MTM
TECHNOLOGIES

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
Best Practices

Developed for:

iBiquity Digital Corporation

8865 Stanford Blvd.
Columbia, MD 21045

Prepared by

Trieu Vu
tvu@mtm.com

July 27, 2006

Contents

I. Introduction

HD Radio has engaged MTM to perform a detailed analysis of the network in relation to the HD Radio streams. Representative sample radio stations were chosen based upon the following criteria

- WAN Link(T1, or RF)
- Manufacturer(Harris, Moseley)
- Protocol(TCP, UDP)

Working and problematic sites utilizing the same configuration, setup, and equipment were also chosen to aid in the understanding, and analysis.

Initially, baseline tests were conducted in the lab with the different audio streams. Then latency and packet drops were introduced and the results were recorded. Armed with the information, field visits to representatives' radio stations over a two week period proved to be invaluable in both discovering and rectifying anomalies, and dropouts.

This paper is a collection of the findings both in the lab and in the real world. It contains rudimentary definitions of networks devices, their functions and role in a network. More importantly, it also contains recommendations on the deployments, its locations, and proper use to ensure the utmost reliability and predictability for an HD Radio network deployment.

II. Networking Component & Functions

To deploy HD radio, a networking infrastructure needs to exist to support the data stream from the Importer -> Exciter. The typical components found in a network infrastructure can be classified as

- Hub
- Switch
- Router/L3 switch

These three components are not all mandatory in the deployment of HD radio. The following section will

- Define its operation, role, and implementation
- Define its limitations

Define the suggested deployment strategy

III. Hub

In the early days of Ethernet, hubs were the means to interconnect all workstations, servers, and printers together. Hubs are simplistic in design because it is a Layer 1 repeater/signal regenerator. It is a "Shared" medium. For example, in a 16 port hub, packets sent from workstation #1, will be repeated/sent to all remaining ports(2-16). The caveat is that only one workstation can speak at time. Should two workstations seize the line at the same time, they would back off a random amount of time, and retransmit.

The limitations of hubs are

- Scalability – an Ethernet hub would work just fine for a network of 1-25 workstations, yet will fail to scale to 100 or even 200 workstations because the 10Mbps bandwidth is being shared amongst the 100 or 200 workstations. Moreover, with the increase in the number of stations, collisions will increase, thereby reducing the chances of transmitting the packet successfully the first time around. With time sensitive application, ie. SNA traffic, this sporadic behavior causes application timeout, and user frustration.

- Shared access – each station sees ALL traffic on the network, regardless if the traffic is destined for it.

Due to the limitation of the shared medium & scalability, hubs were soon replaced by switches in the corporate infrastructure. Having said that, for HD networking, hubs can be deployed at the transmitter site since only a handful of devices will need to be connected and their bandwidth needs are minimal(<500K/s). Should there be a legacy hub in the network inventory, at best, the transmitter site is the only place for its deployment.

IV. Switches

As applications became more robust and heavily reliant on the network for services, shared 10Mbps(hubs) gave way to switches – granting each users their own dedicated 10Mbps bandwidth.

Operation – switches operates at a Layer 2. It has the ability to "see" into the source and destination MAC address. At power up, the switch has a blank forwarding table. It then listens on every port, and with each packet received, it updates the forwarding table with the source MAC address and the port which the packet was received on.

For example, after 30 seconds the forwarding table might look like this:

Workstation	Mac-Address	Port #
A	AA:AA:AA:AA:AA:AA	1
B	BB:BB:BB:BB:BB:BB	2
C	CC:CC:CC:CC:CC:CC	3
D	DD:DD:DD:DD:DD:DD	4
E	EE:EE:EE:EE:EE:EE	5
F	FF:FF:FF:FF:FF:FF	6

With the above forwarding table, when Station A transmit a packets to Station C, the switch looks it up in the forwarding table, and identify the location of Station C as port #3, and forwards the packet ONLY on port #3. This is the primary difference between a hub and a switch. In the case of a hub, the packet would have been transmitted across ALL ports while the switch ONLY transmitted to port #3.

No longer is it a shared medium, because the workstation will only see traffic destined for it – versus ALL traffic in the case of the hub. Therefore, the bandwidth(10Mbps) is now dedicated and available to each workstation/port.

However, one important note – should the switch receive a packet, and the destination address is not present in the forwarding table, that packet will be forwarded onto all ports. Once the destination station responds, then its MAC address will be added into the forwarding table, and normal switching will resume.

Broadcast traffic are always transmitted to ALL ports on the switch. In a Program Automation network for example, Program Automation Servers(PAS) will periodically refresh its MAC address table around every five minutes. The refresh involves clearing out it's ARP cache, and performing new ARP request for every nodes/server. The ARP request by its very nature is broadcast traffic, so every port will receive a copy of these packets.

Broadcast traffic is a part of the normal operation of a network and the protocols that are deployed. In a typical LAN environment, the network will even be able to sustain a large amount of broadcast traffic. However, in the case of a LANLINK/StarLink which is acting as a bridge, these broadcast traffic will be passed from one side to the other, thereby consuming valuable bandwidth for that split duration. The Design Guide for HD networking calls for the separation of between the normal/Office/PAS traffic and the HD traffic for this very reason. Only HD traffic should traverse from the Studio to the transmitter. Switches will minimize the traffic going to the transmitter by only directing traffic that is destined for the transmitter nodes. However, by its nature, switches will also pass broadcast traffic to the transmitter side also. If the LAN broadcast traffic is unwieldy, then the transmitter site will suffer because lack of bandwidth due to unexpected additional bandwidth overhead. To minimize the broadcast traffic or a broadcast domain, a router/Layer 3 switch is needed to segregate the network.

Deployment strategy –

- Switches should be deployed at both the studio and transmitter site to provide dedicated bandwidth, and to remove collisions from the equation. For the transmitter site, a low cost switch – ie. Dlink, Linksys is more than adequate.
- On the studio side, a enterprise switch(ie. Cisco, Nortel) should be deployed to give greater functionality, manageability, security, and visibility into the network.

V. Routers/Layer 3 Switch

A router operates at the Network layer – layer 3 in the OSI model. It looks at the destination address(ie. IP address) and consults the routing table for a matching subnet/route and forwards the packet onto the appropriate interfaces. Each interface on the router is on a different network, therefore different broadcast domain. Broadcast on "Interface Ethernet0" will be localized to "Interface Ethernet0" and its connecting switch. These broadcasts will not leak over to "Interface Ethernet 1".   Therefore in the case of HD networking, to minimize the broadcast over the StarLink/LANLINK, the network should be deployed as follow.

1) Separate subnet for the office network(10.1.1.0/24) for general users, accounting, etc..
2) Separate subnet for the PAS network(10.1.2.0/24)
3) Separate subnet for the STL/LANLINK/StarLink(10.1.3.0/24)

The beauty of this design is that the broadcast traffic is localized to each subnet – Office to the Office subnet, Prophet to the Prophet subnet, LanLink to the LankLink subnet. Yet connectivity is still maintained. Any workstation, on any subnet, can reach any other station on any subnet. Traffic destined for different subnet needs to be routed.

The disadvantage is that additional subnets are required to support the topology. However, the advantage of a predictable, well behaved network far exceeds the initial setup of the router/subnet.

Layer 3 switch is analogous to a router, but only with "Like" interfaces, ie. FastEthernet/GigabitEthernet, but not Ethernet & T1 CSU/DSU.
A router on the other hand, can support a plethora of different interfaces, Ethernet, FastEthernet, GigabitEthernet, ATM, T1, T3, SONET, DSL, ADSL etc..
Though they both vary in Interfaces support, the routing functionality is the same.

In a typical scenario, a router would be used to terminate the WAN connection(ie. T1), and the core switch would perform the Layer 3 routing for the internal network/VLAN, and have a default route pointing to the WAN router.

VI. Equipment Recommendations

The criteria for the switches located at Studio are

- Manageability
  - SNMP v2 support
  - Interface statistics
- VLAN support for broadcast isolation
- 802.1q trunking support
- Layer 3 switching
- Security
- 10/100 wire speed switching/routing

The criteria for the switches located at Transmitter are

- 10Mbps wire speed switching/routing. In reality,10Mbps is more than adequate since the different service modes in HD radio never exceeds 512K/s. *

* These are the minimum criteria. Should manageability, and security features are needed, then a higher end switch is need to support these features.

Studio	Transmitter
Cisco Catalyst 3560-24TS(L2/L3) • 24 Ethernet 10/100 ports and 2 small form-factor plugaable(SFP) ports	Linksys EtherFast Cable/DSL Router with 4 Port Switch (BEFSR41)
Cisco Catalyst 3750-24TS(L2/L3) ▪ 24 Ethernet 10/100 ports and 2 small form-factor pluggable(SFP ports	DLink 5 port 10/100 Desktop Switch(DES-1105)
Catalyst 4507R(L2/L3) Enterprise switch – • 7 slot chassis with redundant power supply, supervisors,	Cisco Catalyst Express 500 – • 24 10/100 ports and 2 10/100/1000 Base-T Uplink(WS-CE500-24TT)
	Cisco Catalyst 2960 – • 24 10/100 + dual Purpose Gigabit Ethernet Uplinks (WS-C2960-24TC-L)

VII. Security

Security at all locations (PA, NJ, NY, CO, OK) is practically non existent. The only measures of security are the telnet & enable password. Even then, the password is the same throughout the whole company/stations so if a hacker gains an entry point into any of the myriad of radio stations, then full access is available. This engagement was not a security audit. A full security audit is highly recommended in order to mitigate any attacks, and minimize entry points of compromise. Having said that, the following security measures should be considered & implemented prior to HD radio going prime time

- Code upgrade – the latest version of code will contain bug fixes & security patches
- Centralized Authentication, Authorization, Accounting(AAA) server – Configure all network devices to authenticate against a centralized database. This could be a radius server in a standalone configuration, or one that ties in with Active Directory. On the radius server, define each station administrators to have their own passwords and rights. Configure their rights & privileges to only have access to their local router & switch, and not any other radio stations. In case of the a compromise on the account, the damaged is limited to the devices that the administrator has rights to versus across the entire network. Moreover, logging can be turned on so that any changes made on the switch & router will be recorded. This can serve as forensic data in case of a network compromise.
- SSH – telnet passes data in clear text. SSH should be configured and used in the management of the switch since the data portion is encrypted.
- SNMP – if SNMP is deployed for the management and polling of statistics, then access list needs to be configured only allowing the network management station access. To prevent further attacks, SNMP read only access can be configured for the statistics polling. No "Read-Write" community will configured.
- Access – access-lists should be configured
  - VLAN -> VLAN
    - Corporate -> HD Radio VLAN = No Access
    - PAS VLAN -> HD Radio VLAN = Restricted Access
    - HD Radio -> PAS VLAN = Restricted Access
    - HD Radio -> Corporate = No Access
  - Telnet/SSH
    - Only certain management station are allowed telnet access into the router & switch

Security - Continued

    - Private VLANs
        - If applicable, map out traffic flow for all HD radio devices and configure access-list accordingly
            - Importer can only talk to Programming automation server & L3 switch
            - LANLINK can only talk to the L3 switch & the other LANLINK
- Unused ports should be turned off.
- Configure BDPU guard to prevent other switches being connected with the administrator consent.
- Configure logging at appropriate levels.
- Configure all router & switch to use Network Time Protocol(NTP) . Use NTP MD5 authentication for secure updates. By synchronizing all routers & switches to common time – forensic troubleshooting with the log files is much easier.
- CDP – Disable CDP to prevent hackers gaining visibility into the networks via adjacencies and neighbors
- Deploy Network Intrusion Detection(NIDS)

VIII. Quality of Services (QOS)

QOS is used to provide differentiated levels of service for selected traffic. In traditional data networking, all traffic and applications are treated as equals and are serviced on a first come first served basis. However, with emerging mission critical applications that requires more bandwidth or a certain Service Level Agreement(SLA), a new model is needed to deliver the expected performance. QOS provides the end to end framework to differentiate traffic and to deliver and guarantee bandwidth to mission critical applications.

QOS is comprised of the following area:

Classification & Marking– what constitute interesting traffic. It could be based on application utilizing well-known ports, flows(source, and destination IP & ports), IP precedence, TOS bits, etc. The idea is to be able to differentiate between the different types of traffic based on a certain criteria.

- **HD networking best practice** – classify traffic on the destination port number. Each model(MP1, MP3) utilizes different port numbers(1011, 1187, 10100). Classification based upon just the destination IP address would not provide the granularity needed to differentiate between telnet traffic, VNC traffic, and critical HD networking data stream.

Once the traffic/application is determined, marking/tagging the packet with an IP Precedence, or a DSCP values will provide a common point of reference to enforce a common QOS policy throughout the network.

- ***HD networking best practice*** –
  - Classify Data Stream(MP1, MP3, etc) with a IP Precedence value of 3
  - Classify Telnet traffic with an IP Precedence value of 2
  - Leave all other traffic(ie. VNC) untagged – IP Precedence value of 0

Shaping & Policing–

- Bandwidth policies – how much bandwidth will be allotted to each application.
  - Percentage based – ie. HD traffic should have 50% of the link speed
  - Hard-coded – ie. HD traffic should have 65K for MP1 mode
  - Default – Best effort – VNC traffic will consume any remaining and available bandwidth once higher priority policies have been met.
    - ***HD networking best practice*** – used hard-coded values based on upon the expected bandwidth consumption table for the different modes.

- **Congestion Management** – how to prioritize traffic in an event of congestion. The method employed and available will be highly dependent on the networking manufacture. For example, Cisco offers
    - Priority Queuing – Always service the higher queue first – this could lead to the starvation of the lower queue.
    - Custome Queuing – Assign different queues to different type of traffic, and service these queue in a round robin fashion.
    - Class Based Weighted Fair Queuing – Traffic are classified into classes, and each class can be guaranteed with a minimum amount of bandwidth.
    - Low Latency Queuing – provides Priority Queuing with Class Based Queueing.
    -

These are some of the queuing mechanism available from Cisco today. Configurations of the QOS policies are beyond the scope of this document since many of the features are dependent on hardware platform and code version. Further explanation and configuration guidelines can be found on Cisco's website at

http://www.cisco.com/en/US/tech/tk543/tk544/tsd_technology_support_protocol_home.html

- ***HD networking best practice*** – Utilizing Low Latency Queuing(LLQ) giving HD traffic the hard coded bandwidth values under the priority class. All other traffic will utilize the default class – which is best effort. The most critical place where QOS must exist is the transition from the LAN to the WAN – usually a router. In the LAN environment, even with 10Mbps – which is old technology, is more than plentiful for HD application. However, in the WAN, 64K is a precious commodity so appropriate QOS policy need to be deployed & enforced to meet the SLAs for HD networking.

IX. Remote Management

Each HD component is different, and operates on different platform. The Importer operates on a Windows, while the Exciter operates on Linux with a proprietary management interface.
The best method of management is to be on the console of the machine itself. Network delays are ruled out completely and are not vulnerable to packet capture/decode. However, this might not be feasible for all HD components since they are spread out between the studio & transmitter.

The Importer & Exporter are on Windows so Remote Desktop(RDP) can be use to manage the machine.

Exgine is based on Linux – Management is provided by a Web interface which in turn spawns off a Java VNC applet. It provides to the end user the same look of the console of the Exgine. However, due to the GUI interface(VNC) one needs to be aware that each VNC session consumes 60K/s. If the bandwidth usage is borderline already, using VNC to remotely manage an Exgine will introduce errors and retransmissions into the network. To minimize such effect, Telnet which only consumes 5K/s should be used. However, since telnet requires Linux familiarity, and intimate knowledge of the operating systems and file locations, it is much harder to administer. Maybe the middle ground is to find the top 5 functions that are frequently used in management of the Exgine via the console, then provide steps/procedures/manuals on how to achieve the same result via command line.