THE BROADCAST BRIDGE

Connecting IT to Broadcast

A Brief History of IP

Essential Guide

EG

ESSENTIAL GUIDES

Introduction from Wheatstone

AoIP, aka Audio Over Internet Protocol, has come of age. The term is on everyone's tongue and in spite of its glib use in daily broadcast technical parlance, not everyone has a thorough grasp of what it really means. And while there are many who use this term knowledgeably as a matter of course, we're sure there are an equal number who might welcome a no-nonsense grounding in the history and technical background behind this now dominant technology.

When the editors of The Broadcast Bridge proposed a multi-part history of Internet Protocol (essentially an "IP-101" series of articles aimed at the broadcast engineering crowd), Wheatstone welcomed a chance to participate. The result is this eBook. While old hands may find much of the information that follows obvious, we are equally sure many would appreciate a chance to brush up on the history, implementation and technical challenges of this worldwide protocol—particularly as applied to broadcast audio. We are proud to be involved in this, and hope the contents will be of use to many of you out there who are tasked with applying these systems to your own particular facilities.

IP protocol for broadcast imposes a more stringent set of requirements than those originally presented in the general distribution of data for the internet. It comes down to timing, and the fact that IP data transmission is neither sequential nor consistent compared to the old analog methods. Internet Protocol digital data takes a continuous signal and breaks it into smaller packages, which are transmitted from source to destination at the whims of the transmission system itself. If the system is operating within its designed capacity, data is transmitted smoothly and evenly.

Wheatstone utilizes intelligent AoIP networking as well as virtual, augmented, and hardware solutions to provide broadcast studios with the tools they need to meet modern challenges.

But if the system suddenly has more data than it can handle at the moment—and we do mean moment—we're talking milliseconds here—data packets may get slowed down, re-ordered, or even lost completely as they wait for the transmission system to accommodate the overload. If the data content is audio (human speech for example) this garbling of packets can result in an unintelligible signal. The technical challenge here is to re-order the errant data correctly so the destination result is a faithful reproduction of the original source content.

This re-ordering takes time. In other words, it requires a buffer, which inevitably introduces delay. If an audio signal is being transmitted in conjunction with a video signal (read Broadcast) the two can end up so out of sync as to be rejected by the target audience. One or both signals must be buffered and successfully recombined for broadcast to work. This is where AoIP comes in.

We hope the following content will give interested readers a basic idea of how all this works in IP networked systems for broadcast, and some understanding of the solutions and techniques being used to make AoIP the defacto networking protocol of choice in audio broadcast today.

Sponsored by Wheatstone

A Brief History of IP

By Tony Orme, Technology Editor at The Broadcast Bridge

Audio Networks

IP networking is taking the radio and broadcast industry by storm, but as a method of distributing data, it has been available since the 1970's. So, what are IP Networks? And why have they become so popular recently?

The history of IP is closely linked to that of the world wide web. Back in 1969 the U.S. Defense Department's Advanced Research Projects Agency Network (ARPANET) funded research into much of the suite of protocols that make up todays internet, specifically IP and TCP.

In 1972 the Internetworking Working Group (INWG) was formed to standardize protocols and by 1973 the University College London (UK) and Royal Radar Establishment (Norway) connected to ARPANET and global connectivity was born, along with the first accepted use of the term "Internet".

Sponsored by Wheatstone

Wheatstone

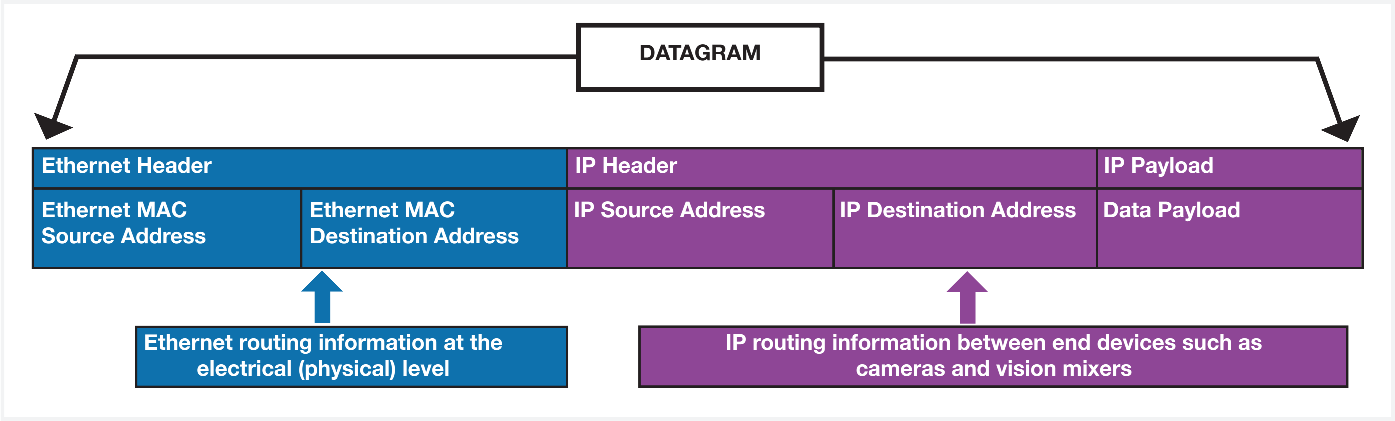
DATAGRAM
Ethernet Header
IP Header
IP Payload
Ethernet MAC Source Address
Ethernet MAC Destination Address
IP Source Address
IP Destination Address
Data Payload
Ethernet routing information at the electrical (physical) level
IP routing information between end devices such as cameras and vision mixers

Diagram showing the Ethernet and IP headers in a datagram.

TCP is Born

1974 witnessed the birth of Telenet, a commercial version of ARPANET and Vincent Cerf and Bob Kahn publish “A protocol for Packet Network Interconnection”, the precursor to TCP. By 1982 TCP and IP emerged as the protocol for the internet and combined with public Domain Name Systems (DNS), established web domains formed to associate names such as .net and .com to IP addresses.

What soon became clear, and probably the single most important reason that has made IP so ubiquitous and adopted throughout the world, is the protocol is independent of the underlying hardware it is being distributed over. For example, IP can reliably work on ethernets’ 10BASE-T as easily as it can work on fiber channels 8GFC. The line-speeds may be significantly different, but the IP protocol works over both equally.

Flexibility is Key

The ramifications of this flexibility soon become clear as we look at distribution between different service providers. A backbone between New York and London may consist of a high-speed fiber running under the Atlantic Ocean. But at either end two different telco’s may connect to the fiber using 100Base-T ethernet. If they both comply with the IP and IEEE 802.x ethernet protocols, they will seamlessly work together.

IP addressing takes this one stage further as one computer can communicate with another computer on the other side of the world without understanding or caring how the routing and connectivity takes place.

Contrast this with broadcast and radio stations using traditional communications. To send an audio circuit from London to New York, the broadcaster would need to understand the routing and connectivity used as well as the massive administrative burden in making it happen. Broadcast stations at each end need specialized connections to the local telco such as 2-wires, balanced pairs or SDI feeds. All this significantly adds to the cost and complexity.

IP Separates Data from Hardware

Inside the station, IP still offers many advantages by simplifying the operational aspects of the system and providing greater flexibility. IP abstracts the data from the underlying hardware meaning any audio, video and meta-data is also distanced from the hardware resulting in a much more flexible system. As new audio and video standards are released to keep up with ever increasing audiences, they can be more easily deployed on an IP system.

It’s worth remembering that ethernet is a distribution system that could be used to transfer audio and video, and to some extent is responsible for the success of IP. However, ethernet has some limitations that make use outside a station difficult. For example, there is no concept of masking, so routing becomes much more complex, and there is no flow control or packet re-send mechanism for lost and dropped packets.

Timing is Critical

One of the major challenges of distributing media over an IP network is that of timing. Broadcast and radio circuits are incredibly expensive due to the guaranteed levels of quality of service they provide. We know an audio balanced 2-wire will have a bandwidth of 20KHz and latency less than one millisecond, but we pay for this through cost and lack of flexibility.

Due to the intricacies of routing technology, IP packets do not always take the same path and can arrive at their destination out of sequence, and the strict limits of sampling clock used in the original analogue to digital conversion is lost.

Packets Must Be Sequenced

Resequencing is straightforward as packets have an incremental counter to identify each packet as it is sent. Therefore, the receiver can easily put the packets in the correct order before they are presented to the digital analogue converter or decoder.

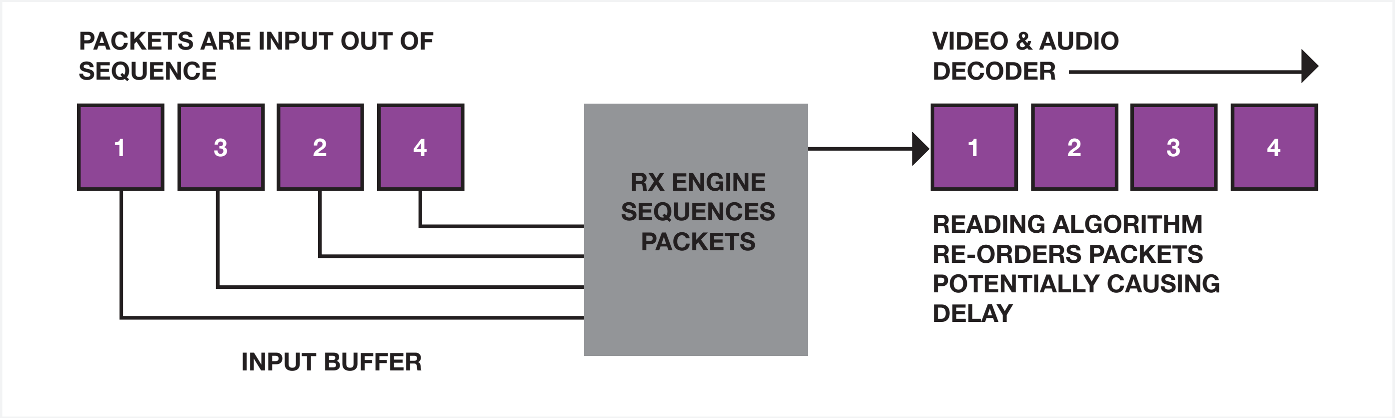
PACKETS ARE INPUT OUT OF SEQUENCE
1
3
2
4
INPUT BUFFER
RX ENGINE SEQUENCES PACKETS
VIDEO & AUDIO DECODER
1
2
3
4
READING ALGORITHM RE-ORDERS PACKETS POTENTIALLY CAUSING DELAY

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

Audio and video streams over IP are sent with User Datagram Packets (UDP) that are encapsulated in IP packets. UDP adds the concept of source and destination ports as well as checksums to the video and audio payload data.

Ports add further granularity to IP addressing to allow a single device with one IP address to receive or send multiple streams. For example, a sound console might only have one IP address, but would have sixty-four separate port numbers for each of its channels.

Unlike with MADI or ASI-MPEG, the piecewise timing relationship in IP is not respected and further eroded by the network itself. Real Time Protocol (RTP) overcomes this by sampling a clock in the encoder to give each payload of audio or video a unique timestamp. The receiver reconstructs the audio or video stream using these timestamps and rebuilds the audio or video stream.

Data Multiplexing

Network latency is a new concept to television and radio. Routing switches direct packets between ports to transfer data around different networks. This is a form of multiplexing but on a much more complex scale. Consequently, blocking can occur within the switch as packets are delayed in buffers leading to variable latency.

Buffers in receivers are used to negate the blocking effect of switches but if used too much will result in increasing the end to end latency. Great care must be taken to keep RTP timestamps within the buffer size.

Watch the Backbone Speed

Some network latency is inevitable, but we work on the principle that the backbone speed of the network is sufficiently large to keep it small. But we must continue to monitor latency and make sure it keeps within tolerable parameters, especially when controlling remote equipment.

IP has brought unimaginable cost-effective flexibility and scalability to television and radio stations. And as adoption of the technology expands vendors will bring better solutions to make best use of the flexibility IP offers so we can continue to meet the ever-increasing demands of viewers and listeners.

Routing IP Networks

Network routing is a phrase that is bandied about broadcast forums liberally. But what exactly does it mean to route an IP datagram? And why is it important for broadcast and radio stations?

In traditional facilities, SDI, MADI and AES are the common forms of signal distribution. We use a combination of X-Y routing matrices and patch-cords to distribute signals around a facility. One to one connectivity is required to maximize transfer of power between source and destination through impedance matching.

If we need to send a signal to more than one destination, such as a sound console output to multiple monitoring loudspeakers, some form of distribution amplifier is used. And using one-to-one mapping provides us with a simple overview of which microphones, console channels and monitoring is routed together.

SDI, AES and MADI Integration is Complex

Transferring signals between different formats starts to get challenging as we interface SDI, MADI and AES signals together. If we need to extract channel two out of a MADI stream and insert it into channel one of an embedded SDI stream, then an array of embedders, multiplexers and de-embedders is required, often making a system extremely complex and difficult to visualize or remedy if an error occurs.

IP routing solves many of these problems for us as audio and video signals are treated as pure data, and one of the greatest strengths of IP is that it neither knows nor cares what type of data it is transporting in its payload.

IP Allows for Mixed Formats

Diverse types of audio and video formats can be inserted into the IP payload, for example, audio sampled at 48KHz with a bit depth of 24bits can be distributed on the same network as audio sampled at 44.8KHz with a bit depth of 16bits. The same is true of video, signals of 720p29.97 can be distributed on the same IP network as 1080i59.94, or 4K-60P-444.

There needs to be a point where the signals are converted to a mezzanine format for processing, but this only happens at a production interface leaving the audio or video signal to be stored and distributed in its native format to maintain the highest quality possible.

In IP networks, terminal devices such as servers, microphones, production switchers and mixing consoles are referred to as “host” devices. A host device must have a unique IP address so a router, or switch knows where to send the signal to.

Another often unappreciated aspect of an IP system is that the network is bi-directional, that is, signals can be sent to and from a host. It’s entirely possible to have a microphone with an IP interface that will output audio across its IP connection, but also receive configuration information from a sound console or control application. Microphones could have configurable pre-amps or equalization built directly into them to further improve quality and provide greater flexibility.

Routers are referred to as layer-3 devices and switches as layer-2. The difference is due to the level at which they transfer data packets, and the numbering system is derived from the Open Systems Interconnection model, a reference model describing how data can be distributed throughout a network. Layer-2 refers to the data-link layer, and layer-3 refers to the network layer for routing.

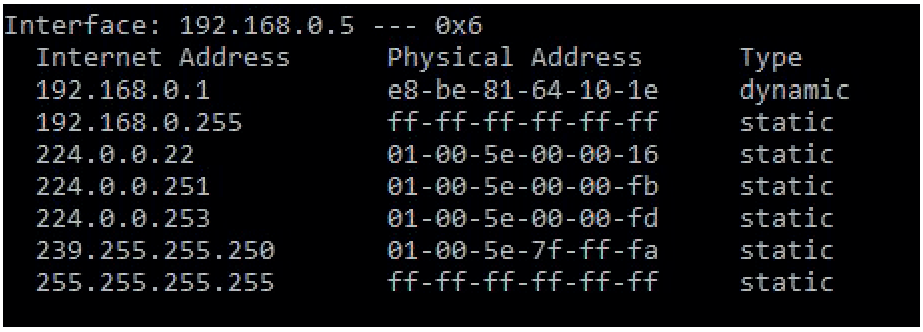

```
Interface: 192.168.0.5 --- 0x6
  Internet Address      Physical Address      Type
  192.168.0.1           e8-be-81-64-10-1e     dynamic
  192.168.0.255         ff-ff-ff-ff-ff-ff     static
  224.0.0.22            01-00-5e-00-00-16     static
  224.0.0.251           01-00-5e-00-00-fb     static
 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

Resolution of IP and ethernet MAC addresses on a typical network.

MAC Addresses are Universally Unique

Ethernet is usually the dominant data-link layer used in broadcast facilities, but other formats are used such as Synchronous Optical Networking (SONET) and WiFi, especially as broadcasters connect to the outside world through telecommunication providers (Telco’s).

Each host device must have its own unique Media Access Control (MAC) address so that the device can be identified. Although IP also has an addressing scheme, it differs from the MAC as its address is not absolutely attached to that device, this is to allow faulty units to be replaced without major configuration. For example, if a microphone has an IP address of 10.10.56.19, and it goes faulty, it can be replaced with another microphone with the same IP address. However, the MAC addresses will be different but the Address Resolution Protocol (ARP) is designed to deal with this.

The Institute of Electrical and Electronic Engineers (IEEE) administers the issuing of MAC address numbers, but it is the responsibility of the manufacturer to make sure that the MAC address is programmed correctly during manufacture. Once installed, the MAC address should not be changed.

Layer-2 Keeps Switching Fast

IP datagrams are encapsulated within the payload of an ethernet datagram and switching of data streams takes place independently of the IP address to keep switching speeds high. Routing at the IP address level requires the ethernet header to be decoded and then the IP header, thus increasing switcher processing times and reducing the time taken to transfer the datagram.

An ethernet switch has many ports to connect to other ethernet switches or host devices such as microphones and sound consoles. If, for example, ten microphones are connected to the first ten ports of an ethernet switch and the eleventh port is connected to a sound console, the switcher will need to know it must send the first ten microphones to the sound console, in effect the switcher is acting as a multiplexer, as ten microphones are routed to one sound console ethernet connection.

The quality and speed of the switch impacts heavily on how effective this multiplexing process is. In broadcasting, we take for granted that the throughput of an X-Y matrix guarantees one hundred percent data delivery, but in IT, this is not always the case.

Non-blocking switches guarantee that the data received on all its ports can be sent on the designated port without packet loss. For example, if a switch consists of sixteen ports, each with an input and output speed of 10Gbps, then the total input speed is 160Gbps, so the switch manufacturers guarantee that all the packets can be sent to the outputs of the ports totalling 160Gbps.

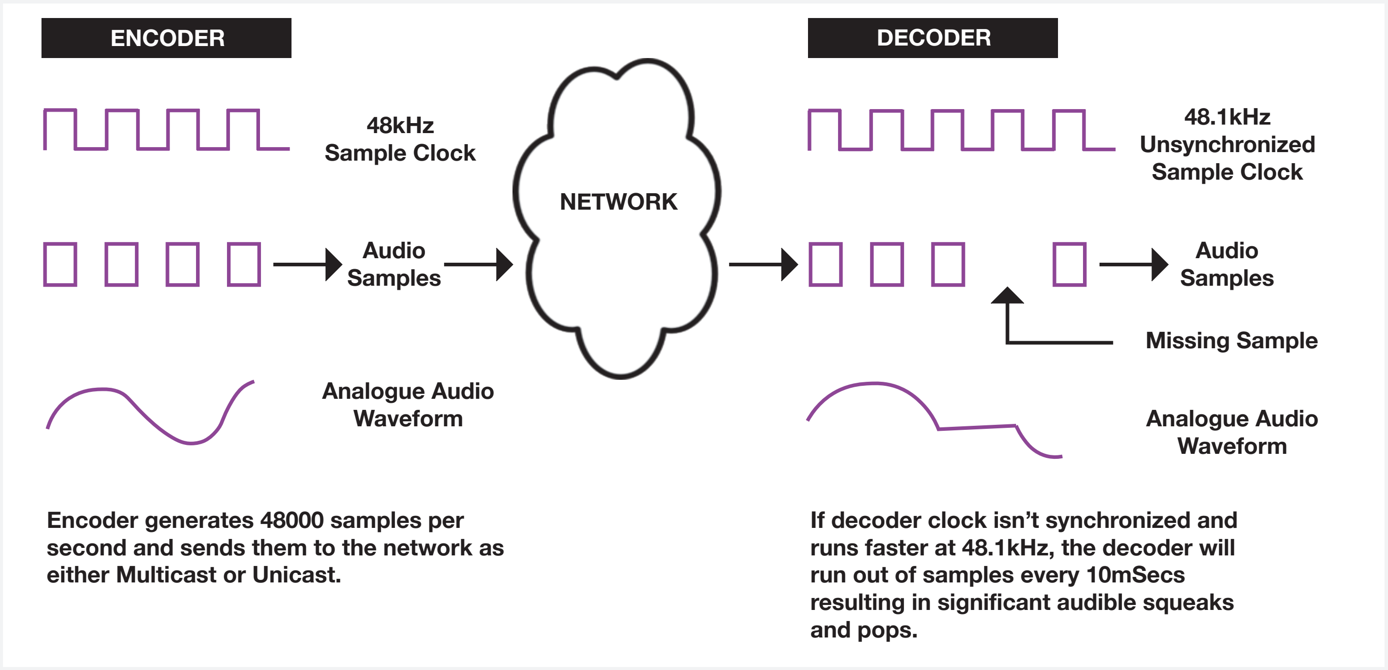

Encoder generates 48000 samples per second and sends them to the network as either Multicast or Unicast.

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.

Non-blocking switches must be used to guarantee accurate timing in IP systems.

Use Non-Blocking Switches

Non-blocking switches are usually only available on high-end switches. The alternative to "non-blocking" is the "blocking" switch. As the design cannot guarantee all packet delivery, some packets will either be lost or unacceptably delayed. Blocking switches will work in some broadcast applications, but a great deal of attention must be paid when designing for data throughput, otherwise packet loss, and excessive delay and jitter will occur.

When the audio or video stream is sent to another building or district, IP routing will be used as the host will need to know where the transmission is going. For example, the sound console might have an IP address of 10.10.76.01 and its destination address might be a transmitter site with address 10.1.100.89, the IP router will have been configured to send all datagrams destined for 10.1.100.89 to the port connected to its telco.

Ethernet switching is generally faster than IP routing and if non-blocking devices are used then packet jitter and delay are greatly reduced. But to distribute signals outside of the broadcast facility, and to maintain maximum flexibility for addressing, IP routing is used, although it's kept to a minimum to help keep transfer speeds high, and jitter and delays low.

Putting It All Together

Building reliable, flexible IP networks requires an understanding of infrastructure components and the interoperability of systems that run on them, especially when working in fast-paced, dynamic studios. Protocol interfacing is relatively straightforward, but as we investigate application level connectivity further, systems become more interesting.

Simplistically, layer-2 ethernet switching provides faster data throughput than IP routers; this leads to reduced packet jitter and delay. But layer-3 IP routers provide greater flexibility, especially when distributing to other districts or cities via telcos (telecommunication providers).

Excessive packet jitter and delay cause problems with signal reconstruction at the receiver. The standard method of removing jitter and delay is to use a buffer; each packet is written to continuous memory, but is then read out in sequence and in time.

However, if a packet arrives too late or with too much jitter, it will not be available to the decoder engine at the appropriate time. A packet that is too early suffers a similar issue.

The receiver buffer only has a finite length. If it's too long, excessive delays in decoding occur, resulting in lip-sync errors. If it's too short, packets will quickly become out-of-date, resulting in splats and pops in audio, and frame freeze and break-up in video.

Administration Challenges

Although IP routing provides a great deal of flexibility, it also poses some interesting administrative challenges. When connecting a Windows laptop to a network, a great deal of configuration unseen by the user goes on between the laptop and the network. The process used in IT networks is called Dynamic Host Control Protocol (DHCP) and provides the laptop with an on-demand IP address.

Sponsored by Wheatstone

The network's DHCP server keeps a pool of IP addresses in its database, and each time a computer connects to the network, it will provide an IP address for it. When it disconnects, the allotted IP address will be returned to the free-pool so it can be allocated when another device connects to the network.

Without DHCP, system administrators would need to actively issue IP addresses. This is relatively straightforward with fixed devices such as desktop PC's or rack servers. However, the system becomes more complex when portable devices are used as they may connect and disconnect to a network many times throughout the day – requiring the system administrator to issue new IP addresses. Clearly this is unworkable.

Broadcast television and radio face similar challenges as each host device must have an allocated, unique IP address. Allocating the same address to two devices gives rise to a phenomenon called IP-ghosting causing all kinds of problems with a network; systems assume and specify that all devices must have unique IP addresses.

IP-Ghosting Must be Avoided

But there is no mechanism within the IP protocol to stop IP-ghosting from occurring, it is the responsibility of system administrators and engineers configuring a network to make sure all devices have unique IP addresses within a domain.

Automated options like DHCP are available, such as the plug-and-play method offered by Wheatstone to automatically configure their BLADEs – distributed IP audio processing modules that provide an array of audio services including mixing, level control and equalization.

Dedicated servers have the advantage over PC servers as they have dedicated hardware that can process audio in real-time with very little delay. Digital Signal Processors (DSP's) with localized near-chip memory and streamlined, pipelined data flows process audio in just a few samples, resulting in very low processing delays.

Scalability and Flexibility

Distributed processing empowers scalability but traditionally, broadcast infrastructures were designed for rigid peak demand to allow for the worse-case-scenario use-case. It's almost impossible to predict requirements years ahead, so systems tended to be over-specified and over-designed, resulting in unnecessarily inflated costs and complexity.

It is possible to use PC rack or cloud servers to provide similar functions to dedicated audio processors, but their data throughput is inherently slower due to the buffering required. Computer servers based on PC architectures are designed for generic processing of data. DSP's specialize in short-loop high-bandwidth processing found in audio algorithms such as filters and gain control.

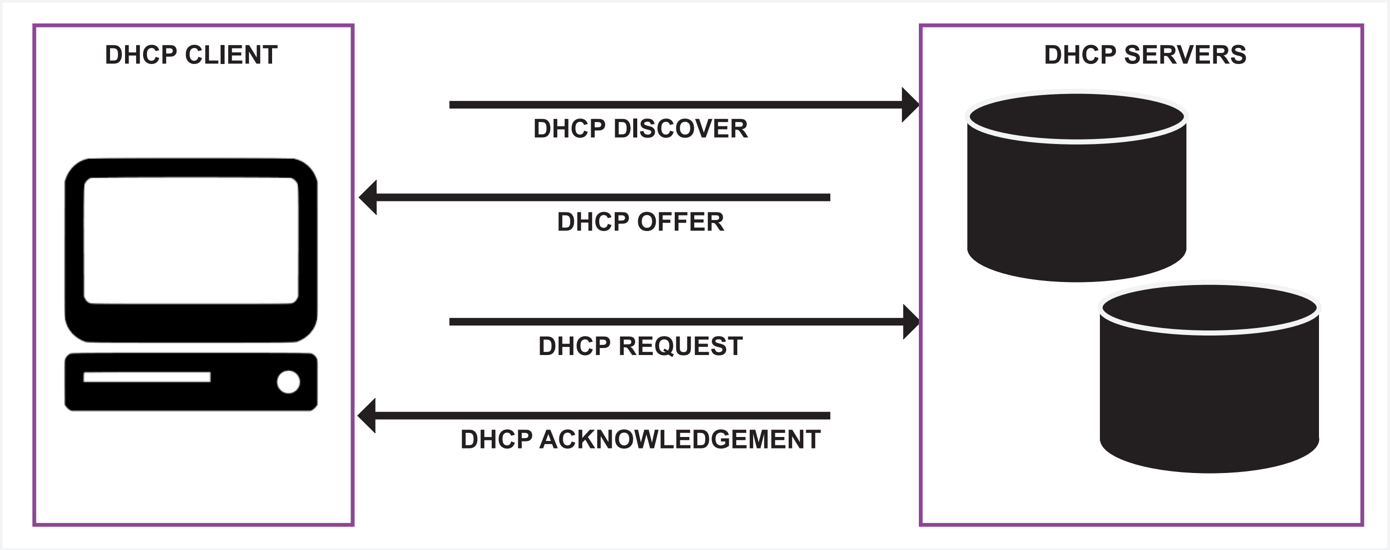

To reduce the administrative nightmare of allocating IP addresses, DHCP is used.

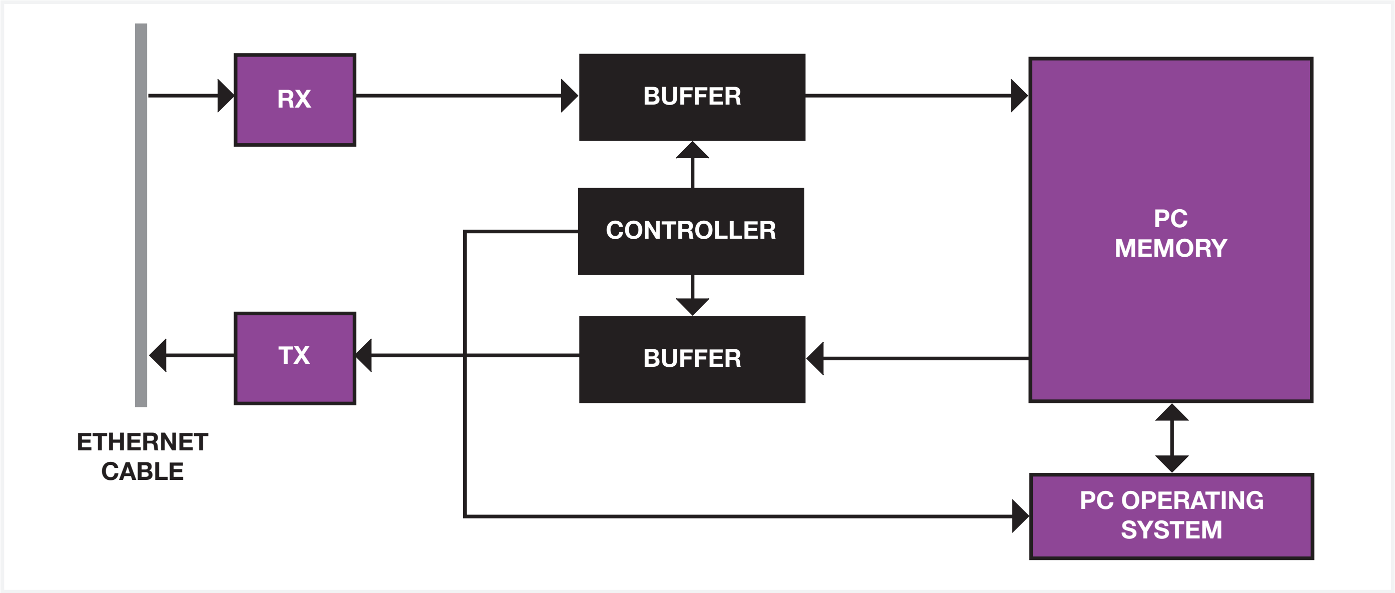

The PC's network interface card causes IP packet jitter due to its use of buffers.

Generic PC's Compromise Speed

PC architectures are designed to facilitate many tasks, and use round-robin task switching to achieve the appearance of parallel processing. Generic operating systems provide subsystems that interface to computer screens, keyboards, and network interface devices; these are all relatively slow and can cause blocking within the processor, especially if audio and video samples are being stored on disk drives.

To overcome these issues, PC architectures adopt buffering strategies. Each slow device stores its data in a memory buffer, and is read in when allowed by the operating system. Programmers have little, or no control over these processes, so data input and output take place at the discretion of the operating system. Although data is processed in real-time, throughput is delayed and can be greatly compromised.

Dedicated hardware systems such as BLADEs have hardware specifically designed to process audio with bespoke operating systems dedicated to maintaining fast data processing and throughput.

Auto Back-Up

Each BLADE uses a common protocol to detect similar devices on a network, so the possibility of IP-ghosting is greatly reduced and multi-studio configurations can be easily established. Sound consoles detect BLADEs within a network and determine their configuration. For example, they might have microphones connected in studio one, but MADI interfaces connected to the BLADE in studio two.

Localized configuration databases in each BLADE store not only their own configurations, but also those of other compliant devices on the network. If a new BLADE is connected to the network, the other devices will recognize it and send their own databases as well as requesting a new one. Using distributed computing in this way means there is no single point of failure for the configuration database and the system automatically backs itself up.

Building an IP network is only part of making a system work reliably. Interoperability plays a huge role in integration, and anybody installing a system should consider this from the outset. Dedicated vendor-specific solutions solve many of these problems and continue to build on the major benefits of scalability and redundancy in IP systems.

Wheatstone – The Sponsors Perspective

IP Audio Is Not A System

IP audio can bring unbelievable adaptability and extendability to broadcast production. But you still need a way to bring that audio in, then manage, process and control it. This is Wheatstone's area of expertise. Our AoIP networks are working *systems* that integrate audio control, processing, routing and distribution in an intelligent manner, programmed by logic suited to your own particular content and workflow. Over the years we have developed an entire family of products that integrate together smoothly and transparently.

Click on the image below to learn more about what we have to offer.

Should you want more in-depth information about our AoIP networked products, we encourage you to download these WheatNet-IP brochures:

- AoIP for Television Brochure
- WheatNet-IP for TV Overview and Planning Guide

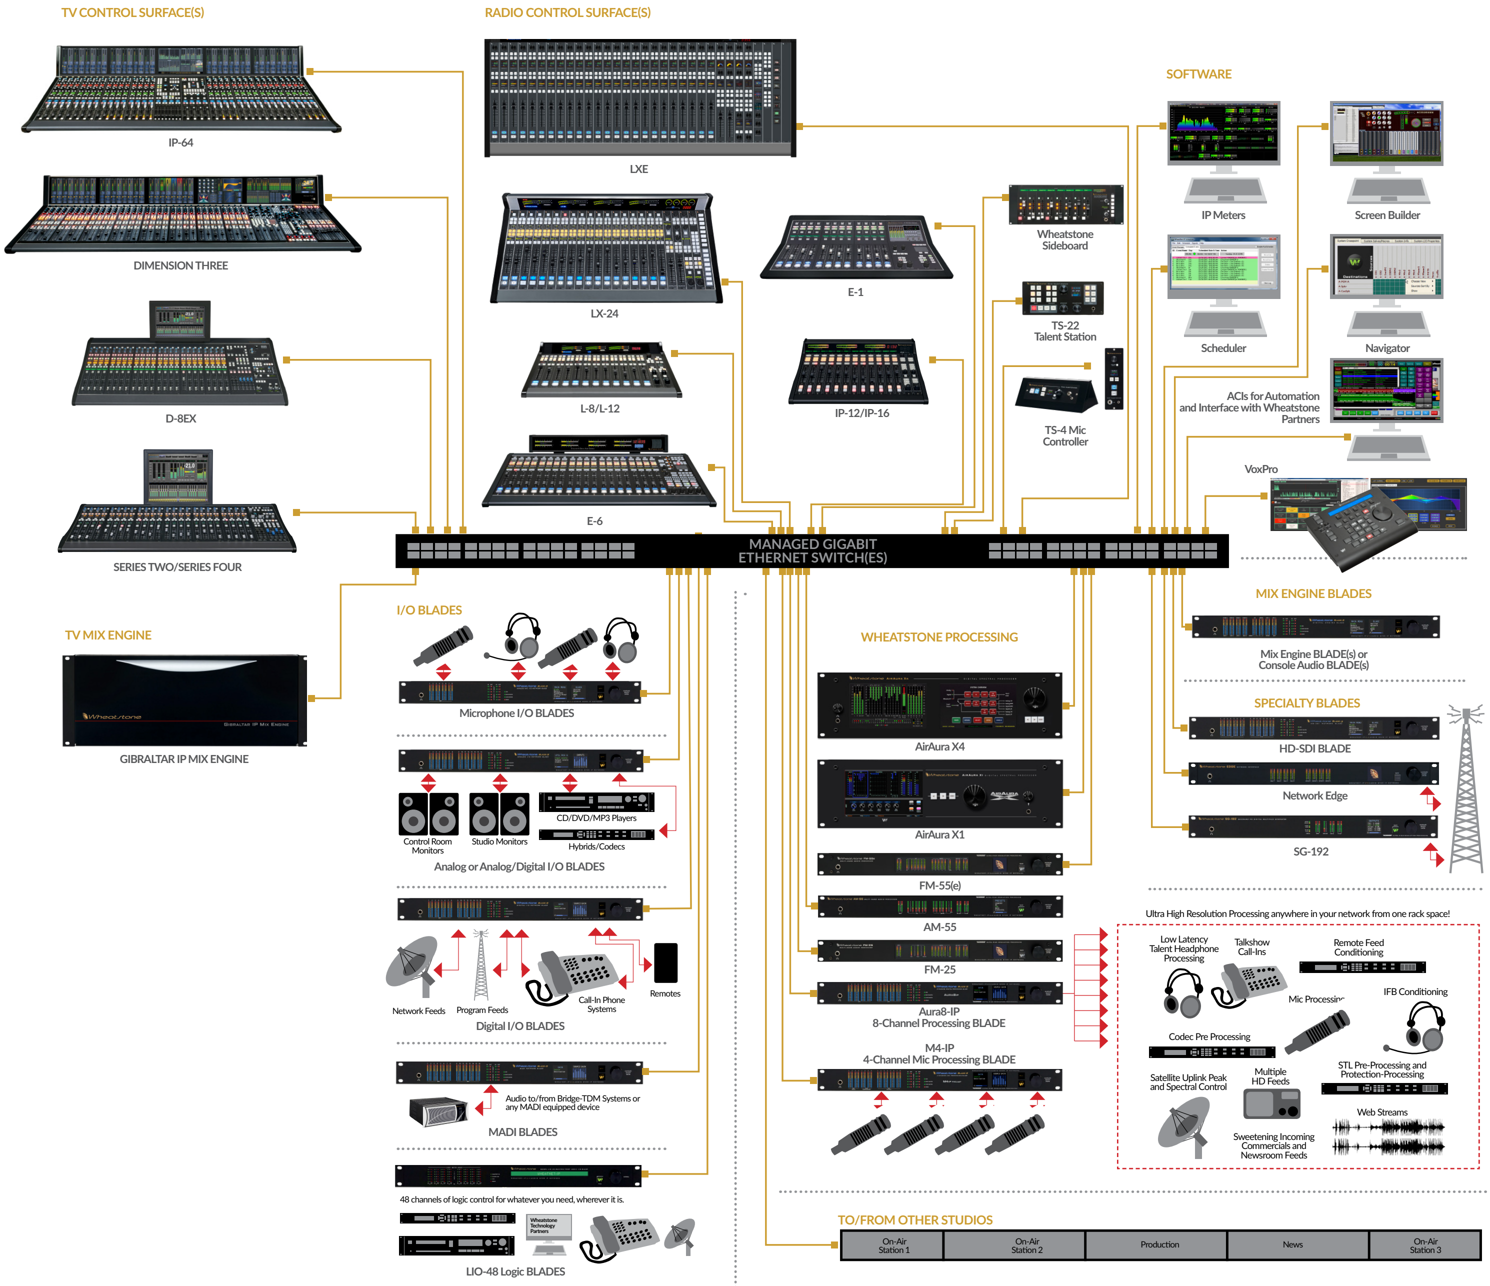

Sponsored by Wheatstone

WP
WHITE PAPERS

EG
ESSENTIAL GUIDES

MEDIA

CS
CASE STUDIES

Find Out More

For more information and access to white papers, case studies and essential guides please visit:

thebroadcastbridge.com

7/2018

Sponsored by Wheatstone

Wheatstone