measure how long individual packets take to reliably get to the receiver, and the maximum and minimum delay of all packets at the receiver.

Video and audio monitoring in baseband formats is well established for levels, noise, and distortions. Television monitors provide subjective visual checks and objective measurements can be taken using waveform monitors. Audio is similar, loudspeakers and headphones provide subjective checks and PPM's, VU's and loudness meters provide objective verification.

IT subjective information consists of determining the user experience; how long does it take for a web page to respond to a mouse click? And how fast will a file transfer? IT networks use packet analysis tools such as Wireshark to look closely at the packets, and IPerf is used to find absolute maximum data rates of network links.

Video and Audio brings a new dimension to monitoring for the IT department. Not only are we concerned with how to measure the video and audio, we must analyze the time it takes for an IP packet to arrive at a destination, and the variance of all other packets in the stream. If they take too long, then the receiver will drop them from their decoding buffer and cause signal corruption.

High level audio and video monitoring will always be important. Evangelists have often proclaimed that in a digital world we don't need audio level monitoring as the signals don't suffer the same distortion and level problems as analog lines. Anybody working at the front end of a broadcast station will tell you the reality is somewhat different.

In the past, broadcast engineers have had the luxury of assuming the underlying network is robust and solid. An SDI distribution system will provide nano-seconds of delay at 3Gbps, and a twisted pair balanced audio system will have similar delays with virtually no dropout.

IP networks are very different. They're designed with the assumption that there will be packet loss and variable delay. As IP networks are resilient and self-healing, it's possible and likely that IP packets streamed across a network will take different routes and some won't get there at all. If a router fails then the resilience in a network will send subsequent IP packets via a different route, often longer than the original. If the first router recovers, then the IP packets could be sent over this shorter link, resulting in packets being received out of sequence.

Significant variation in transmission of packets occurs due to the queueing that takes place in switches and routers. In integrated IP networks transfer of all kinds of data is taking place, from accounts transactions to office files; video and audio is competing with these to get to their destination.

Receiver buffering is a straightforward way of dealing with delay and sequencing problems. A buffer is a temporary storage area of computer memory where packets are written out of sequence and in varying time. The receiver algorithm reads the packets and pulls them out of the buffer in sequence and presents them to the decoding engine.

Buffers are a trade off between delay and validity of data. The longer the buffer the more likely it is to receive packets that have taken a disproportionate time to travel. However, the read-out algorithm has a delay of the time of the latest packet. In effect, the bigger the buffer, the longer the delay.

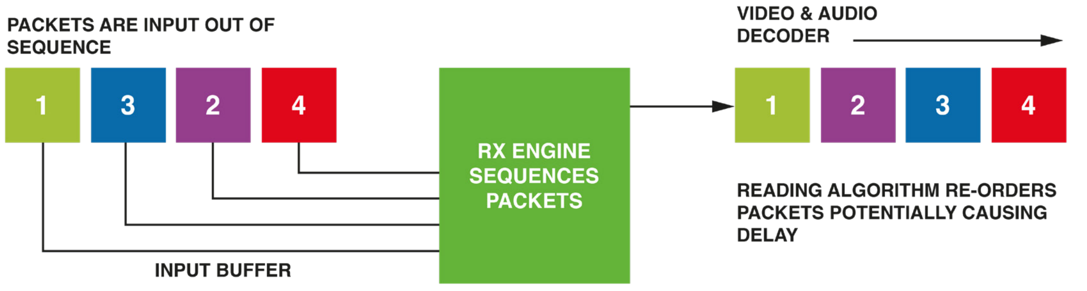
PACKETS ARE INPUT OUT OF SEQUENCE
1
3
2
4
RX ENGINE SEQUENCES PACKETS
INPUT BUFFER
VIDEO & AUDIO DECODER
1
2
3
4
READING ALGORITHM RE-ORDERS PACKETS POTENTIALLY CAUSING DELAY

Figure 1 - Buffers are used as a temporary store to re-sequence packets.

Dropped packets are caused either through congestion in a switch or router, or interference on a network cable. Congestion occurs when too many packets arrive at the router's inputs too quickly and the router cannot respond to them quickly enough, or the egress port becomes oversubscribed. Much processing goes on inside a router or switch, the more features the device provides, the more chance there is off packet loss.

This is one of the reasons IT engineers try and use layer 2 switchers (Ethernet) wherever possible. They use look up tables to decide how to send the frame based on the Ethernet packet header destination address, this is relatively simple and can be achieved in almost real-time using a bitwise comparison in an FPGA (Field Programmable Gate Array).

As a router needs to dig deeper into the Ethernet header or IP packet it requires more processing power and the potential for packet loss increases. This is one of the area's IT engineers tend to quickly gloss over, working on the assumption that congestion occurs infrequently, and when it does TCP and FTP type protocols will fix the problem as they will resend any lost packets.

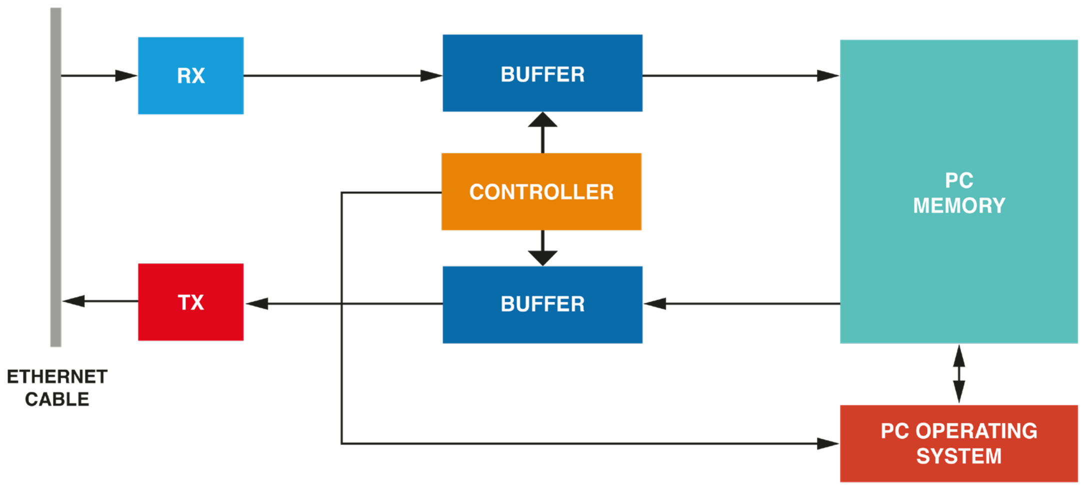

Figure 2 - Computer network interface cards introduce delay and jitter.

In broadcast television, we cannot afford to drop even one packet. ST2022-5 incorporates FEC (Forward Error Correction), but this isn't really designed to take the place of TCP or FTP to fix large error caused by congestion, and relying on it to do so could result in unpredictable results.

Consequently, we are interested in two network measurements; how long individual packets take to reliably get to the receiver, and what is the maximum and minimum delay of all packets at the receiver. On the face of it this sounds like an easy measurement to make using analyzers such as Wireshark. However, PC protocol analyzers rely on receiving data from the NIC (Network Interface Card) and time taken for the operating system to move data from the NIC to the main processor.

NIC's have built in buffers that are used to receive and transmit data to the Ethernet cable or fiber. For transmission, they provide a temporary store should a collision be detected on the Ethernet link, and the packet needs to be transmitted, and for receiving they hold packets until the processor has time to copy them to main memory and process them.

The buffers and operating system incur further delay into the system and make critical measurement very difficult. We cannot be sure whether we are measuring the time taken through the network, or the time taken to process by the measuring systems OS and NIC. This is one of the occasions where a hardware solution gives consistently better results than software tools.