ATTEROTECH

RESIDENTIAL AUDIO

The CobraNet Way

OVERVIEW

As the size and capability of whole-home audio systems has increased, so has the wiring infrastructure needed to support these systems. With many zones of A/V content, hundreds of feet of specialized cabling must be installed in a home to move control and metadata, audio, and more recently, video. This jumble of analog wiring is expensive, both in material and installation costs. Moreover, the wiring in every installation is different, based on the specific needs of the end user.

Wouldn't it be a huge advantage to be able to route control and metadata, audio, and video over a single, inexpensive, and standard wiring infrastructure? This is no longer a dream, but a reality with CobraNet®. The following will explain the advantages of using CobraNet as the basis for a flexible, expandable, and robust A/V control, routing, and distribution system for residential applications.

Three appendices are included in this white paper. Appendix A gives a short introduction to CobraNet technology, and some insight into the signal routing capability available from this powerful protocol. Appendix B addresses some of the issues involved in running CobraNet data concurrently with other network data on a home LAN. Appendix C identifies key issues surrounding the selection of a network technology for residential A/V systems.

WHITE PAPER

DEFINITIONS

CobraNet
A digital media transport protocol to move uncompressed audio and compressed video data over a standard 100BaseT Ethernet network.

Legacy Source
Any source that produces a pure audio output only such as a CD player. Typically, legacy sources support only IR for basic control.

Power Over Ethernet (PoE)
A technology to power remote Ethernet devices through their Ethernet connection using the CAT-5 cable.

Ethernet Hub
An unintelligent network device used to connect multiple Ethernet devices together. Any traffic received from one device is passed to all other devices connected. This generates large amounts of network traffic and is inefficient.

Ethernet Switch
An intelligent network device similar to a hub used to connect multiple Ethernet devices together. Unlike the hub, a switch learns where devices are. When it receives data, it only sends it on to a device or devices it knows need it. It consequently makes more efficient use of the bandwidth than a hub.

Ethernet Router
A router is a networking device that allows separate networks to be connected together. Unlike a switch, a router has 2 domains. The LAN domain is used for local devices. The WAN side connects to an external network. Such a device typically includes a switch on the LAN side for multiple device connection and many contain a firewall to prevent unauthorized access from the WAN domain into the LAN domain.

THE TRADITIONAL SETUP

With multiple audio sources available, from CD changers through internet and satellite radio, to MP3 players, audio is available everywhere. Many homes have one or more of these sources and aggregate them into a whole-home audio system.

A typical analog system, like the one shown on the next page, is built around a central control unit. This unit has inputs for multiple sources and speaker outputs for multiple zones. Control is provided by control interfaces (wired or wireless) in each zone that pass control information back to the central unit. Metadata flows from the source through the central control unit to the keypads, although for legacy sources metadata won't be available.

This type of system has some disadvantages.

- **Cabling** – Multiple cable types are needed. Audio cables for the sources, speaker cables for the speaker connections, and control cables for the control pads. Speaker cabling, in particular, can be expensive especially if you consider the potential distance between the central control unit and the speaker which may be up to 100 feet or more. These cables all need installing, which adds to the cost. If

video signals are also distributed by the same system, coax or CAT-5 dedicated to analog video is required.

- **Expandability** – Expansion is limited to the hardware in the central unit, which has a specific number of source inputs and a specific number of zone outputs. Expansion of a full system cannot be done incrementally, instead requiring the addition of a complete second central unit. The end user has to pay significant incremental costs for what may be only a small upgrade, and is paying for resources which may never be used.

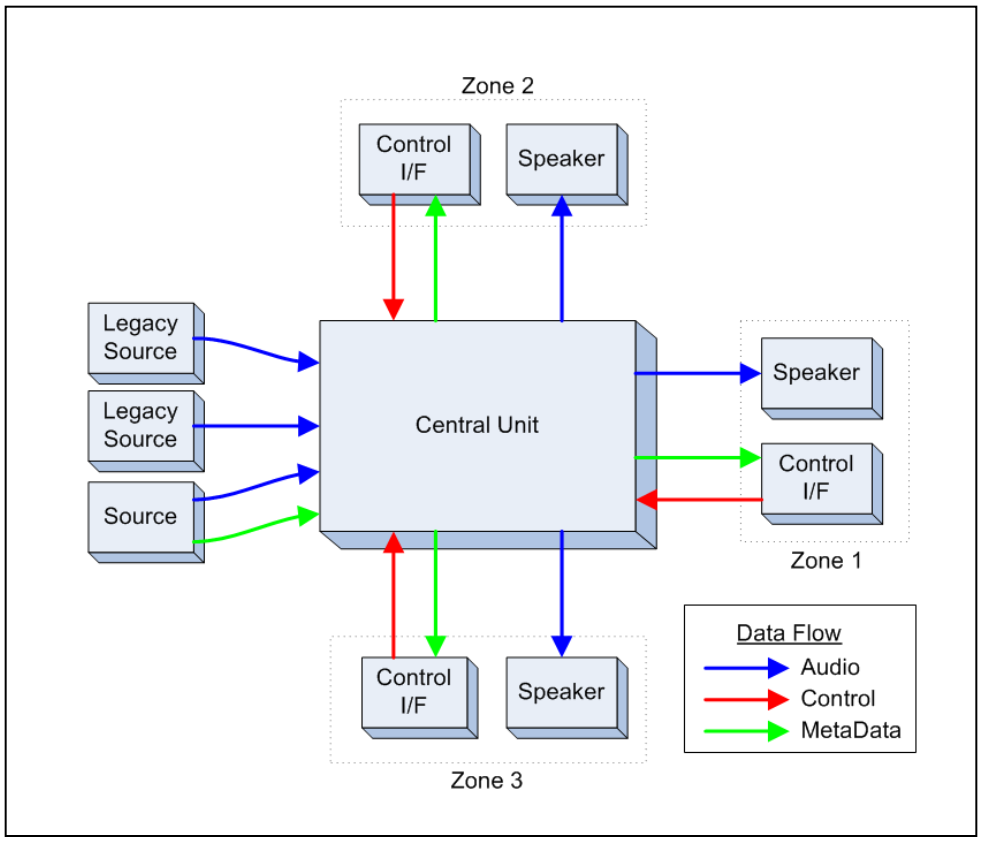

- **Audio Quality** – With long runs of analog cable, there is a great possibility for audio to become degraded with unwanted noise by the time it reaches the output.
- **Flexibility** – Sources have to be connected to the central unit. To keep the audio cables short to ensure good audio quality, sources should to be placed local to the central control unit. Longer cables and more remote sources can be used, but this increases the potential for audio problems. This limits the flexibility of the system.
- **Power** – Long cable runs also affect the power that can be supplied to system loudspeakers. As a cable increases in length, the power losses in the cable increase, reducing the power available to the speakers.

GOING DIGITAL

The analog system can be improved upon by moving to a digital distribution system using CobraNet. CobraNet is an audio transport system originally designed to transfer uncompressed audio in real time across a network with low and fixed delay. But why use it over traditional analog?

CobraNet requires no special infrastructure. It uses standard Ethernet switches and networking components found in all IT systems. Infrastructure costs are kept low. There is also the possibility to use Power over Ethernet (PoE) to power remote devices, as well centralizing some of the power requirements. Many homes have already been wired with CAT-5 cable for home networking and internet sharing. CobraNet can co-exist with existing PC Ethernet traffic, so a CobraNet digital system can piggy back onto the existing cabling (see Appendix B for more detailed information on this scenario).

With CobraNet, audio and video can be delivered digitally right to the endpoint in each zone. This feature removes the need for long analog cable runs. This not only cuts installation costs, but also boosts audio quality as well. There's also no worry about power drops over long cables with the digital system.

One of CobraNet's prime features, being a network-based protocol, is its audio and video routing capabilities. Audio and video sent via CobraNet can easily be routed to new endpoints just by changing a few settings and without the need to alter any of the underlying hardware or cabling. The physical inputs and outputs available on the Central unit no longer limit system expansion. Incremental upgrades thus become a reality. Additional inputs for sources or outputs for new zones can now be added simply and at a lower overall cost to the end user.

CobraNet is also full duplex, meaning an endpoint can both receive and transmit CobraNet data at the same time. Using this functionality, we can cut down on overall hardware costs. A single endpoint can now support multiple sources and act as an output device. More importantly, it can do both simultaneously, and all through just one CAT-5 cable. For example, you may want a CD player in a particular room. The audio output can be fed into the system for distribution to other parts of the house. At the same time, the output from that same endpoint can be used to listen to any audio source including that same CD player.

CobraNet has some additional features that can be used. In addition to incorporating CobraNet for audio and video transport, it can be used for control as well. CobraNet has the ability to transport serial messages and/or IR control commands over the network in addition to audio and video. This removes the need for a separate control bus and all interfaces become part of the CobraNet network: endpoints and control interfaces alike.

By adding a small and low cost microprocessor to the CobraNet interface, there is potential for a truly flexible endpoint solution capable of audio output, source input, and control as well as auxiliary control for lights, screens, curtains, and virtually anything else. There is no reason why an endpoint has to contain all these features. Instead a selection of endpoints could be designed to suit different situations.

Digital audio has some great advantages over analog but how would such a digital system work in practice? Following are three possible strategies to implement a digital audio system. Each has advantages and disadvantages when compared to one another. The right

one to use will depend on the type of system you wish to provide and the functionality required.

THE ANALOG REPLACEMENT SYSTEM

Consider a simple system where analog connections have been replaced by CobraNet.

This type of system is ideal if the end user has only legacy sources such as CD players and AM/FM or satellite radio. These can be connected locally to the central unit or connected remotely via a CobraNet connection. The central server still has limited hardware but there is room for some expandability. This solution also requires an Ethernet switch to connect all the devices. This switch can be embedded into the central unit as shown here or be a completely separate unit mounted in the rack.

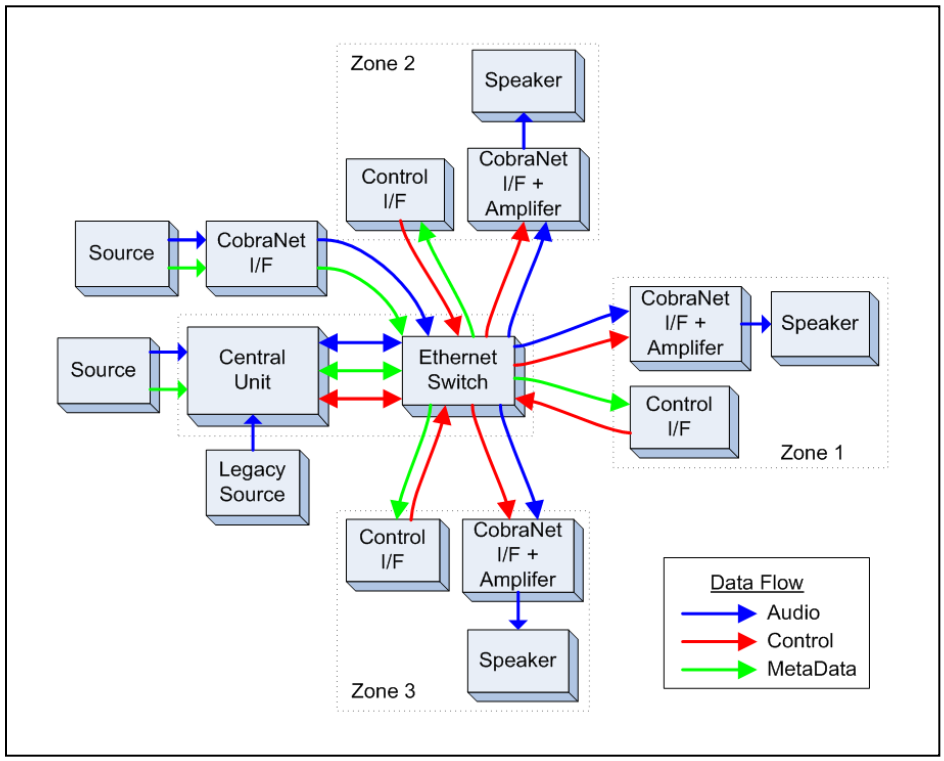

Control of the system is maintained by the central unit with all remote interfaces passing commands back to it for processing. The remote interfaces may be wired or wireless. The control bus could be separate or part of the CobraNet system. Simple control of remote devices such as stop, play, pause, next or previous track, next or previous disc (in the case of a CD multichanger) can be achieved by using a remote IR emitter. Such a system also has flexible zoning but control over those zones would only need to be rudimentary.

Characteristics of this type of CobraNet system include:

- Low endpoint costs
- Requires a central unit
- Requires an installer to build system
- Some expandability
- Some flexibility

A MORE FLEXIBLE, SCALABLE SOLUTION

The second option is significantly more flexible. Here we can have widely distributed sources and endpoints. Utilizing the full-duplex capability, an endpoint can not only receive audio but transmit it as well. This means each endpoint can act as an output, a source, or indeed, both at the same time. Audio sources can then be plugged into the system virtually anywhere.

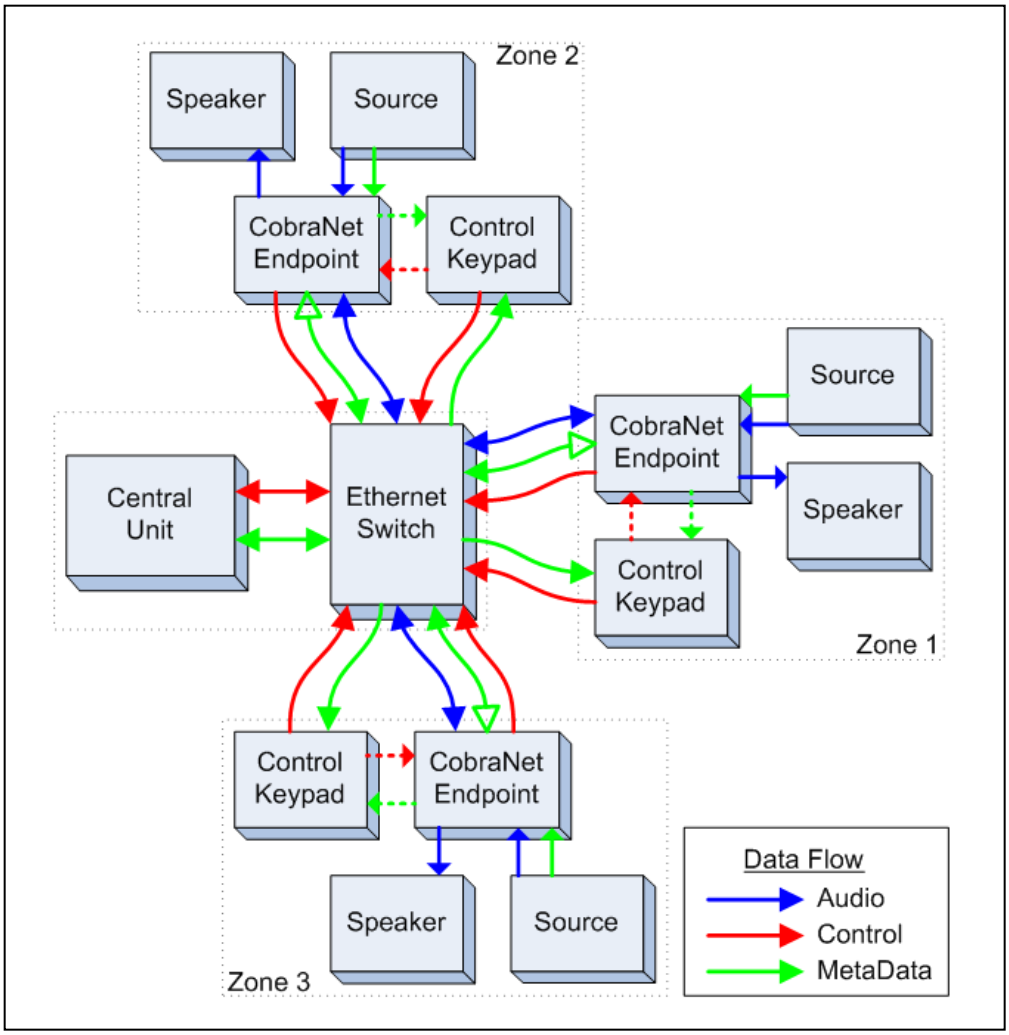

In this system, each endpoint would require some control and setup as it could be treated as a source and/or an output. The endpoint would thus have to contain a low-cost microprocessor and some intelligence. This adds significant functionality. The control system can be integrated into the CobraNet system. Local control would no longer need to be through the Central unit, but direct with the endpoint concerned although overall system control is still retained by the central unit. Additional reductions in wiring are also possible by using a direct link from the control keypad to the endpoint instead of directly back to the Ethernet switch.

Characteristics of this type of CobraNet system include:

- Low to medium endpoint costs – depending on the functionality (LCD display, touch screen versus hard buttons, etc.)

- Requires a central unit – the central unit may optionally have audio and video inputs for legacy sources
- Requires an installer to build system
- Maximum expandability
- Maximum flexibility

CONSUMER ORIENTED, FOR SMALLER RESIDENCES AND APARTMENTS

Both the previous examples would need professional installation. What about a more consumer oriented system? Such a system would need to be simple to connect and configure, whilst maintaining functionality.

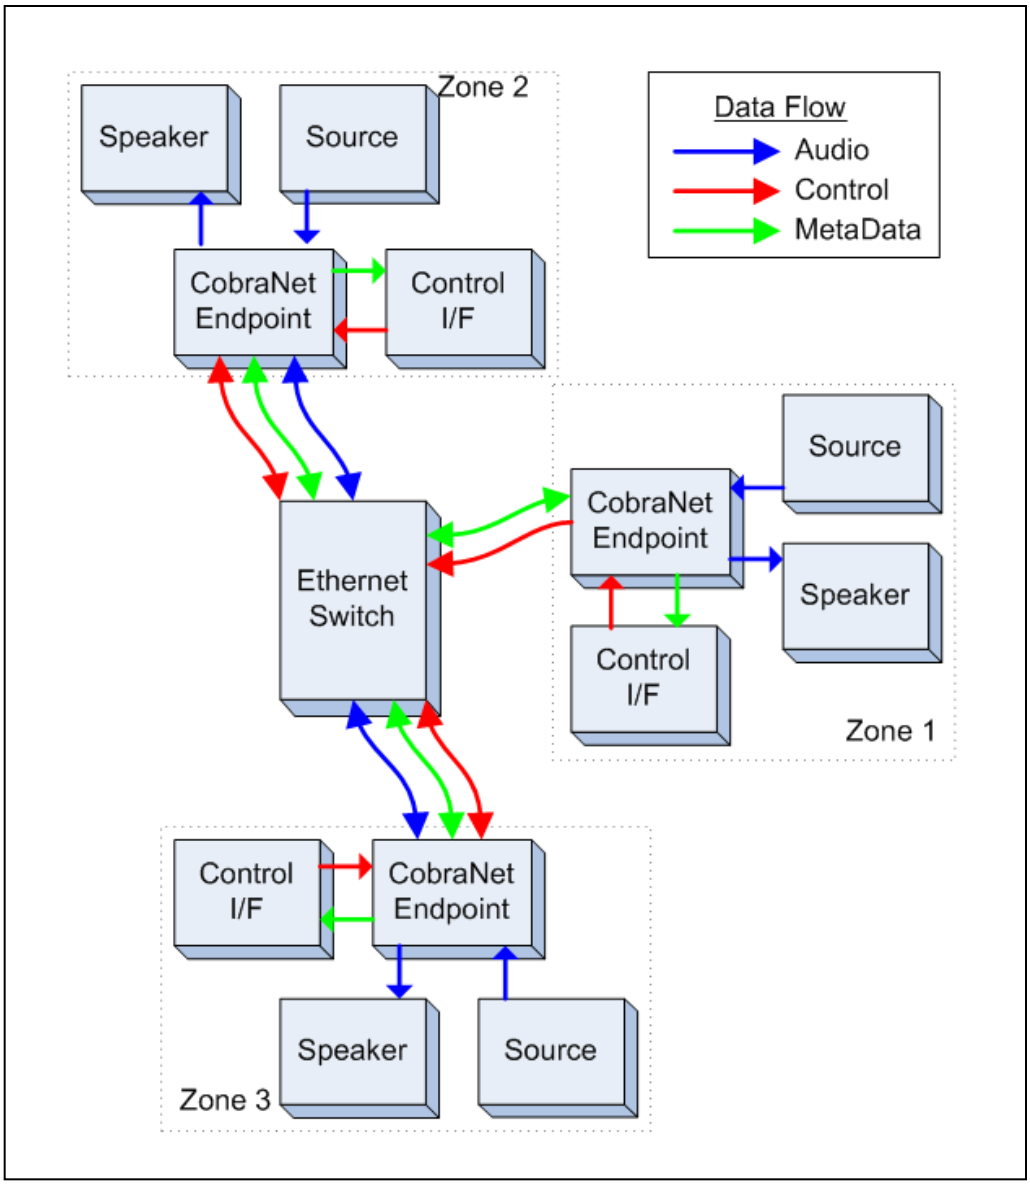

The main difference between this and the previous strategies is that there is no longer a central control unit. All the intelligence has been moved into the endpoint and each can function on its own.

Such a system would be limited to only a few endpoints to ensure overall system configuration could remain simple. However, the endpoints themselves may well include some very advanced features. Endpoints could be sources incorporating connections and control for MP3 players, media servers and internet radio. They could also act as the outputs too.

Each endpoint has local control with very little overall system control required. Whilst the system would be limited to the capabilities of the endpoint, expansion is limited to ensure the system control required can be kept to a minimum.

Characteristics of this type of CobraNet system include:

- High endpoint cost
- No central unit costs
- Installation possible by consumer, assuming CAT-5 is already installed in the residence
- Limited expansion – best for installations of a few zones
- Limited flexibility – overall system control typically limited to independent A/V content in each zone or “party mode” with overall volume controllable at one or more endpoints.

APPENDIX A

INTRODUCTION TO COBRANET

CobraNet is an audio networking technology for delivery and distribution of real-time, high quality, uncompressed digital audio using a standard Ethernet network. It is implemented using a combination of hardware, firmware, and the CobraNet protocol.

NOTE

While this appendix refers to CobraNet as an audio networking technology, Attero Tech offers technology that can deliver video content over CobraNet networks as well, for a complete A/V solution.

Unlike other audio networking or distribution technologies, CobraNet is a true network and exists on standard Ethernet networks using standard Ethernet hardware. Since it is a true network, audio routing is highly flexible between network nodes and can be used in a variety of audio distribution applications.

In addition to the high degree of routing flexibility that CobraNet provides, the technology also incorporates the ability to monitor and control CobraNet devices remotely. This is a key feature that is highly important in fixed installation applications where the audio distribution equipment may not be readily accessible. All CobraNet devices on the network can be controlled and monitored from a central location by sending control commands and monitoring device specific parameters.

CobraNet provides this capability by implementing Simple Network Management Protocol (SNMP). SNMP is a standard protocol typically used for monitoring network devices such as Ethernet switches. In the case of CobraNet, it allows users to communicate with any CobraNet device using standard SNMP tools or a customized user interface designed specifically for CobraNet, such as Attero Tech's Control Center application.

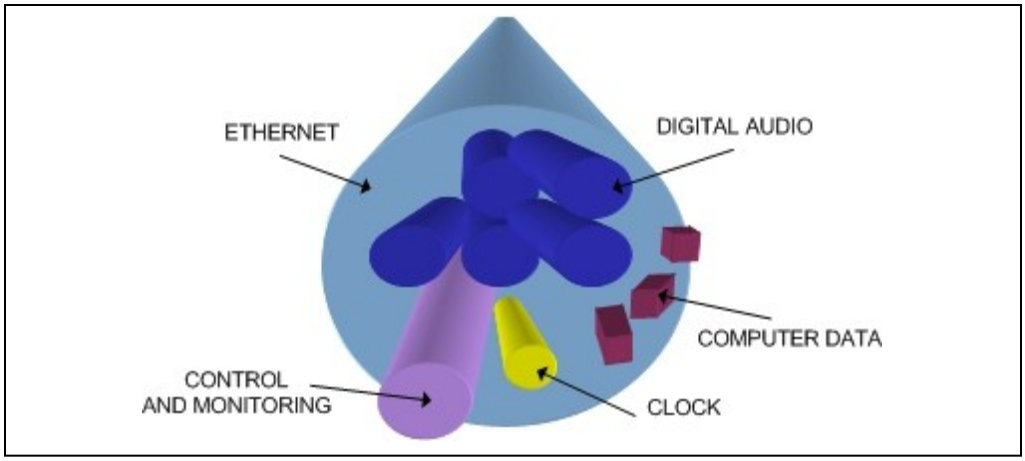

The figure above represents the types of data that coexist on a CobraNet network.

Before a CobraNet system can be configured, it is important to first understand how CobraNet distributes audio between devices.

Audio is sent in "bundles" on a CobraNet system. Each bundle is capable of holding up to 8 logical audio channels. Every CobraNet device has a number of bundle transmitters and bundle receivers. These transmitters and receivers are the mechanism used to send and receive bundles between devices.

For a transmitted bundle, audio may be sourced either directly from the local audio inputs of the device or from internal audio via the on-board DSP, but not both simultaneously. The internal audio from the onboard DSP could have originally been sourced from the local

device inputs, sent from another CobraNet device or even generated by the DSP itself. Combinations of audio may exist within a bundle in any order. Additionally, a single audio source in a device may be used multiple times in a single transmitter bundle or across multiple transmitter bundles.

For a received bundle, the received network audio may be routed directly to the device's local outputs, the internal DSP[1] or simply ignored.

Once the contents of a bundle have been decided, the next step is to pass the bundle to another CobraNet device. To do this, every CobraNet device has up to 4 bundle transmitters. Each bundle transmitter has a transmit mode that must first be selected. This affects how many devices may receive that particular bundle at a time.

The modes are as follows:

- **Unicast** – Used for one-to-one connections. In this mode, only one receiver at a time can receive this bundle. Once a link is established from this transmitted bundle to a receiver, any future requests for that bundle from other potential receivers will fail.
- **Multicast** – Used for one-to-many connections. This mode broadcasts its contents over the entire network. There is no restriction on the number of receivers. However, the downside is that CobraNet packets are distributed to all nodes on the network, whether they need them or not thus creating possible network bandwidth issues.
- **Multi-unicasts** – Another one-to-many mode. Whilst this is the most efficient method for getting a bundle to multiple receivers in terms of network bandwidth, it requires more processing power on the CobraNet device so in this mode there is a maximum limit of four receiver connections (this can be reduced if required). If more connections are required than the limit, the node can be configured to automatically switches to multicast.

NOTE

When a bundle must be transmitted to multiple receivers, multi-unicast transmissions should be used where possible.

Once the mode is selected, to enable a device to transmit the bundle, simply allocate the particular transmitter bundle a non-zero number. Since this number identifies all the network packets sent out by that transmitter, each transmit bundle number must be unique on a network.[2]

Now that the transmitter is set up, it is time to set up the receivers. In order to receive bundles, each CobraNet device has up to eight bundle receivers. To enable a device to receive a bundle, simply allocate one of that device's bundle receivers the same bundle number as a transmitted bundle. By doing so, a virtual link is created and audio should now be passed from one device to the other. It should be noted that no knowledge of a device's network topology or connection is thus required in order to configure audio connections. The only restriction to this is that a device cannot be set up to receive a bundle it is also transmitting.

[1] Not available on all devices – CS496xxx devices only

[2] Bundle numbers range from 1 through 65535. A value of 0 represents an inactive bundle. Numbers 1-255 are reserved for multicast mode transmissions only.

The above case creates a simple, one-to-one, unidirectional link. If more devices are required to receive that bundle, allocate the same transmitted bundle number to a bundle receiver on the other CobraNet devices.

It is also important to note that CobraNet supports simultaneous bidirectional audio distribution in each device. Not only could audio be sent from Device A to Device B but at the same time, should it be needed, audio could also be sent from Device B to Device A. The exact bundle and routing configuration will be determined by the needs of each individual installation. An installation may have multiple units transmitting multiple bundles. The only restriction is the bandwidth available on the network to transfer the audio.

CobraNet does more than just transfer audio/video data. It can be used to pass serial information as well. A feature called serial bridging has been incorporated that allows the passage of serial data between nodes. Each node can pass serial data to a specific node or multicast the data to multiple nodes. A node can also receive data from either a single source or multiple sources. Baud rates, data bits, stop bits, parity, and so on are all configurable. There is also support for multi-drop serial buses as well.

Finally, CobraNet has the capability to alter all of the above options in real time making the whole system completely dynamic. By use of control software, all of the bundle assignment parameters can be configured with no need to change cables, switch out connectors, or pull new wiring. Most importantly, this control capability can be implemented from a single location!

APPENDIX B

COBRANET & PC NETWORKS

Whilst CobraNet is compatible and can co-exist with standard Ethernet traffic from PCs using the same infrastructure, it isn't all plain sailing. On lightly used systems, it is likely there will not be any noticeable problems. However, as audio/video system usage increases and/or PC network traffic increases, network bandwidth quickly disappears. Eventually, there becomes a point where PC operations over the network, such as web access, will begin to slow as CobraNet data is given priority. Fortunately, such problems can be overcome.

Physically separating the networks into two separate pieces, one for CobraNet traffic, the other for normal network traffic, is the ideal solution. This is often the way practical systems are implemented, since the A/V installer has little control over the typical home LAN. Having a separate A/V cabling infrastructure is the best way to prevent nuisance callbacks, and guarantee that changes on the home LAN do not affect the A/V system.

A slightly different approach uses network devices that can separate the traffic internally by implementing virtual networks. These virtual networks ensure that the PC traffic and CobraNet traffic cannot interfere with each other and are kept completely separate even though they are travelling through the same network device. The down side is that Ethernet switches with this capability are more expensive than standard switches.

Separating the audio system from the PC network gives best performance, but what if one or more of our A/V sources is PC or internet based such as internet radio, and needs access to the PC network? Homes aren't likely to want or have a dedicated PC or a dedicated internet connection just for the audio system. In such cases, a bridging device will be needed that prevents CobraNet data getting onto the home LAN while permitting the A/V device full home LAN access. Such a device could be a standalone device or the functionality could be built into an endpoint at very little cost.

APPENDIX C

KEY CONSIDERATIONS FOR SELECTING A RESIDENTIAL A/V NETWORK INFRASTRUCTURE

With the wide variety of both open and proprietary networking technologies available, what are the key aspects to consider in selecting a networking approach for residential A/V routing, distribution, and control systems? Raw speed, latency, routing, standards-based or proprietary, physical infrastructure – how can these be sorted out for the whole-home market space?

The following points identify some of the important requirements for residential A/V systems that drive the choice of a network infrastructure.

- **Reliability** – Systems that do not work as intended every time mean service calls. Service calls erode a dealer's profits and give products a bad reputation. Either of these can be the kiss of death in the residential A/V market space.
- **Latency** – Simply means the time that audio, video, or control data takes to be transported between a source node and a sink node. Although some latency may be acceptable in some residential A/V use cases, the requirement for quick response to user control commands means that the basic network infrastructure should add as little latency as possible.
- **Quality of Service** – Gaps in video, breakup or stuttering in audio, or playback stoppage for "buffering" may be acceptable when streaming YouTube or internet radio content, but is not acceptable in residential A/V systems. Steaming media content viewed in real time puts large demands on the network infrastructure at all times (not just during peaks), so the network infrastructure must be robust under constant high demand.
- **Multiple Endpoint Synchronization** – Since very few residential A/V systems are only one zone, it is important for A/V content to be synchronized when the same content is played in two or more zones. This is particularly important when the same content is played in two adjacent zones or in party mode. While acoustic delays between zones (caused by the physical separation between speakers in different zones) are unavoidable, the network infrastructure should not add to this delay. Therefore, synchronizing the media content delivered to multiple endpoints to within a few milliseconds is a must. Note that each endpoint can implement buffering to achieve basic synchronization, but this increases both latency and endpoint cost.
- **Cost** – Even the wealthiest client will have a budget, and any network infrastructure for residential A/V applications must be cost effective for systems both large and small. In particular, since networked systems tend to be comprised mostly of endpoints (as opposed to one or more large central units), low cost endpoint BOM cost is critical to achieving an acceptable overall system cost. The ability to carry A/V media content on a 100BaseT infrastructure (as opposed to a 1GB or 10GB) also minimizes central unit cost.

The following table summarizes the several approaches and how they stack up based on the criteria above:

Type	Reliability	Latency	Quality of Service	Multi-Endpoint Sync	Cost	Notes
Standard Wired Ethernet	Good	Variable (1-10ms)	Good (unless heavily loaded)	Not supported	Low (100BaseT Medium (1000BaseT) (No standard low cost endpoint solution exists)	No standard low cost endpoint solution exists
802.11b/g Wireless	Average (susceptible to RF interference from other 2.4GHz devices)	Variable (10-100mS)	Average (load and signal strength dependent)	Not Supported	Low (802.11b/g) Medium (802.11n)	No standard low cost endpoint solution exists
Powerline	Average (susceptible to wideband noise on power lines from motors, switching power supplies, etc.)	Variable (100mS – 500mS)	Average (load and distance dependent)	May be supported in proprietary implementations	Medium	No Standard low cost endpoint solutions exists, but there are proprietary solutions
CobraNet	Good	Fixed (5.33mS)	Good (even heavily loaded)	Supported. All endpoints synced to within 20.4uS	Low (100BaseT)	CobraNet LE is a standard low cost endpoint solution. PoE (both 802.3af compatible and proprietary is available)

ABOUT ATTERO TECH

Attero Tech is a leading provider of CobraNet® audio interfaces. These products make it easy and cost effective to integrate a wide variety of audio components such as consumer electronics, microphones, paging speakers, computers, and recording devices into a networked audio system. Attero Tech solutions help AV systems integrators reduce cost, improve audio quality, future proof systems, and meet their most unique design requirements.

Attero Tech is headquartered in Fort Wayne, Indiana. For more information on Attero Tech's full line of products, please visit www.atterotech.com.