



INTRODUCTION TO LIVEWIRE+™



IP-Audio System Design Reference & Primer

Version 3.0, May 2015

About This Manual

This document is your introduction to Livewire+. We explain the ideas that motivated it and how you can use and benefit from it, as well as nitty-gritty details about wiring, connectors, switches, and the like. Since Livewire+ is built on standard IP networks, we also help you to understand general network engineering so that you have the full background for Livewire+'s fundamentals. After reading, you will know the essentials of Livewire+ along with many of the details of the gear that vendors and the network guys that are often hanging around in radio stations will be using.

This document covers topics common to all Livewire+ equipment and standards-based AoIP networking. It is only a part of your full documentation package. You will also have manuals for each specific piece of equipment that are used to build your system. From this document, for example, you will not learn how to install or operate an Element console, but you will understand the nature of the network it plugs into.

(Permit us a small, shameless plug: Focal Press has available an excellent volume entitled *Audio Over IP – Building Pro AoIP Systems with Livewire™*. This was co-authored by Telos founder Steve Church and respected technologist Skip Pizzi. For a truly in-depth study, it can't be beat.)

It was early in 2003 when Livewire™ was first introduced and came to market being new and fresh. Today, Broadcasters and Engineers have grasped this standards based technology to the point where you now find Livewire+ (the new, extensible protocol which incorporates and complies with the AES67 networking standard) incorporated into many other vendors' equipment. Not to mention making Axia the fastest growing AoIP console company in the world!

As always, we welcome your suggestions for improvement. Please feel free to contact us with your comments:

The Telos Alliance
1241 Superior Ave. E.,
Cleveland Ohio 44114 USA
Phone: +1.216.241.7225
Web: www.TelosAlliance.com
Email: inquiry@TelosAlliance.com

Table of Contents

	<i>About This Manual</i>	II
	<i>A Note From The Founder of Telos</i>	III
<hr/>		
1	<i>Livewire+ For Beginners</i>	1
	Why Ethernet?	1
	Compared to AES/EBU	2
	Audio Routing	2
	The Livewire+ Advertising System	2
	Control.	2
	Livewire+ and PCs	3
	Support for Surround	3
	Audio Quality.	4
	Fidelity	4
	Delay	4
	The Pac-Man of Protocols: Internet Standards	6
	Converged Networks	6
<hr/>		
2	<i>What Can You Do With Livewire+?</i>	7
	Make A Snake.	7
	A High-Performance Sound Card Replacement	7
	Build An Audio Router	8
	Build A State-Of-The-Art Broadcast Studio.	8
	Make a Flexible Two-Way Multi-Channel STL.	10
	Create A Facility-Wide Audio Network	11
	Create An Integrated National/Local Radio Network	11
<hr/>		
3	<i>The Components Of Livewire+</i>	12
	xNodes.	12
	Axia IP-Audio Driver for Windows®	14
	iPlay (PC Router Selector)	16
	Mixing Engines	17
	StudioEngine	17
	PowerStation	17
	QOR Mixing Engines	18

Broadcast Consoles	19
Fusion Broadcast Console	19
Element Broadcast Console	20
iQ and Radius Broadcast Consoles	21
DESO and RAQ Broadcast Consoles	21
SoftSurface Virtual Consoles	22
PathfinderPC Routing Control Software	22
Provisions for Redundancy and Back-up	23
Timed Events	23
Livewire+ Audio Router Control Protocol	23
iProbe Network Management Console	24
iProFiler Automated Program Archiving	24
Networked Audio Processing	24
Networked Codecs	25
Networked Telephone Systems	26
Networked Web-Streaming Encoders	27
Networked Time Management	29
Networked Intercom	29
<hr/>	
4 <i>Nuts & Bolts: Making Livewire+ Play</i>	30
Livewire+'s Channel and Name system	30
Channels	30
Text Name	31
GPIO	31
Sources vs. Destinations	31
Examples	32
xNode Configuration & Web Access	38
Front Panel xNode Configuration	38
Plugs & Cables	39
CAT-5 for Audio?	40

5	<i>Designing & Building Your Livewire+ Ethernet System</i>	47
	Cabling	47
	Twisted-pair Cable Categories	47
	Special Care for Ethernet Audio	48
	To Shield or Not to Shield	49
	Unbalanced Connections	49
	More than Four Pairs in a Cable	49
	Patch Panels	49
	Wall Jacks	50
	CAT-6 Jacks	50
	Architecture Options	50
	Simple One-Switch Network	50
	Daisy Chained Multiple-Switch Network	53
	Hierarchical Multiple-Switch Network	54
	Options for Redundancy	54
	Fiber	55
	Radio Links	56
	Designing For Security	56
<hr/>		
6	<i>The Ethernet Switch</i>	57
	Livewire+ Ethernet Switch Requirements	57
	Some Switches We Like	58
	Cisco Switches	58
	Axia xSwitch	58
	Switch Configuration	59
<hr/>		
7	<i>Testing, 1-2-3...</i>	60
	General Ethernet Troubleshooting	60
	Prevention	60
	The Basics	60
	Cable Testers	62
	Sniffers	62
	Diagnosing Problems using Livewire+ Components	63
	xNode Indicator Displays	63

8	<i>Network Engineering For Audio Engineers</i>	64
	Ethernet / IP Networks.	64
	Layering Model	64
	Making Packets	66
	IP and Ethernet Addresses	67
	Subnet mask	67
	Ethernet Switching	69
	Multicast	69
	Prioritization	71
	The Role of TCP	72
	Virtual LANs (VLANs).	73
	Ethernet Switching versus Routing	75
	Livewire+ Networks.	75
	Quality of Service (QoS)	75
	Source Advertising.	76
	Synchronization	76
	Network Standards and Resources.	78
	A Note about Protocol Design	79

9	<i>Frequently Asked Questions</i>	80
	General Questions.	80
	Livewire+, Standards and Other Vendors	85
	PCs and Livewire+	88
	Building Livewire+ Facilities	88
	Ethernet Media	91
	The Internet and Livewire+	92
	Networked Consoles.	92
	Analog Audio & AES On RJ45 And CAT-5.	93

10	<i>Resources</i>	95
	Livewire+/Broadcast95
	Ethernet95
	General Networking and Interest95
	Cabling Information & Standards96
	Cable and Contractor Supplies96
	Cable Testers96
	Ethernet Switch Vendors96
	Network "Sniffers"96
	Ethernet Radio Equipment97
<hr/>		
A	<i>Appendix A: Livewire+ Tech Details</i>	98
	Livewire+ Packet Structures98
	Standard Streams98
	Livestream99
	Surround Streams99
	Network Link Capacity	100
	Multicast Address Translation	100

A Note From The Founder Of Telos

It's been a tradition since Telos' very first product, the Telos 10 digital phone system, that I share a few words with you at the beginning of each manual. So here goes. In radio broadcast studios we're still picking up the pieces that have fallen out from the digital audio revolution. We're not using cart machines anymore because PCs are so clearly a better way to store and play audio. We're replacing our analog mixing consoles with digital ones and routing audio digitally. But we're still using decades-old analog or primitive digital methods to connect our gear. Livewire+ has been developed by Telos to provide a modern PC and computer network-oriented way to connect and distribute professional audio around a broadcast studio facility.

Your question may be, "Why Telos? Don't you guys make phone stuff?" Yes, we certainly do. But we've always been attracted to new and better ways to make things happen in radio facilities. And we've always looked for opportunities to make networks of all kinds work for broadcasters.

When DSP was first possible, we used it to fix the ages-old phone hybrid problem. It was the first use of DSP in radio broadcasting. When ISDN and MP3 first happened, we saw the possibility to make a truly useful codec. We were the first to license and use MP3 and the first to incorporate ISDN into a codec. We were active in the early days of internet audio, and the first to use MP3 on the internet.

Inventing and adapting new technologies for broadcast is what we've always been about. And we've always been marrying audio with networks. It's been our passion right from the start. In our genes, if you will. As a pioneer in broadcast digital audio and DSP, we've grown an R&D team with a lot of creative guys who are open-eyed to new ideas. So it's actually quite natural that we would be playing marriage broker to computer networks and studio audio. What you get from this is nearly as hot as a couple on their wedding night: On one RJ-45, two-way multiple audio channels, sophisticated control and data capability, and built-in computer compatibility.

You can use Livewire+ as a simple soundcard replacement – an audio interface connecting to a PC with an RJ-45 cable. But add an Ethernet switch and more interfaces to build a system with as many inputs and outputs as you want. Audio may be routed directly from interface to interface or to other PCs, so you now have an audio routing system that does everything a traditional "mainframe" audio router does – but at a lot lower cost and with a lot more capability. Add real-time mixing/processing engines and control surfaces and you have a modern studio facility with many advantages over the old ways of doing things.

Ok, maybe this is not as thrilling as a wedding night – perhaps kissing your first lover is a better analogy. (By the way, and way off-topic, did you know that the person you were kissing was 72.8% water?)

While were on the subject of history... you've probably been soldering XLRs for a long time, so you feel a bit, shall we say, "attached" to them. We understand. But no problem – you'll be needing them for microphones for a long while, so your withdrawal symptoms won't be serious. But your facility already has plenty of Ethernet and plenty of computers, so you probably already know your way around an RJ-45 as well. It's really not that strange to imagine live audio flowing over computer networks, and there's little question that you are going to be seeing a lot of it in the coming years.

The 20th century was remarkable for its tremendous innovation in machines of all kinds: power generators, heating and air conditioning, cars, airplanes, factory automation, radio, TV, computers. At the dawn of the 21st, it's clear that the ongoing digitization and networking of text, audio, and images will be a main technology story for decades to come, and an exciting ride for those of us fortunate to be in the thick of it.

Speaking of years, it has been a lot of them since I wrote the Zephyr manual intro, and even more since the Telos 10. Amazing thing is, with all the change around us, we're still here and Telos is still growing in new ways. As, no doubt, are you and your stations.



Steve Church

1


Livewire+ For Beginners

Livewire+ offers a revolutionary change in how studios can be built. But at the same time, it's a natural continuation of general trends and what you already know. This section explains the basics and puts audio over Ethernet into context.

Within the next few years, it is certain that the transition to digital now happening in our studios will be complete, with all audio storage, mixing, processing and routing being digital. This transition is now in full swing. But the sheer number of competing methods for integrating computers, legacy hardware and state-of-the-art audio devices is (and long has been) bewildering. What we need is a connection method that gets the interconnection job done easily, effectively, flexibly, and cheaply. So why not look to the computer and telephone worlds to find the technology? We can then take advantage of the huge manufacturing scale in those industries and can piggyback on the billions of dollars (and Euros, Yen, Yuan...) of R&D going on in those industries.

Why Ethernet?

Ethernet makes overwhelming sense. Today's computers are near universally linked via Ethernet – and telephony is moving that way as well, with VoIP rapidly and continually gaining market share. Even remote controlled stage lighting is transitioning from the XLR-based DMX protocol to Ethernet. Ethernet cables, plugs, cards, and chips are produced in the hundreds of millions so we get tremendous economy of scale. We get patch bays and cords, testers, and all kinds of “structured wiring” components ready-made. Plugs are easy to install and jacks are efficiently small.



But much more important is that Ethernet allows us to combine many channels of digital audio with whatever data transmission we might need on a single cable. This data could be as simple as a start command for an audio player or could be anything that computers and Ethernet do, such as file transfer, e-mail, web communication, etc.

Further, we are in the line of future development. Since its invention over 30 years ago, Ethernet has been constantly evolving. It started as a 2Mbps shared bus over coaxial cable and has grown to today's modern (and very common) 1 Gigabit star and switched system. 10 Gigabit is already widely available on many Ethernet devices and is following the usual curve to low cost as volumes increase. While copper is the most common Ethernet con-

nection, fiber is popular as well and media converters allow the two to be interconnected. Ethernet switches cost \$6000 for 8 ports a half-decade ago; now high-end 24-port switches cost \$500. And they include many advanced features that were unheard of only a few years back.

There are radio links in many varieties, from WiFi for short-range to sophisticated long-range systems. There are satellite links. And LASER links. Ethernet opens the door to a world of options.

Ethernet has proven to be the PC of networking: Initially released with only basic capability – low speed and busbed – it has been expanded to today's fast, flexible, switched architectures.

The combination of huge R&D expenditures, open standards, massive economies of scale, technological evolution, and flexible multi-service packet design is hard to beat. Not to mention the surprisingly appropriate name: Ethernet was named by its inventor, Robert Metcalf. He had been involved in a radio data network in Hawaii called ALOHA. The first Ethernet was a bussed coax that carried data packets similar to the way ALOHA had sent them over the “ether.” As to the origin of ether... for many years after James Clerk Maxwell’s discovery that a wave equation could describe electromagnetic radiation, the aluminiferous ether was thought to be an omnipresent substance capable of carrying electromagnetic waves. In 1887 scientists Albert Michelson and Edward Morley disproved its existence. The ingenious experiment that did so was performed at Case Western Reserve University, just down the street from the Telos Alliance main office in Cleveland.

Compared to AES/EBU

For digital audio transport, AES3 is the main alternative to an Ethernet based system. Invented in the days of 300-baud modems, it was the first practical answer to connecting digital audio signals. But it’s now 30 years old and is showing its age. Compared to Livewire+’s computer-friendly, two-way, multi-channel + high-speed data capability, AES3 looks pretty feeble with its 2-channel and one-way only limitation. Not to mention 50-year old soldered XLR connectors and lack of significant data capacity. AES3 is a low-volume backwater, with no computer or telephone industry R&D driving costs down and technology forward. Your 300-baud modem has been long retired; it’s well time to progress to the modern world for studio audio connections as well.

That said, AES and Livewire+ may comfortably co-exist in your facility. You can use Axia interface nodes to connect from one to the other. If you are using a house sync system for AES, Livewire+ may be synced to that system also.

Audio Routing

Low-cost mass-market Ethernet switches offer us something very interesting: Since their function is to direct packets from port-to-port, we can use them to move our audio signals from whatever source to whatever destinations we want. This means we get a simple, flexible, facility-wide audio routing system, almost for free. Say goodbye to racks of distribution amps or expensive proprietary mainframe audio routers.

An audio source entered into the system from any point becomes available for any number of receiving destinations.

The Livewire+ Advertising System

Livewire+ has an audio advertising system. Every source has a text name and numeric ID. These are transmitted from source devices to the network. Receivers can build lists of all available sources from which users can select.

Configuration is made simple with hardware nodes; you enter the names, numbers, and other configuration information through a web browser via an attached PC. With PC nodes, it is even easier, requiring only the opening of a configuration window to make all necessary source parameter changes.

Control

Most audio these days needs associated control. A delivery system needs a start input at minimum, but could well benefit from a richer control dialogue such as text identifying what is playing that can be sent to the studio mixer and to the HD Radio and RDS encoders. Satellite receivers have control outputs. Telephone systems need dialing, line status, hold, transfer, etc. Even a simple CD player needs ready indication out and start in. Even the simplest source, a microphone, needs to convey on/off status for the air lights. Most conventional controls have been done with primitive GPIO parallel “contact closures.”

As a first step, Ethernet can transport GPIO data, reducing and simplifying cabling, and Livewire+ offers this basic capability to replicate traditional start/stop control. But it continues from there; Livewire+ also supports sophisticated remote operation of studio equipment over Ethernet which allows the network to transport much more advanced information than just simple start commands.

For instance, we can send the song title from a delivery system to a display on a mixing console's fader channel. Control of telephone systems and codecs can follow fader assignment and be accessible from any location. With a high-bandwidth network linking everything and a flexible communication protocol, the door is open to many interesting possibilities. Why couldn't the satellite receiver identify its content with "metadata" tags? Then an automatic system could store a program along with the information about it for later play. An on-air audio processor might respond to program type information to adjust its parameters. Microphones switched-on could activate a logger. And these have all been accomplished with Livewire+ products from the Telos Alliance and our many partners; there are still many possibilities yet to be explored.

Livewire+ and PCs

One of the advantages of a Livewire+ system is that PC-based audio may be directly connected to the network without soundcards. This means radio station delivery systems can use the Ethernet connection they already have to send and receive audio. Soundcard problems such as noise and multiple conversions are avoided – the audio remains in digital form from the PC's files to the network with no alteration or degradation. Received audio may have originated from another PC or from a hardware audio node. Audio sent from a PC may be received by other PCs or hardware nodes.

With nearly all audio in today's radio stations being either played from computers or recorded into computers, isn't this a tremendous advantage? Not only do you save the soundcard, but also the port that it needs at the other end to connect to your console or router. And you can pass control and other information over the same connection.

Support for Surround

Surround sound is a *de facto* standard in the home entertainment industry today. Blu-Ray, Broadcast TV, Streaming Services, and even video games all contain surround audio on a large percentage of content.



The introduction of surround into radio broadcast has also been made possible by advances in multi-channel codec technology making surround a possibility over European DAB channels and the USA HD Radio system. Continued migration to digital streaming services originating from computers makes the surround Internet broadcasts a possibility as well.

A large number of home-entertainment systems with surround sound are in use on a daily basis, and many high-end cars also offer surround audio systems. If you think about it, the car is a great environment for enjoying surround sound since speakers are installed all around the vehicle to start with.

A surround program typically contains 5.1-channels: 2 front, 2 surround, 1 center, and 1 subwoofer. From a strict channel count, this is really 6 channels. The ".1" channel is named that since it only contains 1/10th of the bandwidth that the other, full-range channels do. It might also be required to keep a stereo version of the program (called a downmix if it is derived from the original surround mix) with the 5.1-channel program. Sometimes the surround program might be 7.1-channels, adding 2 more surround channels. So there could be 8 total audio channels required. In TV broadcasting, this is quite common. Using a traditional approach, that's a lot of plugs, cables, router cards, and rack space!

On the other hand, Ethernet has plenty of bandwidth to carry the multiple channels surround broadcasting requires. All eight channels plus associated control can easily be conveyed on one convenient Ethernet cable.

Expanding the audio paths and equipment (consoles, routers, etc.) in a hard-wired facility from a stereo infrastructure to eight channels (or more!) would be either difficult or very expensive. Not so with Ethernet and Livewire+. In fact, there is no additional cost for the core Ethernet switch because the one you need for stereo would also be fine for surround, and audio from PCs can be multichannel.

We have designed Livewire+ with the future well in mind. As of this writing, Livewire+ already provides the infrastructure for over 5,500 modern radio studios. By this simple fact, these facilities are also outfitted with surround capability – ready to implement simply and with low cost.

Audio Quality

We're always asked, "Is Livewire+ like audio on the Internet?" Yes and no. While Livewire+ uses internet transport standards, it is intended to operate over switched Local Area Networks (LANs). Without the limitations of the public internet and with 100% control over all parts of the system, we are able to achieve full studio quality.

So now the question is, "Is Livewire+, audio on the Internet and Audio over IP reliable?" The answer is that Axia uses the same technology that underlies VoIP telephony.

Did you know that more than 70% of the Fortune 100 companies now use VoIP? Or that VoIP PBX systems now outsell the old kind by a wide margin? With these systems, telephones plug into a standard Ethernet/IP network. Contrast this with traditional PBX phone gear — proprietary devices which required you to purchase phone sets and parts exclusively from the company that built the mainframe. You were locked into a single vendor, because the technology that ran the mainframe was owned by the company that made the gear (kind of like the TDM router companies).

IP is growing as a universal transport for almost any kind of signal. You see it now in television studios, business teleconferencing, government communications, banking, etc. And it's hardly unproven, even for applications specific to broadcast studio infrastructure. There are plenty of people successfully using it – now.

Fidelity

Internet streams are usually compressed for transmission over public links with limited, variable bandwidth and low reliability. Livewire+ audio is not compressed – we use studio-grade 48kHz/24-bit PCM encoding. Our audio interface nodes have more than 100dB dynamic range, < 0.005% THD, and headroom to +24dBu. LANs offer a safe, controlled environment where there is no risk of audio drop-outs from network problems and plenty of bandwidth for many channels of high-quality audio without compression.

Indeed, we often hear from Livewire+ users that they notice an improvement in fidelity when they transition from other systems. This is probably due to the direct connection of PCs, the 64-bit accumulator in the Axia mixing engine, and to the careful design of the audio stages of our audio interface nodes, rather than the network itself. But, in any event, the network takes nothing away from audio performance.

Delay

In packet-based systems, delay is an important issue and certainly has an effect on your talent's perception of "quality." Packetizing audio for network transmission necessarily causes delay, and careful design of the system is required to reduce this to acceptable levels. Internet audio delay is often multiple seconds because the receiving PCs need long buffers to ride out network problems and the delays inherent in multiple-hop router paths. However, with fast Ethernet switching on a local network, it is possible to achieve very low delay. To do this, we must have a synchronization system throughout the network. This also avoids sample or packet slips that cause audio dropouts. Internet streaming does not use this technique, so even if it were to have guaranteed reliable bandwidth, you still couldn't achieve the very low delay we need for professional studio applications.

For Livewire+, we generate a system-wide synchronization clock that is used by all nodes. Within each node, a carefully-designed PLL system recovers the synchronization reliably, even in the case of network congestion. Hardware nodes provide this clock and in each system, there is one master node which sends the clock signal to the network. If it should be disconnected, or stop sending the clock for any reason, another node automatically and seamlessly takes over. Or, should you wish, a Livewire+ network can easily be synchronized to input provided by an external “house clock” or the IEEE-1588 clock required for AES67 signals.

In broadcast studios we care very much about audio delay in the microphone-to-headphones path for live announcers. Maximum delay must be held to around 10ms or announcers will start to complain of comb-filter or echo problems. We need to consider that this is a total “delay budget” and that multiple links and some processing will often be in the path. So we’ve decided to have a link delay around 1ms end-to-end for anything in this path, allowing us a few links or maybe a couple of links and a processor before we get into links: one from the mic node to the mix engine and one from the engine to the headphones out node. Thus, 2ms total.

Delay	Effect
1-3 ms	Undetectable
3-10 ms	Audible shift in voice character (comb filter effect)
10-30 ms	A slight echo turning to obvious slap at 25-30ms
30-50 ms	Disturbing echo, disorienting the announcer
> 50 ms	Too much delay for live monitoring

Here are the air-talent reactions to delay in a test conducted by Jeff Goode at WFMS in Indianapolis

In our experience, delays to around 10ms are not a problem. From 10-25ms announcers are annoyed but can work live; anything above 25-30ms is unacceptable.

Another way to think about delay: Audio traveling 1 foot (0.3 meters) in air takes about 1ms to go this distance.

And another data point: A common professional A-to-D or D-to-A converter has about .75ms delay.

But, as is universally the case in engineering, there is a tradeoff – otherwise known as the “if you want the rainbow, you gotta put up with the rain” principle. To have low delay in a packet network, we need to send streams with small packets, each containing only a few accumulated samples, and send them at a rapid rate. Bigger packets would be more efficient because there would be fewer of them and they would come at slower rate. But they would require longer buffers and thus impose more delay. Big packets would also have the advantage that the necessary packet header overhead would be applied to more samples, which would more effectively use network bandwidth.

With Livewire+, we enjoy our rainbow and avoid the rain by having different stream types: Livestreams use small and more frequent packets, while Standard Streams have bigger and less frequent packets.

Livestreams require dedicated hardware and achieve the required very low delay for microphone-to-headphone paths. PCs are not able to handle these small packets flying by so quickly, therefore they use the Standard Streams. These are compatible with Internet standards and can be directly received into the network from PCs running delivery software. The network delay in this case is around 5ms and the PC’s latency is likely to add perhaps 50-100ms more. Since PCs are playing the files and they are not in live paths needing Livestreams, this is not a problem. Our only concern is how long it takes audio to start after pressing the On button, and delays in this range are acceptable.

Standard Streams can also be sent from the network to PCs for listening and recording. Again, this small delay is not an issue – especially given that PC media players have multiple seconds of buffering. However, off-the-shelf PC hardware with a special operating system and software optimized for real time audio is able to handle the fast streams. Indeed, we use this approach for our studio mixing and processing engine.

All Livewire+ hardware devices transmit both stream types and can receive both stream types. There is no inefficiency from having both available because all streams stop at the Ethernet switch and take no system network bandwidth unless they are subscribed to by a receiver or node. Each receiver takes only the one it needs, taking the low-delay version if available, or the higher-delay version if not. The selection happens transparently with no user action needed; users just select the channel they want and audio is delivered by whichever method is appropriate for the equipment they are using.

Livewire+'s low-delay streams are also fixed-delay. The delay is constant, regardless of the system size or anything else. In fact, a source being received at multiple Nodes will have a differential delay of less than 5 μ s - less than 1/4 sample at Livewire+'s 48kHz rate.

The Pac-Man of Protocols: Internet Standards

We use the Internet's IP standard for streaming media called RTP/IP for Standard Streams. RTP stands for Real-Time Protocol. It is the internet's standard way to transport streaming audio and video, just as TCP/IP is the standard for general data. Both use the same underlying IP packet structure, but each has a header and transmission method appropriate to the content.

Since we adhere to Internet standards, your audio may be played by PC players such as Windows Media and VLC, which support standard protocols and uncompressed PCM audio.

Converged Networks

In the telephone and networking worlds, IP has become the "Pac-Man" of protocols, eating up everything in sight. Major networking companies like Cisco and HP are dedicated to the idea that a facility needs only one network for data, telephones, and media. Many companies have jumped on this idea supporting the products that these companies are building today.

Meanwhile, PBX companies like Alcatel-Lucent, Nortel, Mitel, AT&T and Siemens have plunged into IP transport for their telephone products. This is bringing converged networks, serving all needs from PBX companies too.

Skepticism from a few years ago, which questioned whether Ethernet would handle convergence with services like telephone and live media while sharing the network with computers, has been patently disproved since then. (A network technology called ATM was proposed as a better solution. But it was expensive, difficult to administer, and would have required a "fork-lift" upgrade to existing systems. So it never caught on and has pretty much faded from sight for local networks, although it has a role at the core of some Telco networks.) Ethernet's switching, priority mechanisms, and increasingly fast speed has put most concerns to rest, and all the vendors who offer VoIP telephones connect them over Ethernet, not ATM.

Ethernet might just as well be said to be the Pac-Man of local networks too. It has nearly a 100% share of new LAN installations and is the network that all VoIP phone systems we know about use for connection to the desktop.

An Ethernet network being used for Livewire+ audio may be shared with any other data transmissions such as file transfers, web browsing, and the like. An Ethernet system with a switch at the center may have a mix of audio nodes and normal servers, PCs, et cetera. The Ethernet switch directs traffic only to where it is needed. Even on a single link, traffic can be mixed because we use modern Ethernet's priority mechanism to be sure audio packets have first call on the link's bandwidth. A studio audio delivery system could use this capability, for example, to download an audio file from a server while simultaneously playing another, live.

Livewire+ adds to the convergence possibilities in a broadcast facility. In fact, nearly all new broadcast facilities today have computer data, telephone, audio, and control on a single network that uses computer/telephone industry standard wiring; many of our customers have had this benefit for years.

2

What Can You Do With Livewire+?

Imagine everything that you can do with a PC connected to a network: Share files, send and receive emails, chat, surf the web, listen to audio, etc., etc.. PCs and networks are designed to be general-purpose enablers. You have a similarly wide range of possibilities for audio applications using Livewire+. Here are examples, starting with the most simple, and continuing to the most interesting.

Make A Snake



Concert sound guys need to get a lot of audio from the stage to their mixing consoles in the center of the house. They call the multi-conductor cables they traditionally use for this function a “snake”. Livewire+ lets you put such a snake on a diet! A single Ethernet cable connects multiple audio channels. Add a switch at each end and you can have as many nodes as you want. Use Gigabit Ethernet and you can have hundreds of channels. Add fiber optic media converters and cable to extend the distance between units to many kilometers. Maybe you need to get something from here to there?

A High-Performance Sound Card Replacement



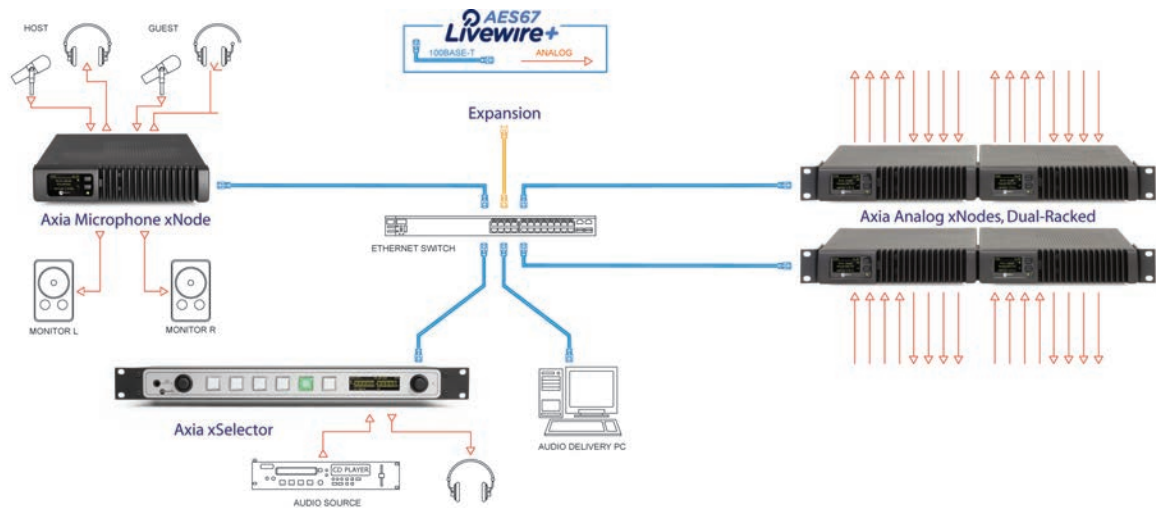
Livewire+ can talk directly to PCs, making the network look like a soundcard to delivery systems, editors, etc. Axia xNodes have excellent audio performance: Balanced I/O with more than 100dB dynamic range, < 0.005% distortion, headroom to +24dBu, etc. They make excellent multi-channel “soundcards” for professional applications. You can position the xNode at a distance from the PC, and you get balanced audio on connectors that are a lot more reliable than mini phone jacks.

With the addition of an Ethernet switch you can feed your audio to multiple computers and/or have multiple I/O boxes – which take us to the next application...

Build An Audio Router

A system with Livewire+ nodes, one or more Ethernet switches, and PC-based routing controller software make an excellent facility-wide audio router. PCs send and receive audio directly to the network without soundcards or audio ports, thus lowering costs and eliminating conversion steps. telephone, codec, streaming encoders, time-management solutions and audio processing equipment from Telos Alliance brands (Telos, Omnia, 25-Seven, Axia and Linear Acoustic) is now able to connect directly, as are numerous products from other broadcast manufacturers such as International Datacasting, AudioScience, RCS, Digigram and BSI, to name a few (for a complete list, see www.TelosAlliance.com/Axia/partners).

To interface conventional analog and AES signals, Livewire+ interface nodes come in a number of versions. One (called xSelector) operates like a traditional audio router X-Y control panel, but with a difference: audio in and out is available on the same box.



A PC-based router control package, called PathfinderPC, is available that makes your whole system, with all of its disparate parts, capable of being controlled and supervised as a single entity. You can control which outputs are connected to which inputs just as if the system were a single-location box.

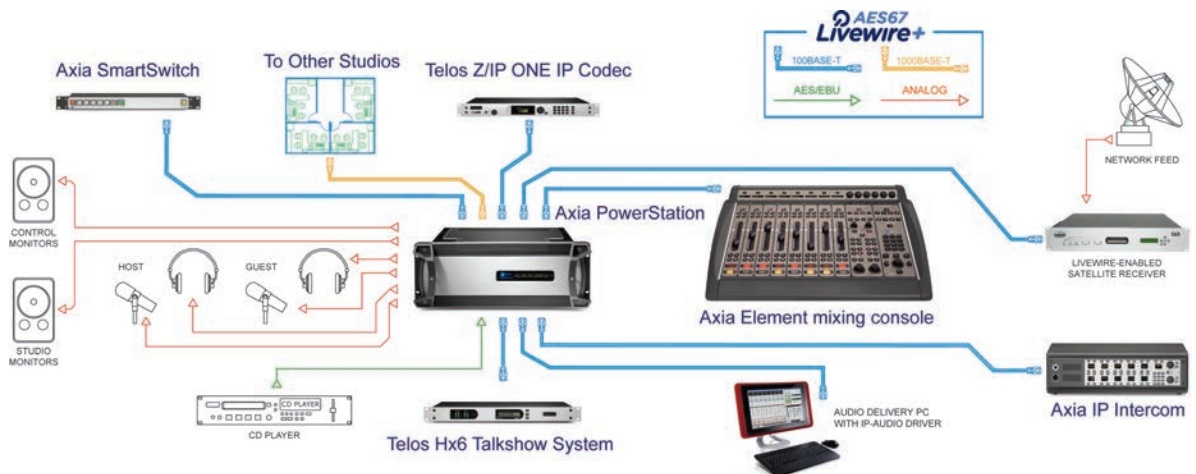
Since there is no requirement for a mainframe, the base cost is low – you can make a small system at very reasonable cost and expand it over time. Indeed, the total cost of a large system will be much lower than older approaches due to the use of commodity switches at the core. Just as using standard PCs to play audio makes much more sense than any proprietary approach, building routers from common computer industry parts makes similar sense. Indeed, this approach gives you a true “audio network” quite unlike other approaches.

Build A State-Of-The-Art Broadcast Studio

Plug an audio processing engine and a control surface into the network and you have a modern radio studio with many advantages over the old way:

- ◆ Simplified and unified cabling for audio, control, general data, and telephone.
- ◆ No multiple conversions. With most studio audio coming from or going to PCs, audio is kept in the networked digital domain. Audio may be monitored on any PC with a player such as Windows Media, VLC player, etc.
- ◆ Integrated data means you are ready for synchronized text and metadata, which is needed for HD-Radio in the USA and DAB in Europe, as well as for streaming audio channels. It will also be possible for audio processor parameters to be controlled by metadata, depending upon source characteristics.
- ◆ Tighter integration with delivery systems means that mixing, scheduling, and playing can work together. For example, song titles can appear on the mixer surface, start and other control functions may be conveyed over the network, and logging can confirm that an audio piece was really played on the air.
- ◆ Troubleshooting and repair are transformed. Extensive diagnostics are available over the same network that connects the audio. A suspect surface or engine may be swapped by re-plugging only one Ethernet cable.
- ◆ Low-cost power. Computers replaced cart machines because they are a lot more powerful, convenient, reliable, and cheap. The technical side of radio broadcasting is tiny compared the computer and networking industries. We get tremendous value by plugging into the massive R&D and production scale offered by the computer world. Leveraging low-cost mass-produced computer components makes the same sense for studio mixing and audio distribution as it did for cart machine replacement.
- ◆ Surround-ready. As one would expect from its flexible computer technology-based origins, Livewire easily handles common audio formats such as surround sound, which is quite prevalent in television.

In the example below, a Livewire+-based system is being used as a studio console. Sources such as microphones and CD players are interfaced to the network with a node in the studio, while sources such as network feeds interface with a node in an equipment room.



Certain peripheral equipment connects directly to the network. Audio from the delivery PC goes to the network via an Ethernet connection and control is also over the network. Telos Zephyrs or Telos telephone hybrids can connect directly to the switch to make the audio and control connections much simpler, only one RJ-45. The network also supports file transfers to the delivery system from a server. The studio operator surface controls a rack-mount mix engine, which has a single Ethernet connection for both control and audio.

Make a Flexible Two-Way Multi-Channel STL

Studio and transmitter sites may be linked with “Ethernet STLs”. Livewire+ nodes provide audio interface to Ethernet point-to-point radios. These are off-the-shelf today from such companies as Dragonwave, Cambium, Ubiquiti and many more. In addition to the audio, anything that can be carried over Ethernet can be conveyed over the radio link, such as VoIP telephones, email, file transfers, and transmitter remote control.



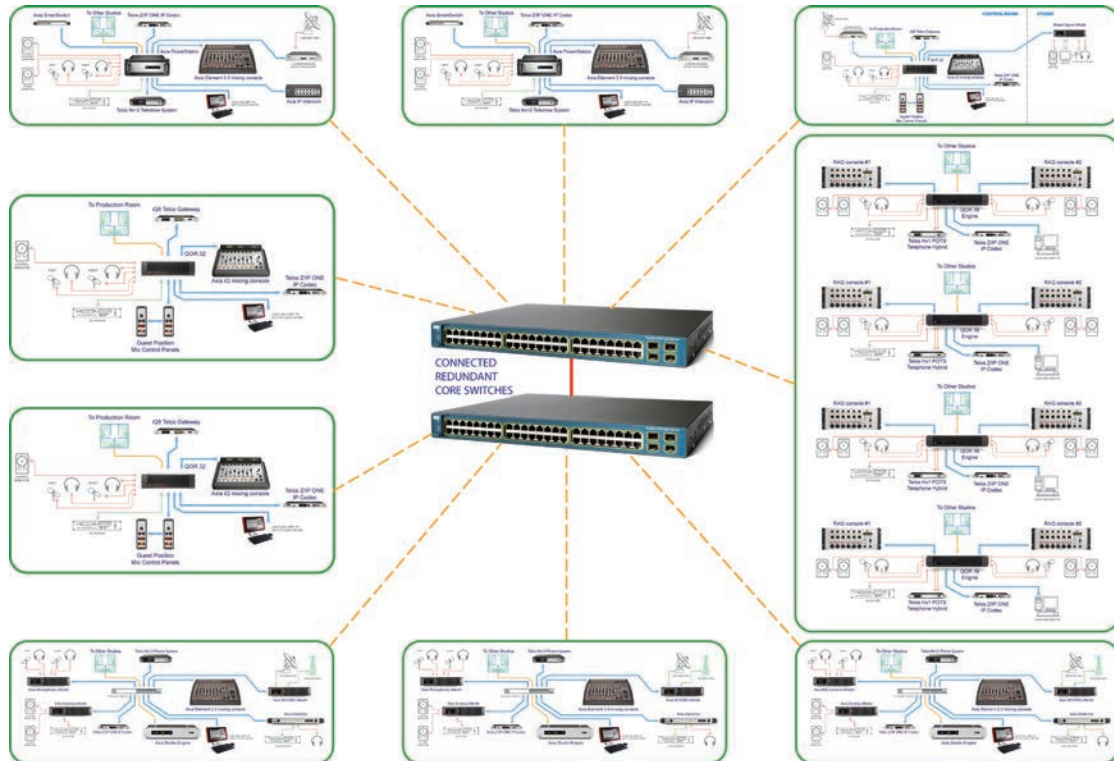
Ethernet radio systems are available that can connect at speeds of 100Mbps and better. This data rate would support more than a dozen stereo uncompressed audio channels in each direction, with capacity remaining for VoIP telephone and facilities control. Seems these radios would make an interesting two-way RPU also. For co-owned stations that are not co-located, these could be an effective way to link studio facilities.

NOTE:

Many of these systems are optimized for speed vs low error rate and therefore may not work. Axia has evaluated several units and we can offer guidance if you are interested in pursuing this option.

Create A Facility-Wide Audio Network That Includes Integrated Studio Consoles

Combine all of the above for maximum power, convenience, and flexibility. You get facility-wide audio routing, state-of-the-art studio mixing, a single wiring infrastructure for audio, computer data, control, and telephone.



Audio processors with Livewire+ ports may easily have multi-channel outputs, such as for simultaneous analog FM, HD Radio, and low-delay monitoring feeds. A single Ethernet would serve for all needed inputs and outputs. With a data capability alongside the audio, it would be possible to control processing parameters depending upon which audio source is active.

Create An Integrated National/Local Radio Network

Imagine a satellite transmitting IP packets. Now live audio, audio can be stored for later play, and identifying data can be delivered. Wouldn't this transform radio networks into something much more interesting, useful, and powerful? Including an Internet-based return path adds a whole new dimension.

3 The Components Of Livewire+

Livewire+ is not only a technology. It is a solution, made for broadcast. Here are the essential pieces that put Livewire+ to work for you.

A Livewire+ system usually has a mix of hardware nodes and PCs with driver software that lets them send and receive Livewire+ audio streams. There will also be one or more Ethernet switches, unless you are making only a very simple 2-box snake or a PC soundcard replacement. This section gives an overview; currently available switches are covered in a later section.

Many equipment manufacturers are offering equipment with Livewire+ onboard, and we expect many more in the future. Of course, this capability is a part of most new Telos Alliance products, meaning you can now get Telos telephone hybrids and systems, Z/IP ONE IP codecs, Omnia and Linear Acoustic audio processors, 25-Seven time management gear and Z/IPStream web-streaming equipment with Livewire+ connectivity.

xNodes

These interface traditional audio to the Livewire+ network. Some are used to interface GPIO such as for starting CD players or lighting on-air signs. Configuration and monitoring is via a networked web browser.



Analog xNode: Four stereo / eight mono balanced inputs and outputs with more than 100dB dynamic range, < 0.005% distortion, headroom to +24dBu. Software controlled gain lets you trim adjust to accommodate different levels. Front panel OLED audio level metering, can be PoE-powered.



AES/EBU xNode: Four AES3 inputs and outputs. An input can be used to sync your Livewire+ network to your house AES clock, if desired.



Mic + Line xNode: Four microphone inputs with very high-grade pre-amps, Phantom power, and four stereo / eight mono balanced line outputs. Like the other xNodes, it can also run on PoE connections.



Mixed-Signal xNode: This multi-purpose xNode provides one switchable Mic/Line analog input, two analog line inputs (dedicated), 3 analog line outputs, 1 digital AES3 input and 1 AES3 output, and 2 GPIO ports.



General Purpose Input/Output xNode: This GPIO interface for parallel closures has eight DB-15 connectors, each with five inputs and five outputs. Interfaces control to CD players, delivery systems, on-air lights, etc. that need simple parallel control. Many Axia mixing engines, such as the PowerStation and QOR mix engine families, also offers identical GPIO functionality in addition to audio I/O and Ethernet switch ports.



SDI xNode: Essential for using AoIP networks in TV plants. Supports up to two separate, independent SDI paths; de-embeds up to 16-channels of audio (8 stereo pairs) from the incoming SDI streams and makes that audio available to the Livewire+ network. Additionally, it is possible to take audio from the incoming streams, shuffle pairs, and create two unique outgoing SDI streams with matched audio/video latency.



X2 Router Node: Emulates the function of traditional X-Y audio router controllers. The OLED on the left presents a list of available audio sources, which are selected with the adjacent knob. The OLED on the right presents a list of possible output destinations. The unit can be used for equipment room monitoring and production studio or newsroom audio interface.




xSelector Node: An X-Y router controller with a twist: it has its own pair of local inputs and outputs. Not only can it be used to choose any network audio stream and route it to the output destination, xSelector can also contribute audio streams to the network. Many users find these indispensable in NOCs, terminal rooms or Master Control racks; they are also an excellent way to outfit news bullpens or other audio contribution stations where only a single input and output are needed.

There are also a number of switch panels that work with Axia PathfinderPC routing control software to enable routing scene changes with a single button press.

Axia IP-Audio Driver for Windows®

This is the software interface between your PC audio applications and the Livewire+ network.



24-in/24-out Driver: This is a driver that interfaces 24 stereo inputs and outputs and is available through one of our delivery system partners. It provides these functions:

- ◆ Interface for audio sent to Livewire+ from audio applications such as delivery systems and other audio players.
- ◆ Interface to receive audio from Livewire+ into applications such as audio editors.
- ◆ A GPIO function to convey “button-press’ data from the network to applications, such as from a control surface fader start button to an audio player.

8-in/8-out Driver: This is an eight-input/output version of the OEM Driver described above, available from our delivery system partners.

4-in/4-out Driver: This is a four-input/output version of the Driver, available for direct purchase from any Axia distributor. Like its OEM cousins, it works with any delivery system and editor that support standard Windows audio. It provides these functions:

- ◆ The Axia Windows® Driver connects PC audio directly to the network via Ethernet and without sound cards.
- ◆ Provides four stereo inputs and four stereo outputs via the Livewire+ network.
- ◆ Audio applications see the Livewire+ network as if it were a standard sound card. A sample rate converter and clock generation functions are included.

1-in/1-out Driver: A single-channel version of the retail 4-in/4-out Driver described above, ideal for workstations or audio capture/editing applications.

If you have a Linux plant, a Linux version of the Driver is available for specific OS distros from our partner, Paravel Systems.

iPlay (PC Router Selector)

The next application in the Livewire+ Windows Suite is an interface to display and select Livewire+ streams – essentially a software version of the Router Selector. The selected audio is sent to any audio application that works with standard Windows sound cards. The Preview function lets you listen directly without another application.

Sources are listed for selection with a mouse click. They may be filtered by category.



There is a capability similar to the radio buttons on the hardware Router Selector. Dragging a listed source to one of the buttons allows it to be used to quickly select a desired source.

In addition, Livewire+ streams are based on standards and can be played with standard Internet audio players such as Microsoft Windows Media and VLC players, with the use of SDP files.

Mixing Engines

With all audio sources in your facility available on a single Ethernet jack, the door is open to new ways of mixing and processing audio signals. We are now able to build a low-cost, but very powerful mixing/processing engine that subscribes to networked audio streams, modifies them and presents the resulting streams back to the network on that same jack. One mixing engine is required for each console mixing surface.

StudioEngine



The Axia StudioEngine is a powerful stand-alone processor for Axia Fusion and Element mixing consoles. It has no I/O of its own, taking all audio from the Livewire+ network via Gigabit Ethernet and performing internal mixing, with mixed output sent back via the same Livewire+ port. The StudioEngine performs all the mixing and signal processing functions that would have been performed in the past by an audio console. Of course, a Livewire+-based routing system may be used with any traditional console, but integration brings many advantages.

Each StudioEngine can perform all the mixing and processing functions needed by even the largest console, with per-channel mix-minus feeds, multiple outputs and monitor feeds, EQ, mic processing, et cetera. There's plenty of headroom to support future features.

The front panel display on the StudioEngine provides confidence feedback. The selector knob allows you to easily perform basic configuration. As with all Livewire+ components, web-based interaction is used for more advanced configuration.

PowerStation



PowerStation is an “integrated”, all-in-one mixing engine for Axia Fusion and Element mixing consoles. It includes audio I/O, GPIO, a console power supply, and a zero-configuration Livewire+ network switch into one easy-to-deploy package.

There are two PowerStation models, a Main and an Aux. Each PowerStation Main provides:

- ◆ 4 Analog inputs and 6 Analog outputs
- ◆ 2 AES/EBU inputs and 2 AES/EBU outputs
- ◆ 2 Microphone inputs with selectable Phantom power
- ◆ 4 GPIO machine-control logic ports, each with 5 inputs and 5 outputs
- ◆ An integrated network switch with 14 100Base-T Ethernet ports and 2 1000Base-T (Gigabit) ports with SFP
- ◆ Heavy-duty power supply
- ◆ Industrial CPU designed for harsh-environment reliability

Simple Networking built in to the PowerStation switch allows daisy-chain connection of up to 4 PowerStation-based studios without the use of an external core switch. Connecting a PowerStation Aux adds auto-switching, redundant backup power — as well as doubling audio I/O and GPIO capacity. If you require more I/O, xNodes can be added at will to expand capacity.

QOR Mixing Engines

Like the PowerStation integrated engine above, QOR mixing engines are self-contained units with DSP mixing engine, audio and logic I/O ports, and a network switch with Livewire+ ports and Gigabit trunking. There are two models, which work with Axia iQ, Radius, DESQ and RAQ mixing consoles. If more I/O than that built-in is required, xNodes can be added to expand those capabilities. Like the larger Axia mixing engines, these also provide automatic mix-minus, IFB/talkback paths and source EQ.



QOR.32 is suited for use with iQ or Radius consoles of up to 24 faders (or 3 frames), but can also be used with smaller DESQ and RAQ consoles where large amounts of local I/O are desired. It includes:

- ◆ 4 mic inputs with selectable Phantom power
- ◆ 16 Analog inputs
- ◆ 8 Analog outputs
- ◆ 2 AES/EBU inputs and outputs
- ◆ 8 GPIO ports, each with 5 inputs and 5 outputs
- ◆ Zero-configuration Ethernet switch with 6 100Base-T Livewire+ ports (4 with PoE) and 2 Gigabit Livewire+ ports (RJ-45 & SFP).
- ◆ Redundant, auto-switching backup power can be added by connecting a QOR Backup power supply.



QOR.16 is designed for personal “desktop” mixer installations and small-studio applications. It has half the I/O count of the QOR.32.

- ◆ 2 mic inputs with selectable Phantom power
- ◆ 8 Analog inputs
- ◆ 4 Analog outputs
- ◆ 1 AES/EBU input and 1 AES/EBU output
- ◆ 4 GPIO ports
- ◆ Zero-configuration Ethernet switch with 6 100Base-T Livewire+ ports (4 with PoE) , and 2 Gigabit Livewire+ ports (RJ-45 & SFP) for connection to a larger network.

Broadcast Consoles

Of course, operators still need to have control interfaces. The Axia consoles (and virtual console software!) described below work with the mixing engines just shown. They, too, connect with a single Ethernet plug.

Designed for the needs of live programming, Axia consoles provide your on-air staff with a familiar and comfortable set of controls in an uncluttered and intuitive format. With a lot of broadcast experience under our belts, we worked carefully to keep the basic functions simple and trouble free, but still include all the sophisticated functions of large traditional consoles supported in a deeper layer.

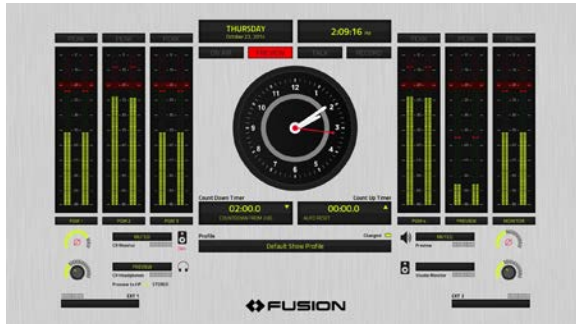
Fusion Broadcast Console



Axia’s most sophisticated broadcast console to date, designed after more than a decade of experience with AoIP console design. Fusion is a modular console, with faders presented in “gangs” of 4 per module. There are a large variety of module types; some that directly-control Telos phone systems, others that provide full-plant intercom services over the same IP network upon which all your other plant audio rides.

Metering, time, timer, and essential status info are clearly presented on an adjacent DVI LCD monitor of your choosing. Pressing the Option

knob on fader channels brings up all the fancy stuff. All sources in your Livewire+ system are listed and available for selection, and there is pan, EQ, L/R select, send bus access, mic processing, headphone EQ, and more.



In addition, there are “expert user” modes that allow engineers to configure very advanced fader/source behaviors, based on channel-strip status. Monitor assignments, bus assignments, even mix-minus sources can all be altered automatically, without operator intervention, based upon channel on/off status or the state of the source’s assignment to the Preview bus – a power-user’s dream. No longer is the operator saddled with the heavy lifting of complicated bus assignments and mix-minus generation, as the console and mixing engine relieve him or her of those tasks.

But it goes further. As you would expect from Telos and Axia, our control surfaces have a smart approach to mix-minus for phones and codecs. Every channel has the ability to provide a mix-minus output automatically. Operators simply select a phone or codec source and the backfeed is automatically generated based on preferences established when the user profile was configured. There is a single button that selects a Phone Record mode when users need to record phones off-air for later play.

High-resolution OLED displays on the overbridge above each fader show the active source, and icons indicate status when needed, along with full-time confidence meters.

Fusion can save profiles for each user, allowing different preferences, layouts and defaults for a variety of shows and talent.

In addition to console functions, Fusion provides controls and displays that interact with phone systems, codecs, editors, PC delivery systems, and other broadcast gear. Together with the StudioEngine or PowerStation mixing engines, Fusion was designed to meet all the console/control needs of the most demanding live and live-assist radio operations.

Element Broadcast Console



Axia’s most-popular, longest-running broadcast console, installed in more than 4,500 studios around the world. Like the Fusion, Element is a modular console that comes in multiple frame sizes and supports as many as 28 faders in a single frame or up to 40 in multiple linked frames. Metering and status info are likewise presented on an adjacent monitor.

Also like Fusion, every channel is equipped with automatic mix-minus, and there is enough horsepower in our mixing engines to provide one for every fader in operation, were such a situation to occur.

Clear, bright LCD displays complement each fader strip and display active source and status icons when needed. Show Profiles can be saved for individual users or shows.

iQ and Radius Broadcast Consoles



Where Fusion and Element are well-suited for large operations, iQ and Radius consoles are intended for medium-sized control rooms. Each come with a matched QOR mix engine (explained in the previous section); they require no external monitor for meters, clocks etc., as this is provided via high-rez OLED displays (iQ) or multi-segment LED displays (Radius). The iQ and Radius main units each contain 8 faders, but may be expanded up to 24 faders using a combination of accessory frames which may also provide features such as built-in multi-line phone system control, user keys which can be programmed to execute GPIO commands or router salvos, etc.

Although visually different the iQ and Radius are functionally the same; you may choose whichever meter style suits your preference. Each should be paired with a QOR.32 mixing engine in order to provide the I/O needed to accommodate a good-sized studio with associated audio and GPIO connections, etc.



DESQ and RAQ Broadcast Consoles

Axia DESQ and RAQ are compact consoles designed for tight spaces, such as dubbing stations, remote road kits, turret-based news assembly booths, or voiceover booths. Each provides two stereo mixing buses and 6 faders to work with.



Like their bigger brothers, DESQ and RAQ feature automatic mix-minus, source EQ and user profiles that can recall frequently-used configurations.

They work with the same QOR mixing engines used with iQ and Radius mixers, but with an added benefit: a single QOR engine can power two DESQs, two RAQs, or a combination of one of each. This allows the construction of networked small rooms with a very high cost-to-benefit ratio, as the cost of equipping each mixing surface with a mix engine drops dramatically, and since these smaller rooms are also now part of your greater network, they can access and use any audio source from any room in your plant — previously a task too expensive to accomplish for small studios.

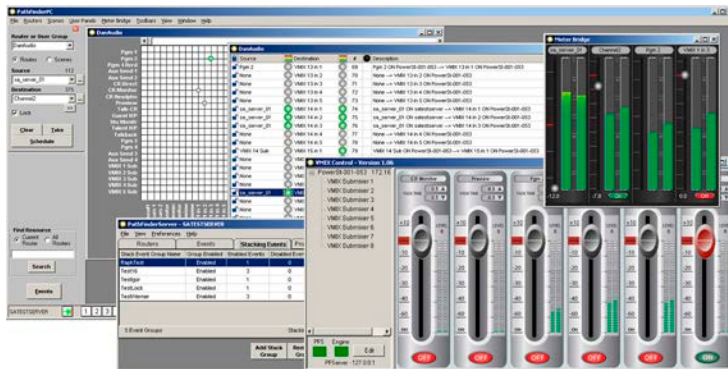
SoftSurface Virtual Consoles



Perhaps the idea of a virtual console – a mixing engine controlled by software only – interests you, for use in giving talent access to mixing capabilities in round-table discussions, in-studio performances or in-person interviews. Axia SoftSurface is such a product, with the ability to run on any Windows-based PC, laptop or touchscreen tablet. Using SoftSurface, talent can take complete control of a StudioEngine or PowerStation mixing engine, using it just as they would a physical console, with access to all the same controls and options. Imagine carrying a 40-fader console around in your hand! Compelling, no?

PathfinderPC Routing Control Software

Distributed Livewire+ systems can be controlled just as if they were traditional centralized audio routers. In this case, you will need a way to control the multiple nodes as if they were a single device.



We offer a PC software package called PathfinderPC, developed in cooperation with a partner, Software Authority, which specializes in router controls. It is a client-server system that serves as a front end for X-Y style router switching. The server communicates with all of the Livewire+ nodes in your system, and offers a common point of control to clients. Multiple clients can connect to the server to provide any number of control points. Each client may be optimized for a

particular style of operation or control features. For example, a master control client will probably be very different from a controller within a studio.

Scenes (presets) can be created and recalled to allow changes to the local studio or to the global network. A “virtual patch bay” function provides an intuitive way to manage routes. The server and clients run on Windows PCs.

Because Livewire+ nodes put audio level information onto the network, PathfinderPC clients are able to display level metering. This is indicated with on-screen cross point icons – green dots indicate the presence of audio. Users may also select accurate multi-segment meters for audio sources they want to check carefully.

You can use PathfinderPC to make “virtual routers,” which can be subsets of the real routers. For example, if a Livewire+ system has 128 different sources and destinations on the network, but you may only wish to use a small number of these points in a particular studio area. You can create a virtual bay that specifically includes only the sources and destinations required by this studio. This virtual router can have its own set of “snapshots” (scene changes). The virtual router also allows you to map multiple points to a single virtual point. For example, you can make a virtual source and destination that contains both the audio inputs and outputs for a particular device, and also the GPIO points. Thus when the route is made, both audio and GPIO is routed simultaneously.

PathfinderPC supports non-Livewire+ routers including video routers and machine control routers. Thus, you can make routing points in the virtual bay which will simultaneously route audio, video, GPIO, and Machine Control. This makes the software ideal as a master centralized router control package. Software Authority continues to expand our list of supported products, and the software is designed to allow us to add support for additional protocols and routers quickly.

PathfinderPC supports the use of tie lines or gateways between routers. For example if a system has both an analog video router and an SDI video router, one or several tie lines can be wired through Analog to SDI converters between the two routers. PathfinderPC will then combine the routing tables and automatically use the tie lines when necessary to get analog sources to the SDI router. The complexity is hidden from the end user. This capability allows Livewire+ terminals to extend an older and already filled router.

An even more powerful program, PathfinderPRO, allows clustering of redundant servers for large facilities. PathfinderPRO allows the multiple servers to mirror data and ensures that, should a primary server be taken offline, a backup server is instantly available to take its place, virtually guaranteeing 100% uptime.

Provisions for Redundancy and Back-up

PathfinderPC includes a silence detector that allows you to place a “watch” on a particular Livewire+ destination. If the audio level falls below a specified threshold for longer than a specified period of time, the system will switch to a backup audio source. This lets you build automatic redundancy into a signal path. If the primary and backup sources and destinations in the silence detector are assigned to different Livewire+ units and these units are wired to different AC power sources, the signal path can be maintained even in the event of a failure of a terminal or power source.

In addition, multiple PathfinderPRO servers can be simultaneously monitoring the Livewire+ network, building redundancy into the control system as well. The Livewire+ system is an ideal tool for building a redundant audio chain. Since every audio unit is an independent device, the server can automatically switch audio to a different unit if the usual one fails. With careful planning, you can even arrange your system so that the primary and backup audio units are connected to different LAN switches which are chained together using the inherent Ethernet redundancy protocols. Thus audio is continuously and reliably passed, even in the event of a LAN switch failure.

Timed Events

PathfinderPC has a simple timed-event system built into the server with which you can program events to happen at specified times. Individual routes or snapshots (scenes) can be triggered at a particular time and date or on a rotating schedule on certain days and times of the week. Events can also be created which will monitor a GPIO and initiate a snapshot (scene change) or route whenever a GPIO condition changes.

For more sophisticated timed operation, external automation systems can access and manipulate the routing tables provided by the Pathfinder server using the protocol translator. Multiple protocols may be simultaneously translated and connection may be on either IP/Ethernet or serial ports.

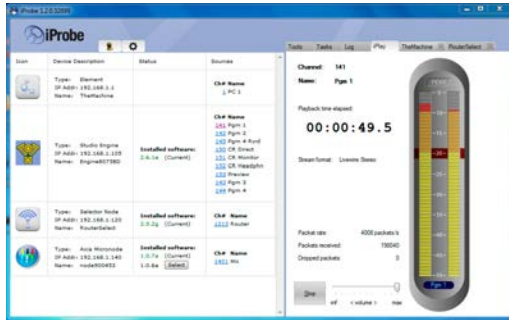
Livewire+ Audio Router Control Protocol

We provide a documented protocol for those who want to develop their own controllers. Applications designed for controlling traditional audio routers can implement Livewire+ Audio Router Protocol directly or use a software gateway between this protocol and their native protocol. The first solution may be preferable, as it enables applications to fully control every Livewire+ unit and is free from potential problems with the gateway reliability. To avoid multiple TCP/IP connections, the gateway solution may be used. In this case, there must be gateway/translator software developed for each protocol that has to be supported.

Livewire+ Routing protocol assumes multiple audio input/output nodes. Every node has its unique IP address, N input ports and M output ports.

We offer a software interface that emulates a traditional router and does the mapping and translation. Input to this module can be either serial port or TCP/IP over a network. Network configuration of Livewire+ devices can be communicated to this program using command line or a text configuration file. There is a TCP/IP server waiting on every Livewire+ node. The client simply writes text commands to the appropriate device.

iProbe Network Management Console



Axia iProbe is an intelligent network maintenance and diagnostics tool that makes managing, updating, and remote controlling a Livewire+ system very easy. There's an auto-documentation feature that generates configuration docs for every device. The organizer lets you group Audio Nodes into logical groups for easy management, upload software to single or multiple devices, make device configuration backups and more. The status of every device connected to the Livewire+ network is immediately available from a single point of control. iProbe can even perform automatic software updates for specified devices,

querying them to determine installed software versions and auto-downloading newer releases from our servers for immediate or delayed update.

iProbe also provides a centralized command interface for each connected network device; individual device Web configuration pages can be quickly accessed, and individual audio sources auditioned directly.

iProFiler Automated Program Archiving

iProFiler logging software simultaneously captures up to eight stereo audio channels to time-stamped MP3 audio logs directly from the Livewire+ network — no audio cards required. iProFiler monitors GPIO channels you select and begins recording when those channels activate. Record mode can be set for logging, skimming, or a combination of both. Logged audio can be stored locally or remotely, and may be auditioned via LAN, WAN, or Internet, if you choose.

Networked Audio Processing

With the addition of Livewire+ connections to the latest audio processing gear, it is possible to route program streams directly from control room to processor without any interim conversions and in fact, without leaving the network at all. Many Telos Alliance audio processors with Livewire+ capability are available now, with plans for more in the works. Analog and AES/EBU I/O is provided for compatibility with legacy audio chains, but with Livewire+ I/O, it is possible to build an audio chain that remains completely within the networked digital domain.



Omnia.11 is, perhaps, the most powerful FM and HD/DRM audio processor available yet. Thanks to an incredibly powerful processing platform and very sophisticated, yet easy-to-use processing tools, Omnia.11 delivers thunderous bottom end, sparkling highs, and crisp, clear voice reproduction. The firmware in Omnia.11 takes advantage of software capabilities never before possible: AGCs, compressors, and limiters analyze music in real time and adjust internal parameters for optimum performance across a broad range of material. Features include touchscreen front panel controls, Ultra-Multiband Limiter System with self-adjusting attack/release, Ultra LoIMD Distortion Controlled Clipper, Single Sideband Suppressed Carrier technology for reduced multipath distortion and many more audio-shaping tools.

ing attack/release, Ultra LoIMD Distortion Controlled Clipper, Single Sideband Suppressed Carrier technology for reduced multipath distortion and many more audio-shaping tools.



Linear Acoustic AERO.software provides enterprise-wide audio and loudness management, deployed on one or more AERO.software servers. As AoIP is being adopted at an ever-increasing rate in TV facilities, AERO.software connects via Livewire+ to provide engineers with trusted Linear Acoustic AEROMAX® loudness control, UPMAX® upmixing, and ITU-R BS.1770-3 metering, guaranteeing both compliance with worldwide audio loudness regulations, and the highest audio quality. AoIP connectivity via Livewire+ or AES67 enables I/O of any flavor and combination or integration with third party video and/or audio encoding software.



Omnia ONE is a single-rack-space, full-featured broadcast audio processor for AM, FM and HD Multicast, and studio audio processing applications. Its features include 24-bit processing, four-band AGC plus wideband AGC, and a four-band limiter. Omnia ONE also contains a fully distortion-controlled final limiter/clipper, an integrated digital stereo generator with advanced peak control, and GPI remote control ports, are of which are accessible and configurable via Web browser.



Omnia VOCO 8 makes maximum use of networked environments, packing eight separate microphone processors into a single unit. Although analog and AES3 I/O are provided, this unit really shines when connected to a Livewire+ network, enabling all inputs and processed outputs to be delivered using a single CAT-5 Ethernet connection.

Networked Codecs

Broadcast audio codecs can also become a native part of Livewire+ networks, offering the same networked control, consolidated I/O and service density described above. Here are some examples now available from Telos:



Zephyr iPort PLUS Multi-CODEC Gateway enables broadcasters to transport eight bi-directional channels of stereo audio, plus eight linear PCM channels, all with associated GPIO commands, across any network with guaranteed QoS, such as T1 and T3 connections, MPLS networks, et cetera. Each iPort contains eight stereo MPEG-AAC codecs, and connects to Livewire+ networks using a single CAT-6 cable for all I/O. The Zephyr iPort makes use of Livewire+'s advertising system in such a way that when two remote facilities are connected via its gateway, each facility's audio streams are visible to the users on the other end of the link. Operators can then select desired streams by name, and consume and control them just as if those streams were local audio sources.



Z/IP ONE IP Codec is the natural product of the work on codecs, streaming, and Audio-over-IP that Telos has been conducting for years, making generic IP networks as reliable and easy to use as ISDN. Z/IP

ONE features lots of proven tech and genuine FhG algorithms to deliver broadcast-quality audio over IP connections from 16kBit to 256kBit via UDP/RTP and TCP/IP via wired and Wi-Fi connections, with network edge traversal (NAT) and built-in directory assistance. With Z/IP ONE as a shared resource of a Livewire+ network, broadcast remotes can remain IP-based from beginning to end.

Zephyr Xstream ISDN Codec, the most widely-used broadcast codec in the world, has Livewire+ capabilities. It has all the capabilities familiar to Zephyr users, with the additional functionality of a Livewire+ interface. With this addition, the codec can become a truly shared resource within the broadcast facility, enabling users in-studio to take control of an idle codec and make use of it on-demand, from anywhere inside the studio complex.

Networked Telephone Systems

Besides the consolidated I/O, one of the advantages of equipment that connects via Livewire+ is the ability to bring control directly into the console. Phone systems can naturally benefit from this type of distributed control. One of the most concentration-breaking operations for air talent has always been phone operation; removing the operator's hands and eyes from the console to operate the phone control next to it can easily disturb the flow of on-air productions. Telos phone systems such as the Nx6, iQ6 and VX VoIP system make use of a console module specifically designed for on-console control of the phone system, enabling operators to keep their eyes on the board while simultaneously taking callers, leading to smoother shows and fewer errors.

Also, phone systems with Livewire+ can now be used as shared devices. The VX system, Telos' multi-line VoIP phone system, has a native Livewire+ connection, and has so much processing power that it can handle an entire broadcast plant and supply a high-performance digital hybrid for every incoming phone line; this mix of features means that VX can easily handle telephone needs for multiple stations and, thanks to its Livewire+ connectivity, can share lines and control between multiple studios.

Even the VSet phones which are used for off-console control of certain of these systems can plug into the Livewire+ network to communicate with the systems' phone-switching base units, and can draw power from PoE-equipped network switches, further reducing wiring cost and infrastructure complexity.

Finally, call-screening and on-air contesting programs, such as those from Livewire+ partners Broadcast Bionics (which come with each Telos phone system) run on computers connected to the Livewire+ network, simplifying connection of yet another control interface.



Telos VX, the world's first multi-studio VoIP phone system for broadcast, is scalable to provide advanced phone interface capabilities to facilities of nearly any size; VX's use of standard SIP (Session Initiation Protocol) means that it can work with nearly all VoIP-based SIP PBX or SIP Telco services. A VX Engine can easily have over 100 SIP phone numbers associated with its configuration, and each VX Engine can have 48 of these numbers active at any time, available for callers to ring-in. Of the 48 active lines, any 16 may be simultaneously engaged in conversation

on-air, with full audio processing. If your in-house PBX is not VoIP-based yet, commonly-available gateways from companies like Patton allow you to connect to traditional Telco lines, including T1/E1, ISDN, and POTS. Making use of the backfeed system that is an integrated part of Livewire+, each VX studio is supplied with its own dedicated Program-on-Hold input and, of course, each caller receives a custom mix-minus generated on-the-fly, with no console operator involvement.



Telos HX6 and iQ6 are cousins, six-line digital Talkshow system with two integrated high-performance digital hybrids that can manage up to six callers. They can connect to analog POTS or digital ISDN phone lines, and are controlled, like the VX, either by on-console control modules, desktop VSet phonesets, or call-screening software. They differ in that Hx6 includes standard Analog or AES/EBU ins and outs along with its Livewire+ connection, while iQ6 eliminates legacy I/O in favor of a single Livewire+ connection, which

transports all audio inputs, outputs, mix-minus and PoH audio, plus control of course, on a single CAT-6 cable. For legacy plants migrating to AoIP, Hx6 makes sense; for a new all-Livewire+ plant, iQ6 is an excellent choice.

Networked Web-Streaming Encoders

Internet streaming of audio has been around for a long time of course (Telos first entered this field in 1999 — you may remember the AudioActive Encoder, the first professional hardware for MP3 streaming), but as consumers increase their use of cellular data networks and WiFi connections to take their entertainment with them, streaming content has finally come of age. A recent industry study showed that Web audio streaming contributed 36 billion hours to listening time in 2014¹, and a Cisco research paper predicts that online video will account for 79% of the world's IP traffic by 2018² — a staggering amount of time spent listening/viewing digital entertainment.

It makes sense, then, for generation of streaming content to be an integral part of broadcast studio networks. The Telos Alliance has fielded a full family of products which take advantage of Livewire+ networks to select, process, encode and deliver streaming audio to listeners, with tightly integrated workflows that require almost no additional infrastructure.



Z/IPStream R/1 is a rackmount hardware device which combines audio processing for bit-reduction with MP3 and AAC encoding, and dual-stream encoding for simultaneous output of high and low bitrate streams, or of different encoding methods (such as MP3 and AAC). The

unit is fed a selected program output from your Livewire+ network; the audio processing section then tailors input audio to match selected output bit-rates using AGC, 3-band compressor/limiter, EQ, low-pass filter and look-ahead final limiter. Then the stream is encoded with your choice of MPEG Layer III, AAC-LC, HE-AAC or HE-AAC v2 formats. Support for Adobe RTMP and RTP streams (including RTP multicast) are also included. Streams can then be “tagged” with “now playing” metadata received from automation systems, and the encoded audio fed to your streaming replication service via a WAN gateway.

1 Research And Markets, “[Internet Music Radio Programmers 2014 – 2017: Music Plays and Monetization Mainstays](#)”, August 2014

2 Cisco Systems Inc., “[Cisco Visual Networking Index: Forecast and Methodology, 2013–2018](#)”, June 10, 2014.



Z/IPStream X/2 and Z/IPStream 9X/2 Streaming Software Encoders for PC workstations have capabilities similar to the hardware device just described, but with added support for “adaptive streaming”, a video technology now making inroads with audio streamers. This is interesting because, using adaptive streaming, technologies such as Microsoft’s Smooth Streaming and Apple HLS formats allow media players to automatically adapt to changing network conditions by encoding at multiple bitrates, whereupon sample-aligned output streams enable listeners’ media players to seamlessly switch between bitrates in response to Internet conditions. Z/IPStream 9X/2 includes multiple processing upgrades, notably “Undo” for restoring poorly mastered source material, as well as 6-band parametric EQ and multi-band processing configurable from two to seven bands.

Taking advantage of its networked environment, Z/IPStream software has an audio replacement capability that allows some streaming audio material – such as local spots – to be automatically replaced with content from secondary source or audio files, all from Livewire+ networked sources of course. The program is cloud-ready as well, and can be monitored and configured remotely, the same as all the rest of your Livewire+ network devices.

Z/IPStream.S4 is a PCIExpress card for PCs that allows offloading of streaming content origination to a local PC. Connecting to the network via Livewire+, Z/IPStream.S4 cards can process and encode audio stored on the PC itself, then deliver it to a streaming provider, or take Livewire+ streams from elsewhere in the facility and perform processing and encoding before delivering them for distribution via a WAN gateway.

Networked Time Management

Broadcasters have long taken advantage of digital audio manipulation in order to bend time – make programming shorter or longer, to fit, with associated pitch-management routines that achieve inaudible results. The devices that perform these functions, operating in the digital domain already, are an excellent fit for the networked Livewire+ environment. They can receive and manipulate audio by way of their network connections; those same connections carry the GPIO commands which trigger START, STOP and DUMP commands. Further integration is accomplished by attaching those commands to reserved keys on the operator’s console, allowing the delay unit itself to be located in the TOC or other remote area, while the operator controls its function without removing attention from the mixing board itself.



The 25-Seven Program Delay Manager is one of the best-known of this type of unit, and is equipped with Livewire+. In addition to its primary function of blocking unwanted words from air, it has the unique ability to capture both the censored and uncensored

audio, and send it via e-mail to the appropriate manager. Because Livewire+ networks can connect, via gateway, to your plant’s business LAN, this is a very handy and easily-implemented feature.

Networked Intercom

In larger radio operations, and nearly all television plants, intercom communications are an essential piece of day-to-day functioning. In the past, whole-plant intercom meant a complicated and expensive central switching matrix, discrete wiring pairs to and from every location imaginable, and audio bandwidth so limited it sounded more akin to a burger drive-through speaker than something employed in a broadcast plant.

With IP, all of this can be radically simplified, and QoS greatly improved as well. Intercom audio becomes no different than any other audio source on your network. Point-to-point or point-to-multipoint Intercom communications can be sent and received using the Ethernet infrastructure already in place for your broadcast audio; GPIO closures for attention tallies or speaker mutes may be sent with Intercom audio in the same way logic commands are sent with audio from playout systems, or to activate recording devices, etc. With an IP-based intercom system, substantial savings are realized because the intercom audio “rides along” with the broadcast audio for free – no additional wiring infrastructure or central card cage required.

Also, a huge performance upgrade is realized. Since AoIP carries full-bandwidth digital audio, the sound from an IP-based intercom system is radically upgraded from the old 5kHz systems; in fact, it can carry full-fidelity 20Hz – 20kHz audio suitable for air. No more need to hold a breaking news story until a reporter can be mic’d; you can now simply put the intercom itself on the air.



Axia’s IP Intercom is just such a system. It’s the world’s first broadcast intercom system that’s IP based. It connects to the same Livewire+ network as your consoles, mix engines and other audio devices. There are drop-in modules for Axia consoles that enable board operators to converse directly with any location in your plant using the board mic and preview speaker already in place (rackmount and desktop stations are available as well, natch).

Further, since PCs can contribute to and consume audio from the Livewire+ network, why not turn them into intercom stations as well? Axia SoftCom software installed on any workstation or laptop turns that computer into an intercom station as well, making it possible for literally every location in your facility to be reached instantly.

4

Nuts & Bolts: Making Livewire+ Play

Now we move to making audio happen. Time to take the gear out of the shipping carton and make it play. This section gives you practical information. Details about the underlying tech are reserved for later.

Livewire+'s Channel and Name system

An advantage of having a data network carrying our audio streams is that we can send identifying information on the same cable and system. Receivers can build tables of available audio, and testers can identify specific streams on a cable. In Livewire+, we have both a numeric and a text ID for each audio source.

Hardware Livewire+ devices are configured either using a networked PC's web browser, or with local pushbuttons and displays. PC Livewire+ nodes will have a configuration window that opens when you click on the application icon. Details for each are in the manual specific to the product, but the general approach is the same for all audio and GPIO.

Channels

Channel numbers may range from 1 to 32767. You assign these to audio sources as you wish.

New units are pre-configured from the factory to start with channel 1, thus a 4-channel xNode will come assigned to channels 1-4. Two new units can be connected to each other with a "cross cable" (described later) for immediate out-of-the-box testing. For your network, you should reserve channels 1-4 for testing and not assign them for routine use. Then, if you plug a new unit into the network before you configure the channels, there will be no problem with conflicts.

In a large system, you will probably want to have a people-friendly naming and numbering system that reflects studio use or location and to help prevent accidental duplication of channel assignments (a big no no by the way). You have plenty of numbers to use, so you don't have to conserve them. For example, the channels associated with Studio 1 could start with 100, Studio 2 with 200, etc. There is no requirement that channels be assigned in order or contiguously from a multi-channel device.

Devices such as telephone hybrids and codecs need audio in both directions, so when appropriate, a single channel contains a "to device" audio stream as well as the usual "from device" audio. In this case, you can think of the channel as something like a telephone number that connects a call with audio in both directions. The advantage of this "bundling" of the two audio directions is that the association is naturally maintained when studio mixers are in the picture. Mixers generate the feed to devices (usually mix-minus, but not always) and automatically assign it to the source channel number, and this association is kept regardless of which fader is being used, etc.

Text Name

The text name may be up to 24 characters and you choose it as you wish. This is what will appear on the xSelector's display screen, console source select lists, etc. Most devices are not able to display all 24 characters, so will truncate to show what they can.

You may wish to include in the name the rack number or room name of where the Node is located, to help orient yourself in the case of a future emergency.

A typical name might be ST1CD2 for Studio 1, CD Player 2.

Our studio mixing consoles (Fusion or iQ, for e.g.) automatically generate return feeds to devices that need them, creating the text name for these in the form "To: name". For example, if you have a source called "Hybrid 1", the mixer will generate an audio stream named "To: Hybrid 1" and advertise it to receivers.

GPIO

There are also GPIO (General Purpose Input/Output) channels and text names. These work in a fashion very similar to the audio source channels and names.

GPIO channels often share the same channel number as an audio source. A typical situation would be when you have a CD player that needs start control from an audio mixing console. The mixer automatically generates the start command and puts it on the channel number you assigned to the audio source. To cause a particular hardware GPIO to output this command as a contact-closure pulse, you configure the GPIO device to listen to this channel. As with back-fed audio, control follows the audio source to whichever fader is being used.

But GPIOs may also be independent of audio sources. In this case, the Livewire+ system provides a pass-through function where outputs follow inputs – sort of like a GPIO distribution amplifier.

Sources vs. Destinations

We've always struggled with terminology when referring to audio input/output from devices such as codecs and hybrids where there's local audio I/O as well as a combined network I/O port. We will try to be consistent within the Livewire+ realm by using the following terms:

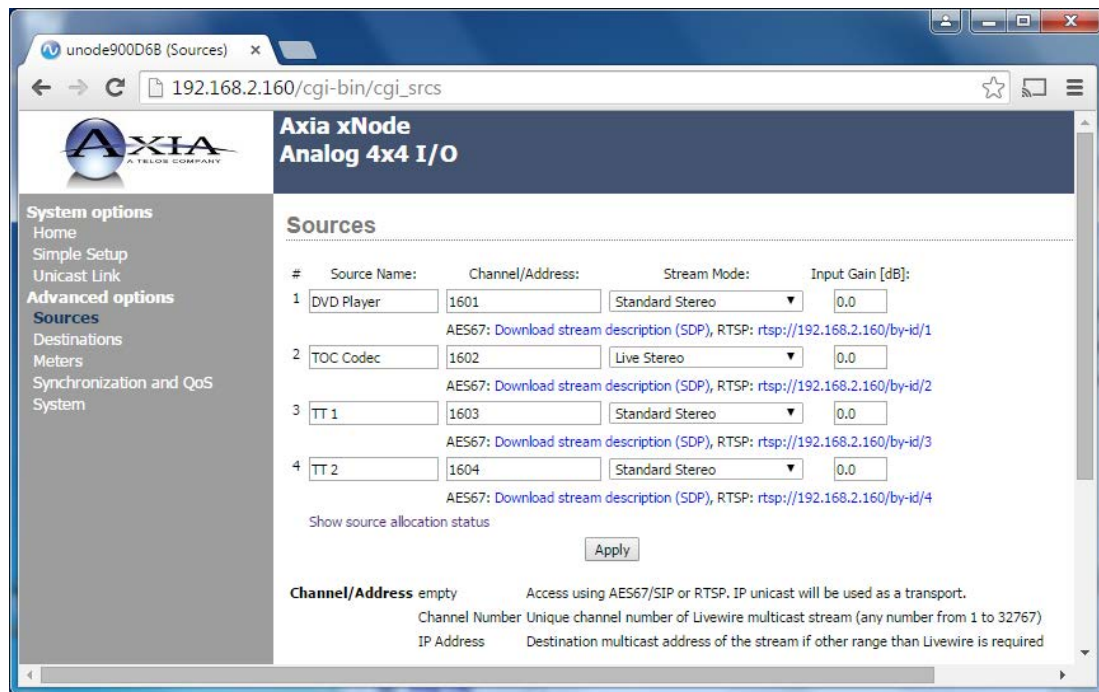
- ◆ **Source** – this is an audio input to a hardware Node and therefore available on the Livewire+ network as an audio stream that can be accessed by other Livewire+ nodes. Of course a mixing engine can generate new audio sources and in this case there is no associated hardware audio input.
- ◆ **Destination** – this is an audio output from a hardware Node and therefore represents playback of some stream from the network. Of course a mixing engine or Livewire+ capable audio device may access a Livewire+ stream and in these cases there would be no associated hardware audio output.

So, to reiterate, sources represent the feeding end of the audio stream equation whereas destinations are just that, one or more destinations where that stream is used.

Examples

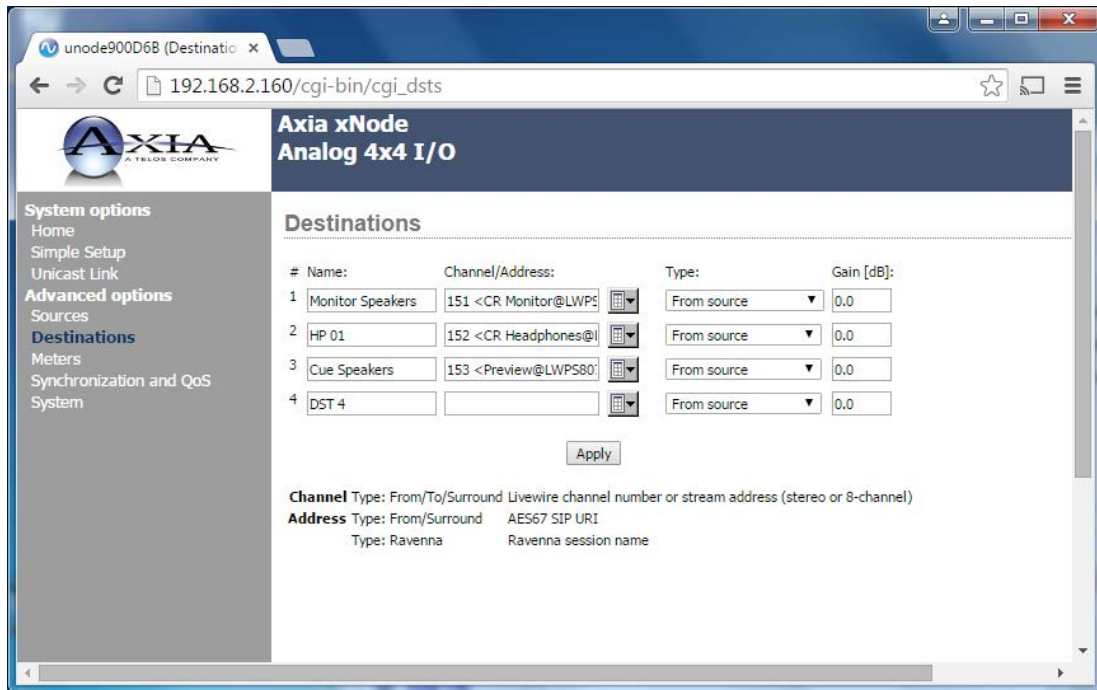
Following Gauss' dictum that "an example is worth two books," let us now turn to some to show you how Live-wire+'s channel and name identification work.

On the next few pages we'll show some of the web configuration pages for the Analog xNode. The first example is the home page that is displayed once you have logged into the node. It simply lets you navigate to the other configuration screens.



The Sources page shown above permits you to configure locally generated sources. It allows you to assign names and channels to the sources that will be generated by this node, and to configure the audio inputs associated with those sources.

- ◆ The **Source Name** entry is where you put the text ID for the node. There are 4 Source Name entries, one for each audio channel. This is the text name for the individual audio source. (xNodes can also split their 4 stereo inputs, turning them into 8 independent mono inputs; in this case, you would have 8 Source Name entries.)
- ◆ **Channel/Address** is where you enter the channel number for each source. Note the instructions at the bottom of the screen for entering AES67 sources; more on this in the xNode User Manual.
- ◆ **Stream Mode** allows you to decide whether each stream should be a Livestream, Standard Stream or Surround Stream. Why would you use one versus the others? A satellite feed will never be in a mic-to-headphones path, so only the Standard Stream would be required. A microphone source is live audio, and so would be a Livestream. You may also select Mono streams or Surround streams (only used if yours is a 5.1 Surround broadcast facility).
- ◆ You can trim **Input Gain** by specifying an amount in dB. This can also be set on the Meter screen, in case you desire to set gain "by eye".

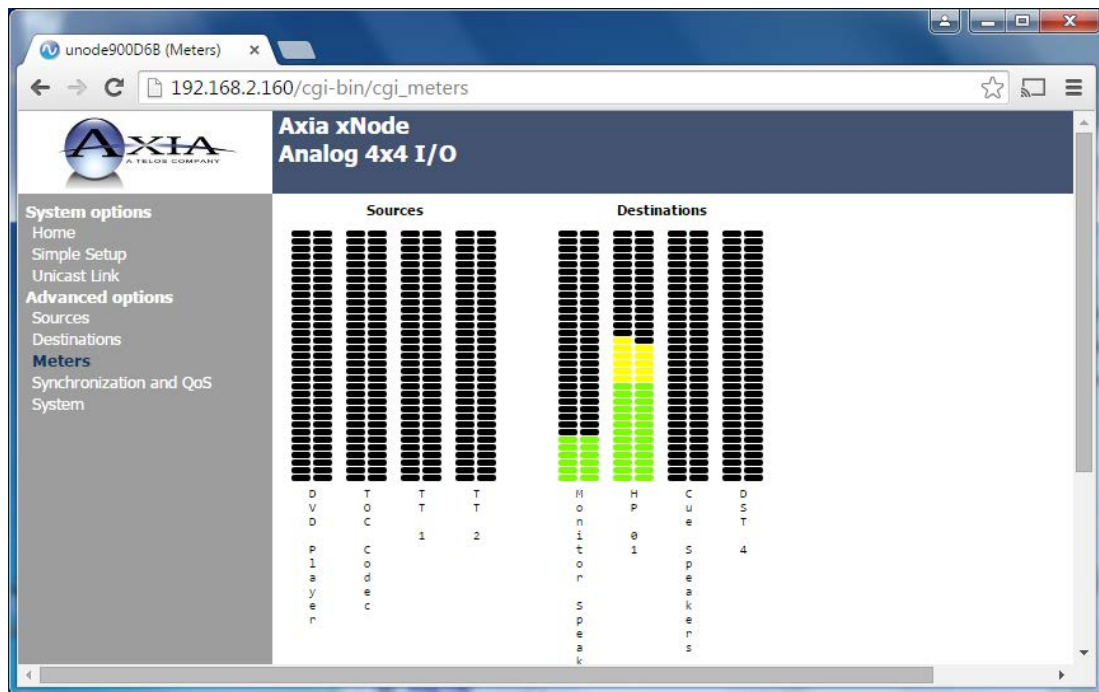


The sample Destination Screen, pictured here, is used to configure the local units' outputs, or destinations. This is where you configure the output channels, with the menu options described below.

- ◆ As with sources, you can enter a text Name for the destination associated with each output port.
- ◆ **Channel/Address** is where you tell the unit which audio stream is to be output from each audio output. As previously discussed, each Livewire+ stream is identified by both a text name and with audio channel number. You can enter the channel number directly, or you can use the button to the right of the channel entry field to open a page with a list of active audio channels; you can choose from among the channels listed within it. Usually the list contains text names of audio streams, but if no text name has been assigned, you will see the device type and IP number instead.
- ◆ **Type** lets you choose whether the Destination feed is one-way (“From Source”, outbound to a monitor or headphone feed, or a record input) or two-way (“To Source”, supplying a backfeed for a phone system or codec).
- ◆ **Gain** again appears in the Analog Node Destination screen to allow you tailor specific outputs to feed High-Z (professional) or 600-Ohm (prosumer) gear.

Hot Tip!

You will only have a complete list of audio sources if you have already configured all your source nodes and have them connected and operating so that they are advertising to the network. You can always just enter the channel number here if you don't yet have your source nodes working, but it will be more convenient if you prepare all your sources before moving on to destinations.



The Meters page for our example audio node lets you monitor the levels for each input and output channel. This is also an alternative to the source page for setting input gain.

The **Synchronization and QoS** page is an advanced feature page, and it lets you set some other values related to system clocking and AES67 operation. These are described in detail in the unit's manual, but to appreciate the context, you will also need to understand more of the Livewire+ and networking basics described later in this document.

unode900D68 (QoS)
192.168.2.160/cgi-bin/cgi_qos

Axia xNode Analog 4x4 I/O

System options
Home
Simple Setup
Unicast Link

Advanced options
Sources
Destinations
Meters

Synchronization and QoS
System

Synchronization:

Clock mode: Livewire clock priority 3 (default) ▼

PTP domain number: 0 default ▼

PTP delay mechanism: E2E - uses the end-to-end delay request-response (default) ▼

PTP clock priority1/priority2: 128 ▼ / 128 ▼

PTP clock sync interval: 0 (1s) - PTP Default profile (default) ▼

Livewire advertisement slot (default 0): 0 (0 - 1000; 0-random)

Live Audio / Clock Streams:

802.1Q VLAN 0 priority frame tagging: 6 (AES67 recommended) (default) ▼ (802.1p priority level: 0-7)

DSCP Class of Service: 48 CS6 (default) ▼

AES67 / Standard Audio Streams:

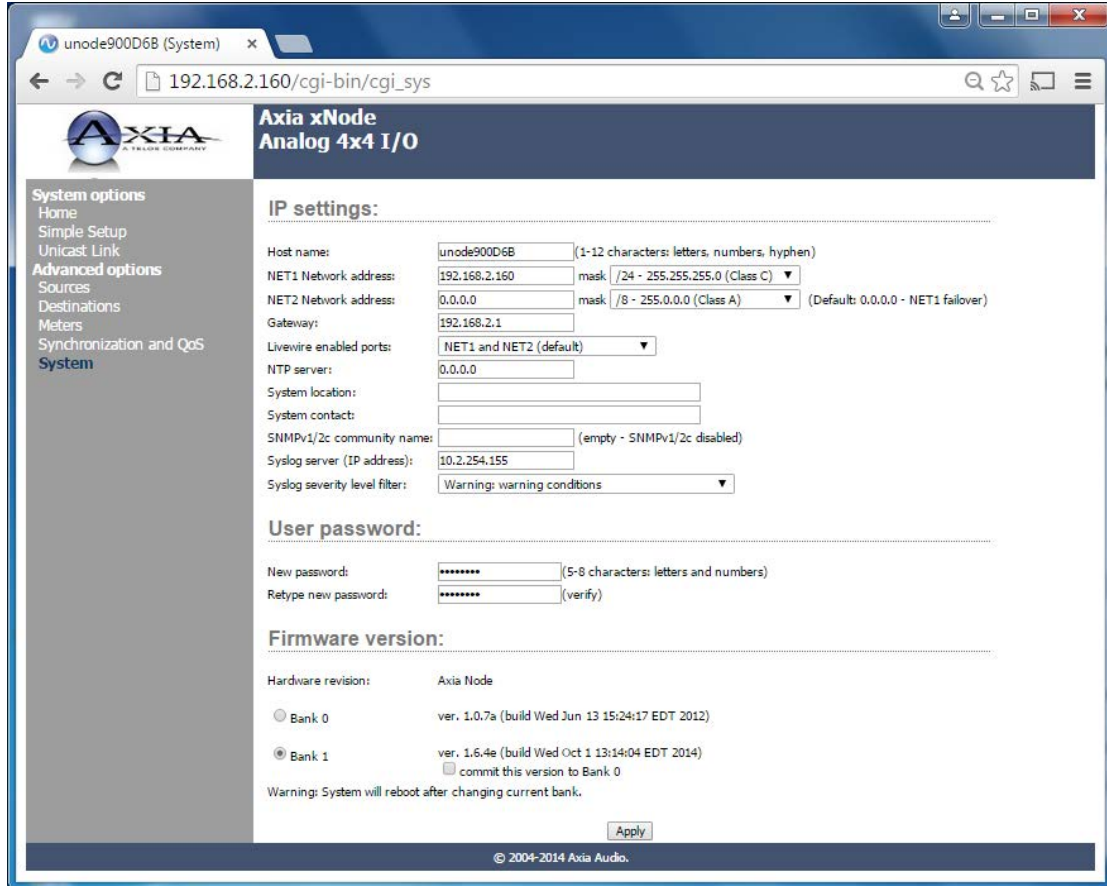
AES67 Link Offset or Receive buffer size [ms]: 20 (15 - 100; default 20)

802.1Q VLAN 0 priority frame tagging: Disabled (default) ▼ (802.1p priority level: 0-7)

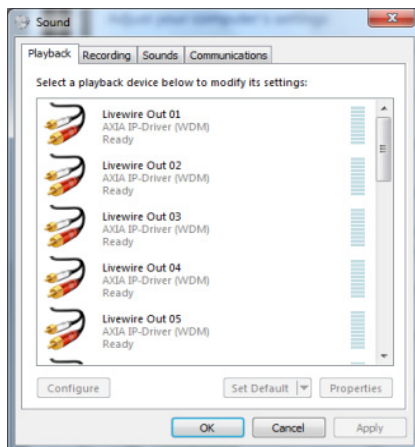
DSCP Class of Service: 46 Expedited Forwarding (default) ▼

Apply

The System page shown below permits checking the IP address and performing software updates.

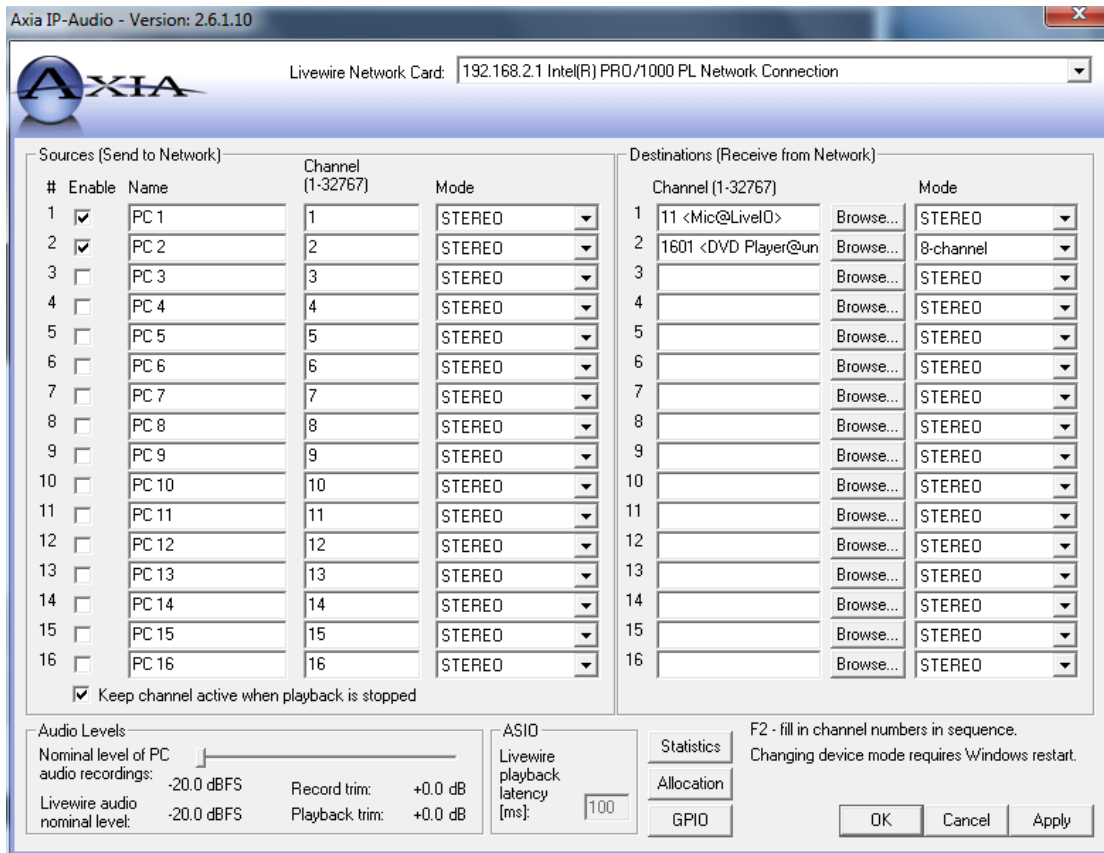


Our samples so far have been from the 4x4 Analog xNode. Next let's look at how sources and destinations are handled by the IP-Audio Driver used on Windows™ computers.



Soundcard Emulation. The IP-Audio Driver looks like a standard sound card to Windows™. Each of the Driver's sources (e.g. streams originated by this computer) shows up as a sound card. Computer sources sent to the network are called **Livewire Out XX** in the Windows™ Sounds and Multimedia Properties Control Panel; inputs to be received or recorded from the network are called **Livewire In XX**. You can define one of these Livewire+ sources as Windows' Preferred Sound Playback Device from the Windows Control Panel as shown here.

Driver Configuration. The Axia IP-Audio driver is configured for sources and destinations much like the Axia audio nodes. The driver is configured from the window shown below which is invoked either by a Start menu shortcut or a Control Panel applet; the various settings are described as well.



- ◆ **Sources and Destinations** – You can see that the node and source naming and channel idea is the same as for the hardware nodes. Any audio channels you want to receive are entered into the Destinations boxes. If you don't know the ID number, you can choose from text lists instead, by clicking on the Browse button.
- ◆ **Livewire Network Card** – A PC running this driver may have two network cards, one for general data and another for audio streams. The Livewire+ Network Card entry lets you associate Livewire+ audio with the appropriate card.
- ◆ **Allocation** – Clicking this brings up a screen that lets you set stream characteristic values. This is covered in greater detail in the Axia IP-Audio Driver manual.
- ◆ **Statistics** – This button brings up a screen with lots of information useful for debugging network problems.
- ◆ **GPIO** – Opens the configuration screen that allows GPIO triggers to be attached to individual Driver audio streams; useful, for example, when you wish your automation system to turn a console fader on when an audio stream begins to play. This is covered in more detail in the User Manual.

When you have finished configuration, the Livewire+ network looks like a sound card to any Windows application that uses the standard wavin/out or ASIO audio interface. In Windows applications where you normally select the soundcard you want to use, you will select a Livewire+ channel instead. In the example, an audio player that has selected Livewire Out01 will put its audio into Livewire+ stream channel 11111 and be available to all Livewire+ devices on the network.

The GPIO xNode is a hardware box with 6 DB-15 connectors, one for each port. Each has 5 inputs and 5 outputs.

GPIO channels may be associated with audio channels or may be independent. If they are independent, they must not use the same number as any audio channel – they share the same “channel space”.

You can monitor the status of each with the indicators on the xNode’s front panel display or configuration Web page.

xNode Configuration & Web Access

As shown in our examples above, you’ll need to configure various parameters in the hardware nodes. Some very basic parameters such as the name and IP address can be configured from the front panel. In fact, for the basic audio snake application we need not even access the xNodes’ web pages. However in most cases we will need to do so, but we will need to assign an IP address first.

Front Panel xNode Configuration

A number of items can be programmed from the front panel of most of the hardware nodes. This is covered in more detail in the individual manuals. However we will cover setting the IP address here since you’ll need to assign an IP address to enter via the web browser, and we’ll cover that next. Here’s how to assign an IP address to a typical audio xNode.

Configuring xNode IP address

Each Livewire+ node must have a unique IP address. The only exception is when two nodes are connected in the point-to-point configuration.

Fast Setup

A Fast Setup routine is built into each xNode to assign an IP address and I/O channel numbers in less than 1 minute. Here’s how:

- ◆ Apply power to the xNode and wait for the boot process to complete. You will see a home screen identifying the xNode’s type and software version.
- ◆ Press the top button to the right of the display twice, to reach the ID Page. You’ll see a Node ID with no value no IP address.
- ◆ Press the lower button, pencil icon, for 10 seconds to enter into edit mode.
- ◆ With the cursor shown, use the top button to increment value and the bottom button to move the cursor to the next position. Press “next” button twice and press the increment button once to give the ID value of 1 to the xNode. (You’ll give a different ID number to each xNode you install.)
- ◆ Pressing the move button one additional time completes edit mode and automatically configures the xNode with IP address 10.216.0.101 and channel numbers of 101 – 108.

That’s actually the end of fast setup! With IP and channel numbers configured, you can attach the xNode to your network and begin connecting sources.

Manual Configuration

If you'd rather assign an IP address manually, you can do so from the xNode's front panel:

- ◆ Press the top button on the xNode front panel until you see the screen that displays the assigned IP address and netmask.
- ◆ Press and hold the bottom Pencil button to enter edit mode.
- ◆ With the cursor shown, use the top button to increment value and the bottom button to move the cursor to the next position. Do this for each digit in the IP address field; at the end, pressing the bottom button to highlight the netmask field.
- ◆ The entire netmask field will be highlighted; press the top button to begin editing the octets or press the bottom button to complete the netmask editing.
- ◆ Pressing the top button one additional time completes edit mode and highlights the selection for the xNode's active Ethernet port. Press the top button again to complete your editing and save your new IP address.

Accessing a Node via a web browser

To access the built in web server from a computer, the computer and node must be connected to the same LAN (or the computer and node can be connected using a "crossover 100 Base-T" Ethernet cable).

Note that the IP range (e.g. the first three numbers of the four numbers of the IP address of the computer and the node must match, or additional configuration will be required.

- ◆ Open your computer's browser and type in the IP address of the xNode, such as `http://192.168.2.10` where "192.168.2.10" is the IP address of the xNode to be configured, into the URL field.
- ◆ You'll be asked for a username and password. Default Authentication is:
 - ◇ Username: user
 - ◇ Password: *(leave blank)*
- ◆ Once logged in, select the Simple Setup button to enter into a simple configuration screen for the xNode. The options available will vary based on the xNode.
- ◆ Enter descriptive text in the Source Name fields which describe what devices are connected to the xNode (e.g. CD player, Turntable, PC Out-1, Aux Input).
- ◆ Enter descriptive text into the Name field under the Destination section which documents what is connected to the xNode's outputs (e.g. Control Room Monitors, Headphone, CF Recorder, STL input).

Plugs & Cables

Livewire+ systems use primarily copper cables, but you can add fiber where it makes sense. We'll start here with copper.

An important goal in the design of Livewire+ was to simplify installations. One of the ways we do this is to let you standardize on a single cable type, plugs, patchfields, etc. This is consistent with the modern way of thinking about cabling in office buildings where a common type can serve different applications. You can use the same connectors and cables for everything in your plant. And for big new installations, outside contractors can install and test the wiring infrastructure for everything. In fact this is one reason why we suggest that Broadcast Engineers should become familiar with the relevant standards such as EIA/TIA-568-A & B.

The 100/1000BASE-T Ethernet we need for most Livewire+ devices requires RJ-45 8-pin modular plugs and jacks. So, we've standardized on RJ-45s for balanced high-level analog and AES connections as well. There are a lot of connectors being used for analog audio these days, so why did we go this route? The reasons are cost, density, compatibility, and convenience. RJ-45 sockets and plugs are a lot cheaper than other choices, both for us at the manufacturing level and for you at the time of installation. Density is an important advantage: We can get only a few XLRs across the rear panel of a 1U rack box and we need two of them for each stereo connection.

xNodes would have to be twice their size to have the same channel capacity as we have now using 1U nodes that fit in half a rack-width. A single RJ can do both channels on one jack and we can fit lots of them on a 1U box. XLRs and DBs need to be soldered, and shells assembled, etc. Molexes need a separate crimp for each wire and are not standard. RJ crimping is convenient procedure compared to the others. And you will already have the plugs, cables, and tools at hand.

The following tables summarize the cable types that could be used in a Livewire+ system and their applications.

Non-Ethernet Cabling Relevant to Livewire+ Systems

Description	Cable
Analog Audio, balanced, high-level	Usually shielded CAT-5, but unshielded with care
Balance 110-Ohm AES3 Digital Audio	Usually shielded or unshielded CAT-5 or 5e

Ethernet Cabling Relevant to Livewire+ Systems

Name	Description	Max Length	Where Used
10BASE-T	10Mbps on 2 pairs CAT-3 copper. Obsolete for new installations.	100 m.	Not used.
100BASE-TX	100Mbps on 2-pairs CAT-5 copper (CAT-6 recommended for Livewire+ to add a safety margin). Most common Ethernet media.	100 m.	Livewire+ nodes, PCs.
100BASE-FX	100Mbps on fiber Multi-Mode Single-Mode	2 km, 20 km.	Livewire+ nodes with ext. media converters
1000BASE-T	1Gigabit on 4 pairs CAT-5e copper (CAT-6 recommended for Livewire+ to add a safety margin).	100 m.	Mixing engine to switch, PCs, switch-to-switch
1000BASE-SX	1 Gigabit on short wavelength fiber, multi-mode	220-550 m.	Switch-to-switch
1000BASE-LX	1 Gigabit on long wavelength fiber, single-mode	5 km.	Switch-to-switch

CAT-5 for Audio?

Using CAT-5 “digital cable” for audio may seem strange at first, but it does make sense. The low capacitance and tight twisting requirements necessary for high-speed networks are good for analog and balanced AES audio as well. Because you have a single cable and connector type for everything in your facility, as your studios evolve, the same cable that was once used for analog may be used for AES, Livewire+ digital audio, general Ethernet data, or whatever else might come along.

Does this work in the real-world? Sure it does – as demonstrated in the many installations that have been done using the Radio Systems “StudioHub+” product family. We use the StudioHub 2 format for our connectors by the way, to allow convenient use of prefabricated connectors by Livewire+ users. More on this below.

Ethernet 100BASE-TX

Livewire+ uses 100BASE-TX copper wiring with RJ-45 style plugs and jacks for connections from audio nodes to switches.


You must use Category 5e cable and accessories or better. If you're interested in "future-proofing" your new facilities, you might also consider CAT-6 in order to be ready for 1000BASE-T. CAT-6 is not much more expensive than 5e and it does perform better, particularly when a run has lots of bends that could disturb the pair relationships within the cable jacket, or has many punch blocks and/or patch cables.

Pin numbering, jacks, cables, and color codes

Modular wall jacks are normally installed so that the pins are at the top of the cavity. This helps to protect the contacts from water when baseboards are mopped and from dust. When the jack is oriented in this position, looking into the jack with the contact pins at the top, the pins are numbered 1 to 8 from left to right. Some jacks may not have all pin positions populated, but the numbering would still begin with the first position. For instance, the "RJ-11" style jack is a 6-position 4-pin jack. Therefore it has pins 2-3-4-5 and pins 1 & 6 are usually absent.

You should take care not to plug an RJ-11 into an RJ-45 jack. It will work to connect the pairs that are supported in the plug, but the plastic part on both sides will push the outer pins on the jack up, and they may not make good connection when the jack is again used for an RJ-45 plug.

Ethernet uses 8-position 8-pin modular connectors. TIA/EIA specifies two standards for wiring RJ-45 style cables. The T568A color code is "preferred" by TIA/EIA but is not so usual in the USA for business installations.



The TIA/EIA T568B color code cable specification has the same electrical connections, but has the green and orange pairs swapped (as shown). This is also known as the AT&T 258A wiring sequence and has been widely used in the USA. It is used by the Radio Systems StudioHub+ system for analog and AES connections, so we recommend it for all new installations.

Either sequence will work just fine if you have it on both ends. In either case, you have a cable with 4 pairs wired straight through, both ends wired identically.

Interesting note: 100BASE-T with the final X being dropped is oftentimes used as shorthand for 100BASE-TX. The 100BASE-T designation officially refers to both copper and fiber formats at 100Mbps rate, with TX the specific designation for copper. The abbreviation in popular use arises from the fact that the copper formats on either side are called 10BASE-T and 1000BASE-T. And that the -T is supposed to stand for "twisted pair" – except here for some reason. Leave it to standards bodies to be non-standard.

Depending on the cable manufacturer, the color conductor of each pair may or may not have a white stripe. The other half of the pair is usually white with a colored stripe, but sometimes can be only white. Both formats are shown in table form below:

TIA/EIA-568-A T568 Wiring Standard (preferred for Livewire+)

Pin	Function	Color
Shield	Protective Ground	
1	Transmit +	White/Green
2	Transmit -	Green
3	Receive +	White/Orange
4	N/C	Blue
5	N/C	White/Blue
6	Receive -	Orange
7	N/C	White/Brown
8	N/C	Brown

TIA/EIA-568-B T568 Wiring Standard (Optional and acceptable for Livewire+)

Pin	Function	Color
Shield	Protective Ground	
1	Transmit +	White/Orange
2	Transmit -	Orange
3	Receive +	White/Green
4	N/C	Blue
5	N/C	White/Blue
6	Receive -	Green
7	N/C	White/Brown
8	N/C	Brown

Something to watch out for: The old telephone USOC wiring code has the pairs in the wrong place, with the wiring in simple one-pair-after-the-other sequence. You'll have a split-pair if you have this sequence – and a lot of crosstalk and interference problems. You need to be sure that the pairs correspond to Ethernet's requirements.

Why does Ethernet have such a strange wiring sequence? Because the center two pins, 4 & 5, are where telephone voice circuits are wired. The designers of the standard thought that some people would want to use a single cable for voice and data, so they kept Ethernet clear of the telephone pins. There is also this: if a user plugs his PC's network connection into the phone jack, he doesn't blast the network card with ringing voltage.

Even though you have two unused pairs in the standard CAT-5 4-pair cable, you should not share the cable with any other service since 100BASE-TX was not designed to withstand additional signals in the cable. The reason for the extra pairs is that you might want to upgrade to 1000BASE-T or some other yet-to-be-introduced service later.

Finally, on this topic, something really nuts... The overall cabling specifications standard and document from TIA/EIA was called TIA/EIA-568-A Commercial Building Telecommunications Standard. Within this were the T568A and T568B pinout standards. Note the dashes and lack of same. Now there is a new TIA/EIA-568-B overall standard, which has the same two pinout standards within. Couldn't these guys have been a bit less confusing?

Crossover 100BASE-T Ethernet Cable

Sometimes you want to connect two Livewire+ nodes directly together without a switch, such as for initial setup, testing or when you want to make a snake. Or you might want to connect an xNnode directly to a PC for initial configuration or as a sound card. In this case, the Transmit of one device must be connected to the Receive of the other.

For this, you'll need the special crossover cable wired as shown below.

Pin	Color	Pin
1	White/Green	3
2	Green	6
3	White/Orange	1
4	Blue	Not Used
5	White/Blue	Not Used
6	Orange	2
7	White/Brown	Not Used
8	Brown	Not Used

Many modern Ethernet switches have ports that automatically sense the need for a crossover function and configure their ports appropriately. So when you are connecting ports from two switches, you probably will not have to use a crossover cable.

1000BASE-T Gigabit Copper

We use 1000BASE-T to connect studio mixing engines to switches. If your Livewire+ network consists of multiple switches, you will also use it to link switches to each other.

1000BASE-T works with CAT-5e, but again we recommend CAT-6, because 1000BASE-T uses the same RJ-45s as 100BASE-TX, but needs all four pairs. Either the T568A or T568B wiring sequence will work, but you will have to be sure all four pairs are wired through and working. Again here, the advantage of choosing one scheme and using it for everything (e.g. T568B on CAT-6) is obvious. 1000Base-T Signal Designations are as follows:

Pin	Color	Function
1	White/Orange	BI_DA+
2	Orange	BI_DA-
3	White/Green	BI_DB+
4	Blue	BI_DC+
5	White/Blue	BI_DC-
6	Green	BI_DB-
7	White/Brown	BI_DD+
8	Brown	BI_DD-

There are no separate send and receive pairs for 1000BASE-T. Each pair both sends and receives with a hybrid at the ends to separate the two signal directions. Thus, there are effectively four paths each way. The signaling rate for 1000BASE-T is the same as for 100BASE-T – which is why it can be run over the same cable.

Nevertheless, 1000BASE-T is more sensitive to certain performance issues owing to the hybrids and twice the number of signals in a 4-pair cable. That's why CAT-5e or CAT-6 is recommended. And you should always use high-quality factory-made patch cables.

You shouldn't ever need a 1000BASE-T crossover cable, but who knows? Anyway, a universal crossover cable can be made (or better, purchased) that works for both 100 and 1000BASE-T; the following table shows the configuration for a Universal 1000Base-T/100Base-T Crossover Cable.

100Base-T Crossover Cable

Pin	Color	Pin
1	White/Green	3
2	Green	6
3	White/Orange	1
4	Blue	7
5	White/Blue	8
6	Orange	2
7	White/Brown	4
8	Brown	5

Audio connections

xNodes are high-density devices, and they provide many inputs/outputs in a fi-width, 1RU package. To address this, we give you two options: RJ-45 connectors for individual pairs of wires, or pigtailed such as those used in pro-audio applications which provide an entire bundle of 8 pre-wired XLR-M or XLR-F connectors while using just a single DB25 hookup to the xNode.



AES/EBU xNode rear panel. Both RJ and DB audio connections are provided

For RJ-45 connections, we use the pin-outs established by the Radio Systems StudioHub+ wiring system, which has become a de-facto standard. Since we follow this standard, Studio Hub wiring components may be used for the analog and AES part of Livewire+ installations. Radio Systems, and other Livewire+ partners such as Xi Audio, offer an extensive line of single “dongle” and multi-pair harness cables pre-wired to connect to a variety of popular studio gear. There are also balanced-to-unbalanced, AES to S/PDIF, and AES to TOSLINK adapters, headphone amps, etc.

For the DB25-to-XLR fantails, cables such as those made by Hosa are generally regarded as gold-standard, although these are available from other makers, too. They come in lengths from 3 – to 7 m. to suit your application.



Pre-made breakouts from Radio Systems, Hosa, Xi and others help quicken connections.

While unbalanced connections can be used for very short runs with unshielded and shielded cables, balanced connections are essential for anything over a few feet in length. The input stage of any attached analog equipment needs to have good CMRR (Common Mode Rejection Ratio) and high-frequency filtering in order for balanced connections to effectively cancel crosstalk and interference. With 60dB CMRR, Axia Livewire+ xNode inputs are designed to be no trouble in this respect.

The pinouts for the RJ-45 style audio connectors are shown below:

Axia/Radio Systems Standard for Analog and AES wiring on RJ-45

Pin	Function	Color
Shield	Protective Ground	White/Slate & Slate/White*
1	L +	White/Orange
2	L -	Orange
3	R +	White/Green
4	N/C (GND)**	Blue
5	N/C	White/Blue
6	R -	Green
7	N/C (15-)**	White/Brown
8	N/C (15+)**	Brown
*Optional	** Used to power “spoke” devices such as balanced-to-unbalanced converters. xNodes do not supply this voltage, but external supplies can be used if needed.	

Off-the-shelf or homemade RJ-45 cables, together with adapter dongles from Radio Systems, Xi, Hosa and others, connect the nodes to audio equipment.

Installing RJ-45s

It would be possible to build a sophisticated multi-studio facility without ever wiring a single RJ-45 plug – you would use modular patch fields or jacks at each end of the long “horizontal” cable with punch-down 110-style connections. Then factory-made patch cords would be used to get from the switch or Livewire+ node to the patch jack. And this might not be a bad idea!

Nevertheless, you’ll probably find yourself installing your own plugs at some point, so here is some advice:

- ◆ If you are making a patch cord, use stranded-conductor cable. Solid is likely to break after some time being plugged and unplugged. However solid cable should be used for backbone wiring as it has less loss.
- ◆ Be sure you are using plugs designed for the cable type you are using. Plugs for solid and stranded wires are not the same.
- ◆ Plugs from different manufacturers may have slightly different forms. Be sure your crimp tool correctly fits. In particular, the crimper made by AMP will only work with AMP plugs. Buy a high-quality crimping tool to help prevent any problems.
- ◆ The outer jacket should be cut back to about 12 mm (.5 inch) of the wire tips. Check to be sure there are no nicks in the wires’ insulation where you cut the jacket (an appropriate tool can be purchased to permit you to do so rapidly without fear of damaging the inner insulation).
- ◆ Slide all of the conductors all the way into the connector so that they come to a stop at the inside front of the connector shell. Check by looking through the connector front that all the wires are in correct position.
- ◆ After crimping, check that the cable strain relief block is properly clamping the outer cable jacket.
- ◆ When checking the cable either with a tester or a real device, wiggle the cable around near the plug to be sure that connector is working reliably with stress.



If it’s your first time crimping your own, you’ll probably need a couple of tries to get it right, but after some experience it will start getting pretty easy. Certain RJ connectors include a small carrier that the wires can be fed into first, and then slid into the connector itself. These are recommended as they speed installation and improved accuracy.

5

Designing & Building Your Livewire+ Ethernet System

As with analog audio installations, Livewire+ set-ups range from the very simple to complex facility-wide installations with hundreds of ports. This section is aimed primarily at those who will be building large systems.

Cabling

Ethernet is balanced and transformer coupled, so has very good resistance to interference and has no problem with ground loops. However, frequencies ranging to tens of megahertz are being used, so care must nevertheless be taken.

In the bad old days, wiring was specific to the task – and often to the vendor. Each telephone, network, and audio had its own cable type and wiring protocols. The idea at standards bodies like the Telecommunications Industry Association (TIA) and the Electronic Industries Association (EIA) in the USA is to define classes or categories of cables and accessories that can be used for all applications specified for that class. With this, you have a vendor-independent way to wire buildings and facilities so that services from many vendors can be supported over time without replacing cabling and connectors. The name for this concept is structured wiring.

The long cables that go from equipment rooms to node locations are called horizontal cables. They usually terminate in RJ-45s, either in patchfields or on wall jacks. Patchcords with RJ-45s at each end complete the system, connecting the nodes and central equipment to the jacks.

(Charles Spurgeon in *Ethernet: The Definitive Guide* says that you should consider wiring to be the essential skeleton for your network installation. He goes on to point out that network cabling skeletons are often hidden in the time-honored place for skeletons: a closet. Rim shot.)

Twisted-pair Cable Categories

Cable categories are key to the structured wiring concept. The cabling specifications for the various categories are in the TIA/EIA-568-A (and B) Commercial Building Telecommunications Cabling Standard. The following categories are defined:

- ◆ **Category 3:** These are used only for telephone and Ethernet 10BASE-T, so are not useful for Livewire+ installations.
- ◆ **Category 5:** This designation applies to 100 Ohm unshielded twisted pair cables and associated connecting hardware whose transmission characteristics are specified up to 100MHz. CAT-5 cables are today's most common because they support Ethernet 100BASE-TX.
- ◆ **Category 5e:** This is enhanced Category 5 cable. The main application is for Gigabit 1000BASE-T. While CAT-5 is acceptable for 1000BASE-T, CAT-5e is officially preferred.

- ◆ **Category 6:** While not strictly necessary except for 1000BASE-T links, we recommend CAT-6 for all new Livewire+ installations. CAT-6 provides significantly higher performance than that of CAT-5e. The main difference is that this cable has a plastic pair separator inside that holds the wires in correct relation to each other. This is what makes CAT-6 larger in diameter than CAT-5 cables.

Belden has a CAT-6 cable called Mediatwist that looks very interesting. This cable has a half-moon shape and the pairs are tightly held in molded channels. This product also has the two wires in each pair glued together so that the twist characteristic is fixed and stable regardless of manufacturing tolerances, cable flexing, etc.

The most significant difference between cables from each category is the number of twists per foot and the tightness with which the twists and the spacing of the pairs to each other are controlled. The wire pairs in a voice-grade Category 3 cable usually have two twists per foot, and you may not even notice the twists unless you peel back quite a lot of the outer insulation. Category 5 is tightly twisted, something like 20 per foot. This results in superior crosstalk performance at higher frequencies.

Another characteristic of twisted-pair cables is the type of insulation used on the wires and the cable jacket. “Plenum rated” cables are more stable with changing temperatures due to their using Teflon rather than PVC insulation. Plenum rated cables are required in air handling spaces in order to meet fire regulations. Teflon produces less smoke and heat in the case of a fire than PVC and does not support the spread of flames.

Special Care for Ethernet Audio

“Normal” data over Ethernet is usually TCP/IP protocol. As discussed later, TCP has a re-transmission mechanism that detects errors and fixes them by requesting and obtaining replacement packets when one has been received with a defect. This mechanism is not used for audio – it can’t be when you need low delay and multiple receivers. So it could be possible that a network could apparently be OK with computer data, yet exhibit errors with audio because TCP is covering-up underlying problems.

A particular concern is to prevent impedance reflections at cable termination points and to not disturb too much the position of the wires inside the cable. Here are some specific recommendations:

- ◆ Use the minimum number of terminations and patches that will support your application.
- ◆ Use patch cables, connectors, and other accessories rated at the same or higher category level as the cable you are using. Generally, your best bet is to buy pre-made patch cables to both save money and time as well as assure reliability.
- ◆ Keep a wire pair’s twist intact to as close as possible to any termination point. For Category 5, this should be to within 1.3 cm (.5 inch).
- ◆ Maintain the required minimum bending radius. For a 4-pair 0.5 cm (.2 inches) diameter cable, the minimum bend radius is 4 times the diameter, or about 2 cm (.8 inches).
- ◆ Minimize jacket twisting and compression. Install cable ties loosely and use Velcro fasteners that leave a little space for the cable bundle to move around.
- ◆ Do not staple the cable to backboards. If you tightly compress the jacket, you will disturb the twists inside and the relationship of one pair to another, which could cause crosstalk.
- ◆ Do not overfill conduits.
- ◆ Avoid stretching the cable. The official recommendation is to use less than 25 lbs. pulling pressure.

- ◆ Avoid close proximity to power cables and equipment that generate significant magnetic fields. The official recommendation is minimum 6.4 cm (2.5 inches) from power cables when the Cat 5 is either inside a conduit or shielded. Care should be taken also with fluorescent lighting fixtures, motors and transformers.
- ◆ The pins on RJ-45 plugs are gold plated. But not all connectors are. For maximum reliability, use connectors with 50m gold plating.

To Shield or Not to Shield

Unless you are in a high RF environment or you intend to run your network cables close to audio cables with equipment that has poor balancing on the inputs, you should be able to use unshielded twisted pair for your Ethernet connections. If you decide to shield, the usual procedure to attach it only at one end applies in order to prevent ground loops.

Unbalanced Connections



Livewire+ nodes have very good common mode rejection. This, coupled with the highly twisted CAT-5 or -6 cable works extremely well in the balanced pro-audio environment. Unbalanced interconnections are problematic however and should be avoided for the usual reasons. If you need to interconnect a Livewire+ node to unbalanced gear we strongly recommend that you use a balanced to unbalanced buffer amplifier or transformer located as close as practical to the unbalanced equipment. There are a number of commercial off-the shelf options to accomplish


this. In particular the Radio Systems StudioHub+ Matchjack series offer plug and play compatibility between the RJ-45 balanced and consumer unbalanced worlds.

More than Four Pairs in a Cable

Back in the 10BASE-T days, it was usual to have phone-type 25-pair cables carrying data signals. But the standards for Cat 5 and better call for individual cables for each connection due to the possibility of multiple disturber near end crosstalk – or many signals adding up to create combined crosstalk at too high a level.

On the other hand, Belden has some papers on their website proposing that their finest cable, Mediatwist, would support even 100BASE-T and analog audio inside a shared sheath. Nevertheless, they offer the cable in only 4-pair versions at this time.


Patch Panels



Patch panels come in versions for rack or wall mount and with varying numbers of jacks. CAT-5/5e cables are punched down at the rear into 110-style insulation displacement connectors using a tool very similar to the one that is used with traditional “66 blocks.”

Wall Jacks

Again 110-style IDC connectors terminate the cable. Then these wired-up “Keystone” RJ-45 jacks are pushed into a hole in the wall plate to complete the job. The diagram presented here shows the simple steps involved in terminating these.



CAT-6 Jacks

CAT-6 cables and their accessories need more care to maintain the twists as close as possible to the end. At left is a high-end CAT-6 jack assembly ready for installation into either a rack-mounted patch-field or a wall jack. This is a shielded version, so the shell is made from metal to maintain the shield all the way to the edge of the jack.

Next, you can see the components that make up the jack disassembled. This is now the non-shielded version, so the shell is plastic, the blue piece in the photo.

At the lower left is a closer look at the part that holds the wires. Assembling one of these can be done in a minute or two. First the wires are put into the slots and the ends are trimmed. Then this piece and the front part of the jack are pushed together. The shell is then placed over these pieces and pushed over them, which draws the wires into the insulation displacement forks and locks everything together.

Architecture Options

There are a lot of ways to build a Livewire+ network. For many people a simple one-switch layout will be perfectly sufficient. Others will want to build sophisticated networks to support multiple studios and perhaps hundreds of audio channels. Fortunately, Ethernet scales easily – and therefore so does your Livewire+ installation.

Here are some examples and ideas to get you started.

Simple One-Switch Network

Common 1U switches can have as many as 48 ports. That is a lot of audio! Most modern switchgear has also migrated to all-Gigabit with 2- or 4-port SFP (Small Form-factor Pluggable) ports for fiber or copper connections. Here's a setup that supports an on-air studio and a production studio.



There is the microphone version of the xNode in the on-air studio and the 4x4 analog line version in the rack. The production studio connects with the xSelector node, which has one send channel and a selectable receive channel.

The Element power supply (not shown in the diagram) includes plenty of GPIOs for starting CD players, lighting on-air lamps, remote mic on-off, etc.

The StudioEngine connects with a 1000BASE-T copper link to one of the two 1000BASE-T SFP ports with a standard copper transceiver module.

The delivery PC connects directly to the audio network with the Axia IP-Audio Driver software. Control for it may be directly over the network or could be with a hardware parallel connection. Servers and additional PCs can be connected to the switch for file storage and delivery systems.

Peripherals such as codecs, telephone systems, and satellite receivers may be connected into the network wherever it is convenient. In the diagram, the Z/IP ONE IP codec and Telos phone system are shown attached directly to the switch; many Telos Alliance devices, as well as those of other manufacturers, can be attached directly with Livewire+ connection ports. Equipment without Livewire+ can be attached directly to an xNode, as is the satellite receiver in the drawing.


You could expand this to two Elements and Engines to support two studios since the switch has two 1000BASE-T ports. An all-Gigabit 100/1000BASE-T switch can support nearly as many studios as you want.



In another variation as seen here, the standalone switch can be replaced entirely with a PowerStation mixing engine, which also reduces separate I/O device count by aggregating audio and logic IO with the network switch into a combined mixing engine environment. The switch contained within PowerStation and QOR mixing engines supports daisy-chaining of up to 4 devices without the need for a core switch, so smallish studio networks can dispense with additional switches nearly altogether.

Daisy Chained Multiple-Switch Network

While one switch can support multiple studios, you would have a single point of failure. Here’s another approach that gives each studio its own switch. This example uses three switches, one for each studio group. This layout style could easily be expanded to any number of switches and studios.



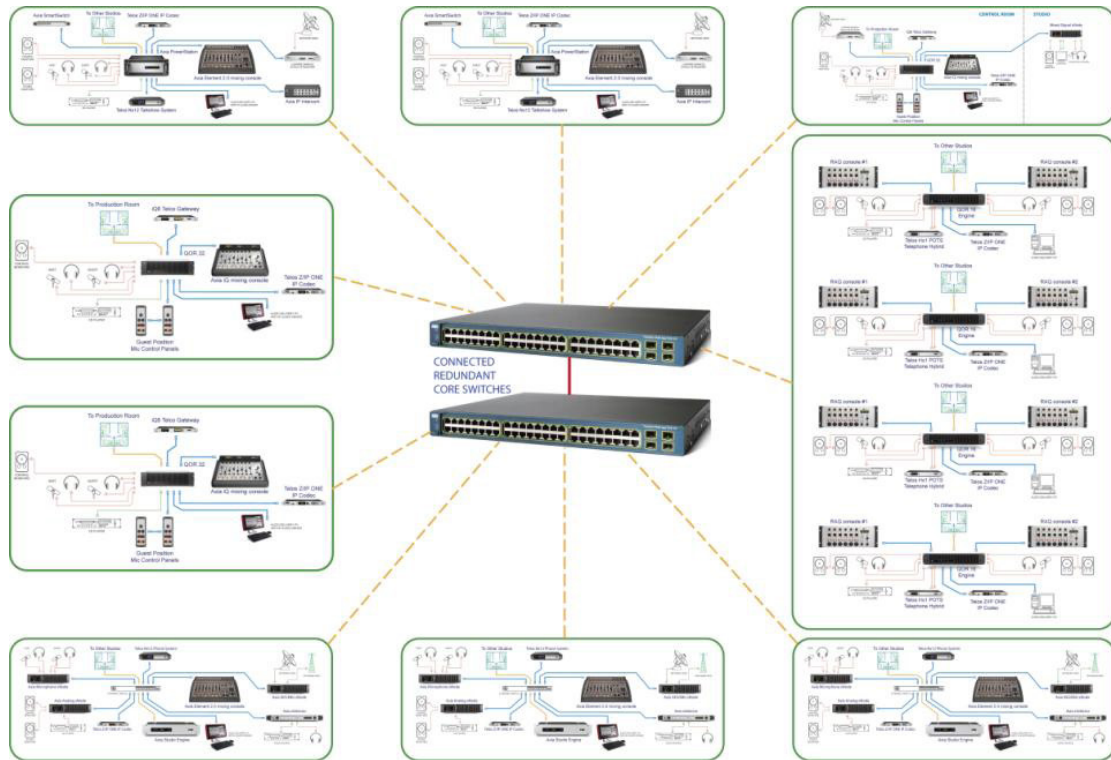
The switches are connected together so that audio sources are shared. A 1000BASE-T link between the switches allows hundreds of audio channels to flow from one group to another. With more than two switches you could have a “circular backbone” with redundant Spanning Tree links (described below) between the switches.

Peripherals that are used in common could be plugged to any of the studio switches, or there could be a separate switch to pick up such feeds. In the example here, one studio takes in a satellite feed for sharing among studios; another hosts a Telos VX multi-studio phone system base unit controlled in other studios by VSet phones. An Omnia VOCO 8 multi-mic microphone processor is likewise shared.

Hierarchical Multiple-Switch Network

This is a layout that could support a very large facility. Dual-redundant Gigabit switches are at the center, and 100/1000 switches are used at the edge with one for each studio or logical group.

Note that the studios themselves are using a mix of standalone switches and Axia PowerStation and QOR mixing



Note that the studios themselves are using a mix of standalone switches and Axia PowerStation and QOR mixing engines with built-in switches, each connected back to the central switching core. In this way, with each studio coupled to an individual studio switch, there is no single point of failure for any studio since each operates as a separate, independent facility.

Gigabit links are used between the edge switches and the center. These could be copper or fiber with a suitable switch.

The physical location of the switches is a matter of taste and trade-off. Putting the edge switches near the studios saves cable runs, but locating all the gear in a central room simplifies engineering activities. So this is a quite reasonable-cost option that provides a lot of power, flexibility, and expandability. Dozens of studios and thousands of audio channels are possible.

Options for Redundancy

Ethernet switching has a built-in scheme for redundancy, called Spanning Tree and standardized as 802.1D. A newer variant is called Fast Spanning Tree. Switches with spanning-tree enabled exchange information with each other about the topology of the overall network. You can have redundant backup links that are automatically activated in the case that a main link has failed. Depending on the switch and layout, it could take as little as a second or as much as a half-minute for a redundant link to be connected.

Link aggregation (sometimes called port trunking) is another method. With Spanning Tree, even if you have two links between two switches, only one of them at a time will be active. But, it's often better to have both active simultaneously because you get twice the bandwidth during normal operation and instantaneous backup should one fail. The link aggregation standard is 802.3ad. To use it, you usually have to specifically enable it on your switch. Incidentally, this is supported on some PC network interface cards intended for servers, so it's not only for switch-to-switch links.

Most Ethernet switches offer a redundant power supply option.

We've been talking here about automatic on-line redundancy, but there is also manual swap-out as a reasonable option. Because RJ-45s are so easy to unplug and re-plug and because switches and other Livewire+ components are much cheaper than traditional alternatives, you can have spare units on the shelf for fast substitution.

Fiber

Fiber optic links can extend the range of Ethernet. Because they are not subject to crosstalk and magnetic interference, they also can solve problems that might crop up in difficult locations with copper cables.

External media converters can be very simply plugged to Livewire+ nodes and switch 100BASE-T ports to convert copper connections to fiber.

When choosing fiber connectors, you probably expect something with "multi" in the name to have more capability than the same thing designated "single". But this is not the case with fiber optics: single-mode cables are better and more expensive than multi-mode. These names refer to how light is contained within the fiber. Single-mode strands are smaller and more carefully control the light so that it doesn't bounce around so much inside, thus are more efficient and permit longer ranges.



This Gigabit unit from Allied Telesyn uses 1000Mbps SC single-mode fiber for up to 75km range. Units supporting ST multimode fiber, for 100BASE-T links, are also available for runs up to 2km. Many are available with PoE capability to power nearby peripherals.

Modern Ethernet switches often have the option to plug a media converter directly into a special socket so that fiber

may easily be connected from switch to switch. This is useful to make high capacity backbone links without any external boxes.



Here is a typical case. The Axia xSwitch shown above at left has two “uplink” ports for use with 1000BASE-T SFP (Small Form-factor Pluggable) fiber or copper transceiver modules. The device to the right of it is a typical modern media adapter in the “SFP/mini-GBIC” size – about the same in width and height as an RJ-45 jack. They come in different flavors, for 1000BASE-SX, 1000BASE-LX, etc. Generally, SX cables have a range to 500 meters, LX to 5km, and LH to 70km.

Radio Links

There are Ethernet IP radios with surprisingly high bandwidth – and at surprisingly low cost. Not all units are capable of achieving true Ethernet performance in terms of error rates, so some caution is in order. Most of these operate in the unlicensed ISM bands, but with modern spread-spectrum technology and elevated directional antennas, interference doesn’t look to present much problem. Licensed radios following the IEEE 802.16 “Wimax” standard are easily available.



Bitrates range to 150 Mbps and distance to 25 miles depending on power level, antenna, and terrain.

For studio-to-transmitter link, remote pick-up, and studio-to-studio applications, these offer multiple channels of uncompressed audio, two-way transmission, and the ability to multiplex VoIP telephone, remote control, and general data. When audio and general data are mixed, the Ethernet switch provides the prioritization function. As with all Livewire+ elements, you can check them with a web browser on a network-attached PC.

We have studied many of these Ethernet IP radios, testing for Livewire+ compatibility and general performance. You should consider these like the Ethernet switch – please let us advise you on the best choice and help with your application.

Designing For Security

You will have 100% security if you keep the Livewire+ system completely isolated from any other network, local or wide area. Those very concerned with protecting the studio system may well want to take this approach.

But there are advantages to sharing with or linking to an office network. You can configure and monitor the system from any connected PC and audio can be monitored on any desktop with access. In this case, separate switches or VLANs (described later) can be used to provide isolation. An IP router passes only the correct packets from one to the other and thus provides a firewall function.

The next step up in connectivity would be to have a network linking co-owned or otherwise affiliated stations. In this case, a network engineer is probably in the picture and he can take the necessary steps to protect your audio.

Connection to the public Internet brings the advantage that you can monitor and configure from a remote site, but you now have much risk from unwanted intruders, viruses, etc. A qualified network engineer should be consulted to be sure you have an appropriate firewall and other protections in place.

In Livewire+ nodes, web and Telnet access are password protected to provide some measure of security. But we do not use exotic techniques like SSL (Secure Sockets Layer), so please understand that our devices were not designed to be exposed to the public Internet without external protection.

6

The Ethernet Switch

This is what makes it all possible. Here are some details on requirements for a capable Livewire+ switch.

Livewire+ streams are “routed” at Layer 3 of the Ethernet model. We recommend a “Layer 3 switch” that includes the required IGMP Querier and snooping functions. Switches range from very simple to amazingly elaborate.

Livewire+ Ethernet Switch Requirements

- ◆ Sufficient backplane bandwidth. It is required to be fully “non-blocking” to handle all ports at full capacity.
- ◆ Sufficient frame forwarding rate. Livewire+ Livestreams, as well as AES67 streams, have small packets at a fast rate. The switch needs to handle this.
- ◆ Correct handling of IEEE 802.1p/Q frame prioritization. Livewire+ audio frames must be given priority without too much delay or jitter. The IEEE standard specifies 8 levels of priority, but few switches support all the levels. Many support only 2 or 4, lumping some of the incoming levels together. We recommend 4 as the minimum for a Livewire+ system.
- ◆ Support for multicast, with sufficient filter entry capacity to cover the total number of audio streams you need. This latter is important, because when the filter capacity is exhausted, switches forward multicast packets to all ports, subscribed or not. This would cause serious problems. You will probably want 256 minimum.
- ◆ IGMP control for multicast. Traffic must be under IGMP control – strictly no flooding of ports with multicasts under any circumstances.
- ◆ Support for both port-based and tagged-frame-based VLAN. This latter is the IEEE 802.1Q standard and is what allows the switch to determine priority on a frame-by-frame basis. Port-based VLAN can also be useful: it lets you “hardwire” a particular port for a single VLAN, useful to be 100% sure an office PC can’t get onto the Livewire+ audio VLAN.
- ◆ If you will use a separate VLAN for Livewire+, the switch needs to have an “IGMP querier” on each one, which also means that you can assign an individual IP number to each VLAN. This is an advanced capability and its absence disqualifies many switches.
- ◆ Management. This is how you get remote monitoring.

The practical bottom line is that you should use a switch that has been selected and tested by the Telos Alliance. When we check a switch, we use a laboratory setup that lets us send frames on a number of ports at a high rate, while switching channels on/off with IGMP, etc. We have a lot of experience with different switches and know what to look for. Using a recommended switch will also help you when you need customer support because we will be familiar with it. You are encouraged to contact Telos Alliance Customer Support or visit www.TelosAlliance.com/switches for our full list of currently supported switches.

Some Switches We Like

There are new switches introduced everyday it seems, with ever increasing performance and falling prices. Please check our website for our latest recommendations.

Cisco Switches



Having said that, the Cisco Catalyst 6500, 4900, 4500, 3560X, and 2960 series switches are some of the units that are fully compatible and qualified for use with Axia. There are literally hundreds of models that can be used with your Axia networked audio system. 24 10/100BASE-T ports, 48 10/100BASE-T ports, with or without 1000Mbps SFP (copper/fiber) transceiver ports — connect them together and have virtually unlimited flexibility and redundancy. All of these units include built-in simple router and IGMP Queriers on every VLAN. They also come in a powered-port version supplying IEEE 802.1af PoE (Power over Ethernet) voltage that can be used with VoIP phones such as the Telos VSet, or even to power Axia xNodes (which are capable of running on either AC mains or PoE, cutting down on power cabling).

Axia xSwitch



The Axia xSwitch is a half-rack 1RU unit with the same form factor as our xNodes. Because xSwitch is designed and built by Axia, it is optimized for Livewire+ IP-Audio applications, comes pre-configured and needs virtually no setup. Fast setup requires only IP address assignment, accomplished in much the same manner as that of an xNode, via the front-panel display or Web management interface. It is an advanced switch designed specifically for AoIP network deployment:

- ◆ Eight 10/100MBit Ethernet ports — 4 with PoE .
- ◆ Two Gigabit for trunking, both with RJ-45 (copper) and SFP (fiber) connections.
- ◆ Supports redundant copper/SFP Gigabit connections with auto-switching.
- ◆ Has 2,000 Multicast groups and 2,000 ARP table entries (8x more than other small-form Ethernet switches) .

xSwitch is an excellent choice for either main switches in small studio installations, or edge switches in larger facilities.

Switch Configuration

Most switches offer three connection options: an RS-232 console port, Telnet over Ethernet, and web over Ethernet. For Axia-supported switches, we offer a configuration “cheat-sheet” that gives you the basics. We also will be happy to pre-configure your switch and test it at the Telos Alliance before shipping it to you. There are two different software images that Cisco switches can be configured with. They are available with either the standard multilayer software image (SMI) which includes advanced QoS, rate-limiting, ACLs, and basic or the enhanced multilayer software image (EMI) which includes a richer set of enterprise-class features and advanced hardware-based IP unicast and IP Multicast routing as well as policy-based routing (PBR). We do not expect or require anyone to be Cisco Certified Network Engineers so we provide you everything you need to know about configuring your network switch based on the software image of your unit. Configuration documents can be found on our website.

7

Testing, 1-2-3...

There are tens of thousands of people installing Ethernet networks every day, and many millions of working installations. So there are a lot of tools to help you. Livewire+ equipment have a lot of diagnostic functions built-in as well.

General Ethernet Troubleshooting

Ethernet is a mature technology, with years of proven reliable service. You are not very likely to see problems in the fundamental technology if you follow the network wiring and layout recommendations.

Prevention

The best way to avoid downtime is to build the network well in the first place. Use high-grade cables, good quality factory-made patch cords, etc. And be careful with the punchdown and plug installation.

More on the topic of patch cords: If you really must make your own, they should be built with stranded wire cables. Solid conductors are likely to crack when flexed a lot, usually right at the RJ-45 plug. From this you can get intermittents and bit errors. Also, as mentioned in the cabling section, be sure you have the right plug for the cable you are using. An RJ-45 plug designed for stranded wire will cut through a solid conductor.

But you know all that. So let's get on to troubleshooting, when despite all due care something still goes wrong.

The Basics

Link Test

A layer 2 test, this checks the connection between the switch and a designated network device on the same LAN. During the link test, IEEE 802.2 test packets are sent to the designated network device in the same VLAN or broadcast domain. The remote device must be able to respond with an 802.2 Test Response Packet. Most switches support this test via a web or command line interface.

Ping

A layer 3 test, a simple and effective way to check basic "reachability" of an IP-enabled device. Ping sends a test packet to a device and waits for an echo response. A Windows PC can do this within the command prompt window. Just enter ping x, where x is the IP number or the domain name (if a DNS server is available) and see the result. If you get the echo, the basic connectivity (including Layers 1, 2, and 3) is OK. Most switches and almost all IP-enabled devices support this test.

Switch Diagnostics

Ethernet switches have many diagnostic tools, ranging from front panel LEDs to sophisticated software monitoring functions. See the switch manual and software description for details for your unit.

Simple Network Management Protocol (SNMP) and Remote MONitoring (RMON) are part of the TCP/IP internet suite. (RMON is built on SNMP so they are closely related.) They offer a way to probe and monitor network equipment operation in a vendor-independent way. For example, an Ethernet port has a standard way of communicating its status that is supposed to be used by all products with these ports.

Almost all sophisticated Ethernet switches offer these, and they are useful tools to monitor traffic, check operation, etc. You can do a lot of this with web and Telnet based communication but SNMP usually offers a deeper look.

You will encounter the acronym MIB, for Management Information Base. This is how information is organized within SNMP. Cisco MIB files are available as free downloads from their website.

To use SNMP and RMON, you will need a software application that presents the information. There are many freeware, standalone programs that can be used for this purpose; Cisco themselves offer an online tool, SNMP Object Navigator, available at <http://tools.cisco.com/Support/SNMP/do/BrowseOID.do>.

A full discussion would be too much for this document, but there is a lot of info that comes with Ethernet switches, and a lot more in bookstores and on the web.

Some Things to Check

- ◆ Switch configuration must be correct. IGMP must be switched on, VLAN parameters set if you are using them, etc. In our experience to date, this is the most common cause of problems. (With the exception of cables, of course.)
- ◆ Ethernet links can be 10, 100, or 1000 Mbps, and full or half-duplex. We always want the maximum rate and full duplex. You can configure the Ethernet ports on some devices for specific modes – but you should not do this. The Auto mode is the correct setting, which will cause the device and node to automatically negotiate to the appropriate condition. If you manually set the mode to full-duplex, the switch – in compliance with a flawed IEEE standard – will set itself to half-duplex (!), leading to many problems. Livewire+ hardware nodes are always set to the auto mode, so this problem will not arise with them, but with other equipment such as PCs.
- ◆ If you want and have multiple redundant links using port trunking or spanning-tree, you have to set up the switch to support them. Taking the default will usually not work.
- ◆ The “activity” LEDs (usually amber or green) on many network cards and switches will be on continuously when any Livewire+ audio streams are present on the link. That is because the logic that drives the LED extends the on time so that you can see it with normal traffic. Livewire+ packets are traversing the network at such a fast rate that the LED never has a chance to turn off.
- ◆ Most Axia Livewire+ hardware nodes have status displays. They provide useful information and should be checked. This is covered later in this section.

Cable Testers

“It’s the cable – it’s always the **@@ cable!” said my first boss. About half the time, he was right. That percentage is probably a bit higher in Ethernet systems. Indeed, a number of surveys have put the “network medium” to blame 70-80% of the time. This being the cables, connectors, and hardware components that make up the signal-carrying portion of the installation.

Wiggling and unplug-plug operations are legitimate and effective troubleshooting methods. But there are plenty of cable testers to help you perform more elaborate checks. These range from simple conductivity testers to sophisticated units that test cables for adherence to the TIA/EIA standards, detect breaks with a Time Domain Reflectometer, and more. Contact info for the main manufacturers of these are listed in the Resources section.



The testers shown here represent something of the range available.

The Fluke DSX CableAnalyzer™ Series can certify that your cable meets the appropriate category requirements with regard to crosstalk, attenuation, etc. and perform a number of sophisticated tests. An adapter can be changed to allow the unit to work with both copper and fiber cable types.



The Fluke MicroScanner™ Cable Verifier is a much simpler and cheaper variant from Fluke that checks for conductivity and correct wiring. It can also tell you the distance to a break with a TDM function and can do tone trace with an optional remote unit.

Finally, the Byte Brothers LAN Tester TVR10/100/1000™ 2-piece set is a basic wiring tester and tone line-finder.

Sniffers



These are software applications that run on PCs and can listen-in on the packets flowing on an Ethernet link. Usually used in conjunction with an Ethernet switch’s port-mirroring function. This lets a designated monitoring port mirror that traffic on any other port you select. Livewire+ audio packets are small in length and very frequent compared to general data traffic so are quite challenging for a sniffer. To be useful, you will need a good one and a fast computer to run it. Very useful, but expensive. Perhaps best borrowed from your company’s network guys’ kit.

Diagnosing Problems using Livewire+ Components

All Livewire+ components have built-in diagnostic tools. For example, audio nodes have a loop-back testing procedure that measures audio noise and distortion. The web interface lets you check a number of internal values.

The xSelector is a useful device for displaying available audio streams and listening to them. It has one channel of send, so is useful as an audio source injector as well.

The Livewire+ IP-Audio Driver has a diagnostic window that tells you a number of things about the system clock and audio streams.

xNode Indicator Displays

xNode displays have indicators of the status of Livewire+ and Ethernet connections, as well as system synchronization as follows:

LINK and ACTIVITY LEDs on each Ethernet and Livewire+ port. A solid amber LINK light indicates the presence of a live Ethernet link to another Ethernet 100 Base-T device. A flashing green ACTIVITY LED indicates that the connected Ethernet segment has Livewire+ traffic present. If the LINK LED is illuminated, and the ACTIVITY LED fails to illuminate, it's likely that the Ethernet switch has not been programmed to pass such traffic through to the port to which this node is connected.



SYNC & MASTER indications in the upper right corner of all xNode front-panel readouts. The Livewire+ system employs a sophisticated master/slave clocking system over the Ethernet network. By default any device may become the clock master, however this can be changed if desired. The system has the ability to automatically change to a different clock master should the current master become disconnected, or otherwise inoperable. This happens transparently without any audio glitches.

- ◆ SYNC indicates the receipt of clock information from another (Master) Livewire+ node, and that the local xNode's PLL is locked.
- ◆ MASTER indicates that an xNode is acting as the master clock source for the Livewire+ network.
- ◆ If neither SYNC nor MASTER are displayed, something is not correct.

8

Network Engineering For Audio Engineers

You don't need to know most of what's in this section to use Livewire+. Just as a beginner can plug analog XLRs successfully together without knowing anything about op-amps, you can connect and use Livewire+ without knowing details about packets. But just as fixing tricky problems in the analog world calls for higher-level understanding, so does an awareness of the internal technology of Livewire+ help you to solve problems and build complex systems.

This section introduces basic concepts – enough for you to get a feel for how data networks work and to understand the lingo so you are ready to ask intelligent questions of network guys and vendors. It also explains a lot of points specific to Livewire+.

Livewire+ is built upon standard components, so if you understand data networking generally, you'll be ready for the specifics of Livewire+ audio networking. Network engineering is a rich topic, abounding with information and nuance, and in constant flux. Fortunately, Livewire+ uses only a small subset that is easy to learn and understand. That is mainly because most of the complexity comes with IP routing and wide-area networks such as the Internet – and we don't use much of that, staying only with the much simpler Ethernet LAN level. Even if you don't know anything yet, you'll get pretty much what you need in the next few pages. If you want to know more, bookstores have shelves loaded with networking advice and information. We offer a few starting points in the Resources section.

As always, Telos Alliance support is at your side to help with any specific practical issues that may come your way.

If you are developing for Livewire+, this will offer only a brief introduction, and you'll want to know more. Please contact us for any of your needs, such as software API documents. We welcome partners!

Ethernet / IP Networks

Layering Model



You need to know layers to know networks. The notion of layers and the open systems they support are central to network engineering. Because layering is a key to enabling multiple vendors for each function, this design has also been a major factor in the growth and operation of the Internet. It's also one of the keys to Livewire+, allowing us to build our professional audio transport application on existing standard lower layers.

For many years, the Open Systems International (OSI) model was the reference paradigm for data networking. For example, the ISDN D-channel communication between nodes and the telephone network is loosely based on this model:

Layer	Function
7	Application
6	Presentation
5	Session
4	Transport
3	Network
2	Data Link
1	Physical

But this proved to be too complex for most practical applications, and an architecture has evolved that is simpler than the OSI model. The next chart shows how that simpler model applies to the IP-over-Ethernet combination we are using.

Layer	Function	Example
5	Application	HTTP (Web), POP (mail), etc.
4	Transport	UDP/TCP/RTP
3	IP Routing	IP
2	Switching	Ethernet Addressing
1	Interface	Ethernet Physical

Layer 1, Physical Interface: This layer is responsible for hardware connectivity, which is provided by Ethernet.

Layer 2, Ethernet and Switching: This layer is Ethernet’s end station addressing and everything related to it. An Ethernet switch is working at Layer 2 because it forwards packets based on Ethernet Media Access Control (MAC) addresses which are unique ID numbers assigned by the Ethernet-capable equipment manufacturer.

Layer 2 does not ordinarily extend beyond the corporate boundary. To connect to the Internet requires a router. In other words, scaling a Layer 2 network means adding Layer 3 capabilities.

Officially, the transmission units comprising header and data are called frames at this layer. At Layer 3, the correct designation is packets. But, since Ethernet frames are almost always carrying IP packets, the word used to describe the combination most often depends upon the context or the author’s preference. Unless we are referring to layer 2 functions, we usually use “packets” because Livewire+ audio has the IP header – and because “packets” has become the usual way to describe this sort of thing generally.

Layer 3, IP Routing: In addition to Ethernet addresses, each IP packet on a LAN also contains source and destination IP addresses. These were intended to be used by routers to forward packets along the most efficient route and link LANs of different types. When the Internet was invented, there were dozens of LAN technologies in use and this was an important capability. Now, IP addressing is used both within LANs as a way to access servers from clients, etc, and to connect to Internet resources offsite.

IP in itself is not a particularly complex protocol, but there are numerous capabilities supplied by the other components of the IP suite. The Domain Name System (DNS) removes the burden of remembering IP addresses by associating them with real names. The Dynamic Host Configuration Protocol (DHCP) eases the administration of IP. Routing protocols such as Open Shortest Path First (OSPF), Routing Information Protocol (RIP), and Border Gateway Protocol (BGP) provide information for Layer 3 devices to direct data traffic to the intended destination.

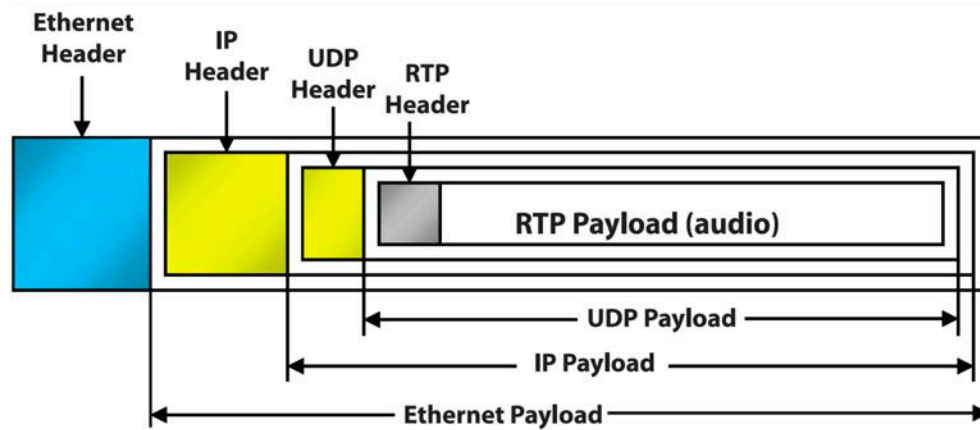
Layer 4, Transport: This layer is the communication path between user applications and the network infrastructure and defines the method of communicating. Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) are well-known examples of elements at the transport layer. TCP is a “connection-oriented” protocol, requiring the establishment of parameters for transmission prior to the exchange of data and providing error recovery and rate control services. UDP leaves these functions to the application.

Layer 5, **Application**: Web browsers, audio editors, and email clients, for example. And our Livewire+ audio.


Applications developers decide on the type of Layer 4 transport necessary. For example, database or Web access require error-free access and use TCP, while live streaming media use Real-Time Protocol layered on top of UDP/IP.

Making Packets

Livewire+ Standard Streams use all of the recommended Internet protocols and are constructed in the usual layered fashion. Here is one representation of the packet structure:



In the next graphic, below, you can see this structure in more detail. This is the way network engineers usually visualize a packet. It's not important to know what each of the fields means; the idea is for you to see how a packet is constructed generally. Each of the horizontal bars are 4 bytes. At each layer, devices are operating only with the information contained within the associated header. An Ethernet switch only cares about the layer 2 headers and everything else is just payload.



An IP router only “sees” the layer 3 header and doesn’t care about the lower-level transport. Applications don’t care about headers at all – they just deliver their data to the network and expect to get it back at the other end (There are, however, exceptions, such as sophisticated Ethernet switches that can inspect layer 3 headers for some advanced functions).

IP and Ethernet Addresses

As with everything connected to IP/Ethernet networks, Livewire+ devices require both IP addresses and Ethernet MAC (Media Access Control) addresses.


IP Address

IP addresses are four bytes long and are written in “dotted decimal” form, with each byte represented decimally and separated by a period. For example, in the IP address 193.32.216.9, the 193 is the value for the first byte, 32 for the second, etc. Since a byte can hold values from 0 to 255, this is the range for each decimal value. Host IP addresses are assigned to your organization by your internet service provider and parceled out to individual host computers by your network administrator. He may give you this number to be entered manually, or could opt for DHCP (Dynamic Host Configuration Protocol) to let your computer get the address automatically from a pool. Because Livewire+ devices are permanently attached and because it is more convenient to know the IP address attached to a particular node and perhaps assign them in some kind of logical pattern, we do not support DHCP for our hardware nodes. Therefore, you will need to enter an IP address into each node.

In addition to the address, there are a few more numbers to enter into an IP configuration:

Subnet mask

Subnets allow a network to be split into different parts internally but still act like a single network to the outside world. There are three logical parts to any Internet address: the main network address, the subnet address, and the particular device address. The mask marks the dividing point in the address between the subnet part and the device (host) part. What is meant here by “network” and “subnet” depends on your internet provider. A network in this context could mean all of the address space allocated to the provider, and the subnets could delineate the individual customers. Or the network could be all the addresses allocated to a university or major corporation and subnets could divide the address space to correspond to departments. Network addresses are assigned by IANA, the internet names and numbers authority, while subnets may be changed without any official approval.



The mask is written in the same dotted-decimal form as IP addresses. In the example a very large network supporting 64k hosts is divided into 64 subnets, each with 1k hosts. The subnet mask would be 255.255.252.0, which is just another way of writing the binary ones and zeros value shown above.

As a practical matter, you usually just take the number given to you by you network administrator or service provider and enter it.

Gateway address: This is the IP address of the device that passes traffic out of your local network to the internet. This is usually a router.

DNS server address: This is the address of the computer that provides name look-up service, translating text domain names like www.TelosAlliance.com to IP address numbers.

In careful language, devices that attach to the Internet and have IP addresses are called hosts, a name that probably made sense in the early days. (They “host” the IP stack and interface.) And Ethernet-connected devices are officially called stations to keep the radio/ether analogy going. But what do you call something that is both host and station, as almost everything being “Host” doesn’t sound very natural for our audio devices and “station” would be very confusing, indeed. As you’ve noticed, we usually just say “Livewire+ node” in the context of our audio equipment, which should be clear enough. They won’t be nodes, will they? Unless something better comes along, we’ll probably say Livewire+ device. As to “host” and “station” for other devices, we’ll just use connected PC or some variant, thank-you very much.

Ethernet Addresses and Address Resolution Protocol (ARP) : Machines that use IP and are connected to an Ethernet have two addresses, IP and Ethernet MAC. While the IP address is user-determined, the Ethernet MAC address is usually programmed into the network card or interface by the manufacturer.

You will probably never have to deal with them directly, but who knows? Ethernet addresses are 6 bytes long and are written in “dashed hexadecimal” form like this: 5C-66-AB-90-75-B1. (Sometimes colons are used as the separators.) Hex notation is just another way to write binary values. Single digits range from 0 to 9, A, B, C, D, E, F and byte values from 00 to FF. The value FF means all the bits in a byte are 1s and is equivalent to decimal 255. While this notation may seem strange at first sight, it is very useful to programmers, who need to think in powers of two.

There is a unique Ethernet MAC address for each and every network adapter ever made in the world. IEEE handles the allocation among manufacturers and each manufacturer is responsible to ensure that they make no two alike within their assigned range. I (Steve) used to feel bad about all those wasted addresses from obsolete and thrown-away network cards – guess that’s the Protestant USA mid-westerner in me – but supposedly 6 bytes is enough that each of Earth’s grains of sand could have its own address, so not to worry.

There is a need to translate between IP and Ethernet addresses. Consider a server sending data to a machine it knows only by IP address. To communicate, it has to generate an Ethernet frame including the Ethernet destination address corresponding to the desired IP address. To do this, every IP-based device has an ARP module, which takes an IP address as input and delivers the corresponding Ethernet address as output. It maintains a local table with the associations. When it encounters one it doesn’t yet know, it broadcasts an ARP query packet to every device on the LAN and the device that owns the specified IP address responds with its Ethernet address. If there is no owner, the packet is presumably intended for an off-site device and is sent to the gateway address of a router. How does the transmitting device find the router’s Ethernet address?

With ARP, of course. Entering `arp -a` into Windows’ command prompt will give you the current list of IP addresses and associated Ethernet addresses – the ARP table for that machine.

Multicast Addresses: All of the above discussion was only relevant to the usual Unicast situation that is used for web surfing, emails, file transfers, etc. We also use it in Livewire+ for configuration and control, such as when a web browser is connected to a hardware node. But audio is multicast because we want it to be available to multiple destinations. The principle is simple: rather than specifying a specific destination, a special “virtual” multicast address is used that is not assigned to any particular device. Audio nodes can listen-in in a party-line fashion by receiving any packets at this address.

Our audio streams are multicast at both Layer 2 and Layer 3, using the set-aside multicast addresses at each layer. The Livewire+ channel number is automatically translated to the appropriate addresses at both layers internally.

Livewire+ uses the IP address range starting from 239.128.0.0. This choice is based on the assigned numbers from the IANA (Internet Assigned Numbers Authority) allocation of this range for use within organizational and site specific scopes. These addresses are to be used for multicast applications that are not used across the global Internet. Since our application will be used within a single facility on a single switched LAN, this range is appropriate.

Over 8 million unique IP multicast addresses are available with each address mapping to a globally unique Ethernet multicast address.

Even so, IP is relatively stingy with its multicast space. Ethernet has set aside half of all destination addresses for multicast - 140,737,488,355,328 addresses, which should be enough for even the very largest broadcast facility! The designers clearly had big plans for multicast that have not yet been realized.

The distinction is made in the first transmitted bit of the 48-bit address that divides the total available address space in two: a 1 in this position signifies a multicast.

Ethernet Switching

Ethernet switching has caused a revolution in data networking. With switching, each device owns all the bandwidth on its link. No sharing and no collisions. Incoming frames are forwarded only to the nodes that need them.

Despite their amazing power, the invention of switching was more akin to falling off a log than sawing one in two. The switch builds up a table of what addresses are attached to what ports, which it does by merely by examining the source addresses of sent packets. When frames come in, the switch looks into the table, discovers what port owns the destination and forwards the data only to that port. In the rare case that no entry exists for an address, the frames are “flooded” or broadcast to all ports to be sure the intended recipient gets it. If a connection is unplugged or there is no data for a long time, the entry is removed. Pretty simple, eh?

Multicast

The operation described above is for the common Unicast, or point-to-point, communication that you have for typical traffic such as web, email, etc. But Ethernet supports three communication types:

- ◆ **Unicast** means point-to-point. The usual mode for traffic.
- ◆ **Multicast** means that multiple receivers may “tune in” to the transmission from a source so that a selected subset of nodes is served.
- ◆ **Broadcast** means that packets are sent to all receivers, which is quite common on Ethernets. Microsoft file sharing, for example, advertises the PCs on a network this way. ARP uses this to get a query to all machines on the network.

We use multicast for Livewire+ audio streams because we want to emulate distribution amps and audio routers, with multiple receivers being simultaneously able to listen in to a source. The automatic procedure described above does not work for multicasts because they are not associated with a particular output port and node. Fortunately, switches offer a way to control these one-to-many streams. A multicast Ethernet frame has a “virtual” destination address that is just stopped inside the switch if there are no interested receivers. When receivers want to tune-in, they send a message to the switch telling it to turn on the stream to their port.

The switch knows what frames are multicasts because the destination address belongs to the set-aside multicast pool.

Livewire+ uses one Ethernet/IP multicast address for each audio stream. These are derived automatically from the channel numbers you assign. Streams are multicast at both Ethernet and IP layers using the assigned multicast addresses at each.

IGMP (Internet Group Management Protocol): We need some way to tell the switch what streams go to what ports – that is, a way to control multicast switching. IGMP was designed for just this purpose.

IGMP is part of the IP suite and is a Layer 3 function that was designed to communicate with IP routers to control multicasts. But switch manufacturers started to implement “IGMP snooping” on the messages between hosts (computers) and routers as a way to control multicasts at Layer 2. In recent switch implementations of IGMP, this is taken further and a router is not necessary as long as a switch is configured to support IGMP with the “Querier” feature enabled. We want this because there is often no router in the system. Even were there to be one, better to have this capability in the switch as a back-up.

IGMP uses three types of messages to communicate:

- ◆ **Query:** A message sent from the querier (multicast router or switch) asking for a response from each host belonging to the multicast group. If a multicast router supporting IGMP is not present, then the switch must assume this function in order to elicit group membership information from the hosts on the network.
- ◆ **Report (Join):** A message sent by a host to the querier to indicate that the host wants to be or is a member of a given group indicated in the report message.
- ◆ **Leave Group:** A message sent by a host to the querier to indicate that the host has ceased to be a member of a specific multicast group.

An IP multicast packet includes the multicast group (address) to which the packet belongs. When an IGMP client connected to a switch port needs to receive multicast traffic from a specific group, it joins the group by sending an IGMP report (join request) to the network. (The multicast group specified in the join request is determined by the requesting application running on the IGMP client.) When a networking device with IGMP enabled receives the join request for a specific group, it forwards any IP multicast traffic it receives for that group through the port on which the join request was received. When the client is ready to leave the multicast group, it sends a Leave Group message to the network and ceases to be a group member. When the leave request is detected, the appropriate IGMP device will cease transmitting traffic for the designated multicast group through the port on which the leave request was received (as long as there are no other current members of that group on the affected port).

Thus, IGMP identifies members of a multicast group and allows IGMP-configured hosts (and routers) to join or leave multicast groups.

The function of the IGMP Querier is to poll other IGMP-enabled devices to elicit group membership information. The switch performs this function if there is no other device, such as a multicast router, to act as Querier. The switch automatically ceases Querier operation if it detects another Querier. A switch with IGMP querier capability will become a Querier in the absence of any other Querier on the network. If you disable the Querier function on a switch, ensure that there is an IGMP Querier (and, preferably, a backup Querier) available. If the switch becomes the Querier, then subsequently detects queries transmitted from another device on the same VLAN, the switch ceases to operate as the Querier for that VLAN. In the above scenario, if the other device ceases to operate as a Querier, then the switch detects this change and can become the Querier as long as it is not pre-empted by some other IGMP Querier.

In a Livewire+ system, it is the responsibility of the audio nodes to generate the IGMP messages.


Prioritization

Within a link, we sometimes want to have audio mixed with general data. This happens, for example, when a delivery PC is playing audio and downloading a file at the same time, or when our mixing engine is sending and receiving audio and control messages simultaneously. To be sure audio always flows reliably, we take advantage of the priority functions that are part of the switched Ethernet system.

Compared to the original, modern Ethernet has an additional 4 bytes of data inserted into the frame’s header. One field provides a 3-bit priority flag, which allows designation of eight possible values.

Priority Level	IEEE Recommendation	Livewire+ Assignment
7	Network Control	
6	Reserved	Livewire+ Audio
5	Voice	Telephone Audio
4	Video Conferencing	
3	Call Signaling	Livewire+ Control & Advertising
2	High Priority Data	
1	Medium Priority Data	
0	Best Effort Data	

Highest-priority packets have first call on the link’s bandwidth. If high-priority packets are in the queue and ready to go, the lower-priority ones wait. If there is not enough bandwidth for both, low-priority packets will be dropped – but this is not a problem, as you will soon see.



The adjacent graphic shows only two queues, but the idea is the same for four or eight. Switches used for Livewire+ must support a minimum of four queue and priority levels. Some low-end switches have no support or may have only two queue levels.

If you have multiple switches in a hierarchical configuration, the priority information is carried automatically to all the switches in a system.

This prioritization scheme works only within a facility’s local area network. Because it is at the Ethernet layer, it has no effect past the router boundary into the internet. However, we also set the priority bits in the IP header to match the Ethernet priority so that as LAN switching evolves to use more Layer 3 intelligence, our packets will be ready.

The Role of TCP

TCP is a key to sharing high-priority audio with best-efforts data on a single network link. Because the acronym TCP/IP is so often written, many people think that they are necessarily and always joined. This is certainly not so. IP is independent from TCP and may well be used without it. For example, RTP/IP is specified for streaming media, and UDP/IP is used for a variety of transmissions, such as DNS, the Internet's name look-up service.

TCP has two functions: Ensuring reliable transmission and controlling transmission rate.

Routers and switches may drop packets when there is not enough bandwidth to transmit them or when they are overloaded.

They also do not guarantee to deliver packets in the same order as they were sent. And there is no protection for bit errors from signal corruption. None of this is a mistake or oversight in the design of the Internet. The inventors knew what they were doing: they wanted control of any needed correction process to be as close as possible to the endpoints, consistent with the general internet idea to move as much as possible from the center to the edges.

Certainly we need 100% reliable transmission for most data files – even a missed bit could have bad consequences. TCP gets this done by using a checking and re-transmission approach. Whenever TCP detects any corrupted or missing data, it requests another copy to be sent and holds any data it might already have in its queue until the replacement has arrived. Packets are numbered by the sender so that they can be delivered to the application in correct order. The application always gets good data – but it could be after significant delay.

Transmission rate control is essential for most internet applications because the bandwidth of the many transmission “pipes” from sender to receiver are almost always different. And the available bandwidth to a particular user is constantly changing as the demands from the many users sharing the net ebb and flow. Think of the common case that you are working at your laptop with a cellular data modem, connected by VLAN to your office server. The server and its local network can certainly send data faster than your modem can take it. And the available bandwidth on the public part of the net is varying. So something needs to slow the sending rate to match both the network and your modem's ability to receive. That function is performed by TCP. This is called flow-control. While the details are complicated, the principle is simple: a TCP sender monitors the condition of the buffer at the receiver so it knows how fast the data is arriving and can adjust its transmission rate to maintain the correct average buffer fill.

TCP also has a function called congestion-control. While it also controls rate, it does it with a different mechanism and for a different reason. The re-transmission procedure we discussed earlier addresses a symptom of network congestion, but not its cause – too many sources trying to send at too high a rate. To treat the cause of congestion, we need to have some way to throttle senders when needed.

TCP's congestion control is unusual in that it is a service to the network at large rather than to the individual user. It was conceived as a way to fairly ration network bandwidth to all users. To do this, TCP monitors dropped packets, assuming that lost packets signal congestion. When a new connection is established, a slow-start function causes the rate to start low and ramp up until a lost packet is detected. Then the rate is cut in half and the ramp up begins again. In this way TCP is always probing for the maximum available bandwidth and always adjusting its transmission rate to match. It's really a very slick technique, one that is very well suited to getting the fastest transmission of bursty data over a shared link.

We've gone into a lot of detail on TCP because it is one of the keys to being able to share Livewire+ audio over a network link with other general data. The Ethernet switch handles congestion in a similar way to the routers in the internet – when there is too much traffic, it drops packets. But we have something very important: Priority. Audio packets are assigned higher priority than general data. So they are never dropped before all TCP packets are. The usual condition is that some percentage of the link is filled with constant audio streams and the remaining capacity is left for data. For example, an 8-audio channel Livewire+ node with all channels active will take about 40% of its 100BASE-T link, leaving 60% for data. But, we could have one or we could have a dozen audio streams active on a link – and this number could well change over time. TCP automatically finds how much bandwidth it can use and adjusts its rate naturally to match.

You might be thinking, “All well and good, but what if we put too many high-priority packets into the link? Won’t we have drop-outs then?” Yes, we would. But we never allow this to happen. Remember that each Livewire+ node knows about the link attached to it because it “owns” it. The link from a node to a switch is full-duplex point-to-point with no sharing. The node knows how many streams can fit and never is allowed to send more into or request for reception more than can be supported by the link.

All of the above applies to a shared link, such as for a delivery PC that needs both audio and data. It is the Ethernet switching function that allows the overall network to be shared, since general data never even gets to a port connected to a Livewire+ node.

Virtual LANs (VLANs)

This is a technology that came to Ethernet along with switching. It is a way to have “virtual” isolated LANs, while using common hardware.

Remember those Broadcast packets? They go to all devices, even with an Ethernet switch in the picture. If there are a lot of computers on the network, there could be a lot of traffic generated by these transmissions. VLANs can be used to contain broadcast packets, since they are not propagated outside of their assigned VLAN.

VLANs can also be used for security. If the Livewire+ network is on a different VLAN than the internet, a hacker would not be able to gain access to your audio streams or send traffic on the audio network.

In a Livewire+ network that is shared with general data, VLANs offer protection against a computer that could have a problem with its network software or interface card. The Ethernet switch can be configured so that the ports to which general computers are connected are not able to forward packets outside of the assigned VLAN, so can never reach Livewire+ audio ports.

Finally, VLANs protect against the rare case that an Ethernet switch has not yet learned an address and has to flood all ports until it knows the specific destination.

All Livewire+ devices allow choice of VLAN. We recommend:

- ◆ If you have a separate network for Livewire+ audio, you can just stay with the default VLAN 1 and pay no more attention to this topic.
- ◆ If you have your Livewire+ network connected to the Internet, or shared with a large group of office computers, use the default VLAN 1 for general data and VLAN 2 for Livewire+ audio and control.

A router must be used to bridge the traffic between VLANs, while providing a “firewall” function. So if you have PCs on the Livewire+ network that will be used for audio and web surfing, etc, you will need to provide this bridge. You will also need this to access Livewire+ devices on VLAN 2 with PCs connected to VLAN 1 for configuration, monitoring, etc.

A router that bridges VLANs is sometimes called a “one-armed” router because it has only one Ethernet port, rather than the usual two. But you can use the same router that you have for your Internet link to provide this function. Or maybe better: Some sophisticated Ethernet switches provide an internal routing capability that can be used to bridge VLANs. Simpler, and saves boxes.

Tagged vs. Port-Based VLAN Operation

When the VLAN information embedded in the Ethernet frame is used to direct the switch, this is called tagged VLAN operation. With Livewire+ devices, when you configure a VLAN value, the device will transmit Ethernet frames with the embedded value you specify. But some devices are not able to do this. That means the switch itself has to insert the tag – a procedure called port-based VLAN. In this case, all frames that enter from a particular port are tagged with a certain value, defined by switch configuration. To enable this, you must configure the switch appropriately.

In Windows 7, you can install Advanced Network Services to get VLAN support; it should also be available, when using compatible NICs, in Windows 8 and Windows Server 2012. Even with this capability, however, it is still more efficient to handle VLAN via the switch. Because the PC uses an Access port, we simply assign ‘switchport access vlan 100’ (for example) and then the PC participates in VLAN 100 with no settings needed on the PC.

In Windows 7, you can install Advanced Network Services to get VLAN support; it should also be available, when using compatible NICs, in Windows 8 and Windows Server 2012. Even with this capability, however, it is still more efficient to handle VLAN via the switch. Because the PC uses an Access port, we simply assign ‘switchport access vlan 100’ (for example) and then the PC participates in VLAN 100 with no settings needed on the PC.

There is one special case: Frames tagged with VLAN=0 are called priority frames in 802.1p standard. They carry priority information, but not the VLAN ID. The switch will translate to whatever VLAN is default for that port. This is useful if you want to use port-based VLAN assignment at the switch, rather than tagging from the Livewire+ device.

Many switches allow a combination of port and tagged VLAN. In this case you assign a default value to the port and frames either with no tag or with tag=0 go to this default VLAN, while tagged frames override the default.

It would be possible to use port and tagged VLAN in combination. For example, you use Livewire+ node configuration to put all your audio devices onto VLAN 2. Since the debut of Windows 7, there has been VLAN support inside Windows, but an easier method is to use your switch’s port-based assignment — simply set a port to be always VLAN 2, then plug your PC into this port.

Some switches have other options for assigning VLANs. Assignment could be “hard-coded” to MAC addresses with a configuration set-up. Or layer 3 protocols (TCP, RTP, etc) could be detected and used as a way to make VLAN assignments. These may have their place, but since Livewire+ devices provide the tagging, it doesn’t seem that these methods make much sense. The less you have to configure the switch, the better.

Ethernet Switching versus Routing

Both switches and routers examine packet addresses and send them appropriately on their way. So what is the relationship between these technologies? Why and where would you use one versus the other? Routing works at Layer 3, where IP information resides, while Ethernet switching works at Layer 2. Routing is a much more complex operation than switching, where multiple paths from one site to another are the norm, and it is the job of the router to find the optimum route (get it?), which may well be changing from minute-to-minute. A comparison of the two side-by-side:

Ethernet Switch vs IP Router Comparison

	Switch	Router
Layer	Layer 2/Ethernet	Layer 3/IP
Function	Determines to which port the addressed node is connected and switches incoming frame to it	Finds the best route from among many and forwards packet to next router along the path
Terminology	“Switching”	“Forwarding”
Technology	Simple table look-up in hardware	Complex dynamic best-route determination in software
Standards	IEEE	IETF
Ports	Many, connecting mostly to end nodes	A few, connecting to networks and Telco lines

As do switches, routers also support multicast and prioritization. But routing is a much more precise and efficient way of controlling destinations for our mission-critical audio. This is why Livewire+ networks user Layer 3 routing.

Traditionally, routers did their work with software, while switches had dedicated hardware chips. Now there is something called Layer 3 Switch, a hybrid of traditional routers and Ethernet switches. Layer 3 switches perform their forwarding – whether Layer 2, Layer 3, unicast, multicast, or broadcast – in hardware. Software handles network administration, table management, and exception conditions.

Since we first introduced studio AoIP in 2003, the cost of such devices has fallen dramatically, making it even more economically feasible to base studio networks on Layer 3.

Livewire+ Networks

So now we are ready to consider all that has gone before in the context of Livewire+. And to begin the discussion technologies specific to Livewire+.

Quality of Service (QoS)

An important concept in a converged network is Quality of Service. When general data is the only traffic on a network, we only care that the available bandwidth is fairly shared among users and that the data eventually gets through. But when our studio audio and general data are sharing the same network, we need to take all the required steps to be sure audio flows reliably.

Our method for achieving QoS is system-wide, with the following components each contributing a part of the whole:

- ◆ Ethernet switch. Allows an entire link to be owned by each node. Isolates traffic by port.
- ◆ Full-duplex links. Together with switching, eliminates the need for Ethernet’s collision mechanisms and permits full bandwidth in each direction.
- ◆ Ethernet Priority assignment. Audio is always given priority on a link, even when there is other high-volume non-audio traffic.

- ◆ IGMP. Ensures that multicasts – our audio streams – are only propagated to Ethernet switch ports that are subscribed.
- ◆ Limiting the number of streams on a link. Audio nodes have control over both the audio they send and the audio they receive, so they can keep count and limit the number of streams to what a link can safely handle.

The result is rock-solid QoS, combined with the ability to share audio and data on the same or interconnected networks.

Source Advertising

Audio source nodes advertise their streams on a special multicast address. Receive nodes listen to these advertisements and maintain a local directory of available streams. The advertisements are sent when the streams first become available and at 10-second intervals after that. (Actually, only the data version number is sent every 10 seconds. The full data is advertised only upon entering the system, on any change, and on explicit requests from those having detected the data version number increase.) If a node's advertisements are not received for 3 consecutive periods, it will be assumed to be removed from service. There is also an explicit "stream unavailable" announcement.

Receive nodes maintain a local table of available streams and their characteristics, updated as any new information arrives. Sources are cleared from local tables when an explicit message is received announcing that a stream is no longer available, or when three consecutive advertisements have been missed.

A receive node may be configured to be permanently connected to particular multicast streams, or users may select audio sources from a list. The list may display all available sources, or a filtered subset.

Synchronization

You may ignore this matter completely – and your Livewire+ system will work automatically "out of the box". But there are times when you might want to modify the default behavior of the clock sync system, so here is some detail on how the system works.

Livewire+ (and AES67) streams need careful system-wide synchronization in order to have small buffers for low-latency streams. If we did not have a distributed way to derive a bit clock, we would eventually have buffer over or under-flow, resulting from the input and output node clocks being not exactly the same frequency.

A PLL (Phase Lock Loop) in each Livewire+ node recovers the system clock from multicast clock packet that is being transmitted at a regular interval. At any given time, one Livewire+ hardware device is the active system clock master. In the event the master develops a fault or is removed from service, the local PLLs in the nodes are able to "ride out" the brief interruption and there will be no problem with operation.

All nodes are capable of being a clock source, and an arbitration scheme ensures that only the unit with the highest clock master priority is active. Clock mastership priority may be set by the user, or left to the default case of all being equal priority.

When the clock goes away for 3 consecutive periods, all capable units begin transmitting clock packets, after a delay skewed by their clock mastership priority.

When a unit sees clock packets from a unit with a higher mastership priority on the network, it stops its own transmit of clock packets.

You can specify the clock mastership priority behavior. The clock mastership can be made predictable, rather than end up being any node in the plant – maybe the one down in an out of the way equipment closet.

Each node has a clock mastership configuration setting that can range from 0 to 7.

- ◆ ‘0’ means never - slave only
- ◆ ‘7’ means always - forced master (don’t use multiple forced masters in a system)
- ◆ Factory default is 3. So all units have equal priority out of the box, and the following is used to break ties (in descending order): lowest Livewire+ audio transmit base channel, then lowest IP address, then lowest Ethernet address.

Livewire+ xNodes have a status display on their front panel that shows “MASTER” when that unit is the clock master; else the unit displays “SYNC” to indicate that it is a slave synchronized to the master clock.

Jitter in the timing and PLL functions ultimately set a lower bound on output buffers and therefore audio delay. And any drift in the time calculation produces buffer pointer drift. Further, jitter in the derived A-to-D and D-to-A bitclocks causes sampling uncertainty that can generate unwanted noise in the audio.

The LAN network is a “noisy” environment with packets of various kinds and lengths being numerous and unpredictable. Thus, the PLL system needs to be quite smart so as to generate a reliable, consistent, low-jitter output, while not being confused by dropped or jittered time packets.

Our method for handling this PLL problem is the subject of a patent, to give you some idea of the novelty and complexity involved in solving it.

Synchronizing to AES3 Systems

To avoid passing audio through sample-rate-converters, we recommend that Livewire+ be synchronized to your AES master clock, if you have one. Our Livewire+ AES node provides this function, recovering the clock from an attached AES input and creating a Livewire+ sync packet. The AES node must be active clock master.

Here’s an interesting sync application for Axia AES/EBU xNodes: Two AES xNodes can be used as a way to synchronize two AES systems located apart, but with an available IP path between them. One becomes the master, connecting to a Livewire+ AES input. The slave attaches to a Livewire+ AES output and is configured to recover clock from it. Ingeniously simple, no?

Synchronizing to AES67 Systems

AES67 standard requires device to use the IEEE-1588 network time protocol standard. Synchronization related options are available on the xNodes’ QoS web page, accessible with your PC’s browser.

Synchronization:

Clock mode:	<input type="text" value="PTP/IEEE 1588 ARB clock class 248"/>
PTP domain number:	<input type="text" value="0 default"/>
PTP delay mechanism:	<input type="text" value="E2E - uses the end-to-end delay request-response (default)"/>
PTP clock priority1/priority2:	<input type="text" value="128"/> / <input type="text" value="128"/>
PTP clock sync interval:	<input type="text" value="-3 (1/8) - AES67 Default"/>

- ◆ Clock Mode selects the protocol required by the standard. xNodes offer a few options depending on other system requirements:
 - ◇ PTP/IEEE 1588 ARB clock class 248 mode allows the node to be a master when no better clock is available on the network. Since xNodes are not locked to GPS or any other high precision source, only Arbitrary (ARB) Clock class is provided.
 - ◇ PTP/IEEE 1588 slave only is recommended for networks with high precision grand masters present.

- ◆ **PTP Domain Number:** “0” should be selected unless there is a need for multiple clock domains on a network segment. Choosing a different domain is a useful option for doing tests without affecting other components present on the network.
- ◆ **PTP delay mechanism** specifies how the xNode determines network round trip latency. This setting **MUST** match the master clock setting.
 - ◇ In practice, **End-to-End (E2E)** is always used by ordinary clocks and boundary clocks.
 - ◇ **Peer-to-Peer (P2P)** mechanism was defined to limit amount of traffic reaching the master clock, but is rarely used.
- ◆ **PTP clock priority1/priority2** numbers provide control over the master clock selection. Choosing a lower priority number makes the clock a preferred choice over devices with a higher number. The range is 0-255, and 128 is the default and middle value.
- ◆ **PTP clock sync interval** determines how often clock packets are sent. The faster those packets are sent, the faster and more accurate synchronization can be achieved. Some IEEE-1588 devices do not support lower clock sync interval than 0.5s, so this setting may need to be adjusted depending on what other networking equipment is used. AES67 recommends an interval of 1/8 (8 packets per second) . This setting only affects behavior of the master clock, and does nothing when the xNode is a slave.

Other xNode AES67 settings are discussed in detail in the xNode User Manual.

Network Standards and Resources

We use standards throughout. Ethernet, as well as the AES67-2013 specification which defines the way that AoIP devices interchange audio between manufacturers, is standardized by the IEEE and information is available on their website at www.ieee.org. Internet Protocol and associated technologies are standardized by the Internet Engineering Taskforce (IETF) and much can be learned from their website at www.ietf.org. Documents are a free download. Bookshops are full of books on Ethernet, IP, and networking and we offer a list of suggested reading.

Livewire+ operates at both Ethernet and IP network layers, taking advantage of appropriate standards-based resources at each layer.

Here are the resources we are using at the various layers:

- ◆ **Layer 1:** IEEE Ethernet Physical
- ◆ **Layer 2:** IEEE Ethernet switching, IEEE 802.1p/Q prioritization, IEEE 802.1p multicast management
- ◆ **Layer 3:** IETF IP (Internet Protocol)
- ◆ **Layer 4:** IETF RTP (Real-Time Protocol) , IETF UDP (User Datagram Protocol) , IETF TCP (Transport Control Protocol) , IETF IGMP (Internet Group Management Protocol)
- ◆ **Layer 5:** IETF NTP (Network Time Protocol) , IETF DNS (Domain Name Service) , IETF HTTP/WebIETF ICMP Ping, IETF SAP/SDP (Session Announcement Protocol/Session Description Protocol, in the Windows PC Livewire+ Suite application)

Network Time Protocol (NTP)

This is the Internet’s standard for conveying time. There are a number of servers on the net that users can connect to in order to retrieve accurate time. There are also boxes from manufacturers such as EXE that receive radio time signals and translate them to NTP packets. Livewire+ does not need NTP, but some peripherals do. For example, Axia studio mixing surfaces and Omnia processors use NTP to automatically synchronize to the correct time.

A Note about Protocol Design

There is no question that among network protocols, the Internet has been an impressive success. One of the reasons for this was the approach its designers took – and still use – when inventing its protocols. These are outlined in the IETF RFC 1958 document. We’ve taken the principles to heart in the design of Livewire+. Here they are, in priority order, and with our comments in parenthesis:

1. Make sure it works. Make prototypes early and test them in the real world before writing a 1000-page standard, finding flaws, then writing version 1.1 of the standard. (The Telos Alliance is a practical commercial outfit, not an academic or governmental organization. We had two years of extensive lab tests of prototypes in two locations, and then real-world field tests at radio stations before locking the core tech down. Our prudence has now been proven out around the world in thousands of broadcast facilities.)
2. Keep it simple. When in doubt, use the simplest solution. William of Occam stated this principle (Occam’s razor) in the 14th century. In modern terms, this means: fight feature creep. If a feature is not absolutely essential, leave it out – especially if the same effect can be achieved by combining other features. (We believe firmly in this principle. We tried very carefully to add nothing unnecessary.)
3. Make clear choices. If there are several ways of doing the same thing, choose one. Having multiple ways to do something is asking for trouble. Standards often have multiple options or modes or parameters because several powerful parties insist their way is best. Designers should resist this tendency. Just say no. (It was just us – and we did say no. No committees or politics to cause bloating.)
4. Exploit modularity. This principle leads directly to the idea of having protocol stacks, each of whose layers is independent of all the other ones. In this way, if circumstances require one module to be changed, the other ones will not be affected. (We built Livewire+ on all of the available off-the-shelf lower layers.)
5. Expect heterogeneity. Different types of hardware, transmission facilities, and applications will occur on any large network. To handle them, the network design must be simple, general, and flexible. (We had to accommodate both dedicated hardware audio nodes and general-purpose PCs being used as audio nodes.)
6. Avoid static options and parameters. If parameters are unavoidable, it is best to have the sender and receiver negotiate a value than defining fixed values. (These were avoidable – we don’t have any such negotiated parameters. We do have the receiver selection of stream types, but this is simple one-ended selection.)
7. Look for a good design, not a perfect one. Often designers have a good design but it cannot handle some weird special case. Rather than messing up the design, the designers should go with the good design and put the burden of working around it on the people with the strange requirements. (Steve, Mike, and Greg’s mantra! Make it work, make it solid, build enough flexibility to get the job done – and no more.)
8. Be strict when sending and tolerant when receiving. In other words, send only packets that rigorously comply with the standards, but expect incoming packets that may not be fully conformant and try to deal with them. (We told the s/w guys to do this. Hope they did!)
9. Think about scalability. No centralized databases are tolerable. Functions must be distributed as close to the end-point as possible and load must be spread evenly over the possible resources. (We kept very close to this idea – which is the main spirit of the Internet. We don’t have any central databases or other pieces along these lines. We have a fully distributed system. If one part fails, the others keep going.)
10. Consider performance and cost. If a network has high costs and there are cheaper variants that get the job done, why gold-plate? (Compare the power and cost of our solution with others. Using simple off-the-shelf commodity parts was the guiding principle for our work.)

9

Frequently Asked Questions

We know there will be questions. Here are some we've already heard, and some we imagine.

General Questions

How do I know that Audio over IP will be reliable?

Axia gear with Livewire+ uses the same technology that underlies VoIP telephony. Did you know that nearly 75% of the Fortune 100 companies now use VoIP? Or that VoIP PBX systems now outsell the old kind by a wide margin? With these systems, telephones plug into a standard Ethernet/IP network. Contrast this with traditional PBX phone gear — proprietary devices which required you to purchase phone sets and parts exclusively from the company that built the mainframe. You were locked into a single vendor, because the technology that ran the mainframe was owned by the company that made the gear. (Kind of like the TDM router companies.)

IP is now accepted as a universal transport for almost any kind of signal. You see it in television studios, business teleconferencing, government communications, banking, etc. And it's hardly unproven, especially for applications specific to radio studio infrastructure. As of April 2015, over 5,500 studios around the world - many in mission-critical, 24/7 broadcast applications in major markets like New York City, Chicago, Paris, Rome and Bangkok - have been built using Axia IP-Audio infrastructure.

That's a lot of customers. Are they happy?

Very! Axia systems are faster to install than traditional routing setups, work reliably and are easy to reconfigure. Why not talk to the people actually using it and see what they have to say? We'll be happy to provide you with a list of references upon request.

Can the network be used for general data functions as well as audio?

Most certainly, should you choose to do so. The Ethernet switch naturally isolates traffic. You may even use one link for both audio and data, since the audio is prioritized. This will probably be the case when a PC is connected to the network – you will sometimes want to download files, receive email, etc. in addition to the audio stuff. Switch selection is important, though, and you must use one tested and recommended by us. You could have two networks and link them as described below.

Of course, we would never mix on-air audio and business functions or open ourselves up to hacking.

Can I make this a completely separate network?

Yes, we understand and agree. You have a few choices:

- ◆ Have a completely separate and isolated network for Livewire+. Take advantage of Ethernet, but don't combine any internet or business functions with studio audio.
- ◆ Have two physical networks and link them with an IP router. Correctly configured, the router provides a security barrier.
- ◆ Share the network hardware for audio and general functions but isolate Livewire+ to its own VLAN. Again, an IP router could be used to link the two networks.

I've heard that Livewire+ networks cost a lot less than traditional studio systems. What did you leave out?

Nothing. Our cost savings compared to traditional routers are achieved by using standard, off-the-shelf switching hardware rather than custom-built solutions. It's a lot less expensive to use standard Ethernet for signal switching and transport than it is to construct a customized cross-point routing switcher, with its cards, frame and peripherals. This is the same principle that has driven almost all stations to use PCs for audio playout and editing – they are a lot cheaper and more powerful than any broadcast-industry specific machine could be.

Another way Livewire+ saves money lies in the way PC audio is handled. With a traditional router, PC audio must be brought in through a router input card or console module; bringing multiple channels of audio into the system in this manner (from workstations or digital delivery systems) can significantly increase the overall cost of the system.

Instead, we wrote an IP-Audio Driver for Windows PCs that looks just like a sound card to the OS, but streams audio in and out of the computer's network card instead. Or, if you need the realtime MPEG compression or time compression features of a high-end sound card, our partner AudioScience makes an audio card with a Livewire output that plugs directly into the Axia network. Either of these approaches eliminates the cost of the I/O needed to get audio into the switching network. So Axia clients usually realize several thousand dollars worth of savings over and above the cost of the sound cards themselves.

Do Livewire+ networks have any single points of failure? Is there a central 'brain' I can lose that will take the system down?

Livewire+ networks are distributed, with no central box. Ethernet networks can be designed any number of ways, including those that are fully-redundant and self-healing. Normally, our clients build larger facilities with "edge switches" serving each studio, connected to a redundant core. Each studio is able to operate stand-alone.

I've heard that with AoIP, latency increases whenever you add inputs. The more sources you add, the higher the delay. Is this true?

No, Livewire+ latency remains fixed at the same low value regardless of the channel count. You can run a system with a thousand channels and the latency will be the same as for a single stereo stream. Indeed, the delay is so consistent that channel-to-channel phase shift is less than 1/4 sample. The total latency of an analog input to analog output using the Axia Livestream format is about 2.75 milliseconds:

- ◆ The time through the A/D and D/A converters is about 1.5 ms.
- ◆ The network transit time is 1.25 ms.

To put this into perspective, modern, high end audio processors typically clock in with around 10 ms. delay (and talent regularly monitors those on-air).

Does Livewire+ route logic with audio, too?

Of course! IP is great for data, no? Logic commands from external devices enter the network using GPIO ports on xNodes and mixing engines. The logic data is then “bound” to the audio stream, and is routed with it to whatever console the source is loaded on.

Devices equipped with Livewire+ interfaces (like Telos Zephyrs and phone hybrids, Omnia audio processors, 25-Seven profanity delays and IDC satellite receivers, for example) supply audio and control logic directly from the device to the Ethernet switch over a single CAT-5e connection, further simplifying in-studio wiring and making Livewire’s audio+logic routing even more convenient.

Are there any problems with delay of control commands over the network? I’ve heard of other systems using TCP/IP that have problems in this respect.

No, Livewire+ control latency is very small – no more than 50ms for hardware GPIO closures from Surface button pushes. We are using a special network protocol we invented called R/UDP (Reliable UDP) rather than TCP/IP, in part to be sure control delay is low.

What about Program Associated Data? Is Livewire+ compatible?

Yes. Devices that generate PAD plug into the Axia network; the information they supply is sent along with its associated audio, and any devices that need it can also plug into the network and retrieve it. This means that you can send audio and PAD together, without incurring extra costs for separate audio and data networks.

Mix-minuses are always tricky and I’ve heard that Axia consoles are good at mix-minus. What’s so different about the way you do it and how are they generated?

As a part of Telos, you might imagine we’ve studied mix-minus for a long time. And we’ve always been amazed at how complicated and confusing it is to set up correctly. With today’s radio shows relying heavily on phones and remotes, something needed to be done.

Our consoles generate mix-minuses automatically, on-the-fly, without any intervention needed from talent. The way it works is simple: when a caller is on the air, he hears the main Program feed, minus himself. When off the air, he hears a special “off-line” phone mix that can contain audio sources: pre-fader audio from the host mic, other phone callers, etc.

Best of all, mix-minus settings for audio sources such as phone hybrids and remote codecs are assigned to the source itself, not the console fader — so when a source that needs a mix-minus is loaded onto any fader, on any console, the mix-minus settings are automatically loaded too.

At the physical level, mix-minus is easy, too. Livewire+ carries audio in both directions, so one RJ-45 covers everything.

How many mix-minuses can your consoles have?

One for every fader! There’s a lot of processing power in our mixing engines, enough that Axia consoles can provide automatic mix-minuses simultaneously for every fader on the console. You’ll probably never need to have 40 mix-minuses running at once, but isn’t it nice to know that you could?

Can I use Livewire+ without Axia consoles?

Yes, of course. You could just use it as a snake or router system and connect whatever consoles and other equipment you like.

How does Livewire+ compare to other audio networking systems?

The original Livewire protocol debuted in 2003 as the first audio networking system to allow real-time uncompressed digital audio to be conveyed over standard Ethernet hardware, so it's accurate to say that all other AoIP networking systems sprang from our work.

AES67-compliant Livewire+ is extremely low latency, which is especially important for broadcast facility operation, where live monitoring and cascaded links are common. Livewire+ includes all the technology you need for practical studio application: Switches are controlled, sources are ID-ed and advertised to receivers, GPIO over the network is covered, etc. Finally, Livewire+ connects directly to PCs – no soundcard or other hardware is required.

Livewire+ is a not just a technology, but rather a get-the-job-done solution. We offer you all the pieces you need to build a modern broadcast studio. Interfaces, mixing engines, consoles, PC drivers. We are experienced broadcasters, so we know how to support radio studio applications.

As you compare Livewire+ to other systems, be sure to note the number of devices that work with those systems, and the number of technology partners in-hand. Livewire+ is a mature system, with dozens of products and partners that contribute to a robust, connected studio equipment ecosystem. Remember, the more devices connected via Ethernet, the less infrastructure cost and the easier systems are to deploy, configure and manage.

What is the best audio format to use with Livewire+ systems?

Our networks don't care what format your audio is in. During the playout process, your playout software may uncompress any compressed-format files (MP3, MP2, apt-x, etc.) and present them to the Axia IP-Audio Driver as PCM.

Similarly, compressed audio, such as Dolby Digital or Dolby E, can be carried across a Livewire+ network without being decoded, just like it is handled in an AES or SDI stream.

What this means is that all audio that moves within the Livewire+ system can be the same — 48 kHz, 24-bit studio-grade audio as either linear audio or compressed audio data, and the switches we recommend have enough bandwidth to carry 10,000+ channels of uncompressed, real-time stereo audio simultaneously.

Is it possible to run out of bandwidth?

No. The backplane of a modern Ethernet switch can handle full duplex traffic on all ports simultaneously without any packet loss. And since our component links are designed so that they never exceed any port's capacity, we never exceed the switch capacity. The way we prevent port overload is simple: we "own" each port. Every xNode audio interface is plugged into an unshared 100Base-T port on the switch. Even when all of an xNode's inputs and outputs are active, we are still well under the bandwidth of the ports, and the switch is completely under control. Because the switch has the backplane capacity to handle all ports fully loaded, the system performance doesn't change from one to thousands of audio channels.

Let's explore the issue of switch capacity a little further. We know how much capacity is required per port for each node, and we know that a node will never produce or consume more than 16 stereo streams total. But what about the studio mixing engine? To support a large console with a lot of buses, inputs, mix-minus outputs, etc., you may have 40 or 50 simultaneous signals (or more). Because this could exceed the port capacity of a 100Base-T port, the mix engine is connected via Gigabit Ethernet only. Using Gigabit for the engine, we could support a 200 fader console with 200 outputs and still have room to spare!

Are you sure this is robust enough for 24/7 operation? My Windows networks always have downtime.

Livewire+ equipment is based on tight, embedded hardware and software. The Ethernet switches we recommend are fully professional devices with high reliability and options for redundancy.

Can I connect two studios across town?

Yes, but not the way you're probably thinking. Remember that Livewire+ audio is uncompressed 24-bit, 48kHz, so each stereo stream is 2 Mbps. To get this done over a reasonable phone line, you'd use the Telos Zephyr iPort PLUS MPEG Gateway. This unit can be used to transport multiple channels of stereo audio across any network with guaranteed QoS, such as T1 and T3 connections, or MPLS networks. It contains eight bi-directional stereo MPEG-AAC codecs in one box and connects to your Livewire+ network using a single CAT-6 cable for all I/O.

GPIO and metadata "ride along" with the audio as well, so Livewire+ audio streams transmitted across town via a Zephyr iPort link look like "local" audio sources to the receiving end, and can be assigned to console faders and used just as if they were originating across the hall, with no visible difference in operational behavior.

If my favorite delivery system isn't on your list of partners, can I still use it with a Livewire+ network?

Yes, the same way you do it now: just plug the delivery system's outputs into our inputs, and send the contact closures into our GPIO ports. (Then ask your delivery system provider when they're going to become an Axia partner!)

Our facility requires mono audio, not stereo. Are Livewire+ systems capable of mono streams?

Yes. With Livewire+, all audio is data, so Livewire+ can route mono, stereo and even 5.1 Surround streams, too. Our xNode audio interfaces can be configured for true mono, stereo, or surround I/O. This is built-in to all of our Analog and AES/EBU xNodes.

Is a Livewire+ network more difficult to install than traditional routing systems?

In fact, Livewire+ is much easier to install and commission, because everything connects using off-the-shelf Ethernet cables, which carry multiple uncompressed channels of stereo audio. 100Base-T links can carry 25 stereo or 50 mono audio channels simultaneously; Gigabit links can handle 250 stereo or 500 mono channels each. This labor savings from this link aggregation and resulting elimination of expensive multi-pair cable for studio interconnects can be significant.

Even our audio connectors are designed to promote fast, inexpensive installation. Livewire+ devices use the Radio Systems StudioHub+ RJ-45 standard for I/O jacks; a huge variety of adapters are available for all kinds of devices. Tally up the savings in labor realized from not having to purchase and hand-solder hundreds of XLR and RCA connectors, and reduction in complexity becomes even more impressive.

In fact, due to the reduction of cabling and the quick connection of devices, clients tell us that installation goes 30% to 50% quicker than wiring studios the traditional way.

Livewire+, Standards and Other Vendors

I hear you've renamed Livewire to Livewire+™. What does this mean?

You probably know that the Telos Alliance has been involved in designing the architecture of AoIP interoperability since the very beginning (when some folks claimed it wouldn't even work!).

First, we promoted interoperability using the standard IP networking technologies built into Livewire, beginning in 2003. Then we became sustaining members of the X.192 project committee, which wrote the AES67 standard ratified by the AES67 in September, 2013. By December 2013, we had already implemented AES67 in our Axia xNode audio interfaces.

But we didn't just stop there. We've continued to fully integrate AES67 into Livewire+, in order to give Axia clients the best possible experience. Not only that: we've structured Livewire+ so that it can be extended to comply with future standards.

This represents a huge leap forward — so much so that Livewire technology has evolved to become Livewire+.

I heard that Livewire+ is actually the basis for AES67. Is that true?

As part of the X.192 working group, Axia and the Telos Alliance worked with other companies to help define the standard. During this process, we made available, at no charge, parts of our own patented technology to help speed development of the standard.

Does that mean that Livewire+ actually is AES67?

Not exactly! What it does mean is that much of Livewire+ is the basis for AES67. You might say that AES67 has Livewire+ in its DNA.

So what does the "plus" in Livewire+ mean, exactly?

It means three things.

1. Livewire+ is completely AES67-compliant. Compliant means that Livewire+ fully complies with all parts of the AES67 standard.
2. Livewire+ is future-proof, because it's extensible. Future standards (such as might come from the AES X.210 group) can be included once ratified; Livewire+ can never become obsolete.
3. Choosing Livewire+ means you don't have to wait for interoperability. Livewire+ represents the world's largest grouping of AoIP technology partners – over 80 software and hardware manufacturers and integrators whose products work together, with integrated audio, source discovery, logic control and data — right now, without the need to wait for additional standards.

You say Livewire+ is "AES67-compliant". Other companies say their gear is "AES-67 compatible." What's the difference?

The difference is subtle, but important — compliant is not the same as compatible! "Compliant" means that Livewire+ fully complies with all parts of the AES67 standard. It's baked-in, if you will.

"Compatible", on the other hand, means that someone's proprietary tech can co-exist with the AES67 standard. In essence, compatibility is like two different people sharing one body: the proprietary tech is patched with add-on code to get by, but doesn't fully comply with all aspects of the standard.

What does the AES67 standard contain?

According to Mark Yonge, AES Standards Manager:

“This standard defines an interoperability mode for transport of high-performance audio over networks based on the Internet Protocol. For the purposes of the standard, high-performance audio refers to audio with full bandwidth and low noise. These requirements imply linear PCM coding with a sampling frequency of 44,1 kHz and higher and resolution of 16 bits and higher. High performance also implies a low-latency capability compatible with live sound applications. The standard considers latency performance of 10 milliseconds or less. This standard provides comprehensive interoperability recommendations in the areas of synchronization, media clock identification, network transport, encoding and streaming, session description and connection management.”

(Reference: <http://www.aes.org/standards/blog/2013/9/aes67-2013-audio-over-ip-130911>)

If AES67 does all this, why does Livewire+ exist?

Good question! The AES67 specification is a good start toward interoperability, but is actually a subset of the many functions that Livewire+ performs today. When we developed Livewire back in the early 2000s, we had to synthesize the critical links between networking technologies, because a full standard didn’t exist yet — and earlier networked audio systems lacked critical functionality. We developed a way to make GPIO logic “ride along” with the audio streams, and a way for available sources to “advertise” their availability to all the networked devices that operators might want to use.

We also recruited [partner companies](#) whose products are respected and widely-used in the radio industry. Then we shared our technology with them, so that station engineers could connect as many audio devices as possible directly to their audio network. Not only is native connectivity an elegant solution, it reduces total system costs by removing the need for those extra I/O devices.

Ten years later, the industry finally adopted a standard for AoIP audio transport in AES67. The goal is for every studio audio device to eventually click together with CAT-5 and share audio. But along with that shared audio, there’s a whole world of other functionality that broadcasters expect — like device start/stop functions, monitor mutes, on-air tallies, the ability to control peripherals from the console, the ability to know when an audio source is live and ready for air, the ability for playout systems to control fader on/off functions and more. Those are functions that AES67 alone doesn’t provide for, but Livewire+ does. And now that the first AES67 devices are appearing in the marketplace, broadcasters have quickly found that they also need to support those additional capabilities in order to provide an integrated control experience for the user — otherwise they’re no better than AES3 streams, with serial GPI cables running alongside.

With Livewire+, you can have your cake and eat it, too. Over 5,500 Livewire+ powered consoles and 60,000 devices are in use around the world daily. Livewire+ provides audio with integrated control and discovery right now. And, because Livewire+ is extensible, broadcasters using it can enjoy its rich control architecture and integration with over 80 partner products, exchange audio with AES67 devices, and be assured that when an AES standard for control is ratified, it will be part of Livewire+ as well. Livewire, the original AoIP technology for broadcast is still on the leading edge, and is the most successful by a wide margin.

Does Livewire+ still work with RAVENNA?

Yes. Axia and RAVENNA have been partners for several years, both working together to define the AES67 standard. Going forward, standards-based Livewire+ will continue to be backward-compatible with the RAVENNA protocol.

The Telos Alliance and ALCNetworX were a big part of the X.192 project. Isn't AES67 just a synthesis of those two systems?

Not at all. The X.192 Task Group consisted of over 100 members from a wide variety of equipment manufacturers — including some who are direct competitors. It wasn't a "rubber-stamp" project! AES67 is the result of 3 years of collaboration by these many contributors, ensuring use of the best ideas, no matter where they originated.

You say that Axia introduced AoIP to broadcast in 2002, and you've got 5,500 studios on-air. So why is AES67 even needed?

When Steve Church, Greg Shay and their team were developing Livewire, they had to invent tech that didn't exist before. One critical piece of tech was network clock sync. Problem was, the Ethernet standard in place at the time had no criteria for high-precision time-synched audio.

Why is this so critical? As Telos Alliance Chief Science Officer Greg Shay explained in his excellent paper "Taking the 'Sting' Out of Evolving Digital Audio Networks" (presented at NAB 2013 and downloadable from www.TelosAlliance.com): "Accurate timebase recovery is directly related to, and essential for, low latency (low delay) of the audio going over the network."

In other words, if you want networked, real-time, broadcast quality audio without jitter and delay, you've got to have all network devices synchronized to a network master clock. So we invented the first distributed high-precision clocking system for Ethernet, and debuted it in Livewire.

Although Axia has always been happy to share our tech with software and hardware manufacturers, other, open methods of Ethernet clock sync emerged, one of which became the IEEE-1588 synchronization standard — which is an integral part of AES67.

I need new studios. Should I wait to purchase equipment until all the manufacturers adopt AES67?

Not unless you're prepared to wait quite a while. Implementing a new standard always takes time, as manufacturers adapt existing products, design new ones, and release software updates.

In the meantime, Livewire+, with over 5,500 studios on-the-air worldwide, is working today, and has a roadmap for future compatibility of existing hardware. This assurance means that the studio you build with Livewire+ today will be compatible with future AES standards for AoIP — your studio gear can't become obsolete.

What changes, if any, will be required for my existing Livewire network to support AES67?

Compliance with AES67 is a part of the latest software present in Axia analog, AES/EBU, microphone and mixed-signal xNode AoIP interfaces, and Linear Acoustic SDI xNode devices. If you have these products already, simply download the update packages and apply them using a standard Web browser and PC, or use Axia iProbe software to "push" the update to your devices en masse. After this update, you'll be able to generate and consume native AES67 streams.

And of course, the Telos Alliance plans to implement AES67 capability within more products going forward — your assurance that studios built with Livewire+ are future-proof.

Will I need to change my Ethernet switches to work with Livewire+ and AES67?

No! Livewire+ and AES67 coexist on the same switch fabric.

What are the advantages of specifying Partner equipment that works with Livewire+, versus gear that supports AES67 alone?

The whole idea of AES67 is to open up more equipment choices when you buy. The AES67 standard enables equipment from a variety of manufacturers to exchange “agnostic” audio streams without fuss. However, the AES67 specification provides interoperability standards for audio only — not for control or data exchange. Nor does AES67 equipment have the ability to automatically discover available audio, and present it to the operator for use. Most engineers agree that these abilities are a big part of the AoIP advantage. So, if control and source advertising is important to you, choosing equipment with Livewire+ capability will definitely be to your advantage.

Are you planning to share information so that other vendors can make gear that directly plugs to Livewire+?

Yes, and in fact we have done this for several years. Software vendors for PCs can use our driver to easily make their applications compatible; many delivery system providers already have built in Livewire+ compatibility. Through our Livewire Limitless License program, we provide technical data and schematics that allow hardware vendors to build Livewire+ interfaces into their products as well; this is why Livewire+ networks have the largest ecosystem of devices that connect directly — nearing 100 partners as of this writing.

Of course, you can use whatever equipment you want via the analog and AES xNodes.

PCs and Livewire+

Tell me about your “sound card” driver for workstations.

The official name is “Axia IP-Audio Driver for Windows”. It makes the Livewire+ network look like a sound card to a PC Windows application. Most audio applications should work unmodified.

We’re a Linux facility. Will the Axia Driver work with Linux?

Yes, we’ve partnered with Paravel Systems, developers of the Rivendell Radio Automation System for Linux. They can compile Axia IP-Audio Drivers to work with your Linux PCs. To find out what distros are currently supported, contact them directly at www.paravelsystems.com.

Building Livewire+ Facilities

I’ve got a large facility. How many studios can I interconnect?

There is no practical limit. Livewire+ networks can handle as many as 10,000 audio streams at a time; you may have as many studios and audio channels as your Ethernet switch fabric can support. Switches come in all sizes, some with hundreds of ports. And multiple switches may be cascaded to expand ports. For large networks, we recommend that you use a switch per studio to isolate any problems to a defined area. These are then interconnected with a backbone. Switches may be physically associated with each studio or may be all in a central location, as you prefer.

What about for smaller stations? This all sounds pretty sophisticated for a simple set-up.

Look at Ethernet for data applications. You have everything from a single PC connected to a printer to a few PCs in a small office tied to the Internet and a couple of printers to huge campus networks with thousands of nodes. This is one of the reasons we went with Ethernet — you can use it for big and small facilities. The technology and economics naturally scale to suit the application size. We figure, in fact, that small stations may benefit the most as they gain routing capability at a very modest cost.

This seems like a lot of IP to keep track of. What administration tools does Livewire+ have?

All Livewire+ devices have a web browser control and monitoring capability. Keep the IP numbers in a “favorites list” and you can easily check them. Or make your own web page with all the links.

An off-the-shelf tool is the Axia iProbe Network Management Console. iProbe is an intelligent network maintenance and diagnostics tool that makes managing, updating, and remote control of your Livewire+ system very easy and intuitive.

What Ethernet switches will I need for my Livewire+ network?

You won't need any bleeding-edge products from the skunkworks, if that's what you mean. In fact, you may not need a third-party switch at all! Every Axia console has a custom-built, zero-configuration Ethernet switch built in – perfect for standalone studio deployment, but these switches may also be daisy-chained to create larger studio networks. If you just need a few local switch ports on the edge of your network, or if you're building a smaller studio, our standalone xSwitch might be what you're looking for – it's pre-configured to install in under a minute.

Of course, facilities with multiple rooms and large routing networks require core switches with larger capacity. For these systems, we've qualified a number of switches from Cisco, the industry leader in switchware. You can see the selection at www.TelosAlliance.com/switches.

You built your own standalone network switch? Why?

Yes, it's called xSwitch, and it's the first of its kind: an Ethernet switch custom-built for the needs of broadcasters using Livewire+ IP-Audio networks. Here's why we developed it.

Up until now, configuring the network switch has been the most tiresome part of setting up an IP-Audio network. To eliminate this, we built an Ethernet switch (using industry-standard chipsets found in high-performance, brand-name switches) and pre-configured it for use with Livewire systems. So, no configuration, no programming, and virtually no setup. Just assign an IP address, plug it into your network, and it's ready to go.

Do switches need periodic replacement as new features are introduced?

No. Ethernet is a standard, IEEE 802.3. Livewire+ gear works with any switch or router that supports the standard. You might want to upgrade to a larger or more powerful switch for some reason in the future; for example, if you were to add more studios to a cluster. Or maybe you would like to change from copper to fiber for some kind of remote uplink connection. Or you might want to replace an older switch at some point. Thanks to IEEE 802.3, the replacement switch will simply plug into the network and function, but there is certainly no requirement to do this to support new product features.

Why can't I just use my favorite switch with Livewire+?

People ask us this question every so often. When clients require a third-party switch, we recommend Cisco because their reliability, feature sets and performance are the best we've found. They also offer a wide range of switches at all price points to meet individual users' needs.

How come other manufacturers' switches don't measure up? This is mostly due to individual manufacturers' differing implementation of the same “standards”. For file transfers and e-mail, these differences are immaterial. But for VoIP and, most especially, AoIP, these implementations become more important. It's quite possible for a given switch to “work” with just a few xNodes attached, but when a more robust test is applied, that same switch can fail.

For example, we found a nice, inexpensive switch from a well-known manufacturer which, on paper, met all specs and worked with small systems in the lab. However, it turned out to not actually meet its own published specs when deployed for rigorous service in a large Livewire+ system.

For this reason, we recommend and qualify only selected Cisco switches for clients whose network designs require a large switch fabric. And for more moderately-sized networks, there are Telos Alliance-designed switches built into our PowerStation, QOR and xSwitch products — which employ some of the very same high-performance chipsets found in Cisco switches.

Do I have to be an IT expert to run a Livewire+ system?

We're engineers, and we like to talk tech. Sometimes, we talk about the tech more than we need to! But in this respect, an IP-Audio network is like a car: you don't have to understand how the engine works in order to drive it. Just connect two pieces of gear together with CAT-5e and they will talk to each other — like plugging a mic into a mixer. We've even created an Ethernet switch just for Livewire+ – xSwitch – that requires zero programming. Plug it in, assign an IP address, connect audio devices and start passing audio, with no other setup required. The Livewire+ protocol takes care of routing the audio without any need for intervention from you. And the equipment interface is all web-based with GUI control. It works intuitively, and you don't have to know anything about the tech inside to make it work.

That being said, another of the advantages of Ethernet and IP is that bookshop shelves are full of well-written books that can explain any aspect of standards-based networking at any level of detail you might want.

What levels of redundancy are available in Livewire+ systems?

We let you choose the amount of power and network redundancy your plant demands. If you like you can design a system with power and network connection redundancy from the core, all the way to the edge – you decide what's best for you.

For instance, our StudioEngine mixing engine includes standard redundant, auto-switching power supplies that are field-replaceable in seconds. PowerStation and QOR.32 mixing engines feature auto-switching backup power as an option. xNode audio interfaces can be connected to both AC and PoE-enabled network switches, for automatic failover. xNodes and xSwitch components feature twin network interfaces, to protect the integrity of network links. And we're glad to help you optimize your network switch fabric for maximum redundancy, too; just ask us and our Support experts will happily assist in your network planning process.

How do analog sources become part of the network?

With Livewire+ xNodes. These come in variants for line and microphone applications.

What about AES?

There is also an xNode that interfaces your AES audio to the network. This is a direct bit-to-bit procedure with no conversion of any kind. You can also sync a Livewire+ system to an AES master clock.

Is there a way to import audio from SDI links as well?

We have an xNode that interfaces the audio from SDI streams, both HD- and SD-, to the network. Livewire+ systems can be locked to the reference clock of an SDI input stream, which allows for handling coded audio without destroying the audio. Video compensating delay is included to ensure that the video and audio are synchronized at the outputs.

Do I have to use CAT-6 to connect everything? That could get pricey.

CAT-6 is used only for heavy-traffic network segments, like connecting one studio to another, or connecting switches to each other. All other equipment is connected with common, inexpensive CAT-5e cable.

You said I can get RS-232 data through the system. How is that done?

Using 3rd party devices, such as from Lantronics, your serial data can go anywhere across the network and be used where it's needed.

Most companies recommend that I bring them on-site to help install and configure their systems. Do I need your help to install a Livewire+ system?

With those other guys, you'd better hire their systems engineers. With us, it's much easier! If you know how to use a Web browser and plug a telephone into the wall, you've got all the skills needed to install and configure your new network. And Telos Alliance 24/7 Technical Support is there to help if you need it, too, around the clock, every single day of the year.

If you still decide you'd like on-site installation services, we have those available as well and will be happy to talk with you about it.

I've been hearing about IPv6 a lot. Will this affect how Livewire+ networks operate?

Internet Protocol version 4 (IPv4) is the IP addressing structure that powers the Internet. The Internet Engineering Task Force (IETF) has designated IPv6 as the successor to version 4, due to its larger address space, which offers more flexibility in routing traffic and allocating addresses, providing hundreds of billions more possible IP addresses than IPv4.

There are no compatibility issues between products using IPv4 or IPv6 addressing. Both addressing schemes can co-exist on the same network and interoperate smoothly. Axia systems employ Multicast streaming for audio routing. This is fully developed within IPv4, but it is not widely used under IPv6. For this reason, Livewire+ devices will continue to use IPv4 until such time as IPv6 will provide the same consistent performance. We monitor the technology standards very closely and we plan to move to IPv6 when Multicasting is widely implemented.

There is one case in which IPv4 could be a limitation: If you plan to have several hundred billion devices on your network. If you have plans to build such a large network, call us. We'll gladly implement IPv6 for you right away!

Ethernet Media

Are optical audio links supported?

Livewire+ is fully compatible with copper and fiber connection types. A common configuration is for switches dedicated to studios to be connected with copper connecting nodes, engines, surfaces, etc. A fiber backbone can then connect the switches in order to share audio among the studios.

What Ethernet rates do you support?

Switch-to-switch links may be any supported Ethernet media, Gigabit 1000BASE-T recommended. Our processing engines use 1000BASE-T. PCs may use Gigabit 1000Mbps or 100Mbps, copper or fiber. xNodes and other devices connect with copper 100BASE-T links. Media converters allow the use of fiber on nodes, such as for extended-range snake applications.

What happens if someone accidentally unplugs a cable?

Livewire+ nodes “advertise” the presence of their audio streams to the entire Livewire+ network. So if someone unplugs a node, the sources attached to it will be offline. But all you have to do is plug that node back in, and the node will “advertise” that the audio streams are available again. Within 10 seconds, all destinations that need those sources will be back up and running.

Additionally, among the many features of our PathfinderPC routing control software is a silence detect function that can be programmed to switch to another feed should one stop working for any reason.

The Internet and Livewire+

What about hooking up over the internet? With my studio audio in IP form, can I just plug a port from the switch into an Internet router?

As Internet bandwidth becomes ever more plentiful, arguments for using it for audio transmission become more convincing. A gateway device, like the Telos Zephyr iPort PLUS, can perform compression from Livewire+ PCM to a lower-rate bitstream using an MPEG AAC codec.

The main problem to be overcome is the Internet’s lack of any Quality of Service guarantees; a “net storm” that starves bandwidth and drops audio might not be a big deal to a kid at home with his computer, but it sure wouldn’t be good for an important on-air feed. Private networks with reserved capacity are one answer. Another could be the “resource reservation” and “differentiated services” technologies that has passed out of the laboratory and might eventually be implemented by Internet Service Providers – at a cost, of course.

For now, “nailed up” IP links which support QoS offer us the most reliable audio transmission. Our Zephyr, Z/IP ONE and Zephyr iPort PLUS codecs all support direct Livewire+ connection, so you can use them to get a remote link into your Livewire+ network that way and effectively have the same result.

Can a Livewire+ network can catch viruses?

There are no general purpose operating systems in Livewire+ devices, so the answer is “No.” You can keep computers attached to your Livewire+ audio network safe by keeping it isolated from data networks.

Networked Consoles

How many mixing consoles can a single mixing engine handle?

Each mixing engine provides a huge amount of power and flexibility to each control surface, so it’s a one-to-one ratio.

The exception is our DESQ and RAQ utility mixers, which are designed for less demanding work in locations such as VO booths, road kits, and personal studios. Our QOR mixing engines can each support two DESQ or RAQ consoles, or one of each.

Just how powerful are your mixing engines?

The amount of processing power in modern CPUs is staggering. With optimized software design, a single Core chipset can outperform the largest multi-DSP consoles and routers.

Additionally, we’re running a pared-down and highly-optimized version of the Linux operating system and our engine processing application code is carefully “written to the metal”, making our mixing engines very powerful indeed.

For instance, voice dynamics are a standard feature of all Axia consoles, a feature embedded in each mixing engine — and it was developed by Omnia Audio’s Frank Foti, who knows a little something about audio dynamics processing. Other standard features include per-source EQ and panning, headphone EQ and a one-touch off-air recording function, all of which can be set and saved to instantly recall each jock’s favorite settings. Fusion, Element, iQ and Radius consoles all feature four stereo program buses, and four stereo aux sends with two stereo aux returns for off-air production, and all provide automatic mix-minus for any and all phone callers or codec inputs. That’s a lot of processing, and Axia mixing engines are powerful enough to handle it all.

I like Axia consoles’ features and design, but I’m not ready to commit to Livewire+ for my full facility. Can I just use one surface and a mixing engine as a drop-in console replacement?

Sure, you can. Take a Fusion console and PowerStation mixing engine, for example, plug in your audio I/O and you have a stand-alone console that interfaces via analog or AES to your other equipment.

Analog Audio & AES On RJs And CAT-5

You recommend an outer shield for analog audio. Why?

As a precaution. Shielded cable protects against RF and eliminates any possible crosstalk between cables in multi-cable bundles.

Is there any crosstalk between the pairs within the CAT-5 cable?

As long as your circuits are balanced, there is almost no left/right crosstalk inside the cable. With a balanced input circuit that has 50 dB CMRR (Common Mode Rejection Ratio), separation will be greater than 90 dB.

So, must all the audio and digital signals be balanced?

Generally, yes, or crosstalk will degrade. Unbalanced connections can be used for short runs only and preferably with separate cables for left and right if you care very much about stereo crosstalk. Radio Systems makes small devices that adapt unbalanced RCAs to balanced RJs for their StudioHub system that could be used to convert any unbalanced sources you have. AES3 digital audio signals are always balanced and require no conditioning.

Is CAT-5 OK for AES3 digital audio?

A 1997 report, Review of Cables for AES/EBU Digital Audio Signals, conducted by the BBC Research and Development Department, concluded that CAT-5 shielded twisted audio pair cable “offered the highest performance of all the cables tested here.” Their tests included coaxial cables and special cables specifically designed for digital audio. They preferred CAT-5 cables for their consistent performance and because they have the flexibility to support other signal formats.

CAT-5 cables are engineered for data rates up to 100 Mbps to support networks such as 100BASE-T. Since AES3 signals have a bandwidth of about 3 Mb/sec (depending on sample rate), AES3’s requirements are well within the CAT-5’s guaranteed performance parameters. Dependable error-free transmission is possible at lengths up to 920 meters (over 600 ft). CAT-5 cables perform well for AES3 because they are engineered to have characteristic impedance of 110 ohms and extremely low capacitance – in the 12 pF/ft range. This yields low signal reflection and excellent high frequency response, thus lowest error rates.

The information presented above refers to 100-ohm balanced AES connections used in radio plants. In TV-land, most AES connections will be unbalanced, 75-ohm signals, and use the same type of cable as SDI audio.

Is CAT-5 OK for analog audio?

Sure, it is! The low capacitance, needed for data's high velocity and wide bandwidths, yield exceptionally flat analog audio frequency response, even over very long cable lengths. The tight, controlled twists are good for hum and crosstalk rejection. Steve Lampen, a senior audio video specialist for Belden Wire & Cable writes, "Digital cables make the absolute best analog cables. You can go farther with flatter frequency response than with any cable designed for analog". Steve Lampen of Belden graciously wrote us a paper on the subject entitled "The Axia Guide to Choosing Category Cable"; you can get it from the Telos Alliance website.

Also see Belden's web site for more interesting and revealing papers on the subject of using CAT-5 and CAT-6 cables for analog signals.

10 Resources

Networking is a field well covered by books and web sites. There's plenty of information out there. Here is a selection of some resources we've found useful.

Livewire+/Broadcast

The Telos Alliance – www.TelosAlliance.com

Contact us at: inquiry@TelosAlliance.com or by phone at +1 216 241.7225

Steve Church & Skip Pizzi, Audio Over IP: Building Pro AoIP Systems with Livewire™, Focal Press, 2010

The Bible of Livewire, written by its co-creator and Telos founder Steve Church, with NAB Senior Director, New Media Technologies Skip Pizzi. Pretty much the definitive resource for Livewire+, with chapters covering everything from basic network topology to advanced control protocols.

Ethernet

IEEE - www.ieee.org

The standards body for Ethernet. The documents are now a free download, but will cost you a lot of paper and toner – the basic Ethernet standard is 1,268 pages!

Charles E. Spurgeon, Ethernet: The Definitive Guide, 2nd Edition; O'Reilly & Associates, 2014

www.ethermanage.com

Living up to its title, it is pretty definitive on basic Ethernet topics.

General Networking and Interest

IETF (Internet Engineering Task Force) - www.ietf.org

The Internet's main standards organization. Look for the RFC (Requests For Comment) documents to see in detail how the Internet is built.

Andrew Tannenbaum, Computer Networks; Pearson Education/Prentice Hall, 2003

Our favorite general networking book. Popular college textbook covers it all, including multimedia, with a breezy style and at just the right level of detail: enough to be useful, but not so much as to be overwhelming.

J. Naughton, A Brief History of the Future; Overlook Press, 2000

Not really so interesting for audio and Ethernet, but still worth reading for perspective. This history of the Internet tells how it happened in a friendly – even charming – way. Lots of stories and anecdotes. We particularly love AT&T's repeatedly making clear that digital communication had no future. (Something a lot like what we heard from certain quarters regarding the future of computer networks for studio audio.)

Cabling Information & Standards

Cabling Installation & Maintenance - www.cablinginstall.com

This magazine, targeted to cabling contractors, is a good way to keep abreast of the latest TIA/EIA cabling specs. It is also a great source for innovative cabling accessories, testers, and installation techniques.

Jim Abruzzino; Technician's Handbook to Communications Wiring; CNC Press, Chantilly VT, 1999.

This book is concise yet contains a lot of great information including proper technique for working with CAT-5 cable and connectors. Small enough to keep with your toolbox.

Cabling Design - www.cabling-design.com — Cabling tutorials

TIA - www.tiaonline.org — Standards organization for cables

Global Engineering - www.global.ihs.com — Sells the TIA/EIA cabling standards

Cable and Contractor Supplies

AMP - www.amp.com — RJ plugs and tools

Belden Cable - www.belden.com — Leading cable supplier

Hubbell Premise Wiring - www.hubbell-premise.com — Devices for CAT-5, etc.

Panduit - www.panduit.com — Marking and installation products

Corning - www.corning.com/opcomm/ — Fiber optic cabling and components

Siemon - www.siemon.com — Punch blocks

Cable Testers

Fluke - www.flukenetworks.com — Full range of testers

Keysight - www.keysight.com — Top-end tester formerly made by Agilent

Triplett / ByteBrothers - www.triplett.com/byte-brothers/ — Low-end testers

JDSU - www.jdsu.com — Cable testers, Ethernet assurance and fancy sniffers, too

Ethernet Switch Vendors

Cisco Systems - www.cisco.com

Network "Sniffers"

Wireshark - www.wireshark.com

Ethernet Radio Equipment

Here's a list of some units we or our clients have successfully used

Cambium Networks PTP Series - www.cambiumnetworks.com

Carlson Wireless - www.carlsonwireless.com (various products)

Dragonwave Horizon Quantum - www.dragonwaveinc.com

Exalt Communications - www.exaltcom.com (EXr Series, and others)

Harris - rf.harris.com – RF7800w

Routerboard - www.routerboard.com (Formerly Mikrotik, SXT Series, for 6 channels or less)

RAD - www.radusa.com (Airmux-400)

SAF Tehnika - www.saftehnika.com (various products)

Ubiquiti - www.ubnt.com (airMAX NanoBridge M series)

Appendix A: Livewire+ Tech Details

You don't need to read any of this unless you want to know about the internal details.

Livewire+ Packet Structures

The speed of the link, the bit requirements of the header and payload, and the number of samples that are combined into a packet determine link capacity. The facts are the more samples that are combined, the less the header overhead per packet, and the higher the efficiency and capacity.

There is a fundamental tradeoff: When we have more samples per packet, we have more capacity – but at the expense of more delay. Good design means finding the best compromise.

The sampling rate and the number of samples that are combined into a packet determine delay:

$$\diamond \text{Packet-time} = 1/\text{sampling-rate} * \text{samples-per-packet}.$$

There is one packet send buffering; three packets receive buffering, and the switch latency, therefore:

$$\diamond \text{Link-delay} = \text{packet-time} * 4 + \text{switch latency}$$

Standard Streams

Standard Streams are compatible with Internet standards. They use large packets so as to be very efficient with both computer resources and network bandwidth.

Standard Stereo Stream Packet Format

Function	Bytes	Notes
Interpacket Delay	12	This is not actually transmitted, but must be taken into account for network bandwidth calculations
Ethernet Header	26	Includes the VLAN/priority fields
IP Header	20	Standard
UDP Header	8	Standard
RTP Header	12	Standard
Audio	1440	240 samples @ 48kHz, 24-bit, stereo
Audio (variant)	720	120 samples @ 48kHz, 24-bit, stereo

Total bytes per packet = 1440. Core delay = 5ms. (720 and 2.5ms with the variant format)

An Ethernet frame's maximum data length is 1500 bytes, so you can see that we have chosen to pack the Ethernet frame to nearly the maximum possible. There are two reasons for this: 1) the frame rate is lowest possible to put the least burden on PC receivers, 2) the header overhead is applied to the most data so the proportion of capacity devoted to audio vs. overhead is highest.

Livestream

Livestreams are specialized for low delay, so we can pack only a few audio samples into each packet. Because they are smaller, less buffering is needed and that means the time delay is lower.

Livestream IP Packet Format

Function	Bytes	Notes
Interpacket Delay	12	This is not actually transmitted, but must be taken into account for network bandwidth calculations
Ethernet Header	30	Includes the VLAN/priority fields
IP Header	20	Standard
UDP Header	8	Standard
RTP Header	12	Standard
Audio	72	12 samples @ 48kHz, 24-bit, stereo

Total bytes per packet = 72. Core delay = .25ms.

The header load for RTP/UDP/IP is 40 bytes per packet, which is a significant piece of the network bandwidth given that our audio is only 72 bytes. Most of the time this is of no consequence, since we have plenty of bandwidth.

However, there are some situations where a “lean and mean” approach makes sense.

Surround Streams

By now, you realize that Livewire+ inherently carries multiple audio streams and surround mixing is a built in feature in the StudioEngine mixing engine for Fusion and Element consoles. This makes Livewire+ and Axia fully compatible with surround formats for television audio. It would also make for a very nice surface to control your personal surround sound system in your living room.

Surround Stream IP Packet Format

Function	Bytes	Notes
Interpacket Delay	12	This is not actually transmitted, but must be taken into account for network bandwidth calculations
Ethernet Header	30	Includes the VLAN/priority fields
IP Header	20	Standard
UDP Header	8	Standard
RTP Header	12	Standard
Audio	1440	60 samples @ 48kHz, 24-bit, stereo

Total bytes per packet = 1440. Core delay = 1.25ms.

Surround streams are actually multiple stereo streams carrying the 5.1-channel audio plus a 2-channel downmix simultaneously. Surround streams carry 8 speaker channels in the following order: Left front, Right front, Center, LFE (low frequency enhancement), Left Surround, Right Surround, stereo left, and stereo right.

Network Link Capacity

Each Standard Stereo Stream has a bitrate of 2.304Mbps. A 100Mbps link can therefore carry 43 such channels at full capacity and a 1000Mbps link can carry 430 channels.

Each Livestream has a bitrate of 3.776Mbps. A 100Mbps link can therefore carry 26 such channels at full capacity and a 1000Mbps link can carry 260 channels.

In practice, links to hardware nodes will have a mix of Standard Stereo Streams, Livestreams, and control data. Our biggest xNnode has 4 stereo (8 mono) channels in and 4 stereo (8 mono) channels out, so there is plenty of link capacity. PCs use the more efficient Standard Stereo Streams, so again there is plenty of capacity to handle both audio and simultaneous file transfers, etc. Our Studio Mix Engines connect with 1000Mbps links, so the sky is the limit there.

All of the above has been concerned with per-link bandwidth. The system bandwidth is effectively unlimited with appropriate switches.

Multicast Address Translation

Livewire+ streams are multicast at both layer 2 and layer 3. The Livewire+ channel number is automatically translated to the appropriate addresses at both layers internally. You might want to know the translation algorithm because maybe you or your network engineer might need to check packets with a “sniffer” or Ethernet switch diagnostics. So here are the details.

Livewire+ channels range from 0 to 32767. Audio streams are mapped into IP and Ethernet multicast addresses using the channel numbers for the lower 15 bits as follows:

IP Address	Type
239.192.000.0/15	Livestream and Standard Stereo Streams
239.192.128.0/15	4 addresses are our system defaults, the others not used (left for expansion)
239.193.000.0/15	Back Standard and Livestream Stereo Streams
239.193.128.0/15	Not used (left for expansion)
239.194.000.0/15	Not used (Livewire v1.0 used for Ethernet Livestreams)
239.194.128.0/15	Not used (left for expansion)
239.195.128.0/15	Not used (left for expansion)
239.196.128.0/15	Surround Streams
239.251.000.0/15	Not used (left for expansion)
239.251.128.0/15	Not used (left for expansion)

The following special addresses are assigned:

IP Address	Function
239.192.255.1	Livestream Clock
239.192.255.2	Standard Stream Clock
239.192.255.3	Advertisement Channel
239.192.255.4	GPIO (UDP port 2060)

These all are within the range specified for “Organization-Local Scope” use by IANA – the Internet Assigned Names and numbers Authority. Routers do not propagate traffic on these addresses to the Internet; they stay contained within LANs. (We also set the “link local” bit and TTL=1 in the IP header to further ensure that streams stay local.)

The range supports our 32k channels, with up to 120 stream types per channel. We are only using four types now, but there is plenty of room for growth.

Our motivation for mapping each type to a contiguous block rather than having the type in the lower-order bits is to allow configuration of switches and routers on a per-type basis by specifying an address range. This direct mapping of channels to addresses also makes sniffing easier: it is simple to know where to look for a particular audio stream.

IP addresses are mapped into an Ethernet MAC layer multicast, according to a de-facto standard process for this procedure. This process is as follows:

- ◆ Using the Class D address, identify the low order 23 bits of the class D address.
- ◆ Map those 23 bits into the low order 23 bits of a MAC address with the fixed high order 25 bits of the IEEE multicast addressing space prefixed by 01-00-5E.

Example:

- ◆ Assume: Channel = 80
- ◆ Assume: stream type is Standard Stereo Stream
- ◆ Then: IP address = 239.192.0.80 (dotted decimal)
- ◆ And then: Ethernet MAC Address = 01-00-5e-00-00-50 (dashed hex)

Ethernet addresses are transmitted most-significant byte first, but least-significant bit first within the byte, so in our example it is the 1 in the leftmost MAC address byte 01 that signifies a multicast address.

