



STL Requirements and Network Recommendations for HD Radio Implementations

Richard Hinkle, VP Engineering & Technical Services, Broadcast Electronics, Inc.

Abstract

The purpose of this document is to provide engineers with information necessary to plan and implement HD Radio solutions. The nature of HD Radio signals makes it necessary to have exceptional STL connections and planned network topographies in order to provide the most reliable performance possible.

HD Radio Components

HD Radio provides broadcasters with a way to deliver programming digitally, but that's just the beginning. In addition to main program audio, the underlying design of HD Radio allows for a variety of additional advanced application services most notably the ability to deliver additional channels of programming using the same signal—an application called “multicasting”. Listeners with HD Radio-enabled receivers can listen to a station's main program service (which is delivered concurrently with the conventional analog FM signal) or adjust their receivers to listen to the secondary program services, commonly referred to as HD2 and HD3.

By the time an HD Radio signal reaches a listener's receiver, the complete payload is made up of multiple parts:

MPS (Main Program Service): This is your main HD Radio signal. This will carry the same audio as analog FM.

SPS (Secondary Program Services): These are your additional channels of audio programming, commonly called multicast channels.

PAD (Program Associated Data): Data associated with program playing on air, including artist/title information. PAD can be specific to either the Main Program Service (**MPS PAD**) or to a Secondary Program Service (**SPS PAD**).

Additionally, the HD Radio specification allows for the delivery of other types of data including information like traffic information. This white paper will deal specifically with the integration and delivery of MPS, SPS and PAD and the common questions and issues facing broadcasters as they prepare to “go HD”.

Architecture Considerations

While MPS audio can be delivered from the studio to the transmitter site by conventional means as AES serial digital audio, SPS and PAD are delivered as data only over an Ethernet connection. This transition to the world of data as opposed to audio forces engineers to at least venture into the world normally consigned to network and IT specialists. Discourses in Information technology are beyond the scope of this document, but you will need to know some basic terms to carry on a conversation:

Duplex: Bidirectional communication, used to describe a link between devices that allows for two-way back-and-forth communication.

Simplex: Unidirectional communication, used to describe a link between devices that limits data flow to one-way, send-receive communication.

UDP (User Datagram Protocol): A protocol that is part of the IP (Internet Protocol) Suite, UDP is a broadcast protocol that lacks provisions for guaranteed packet delivery or the ability to request packets to be resent. UDP communication can be bidirectional (duplex) or unidirectional (simplex). “Broadcast protocol” as described here means that the packets will be sent to all devices on a network, not just the one intended.


TCP (Transmission Control Protocol): A connection-based protocol that allows for more reliable communication between devices. TCP sends data to a destination in numbered packets, and requires the destination device to acknowledge delivery of each packet. The trade-off is in overhead and speed as TCP is slightly costlier in bandwidth and slower than UDP.

Bandwidth: In this context, bandwidth refers to data transmission rates or capacity when communicating over certain media between devices. In this paper, we will differentiate between **average bandwidth**, which is the actual average bandwidth the HD Radio payload requires as measured over a 24-hour period, and **provisioned bandwidth**, which is the maximum bandwidth peak including data, data bursts and overhead.



HD Radio Fundamentals

When HD Radio was first introduced, broadcasters were limited to MPS (Main Program Service) audio at 96 Kbps, broadcasting the same programming as the FM analog signal. The HD Radio specification also included provisions for very low 400 bps worth of MPS PAD (Main Program Service Program Associated Data), which was primarily limited to Artist/Title information.



- Fig. 1: Basic HD STL Configuration, No Multicast -

Main AES audio and this small data payload were delivered to the transmitter site where the exciter handled all the HD Radio generation. Today, almost all standard STLs can handle this implementation.

STL requirements for basic HD Radio systems are quite simple, and practically unchanged from pre-HD Radio requirements.

AES Audio Transport Capability:

- ❖ 44.1 kHz AES Audio = 1.411 Mbps (uncompressed)
- ❖ 32 kHz AES Audio = 1.024 Mbps (uncompressed)

Ethernet Data Transport Capability:

- ❖ MPS PAD – 400 bps Simplex (unidirectional) UDP

Network requirements for these original systems are practically nonexistent. In most cases the studio automation data interface connected directly to STL, but even in cases where the data interface connected to the STL using existing/shared network connections, impact on other network traffic was negligible because of the low data rate.

As HD Radio evolved and Multicasting was introduced, the structure of the HD Radio signal became more flexible and more complex. Currently, there are three modes of HD Radio operation:

- MP1: Standard Hybrid HD Mode of Operation: 96 Kbps available to partition between main and multicast channels
- MP2: Extended Hybrid HD Mode of Operation: 108 Kbps available to partition between main and multicast channels
- MP3: Extended Hybrid HD Mode of Operation: 120 Kbps available to partition between main and multicast channels

Note: *If you decide to broadcast in one of the two extended modes, you must notify the FCC that you are operating in MP2 or MP3 modes because it is considered experimental. Notification can be made in form of informal letter to the commission.*

There are two different ways to implement Multicasting: I2E (Importer to Exporter/Exciter) and E2X (Exporter to Exgine), the primary difference being the location of the individual hardware components. Link and network requirements are different for each approach, with benefits and tradeoffs for each.

Multicasting – I2E (HD Generation at Transmitter Site)

When multicasting was first introduced, engineers were working within a framework that placed all HD Radio generation equipment at the transmitter site. To allow for multicasting, a piece of equipment called an Importer was added to the system which required a duplex (bidirectional) link. The introduction of multicasting increased the flexibility of HD Radio, but also increased the complexity.

Studio Site Equipment: Importer

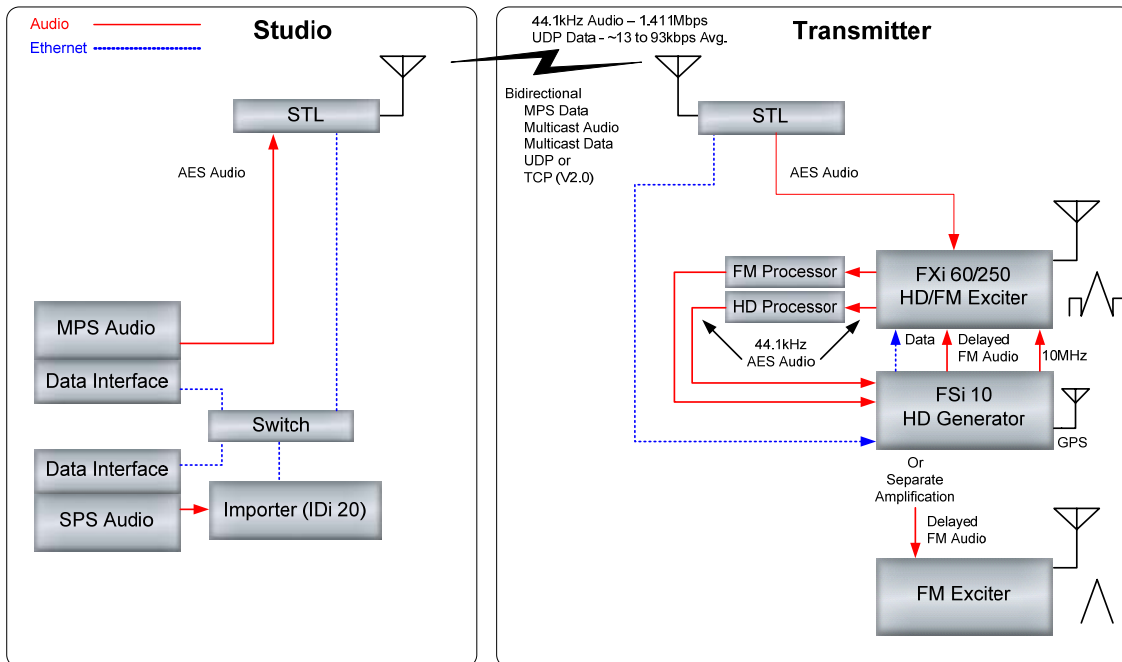
An **Importer** is a Windows-based computer that:

- ❖ converts secondary program audio to a data stream
- ❖ multiplexes all secondary audio data and program associated data into a single data stream
- ❖ sends that data stream to your Exporter/HD Radio signal generator

The Importer allows the further introduction of MP2 and MP3 modes. It adds the capability to adjust the bandwidth allocation of MPS and multicast audio, a process called **bandwidth provisioning**. MPS audio can be scaled to between 48 Kbps and 96 Kbps, and up to two Multicast channels can be added between 12 Kbps and 48 Kbps



each. The Importer is the point of control for the HD Radio bandwidth provisioning process.



– Fig. 2: Multicast Architecture, HD Generation at Transmitter Site –

Transmitter Site Equipment: HD Generator (FSi 10 or XPi 10)

After the Importer integrates multicast audio and data into a single Ethernet stream, the coded data stream is delivered to the HD Generator, either an FSi 10 or an XPi 10 depending on system architecture. The data stream is sent over a duplex link using either UDP or TCP (coming in version 2.0). The HD Generator encodes the main program audio (which is delivered to the transmitter site as AES audio), and integrates it with MPS PAD and the combined multicast data stream from the Importer (which includes multicast audio and PAD). This process creates a single properly-provisioned HD Radio signal payload.

The HD Radio Generator also delays your conventional FM audio so that as listeners switch between analog and HD Radio the audio is synchronized.

Transmitter Site Equipment: Exciter

Also at the transmitter site is the HD Radio/FM Exciter. The exciter digitally up-converts the complete HD Radio payload to the FM band and corrects for non-linearities in the

transmitter. The FXi series of exciters from Broadcast Electronics is capable of HD Radio-only, FM-only and simultaneous FM+HD Radio operation.

STL Considerations (I2E)

STL transport capability requirements for I2E systems are more demanding than pre-HD Radio requirements.

In systems where the main HD coding is taking place at the transmitter site, you **MUST** have an STL capable of a duplex (bidirectional) Ethernet link between the Importer at studio site and the HD Generator/Exporter at the transmitter site. Today, that bidirectional link must support UDP, but with the 2.0 release of the Importer software from iBiquity this will change to TCP.

Only the MPS PAD, SPS Audio, and SPS PAD are sent via this bidirectional link, meaning that if the Ethernet connection is lost, only PAD and multicast services are lost.

The Ethernet data capabilities of the STL should be capable of at least the provisioned data rate which ranges from ~17 Kbps to ~156 Kbps depending on the specific configuration. Keep in mind that much higher data bursts are inherent to the process.

Audio should also be considered when calculating the necessary overall bandwidth capacity of your STL. Bandwidth for AES audio can be calculated with the following simple formula:

$$(C \times R) S = B$$

C = Number of channels – for stereo, this is 2

R = Resolution in bits – standard AES audio resolution is 16 bits

S = Sample rate in bits

B = Bandwidth

Using this formula, stereo 44.1 kHz audio would require 1.4112 Mbps:

$$(2 \times 16) 44100 = 1411200 \text{ bps} = 1.4112 \text{ Mbps}$$

Table 1 shows the data requirements of just the multicast data sent by the Importer to the transmitter site. The average bandwidth column indicates (in Kbps) the average requirement to deliver the multicast data. The provisioned bandwidth column provides



some guidelines that indicate the required capability to pass the multicast data without losing packets. If other traffic is being delivered over this link, that bandwidth should be taken into account as you plan your system.

Data Rates and Provisioning Required for Modes and Services (Importer to Exporter or HD Generator)					
Interface	Direction	IP Protocol	Service Mode	Average Bandwidth Kbps	Provisioned Bandwidth Kbps
Importer to Exporter (XPi 10) or HD Generator (FSi 10)	Bidirectional (Duplex)	UDP	MP1, SPS1 = 12Kb	13.0	17.3
			MP1, SPS1 = 32Kb	34.9	46.6
			MP1, SPS1 = 48Kb	43.5	58.0
			MP2, SPS1 = 12Kb	21.5	28.7
			MP2, SPS1 = 32Kb, SPS2 = 12Kb	57.0	76.0
			MP2, SPS1 = 48Kb, SPS2 = 12Kb	65.2	87.0
			MP3, SPS1 = 24Kb	36.5	48.6
			MP3, SPS1 = 32Kb, SPS2 = 24Kb	69.4	92.5
			MP3, SPS1 = 48Kb, SPS2 = 24Kb	77.7	103.6
		TCP	MP1, SPS1 = 12Kb	16.3	27.2
			MP1, SPS1 = 32Kb	37.6	62.7
			MP1, SPS1 = 48Kb	53.8	89.6
			MP2, SPS1 = 12Kb	29.8	49.7
			MP2, SPS1 = 32Kb, SPS2 = 12Kb	65.2	108.7
			MP2, SPS1 = 48Kb, SPS2 = 12Kb	80.6	134.2
			MP3, SPS1 = 24Kb	42.3	70.4
MP3, SPS1 = 32Kb, SPS2 = 24Kb	78.2	130.3			
MP3, SPS1 = 48Kb, SPS2 = 24Kb	93.3	155.5			

- Table 1: Data Rate Requirements, Importer to Exporter/HD Generator -

Provisioned bandwidth for multicast data is calculated using the following formula:

$$X * (HD + Other) = P$$

X = Protocol overhead constant - 1.33 for UDP; 1.67 for TCP

HD = Combined average measured multicast bandwidth in Kbps

Other = All other WAN traffic in Kbps

P = Minimum bandwidth recommendation in Kbps



Summary of STL Bandwidth Requirements for I2E:

AES Audio Transport Capability:


- ❖ 44.1 kHz AES Audio = 1.411 Mbps (uncompressed)
- ❖ 32 kHz AES Audio = 1.024 Mbps (uncompressed)

Ethernet Data Transport Capability:

- ❖ Multiplexed PAD and SPS Audio: ~17 to ~155 Kbps duplex (bidirectional) UDP
- ❖ “Other” Ethernet capability as demanded by system configuration

Data Transport Considerations (I2E)

Data transport considerations are far more important with the addition of multicast services because of the increased data that is transported via the network. Encoded data is sent from the Importer in 1.48 second intervals, with each data burst referred to as a “frame”.



– Fig. 3: Importer to Exporter/HD Generator Data Capture –

Figure 3 captures graphically the nature of these data bursts as captured by the Ethernet sniffer software EtherReal. Bursts are delivered every 1.48 seconds, so your network must be capable of properly spooling and transmitting the data while waiting for the next burst.

Currently, all traffic is sent using duplex (bidirectional) UDP. Care should be taken when designing your network. When considering the data requirements of the system, include all anticipated traffic, not just the HD Radio component. Factor in at least an additional 33% capability to allow for overhead related to UDP communication.

(Provisioned bandwidth entries in Table 1 include this overhead calculation)

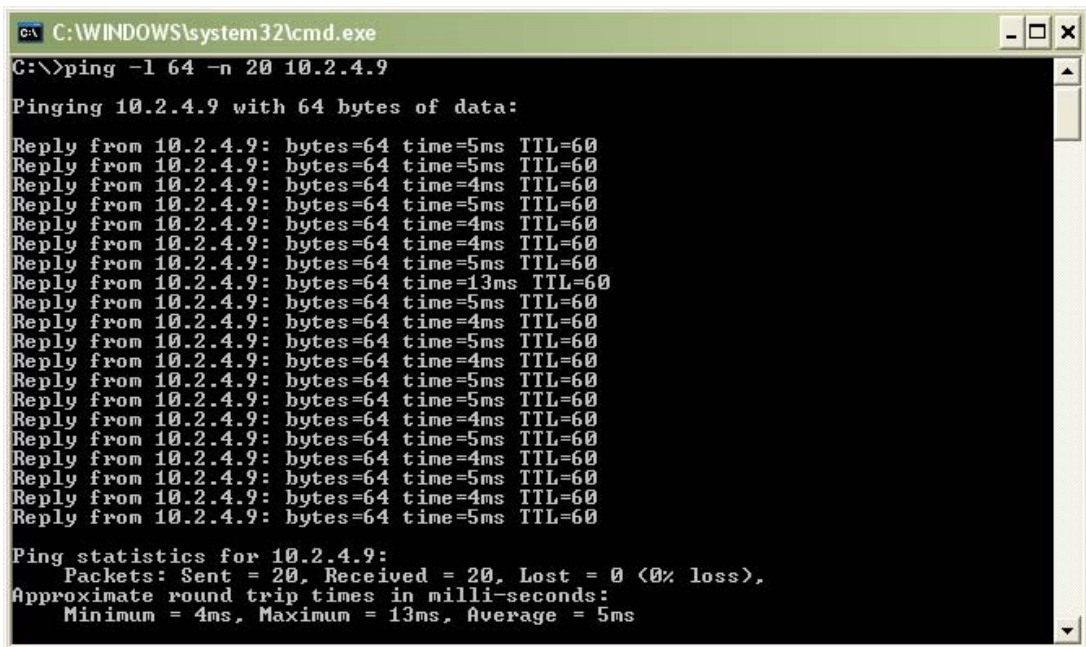
Additionally, the BER (Bit Error Rate) of your STL should be 10^{-5} (0.01%) or better. Higher packet loss will result in multicast service drop-outs. BER is difficult to measure

in the field, but STL manufacturers should provide test data showing BER for their products and this data should be taken into consideration while designing an HD Radio system.

Network link latency should be less than 100 ms. Latency can be measured using the Windows PING command (see Fig. 4).

Ping with 64 bytes of data (while running traffic) 20 times using this command line:

```
ping -l 64 -n 20 [address of HD generator]
```



```
C:\WINDOWS\system32\cmd.exe
C:\>ping -l 64 -n 20 10.2.4.9
Pinging 10.2.4.9 with 64 bytes of data:
Reply from 10.2.4.9: bytes=64 time=5ms TTL=60
Reply from 10.2.4.9: bytes=64 time=5ms TTL=60
Reply from 10.2.4.9: bytes=64 time=4ms TTL=60
Reply from 10.2.4.9: bytes=64 time=5ms TTL=60
Reply from 10.2.4.9: bytes=64 time=4ms TTL=60
Reply from 10.2.4.9: bytes=64 time=4ms TTL=60
Reply from 10.2.4.9: bytes=64 time=5ms TTL=60
Reply from 10.2.4.9: bytes=64 time=13ms TTL=60
Reply from 10.2.4.9: bytes=64 time=5ms TTL=60
Reply from 10.2.4.9: bytes=64 time=4ms TTL=60
Reply from 10.2.4.9: bytes=64 time=5ms TTL=60
Reply from 10.2.4.9: bytes=64 time=4ms TTL=60
Reply from 10.2.4.9: bytes=64 time=5ms TTL=60
Reply from 10.2.4.9: bytes=64 time=5ms TTL=60
Reply from 10.2.4.9: bytes=64 time=4ms TTL=60
Reply from 10.2.4.9: bytes=64 time=5ms TTL=60
Reply from 10.2.4.9: bytes=64 time=4ms TTL=60
Reply from 10.2.4.9: bytes=64 time=5ms TTL=60
Reply from 10.2.4.9: bytes=64 time=4ms TTL=60
Reply from 10.2.4.9: bytes=64 time=5ms TTL=60
Ping statistics for 10.2.4.9:
    Packets: Sent = 20, Received = 20, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 4ms, Maximum = 13ms, Average = 5ms
```

- Fig. 4: Using the Windows PING Command -

With the anticipated release of version 2.0 of the Importer software, TCP will be supported instead of UDP. A duplex (bidirectional) link will still be required, but TCP has some distinct advantages over UDP. Unlike UDP, TCP is not a broadcast protocol, which means it includes the provisions for verified packet delivery and buffering. TCP allows for higher packet loss, but requires higher bandwidth allocation.

Factor in an at least an additional 67% capability to allow for overhead related to TCP communication and be sure to include all anticipated traffic when considering the data

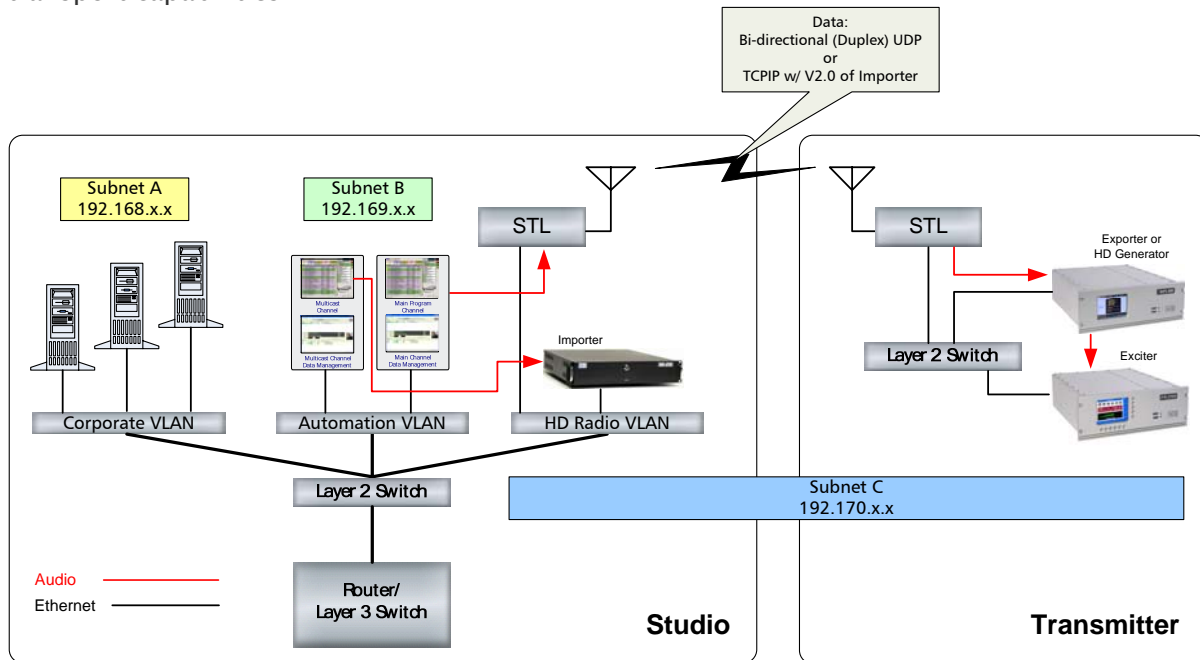


requirements of the system, not just the HD Radio component. *(Provisioned bandwidth entries in Table 1 include this overhead calculation)*

When using TCP, network latency should still be less than 100 ms, but use of TCP allows for a higher STL BER, up to 10^{-3} or 1.0%. Higher packet loss will result in multicast service drop-outs.

Network Considerations (I2E)

Networking is as important as the data capacity of your design. UDP is a broadcast protocol. Once you start sending the large data bursts, it can easily interfere with other components on a network, and other devices on the network can interfere with your transport capabilities.



– Fig. 5: Multicast Network Architecture, HD Generation at Transmitter Site –

One solution is to put the transport devices on a different subnet. Installing a layer 2 switch, you can create separate virtual networks, and isolate their traffic. Of course it is also possible to put all network components on a separate network connected to the automation network through the use of a managed router.

As you design your network, here are some basic guidelines to follow:

- ❖ Use Static IP Addresses for Importer, Exporter, & Exciter

- ❖ HD Equipment should be on own subnet
- ❖ Subnet MUST be separate from rest of facility through use of VLANs or physically separated networks
- ❖ Use a Switch or a Router for communication between VLANs or network segments
- ❖ Utilize IT resources in System Planning
- ❖ Utilize shielded CAT6 cable at the transmitter site

Summary (I2E)

As you plan your I2E HD Radio system, evaluate your STL for data just as critically you did for audio. Remember to allow an extra 33% overhead when using UDP, and if possible prepare for TCP implementation by allowing for an extra 67% overhead. Don't forget to include "other" network traffic demands when putting together your numbers.

With the I2E configuration, audio processing, HDC audio coding and HD carrier generation remain at transmitter site. In the event of an Ethernet link failure, only PAD and multicast services will be affected. Link failure will NOT affect Main FM audio or Main HD audio in the I2E configuration.

HD equipment should be networked on a subnet or separate network, with each device assigned a static IP address. High-quality CAT6 cable and managed switches or routers should be used.

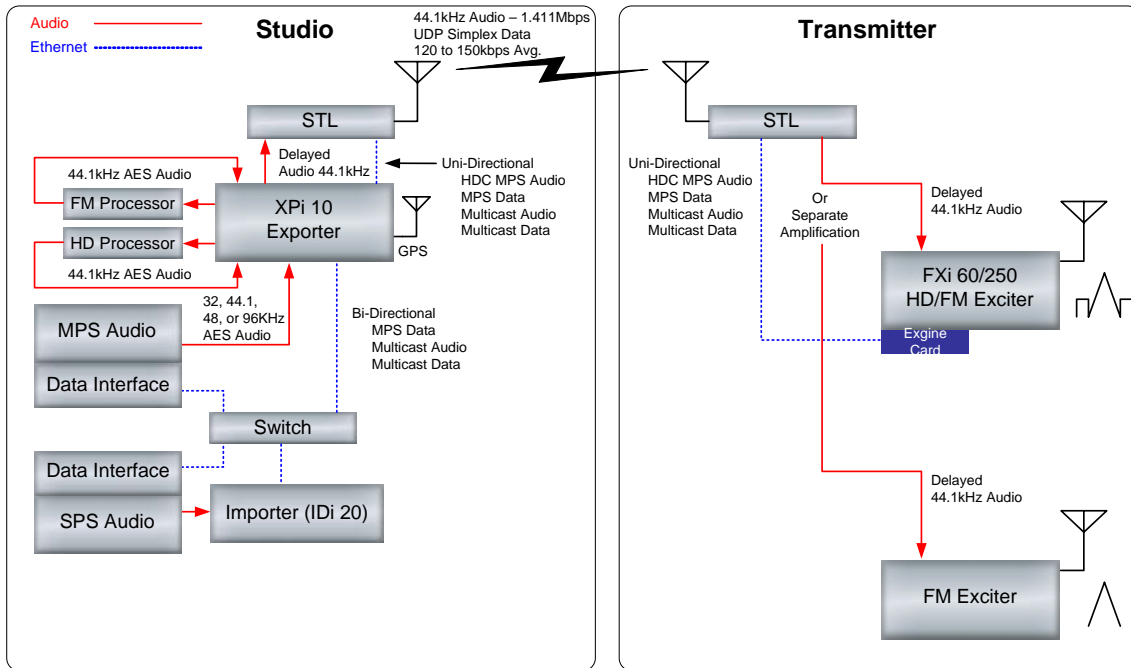
Multicasting – E2X (Exporter to Exgine)

E2X (Exporter to Exgine) allows multicasting in cases where bi-directional communication with the transmitter site is not an option and works well if the simplex (unidirectional) connection to the transmitter site is exceptionally stable. Under this architecture the HD coding for multicast AND primary audio is done at the studio site. This means that failures or deficiencies of the network connection can affect your main audio programming, not just the multicast programming as with I2E. All of the same basic capabilities of I2E (MPS, multicast, etc.) are available with the E2X architecture.

In addition to all audio processing moving to the studio site, key changes in this architecture are the addition of the **Exporter** and the **Exgine card** installed in the HD exciter. The Exporter serves double-duty, not only multiplexing multicast programming into a single data stream, but also HD encoding all audio for delivery over an Ethernet



connection to the transmitter site. This is the only architecture allowing you to send HD Radio data to the transmitter site over a simplex (unidirectional) UDP connection.



– Fig. 6: Multicast Architecture, HD Generation at Studio Site –

Data is still sent out at 1.48 second intervals, but the data rate is higher. The average simplex UDP data stream increases with E2X to approximately 120–150 Kbps due to the addition of the Main HD Radio programming to the data stream. Analog audio goes to the transmitter site as before. In fact, analog audio can be transmitted by a second STL, as long as the audio reaches the transmitter site at approximately the same time.

Studio Site Equipment: Importer

The function of the Importer is identical in both I2E and E2E implementations. The Importer:

- ❖ converts secondary program audio to a data stream
- ❖ multiplexes all secondary audio data and program associated data into a single data stream
- ❖ controls the HD Radio bandwidth provisioning process

Instead of sending the multiplexed data stream directly to an HD Generator, in this configuration the data stream is delivered to an Exporter.



Studio Site Equipment: Exporter

The Exporter codes Main program audio, delays FM analog audio, and integrates all HD Radio audio and data into a single simplex (unidirectional) UDP Ethernet stream. This complete HD Radio payload is then sent to the HD Radio Exciter at the transmitter site.

Transmitter Site Equipment: Exciter

Enabling simplex (unidirectional) communication with the Exporter based at the studio site requires the simple installation of an Engine card in the exciter. Once properly upgraded, the exciter digitally up-converts HD carriers to the FM band, corrects for non-linearities in the transmitter, and generates the OFDM HD carriers.

No GPS signal is required at the transmitter site in this configuration—GPS timing is integrated at the studio site into the multiplexed signal that is transmitted to the Exciter.

STL Considerations (E2X)

When the main HD coding is taking place at the studio site, your STL must be capable of a simplex (unidirectional) Ethernet link between the Exporter at studio site and the Engine-enabled Exciter at the transmitter site.

MPS audio, MPS PAD, SPS Audio, and SPS PAD are sent as a combined data stream via this unidirectional link, meaning that if the Ethernet connection is lost all HD Radio services are lost. The Ethernet data capabilities of the STL should be capable of at least the average provisioned data rate which ranges from ~160 Kbps to ~280 Kbps depending on the specific configuration. Keep in mind that high data bursts are inherent to the process. Audio should also be considered when calculating the necessary overall bandwidth capacity of your STL.



Data Rates and Provisioning Required for Modes and Services (Exporter to Engine)					
Interface	Direction	IP Protocol	Service Mode	Average Bandwidth Kbps	Provisioned Bandwidth Kbps
Exporter (XPi 10) to Engine (FXi 60/250)	Uni-Directional (Simplex)	UDP	MP1	119.7	159.5
			MP2	132.1	176.1
			MP3	149.3	199.0
	Bi-Directional (Duplex)	TCP	MP1	139.3	232.0
			MP2	155.6	259.2
			MP3	167.8	279.5

-Table 2: Data Rate Requirements, Exporter to Engine-

Table 2 shows the data requirements of the HD Radio data sent by the Exporter to the transmitter site. The average bandwidth column indicates (in Kbps) the average requirement to deliver the combined payload. The provisioned bandwidth column provides some guidelines that indicate the required capability to pass the data without losing packets. If other traffic is being delivered over this link, that bandwidth should be taken into account as you plan your system.


Overall, STL transport capability requirements for E2X systems focus on Ethernet transport since in this configuration, all HD Radio services are delivered as a single data payload.

Data Transport Considerations (E2X)

The reliability and data transport capacity of the STL is far more important with the addition of multicast services because of the increased data that is transported via the network. Encoded data is sent from the Exporter in 1.48 second intervals, with each data burst referred to as a “frame”.

Figure 7 captures graphically the nature of these data frames as captured by the Ethernet sniffer software EtherReal. Bursts are delivered every 1.48 seconds, so your network must be capable of properly spooling and transmitting the data while waiting for the next burst. Data bursts still occur at the same rate, but the impact is greater due to the increased amount of information being transmitted.





–Fig. 7: Exporter to Exgine Data Capture–

Even operating at peak efficiency, the 1.544 Mbps capability of a T1 is not sufficient to transport all HD Radio data **and** digital audio for the analog channel. If you are transporting 44.1 kHz audio (which requires 1.411Mbps), that leaves only approximately 137 Kbps—lower than the lowest recommended provisioned bandwidth recommendation.

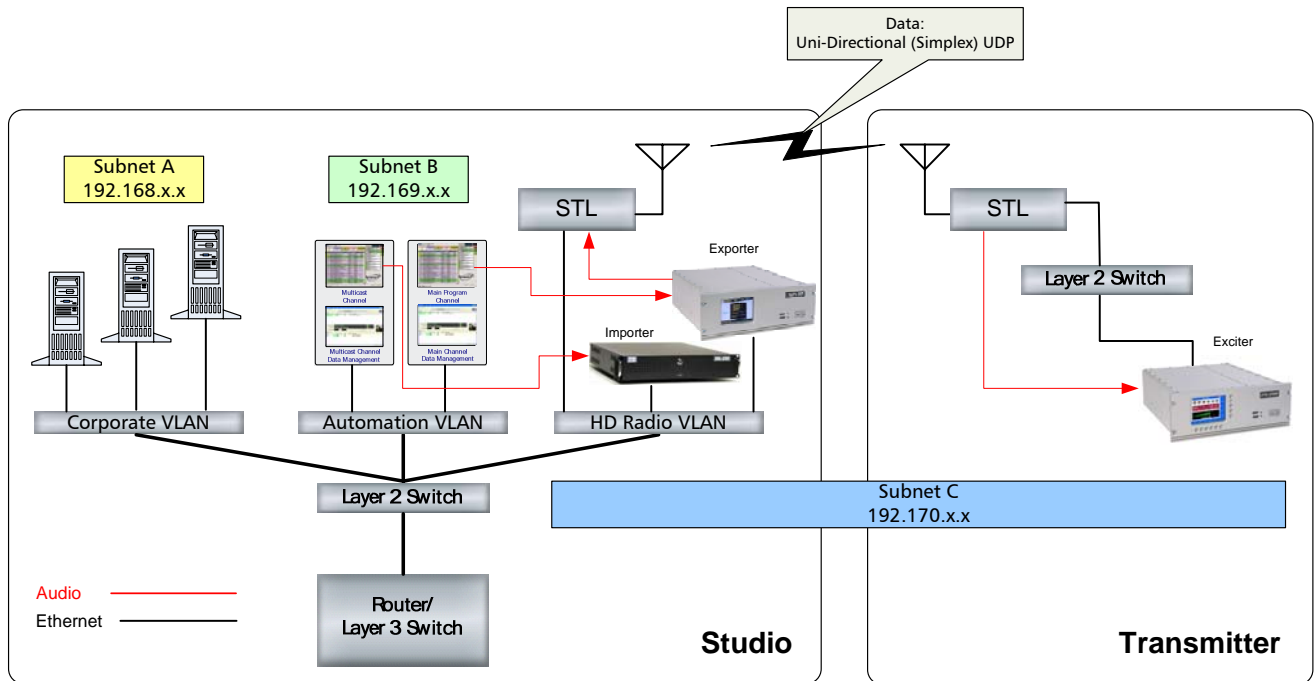
Currently in the E2X configuration, all traffic is sent using simplex (unidirectional) UDP. Care should be taken when designing your network. When considering the data requirements of the system, include all anticipated traffic, not just the HD Radio component. Factor in at least an additional 33% capability to allow for overhead related to UDP communication.

Using a simplex link, STL reliability is a key factor. The BER (Bit Error Rate) of your STL should be 10^{-6} (0.001%) or better. Higher packet loss will result in complete service drop-outs.

Network Considerations (E2X)

As with the I2E configuration, physical or virtual separation of network segments is still recommended. Since data bursts in this configuration are even more extreme, it is even more important to properly plan and manage your network to minimize the risk of overloading your network and risking packet loss.

The same network guidelines as explained for I2E systems should be followed. HD equipment should be networked on a subnet or separate network, with each device assigned a static IP address. High-quality CAT6 cable and managed switches or routers should be used.



- Fig. 8: Multicast Network Architecture, HD Generation at Transmitter Site -

As you plan your E2X HD Radio system, evaluate your STL for data just as critically you did for audio. Remember to allow an extra 33% overhead when using UDP, and if possible prepare for TCP implementation by allowing for an extra 67% overhead. Don't forget to include "other" network traffic demands when putting together your numbers.

With the E2X configuration, audio processing and HDC audio coding takes place at the *studio* site while HD carrier generation remains at the transmitter site. In the event of an Ethernet link failure, all HD Radio services will be affected. If only the data link fails it will *not* affect Main FM audio in the E2X configuration as it may be transported separately. If both the audio and HD services are sent over a common link and that entire link fails it will of course affect both services.


Networking basics

In Information Technology, a network is a series of points or nodes interconnected by communication paths.

The most common forms of Ethernet networks are 10Base-T, and 100Base-T, capable of 10 Mbps and 100 Mbps respectively. The faster 1000Base-T, capable of a nominal speed of one Gigabit per second is becoming more common. For HD Radio purposes, devices capable of at least 100 Mbps are strongly recommended.

All three of these options use shielded or unshielded twisted pair cable, connecting all nodes through a central point. Category 6 cable is recommended. Category 6 offers superior transmission performance, and twice the available bandwidth of category 5e. Shielded cable is preferred due to its higher resistance to RF interference.

Imagine the computers on your network as a bunch of loud people in a large un-moderated meeting room. Only one person can talk at a time, because communication consists of standing up and yelling at the top of your lungs. People are allowed to start communicating whenever there is silence in the room.

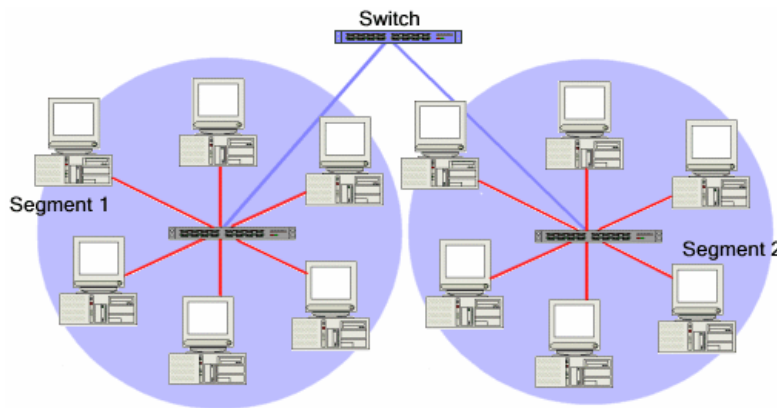


If two people stand up and start yelling at the same time, they wind up garbling each other's attempt at communication, an event known as a **collision**. In the event of a collision, the two offending parties sit back down for a semi-random period of time then one of them stands up and starts yelling again. As the number of talkers and the amount of stuff they talk about increases, the likelihood of collisions occurring increases geometrically.

This huge meeting room is like a network hub. In data communications, a hub is a place of convergence where data arrives from one or more directions and is forwarded out in one or more other directions. In an Ethernet network, each workstation or server is located at the end-point of a cable connected to the hub.

What is a switch?

Simple network hubs have no traffic managing capabilities. As such, they are unsuitable for HD Radio applications due to the critical nature of our communication...we need a better solution to give us the best chance for successful delivery.



A switch is a network device that separates multiple segments on a network, but is smart enough to know if a particular packet needs to be sent to the other side.

Think of our large meeting room. If we split the room in two with a divider that has a single door, the two rooms would be quieter, and the chances of talking over the top of someone else would be reduced. By putting someone in that door that could relay information to the other room as needed, communication would be maintained, but overall each individual room would work better.


What is a router?

Routers perform the same basic function as switches, separating networks into multiple segments. But while switches automatically learn how to pass data between networks, routers must be configured manually by a network administrator. An added function of a router is the ability for the administrator to establish how the data gets from point A to point B...the "route" the data takes between connected segments. While this involves an additional layer of complexity, it also provides additional control as you configure your network.

Communicating on the network

To communicate over a network, computers send packets or datagrams of information to each other. Each packet contains about a thousand bytes of information. To send a message that is longer than a thousand bytes, the computer breaks it down into packet-sized chunks, puts a sequence number on each chunk, and then transmits one chunk in each packet. The computer on the other end receives one packet at a time and uses the sequence information to put all the chunks back together into a single message.

Just like when you write a letter, both parties must speak the same language to have a meaningful exchange. In network terms, this language is called a protocol. Both computers must be configured to use the same protocol in order to communicate.



The IP (Internet Protocol) suite is the closest thing that the computer world has to a universal language. It is also routable, meaning it's easier to connect multiple segments over a wide area network. The two most common transport protocols, or languages, in the suite are TCP (Transport Control Protocol) and UDP (User Datagram Protocol.)

Using TCP, applications on networked computers can create connections with each other to exchange data. TCP guarantees reliable and in-order delivery of data by using receipt acknowledgments and error checking. UDP does not provide the same error checking and ordering guarantees that TCP ensures, but tends to be faster and more efficient.

Network Addresses

Using our letter analogy, when you send information you not only have to use the same language, you need to correctly identify the address of the intended recipient. Since we're working with transport protocols, we also need to define some aspects of the route our letter should take. Computers on a network are configured with a unique address to identify senders and receivers, and a subnet mask and a gateway to help determine routing.

An IP address is a 32-bit number that uniquely identifies a host (a computer or other device, such as a printer or router) on an IP network. Reading binary, a computer sees a typical address as:

11000000101010000111101110000100.

Since most people don't read binary as well as computers do, this number would be hard for us to understand and remember. To make it more manageable, we first divide it into four parts of eight binary digits known as octets. Now we see:

11000000.10101000.01111011.10000100



11111111.11111111.11111111.00000000

Lining up the IP address and the subnet mask together, the network and host portions of the address can be separated. The subnet mask value of 1 indicates the network portion of the address, and the value of 0 indicates the host portion.

255	.255	.255	.0
11111111	.11111111	.11111111	.00000000
192	.168	.123	.132
11000000	.10101000	.01111011	.10000100

In this case, the first three octets are the network address, and the last octet is the host address. All decimal subnet masks convert to binary numbers that are all ones on the left and all zeros on the right. Using some other common subnet masks:

255	.255	.0	.0
11111111	.11111111	.00000000	.00000000
192	.168	.123	.132
11000000	.10101000	.01111011	.10000100

Or:

255	.255	.224	.0
11111111	.11111111	.11110000	.00000000
192	.168	.123	.132
11000000	.10101000	.01111011	.10000100

This last example specifies the first 19 bits for the network address, and the last 13 bits for the host address.

By specifying different subnet masks for different computers, we are essentially creating new subnetworks or VLANs (virtual local area networks). As far as the computer hardware is concerned, different subnet masks indicate separate networks even if everything is physically linked through the same switch or hub. This is an important way to segregate and isolate traffic from different functional networks.

Gateways

When a computer tries to communicate with another device, it compares its own address and subnet mask to the address and subnet mask of the destination host to determine whether the destination is a local host or a remote host. The computer resolves and compares the network address of the two subnets...if they match the destination



computer is on the local network. If they are different, the destination computer is on a remote network.

If the result of this process determines the destination to be on the local network, the computer will simply send the packet on the local subnet. If the result of the comparison determines the destination to be on a remote network, then the computer will forward the packet to the switch or router that connects it to other subnetworks, called the default gateway. It is then the responsibility of the gateway to forward the packet on to the correct network.

Bridging networks

To summarize, when a device sends a packet using either TCP or UDP, it compares its own address to the address of the destination device. Based on the subnet mask, it can tell whether the destination device is on the local network, or on a remote network. If the destination device is on a remote network, it sends the packet to the router configured as the default gateway, which sends the packet along to the remote network.

Using switches and routers, engineers can isolate network segments by manipulating these IP settings.

