


CALREC AUDIO NETWORKING PRIMER



Introduction to professional audio networking

Published in association with



Keeping audio professionals up to speed in the fast-changing world of sound and image.

Calrec Audio Ltd

Nutclough Mill
Hebden Bridge
West Yorkshire
England UK
HX7 8EZ

Tel: +44 (0)1422 842159
Fax: +44 (0)1422 845244
Email: enquiries@calrec.com

calrec.com

No part of this manual may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and scanning, for any purpose, without the prior written consent of Calrec Audio Ltd.

Whilst the Company ensures that all details in this document are correct at the time of publication, we reserve the right to alter specifications and equipment without notice. Any changes we make will be reflected in subsequent issues of this document. The latest version will be available upon request. This publication is for International usage.

Calrec Audio Ltd reserve the right to change specifications without notice. E & O.E.

The established policy of Calrec Audio Ltd. is to seek improvements to the design, specifications and manufacture of all products. It is not always possible to provide notice outside the company of the alterations that take place continually.

Despite considerable effort to produce up to date information, no literature published by the company nor any other material that may be provided should be regarded as an infallible guide to the specifications available nor does it constitute an offer for sale of any particular product.

Apollo, Artemis, Alpha, Sigma, Omega, Zeta, Hydra Audio Networking, Hydra2 and Bluefin High Density Signal Processing (HDSP) are registered trade marks of Calrec Audio Ltd.

All other trade marks are acknowledged.

© 2013 Calrec Audio Ltd. All Rights Reserved.

AUDIO NETWORKING PRIMER

FOREWORD

FOREWORD

This Primer is designed, like the Audio Primer that preceded it, to explain often complex concepts in straightforward, approachable terms. With this Network Primer, our aim was to explain the background and technology behind data networks with specific reference to its application in broadcasting, such that a broadcast engineer with no formal training in the field of computer networking technology might gain a clearer understanding of the subject.

One semantic problem should be dealt with right away: 'network' is a word that already had a meaning in broadcast long before the rise of affordable computer technology — particularly in the USA, where it usually refers to a broadcasting company itself. To be utterly clear and unambiguous: this Primer aims to explain the increasing use of data networking technology in broadcast applications, examines the benefits that this technology can offer forward-thinking modern broadcasters, and considers where it may take the world of broadcast in the future.

Hebden Bridge, Spring 2013

NETWORK PRIMER

CONTENTS

Introduction	9	Chapter Four	29
Introduction	11	Calrec's Hydra2	30
Chapter One	13	Reliability Vs Compatibility	30
The Benefits of Networking	14	Latency & Redundancy	31
Interoperability – The Holy Grail	14	H2O & Network Management	32
Routing and Cost Benefits	16	Compatibility with Other Standards	33
Networks and Resilience	18	Chapter Five	35
Networks and Contingency Planning	19	Control Protocols & Networking	36
Networks and Control Protocols	19	SW-P-02/08 & RAP	37
Chapter Two	21	CSCP & Automation	37
Some Technical Background	22	EMBER	38
Layer 3 Protocols	23	The Future of Broadcast Networks?	38
Layer 2 Protocols	23		
Layer 1 Protocols	23		
Chapter Three	25		
Routes to Interoperability	26		
AES-X192	26		
The AVnu Alliance & AVB	26		
ALC Networx & Ravenna	27		

NETWORK PRIMER

INTRODUCTION

INTRODUCTION

The word ‘network’ has a long history; it was first recorded in the mid-sixteenth century, but its Germanic roots hint that it may well have been in use much earlier than that. Like much of modern English, it has gone through a variety of ever more abstracted meanings, from a description of a physical object to more figurative, intangible concepts. What began as a way of describing an invention for catching fish was later applied to similarly interconnected physical systems, including those of canals, railways, telephone wires – and eventually computers.

By the early 20th century, the word had come to refer to systems of related entities that were no longer physically connected, as with radio and TV transmitters belonging to a single broadcaster, or groups of business colleagues. In the last few years, we’ve had ‘wireless networks’ (which medieval speakers of English might have regarded as an oxymoron) and ‘social networks’ composed of users that ‘connect’ solely via data transmissions over the Internet; itself a network of physically disconnected, but nonetheless linked computers.

A similar process of abstraction has been taking place in recent years in the world of broadcast technology, transforming the hard-wired, localized studios of the past into more flexible, networked systems. Forty years ago, television studios consisted of cameras and microphones hard-wired into discrete hardware vision mixers, patchbays and audio mixing consoles which routed to specific video tape machines in one location. In such systems, a separate physical connection is required for each audio channel, whether from a microphone, mixing console, or recording device. Modern networked

broadcast systems offer more flexibility; all of the hardware is permanently connected to a data network, and the precise nature of the interconnections between the equipment can be redefined and/or reassigned at any time under software control, remotely if required.

Such ideas are not new, and small-scale proprietary networks of this type have existed in broadcasting for many years. But over the past decade, the declining complexity and improving cost-to-benefit ratio of implementing large-scale, networked broadcast systems and the ever-widening scope and capabilities of the technology, control protocols and equipment designed to work with such systems, has tempted more and more of the world’s forward-thinking broadcasters to move to networked systems.

At the same time, mixing equipment used in broadcast studios has undergone massive changes, playing far less of a central role. Once, mixing desks were the crucial nodes of a broadcast studio, through which all signals were routed and processed; now large consoles in networked studios are more usually assignable control surfaces, mere ‘clients’ connected to associated processing and routing units elsewhere on the network.

This further level of abstraction offers many benefits for broadcasters, as we shall see, including the ability to move projects swiftly from one studio to another when required by simply reassigning connections, or to control certain aspects of the mixing process remotely, without having to be in front of the console moving faders.

CHAPTER ONE

THE BENEFITS OF NETWORKING

THE BENEFITS OF NETWORKING

A networked broadcast studio, editing suite or transmission station isn't necessarily more efficient than a hard-wired one — indeed for smaller broadcasters, the costs associated with setting up a network can outweigh the benefits. But introducing networking into a large-scale broadcast environment can benefit the whole system; networked equipment is both more accessible and more flexible than its hard-wired counterpart.

To understand this, consider what audio patchbays did for arrays of hard-wired studio equipment. Studios function perfectly efficiently when equipment is connected directly, but time is always lost whenever new equipment is connected and its output needs to be made available to other parts of the system. Introducing patchbays to studios initially costs money and the time to wire everything up to the patchbay, and some studios chose to save themselves that time and expense.

However, once a patchbay has been integrated into a studio, any input can be routed to any output, producing significant long-term savings. Introducing new equipment to the system and giving it the same routing flexibility becomes a simple matter of connecting it to the patchbay.

When MIDI patchbays were invented, the routing of signals also became remotely controllable, and this greater flexibility and remote controllability is another benefit of networked studios.

But in the 2010s, data networks have sufficient bandwidth to do much more than route audio around; today it's theoretically possible to send broadcast video, audio and hardware control signals over a modern network. As we shall see, the technology is not yet quite at the

stage where all of these signals interface seamlessly with one another, but progress is moving swiftly in that direction.

Furthermore, because network technology has been developed to very high standards in the IT world, and because IT infrastructure is now so widespread, it is now possible to develop affordable broadcast audio networking technology that builds on existing IT technologies. This makes the cost of networking broadcast systems relatively affordable, and also introduces the idea of making them remotely controllable.

THE BENEFITS OF NETWORKING

Interoperability — The Holy Grail

However, as with many areas of human activity, theory and reality are often quite different. Just because data networks are now interconnected across the world and there are ways of adapting IT networking technology to carry audio and video data, it doesn't necessarily follow that a live broadcast stream can simply be fed into an Ethernet router in Salford, UK and emerge unscathed in an edit suite in Shenzhen, China; and nor would you be able to guarantee any of the access control or data security that modern broadcasters demand.

In all of the modern world's major population centers, we take access to electricity for granted, yet the irritating adaptors and transformers we carry with us from continent to continent are a small reminder that even today, the supply is far from standardized worldwide.

International data networks have been with us for far less time, so it should be no surprise to learn that IT infrastructures which were originally designed to handle office-based data transport are far from optimal for routing, mixing, processing and controlling real-time multichannel audio and video, or the equipment that uses them.

That's not to say widespread Ethernet-based networking technology can't be developed to form the basis of broadcast networks — indeed, as we have already covered, this is already a reality. The problem is that there are already many proprietary ways to do this, and there is no one standard for handling audio and video together across a network other than embedding the audio with the video, de-embedding it when processing or mixing is required, and re-embedding it afterwards.

The Holy Grail is a networking standard that allows the use of a single high-capacity network for all of a broadcaster's infrastructure: IT, phones, intercoms, broadcast audio and video (analogue or digital), with equipment monitoring and control protocols, together with some kind of system management, with all the equipment being able to talk to all of the control protocols on the network irrespective of its manufacturer.

This shining goal is known as 'interoperability.' The good news is that cross-manufacturer standards to facilitate it are in development, and we will look at two of the most promising, AVB and Ravenna, later in this primer (see Chapter Three).

In the meantime, such monitoring and control aspects are present only in some of the available proprietary standards (few of which are mutually compatible) and the ability to network video signals and hardware is still missing. That said, the majority of broadcast mixing consoles can route multi-channel audio via their networking protocols, and some of these can also pass control data so that third-party hardware can be controlled from the consoles.

You may well ask what difference it makes whether audio is being routed around a studio via CAT5 network cable, fiber-optic or co-axial links or even analogue tie-lines? It's also fair to consider how much would it cost to add a network infrastructure to traditional connectivity in an existing studio for such little apparent benefit, much less replace one with the other? Surely retaining the existing infrastructure and having patchbays to take care of the routing is more cost-effective?

THE BENEFITS OF NETWORKING

Routing & Cost Benefits

In truth there are a number of benefits associated with a networked approach. Firstly, modern mixing consoles can now take control all of the audio routing in your studio, allowing broadcasters to save themselves the expense of a separate audio router or patchbay, together with all the connections and wiring to it (more on this point towards the end of this chapter).

Furthermore, from a future-proofing perspective there's no question that it's now prudent to connect the constituent parts of a broadcast workflow via a network, certainly on the audio side and especially in the case of new-build projects that are not adapting existing infrastructure.

This is particularly the case for broadcast complexes with large numbers of mixing consoles, control rooms and studios. Using a network and a proprietary audio routing protocol/management system, it's possible to route microphone sources from a wallbox in one studio into the control room of another in seconds, or assign the mixer in one control room the task of mixing the combined output of two or more studios, and then return it to being dedicated to a single room again when the job is complete.

Even with patchbays that allow flexible interconnection between studio and control room, this would be a challenge. What's more, achieving such 'super-studios' via traditional audio connections, whether analogue or digital, requires the running and temporary installation of a lot of extra, expensive cables and looms, and increases the risk of on-air faults.

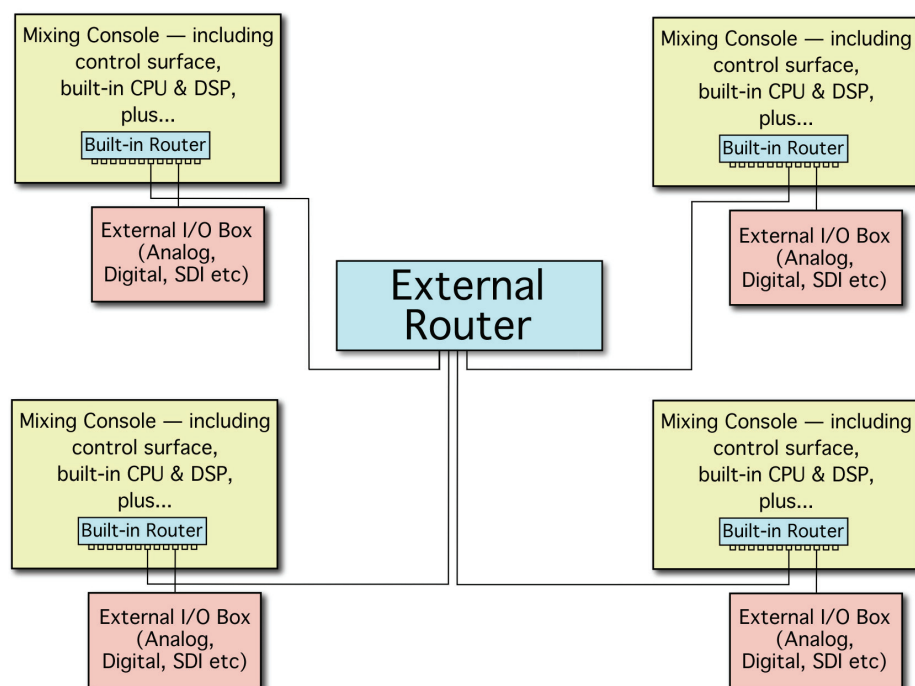
But in a new-build studio designed around a network from the outset, the cost of the network interconnection cabling is minimal and all the hardware is already

connected to the network. The routings simply have to be reassigned by control software. A few clicks of a mouse, and the work is done — or undone.

Such flexible workflows are increasingly the everyday stuff of modern broadcast and the genie is out of the bottle. The next logical step to take once broadcast audio is networked is to connect one audio console's router to another, making it very simple to transfer all of the audio being received at one console, or even just at one I/O box to a console, to a completely different studio for mixing.

Given the bandwidth of today's networks, which typically allow many hundreds of audio channels to be passed down a single connection, there's no need to stop at interconnecting a pair of consoles and their associated I/O. Why not link many consoles together with a stand-alone audio network router, and thus allow several mixers to freely swap audio channels? This is the basis for a star network like the one shown below (Figure 1), with several consoles connected to a stand-alone router.

FIGURE 1 - STAR NETWORK



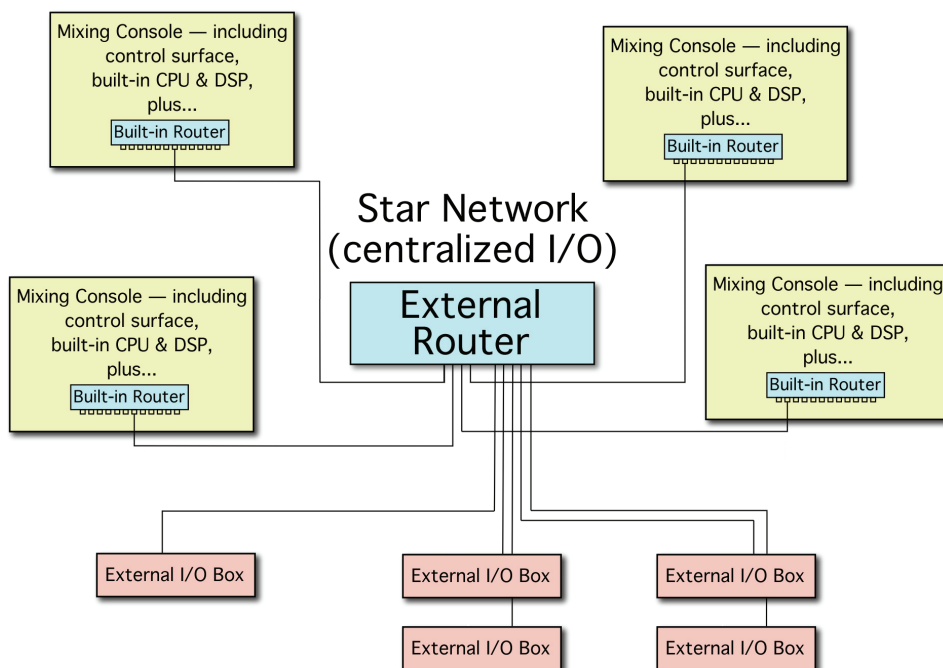
THE BENEFITS OF NETWORKING

In a broadcast complex structured like this, sound stages or studios are (quite literally) no longer tied to a single control room. It's very simple to take multi-channel audio being received from one studio and mix it in another, or to route that audio to another studio to create different mixes for (say) international versioning or commentary. On a rolling news show, the production team in one control room can be mixing the live audio from the news studio and can hand the audio from that studio over to the incoming team in a different control room and go off shift.

Or the team in one control room can switch from mixing the output of one studio or sound stage to working on the output from a different one in seconds. Such things can be done with analogue or digital tie-lines, but a vast amount of expensive wiring is required. This is not a concern with networked audio given that you can route several hundred channels of high-resolution audio down a single inexpensive Ethernet-style network cable.

But even a star network is only the beginning once multiple consoles are networked. In the example above (Figure 2), each console still has its own dedicated I/O interface (or more usually in this day and age, several interfaces, each handling different audio output formats). This is still quite a traditional structure that owes much to the days when the I/O was a fixed part of individual consoles. A network permits something more like the modified star structure shown below. Here an assortment of I/O interfacing boxes in a central location is shared commonly by all of the consoles in a broadcast complex, and is connected to them via the central stand-alone router.

FIGURE 2 - MODIFIED STAR NETWORK

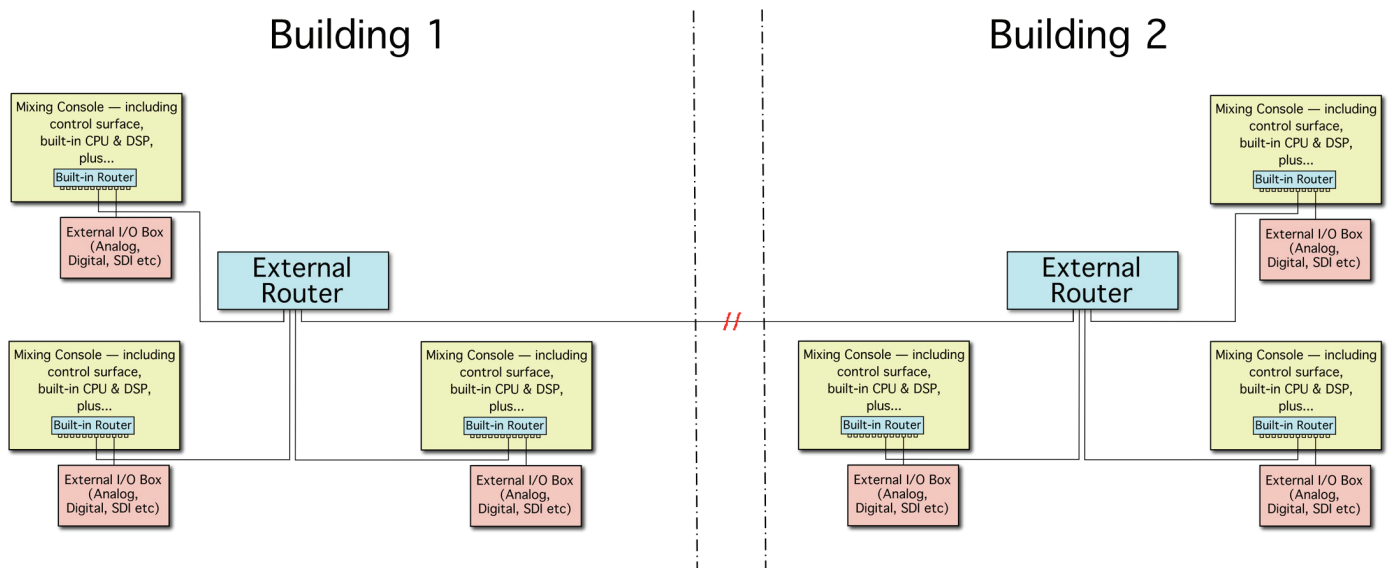


Packing hundreds of channels of audio into a single high-bandwidth networked data connection, where connections can be easily made and reassigned, encourages the construction of complex workflows and network topologies that would be difficult to achieve with standard audio connections.

Consider that it's not unusual for the cost of wiring a national broadcaster's transmission control centre to exceed several million dollars, and it becomes clear that audio networks can offer significant savings.

THE BENEFITS OF NETWORKING

SPLIT ROUTER CORE



Networks & Resilience

For the same reason, it's considerably less costly to design failsafe systems if your broadcast audio is part of a network. In broadcast, the failure of mission-critical connections live on-air is unacceptable and modern broadcast control centers like to factor redundancy into their designs so that every connection has a backup which can easily be pressed into service in the event of a failure.

Doing this with traditional analogue or digital connections requires twice the amount of expensive cabling and a lot of complicated cable splits. With networked audio, the entire output of a master control room or transmission centre, which may include thousands of channels of audio, can be duplicated on a few Ethernet-style IT cables. Many broadcasters choose to mirror their resources across a network in this way, assembling identical equipment in physically or geographically separate buildings which may be networked

together, but may also continue to function independently in the event of one half of the network becoming unusable, as in the diagram below which shows one such 'linked star' network.

In this way broadcasters can be said to achieve 'resilience' in their systems more easily and affordably than those employing traditional designs. Several leading broadcast audio mixing console manufacturers also take advantage of this technology to offer highly resilient network structures with redundant hardware as well as routing.

In such systems, the fundamental hardware components of the audio console can be duplicated, and the duplicates placed in separate physical locations. In the event that one of the locations is rendered inoperable by power failure, fire, flooding or some other unforeseeable natural or manmade catastrophe, the components in the other

location can take over and be switched into operation seamlessly over the network, ensuring continuity of operation.

Calrec consoles offer this way of working as part of a standard setup, as all of the fundamental hardware components that drive its latest generation of consoles (such as the main processor, router, and DSP cards) are supplied in pairs as standard. Placing one set of the pairs in a different location and linking them via the audio network is a simple matter, and constructing similar resilient structures with consoles from other manufacturers, including Lawo and Stagetec, is equally straightforward and increasingly commonplace. Once again, it stands repeating that offering such resilient workflows without networked audio would be incredibly complex and expensive.

THE BENEFITS OF NETWORKING

Networks & Contingency Planning

Networks also make contingency planning easier and more flexible. For example, what do studio owners do if routine maintenance needs to be performed on one studio or control room, and the programme that usually transmits from that studio is due to be broadcast? Or, in a commercial broadcast complex, what happens when a last-minute booking is received from an important client and needs to be accommodated without disrupting the usual assignment of studios and control rooms to other clients?

In a networked broadcast complex, the output of any studio can easily be assigned to a different control room, or conversely a familiar control room can be used to mix the output from a different studio.

There are other challenges to overcome, such as ensuring that settings made on one console can be seamlessly ported over to the replacement studio's console along with all the routings and channel assignments, but manufacturers are working on ways of making this possible over a network too — see Chapter Five for examples of how Calrec is approaching it.

Networks & Control Protocols

One of the benefits of using networking technology for audio routing is that the router in a broadcast audio console can then do the job of an expensive stand-alone studio audio router, allowing broadcasters to dispense with the latter and save money. However, over the years, stand-alone routers have developed the ability not only to pass and route audio and video, but also to exercise limited control, via specially developed protocols, over some of the hardware they were connected to.

For example, video switchers can often exercise basic control over gain settings on an audio console, or control source and destination pairs on an associated audio router. More capable control protocols, such as that from Ross Video or Calrec's development of that protocol (known as the Calrec Serial Protocol or CSCP), permits the control of audio channel levels, and CSCP even allows pan settings, buss assignments and routings to be controlled remotely. Many broadcasters would miss these capabilities if no stand-alone router is included in the system.

Consequently, some console manufacturers (including Calrec) have begun to add support for these hardware control capabilities into their networking protocols. Whilst it can be argued that this is merely adding functions to networked devices that were already available in a pre-networked era, the attraction of incorporating these control abilities into networked devices is too good to ignore. The degree to which control protocols can be integrated into networks varies from manufacturer to manufacturer at present, but the intention is to include control protocols in the interoperable network standards of the future.

We've now covered most of the benefits that networks can offer broadcasters, and it's time to look more closely at the exact state of play in the market today. In the absence of true interoperability, what options are out there for networking broadcast equipment, and how far do they extend?

CHAPTER TWO

SOME TECHNICAL BACKGROUND

SOME TECHNICAL BACKGROUND

As we've discussed, a number of proprietary standards have been developed for transmitting audio over IT networks. Some are better than others, but all have been designed to deal with the fact that the requirements for delivering broadcast audio over a network are considerably more stringent than those for IT-related data.

IT networking protocols for data transfer (such as the ubiquitous Ethernet) are asynchronous, meaning that the order in which data arrives is not held to be so important as long it arrives eventually. However, a broadcast audio feed usually consists of many channels of high-resolution audio, all of which must be kept in sync with respect to each other and which have to be delivered in real time to avoid dropouts. Looking at the technicalities of data transport over a network, we can categorize audio networking protocols in terms of how closely (or not) they resemble IT networking data standards.

Modern electronic networks are often described in terms of a notional model of up to seven layers of increasing complexity that can be used to integrate communications protocols into real-world applications. For audio networking, the most important of these are the first four layers (which are also the most fundamental).

Layer 1 describes the basic electrical standards and voltages used to transmit data over a wired or wireless network, such as an Ethernet network.

Layer 2 describes the most basic unit of data used on the network; in an Ethernet network, this is the 'frame' containing the electronic data. Ethernet Frames include source and destination MAC addresses to

NETWORK PROTOCOL LAYERS

Layer 3 Protocols

Encapsulate audio data in standard IP packets. Layer 3 protocols are not designed to traverse the internet.

Layer 3 audio networking products include:

- Audinate's Dante
- QSC's Q-LAN
- Wheatstone's WheatNet-IP

Layer 2 Protocols

Encapsulate audio data in standard Ethernet frames. Many use standard Ethernet hubs and switches.

Layer 2 audio networking products include:

- AES51
- CobraNet
- Digigram's EtherSound

Layer 1 Protocols

Use Ethernet wiring and signalling components but do not use the Ethernet frame structure.

Layer 1 audio networking products include:

- Aviom's A-Net
- SuperMAC, an AES50 implementation
- Riedel's RockNet
- Calrec's Hydra2

identify the source and destination device for data being transmitted.

Layer 3 adds the IP subnet structure used by all Ethernet networks (and Internet servers) to uniquely identify network devices across the globe, and packages the data being transferred in IP packets. These are all numbered to ensure that all of the data arrives in the right order and can be accounted for.

Layer 4 adds the ability to check the arrival of these packets has occurred in the correct order, without losses or duplication.

In practical terms, all of the audio networking technologies currently on the

market are either Layer 1, 2 or 3 protocols (and all of the Layer 3 protocols contain Layer 4-style data verification capabilities). However, it would be misleading to suggest that Layer 1 protocols are the most basic and Layer 3 the most feature-rich.

Certainly, Layer 3 protocols conform more closely to the defined standards of Gigabit Ethernet (the most common network standard) than others, including Layer 4-style packet checking spliced into Layer 3-style IP packets. These packets sit in turn within an overall Layer 2 Ethernet Frame structure and adhere to the basic Layer 1 electrical definitions of Gigabit Ethernet.

SOME TECHNICAL BACKGROUND

Layer 3 Protocols

Thus the structure of the data in Layer 3 protocols, which include Audinate's Dante, Ravenna from ALC Networkx, Axia's Livewire and QSC's Q-LAN, most closely resemble that passing over a standard Gigabit Ethernet network. As a result, they can multicast data to multiple IP addresses simultaneously, as on an office network, and they can pass data via connected Ethernet bridges and routers. This potentially allows the data to be passed over a wide geographical area and not to remain locked within one Local Area Network (or LAN).

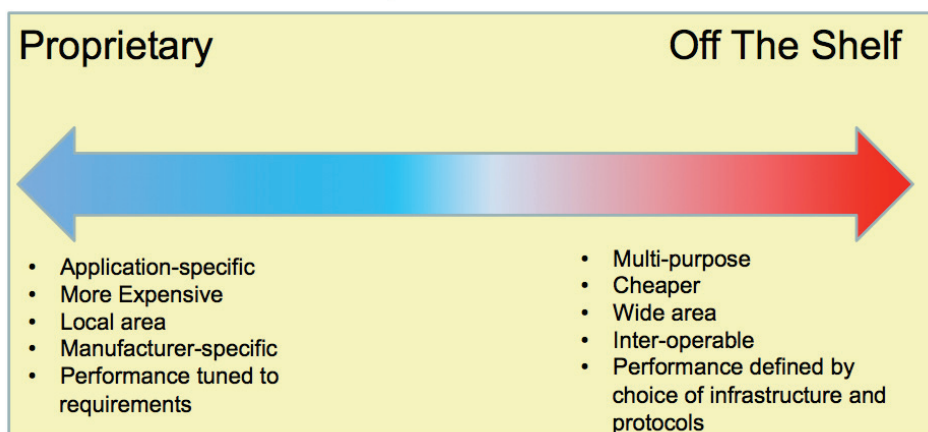
Layer 2 Protocols

The data structure of Layer 2 protocols, which include the IEEE's Audio Video Bridging standard (AVB), Calrec's original Hydra protocol, Peak Audio's Cobranet and Digigram's EtherSound, less closely resemble standard Ethernet data. These protocols dispense with the IP packet structure and thereby lose the ability to be routed to other standard LANs. However, they still use the Ethernet Frame structure and can therefore still be routed within their network via off-the-shelf Ethernet hubs and switches.

Layer 1 Protocols

Layer 1 protocols, which include Riedel's RockNet, Aviom's A-Net, Gibson's MaGIC, and Calrec's Hydra2, have the least in common with IT-style network data. They are geographically limited and also have to use proprietary routing hardware. However, many Layer 1 audio protocols offer similar routing and real-time verification capabilities as Layer 3 protocols — but they do it by means of self-developed, proprietary means. Although compatibility with off-the-shelf networking hardware is lost in a Layer 1 protocol, dispensing with the 'higher-layer' data structures allows the development of very efficient, robust, high-performance,

SPECTRUM



low-latency protocols. When coupled with the hardware required to use them, they are arguably better suited to professional broadcast applications (albeit usually more expensive to implement).

To summarize; 'higher level' protocols offer far greater compatibility with standard networking formats, and allow the use of standard, affordable networking hardware. This can make installation more cost-effective and usable over a wider area, but it can also mean that these protocols are less efficient and higher in latency.

Moreover, because the data in 'higher-layer' networks is usually passed via non-proprietary hardware which is not specifically designed to carry audio data, the reliability of these networks can be lower, and therefore less attractive to broadcasters who need their infrastructure to be as robust as possible.

The 'spectrum' diagram above is a reasonable summary in graphical form, with Layer 1 protocols at one end (more expensive to implement, more application-specific and geographically limited) and Layer 3 protocols at the other (cheaper, with the potential for use over a wider

geographical area, and more interoperable, but with a performance that is entirely dependent on the quality of the hardware and infrastructure being used). Most Audio over IP protocols or Internet Audio streaming standards would fall on the right side of this diagram, being cheap to implement, but may fall below the standard of reliability required by professional broadcasters.

However, compromises that marry the wider compatibility and greater interoperability of Layer 2 and 3 protocols with further standards designed to improve reliability are under development.

CHAPTER THREE

ROUTES TO INTEROPERABILITY

ROUTES TO INTEROPERABILITY

In Chapter One, we touched on the idea of interoperability — the concept of data being shared freely between video and audio equipment. Over the next few years we expect to see a lot of manufacturers in the broadcast market producing equipment which will interface with common transports, although there is still much work to do. A growing number of international broadcast manufacturers are working together to encourage the development of true interoperability between different systems, but as is often the way when an industry tries to establish standards, there are already several different approaches.

AES-X192

An AES standards task group called SC-02-12-H has been formed to develop an interoperability standard for high-performance professional digital audio IP networking. This project has been designated AES-X192 and is partially inspired by an EBU initiative called N/ACIP which published interoperability recommendations for audio over wide-area IP networks.

The scope of this AES initiative is on higher performing networks which allow high-quality, high-capacity and low-latency digital audio transport. There are a number of systems shipping and under development which offer the targeted capabilities — AVB, Dante, LiveWire, Q-LAN and Ravenna — and the aim of this initiative is to identify common approaches and protocols and to suggest and standardise a means for interoperability between the systems.

The two organisations leading this charge in broadcasting are the AVnu Alliance and Ravenna.

The AVnu Alliance & AVB

The AVnu Alliance aims to promote Audio Video Bridging (AVB) as a brand with a view to establishing complete interoperability between manufacturers, and is developing an ecosystem with a clear accreditation scheme (ie. an 'AVB Approved' label). Founded by a handful of companies including Harman, Cisco and Intel, the Alliance has members including Avid, Axon, Beyerdynamic, Calrec, Dolby, Focusrite, LabX, Meyer Sound, Riedel, Sennheiser, Shure and Yamaha.

AVB is a layer 2 protocol that supports various channel and latency options on 100MBit or Gigabit Ethernet. The unique innovation of AVB is that it is designed to make use of a number of extensions to the Ethernet standard (referred to collectively as IEEE 802.1), that are designed to support real-time streaming services. This means that bandwidth can be reserved through network paths to lock down switching resources and ensure streaming-friendly packet queuing and forwarding behavior.

AVB is not an 'audio over IP' protocol — this is a common misconception. In fact, as a Layer 2 protocol, it uses Etherframes rather than IP packets to transport data. As explained in the last chapter, this means AVB networks are geographically limited to their local network, and cannot extend across routers or bridges, unlike Ravenna (of which more in a moment). Furthermore, the benefits of the IEEE 802.1 extensions are only felt if the infrastructure explicitly supports them, which requires specially manufactured AVB switches and hubs to be made available.

The trade-off is the guarantee that what you put in is what you get out, and this reliability and predictability is attractive to broadcasters. Other useful

functions include AVB's DECC (Discovery, enumeration, connection and control) protocol, which presents a view of all AVB devices on the network as soon as a connection is established to it.

If AVB achieves a commercial critical mass, it promises a convenient technology for connecting various devices from different manufacturers across a common network infrastructure. More fundamentally for broadcast audio, if correctly managed this may be shared with other data services without risk to priority audio services.

ROUTES TO INTEROPERABILITY

ALC Networkx & Ravenna

Proposed by ALC Networkx at IBC 2010 as “a technology for real-time transport of audio and other media data in IP-based network environments”, Ravenna is an open technology standard without a proprietary licensing policy. As such, it encourages partners to participate in ongoing development.

The aim of Ravenna's developers is to adapt standard network protocols for use primarily in the professional broadcast market. It makes use of an internet streaming technology which is controlled by Internet Engineering Task Force (IETF) protocols such as RTP, SRP and SIP (these are the protocols used by the likes of Spotify, YouTube and internet radio).

Ravenna is an open standard layer 3 protocol. Although intended for an Ethernet infrastructure, its use of IP packets abstracts it from the underlying network fabric, extending its reach beyond LANs to public networks, and even the internet. In other words, there are no geographical limits to this technology — the use of standard protocols makes it possible to make use of existing IT-style IP infrastructure. This is a clear and obvious benefit. Synchronization is achieved through IEEE 1588-2008 (PTPv2 Precision Time Protocol), another standard protocol which provides the ability to synchronize local clocks to a precision defined by the AES-11 standard.

ALC Networkx has attempted to address the issue of interoperability by encouraging the formation of a Ravenna ecosystem, and they have already signed up a healthy collection of well-respected manufacturers as developers, including AEQ, Digigram, Genelec, Lawo, LSB, Merging Technologies, Neumann, Schoeps, Sonifex, Telos and Linear Acoustics.

These companies are developing the existing protocols, and at NAB 2012 proved the validity of this approach, when ALC Networkx and Axia demonstrated mutually interoperable modes of Ravenna and Livewire (Axia have been selling a similar network technology, Livewire, since 2003).

CHAPTER FOUR

CALREC'S HYDRA2

CALREC'S HYDRA2

Over the last few pages, we've looked at the most successful industry-wide attempts to create a cross-manufacturer broadcast audio networking standard, including the efforts of the AES-X192 working group, the Ravenna-aligned manufacturers and the AVnu Alliance.

However, some proprietary broadcast audio networking standards have already evolved into fully featured protocols, albeit mostly manufacturer-specific ones.

Calrec's Hydra and Hydra2 networks were developed to address specific industry needs and have evolved to meet broadcasters' rapidly changing requirements. At this point in our exploration of network technologies we will refer to these technologies to illustrate some of the challenges which are singularly characteristic of broadcast infrastructures.

Reliability Vs. Compatibility

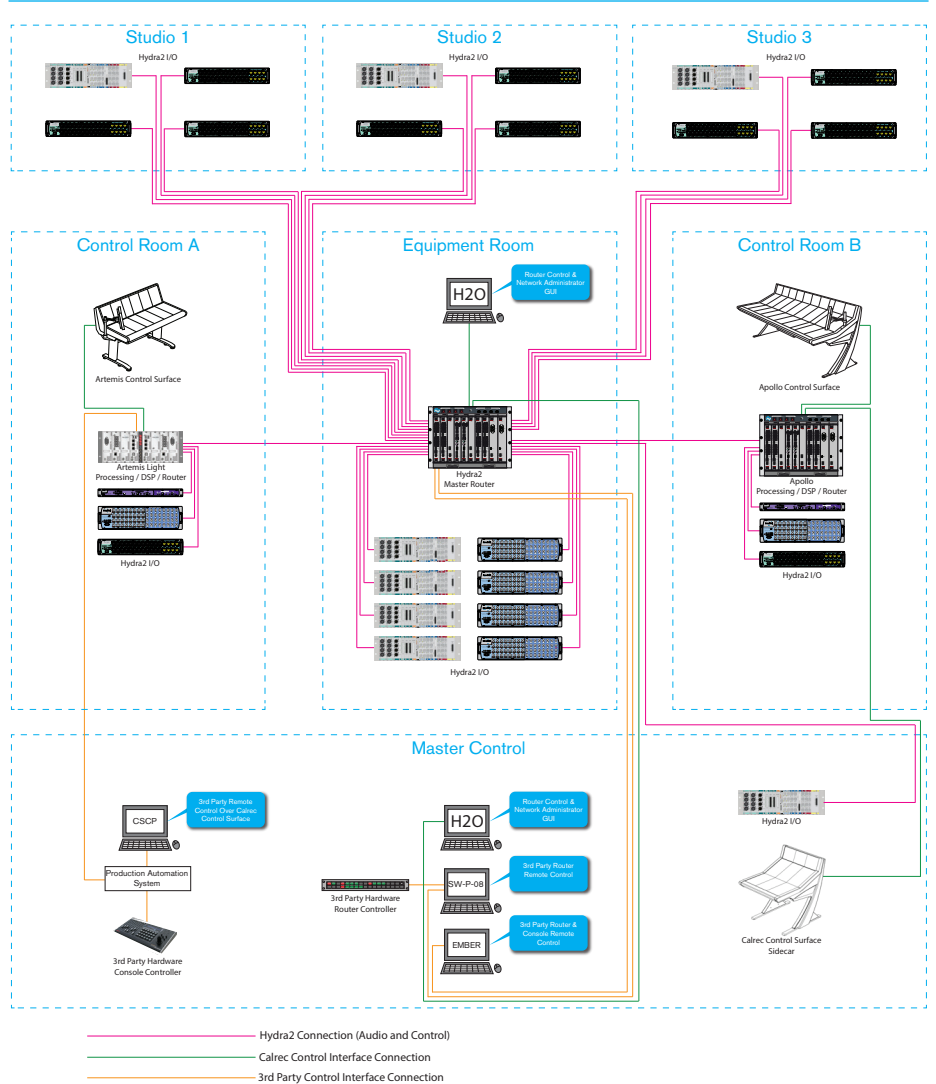
As explained in Chapter Three, it is both a strength and a potential weakness of the Layer 2 and Layer 3 audio networking protocols that they are able to make use of off-the-shelf hardware. Strengths include the fact that Layer 2 protocols can make use of standard, affordable network cabling and hubs; Layer 3 protocols can use off-the-shelf IP-based routers and bridges, like Internet data traffic, and may therefore be multicast over a wide geographical area.

However, the downside is that predictability — and therefore reliability — may be compromised when data is routed via such industry-standard Internet hardware. And put bluntly, most responsible broadcasters are not prepared to run the risk of their content being delivered with the reliability of a YouTube video.

Proprietary protocols like Calrec's Hydra also exhibit this compromise. The most recent manifestation of Calrec's standard, Hydra2, is a wholly proprietary Layer 1 protocol that is unable to use industry standard hardware above standard networking cable. As such, all of its routing and control hardware and software has to be specifically made for use with Hydra2, which makes it relatively expensive. However, the advantage is hugely improved reliability, efficiency and

predictability, and for many broadcasters this is a worthwhile trade-off. Modern professional broadcast audio networks, such as those installed in a multi-studio broadcast centre (below), may be routing tens of thousands of audio channels simultaneously. The ability to do so efficiently and deterministically, without any bandwidth restrictions, whilst keeping latency low and with plenty of scope for swift recovery in the event of component failure, is highly valued.

MULTI-STUDIO BROADCAST NETWORK



Latency & Redundancy

In part the low latency of Hydra2 derives from the data structure chosen for it. As a Layer 1 protocol, its data is not packaged in Ethernet Frames. Instead, audio is sent in 512-sample chunks with no headers, but with some capacity for control data. This method of data-packing is highly efficient, and because the start of the data is coincident with the leading edge of the audio sample clock, it contains implicit synchronisation information. The predictable arrival times of the data almost completely removes the need for buffering which reduces network latency — Calrec reckons with an 11-sample latency for an AES3 signal passing from input to output across a Hydra2 network.

The Layer 1 proprietary nature of Hydra means that proprietary routers had to be developed — but again, this has a performance advantage. Each Router Core (as Calrec terms them), has a matrix that works like a synchronous TDM router, and can route one input to all 8192 outputs simultaneously if required.

Redundancy is another important issue for most broadcasters. On Hydra2 networks, it's achieved by duplicating network hardware. Each element in the router core has a primary and a secondary circuit card, and all of the physical network connections are duplicated too, creating, in effect, two complete networks. Various mechanisms monitor the health of each component and, in the event of a failure, deploy the redundant partner.

This architectural flexibility allows the kind of highly secure operation based around a physically split Router Core located in separate facilities, as described towards the end of Chapter One.

When multiple Router Cores are connected, as in a multi-studio broadcast

centre with many control surfaces, one Router Core is always designated the master controller. This avoids unpredictable behaviour when more than one console attempts to use the same network resources. Although attached devices are aware of the entire network, if they wish to make routes through the network fabric, they must apply to the master controller. This will respond by making the routes - or in the case of conflicts, arbitrate in a predictable manner.

H2O & Network Management

Having central control of the network has further benefits. It provides a single point of contact if, for example, the broadcaster wishes the network to be integrated with an external broadcast control system.

Master controllers which can be accessed through a browser can therefore be managed or edited remotely, such as in Calrec's browser-based Hydra2 management software, Hydra2 Organiser (H2O).

With software like this, the interconnections and routings on the network may be created, assigned and redefined on a channel-by-channel basis. Hydra2 uses graphical constructs called Hydra Patchbays. These are simply virtual routing matrices defined in H2O that determine how signals flow between the hardware connected to the network. Hydra Patchbays are a powerful tool for control room and studio resource management, allowing network administrators to put control rooms 'on-air' and to manage the sources available to them.

H2O provides a variety of other network management services, including the management of access rights, which can confer or deny right of access of consoles to selected network resources, and the protection of in-use ports from inadvertent overpatching by other network users. Both of these are front-line requirements for broadcasters and commercial facilities that may rent out different studios to different (possibly competing) broadcasters.

Imagine one studio complex with three studios, which can all be assigned to work together, or divided into three independent studios rented to different clients. In cases like this, the master controller must apportion the network so that the client renting one studio cannot see the other studios, consoles and routers on the network, much less access them. And at the conclusion of the rental period, the network has to be reconfigurable so that the three studios, and their Router Cores and control surfaces, can access one another again. All of this is possible using the network management features of Hydra2 via H2O, and was implemented in response to requests from broadcast studio complexes who wanted these facilities.

Broadcasters also contributed to the development of another contingency-related network feature in Hydra2 — alias files. The requirement came about in networked multi-studio broadcast complexes where clients suddenly had to move from one studio to another due to unforeseen maintenance requirements or double bookings. Taking the desk settings across to a new studio in the form of a desk memory is easy — the settings files are small enough to be portable on a memory stick — but when the memory is loaded in the new studio, clients would find that all the input and output assignments for the console were wrong, because they had been assigned for the original studio, with an often radically different set of wallboxes, outboard and interconnections.

To make desk memories more portable on a Hydra2 network of connected consoles, Calrec worked with a major European broadcaster to develop the concept of console assignment alias files, whereby I/O assignments are made not to specific network ports but to aliases of ports. Another file, specific to that console, then maps the aliases to the right ports for that desk and studio. This way, if a desk memory is brought in from another studio, or a show is moved from one studio to another, the local console alias file will map the inputs and outputs to correct connections in the relevant studio, irrespective of how the memory has been set up.

Compatibility With Other Standards

Ways to interface Hydra2 networks with other network protocols are planned. As mentioned briefly in Chapter One, Hydra2 already passes several hardware control protocols (of which more in the next chapter) and as we have seen, there is a widespread desire amongst broadcasters to make their equipment interoperable. Responsible manufacturers will all play a role in this. As a manufacturer, Calrec doesn't view the work of the AES-X192 group, ALC Networkx or the AVNU Alliance as replacements for Hydra2 — it is hard to see them having the latency, determinism and capacity of such a purpose-designed protocol as Hydra2. However, they can be regarded as excellent potential companion technologies, and adding a Ravenna or AVB interface to a Hydra2 network could provide a flexible and low cost way of connecting to intercom systems, plant routers, hard disk recorders, speakers and even microphones. Calrec supports these efforts, has recently become a member of the AVnu Alliance, the stakeholder group for AVB development.

CHAPTER FIVE

CONTROL PROTOCOLS & NETWORKING

CONTROL PROTOCOLS & NETWORKING

In the course of this Primer so far, we've looked at the current efforts to create common standards for interfacing broadcast equipment over a data network, and to provide some measure of interoperability between audio and video equipment.

But recent protocols and standards such as Ravenna, AVB, and X192 are only the latest in a series of efforts that have been made through the decades to get video and audio hardware talking to one another. Long before the advent of data networks and broadcast facilities connected with CAT5 cable, video switchers were able to exercise basic control over gain settings on some audio consoles, or control source and destination patchings on connected audio routers via simple serial control protocols. Most of these predate the computer networking era, but have survived into the modern age because they are simple to implement and provide basic and convenient interfacing between broadcast audio and video hardware...even if their capabilities are limited compared to the kind of interoperable standards now under development.

Calrec's current product range, for example, supports the basic Ross Video hardware control protocol (more on this in a moment), as well as the even more venerable but widely supported SW-P-02 and SW-P-08 serial protocols (also known as General Switcher/General Remote), originally designed by ProBel to allow their own products to control switching and routing assignments on their routers, and now used industry-wide.

This approach makes sound commercial sense. By incorporating support for existing control protocols into their audio networking standards, manufacturers that advocate the use of networks ensure that broadcasters can still use the convenient

control protocols they're used to, even after a move to network-based workflows. For example, for broadcasters used to controlling stand-alone audio routers from their video switchers via protocols like SW-P-08, support in current and future networking standards allows broadcasters to continue directing audio from video hardware in a way familiar to them...even if the stand-alone audio router has been completely dispensed with and audio is now being routed via the network.

Broadcast audio hardware manufacturers currently provide compatibility with audio/video control protocols in three main ways:

- 1) By means of existing widely used but rather limited generic protocols (such as SW-P-08/02);
- 2) By adapting or extending existing control protocols, thus building on existing standards while widening the capabilities on offer (for example, Calrec's extension of the Ross Audio Protocol (CSCP));
- 3) By promoting and becoming actively involved in the development of newer control protocols, such as EMBER, whose abilities are still being extended and defined.

CONTROL PROTOCOLS & NETWORKING

SW-P-08/02 & RAP

SW-P-08/02 is a great example of a limited control protocol whose use has become widespread because of its simplicity, similar to how MIDI has been survived in use in the world of musical instrument interfacing despite being decades old and thoroughly superannuated in terms of its feature-set.

In truth, the lack of a need to pay to license SW-P-08/02 from ProBel and the low cost of implementing it in hardware probably play also a significant part in its widespread commercial adoption.

However, SW-P-08/02 is distinctly limited, and cannot achieve much beyond allowing a video switcher to direct one input on an audio console's router to another output; even adjusting the channel gains on a connected mixer is beyond its capabilities. Whilst many manufacturers support it, the demands of today's modern broadcast control centres require more than the ProBel protocol has to offer. Ross Video's audio control protocol, known as RAP, is more fully featured, allowing control of fader levels and a few basic audio parameters.

CSCP & Automation

Calrec supports both of these control protocols in its Hydra2 audio networking standard, which passes RAP and SW-P-02/08 data down its connections. But the company has also extended the RAP protocol to create the Calrec Serial Protocol, or CSCP, which allows control of channel levels, audio channel pan settings, and buss assignments and routings, which is transmitted over Hydra2 networks. Through collaborations with other video hardware manufacturers such as Grass Valley and Sony, support for CSCP has gradually been extended, allowing broadcast control systems such as Grass Valley's Ignite, Ross's Overdrive, or Sony's ELC to integrate audio consoles into their workflows.

This is attractive to many broadcasters as it offers the kind of automatable integration of video and audio systems that they have been calling for, permitting greater levels of production automation. However, the list of audio parameters that you can control from video equipment via the Ross or Calrec protocols still falls short of complete interoperability, and remains far from exhaustive.

CONTROL PROTOCOLS & NETWORKING

EMBER

Perhaps the most promising development in the world of control protocols is the kind provided by LSB's EMBER: a fully-featured, open source protocol whose capabilities closely reflect the requirements of modern broadcasters.

The still-developing EMBER communication protocol was created by the German company LSB, manufacturer of the excellent Virtual Studio Manager broadcast control software, and provides one of the foremost examples of how video and audio integration could proceed in the near future. A proprietary protocol, but one open for use by third-party manufacturers, it's being supported by leading audio hardware manufacturers such as Lawo, Stagetec, Studer and Calrec, and recognises in its design the fact that broadcasters ideally want to be able to hook video equipment to audio equipment via a network, turn everything on, and simply begin controlling one from the other.

Using a standard client/server model an EMBER server, will first describe all of the equipment on the network it's connected to, and all of the parameters in that equipment which it can address, and then facilitate communication between the EMBER client in broadcast control software such as VSM and the EMBER server on the broadcast network. When a piece of EMBER-compatible audio equipment is connected to a broadcast network running an EMBER server, the controllable parameters of that device are automatically described and made addressable over the network via the EMBER protocol.

Similarly, if you hook some EMBER-enabled broadcast control software to a network containing various types of EMBER-compatible audio and video

equipment, the software can begin controlling the equipment over the network without any prior knowledge of the network and its addressable components and parameters.

The list of controllable parameters that are addressable via EMBER is growing as customers request more features and participating manufacturers widen their support accordingly. Unsurprisingly, because the protocol is still being rapidly developed, the limits of EMBER are being tested with modern demands that go way beyond the simple control of gain and pan settings offered by older protocols. To take a recent example, the ability to control the addition of Dolby metadata into an embedded audio stream has been added to Calrec's implementation of EMBER.

The Future of Broadcast Networks?

Like all aspects of broadcast audio networking described in this Primer, development of EMBER is still proceeding apace — but even with its capabilities today, it gives a tantalising glimpse of what might be possible in the near future, when interoperable standards are more fully developed. If this Primer has a message, it's to emphasise that networking technology is on the brink of delivering unparalleled audio and video integration to broadcasters in many areas, even if some aspects still remain just out of reach. Even half a decade ago, such claims would have been dismissed as misty-eyed exaggeration — but today, as we have seen in this book, the scale of developments on all fronts leading to an interoperable future can no longer be denied.

Calrec Audio Ltd

Nutclough Mill
Hebden Bridge
West Yorkshire
England UK
HX7 8EZ

Tel +44 (0)1422 842159
Fax +44 (0)1422 845244
Email Enquiries@calrec.com

calrec.com

(926-185 Iss.1)