The Local Oscillator

The Newsletter of Crawford Broadcasting Company Corporate Engineering

DECEMBER 2024 • VOLUME 34 • ISSUE 12 • W.C. ALEXANDER, CPBE, AMD, DRB EDITOR

Another One Gone

Here we are at the end of yet another year. Seems like it happens more and more quickly these days, no doubt a function of my age, but I also think that technology has a lot to do with it. Things just move faster in this day and age. We can achieve a lot more in a shorter time than we could in earlier days, and just seems to me that those achievements should have taken a lot longer.

Relax. I'm not going to give you a year in review recap here. Past issues of this newsletter are readily available for you to do that on your own if you have a short memory. But I will say to our first-rate engineering crew, well done! I'm proud of each and every one of you. There's not an engineering/tech team in the radio industry that can top you and what you routinely do every day. It's a blessing to work with each of you.

I will, however, give you just a little peek behind the curtain, and you can pay attention to that man pulling the levers. His name is Zetta, and he's coming to a radio station or cluster near you in the coming year.

Last spring, we got word from RCS that Nexgen, the playout system that we have used since 2003, is being sunset at the end of 2025. It will continue to function and have some support, but there will be no more updates and the support clock will be ticking down. We knew this had to be coming at some point.

Nexgen has its roots in the old Prophet Wizard system and still uses some of those software components. I remember being awed by the Nexgen platform when I first looked at it those 21 years ago. Our deal with Dalet Digital Media Systems was winding up and we either needed to stay with it in its new iteration or move to another platform. It was still early in terms of playout systems – they really hadn't been around all that long, and of course the internet was mostly a thing of the future.

I had a parade of demos come through our conference room. All the big names of that day: Scott Systems, Audemat, Nexgen, even Dalet plus one or two others whose names I can't recall at present because they're probably long gone. I had one rule, and everything else was subjective after that. If the demo froze, locked up or malfunctioned, that was disqualifying. I figured if a producer couldn't make the sales demo work perfectly, chances were good that there were problems with the deployed system as well. We did have one lockup, and they were off my list. Otherwise, each had the opportunity to impress me. They all had features and functions that were attractive, but Nexgen was head and shoulders above the rest. Very clearly it had been developed by radio people for radio people. That remains true of RCS today.

With the announced sunset of Nexgen last spring, I had the opportunity to repeat the parade of demos… or not. I chose not. Why? Several reasons. One is our long experience with RCS. We know them and they know us. That counts for something. It counts for a lot, actually.

Another reason was that we could convert from Nexgen to Zetta without having to reload our entire library in each market. That would have been a nightmare and would have taken a very, very long time to achieve.

A third reason stemmed from the demos we had seen of Zetta. Amanda and I took some time to look in depth at Zetta at the spring NAB convention, and we were blown away. If Nexgen was, as I thought in 2003, a space shuttle, Zetta in 2024 is a starship! I can't begin to describe all the great features and functions that it has, things we badly need. Clearly RCS has listened to the radio industry and paid attention to its needs and developed a

product that gives users the options that they need for any format on any station.

Perhaps the greatest intended feature is Zetta is its remote capability with Zetta2GO, which can be run on any browser anywhere on the planet. This feature addresses the needs that arose out of COVID, when station personnel either couldn't by local/state edict or were unwilling to come into a studio to work. That period passed rather quickly, in a few months for the most part, but its lessons carried forward. It showed us some things that we see reflected in our industry as well as in business as a whole to this day, namely that we can function just fine if properly equipped without working in a studio or office. Those same features of Zetta can be leveraged to permit a single operator to manage the operations of multiple radio stations from anywhere, That has value.

Finally (but far from ending the list of features), Zetta has amazing optional disaster recovery features. If an asteroid were to crash into one of our studio buildings, with internet service at a tower site, a station that subscribes to that option could be back on the air in minutes and operate with full functionality while the studio operation is restored… or not.

So what's the plan? In 2025, we plan to convert all our larger multi-station markets to Zetta, tentatively in the following order: Denver, Birmingham, Buffalo, Detroit and Chicago. The three California markets will follow in early 2026. Why start with Denver? It's where we always start with new technology. It gives Amanda and me a chance to be involved from the outset in the installation/conversion process and really learn the product. We have already been through the Zetta and Zetta for Engineers online training courses, so we're somewhat familiar – conversational anyway – with Zetta and know what it can do and how to get things done. We'll put our market engineers through those same courses in the coming months.

The installation/conversion process will include several weeks of on-site work by an RCS field tech. It's not clear to me how all this will work and how we'll stay on the air through the conversion process, but that's why we're starting in Denver. Downstream markets will have the advantage of knowing all this before their turn comes. We'll (hopefully) learn from our mistakes and figure out what we can do ahead of time to make the process easier, faster and more seamless.

I will, as we progress, keep our engineering team posted on what we learn and advise what they can be doing to prepare for their own installation/conversion, and as mentioned above, we'll hook you up with some specific and valuable product training.

In the meantime, I would send you to the following link where you can get a look at Zetta: https://www.rcsworks.com/zetta-a-walkthrough-for-new-users/. I think you'll like it.

Dashboard Candy

We have over the past year added some dashboard eye candy in two of our markets, Chicago and Buffalo. Chicago was doing this already using Arctic Palm, but we made the move to TRE+. We did the same thing in Buffalo and just recently got the new importer installed and the whole platform running. As reported in prior issues of these pages, we now have not only dynamic, interleaved RDS messaging running in Buffalo but also now HD graphics.

This month we add Detroit (and specifically WMUZ-FM) to our list of stations with this capability. We have a new-this-year exporter/importer in place, so that major equation variable has been solved. All that we need is the TRE+ RDS and HD Advertiser packages and some means of generating dynamic RDS.

I was surprised to learn that we did not have dynamic RDS in place on WMUZ-FM. They use AES192 for composite from the Omnia.11 processor to the Nautel exciter, and they have long used static RDS generated in the Omnia in that mix. We could have used dynamic RDS with the built-in Nautel RDS generator, but it does not support the DPSTEXT= function whereby metadata is sent to both DP and RT fields, which is not only needed for universal display on all RDS-equipped radios but also recommended as a best practice by the NAB for that very reason.

To address this issue, we have ordered an Inovonics RDS generator, and we'll both shut off the static RDS in the Omnia and connect the Inovonics generator to the external subcarrier port on the GV40 exciter. That is the final piece of hardware needed to do what we need.

In the coming days, Mike Kernen and I will get TRE+ installed and running and we can begin testing our new dashboard capabilities. We may in the months ahead implement the same capabilities to other FM stations, probably in Birmingham. We've got a lot of dashboard competition these days, so we've not only got to sound great but also look great.

As we wrap up 2024 in these pages, I wish each of you a blessed Advent Season and a happy and prosperous New Year!

The New York Minutes
By
Brian Cunningham, CBRE
Chief Engineer, CBC – Western New York

Hello to all from Western New York! So far this winter we have indeed been fortunate in the fact that we have not experienced any snowfall here in the Buffalo/Rochester region, which is quite uncommon for this time of year. We did not experience any type of winter weather until the morning of Thanksgiving, which were very light flurries that went away about mid-day.

Forecasters are predicting heavy lake effect snowstorms for the upcoming weekend as I write this, with totals expected to be over four feet in some areas (mainly in the south towns, where our transmitter sites are located for WDCX-FM and WDCZ).

For the most part, our transmitter sites have been winterized and prepared for the long harsh winter months, so I am not too concerned about the possibility of a several-day lake effect snowstorm. Our generators have been tested and are ready for perhaps an extended run, in the event of a lengthy power outage.

You may recall last year we had several record-setting snowstorms, dumping over eight feet of snow just south of Buffalo over a three-day period. We were able to ride these storms out with no off-air instances. Those who could not make it in due to the travel restrictions were able to log into our network from home and perform their duties remotely.

In last month's report, I wrote about some interference that we were experiencing on WDCX-FM that was sporadic and could only be heard on some radios. I have gone through our audio system with a fine-toothed comb and am pleased to report that our equipment is not the cause of the audio anomalies that some can hear. Everything from the NexGen output to the input of the exciter is perfectly clean and noise-free.

We did discover that the interference was localized within several miles of the studio location in Amherst. If you travel to the north towards Niagara Falls or southward through Buffalo, the noise goes away, so there is something in the vicinity of the studios that is injecting noise on our frequency.

I have an IFR communications service monitor that is used for repair/confidence monitoring of our radio signals, and one of the options is a "Look and Listen" feature that allows you to "see" the waveforms generated in the FM band and hear the audio being broadcast. I attempted to take the IFR out into the field to see if I could determine the cause of the interference but was unable to do so as the audio portion of the "Look & Listen" feature would not work; no audio output was noted at all. I will have to send the communications monitor back to the mothership for repairs before I can continue to track down the source of the interference.

We at first speculated that the noise was being generated or caused by the audio settings in the Tieline Gateway (the head-end of our STL system) or the Bridge-It XTRA located at the transmitter site. Within hours of last month's report being made public, I received a telephone call from the technical support department of Tieline offering any assistance they could provide in helping to narrow down the noise. Folks, that is outstanding customer support! In my 50+ years in broadcasting, I cannot recall any equipment vendor providing this level of technical support, with a true concern for isolating and finding a remedy for the problem.

Although the Tieline products were found to be clean and problem-free, they were there to assist in any way possible. My hats off to everyone in the Tieline support department, true professionals in every way! Thank you, guys, your assistance was truly appreciated!

On Friday night the 8th of November, I called into the Burk remote control to check the WDCX(AM) night antenna readings, as I do each time there is no live, in-person operator at the studio. I found the readings to be way off, so a visit to the site over in Clarkson, NY was warranted early the next morning.

I arrived early on Saturday morning and found that the RF switch at tower #4 that switches the pattern between day and night modes had failed. This is the third switch I have had to rebuild this year. Early in the summer, we lost the switch at tower #5, and not too long thereafter, the switch at tower #6 failed.

We purchased this station in 1999, so these switches have been operating twice daily since that time. I have no service records that precede the purchase date, I can only suspect that these switches have been in operation since the site was built, so I can understand the cause of the recent failures, and attribute it to age.

The tower #4 switch rebuild went very smoothly, as I now have the procedure down to a science. The RF switch located at the reference tower (#3) was rebuilt several years ago, and I recall it was in the dead of winter (with over three feet of snow on the ground). Walking between that tuning house and the transmitter building was treacherous due to the heavy snow, and something I do not ever want to repeat if I can help it!

The switches at towers #1 and #2 have not shown any issues, but I am seriously thinking about rebuilding them before they do fail. You know how it goes – Murphy's law comes into play: they WILL fail at the most opportune time!

Speaking of failures, we had the HVAC unit fail recently at the WDCZ transmitter site in Hamburg. Early in the summer, the motor failed in the condensing unit located on the roof of the transmitter site. Solly Industries, our HVAC contractor for nearly 30 years, was called in to investigate the recent failure also and found that the fan motor in the indoor unit had locked up and burned. The technician did not have a replacement motor on his truck, so they had to obtain a replacement and come back the following day to complete the repair.

We will be looking to replace this HVAC unit in next year's budget, as it is operating at present well beyond its life expectancy. It was installed in 1999, according to records found at the transmitter site, and life expectancy of these units is estimated at or near 20 years of service. We'll keep our fingers crossed that it will hold out for at least another year.

That about wraps up another month here in the great Northeast, and until we meet again here in the pages of *The Local Oscillator*, be well, have a very merry Christmas and a happy New Year!

The Motown Update
by
Mike Kernen, CSRE
Chief Engineer, CBC–Detroit

Gratitude

As we pass through the political season and find ourselves at Thanksgiving, I've been keen to notice what makes me grateful. I've always felt gratitude is important and bolsters positivity – we learn to focus on good in our lives.

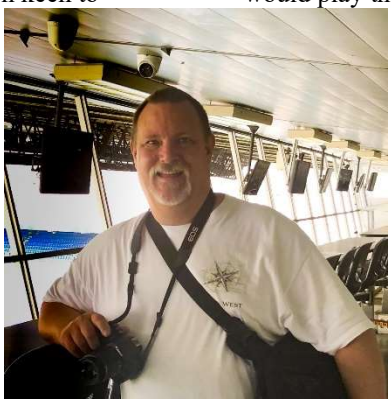

Certainly, I'm grateful for my incredible wife, our close family, and my terrific group of friends. I'm grateful for the Crawford family too, and the opportunity to earn my livelihood at their Detroit based stations. I feel gratitude for my group of coworkers, most of whom I'm fortunate enough to call friends. I'm grateful for my health and my abilities.

I can't end a list like this without saying how grateful I am that for the first time in my lifetime I've been able to call the 11-1 Detroit Lions the number one team in the NFL! My dad loves it and has taught me to love it too. My son played junior and varsity football in high school. My grandfather always said that he didn't understand why anyone would play that game twice.

AM Radio for Every Vehicle Act

My son is a firefighter in Port Huron, Michigan, which is a city of about 29,000 people about 30 miles north of the outer edge of Detroit's suburbs. His job is to drive a firetruck to emergency scenes and operate its pump and ladder systems. He is very enthusiastic about his work and has saved many lives while doing it. It makes me reflect on the importance my work in radio has; what we do to make people's lives better.

While we don't have as direct an effect on people as would a fire and rescue worker, we do have an effect, and it's a vital one. Radio serves the community by providing listeners with information. From something as immediate and urgent as public

safety warnings, to inspirational and motivational content. Radio is a part of people's lives, their communities, and crucially, their safety. I hope you'll take the time to visit this website: https://www.nab.org/amToolkit/default.asp and encourage congress to keep AM radio in our cars where we consume and rely on radio the most. It's unthinkable that even one person could miss vital life safety information because their car's manufacturer saved a few cents on their radio.

Fire the Firewall

Every month I spend untold hours on IT. That's fine with me – our industry has leaned into IT for just about every system and subsystem we rely on and many things we don't.

I love computers and have been drawn to them for as long as I can remember. I remember my high school computer lab and programming with punch cards. I remember the first PC I ran across in radio, too. It was there for the sales department – I can't recall what they did with that first generation IBM, its green monochrome screen and 16 kilobytes of RAM memory. What really made an impression on me was the mainframe we had that had a 10-megabyte hard drive the size of a washing machine.

Today, of course, we always have computers infinitely more powerful in our hip pockets. Our reliance on data is ubiquitous. Recently the need for a new network firewall became evident because the current software reached end of life. The firewall is essentially a router, providing our gateway to the internet and a pathway between our networks. Replacing said device requires careful coordination because just about everything and everyone uses the network heavily. We also use the internet for some of our studio to transmitter links (STL) and our websites' 'listen live' feature.

Careful and complete programming of the new router should allow for a simple, lightning-fast swap. Unfortunately, the new router would be the same Dell single rack unit server class machine that the old one was. This meant we'd remove the old computer and rebuild it with a new operating system and new firewall software. This would be akin to changing a tire while driving down the road!

Foreseeing this would be a problem in several markets, Cris Alexander, Crawford's Director of Engineering, purchased a tiny self-contained firewall to bridge the gap while we rebuild the Dell hardware.

The micro firewall made its way to Detroit but refused to work on our network. Possibly because prior stints in other markets had forced programming adjustments that were incompatible with our market's topography. Factory resetting the unit was called for and easily accomplished, leaving much setup to be redone. Gone were such niceties like routing between networks, and the firewall blocked all communication from everywhere outside.

With much needed and appreciated help from my counterpart in Birmingham, we got the backup firewall installed and operational. I even created another network I plan to use for access and configuration of the firewall while keeping the interfaces I needed for our standard networks isolated to avoid duplication of IP addressing.

Once the firewall was set up, I downed the switch ports connected to our existing firewall and brought up the ports that I had recently connected to our temporary firewall. This worked beautifully and no one was the wiser… except… For some reason, some network attached devices didn't want to talk to their internet attached services. I'm not sure why, but these required power down restarts to begin to work with the new firewall hardware. I can only guess this had something to do with their limited operating systems' ability to recognize the need to flush an ARP cache.

Building the permanent firewall using our old hardware was simple. I wiped its hard drive, created a bootable flash drive, and installed the new firewall software. Restoring a backup of the temp unit's settings finished the installation. I also relocated the unit to our TOC from the telephone/electric closet using a VLAN for the internet connection back to the ISP handoff point.

The new firewall is running smoothly, and I dare say it is much quicker than the old one. It also will be supported by the manufacturer, keeping us up to date with security and features for the long term.

VPN Woes

A few hours after installing the temporary firewall, I realized one of our two VPNs was not functional. We use our VPN for secure access to internal resources and PCs that our radio hosts rely on to do their shows remotely. For example, our call screening software runs in the talk studio and when our afternoon host works from home, he accesses it remotely via our VPN.

These VPNs our created by using Raspberry Pi hardware and twice now I've had one corrupt the micro-SD card that it uses as a hard drive. I was confident that rebuilding the Pi would be a simple task since I had backup files for all users' VPN tunnels, but what I didn't know was that reinstalling

the VPN software would change the private key, negating every tunnel's encryption key.

Once I finished building up a new micro-SD card for the Raspberry Pi, I had to redistribute tunnel keys to every user, in many cases installing them myself. Lesson learned; I now have clones of my micro-SD cards!

News from the South
by
Todd Dixon, CBRE
Chief Engineer, CBC–Alabama

And then there was one …

The PfSense firewall journey continued this past month in our Chicago and Detroit markets. As of this writing, the bulk of the firewall work in these markets is complete.

I ran into a problem in both markets due to something I had not had to deal with in Denver, Los Angeles or Northern California. In those three markets, I had tweaked PfSense on our temporary box based on the network configuration that was working in our Birmingham market. Basically, all of the bumps and bruises I had experienced in setting up our initial install here were pretty well baked into those installs.

PfSense, as a package, comes as a bare template, and anything that you want it to do has to be added to it. Some of the simple things that you add are all your local network basics. The more complex things that make the firewall a "fit" for your network are the rules that allow or deny network/internet traffic onto your devices.

Of course, I had slept several nights since battling that here in Birmingham, but so far, all the markets had been using a tweaked version of our working install.

The initial installs in both Chicago and Detroit were both done on vanilla installations of PfSense. Rick Sewell had a redundant backup firewall already in place, so we just installed PfSense on it. I had created a configuration that should have mimicked his firewall there with the six network interfaces that he required here in Birmingham by adding two USB network interfaces to our four-NIC backup PfSense firewall. Of course, there was no way for me to actually test the configuration for internet traffic, as that could only be done on their network in Chicago with their actual service providers.

While we worked our way through the install, getting IP addresses, networks and port forwarding rules in place, each time we attempted a switchover to the new firewall, internet traffic came to a grinding halt. Each time, we'd painstakingly check through settings pages and I would compare them to our working Birmingham firewall to make sure everything was in place. Over the course of several days, we couldn't figure it out and I told Rick I needed a mental break to gather my thoughts on the problem.

Around the same time, Mike Kernen in Detroit had received the temporary PfSense hardware that had been passed between markets, but he began by installing a bare bones instance of PfSense on it and not using the install that was already on it. He did the install without much help from me and only required a couple of questions about how to setup port forwarding rules. Once I helped him with a couple of them, he finished the rest of his on his own.

Each time he transferred his cabling to the new install, internet traffic simply wasn't getting passed. So he was having the same issue that Rick was having in Chicago. What was going on here?

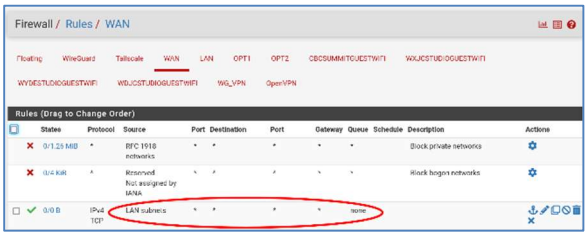

Figure 1 - This simple rule was all it took to make both firewalls pass data.

A number of you likely have been thinking to yourself while reading this, "How dense can this guy be? It is obvious what is going on." It is certainly okay, as I was really asking this question myself as well…

While talking through this with Mike, it hit me – PfSense needs a rule on the WAN interfaces to pass traffic requests from the LAN subnets. A quick look at our WAN interface showed that the rule had been on our firewall here and been passed on to the earlier markets, but both of these markets started with fresh installs not based off of what we had created here. Once the rules were put into both of their firewalls, they simply worked as they should with basic tweaks here or there.

So… the PfSense lesson to the uninitiated: every interface, whether WAN or LAN, requires a rule in order to pass traffic.

Chicago's firewall is working as it should with only a small issue that we have to work out, but all the other primary services that were provided with their previous ClearOS firewall install are working without a hitch. Rick Sewell was really patient with me as I tried to work through this primary issue and then a couple of more minor issues with him. Now that we have the initial install pretty feature complete, he should be able to back up the configuration on it and simply install PfSense on the other and restore that configuration on the other server so that his firewall redundancy will still be in place. He also has another firewall at one of their other sites that we're going address that needs to be upgraded to PfSense as well.

Mike's firewall in Detroit is working well. The only thing that didn't end up making the trip with his new PfSense firewall setup was one of his Wireguard VPN Raspberry Pi units. In a complete coincidence, the microSD card on that unit gave up at about the same time the firewall came online and a new Wireguard install had to be reconstructed. After spending some time trying to resurrect their install by trying to get the Wireguard configuration files copied, he simply set up a new install and had to reconstruct all his client configurations and get them on their machines.

Currently, the only market left that needs a PfSense install is Buffalo. They have a similar setup to Chicago in that when they moved into their new studio location last year, they put primary and secondary firewalls in place in the event of failure. Taking every install experience that we have had into account, we should be able to get their installs done in fairly short order and then this project can officially end...or will it?

PfSense Add On Functionality

Why would we let the PfSense fun end with just basic installs? The PfSense firewall has a number of packages that can be added onto it that increase its functionality.

There are several network traffic monitoring add-ons that help you get a better picture of who is on your network, or more importantly, what is on your network. Everything from intrusion detection to web proxies that cache internet web pages that are visited frequently in order to save valuable internet bandwidth. Two other services that can be added to the firewall are Wireguard VPN and Tailscale.

Most of us are familiar with Wireguard as we are already using it in all of our market environments. Of course, the advantage to running it on our PfSense Dell hardware is that it is not running off a Raspberry Pi using a microSD card. I'm not throwing any shade on the Raspberry Pi. They certainly have no problem with the minimal resources that running Wireguard requires, and they are really almost an afterthought once the service is running on them – except when it isn't (reference Detroit above) and must be rebuilt.

I downloaded the package on our Birmingham install and got a Wireguard tunnel running and it worked great. The resource load for running the service doesn't even create a sweat for the Dell PowerEdge machine.

The only issue I came across, which would be a problem for all of our markets, is that you have to add peer clients manually (including their public keys and pre-shared keys). There is simply no way to export the peer configuration files out of PfSense that I could see. So the user .conf files would also have to be manually set up. Yuck. It is a far cry from our current setups, where users can be added with a simple command and their configurations can be exported with a QR code. It is a package to watch with each PfSense upgrade to see if these pleasantries will be added to a future release.

Another package that provides a service that I think we could get use of instantly is the Tailscale package. The Tailscale service is a way to provide point-to-point Wireguard-encrypted VPN service between an unlimited number of machines. Tailscale is a company, and they provide both personal and business plans, but they provide a personal free tier that allows three users to connect up to 100 devices (what they call a node). Their business plans start at $6 per month per user and go up as more network connectivity is needed.

Setting up Tailscale on PfSense would be a single node, your computer would be another, your phone another, your iPad another, your computer at a tower site another and I think you may be starting to get the picture. All that the Tailscale company is providing is a "coordination" server that sits on the

internet so that a “mesh” network can be created that allows each node to talk to another (or not talk to another as might be the case). The server is only exchanging the public keys with the other nodes so that they can talk to each other directly, all of the other traffic is handled between the nodes themselves.

One of the great features is that a node can be set up as an “exit” node. This might be the PfSense node, or maybe a computer at your tower site that would allow you to exit into your network at the studio or exit to your network at the tower site, thereby gaining access to those networks while completely bypassing your router. Both the coordination server traffic and the inter-node traffic are riding on encrypted Wireguard VPN, so it is completely secure in case you have concerns about a “company” directing your data to its destination.

While not a package in PfSense, there is an open source version of Tailscale called “Headscale” that basically runs the Tailscale coordination server on your hardware with your static IP address. This opens up the possibility of having no limit on the number of nodes and having control over the coordination server, yet it still works the same as Tailscale. It would just take some extra work to get it functioning in PfSense.

While your eyes aren’t completely glazed over from the PfSense talk I’ll stop the computer talk for this month. I hope that all of you have a blessed Christmas season. Until we visit again, may God bless the work of your hands.

Tales From Cousin IT
by
Stephen Poole, CBRE, AMD
CBC Corporate IT Specialist

Happy holidays, another year come and gone! I hope that you and yours have a blessed Thanksgiving and Christmas! And we also shout "Hallelujah" because the election is finally over. I've had my phone set on "Do Not Disturb" for months simply because I was tired of getting repeated texts from politicians begging for money.

However, my joy is muted because I'm writing this in Dwarf Mode™. (Khazad Ai Menu!) The lift cylinder in my desk chair has worn out. I'll pump it up, lean forward and then: "hiss/clunk" ... I'm barely looking over my keyboard. Rinse, repeat, wipe hands on socks. Because I am tough, I shall persevere. (I shall also order another chair. I ain't that tough.)

Anyway. I always look over some older *Local Oscillator* issues before I write for each new issue. It's interesting to see where we've been and what we've accomplished over time. Has it really been three years since I replaced the "trigger" board in WXJC's NX50 transmitter? Yep (2021). And in December of 2020, I recounted how we'd lost a guy wire on tower #5 at Tarrant. Todd introduced the WireGuard VPN to everyone, while Cris was busy designing a three (three!) station directional array at the KLZ transmitter site in Denver(!!).

More Pi, Please

We've previously written (that's a plural "we," not royal; Cris, Todd and others) about the Raspberry Pi in these pages. These little computers are amazing and, as long as you respect their limitations, will give excellent service in network, control and I/O applications. I mentioned last time that Cris and I are working on – among other things – using Pi units to provide quasi-SNMP connectivity for older equipment.

For network applications, the Pi is pretty much ready to go out of the box. We run WireGuard on Pi units in several locations. It's great for that. If you have problems, you can ask for advice in the Raspberry Pi forums; I've found the folks there to be very helpful.

For control applications, you will want one or more "Hardware Attached on Top" (HAT) boards. Digikey and Mouser stock many of these. You'd buy a relay board for control applications; to read voltages, you'll need an analog-to-digital (A/D) converter board. Most communicate via the Inter-Integrated Circuit protocol, or "I2C," usually pronounced "Eye-Two-Cee."

In this issue, I wanted to cover just what this is and how you can troubleshoot common problems. To keep this brief, I'll stick with the "classic" Pi 3, 4

and 5 families. Smaller Pi units are available, as are clones of the Pi that come in the same 85.6 mm × 56.5 mm (3.37 in × 2.22 in) package. These will have different requirements.

Be warned that Cris, Amanda and I have not had a good experience with some of these clones. They seem to be OK for standalone network use, but we've had problems getting one of them to work with a Pi HAT. To be blunt, we ran in circles for a while trying to figure that one out. I'd recommend sticking with the Genuine Pi from a reputable vendor.

I2C

I2C is a 2-wire protocol: one wire for clock, the other for data, providing serial communications and control. We shan't be naughty and refer to the arrangement as "master/slave," though. The preferred terminology is "controller/target" nowadays. However you say it, it's a very common protocol and one that's fully supported by the Raspberry Pi. Subordinate protocols include the System Management Bus (SMBus), Power Management Bus (PMBus) and the VESA DCC (Display Channel Control) bus. Google it or go to Wikipedia for more info, but these are simply different flavors of I2C. We'll concentrate on basic I2C with the Pi models 3-5 here.

An I2C setup can have multiple controllers and targets. You will most likely be running a single controller (the Pi) with one or more HATs to interface to the outside world. Multiple connections on a single data/clock bus are allowed because the chips stay quiescent unless and until they're actually addressed. Open-collector/drain and tri-state I/O is used to avoid loading the bus for inactive devices.

I2C On The Pi

You first need to enable the I2C interface on your Pi, if you haven't already. Run "sudo raspi-config" in a terminal or choose "raspi-config" in the GUI desktop, then choose "Interface Options" (see Figure 1, using the command-line interface). In the next window that appears, select "enable I2C." The wording and selection number may differ between Pi models, but "I2C" is what you want. Enable it, reboot and then you're ready to test.

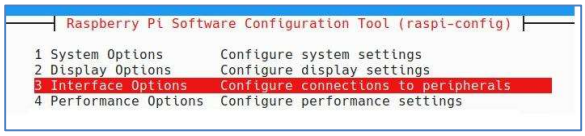

Figure 1 - Via 'sudo raspi-config' in a terminal, enabling I2C

I2C is designed for relatively low-speed control applications, and specifically for intra-board communication between chips. It's not suitable for running long distances. Simply put, if you have a Great Idea™ for putting the Pi over here and the HAT over there, with a lot of shielded cable to interconnect things, it probably won't work. The data and clock lines are unbalanced and are subject to interference. Plus, the total capacitance on these is limited to 400 pF, and the lower, the better.

You can stack several HATs on the same I2C bus; just keep in mind that each is pulling power

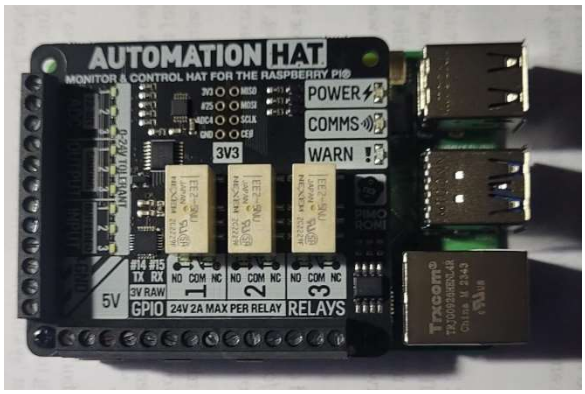

Figure 2 - The Pimoroni Automation hat mounted on a Pi 5.

from the Pi, and it's not unlimited. You also must use separate hardware addresses for each. Refer to the docs for your board(s), but this is typically set via jumpers or DIP switches. Using the same address for two different HATs will gain what the Neckbeards call "undefined behavior," which is a fancy way of saying, "Won't work, Cuz."

The Raspberry Pi 3, 4 and 5 default to a 100 kHz clock/baud rate. Unless you're doing something way out in left field, you shouldn't change that. On my Pi 3B+ test unit, I can see the baud rate in the binary file at /sys/bus/i2c/devices/i2c-1/device/of_node/clock-frequency, and you can supposedly change with a boot-up option. Let's move on.

Programming

We'll have more to say about this in the future. At present, we're using Python for most work; I've also written a few small modules in C. You'll find gladsome descriptions online that claim, "Python is easy to learn!" Well ... yes and no. It's easier than C++ or Rust, but it has its quirks. Working with binary can get very tedious and it's easy to make mistakes. The syntax is picky. Indentation is used to

mark the beginning and end of a "block" of code and it's easy to misalign something.

Seeed Studio (seeedstudio.com – note 3 e's!) has a bunch of tutorials and "how-tos." I'd recommend that you start there. Better yet, for their own add-on boards, they have a lot of ready-to-run examples to play with.

Troubleshooting

First make sure that everything is set up properly. Did you confirm the addresses? Are you "talking" to the correct bus? From my own experience, this is by far the most common failing. There are two parts to each address: the bus number and the target device address. Bottom line, if the HAT is on bus 20h, you'll get nothing if you send commands to bus 21h.

As already mentioned, the I2C bus is quite simple: all attached devices are supposed to present a high impedance unless and until they are specifically addressed and/or told to send something. Ergo, one obvious check with multiple HATs is to try them one at a time when running down comm failures.

Figure 3 - This Pi-4 and relay HAT is the one we're using in Denver to test SNMP functions for legacy transmitters.

The software and documentation provided with most relay and data boards varies from decent to horrible. Seeed Studio, already mentioned, gets a nod from me on this. For other makes, you'll need to do some Googlin'. I found this to be the case with Pimoroni, a company in the UK that makes some nice-looking boards (see Figure 2, a pic of their so-called "Automation Board," which includes everything but a coffeemaker: relays, A/D, and I/O on a single HAT).

If you suspect the board itself, unless you have a top-shelf data/protocol analyzer (and a DEEP knowledge of the I2C protocol and a LOT of patience), the best way to run down a problem with a "HAT" is to swap in a known-good replacement. Fortunately, most of the available boards are relatively inexpensive, so I usually order two of everything.

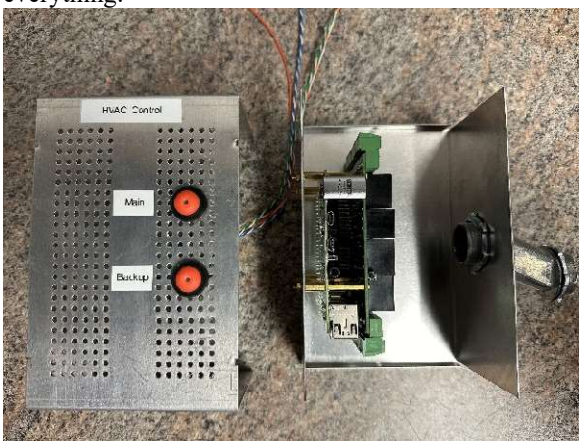

Figure 4 - A work in progress, this Pi-4 and 8-relay HAT will be used for main/backup HVAC control at the KLTT transmitter site.

One final note: people have a tendency to blame the processor and/or software if something isn't working. If you insist, take an oscilloscope and look at the I2C "clock" line; if you see pulses at approximately 100 kHz, I2C is enabled and the Broadcom CPU is trying to communicate. From my experience with the Pi, though, it's usually going to be something silly, like an address jumper that's set incorrectly.

Final Thoughts

We work for the best company in broadcasting. As I write this, some friends have lost their jobs with Cumulus and the trades have been reporting that other groups have been selling stations. I thank God daily for Mr. Crawford (both Senior and Junior!), Frank and Brett, Cris (of course!), Todd, Amanda, Jack ... way too many to name here. And when I say I hope you have a merry and blessed Christmas, that's not just idle chatter. God bless each and all y'all (a little Southern lingo for you, there).

Until next time, keep praying for this nation!

The Chicago Chronicles
by
Rick Sewell, CSRE, CBNT, AMD
Engineering Manager, CBC–Chicago

To Pre or not to Pre?

Over the last year I have had the opportunity to sit on webinars and seminars about air chains and processing. These have been conducted by industry experts who have been at making audio sound good on radio while I was still working in the production and programming side of the business. So, I definitely want to listen to what they have to say.

In the meantime, our engineering director, Cris Alexander has also been taking part in similar seminars. We have had discussions about some of the long-held ideas about pre-processing before STL equipment.

Certainly, there have been differing ideas on the value of pre-processing for a long time. I have always sided with the idea that "subtle" limiting has great value before sending audio into the STL chain to protect from board operators not watching their levels. Even more so with stations having walk-away time and commercials being automatically loaded into automation systems. The levels of materials can be all over the place.

Back up in time, to the advent of PPM encoders now being in the air chain. It became very important that stations have very consistently strong levels going into the PPM encoder for the audio to properly encode the PPM watermarks so that PPM monitors pick them up in varying listening environments. The idea of pre-processing before the PPM encoder was now becoming a standard idea if you didn't put your PPM encoder as an insert to your main processor.

The real drawback to having multiple processors in your audio chain is that compression/limiting in multiple stages can work against each other, and you end up having a "pumping" of audio levels. This is something that you do not want and can give you a product that is unlistenable or at best, a subtle listener irritation.

In mid-November, the conversation came up again and we decided to move away from pre-processing. We had been using the built-in EQ sections of the Wheatnet blades to be the pre-processing in our audio chain is pretty much entirely AoIP after we replaced our profanity delays with Eventide delays that have AoIP cards.

It didn't take long to make the move since there was no physical wiring to be moved. We just used the crosspoint map in the Wheatnet Navigator software no one was the wiser that we made the move.

We began to notice that there are some results that we didn't expect. The pre-processing was certainly masking some issues that we had that needed to be corrected. The first was that there are some operators that are not watching their levels very well. It was also noted that we have problems with source audio in the automation system. Again, the levels of audio files are wildly different. We are investigating this to see if this is the local production staff or the auto loaded spots from outside the building.

I wasn't surprised that we would see those type of issues. What I didn't suspect was that we would see some problems with our main processors. These were most likely being hidden by the pre-processors in the Wheatnet blades.

It didn't take long before I heard from staff members that one of our stations was experiencing sudden drops in levels. We saw issues on our main and auxiliary processors for this station. The aux processor actually looked like it was in bypass mode, even though it wasn't. The output of the processor was tracking very closely with the studio levels and as mentioned previously these are not very consistent.

At the time of this writing, we are working through the issues. I am also watching the encodability of PPM encoding with the Telos TVC PPM Encoder monitors and the confidence meters in the Nielsen Multi-Encoder Monitors. We may have to consider putting the PPM encoders in as an insert to the main processor.

As we go down this path further, I am sure that we will get most of the problems resolved. The human part of operators and production staff watching their levels, is certainly the most difficult to challenge.

Rocky Mountain Ramblings
The Denver Report
by
Amanda Hopp, CBRE
Chief Engineer, CBC - Denver

Fun Times at KLDC!

Wow, is 2024 almost over? As I write this, we are preparing for Thanksgiving. While always fun to celebrate with family, it marks the start of a crazy season. Thankfully, around Denver, things have been slow on the engineering front.

We had our first big snow at the beginning of the month. I think when it was all said and done, we had about two feet at our house. During that storm, I noticed KLDC's microwave link went down. I did not worry about it too much as I assumed it was snow related – the STL seamlessly switched over to the internet backup.

On that Saturday, after things did not come back up, I went out to the site. There was still quite a bit of snow on the ground, but I looked around to the best of my ability. I tried a power cycle of the unit and found it coming and going several times before it finally just went. I could not really get to the tower as it was a bit unsafe, so I called it a day.

The electrolysis-blackened RJ45.

The following Monday I was able to make it back to the site. Things had dried out quite a bit. We quickly found the RJ from the cable that goes to the building, which plugs into the Kintronics CAT6 isolation coil was black. Not good.

There was some moisture in the connector box, which was a bit odd. We are not entirely sure how that happened. It was probably condensation. But clearly water got in and electrolysis started where the RJ made contact burning, damaging the contact fingers. There is 24-volts PoE power on those contact fingers.

We were able to cut the connector off and found no damage to the cable itself. I put a new connector on, and things came right back up.

We still need to find a way to protect the connections from moisture. We are thinking of trying dielectric grease.

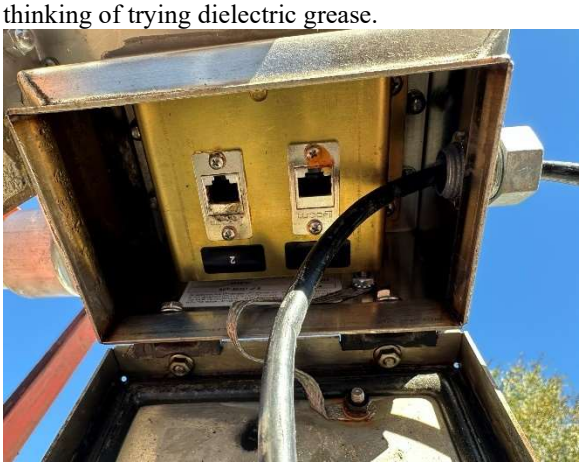

Looking up into the connector port of the Kintronics CAT6 isocoil. Things were damp in there but dry in the opposing compartment on top.

Office Upgrade

We are finally done with the office/studio upgrade. As part of our lease renewal, we got a new paint job as well as new carpet throughout our entire suite. They did it in three phases: open areas, offices, then studios. It proved to be a chore.

With my online training completed, I'm certified!

Thankfully, many used it as a time to clean house and purge things no longer needed. I even went back through some stuff and found items that were outdated and no longer needed. It feels good to be able to get rid of stuff that is no longer needed.

Coming Up

There is still some work to be done outside. KLTT has some copper strap that has been exposed over the years and the horses that wander snagged it and it eventually got ripped up. We have what we need to fix it all, but we have been waiting on conditions. It is a fine line between fire danger and too muddy to work. I hope that sometime in early December I can take some time with my dad to go out and do the work. It is not a one-person job. We also need to do a tack weld on the Austin ring transformer where we recently had to replace a union.

Christmas is fast approaching! I pray you all had a great Thanksgiving and have a blessed Christmas!

KBRT • Costa Mesa - Los Angeles, CA
740 kHz/100.7 MHz, 50 kW-D/0.2 kW-N, DA-1
KNSN • San Diego, CA
1240 kHz/103.3 MHz, 550W-U
KCBC • Manteca - San Francisco, CA
770 kHz/94.7 MHz, 50 kW-D/4.3 kW-N, DA-2
KLZ • Denver, CO
560 kHz/100.7 MHz, 5 kW-U, DA-1
KLDC • Denver, CO
1220 kHz, 1 kW-D/11 W-N, ND
KLTT • Commerce City - Denver, CO
670 kHz/95.1 MHz, 50 kW-D/1.4 kW-N, DA-2
KLVZ • Brighton-Denver, CO
810 kHz/94.3 MHz/95.3 MHz, 2.2 kW-D/430 W-N, DA-2
WDCX • Rochester, NY
990 kHz/107.1 MHz, 5 kW-D/2.5 kW-N, DA-2
WDCX-FM • Buffalo, NY
99.5 MHz, 110 kW/195m AAT
WDCZ • Buffalo, NY
950 kHz/94.1 MHz, 5 kW-U, DA-1
WDJC-FM • Birmingham, AL
93.7 MHz, 100 kW/307m AAT
WCHB • Royal Oak - Detroit, MI
1340 kHz/96.7 MHz, 1 kW-U, DA-D
WRDT • Monroe - Detroit, MI
560 kHz/107.1 MHz, 500 W-D/14 W-N, DA-D
WMUZ-FM • Detroit, MI
103.5 MHz, 50 kW/150m AAT
WMUZ • Taylor - Detroit, MI
1200 kHz, 50 kW-D/15 kW-N, DA-2
WPWX • Hammond - Chicago, IL
92.3 MHz, 50 kW/150m AAT
WSRB • Lansing - Chicago, IL
106.3 MHz, 4.1 kW/120m AAT
WYRB • Genoa - Rockford, IL
106.3 MHz, 3.8 kW/126m AAT
WYCA • Crete - Chicago, IL
102.3 MHz, 1.05 kW/150m AAT
WYDE • Birmingham, AL
1260 kHz/95.3 MHz, 5 kW-D/41W-N, ND
WYDE-FM • Cordova-Birmingham, AL
92.5 MHz, 2.2 kW/167m AAT
WXJC • Birmingham, AL
850 kHz/96.9 MHz, 50 kW-D/1 kW-N, DA-2
WXJC-FM • Cullman - Birmingham, AL
101.1 MHz, 100 kW/410m AAT

Crawford

Corporate Engineering

2821 S. Parker Road • Suite 1205

Aurora, CO 80014