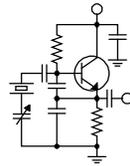


The Local Oscillator



The Newsletter of Crawford Broadcasting Company Corporate Engineering

MARCH 2026 • VOLUME 36 • ISSUE 3 • W.C. ALEXANDER, CPBE, AMD, DRB EDITOR

A Little AI Fun

Eagle-eyed readers will recognize that the columnist photos in this edition are a little different than they have been in the past. I was experimenting with AI last month and found that I could use it to produce some fun and really creative caricatures from photos. Naturally, I had to run our whole *Local Oscillator* crew through that engine, and what you'll see in these pages are what I came up with. So enjoy the new pics! Any resemblance to living persons is completely intentional!



NanoVNA

A couple of months ago, the Alabama Broadcasters Association asked me to do an engineering webinar on AM directional antennas. I was honored by this invitation, and I did that webinar on the 17th of last month. It was, for the attendees, probably a lot like drinking from a fire hose, a lot of information in a one-hour presentation, but it was well received.

In that webinar, I provided a helicopter view of the purposes of AM DAs, along with a brief look at how they are initially adjusted and proofed (both conventional and moment method). I then continued by talking about what to do when a monitor point of a conventionally proofed array goes high. Much of this was really a plug for the advantages of moment method proofing, and there are many.

Then I got into some troubleshooting tips and techniques – what should you do when (not if) the array goes out of adjustment with one or more parameters out of tolerance on the antenna monitor or common point meter. There are just a few possible culprits in such a situation.

The antenna monitor could be lying to you. I've seen that happen many times, more with the old

AM-19 analog monitors than the 1900 digital monitors, but it can happen to either one.

The antenna monitor can be telling the truth, but there can be a problem with one or more of the samples, either base or loop. How can that be determined?

We can have a problem of some sort on the tower, such as some foreign object in the base insulator, a lightning damaged isocoupler, one or more broken guy insulators shorting the tower to a guy wire segment, or even corona damaged fiberglass rod insulators.

We can have a bad component – dollars to donuts it's a mica capacitor if it is a component. Those are the most common failed component in my experience. Sometimes those are painfully obvious. Other times, not so much.

We can have a pattern switching problem, an RF contactor not fully seating in the correct position or some control problem.

Or we could have a bad transmission line, maybe a wet line or one that has internally arced and is effectively shorted.

How do we figure out what the issue could be? Some are obvious and can be found by a visual inspection, by swapping inputs on the antenna monitor, or by looking at indicators on the control system. But what about hidden issues, such as issues with sample transformers, isocoils, sample and transmission lines?

We are fortunate in this company to have some really great test equipment. For many years, I used a General Radio bridge to tune and troubleshoot antenna systems. A couple of decades ago, vector network analyzers (VNAs) came into vogue in AM work, pioneered by the late Ron Rackley, and we bought a four-channel HP analyzer that weighs about

100 pounds, along with a magic bag of tricks (a Pelican case, actually) with power attenuators, directional couplers and cables/adaptors of all kinds and a 100-watt broadband amplifier.



The NanoVNA SAA-2N handheld VNA

Setting that thing up was a chore, but Ron's method had us set it up in the transmitter building and use the sample lines to convey the signals to and from the tower bases, meaning that we didn't have to lug that big thing out to the towers. That was a game changer and greatly sped up directional antenna work, letting us see in great detail not only leg values but Smith charts of the various terminated networks and transmission lines. We used that setup to make our antenna systems IBOC/HD compatible.

A few years later, we invested in an Agilent two-channel analyzer that was about half the size and a quarter of the heft of that HP analyzer. We used that device to proof all the sample systems for our arrays on which we ran moment method proofs. It was much easier to set up and use, but it was still expensive to ship around, and it takes up some space.

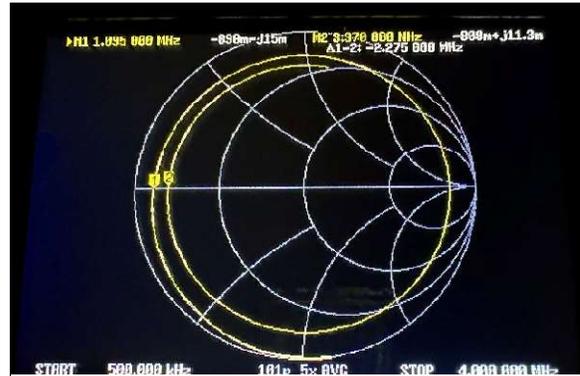
For Christmas last year, my bride bought me a little handheld "NanoVNA" that works from 50 kHz to 3 GHz. It wasn't expensive, less than \$100, and it has a touch screen and USB interface that works with an optional (and free) PC app. It's a two-port analyzer, so you can do S11 and S21 measurements. I had fun using it to tune my 80/40-meter wire antenna over the holidays and wondered how it might function for AM antenna work.

I did a search and found that it can be purchased with N connectors rather than the little SMA connectors the one I have uses. For \$115 total, I ordered one of those, pictured here, and tried it out.

It comes with a set of calibration loads (open, short and 50 ohms). It also comes with some

test cables that give you some flexibility for making connection to devices under test.

For a quick test, I connected it to my notebook computer, ran the app and set it up for a 500 kHz-4 MHz sweep and did a calibration. I then connected it to a spare piece of RG-8X coax that we have running up to the roof from our Denver TOC. The results, depicted on the unit screen, are pictured below. I got a much higher resolution on the app, learning that the length of that piece of coax is 192 feet and that its characteristic impedance is 48.9 ohms.



Screenshot of my RG-8X test sweep

I haven't yet done an A/B test against the gold-plated Agilent VNA yet, but I will at some point and let you know the results. What I do know is that this little \$115 device would be perfect for troubleshooting AM antenna system problems. You could quickly and easily determine if a sample system and transmission lines are in good shape, and if there is a problem, the app has a TDR function (that I haven't yet tested) that might help you find the issue.

You probably wouldn't want to connect the little VNA to a common point or anything with an antenna connected or coupled to it, even a sample transformer if the tubing through it was connected to the tower – the maximum input level is +10 dBm, and it would be really easy to exceed that just picking up RF out of the air. I'm thinking it may well work okay with my 100-watt amplifier and power attenuators. Something else to test it with.

Back to the webinars that I did last month, I told the participants about this little device and said that they all should get one. What a difference that would make for someone trying to figure out an array problem or even tune one up!

The New York Minutes
By
Bill Stachowiak
Chief Engineer, CBC – Western New York

Greetings from Buffalo!

I decided to not continue the Rand building saga this month and tell you instead about another situation I encountered and had to figure out how to handle.

As broadcast engineers, we are frequently presented with problems to solve. Back in the 1980s, early in my career, we didn't have the luxury of purchasing pre-made boxes to handle pretty much any situation that might require some special interfacing.

At WPHD, we had a Motorola two-way radio that wasn't located at the studio. We had a telephone line that ran up to another location where the radio was. The way this worked is that when you pushed the microphone PTT button, a DC voltage was applied to the phone line. This was just a dry copper pair that had complete conductivity all the way to the remote site. The DC voltage pulled in a relay, which turned on the transmitter. Audio was superimposed on the line, similarly to how a POTS line works with a standard telephone.

Eventually we put up another antenna at the site to accommodate a Marti system. We returned the audio on a 15 kHz equalized phone line. The Marti

receiver had two channels. We wanted the ability to switch the receiver between F1 and F2. I devised a solution that worked as follows.

The receiver required a constant closure to hold the receiver on F2. I took advantage of the existing two-way remote control to accomplish this. I had to design an interface that would toggle a relay on and off with each actuation. I won't go into all the design details here, but basically it used an opto-isolator and a flip-flop with some debouncing components to prevent random trips.

I put the diode side of the opto in series with the existing relay coil. I also connected a diode across the opto in reverse polarity so that in normal operation, the diode would essentially short across the opto and allow current to flow as it would normally.

We put a triple-pole switch at the studio that used two poles to reverse the polarity of the phone line and the last pole to actuate the PTT. This worked great. The operator would just push the button to toggle F1 or F2. Unfortunately I didn't have any status coming back, but it solved the problem. Next month I'll continue the Rand building story. There is quite a bit more to tell, so stay tuned!



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

AM Tower Enclosure Cleanup

Since the day I first laid eyes on our 560 AM tower site in Monroe, MI, I knew I had work to do. The site had been neglected and, to be completely fair, it's old, too. The radio station dates to 1956, and I think the site itself to 1959. For many years the transmitter building served as offices and a studio facility, and when the staff left, it became a vessel for garbage, empty boxes, useless retired equipment, and lots of couches. Out of sight out of mind, as it goes. It was also a haven for mice.

Over the past five years, dumpsters have carted away the junk, we've torn down the dilapidated generator shed, the building has been painted, a new roof installed and the place has been thoroughly cleaned. Much more has been done too, if you've been following my column, you'll remember reading about the collapse and subsequent reconstruction of tower 4.

Every summer with intrepid resolve, yours truly and my accomplice Steve (mostly it's Steve) head to the towers with our tractor-mounted sprayer filled with agent not-so-orange and a flamethrower to engage in herbicidal warfare. Alright, we ditched the flamethrower, but chainsaws, loppers, weedwhackers, and PPE were all employed to limited and short-lived effect. Once tower 4 was completed, I knew I couldn't tolerate the subtropical rainforests inside the other three tower enclosures any longer.

For 2026, I planned to have all the woody and herbaceous green squatters permanently evicted, two layers of geotechnical fabric and a thick layer of gravel installed at each of the remaining three towers. Also on my hit list were the omnipresent ankle twisting subterranean marmots (woodchucks).

While the ground was frozen and the plant life was still dormant, I had a crew roll in and convert what was once a chaotic impenetrable nuisance to a clean, walkable, professional looking environment. This summer, we will still spray the fencelines, but no more hacking our way through vines and shrubs to get to our towers.



945.0 MHz STL Fault

Also, down at WRDT, while spot checking some things, I noted a low level hum on our RF STL audio. Normal audio is provided by a Tieline Bridge-IT CODEC, and both backup audio paths are monitored for silence. It's good to take a listen to your backup audio sources from time to time because hum is not silence and the machines we've tasked with monitoring them are happy to count hum as a signal.

Between tower 2 where the STL receiver lives and the transmitter building, there is about 400 feet of underground cable which is unshielded. A bit of snooping revealed a broken wire on an isolation balancing transformer. Open one side of a balanced audio circuit and you'll get no hum rejection and 6dB signal loss.

Strobe Fault

For several months I've been dogged by an intermittent faults on our Detroit tower light monitor, pointing to an issue with our strobes. Replacing and moving boards around that sense the LED beacon strobe's correct operation didn't solve the issue.

Since we had to wait for these circuit boards to make their round trips to Seattle, I decided to fly the drone up to verify 360-degree light output on each fixture. Finally, we settled on water intrusion into the SO cable that connects all of the lights to the controller. After ordering and receiving new the new cable, it was deep winter and finally, we found a break in the weather that allowed us to install the new cable. So far, not one alarm!

UPS Batteries

Each of our TOC racks has a dedicated discrete UPS protecting the loads in its rack. Since they were all purchased at around the same time, we literally got battery replacement alarms on most of them within a few days of each other. I decided to replace all the batteries at the same time. Surprisingly, Home Depot sells compatible batteries for very good prices. Doing it yourself, a 2,200 VA



Figure 1 - The tower 2 base area at WRDT all cleaned up.

UPS can be reloaded with new batteries for around \$89 compared to buying replacement batteries from a computer retailer for around \$275, plus Home Depot ships for free!

Miscellaneous Ramblings and Goings On

February was filled with email issues with varying causes. People's inboxes filling up, their hard drives filling up (with email), mail from our voicemail system suddenly going to SPAM. Of course, you don't find any of these things before spending time inspecting the entire support structure and every detail in between.

I've also spent some time with Linux Mint Cinnamon. I wanted to see if Linux might just be an answer to a question I didn't know I had. Maybe one like, "How do I configure a computer that cannot run Windows 11 to still be useful?" Could it be stationed at a transmitter site providing a site PC I could count on? In other words, one that wouldn't install some update and be waiting for me to tell it what language I speak when I arrive onsite because I couldn't get in remotely? Maybe Mint would be a great free OS to use on virtual machines?

I haven't tried it with VMWare, but I plan to. Also, it could be useful as a high-availability audio player platform for use with Zetta Cloud, which we are just now digging into. Todd Dixon, a longtime Linux Fedora user, is miles ahead of me with Linux, and considering its stability, cost, security, and with Bottles installed, the ability to run windows apps, it could be a great thing. I loaded our test machine with Bottles, and it runs Burk Autopilot without issue.

I keep a whiteboard in the engineering office where we track anything we send out for repair or return. We sent a power supply out for repair on September 15th. I noticed that it hadn't come back but knew we paid for the repair. A phone call to the vendor got them digging and they found their error. Having this on our whiteboard got things going and the supply was shipped today – after 6 months!

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

Cue Points...and the Death of Us All

In the late part of the summer last year, Justin Flores and Frank Franciosi began hatching a plan for some new Sunday programming. We already had some programming from our Detroit market that was being implemented into our Sunday logs, "Sunday Praise" was essentially the music in each market's library and some good news and pay-it-forward stories from callers from each market. The new plan involved having a single, curated set of music that would be used, both old and new, and also some preaching segments from pastors that had gone on to glory and those that are still ministering to this day.



To say that I was proud of our Birmingham market would be an understatement. Originally, we carved out about 1,300 songs that we thought might be used for this format, and we had six to eight people that got into each individual song and made sure that we could separate all of them by several categories and also with "classic" and "contemporary" tags.

It was when we began to make copies of the songs from the Zetta libraries that we realized that the song metadata was not carrying over with the copy. As I'm writing this in hindsight, I should have realized early on that metadata was going to the fly in the ointment.

The solution to the initial metadata issue was upgrading to a later version of Zetta than we were on.

I'm skipping a lot of detail and about 30 calls to RCS tech support, but the format ended up being changed several times, and as it got more focused around November, it became evident that this was going to be a challenge.

The initial thought for the format was that it would likely be just Detroit and Birmingham. It turned out that most of our markets would be able to use these 26-minute audio segments that contained music, preaching and voice tracks to connect them all in the place of giving bonuses to the normal programming that played on their stations during avail timeslots.

The astute among you will recognize the second challenge immediately – the songs in those segments would need to be tagged with metadata so that they could be registered as a play on our streams.

Once Justin Flores had produced a handful of the segments in December, Jack and I started doing a lot of testing on them to see how we could push that metadata out to all of the markets. Mike Kernan and I were getting conflicting information from RCS, depending on the tech we talked to about what Zetta would and wouldn't do and the tools that were available to handle the process.

We finally concluded that we could place multiple cue points in the audio segments that they could give a `LiveMetadata.Send` execute command where the song title and artist could be filled in manually. It was tedious work.

The plan was for Justin to produce eight of these segments that each had between five and eight songs in them. Even for Vito in Detroit, who I can confirm has the spiritual gift of not needing sleep, it took about 45 to 50 minutes to process all of the segments with the correct metadata. In Birmingham, the cue points weren't necessary, as all of the audio with metadata was available for local playout in our system, but in every other market, the cue points would have to be inserted by someone like Vito since the cue points did not stay with the audio when it was saved outside of Zetta.

At the end of a Zoom call in January, the consensus was that we'd have to find a way to automate the cue point insertion or we'd have to schedule the streams to go dark when we played the programming. The format was slated for air in Birmingham and Detroit on January 18th; other markets got pushed back until March 1st to give us some time to figure out getting things automated.

As a result of that call, Cris purchased a

Zetta module called "Publish" that can be used in the show and mini-log tabs of the Zetta library that creates a log file of all of the elements. Even though it is a .log file, it is formatted like an XML file.

It took several more calls with RCS to get to the bottom of how that log file could be used, but one thing was for certain – Zetta had no way of using it to repopulate the associated audio with cue points at another Zetta location, and it seemed we were stuck.

At this point, we were six months into a lot of work and time spent to make the format work, and the idea of it not working as we intended didn't sit right with me. I am not a coder, but I figured that there had to be a way to write some type of script that would do this cue point work in a more automated fashion and might possibly cut down the time to something more manageable for the other markets.

So... I registered for a free account on Anthropic's Claude AI. It is supposed to be the best AI for coding, and since a coding dummy was going to be chatting back and forth with it, I wanted to be working with the best.

To be honest, after a lot of initial explaining and a couple of sessions, the AI was able to make sense of parsing the log XML file and knowing what songs and artists were in the audio.

Even with the early success with Claude, I had gotten to a place where I simply wasn't sure how the script was erroring out. Surprisingly, when I actually wrote those words out in the chat, the AI chatbot created a debug script that I could run and determine where the error was occurring. It ended up removing some superfluous code and eventually the "knowledge" it gained from the debug script got the main script to where it was able to add the cue points to the audio based off of the log file that Zetta's publish created.

After all of that was done, it also created a PDF file that ran through how to install python and all of its extra requirements to make the program run on computers in other markets. I polished the script up with a few more sessions by having it add a loop in the script so that the user could do the next audio segment without restarting the entire script. It wasn't perfect, and it did take some initial user setup as it worked on based on X-Y positions on the screen, but once it started to run, it could add all of the cue points for each segment in about 1 minute and 45 seconds.

I had been regularly reporting to Cris about my progress with it. It was about that time that I sent it all to him to try out. Cris does code, and after working through my version, he almost immediately checked the script and sanded away several of the clunky edges that I had sent him. He and Amanda

spent some time making sure they knew how it worked, and once they had it down, they began to place it in our California and Detroit markets as well.

So we made it happen ahead of the deadline, and all of our markets will be able to use the audio for their programming as they see the need. I know it appears like I may have named dropped at least 20 people in the company, but it was really a concerted effort of a lot of people to pull it off and make it work, and I couldn't be happier with the results and the teamwork it took to get there.

I might add that in the midst of the time that I was working with Claude AI, Mr. Crawford sent an email that included excerpts from Stephen's column from an earlier Local Oscillator. I was a little concerned using AI for this work, but I knew that I couldn't have created it with my own current coding skills within the given deadline.

After considering it for a while, I think of AI like I do the theory of evolution, I don't believe that any species on Earth evolved from a single paramecium that somehow got electrocuted in a primordial soup-macro evolution. I also don't give a passing thought to the existence of micro evolution – ask any breeder at the Westminster Kennel Club Dog Show if inter-species evolution exists.

It's the same with AI for me. I don't think macro-AI can continue to move forward at the pace that it has moved. This era of globalization we've been living in seems to be breaking down, and TSMC in Taiwan is the only company that makes CPU chips performant enough for AI usage. Frankly, the manufacturing process from raw material to finished product is the most complex on the planet. Other chip fabs are on the horizon, but it is going to take some time to build the same processes with the same level of competency that TSMC does. The big AI players are swapping money between each other and creating a real false economy amongst themselves. So I think of that big AI like I do big evolution.

What I do believe exists is micro-AI. Warts and all with its inherent mistakes, it is creating real efficiencies when it is being guided and overseen by

people with real experience. It's a tool that can be used (like I did, but way better) and has a number of use cases where it can be used to multiply one's efforts, similar to what computers did at the dawn of the PC age.

A case in point: recently, an Anthropic AI researcher used their latest Claude Opus 4.6 AI code base and set it up on a server with 16 separate Docker instances as "agent-teams." The task given to this team of AI agents was to build and code a functional "C" programming compiler from scratch. The agents were given a shared "git" software repository. From there, the agents managed their own merge conflicts and acted autonomously.

In two weeks, they had created a compiler that included 100,000 lines of code that was able to compile a bootable Linux kernel on several different CPU architectures. The total cost to run the agents over those two weeks with over 2000 Claude code sessions was \$20,000. Take that in for a second. What would that have cost a company with real software developers and project engineers?

Obviously, it was project from a researcher at a major AI company. They likely had a number of resources at their disposal and you still wouldn't be able to convince me that the software created by those agents was perfect, yet it was able to complete the task it was given. It was an AI milestone as it was the first time that an AI had performed a long-running, software engineering project longer than three to four days.

Just like the dog breeds at the Westminster Kennel Club show, AI is getting better, and in the end, I think it could help remove the technical out of the technical, just like it did for me. In my experience, I got better at reasoning with it and figuring out how to make it do what I needed to.

I plan on trying a couple other projects to see about automating some other tasks... not writing in *The Local Oscillator*, though... I do enjoy putting my thoughts on paper.

Until next time, may God bless the work of your hands.

Tales From Cousin IT
by

Stephen Poole, CBRE, AMD
CBC Corporate IT Specialist

Let's start with some local color. Figure 1 isn't a great picture, because I didn't want to get too close; I was afraid I might be eaten. I used the Zoom on my Android smartphone from a safe distance away, a couple of houses over. There's a lesson here: there's little point in advertising if you don't clearly tell people how to contact you.

The dumpster did have a very faded company name and phone number on it, but you couldn't make it out unless you were very close. I didn't linger, though. I didn't want to get eaten.

This weekend, I'll probably pass kids doing car washes to raise money, or folks having yard sales. Red marker scrawled on a pastel orange or yellow posterboard isn't a good choice. For that matter, when I lived in NC, there was a business on US 1 in Southern Pines that had a dark blue letters on the sign out front. You couldn't even read it in daylight; it was barely legible in the dark.



Figure 1 - I'd love to ... but HOW?

In totally unrelated news, Kim Jong Un has been re-elected as Supreme Thing in North Korea (also known as Chosŏnminjujuŭiinmin'gonghwaguk for those who can actually pronounce Korean). In his honor, I'm thinking about submitting a bid for a statue like Figure 2 for my weed-filled flower bed. Spruce up the neighborhood, wot?

I wanted to send Kim an email telling him that if we're really BFFs again now that Trump is

back in the White House, could he please tell his people to stop hacking our servers (more on this in a moment)? North Korea doesn't publish email addresses for government offices (this assumes they even have email, which is iffy, given that most of the country doesn't even have electrical power). I was unsure where to send it. I briefly considered sending a missive through the United Nations to "crazymidget@rc0-dprk.un.org". Briefly.

The ICNB Can't Help Themselves

The Internet is broken. I won't say that it's dead; that would be too dramatic. But over the past couple of weeks, as I've



Figure 2 - In Honor of HIZZONER Kim Jung Un (Wiki:Sumit Surai)

worked on both our corporate webserver and the new mail server, it has become obvious that attackers, on top of Internet congestion from new data centers, are taking their toll. The attacks are coming from everywhere nowadays, though the Good Ol' Boys –

Russia, North Korea, China – still pop up the most often.

The most common package used for network encryption is called "OpenSSL"; the "SSL" stands for "secure socket layer," though nowadays, "TLS" (transport layer security) is preferred. A recent update by the International Congress of Neck Beards™, the ICNB, has decided that if an OpenSSL connection is interrupted, it should be treated as a serious error. Since a recent update to the package, we've been getting weird, unexplained glitches from our servers as a result.

In the mail server's logs, I see things like, "ERROR: unexpected EOF encountered." There's supposedly a setting that will tell the mail server (or its proxy, a module called "nginx") to ignore this, but it didn't help. I have since learned that Zimbra modifies and compiles its own special version of the SSL/TLS stuff. Yay. Googling for suggestions says that the problem is not our server, but the other end of the connection, i.e., the Internet. Happy times!

Even happier times: we apparently have some users who think that I (or Todd, or Amanda, et. al.) sit around with a finger poised over a well-worn red "kill dis email" button and stab it at random while cackling madly. No, in most cases, it's due to hacking attempts. Our servers are set up to temporarily disable any incoming IP address that keeps trying bad passwords. We have no control over the wider Internet.

Adding to our collection of happies is the fact that the lifetime of SSL certificates going to be shortened. So sayeth the ICNB, and that's that. I've mentioned this in previous issues; it is now official. Over the next two years, this will gradually be phased in. Sometime in 2028, they'll need to be renewed at two-month intervals. To address this, we're thinking of switching to a service called "Let's Encrypt," which can be automated to update and install each new certificate. Happier times!

Blocking Bad Stuff

Let's face it, the bad guys are very good at finding ways to get into your systems. The aforementioned "unexpected EOF" thing could possibly be exploited by a really sophisticated actor using something called a "truncation attack." From what I can tell, it's very unlikely and I couldn't find any active exploits that would concern us.

For the record, the OpenSSL Neck Beards added the, "explode if a connection is abruptly terminated" years ago in Version 1. They had to remove it because a bunch of things just stopped working. (Properly closing a connection isn't a high

priority for lots of network stuff. Ask me how I know.) But the ICNB is nothing if not persistent. Now in version 3 of SSL/TLS stuff, they've decided to bring it back. Stuff is blowing up all over the place, but the ICNB is doggedly determined this time. We can only hope they back off again. For now, if you start experiencing weird errors, especially if you've just updated Outlook or some other Internet-dependent program, that may be why.

On to other things. I recently found "invalid URI path" strings in the error logs on our web server. Part of the URL (the gibberish that you see in your address bar) was the text "/bin/sh." This tells me that someone or something was trying to execute a shell script(!). This is extremely ungood. The URL started with "/cgi-bin/," which is a directory for older "common gateway interface" modules. These are rarely used nowadays, but years ago, they were a way for a web page to execute external programs (!! again).

By coincidence, this coincided with an email from Rick in Chicago saying that the POR system was returning a blank page when he tried to log in. Naturally, I sweated that one until I figured out that the problem was simply that a certain log file in the POR directory had been inadvertently wiped. (Probably by me. I had been doing some work the night before, checking on backups.) Recreating the log file solved the issue, but just for safety's sake, I disabled ALL cgi as well.

So We'll Move to the Cloud!

Some people might think that I hate the Cloud because I feel threatened by it. Let's put that silly notion to rest right now: I couldn't care less. There are things about "the Cloud" that I like; the basic problem is that the Internet is overloaded. When there's a Cloud outage, instead of just affecting one server on our premises, it affects everyone, everywhere. Websites go down, email and stored files are unavailable, you name it.

What's hilarious is that AWS's biggest failure (last October) was caused by their own AI tools. Simply put, their Artificial Intelligence made some very unintelligent tweaks that blew up AWS East for quite a while, with cascading effects that rippled worldwide and through a bunch of different Cloud services.

Yes, I have a weird sense of 'umor. You already know that. I also have a well-developed sense of irony. I'm old enough to remember how many times we've moved from "one big server for everyone" to individual PCs, then back to "one big server," over and over again. The idea of One Huge

Data Bank appeals to big thinker IT and CS guys and always has. Personally, I like the idea that MY data stays on MY computer, right in front of me.

As I write this, the city of Bessemer, AL is determined to approve the building of a huge data center in spite of objections from the public. We'll see. It's undoubtedly going to end up in court. My amusement will be complete if the attorneys for the

data center use AI to prepare their briefs. If the community's lawyers start using AI to prepare their responses, I may need supplemental oxygen.

That's enough for this time. We're looking to move to our new mail server sometime in early March, Lord willing. More on that when we're done. Until next month, keep praying for this nation!

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

We have been on Zetta here in the Chicago market for a little over seven months. For the most part, after the initial problems, we have had good success with the software. I really like the way it loads audio so quickly and there is no conversion to a special format necessary.

This has its advantages with how quickly you can get audio into the system. We have clients doing semi-live programs where they email in their next segments just a few minutes before they have to air. Having a system where you can drag in a file like this is good.

The disadvantage I see in this ability with pulling in MP3 files is that not every MP3 file is created equal. We certainly see or should I say, “hear” this issue with our music on the stations.

It wasn't long after we went to Zetta that I was hearing from one of our engineers, James Kelly, that he was having to change the settings in our Omnia 11 processors. He was noting that even after making changes he would hear the levels of the audio move up and down from event to event on the air.

He was spending hours of frustration trying to get the sound correct. At first, he believed there was something wrong with the processor that it showed these rapid changes and questioning the stability of a few of our units.

After a while of this chasing of the problems, I was convinced that he was chasing his tail trying to address the problems with the processor adjustments. The issue was the source material. He started making notes of the bad audio and looking at the source files in Zetta.

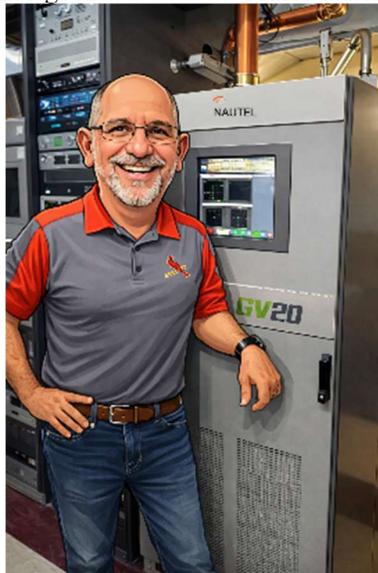
He found that some of these were added recently by programming with them not using good source material. Some were songs that came over in the transfer from our NexGen file server but apparently had some distortion in them from the mass transfer. We were able to go back to the NexGen file server and retransfer them again and this cleaned up some of the issues.

We are still having some bad audio issues, but we are getting through them a file or two at a time. It is important that programming really takes the time to make sure they have the best versions of these files possible. This was true for when we were using NexGen but it seems to be even more true with Zetta.

Another issue that has come up recently is that we have two Zetta control room workstations that have been really annoying our operators. The Zetta software will be up and running and then suddenly shut down. This can be really inconvenient during a live program.

We took the problem to RCS Support. They pointed out that the ESET Anti-virus software was in their words, “beating up Zetta.” We set about creating exceptions in the ESET software, even getting ESET support involved to make sure there was nothing we were missing.

Unfortunately, this didn't seem to fix the issue with these machines. The problem of Zetta suddenly shutting off kept happening. That's when we decided to take off the ESET Antivirus software altogether. This was the best way in my estimation to prove out whether it was ESET that was causing this to happen.



I wasn't surprised, if not a little disappointed, that the problem showed up in less than an hour after removing the ESET software. To me this showed that the diagnosis of ESET being the problem was not correct.

At the time of this writing, we are going to put the machines through the Microsoft and Dell updates. The computers are barely a year old, and you wouldn't expect to have these kind of issues.

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

HVAC Controller

A month or two ago, I noticed our HVAC controller at the KLTT transmitter site wasn't responding. I wanted to switch to the main unit for the winter months since it won't need to cool as much. That main unit has been on the fritz with a compressor down, so it can no longer keep up in the summer with just one compressor running, but the winter should be no problem.

I went to switch it remotely, but it would not respond. I even tried a soft reboot with no change. I didn't want to make a special trip to the site since getting this working wasn't a huge priority. We still had a working unit active.

We were able to get to the site on January 31, and a hard power cycle didn't fix the issue either. We got our spare Pi from the office and went out in early February and replaced it and all seems to be working now. Apparently the GPI from the thermostat somehow got blasted, making the unit "think" that it was telling it to switch to the backup unit all the time.

Omnia Remote Access

I have found that the Omnia.9 units have some weird issues. I can't explain it. It's like they lose their way on the network. On the unit you can enter in the allowed IP addresses, and these IPs are the ones that can gain access. I had that set up but would get an immediate disconnect on the unit when I tried to connect. I could get on with a web browser to the internal web server where you can download the NF Remote software, so I knew the allowed IP address was working.

The next step was to try a power cycle, but I didn't want to take a station down. I actually forgot about the dual site for 810 AM. I could have easily

rebooted one of the units and kept the station on air and sounding great. Once I remembered, I did a hard power cycle. We have all of our Omnia.9's on a smart outlet that is hooked (via SNMP) to our Burk system. After the power cycle, the unit came up with no issues and I regained access with NFRemote.

I am glad this fixed it, and I did the power cycle across all our Omnia.9's. We are lucky that we have all the units on a smart outlet.

Cris has noticed the same issue in other markets. Not all these markets have the benefit of being able to have the units rebooted remotely.

This will be the second time I've had to do this in the past year. If it keeps happening, we'll have to figure out why.

We were having some issues with one of the stations in Denver and I needed to switch it to our backup (analog) audio feed because it didn't seem to switch on its own. I wanted to hop on the Omnia and make sure we were receiving audio to that point. That's how I figured out we had no access.

I will put it on my calendar every couple of weeks to hop on and check the connection to be sure I still have access.

Wheatstone Blade 3 Display Screen

We have some older Wheatstone Blade 3s from the mid-2010s. They've been great. However, over the years, the display has become unreadable. Sometimes a reboot will fix it, but in one of our production rooms, we had two different units that would not self-correct.

I have asked Wheatstone about the issue a few times over the years, mainly when Jay Tyler would show up in January for our SBE chapter meeting. We would go around to our rooms looking



things over with him, and it'd catch my eye. At the time, Wheatstone didn't have a display we could purchase. We would have to send the blades in, which we really couldn't do. It would take a room down for weeks which we did not have.

Now they do have the part available. I ordered two and installing them was a snap.

Gate Opener

The gate opener at the KLVZ site is great. We love the fact that we do not have to get out of the car to open the gate anymore. The only real issue with it is that the range of the opener is not far at all. We typically need to be right in front of the gate for it to open, which is a problem if we're in the truck pulling a long trailer with the tractor and brush hog on it. That leaves us blocking a busy county road while we open the gate.

We purchased a different opener receiver to attach that we hope, once we get the antenna raised up, will allow us to open the gate from the road without stopping. This is the same opener receiver that Steve Minshall recently installed on the gate at KCBC, and it works great there.

UPS Batteries

A few years ago, I created a spreadsheet that has the info of all the UPS units in Denver. I include make, model, serial number, the year it was purchased, the battery replacement number, and when it was last replaced.

I have been working to get on some sort of schedule so that I don't have an "oops" moment and

have one of them go down at an inopportune time. With that, we are also trying to find ways to not have to spend hundreds of dollars on big battery replacements. With these bigger, rackmount units, they typically hold four batteries, and the replacement is the whole cage with the four batteries.

We recently replaced the batteries in a unit at the KLVZ transmitter site. It took some work, some prying, cursing and leaving of DNA samples, but we were able to get the new batteries installed and the unit seems happy. This was much more cost-effective than purchasing the whole cage, and going forward, we will start doing this.

I am slowly making my way through the list starting with some of our oldest units first. It is my hope that by the end of this year, I will have replaced every battery, making them all one year or less old.

Coming Up

I sure hope we start getting snow. Spring is right around the corner, which is crazy because we are still waiting for winter to arrive in Colorado. I don't have any projects on the horizon. Just continue replacing UPS batteries every couple of weeks. I may take some time during the month to do some winter cleaning at the studio.

If it continues with no snow and springlike temperatures, it may also be a good time to get some outside work done. I'm sure the tower bases could use a cleanup as well as around the buildings.

I think that about covers it for this edition. I pray you all stay safe and well.

The Local Oscillator
March 2026

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



Corporate Engineering
2821 S. Parker Road • Suite 1205
Aurora, CO 80014

email address: calexander@crawfordmediagroup.net