10641 Techwoods Circle
Cincinnati, OH 45242
513.786.6638

1370 Reynolds Avenue, Suite 101
Irvine, CA 92614
(949) 852-6660
www.mtm.com

HD Radio™
Networking Implementation Recommendations

Developed for:

iBiquity Digital Corporation

8865 Stanford Blvd.
Columbia, MD 21045

Prepared by

Kurt VanderSluis
kvandersluis@mtm.com

July 27, 2006

Contents

I. Background

Ibiquity Digital asked MTM Technologies to provide recommendations to on the construction and provisioning of the digital links between HD radio studio facilities and their transmitters. MTM makes the following recommendations after studying the nature, resiliency and failure modes of the digital transmissions under adverse network conditions in a controlled lab setting as well as doing a field survey of several radio stations employing iBiquity technology in production settings with varying degrees of success. The investigative team found that success depended on these main factors:

- Adequate provisioning of the WAN link
- Control of extraneous traffic through the link by network design
- Good network infrastructure quality - low packet loss rate and latency including RF links
- Use of receive buffering when available
- Use of TCP for transmission control rather than UDP

II. General Description of the iBiquity Technology

Ibiquity has its technologies running in many configurations of equipment and data streams designed by multiple manufacturers, which can be characterized as follows:

A data source sends audio and data to a receiver which then converts the information into a stream of HD Radio Program Data units (PDUs) and sends these PDUs on to the next device, possibly after combining the received PDUs with data received from another source(s).

Depending on the equipment configuration and the number and kind of digital services offered, there are many ways the PDU frame can be constructed, but in every case, a frame corresponds to 1.48 seconds of content. The amount of digital data carried in an individual frame rate PDU, depending on service configuration, varies from approximately 9,500 bytes to 26,000 bytes.

The transmission of these frames can be characterized as a discontinuous data stream. The bulk of the frame data is sent in a burst at the beginning of the cycle. The remainder of the frame is sent at a time isochronously, corresponding to the time within that frame when that data will be transmitted.
The receiving device has a buffer into which the frame is placed until it is ready to be passed along to the next device or transmitted onto the airwaves. Depending on the specific kind of receiving device, the buffer space available varies from one frame to as

many as twenty frames. For some equipment, the buffer space can be set by the station engineer within a limited range.

Transport between source and destination is provided either by TCP or UDP. In some equipment there is a choice between these two transport mechanisms; in other equipments, there is only one choice. When the equipment offers no choice, the transport mechanism is UDP.

The link between the studio and transmitter, depending on the equipment employed, can be either unidirectional or bidirectional with the transport medium either a digital land link over a private network or over an RF link.

III. Recommendations

Recommendation 1
Provision the WAN Link with Adequate Bandwidth

When using TCP, the WAN link must have a minimum of 40% overhead (reserve bandwidth) in order to function properly. This overhead should be calculated on the total traffic through the WAN Link, which can consist of the following components:

- The IBOC data stream
- Traffic from management utilities such as VNC, telnet, etc
- Any broadcast and/or multicast traffic
- Any auxiliary traffic used for other purposes

The sum total of this traffic should occupy no more than 60% of the provisioned bandwidth. If other traffic is going through the WAN link, the link should have some class of service, QOS or other priority queuing technique employed to ensure that the IBOC traffic has its required bandwidth under all conditions. For more information, please the document, "HD Radio™ Networking Best Practices".

For UDP, the total traffic can be no more than 75% of the provisioned bandwidth.

Recommendation 2
Make the WAN Link a Separate IP Subnet

The WAN Link should constitute its own IP subnet to keep the amount of broadcast and multicast traffic to an absolute minimum. The sending device (the Importer, for example) can either be in the WAN link subnet or in the program automation subnet, but no office, production or studio computers should reside in the WAN link IP subnet.

Recommendation 3
Achieve and Maintain Infrastructure Quality

To successfully deliver the data stream using TCP and 20 receive buffers, the WAN link may have no more than 80 milliseconds of latency (measured on the unloaded link) and may not have more than a 1% error rate. With 3 receive buffers; the WAN link may have no more than 0.01% packet loss and no more than 50 milliseconds latency. See Appendix B for more detail on infrastructure quality

Recommendation 4
Use Receive Buffering When Available

Because all networks are subject to packet loss and congestion, it is prudent to take advantage of receive buffering when available. Receive buffers allow time for the data stream to recover from the occasional loss of transmitted packets. This recommendation only applies when there is a mechanism to recover lost data.

Recommendation 5
Use TCP When Available

Not all equipment configurations allow for it, but TCP is always the preferred mode of transmission control when it is available because of its ability, to recover lost and corrupted packets.

IV. Recommendations - Discussion and Detail

Provisioning

In order for a TCP data stream to function properly under adverse conditions, the link that carries it must have reserve bandwidth above and beyond the data rate of the stream. This is necessary to accommodate the higher data rate that occurs when the stream recovers from lost packets. Through lab experimentation, we found that if the average rate of the IBOC data stream occupies less than 60% of the link's bandwidth, the IBOC stream can tolerate up to 1% packet loss and 80 milliseconds latency. We were sometimes able to transmit the data stream successfully under slightly worse conditions, but the stated guidelines are conservative as well as being very reasonable and achievable infrastructure quality goals. In general, healthy WAN links normally have packet loss rates less than 0.1% and many have loss rates of less than 0.01% (1 packet lost per 10,000 transmitted).

Additional bandwidth beyond the recommended guideline allows operation under poorer conditions, but with diminishing returns. In general, bandwidth should not be used to adjust for a poor network

Traffic Control

Through our field investigations, we found that the only sure way to prevent extraneous traffic from traversing the link is to make the link its own IP subnet. This separates the link from the production network and spares the link from having to carry any broadcast or multicast traffic from the production and/or office network. This was found to be a critical factor during our investigation. This recommendation should be followed without exception.

Any traffic that switches must forward to all ports must be minimized. In addition to broadcast and multicast packets, this category also includes unicast packets sent to MAC addresses that are not in the switch's forwarding table. This is a situation that occurs occasionally when a device from outside the local network sends packets to a device that was recently in the subnet. The device is still in the router's ARP table, but has already been aged out of the switch forwarding tables, which typically has an aging timer than is much shorter than the router's ARP aging timer.

There are steps that the network manager can take to prevent this situation. One is to reducing the router's ARP aging timer to match the switch forwarding table aging timer. Another is to program the switch to block unicast traffic sent to unknown MAC addresses. These steps involve a level of IT trickery that is best avoided, are not available

with all equipment and are not nearly as effective as creating a separate IP subnet for the WAN link.

On those occasions when this special case unicast traffic presents itself, the network utilization offered to the WAN link may be greater than its available bandwidth. We observed this phenomenon in 2 of the radio stations we visited and in each case it caused audio drops.

The Exciter will always be within the WAN subnet. In I2E configuration, the Importer may be placed on either side of the subnet boundary, and in the E2X configuration, the Exporter may be placed on either side of the subnet boundary. Except for such equipment as may be necessary to build the infrastructure, i.e. routers and switches, no other station equipment should be allowed in the WAN link subnet.

Remote Management Traffic

Any remote management of equipment within the transmitter facility will create a traffic footprint that the WAN link must convey in addition to the IBOC stream. This traffic must not be allowed to fill up the reserve bandwidth needed by the IBOC stream for error recovery. It is essential to know how much traffic will be created by the management software you plan to use and to factor this into your provisioning equation.

VNC remote management, for example, uses as much as 40 KB. If you are going to use VNC concurrently with the IBOC stream, you must add the 40 KB VNC footprint to the nominal load of your IBOC stream before calculating your bandwidth provisioning requirements.

TCP v. UDP

In many configurations, there is a choice between using TCP and UDP as the transport control mechanism. In some configurations, UDP is the only choice. When there is a choice, TCP is, in most cases, the better choice because of its ability to recover from packet loss. However, there may be situations where UDP is also a viable choice.

TCP's advantages come from its built-in features:

- Packet sequencing
- Ability to recover lost data
- Ability to adapt data rates up or down according to conditions
- End-to-end data integrity checking

UDP has the following advantages:

- Ability to operate over a simplex (one-way) STL or WAN
- A shorter broadcast delay due to fewer receive buffers (typically 1 or 2)
- Ability to operate with less bandwidth overhead
- Shorter audio drops when they occur (although they may occur more frequently)

Using TCP, audio drops occur whenever the receive buffers are depleted. In most circumstances, the audio stream will not resume until the receive buffers are all restored. Since each receive buffer corresponds to 1.48 seconds of audio, depletion of 20 receive buffers will result in an audio drop of 30 seconds or longer.

With UDP transmission, the loss of a single packet only ruins the audio frame of which it is a constituent and the resulting outage will only last for the duration of that single audio frame - 1.48 seconds. For any constant packet loss rate, one would expect fewer audio drops with TCP but of shorter duration with UDP.

Burst Rate

Because the data flow in the IBOC data stream is not continuous, there has been some concern over the concept of a "burst rate". The bulk of the audio frame is sent in an initial burst at the beginning of the frame. In an example service configuration sending P1, P3 and P4 data, this initial burst would contain nearly 20,000 bytes and be sent at near wire speed in the case of a 10 MB connection. On a 100 MB connection, the data would be sent as fast as the device can send (usually in the vicinity of 40 MB). This burst must pass through the WAN link, which is usually much slower than the burst transmission speed.

In the instance of a 20,000 byte burst on a 10 MB connection, each packet will contain approximately 1500 bytes of data, so approximately 13 full-size Ethernet frames are required to send the data. These 13 frames will be sent at very near the wire speed (10 MB). Each Ethernet frame takes approximately 1.2 milliseconds. While there are gaps between the packets, these are typically very short, less than 200 microseconds. For 13 packets of 1500 bytes each being sent at intervals of approximately 1.25 milliseconds, the 20,000 byte burst takes approximately 16.25 milliseconds to transmit on 10 MB Ethernet. During this burst, the data rate is approximately 9.7 MB. Some concern has been raised that this burst data rate might overwhelm the WAN link, which would only require a bandwidth of 192 KB for the configuration being considered in this example.The concern, however, is a false one.

When faster links are connected via slower links, it is normal for the data to arrive at the bridging device in bursts much higher than the bandwidth of the slow connecting link. A buffer is used to accommodate the incoming data until it can be sent over the slower link. According to the rules of Ethernet bridging (IEEE 802.1), the buffer of a bridge or switch can hold the packet for no more than 2 seconds before transmitting it. After 2 seconds, the packet must be discarded. This rule requires the designers of bridging devices to provide a buffer that can hold at least 2 seconds worth of data at the lower links' speed. For a 192 KB bridged link, this would correspond to approximately 40 Kbytes of data. In the IBOC example that we are considering, (a configuration with the largest possible burst) the burst of data is approximately half of this maximum figure, 20 Kbytes.

While it is possible to overwhelm a slower link with a burst of data transmitted from the faster link, under normal data transmission with infrastructures constructed according to the recommendations outlined in this document, this condition should never occur. In the case of a packet loss rate higher than the 1% allowed, however, it could happen if the data stream got far enough behind and the system was sending at high data rates over more extended periods of time. Under those conditions, the bridge might be forced to discard packets because of the 2 second rule.

As a general note, in almost all cases, including this one, the shortest period of time that needs to be considered in digital networking is 1 second. Since the audio frame is approximately 1.5 seconds and almost all of it is sent during the first second, the burst rate of an IBOC data stream is approximately 1.5 times the average data rate.

Appendix A: Provisioning Calculation Example

IBOC Data stream = MP1 I2E SPS1 = 56 KB
Management traffic = 43 KB
VNC = 40 KB
Telnet = 3 KB
Broadcast and Multicast = 1 KB
Other traffic = 0

Total Traffic = 56 + 43 + 1 = 100 KB

100 KB = 60% of Minimum Provisioning

Minimum Provisioning = 5/3 x 100 KB = 167 KB = 3DS0 (192 KB)

Appendix B: Network Quality Definitions

LAN Network Standard

A healthy LAN has a packet loss rate of less than 1/10,000. An unhealthy LAN has packet loss of greater than 1/1000. A marginal LAN has a packet loss rate greater than 1/10,000 but less than 1/1000.

LAN latency, measured apart from system or application latency, is almost always less than 5 milliseconds and usually below 2 milliseconds. The normal response time of a client-server application on a LAN may be considerably longer, up to 250 milliseconds in the case of a Microsoft Exchange server.

WAN Network Standard

A *healthy* WAN link has a packet loss of less than 1/1,000. An *unhealthy* WAN link has packet loss of greater than 1/100. A *marginal* WAN link has a packet loss rate greater than 1/1000 but less than 1/100.

Good WAN latency is 30 milliseconds or less. A *reasonable* WAN latency is between 30 and 70 milliseconds. A *high* WAN latency is between 70 and 120 milliseconds. Within a metropolitan area, WAN latency should never exceed 120 milliseconds; this value is long even for a WAN link spanning the continental US.

Important: LAN and WAN latency, as expressed here, is the latency measured in the absence of a load on the link. As the utilization of a link approaches the available bandwidth, it is normal for the measured latency to become larger.

If you are measuring latency on a WAN link with appreciable utilization using a ping test, perform the test with a 64 byte packet with at least 10 samples and use the minimum values that you find as an approximate measure of the link latency.

Appendix C: Bandwidth, Utilization, Latency & Packet Loss

In digital networking, bandwidth is the term used to denote the data carrying capacity of a network link. Bandwidth is usually expressed in bits per second or in terms of the designation of a link such as T1, E1, DS0 or OC24. etc., each of which corresponds to a specific data rate expressed in bits per second.

Utilization is the level of a link's carrying capacity that is in use and is expressed either as a percentage of the link's bandwidth or in bits per second. Latency is the round trip time of a packet between source and destination across a link. Latency is a combination of the delay caused by transmission across the link due to packetization and transmission, the response time of the destination node, and any delay due to congestion of the link. When no qualifiers are used, a link's latency usually refers to the round trip time across a link where no other traffic is present and measured using a target node with a system response time small enough to be considered insignificant to the measurement.

All digital networks have latency and packet loss. These phenomena, when they are excessive, can interfere with the transmission of data, particularly streamed data. In streams using TCP, when packets in a stream are lost in transit, the stream momentarily stops to recover the lost data. The amount of time it takes to recover the lost data is governed largely by the link's latency. When the lost data is recovered, the data stream usually has to "catch up" by transmitting data at a higher than normal rate. If packet loss is infrequent and the link's latency is low, the stream can easily recover. As packet loss and/or latency climbs, the recovery process cannot as easily catch up and at some point, packet loss and latency combine so that the buffers in the receiving device dwindle. When they are depleted, the application depending on the data stream suffers a loss or degradation of service noticeable to the user.

The role of bandwidth in the recovery process is that more bandwidth allows the stream more headroom when trying to catch up. If the link can accommodate a higher rate of transmission, the stream can "catch up" more quickly when it gets behind -- larger deficits can be recovered without being noticed by the user.

Latency affects the recovery process as well. Higher latency means that the data loss will not be noticed as quickly, the request for data will take longer to reach the sender, longer for re-transmitted data to reach the receiver, longer for the receiver to acknowledge the receipt of the lost data. And because bandwidth and latency are related (when the available bandwidth is low, latency increases), there is a latency benefit from having adequate bandwidth headroom available.

Appendix D: Ethernet Considerations

Ethernet Collisions and Utilization

Some readers may have become concerned about collisions in the previous discussion of burst rate. In this discussion, it was noted that the burst rate of the audio data frame would be, for a very brief duration, 9.7 MB, or 97% of the Ethernet bandwidth.

There is a very widespread and long-standing misconception that for Ethernet, collision rate rises with utilization and that Ethernet may never be used at near capacity for this reason. This is not exactly true. In Ethernet, the collision rate is actually proportional to the number of simultaneous transmitters. In the current era of networking, high utilization does not correspond to high collision rates for the following reasons.

First, in the 100% switched network, which is common in the present era, each network segment contains only 2 possible transmitters, the switch port and the networked device. Second, when these connections are full duplex, as is common, there should theoretically be no collisions. The presence of any collisions usually indicates a duplex mismatch between the switch and the device and/or a hub device downstream. Third, high data rates do not necessarily correspond to a high number of simultaneous transmitters.

In the early years of Ethernet on twisted pair (which is where this misconception originated), routers and bridges were relatively expensive. It was common to find large networks with dozens or even hundreds of networked nodes within a single collision domain. Also during this time period, the devices were very slow relative to the wire speed; few of them could transmit at speeds higher than about 20% of the10 MB bandwidth. In these networks, there was a correlation between utilization and collision rate because higher utilization was almost always caused by a high number of simultaneous transmitters.

To illustrate this concept, consider a broadcast ping on this historical network. All of the devices on the network would simultaneously try to respond to the ping. The number of resulting collisions might well exceed the number of successfully transmitted packets. Although the corresponding utilization would be low, the collision rate would be very high.

Consider on this same network the case where most of the devices have no traffic to send or receive and 4 devices are drawing data from 4 different servers, each trying to use 20% of the Ethernet bandwidth. Here the data rate would be high, but the collision rate would be very low.

While it was true that, during this era, there was often a strong correlation between utilization and collision rate, the correlation depended on usage conditions that no longer exist.

Ethernet Speed and Duplex Issues

When configuring devices and switch ports, a general rule of thumb is that for any given port, both sides of the connection should have their duplex and speed set manually -or- both sides

should have their speed and duplex set to auto-negotiate, which is in almost all cases the factory default setting.

The reason for this is a rule of auto-negotiation that is not well known. When an auto-negotiating interface encounters a device that does not auto-negotiate, it is obliged to set itself to half-duplex. Many network administrators, unaware of this rule, manually set the end node to 100/full and do not configure the corresponding switch port accordingly. The switch port will set itself to 100/half, causing a duplex mismatch. In networks with unmanaged switches, it is not possible to set the switch ports individually, and in these networks, the devices should never be set to full duplex.

The consequence of duplex mismatch is network performance degradation that increases exponentially with packet rate. Network administrators should be on the lookout for this condition because it is a frequently made configuration mistake.

With managed switches, the mismatch is discovered by examining network error statistics. For the switch port, the duplex mismatch will result in alignment and CRC errors. On the device, the mismatch will result in late collisions. While ordinary collisions do not cause much concern unless they rise above 10% of the packets transmitted, late collisions almost always indicate a problem. If the number of late collisions exceeds 1 per million packets transmitted, the problem should be investigated. On a switch port, the same problem is indicated if the sum of the CRC and alignment errors is greater than 1/10,000.

On unmanaged switches, and in the absence of interface statistics on the end node, the problem can only spotted with a protocol analyzer.

When the high error rate is spotted and a duplex mismatch is suspected, the first step is to check the duplex settings - this is by far the most common cause of late collisions, CRC and alignment errors in networks running on certified cabling installed by professionals. If a duplex mismatch is not found, there are 2 other conditions that may be the causing the errors.

The first alternate cause is inferior cabling. In addition to cabling that was improperly or poorly installed, cabling that was certified at the time of installation can be compromised and below certification due to subsequent damage or mishandling. Using a cable tester, verify that the cable meets the installation specification designated (use Cat 5 if this is unknown) with regard to length, noise, signal attenuation, impedance and cross-talk

A second, less common cause is that slight discrepancies in vendor implementation of the auto-negotiation process may result in an incompatibility between the NIC and the switch port. This rarely happens when using recently made equipment but was common for equipment made in the years soon after the innovation of full duplex Ethernet (1996). To resolve this issue, replace the older device(s), either the NIC or switch or both. If that is not possible, manually set both sides to half-duplex. In cases where the NIC cannot be replaced, set both the switch port and the device to half duplex. It should be noted that in some instances, setting both sides to 10/half often results in the best network performance.