The Local Oscillator

The Newsletter of Crawford Broadcasting Company Corporate Engineering

MARCH 2025 • VOLUME 35 • ISSUE 3 • W.C. ALEXANDER, CPBE, AMD, DRB EDITOR

Zetta – The Adventure Continues

The conversion from Nexgen to Zetta in our Denver market is complete. All that remains are ongoing tweaks as the staff and Amanda learn more about the platform and how to use it.

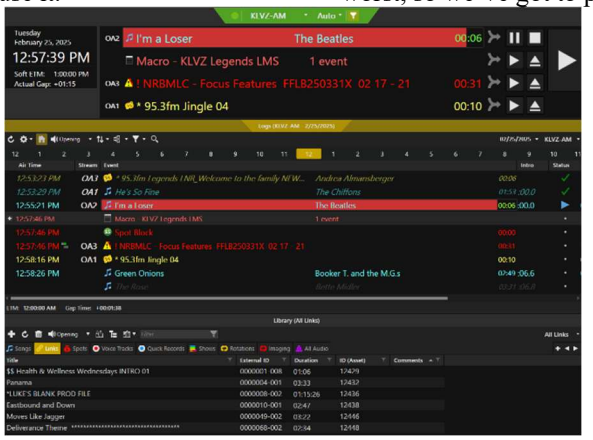

Since last month's issue, we have learned a good bit more about Zetta. Amanda and I got some one-on-two admin training, taking us pretty deep under the hood. We learned about backup and maintenance routines and options, and we got the site replication and Zetta cloud backup working. We have yet to test the cloud backup, but we should find some time to do that in the coming weeks. The time to learn how to use your lifeboat is not after the ship has sunk!

Operators and producers are learning more and more about Zetta, too. One thing that has bitten them several times in recent weeks is the "segue point" in audio elements. Apparently, it's fairly easy, probably from clicking when you shouldn't, to accidentally set the segue point somewhere that it shouldn't be. The results are predictable – the file plays to that segue point and then the next log element starts. That has resulted in a few overnight train wrecks. Confirming the location of that segue point is an important last step in loading or creating an element!

We've also learned a good bit about TREPlus and uncovered some shortcomings that have since been corrected by BE. As a result, our exports, both HD PSD and FM RDS, are looking sharp. I will concede that metadata exports are a secondary consideration, but they are part of the overall Listener Experience™, and if the information displayed on the radio or dashboard is wrong, misaligned with the audio or missing altogether, it's a listener irritation at best and a tune-out factor at worst, so we've got to pay attention to it.

I'm sure I irritated the installer with my insistence that the exports work perfectly, but I know the potential the exports have to complement the programming or irritate the listener, and in some markets, we have money tied to the metadata, so they have to be right. With Zetta talking to TREPlus and TREPlus humming right along, we are getting correct title-artist on the RDS and title, artist, album title and format on PSD, all time-aligned with the audio.

So where do we go from here? Birmingham! Preliminary (remote) conversion work will begin the first week in March, with on-site work taking place throughout this month and wrapping up the week of the 21st. Todd and Jack have been working to get things ready, including updating all the PCs to Windows 11 and cleaning up the Nexgen database and clocks. They also got their EAS units wired up to Wheatstone blade LIOs and the LIOs and SLIOs created and routed so that everything is ready to accept command outputs from Zetta when the time comes.

We also got TREPlus installed in Birmingham, configured and running, ready to accept metadata exports from Zetta. The clients are up and running and we've gone about as far as we can until we have Zetta running.

On the heels of the Birmingham conversion, the Buffalo conversion will start with offsite work the week of the 17th and onsite work the following week. As in Birmingham, Brian and our Buffalo crew are making preparations, getting EAS LIO wiring done and cleaning up the database and clocks. Since we already have TREPlus running, we don't have to do anything special in advance for that, but we will have to change the preset on the import once we get Zetta running.

Note the difference in brightness of the OLEDs that have been replaced (right) with those that haven't (left).

After that… tentatively we have Detroit next in the sequence. As in Buffalo and Birmingham, Mike and Steve have been busy getting things ready on their end. The EAS wiring is done and the LIOs and SLIOs are routed, and the operations and production staffs have a few weeks to get things cleaned up in the database and clocks.

We had planned to push our California markets into 2026, but we have since pushed them up to this year. Those two markets are right after Detroit, probably in late April or early May. We anticipate just one week in each of those two markets, so they should go quickly. That leaves Chicago, which will be a five-week onsite affair. That will probably take place in June.

We'll continue to polish the Denver Zetta system and pass along any tips and tricks to the downstream markets.

LXE OLED Screens

It's no secret that OLED screens can become dim over time. We have certainly observed that in some of our installations.

A couple of months ago, Jay Tyler and Kelly Parker from Wheatstone were in our Denver facility and they saw that some of the OLEDs were really dim and difficult to read. Jay asked if we had reduced the maximum brightness and set the timeout on the displays to extend their lives, and of course Amanda told him that we had. But still, after several years of the LXEs in service, some of the screens were really dim.

That got us thinking about the issue, and we looked at all four of the studios with LXEs in them. It was immediately apparent that one room (KLZ) was in the worst shape with all screens really dim, and a second (KLTT) was not quite that bad but was still dim. The other two rooms were still fine, good as new or close to it. Why would those two studios be good while the other two were dim? It occurred to us that the two that had dim displays had big windows; the other two had no windows (although KLVZ has some indirect natural light through the talk studio, which has a window-wall.

The OLED boards come out and the displays aren't too hard to replace, but it requires a delicate touch!

Here in Colorado, there isn't a lot of atmosphere between us and the sun and as such, a lot of UV radiation reaches the ground. We have to take special precautions with a lot of different things. We've had CAT5e cables connecting roof- or tower-mounted microwave radios just about turn to dust in a couple of years. We specify UV rated cables for anything that goes outside here.

Apparently, those UV rays are affecting those OLED screens indoors. The window walls in KLZ and KLTT are heavily tinted, but some UV still gets through, and we believe that is what is behind the premature failure of the LXE OLED screens in those rooms.

The fix: periodic replacement of the OLED screens.

Mike Kernen found a wholesale source of exact replacement OLED screens and tried a couple out a few months ago. Once he figured out that they would work, he passed the information on to me. Amanda ordered enough to completely replace all the screens in the KLZ LXE. Last month, she started working her way through the modules, and the contrast is readily apparent. Even at (deliberately) reduced brightness, the screens are bright and clear.

If any of you are experiencing LXE OLED screen brightness issues, drop Amanda a note. We can order them for you. They're about $9.00 each, which is a fraction of what Wheatstone wants for them. It may well be below their cost.

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

Hello to all from Western New York! February will go down as one of the worst winter months we have experienced in years. Aside from the daily snowfalls, temperatures have hovered in the teens for practically the whole month, and the wind chill has been mostly in the negative zero and below range.

I am so ready for this winter to be over, and I am actually looking forward to getting out at our tower sites this spring and performing some outside maintenance work.

Even with the extremely cold temperatures, snow & wind, we have not experienced any downtime at any of our transmitter sites, which is truly a blessing. Looking back over the past several winters, we have had to rely on our standby power generator to keep us on the air each year, due to downed poles and power lines caused by the severe winter weather, especially at our Boston, NY site where the WDCX-FM transmitter is located. I am praying that trend continues until the spring flowers start blooming and the warm temperatures prevail once again.

Our audio path for WDCX(AM) in Rochester has been experiencing a strange anomaly which started mid-January. The issue was, every day at about the same time, the audio would drop for about two minutes, dead silence. The weird part was the drop time would advance 23 seconds each day. Trying to determine the cause of these daily drops was quite the chore. Where do you start?

To fill you in on our audio delivery method from our studios in Amherst, NY to the WLGZ-FM transmitter site in Rochester, NY, we send audio out via the internet using Spectrum Business Class Internet on both ends. At the studios, we use the Tieline Gateway to send audio to the transmitter site, where a Tieline Bridge-IT XTRA receives that audio and feeds program audio to both the AM STL and the FM translator.

I began by contacting Spectrum to see if they could identify any issues with our internet service at each end. They replied that they did not see anything specifically that would cause the daily drops but performed a hard reboot of the cable modems at each location. The following day, after the modem reboots, the issue remained, the audio dropped at precisely 23 seconds later than the previous day's drop. Hmmm.

My next step was to download the event logs from the Tieline Gateway and Bridge-IT XTRA and email the logs to Tieline support to see if they could possibly identify the cause. About a week later, I heard back from the technical support engineers at Tieline, who replied that the service drops were caused by a "Keep Alive" timeout, and recommended that the Keep -Alive value be changed to "0" in the Bridge-IT XTRA. I made the recommended change, but the problem did not go away.

I waited for a few days, then downloaded the event logs and sent them again to the Tieline support team for evaluation. After a few days, they came up with the solution, the firmware running on the Bridge-IT XTRA does not allow the Keep-Alive setting to be disabled, which was causing the daily drops. Once I uploaded the latest firmware (2.18.146), the problem was resolved.

I am again impressed with the technical support crew at Tieline, and their diligence in helping resolve any issue with Tieline products. I wish all our manufacturers shared their service philosophy and dedication to complete customer satisfaction.

Another issue we are working to resolve pertains to our streaming encoder for WLGZ in Rochester. Early this year, we installed a new streaming encoder computer, a Dell Inspiron with an i5-13500 processor (2.5 GHz) and 16 GB RAM, along with the newest SoundBlaster Audigy Fx V2 sound card. This computer does not perform any other function; it's dedicated for the streaming function only.

Several weeks after the installation, it was noted that there were digital blips and skips on the Apple App stream as well as the Alexa stream. The Google app stream would shut down after a minute or two of starting the stream, the app would stay open, and the stop button turns into a play button.

I provided SoundStack with all the streaming encoders' information and current settings, along with all driver versions currently installed. Todd VanTasel, our support engineer with SoundStack, has been working on finding the cause of the issues and a solution, but has not come up with any theories as of this writing.

Todd reports that he indeed is hearing some artifacts on the ingest server, which is the system that our streaming encoder is connected to. He asked if we were utilizing any of the processing options in the Audigy dashboard, which could cause the anomalies we have, but all processing is done through the Omnia external processor, which has a clean, noise-free input/output. Hopefully, by next month we will come up with a cause and solution for this issue.

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, and happy engineering!

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

Frozen

I made a decision long ago never to complain when it's hot outside. I can deal with heat, but I'm a person who hates being cold. February tends to be the coldest, and yes, the most miserable month here in Michigan. My colleagues in Denver, Chicago, and Buffalo would argue that they have it worse, but I'll tell you from my only perspective, that it's cold, gray, and yucky here in Detroit.

Zetta Preppers

There has been much work to do in preparation for the upcoming upgrade to Zetta. For us, it has been in our crosshairs for a couple of years, and we've done such things as replacing our old desktop tower formfactor hardware with Dell rack mount machines for TOC and micro formfactor for prod studios. Not only are these a much better fit for the racks in our TOC, they are much more robust, heavy duty, and professional. For a short time, Dell offered them preloaded with Windows 10, which worked perfectly for our NexGen installation, but now they need to move to Windows 11. In fact, all computers running NexGen needed to be upgraded to Windows 11.

Windows 11 is a fairly easy upgrade, but each machine takes about two hours. Once you've finished, you have to call RSC support and have them run their "tweaks." It's a good idea to verify that all Windows updates have been run, too. Also, you need to check that each of your sound devices has the correct sample rate configured. Windows likes to change the recording sample rate from 48,000 to 44,100, which causes obvious problems, and this happened on every machine I updated.

RCS also directed me to switch Core Isolation off. You can find this setting buried deep in Privacy and Security Settings, then click Open Windows Security, then Device Security.

Windows says "Core isolation provides security features designed to protect core processes of Windows from malicious software by isolating them in memory. It does this by running those core processes in a virtualized environment," which is incompatible with NexGen and presumably Zetta.

As much as I like the new Dell 1RU computers, they did bring with them some challenges. As manufacturers move to modernize their I/O, they've understandably dropped things like VGA video and PS/2 keyboard and mouse ports, which caused a gigantic problem for us in TOC. Our KVM (Keyboard/Video/Mouse) switching equipment was old-school. It consolidated 16 different computers onto two workstations. Those workstations were VGA and PS/2, as were the KVM switch and the computers connected to it.

While making the transition to the new Dell hardware, I found myself using adaptors for video and having piles of keyboards and mice stretched all over the place. This wasn't good and it proved especially problematic when having to access several computers while troubleshooting – "Ok, switch to this machine. Now which keyboard do I use? My mouse isn't working! Oh, that's the wrong mouse! What computer am I looking at? Arrgh!!"

All this came to a head while updating NexGen computers for the upcoming Zetta upgrade. As I said, we needed to install Windows 11 on all the computers that would run Zetta, and that meant running studios on ECR (Emergency Control Room). ECR runs a studio on its local database, allowing the back-office servers to be maintained or upgraded. Switching back and forth to ECR, updating Windows, running BIOS updates, etc. requires rebooting, sometimes several times, and not being able to see what you're doing and searching for the right input devices is frustrating!

I decided to upgrade and straighten out our KVM switching situation. I've had experience with big KVM systems. One very large Raritan system I have installed allowed for in-studio host switching, integrated KVM extension, and supported 48 user stations and 192 hosts! It ultimately had the same issues. It wouldn't support widescreen or even 1080P resolutions, was natively PS/2, stubborn about USB adaptors, and not upgradable. Like the Raritan, our old IOGEAR stuff was straightforward and gave us good service. It was just outdated. It seemed like replacing it with IOGEAR's newest KVM would be a good plan.

For the 16 computers we have in TOC, I ordered 2 IOGEAR MiniView Ultra+ KVM kits. The kits include cables and of course the 8x1 KVM switchers. Although they are stackable, I decided to have two workstations in the TOC, so 8 computers on one, 8 on the other. This allowed me to move one of the switchers down the row a bit and since we already had a rackmount monitors and keyboards in both locations, I didn't have to use the stacking function which was good, since the included stacking cable was too short. I might change my mind on this should I find a longer cable, but so far, my attempts to extend the existing cable have not been successful. The video worked, but I'm pretty sure that there is a rollover of some of the data wires. I've also had to order a few more adapters because some of my new Dells only have USB-C for both keyboard/mouse devices and display port (DP) video. Save for it shipping with one bad video cable, it's working quite well and nice not to have to hunt for the right keyboards and mice!

An HD Radio screen showing the WMUZ-FM logo and PSD from TREPlus.

Growing TRE

Broadcast Electronics' TREPlus "middleware" product has been installed and providing us with a platform to help us monetize and enhance our RDS and HD Radio data services.

Our Digital Media Director, Pete Presnal, has been putting it through its paces and getting all our advertisers set up with their messaging and images. The goal is to interleave messages on the radio receivers with current programming, current weather information, and ad messages to enhance the effectiveness of their campaigns.

We're using both the Xperi HD Radio Artist Experience as well as RBDS technology to bring this to the listeners' receiver. It's a big job, and Pete and I have been watching and learning it, growing it into a routine amplification of both our station and advertisers' messages.

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

RCS Zetta-Birmingham Edition

By the time that most eyes have seen this article, our RCS Zetta installer/trainer will be on site and beginning our conversion to our radio playout system future. He is slated to be on-site for about three weeks starting on March 3rd. He has already been doing some remote work on our file servers over the course of the last week and a half.

As I mentioned last month, the lion's share of this early work is related to making sure that the conversion between the Nexgen and Zetta databases occurs. It also includes installing all the relevant software pieces like the Microsoft SQL server to process the database and Zetta's underlying management software that all the workstation clients count on in order to make the magic happen.

During the month leading up to the upgrade to Zetta, Jack and I were busy making sure that all of the workstations are ready to go. As Cris mentioned last month, the metadata exports needed real help, so Birmingham is jumping on the TREPlus bandwagon with the rest of the company. Cris got a lot of that work pre-done so that we are ready to hit the ground running to deal with Zetta's metadata shortcomings. We also got our Sage ENDECs prepared by getting our Wheatstone LIO/SLIO wiring done so that we'll be able to handle all of our EAS Required Weekly Tests.

Every Wednesday, I've been scheduling an after-lunch video training session for all of our staff that are using the software. We've done four of them so far and I've basically been using videos from the RCSWorks YouTube channel that cover some of the basics. Most of the videos on the channel are from their "Live" events that they host every couple of weeks that cover various aspects of the Zetta software.

So far, we've covered basics in the videos, like setting up your best layout, the logs module, the multi-track editor, and importing and exporting audio. Although there is a video for it, we'll probably ask our trainer to go over Zetta2Go while he is here. I am really looking forward to getting this done. I see so much potential to make us more reliable, efficient and to sound better than we already do.

Doctor My Eyes have seen the years…

One of the things that Zetta offers its users is a cleaner, more modern desktop layout. Since you have the ability to pull a number of modules up at once and organize them to your individual preference, one of the things that I asked for and purchased was 14 27" Viewsonic monitors. If I said that each monitor was only $115.00, would that ease the shock a little about the purchase? I try and keep up with the technology, and more particularly, the prices that different peripherals cost, and I simply couldn't believe that they were that inexpensive.

There were a lot of things coming together with this purchase. Almost all our monitors were 19" diagonal and most were set up on the far side of our LXE consoles. We have about 15 employees total in Birmingham that use Nexgen/Zetta on a regular basis and only three of them are under 30!

Everyone has been really thankful for the upgrade, and I can only hope that it will be a wise purchase for us as we begin to convert over to Zetta.

Of course, you wouldn't think that changing out a set of monitors would be an opportunity to rewire anything in the studio, but since the monitors were so much bigger, we had to reconsider the way some things were run and reconfigure a couple of our production rooms. We got the chance to remove some old wiring and neaten up existing wiring. I didn't see any evidence of any on our studio furniture, but the installed "grow lights" can make wiring look like it's out of control.

You might be asking what we are going to do with the monitors that we replaced. Our production room Wheatstone E series surfaces had 15" monitors attached to them that pre-date my employment and were the original monitors that were purchased when we originally switched to RCS Nexgen, so those three monitors got swapped out with some bigger models. We have a couple of CRT

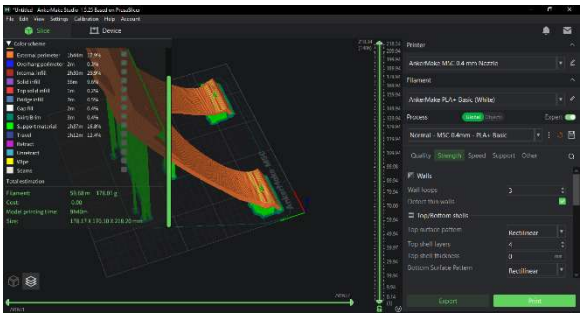

Figure 1 - A screen capture of the slicing software, how the model was angled and the supports.

monitors that we are going to swap out at several tower sites. Another thing we're going to do is mount three in our Technical Operation Center (TOC) in our racks. The plan is to output our Burk Auto Pilot screen, SNMP monitoring and a yet to be determined other screen that will give us some visual station status prompts.

Figure 2 - The fitment of the 19" monitor in the rack (minus some necessary cleaning).

Jack and I put our heads together and thought through how we might mount these monitors in the racks. What we came up with was pretty good. Since most monitors have VESA standardized mounting holes (generally spaced in 75 mm or 100 mm squares), we created a VESA mount that is specifically designed for four similar model Dell monitors that is specifically two mirrored parts that meet in the center with tabs that insert into the other and are screwed into the monitors VESA mounting holes. When I say "we," Jack has a background in design and he jumped onto AutoDesk's Tinkercad.com website and whipped up a design that could be output to an STL file that I used to 3D print the mount. After about 45 minutes of work, he handed me the prettiest design I had ever seen and nonchalantly stated, "Eh, TinkerCAD is like Adobe Illustrator, just in 3D."

Figure 3 - The actual VESA rack mount, the 2 pieces join together in the center of the monitor.

Since a 19" rack mount solution would be too big for most 3D printer beds, we broke it into two sections and I had to turn the model on its ears and angle it up.

After using my slicing program to convert the STL file in to lines of g-code (a set of x-y-z positions that the tool head gets moved around to), the model required printing it with some "supports," which are not structural or in any way a part of the final print, but are designed as minimal help to simply keep the model from falling over while it continues to print each layer. Each part took about 9 hours and 15 minutes to print and ended up using about 165 grams of PLA plastic filament.

Of course, the proof was in getting it put in place and it absolutely fit like a glove even though Jack claimed his "math" was off just a hair. In case you're wondering, the 330 grams of filament were probably about $7.50 total. Jack has already designed version 0.2 that cleans up his math (Todd rolls his eyes).

Until we visit again next time, may the Lord bless the work of your hands.

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

If someone doesn't cover up that stupid groundhog's hole, I might have to go reptile. Just like everyone else, we've had very cold weather here, at times down into single digits. For Alabama, that's unacceptable. Things freeze. Roads get slippery. Faces turn red (well, even redder, and don't ask about necks).

During one of the extreme cold snaps, we had a water main break. I suspect that it wasn't buried very deeply when the lines were run several years ago. Because ... well, it's Alabama. So, one Saturday in February, I was without water all day. When the water came back on, of course there was air in the lines. My faucets were hissing and spitting (also of course), but the toilet in my office bathroom, really, REALLY didn't like it. I went to flush it. It said, "BOOM" and mist shot out from under the lid. The bowl filled instantly and flushed just as quickly. I was quite impressed.

I still haven't gotten all of the sediment out of the water. For the first few days after the repair, it was cloudy, almost milky-looking. It tasted nasty, too. For now, I've been making do with bottled water, which is kind of ironic, given that I have one of those expensive whole-house water filtration systems. I expect it's going to need servicing. What's really scary is that the goop in my water is what's making it past that high-tone filter. I don't want to know what the filter looks like.

SNMP

All y'all need to be warned that Cris and I are about to run wild on this Simple Network Management Protocol stuff. We had earlier figured out how to do a SET command to trigger a relay. I have since successfully programmed a Raspberry Pi to accept a GET request and return a RESPONSE. This will allow us to basically retrofit older equipment with SNMP capability. Niiiiiice.

Getting to this point took quite a bit of work, too, partially because of the Basic Encoding Rules stuff that I mentioned last time. In this issue, I'm going to throw a whole bunch of gibberish at you. It might help you troubleshoot any weird problems that you might experience, whether you're using an SNMP-capable remote control (such as the Burk Arc Plus+), or a monitoring program such as PRTG.

Wireshark

This is our standard go-to when we want to examine network data. In this case, the capture filter that I use is simply "host [IP address of the SNMP sender]." When done capturing, if it's grabbed a bunch of unwanted stuff, I enter "snmp" in the display filter area. (Remember that Wireshark's capture filters and display filters use completely different syntax!)

Figure 1 is a capture of SNMP traffic on our network in Birmingham, courtesy of Todd Dixon. I've zoomed in on the first two packets, which represent an SNMP GET and RESPONSE between PRTG and WDJC's Nautel NX50 transmitter. In this image, I've selected the response packet (the 2nd line). Down below, I can click on "Simple Network Management Protocol," then start clicking on the drop-downs until I can see everything (as is shown here).

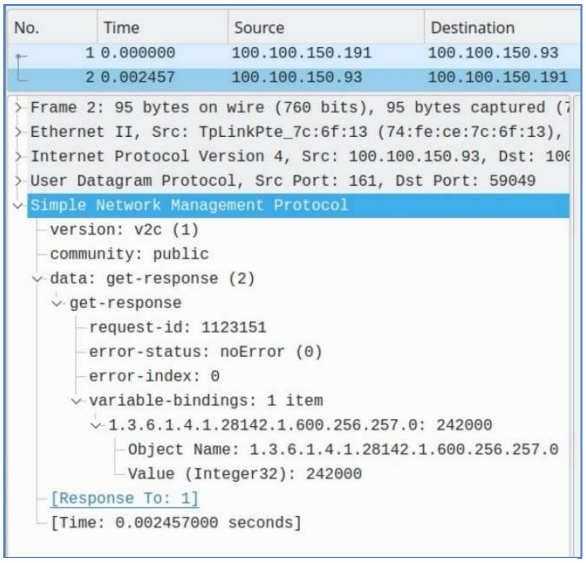

Figure 1 - A Wireshark capture of WDJC's NX50 talking to PRTG.

I'm going to do a more in-depth study of some other packets in a moment. For now, note that the RESPONSE here is returning the current forward power of WDJC's transmitter, which shows as 242,000 watts! Todd should get a raise! (Heh.)

Other items of interest in the display are the Request ID, which is unique for each request, and the OID ("Object Name" in Wireshark's display). This is a great way to check if you've told Burk to send the correct dotted OID string. The vendor code, 28142, is Nautel's registered identifier.

The Program Data Unit (PDU)

If you have acronym-poor blood, prepare to be nourished. In SNMP-land, the

"PDU" is the entire set, request or response, filled with SNMP information. In the Wireshark display, you're going to get all of the network headers; I trimmed the actual data that usually appears to the right in my setup. But the PDU comes after the IP and MAC addresses, the port numbers, and any of the other network-specific stuff. The PDU starts immediately after the UDP header.

See Figure 2. This is a different type of capture; my program, running on a Raspberry Pi, can save the actual received and returned PDUs for debugging. The incoming data from the Burk is on top; the bottom half is the Pi's response.

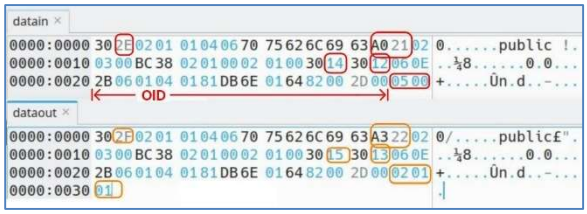

Figure 2 - A capture of SNMP GET and RESPONSE from KLTT, Denver.

All of the data, both here and in Wireshark, is in base 16, or "hexadecimal." If you have trouble with it, many calculator programs (including KCalc, which is part of KDE on Linux, which I use) can directly convert and do mathematical operations between hexadecimal, decimal (i.e., "normal" base-10 numbers) and binary. The fields in the PDU follow Basic Encoding Rules, which is a series of blocks, one immediately after another, that follow this format:

Type Length Value

... or, "TLV" if you're a neckbeard. All values are 8-bit bytes in our case and the only Types we worry about are:

- 01 - Boolean (true or false)
- 02 - Integer
- 04 - Octet string
- 05 - NULL (0 or no value)
- 06 - Object Identifier (OID)
- 30 - A "Sequence" (a block of data)
- AO - "GET" command
- A2 - "RESPONSE" to "GET"
- A3 - "SET" command
- A4 - "TRAP"

First: The overall "wrapper"

See the "30 2E" (top) and "30 2F" (bottom) at the start of the PDUs in Figure 2. This encompasses the entire SNMP block, including the request ID, OID, and any returned values. The GET request is 2Eh (46 decimal) long; the RESPONSE in the bottom box is one byte longer (2Fh, or 47 decimal).

Next is the community name, in our case, either "Public" or "Private." This is type 4 (octet string), 6 or 7 bytes long, followed by ASCII text. Immediately after the community string is ...

The SNMP Command Block

... which I've circled. The top block starts with A0: GET. The bottom uses A2, for RESPONSE. The GET command has a length of 21h (33 decimal) bytes. The bottom block adds one additional byte for an integer return value (1) in this case, so the length is incremented by 1 (22h, or 34 decimal).

The request ID is next. It does not change between the GET, SET and RESPONSE. You might check this in Wireshark, which helpfully shows "Response to [number]" to help you keep track. SNMP managers and agents can't guarantee a nice sequential order. It's possible to get several GETs, followed by a mixed-up string of RESPONSES. Look at the Request ID when troubleshooting.

Object Identifier (OID)

Now the real fun starts. You'll note (at the top and bottom of the image) that we have two "30" blocks. This is because you can actually have several OIDs nested in a single PDU. In the top box, "30 14" simply means, "one or more OIDs follow, 14h (20 decimal) bytes long." There's only one OID in this case, so the next block, immediately afterwards, is "30 12" -- the actual OID, 12h (18 decimal) bytes in length. The lengths of the two "30 blocks" in the bottom box are simply incremented by 1, as already mentioned.

The OID itself starts as shown, type 6. The OID doesn't change between GET and RESPONSE, so you see "06h 0Eh" in both the top and bottom boxes. 0Eh is 14 bytes in decimal.

Decoding an OID can (of course) be done, but to someone unfamiliar with bitwise cramulation (i.e., in this case, stuffing more than one value into a single byte), it can be confusing. Adding to the joy is that the numerical values are "Faux Base128." That's not a misprint; Base 128, or 80hex. Since Wireshark is nice enough to show the actual OID values in the usual dotted form ("1.3.6..."), I knew what I was supposed to get when I started decoding. It took a while to get it right, though, and the available code samples online were needlessly (ridiculously) (hilariously) complex.

(My hat's off to Nathan Burk at Burk Technologies, who I'm told is the guy who figured all of this out for the ARC Plus units!)

One final note: the variable value is after the OID; it is not part of the OID. These are the final values in both the top and bottom of figure 2. In the top box, "05 00" means, "a NULL of zero length" -- i.e., an empty value. The bottom box has "02 01 01," which means, "an integer, 1 byte long, value: 1." As mentioned several times above, adding that one byte required me to increment the length values at several points in the PDU for the RESPONSE.

Finally ...

It's a good thing that God loves me, even though I don't deserve it. I hope you continue to be safe, blessed and happy as we work through 2025. Until next time, keep praying for this nation!

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

New Cambium Install

In the past month we have had a tower crew come to work on several projects. This has been challenging as we have had to deal with weather that is either way too cold or just very windy or too windy to hoist a microwave antenna up a tower. When you have plus-25 mph wind gusts, the dishes make better kites than antennas.

One project was a planned 11 GHz Part 101 link between our Lansing and Beecher sites. The Beecher site is a stretch to make work from our Hammond office/studio site. So, to get data back forth we have to use the link between Lansing and Hammond and then the link between Lansing and Beecher. Even the link between Lansing and Beecher is no easy stretch as we are looking at an 18-mile hop.

For years we have been struggling to make a 5.8 GHz, unlicensed system work on this stretch. For the most part over the last couple of years it's not worked more than it has. So, we decided to go in a different direction with a licensed system, in this case a Cambium link, as we have had great success with them, both locally and across the company.

For the most part, the installation has gone well. However, we have found the receive level is about 10 dB down from the expected level. This is frustrating, and this is not the first time that this has happened with a Cambium installation. Our link between the Hammond and the Burnham transmitter site was installed a few years ago and also was 10 dB down from the expected signal. In fact, we asked the current tower crew to look at that link to see if we could improve the performance of that installation. We only got a slight improvement.

We are trying to figure out the reason these two systems are not a peak performance. For the Burnham link I am looking at a rather tall tree along the path about block away from the studio site. For the new link I am just beginning to investigate. We believe the link to still be useable for codec audio and data between sites.

Tower Lights

We also had the crew take a look at the tower light issue on the Lansing tower. Our top beacon flash head was not working correctly in the day mode. After analyzing the wiring with a meter, the suspicion was there was a short in the wiring of

the cable heading to the top of the tower to the flash head. Having the crew in town to install the Cambium system seemed like a good time to have them help us diagnose the problem.

It turns out that the problem was not up on the tower. When we cleared out the pass-through holes in the wall for the new Cambium wiring, we cleared out the spray foam insulation in these holes. One of the tower crew members got on a ladder to begin the cable inspection near wall entrance when a squirrel popped out of the wall and caused a bit of commotion.

Worse than the startled the tower crew received was the fact that we now knew the culprit for our shorted wire. The squirrel had chewed through the insulation of the wire, and this caused the issue with the tower lights.

We also found other cables with insulation being chewed, specifically some of the CAT6 cables. Worse, the 208-volt cable for the antenna heater was heavily damaged and was a candidate for starting a fire in that wall.

So, it turns out the tower light outage was a blessing in disguise, making us aware of the situation with the antenna heater cable. The cables were repaired with either splices, or in the case of low voltage CAT6 cables, with electrical tape to replace the outer jacket over the copper clad shielding. All is working properly now.

The squirrel is now homeless.

Valley News
By
Steve Minshall
Chief Engineer, KCBC -- Modesto

A few months back, Amanda wrote about her adventures in renewing the Austin Ring transformers at one of her AM transmitter sites. This corresponded closely to my finding deterioration of the ring transformers at my site.

Over the years, I have worked at numerous stations with ring transformers but never gave much thought to their maintenance needs. When I recently found the instructions for the transformers, I was surprised to see that it was recommended to paint them every two years. Since the transformers at my site were installed about 39 years ago, they were a bit overdue.

The "transformation" from start to finish.

Amanda and Cris had done all the research and footwork to find the proper procedures and materials for renewing the transformers. Cris sent me a package with rolls of cotton fabric and quart cans of red and gray Glyptal varnish. Glyptal varnish is great stuff that's used extensively in electrical work. It's also used, with some controversy, to coat the insides of engine blocks and intake runners of cylinder heads in automotive racing engines. It is really tough material.

The procedure for repair is to coat the transformer winding with red varnish, then wrap it with cotton fabric and soak the cotton with more red varnish. Once the red varnish is completely dry, several coats of gray varnish are applied. I used a 2-inch chip brush for the red varnish and a mini-roller for the gray. I discovered, quite by accident, that a 4-inch roller on a 6-inch frame is perfect for the job.

I did this work in winter, so I had to pick days that were dry and sunny. While two coats of gray are recommended, I didn't think it looked finished, so I ended up applying four coats. The result looks close to the factory finish.

One thing that always bothered me about the ring transformers at the site was that they weren't centered well – the primary coil was droopy. Years ago, I tried to adjust the couplings on the conduit supporting them, but they were well pinned with Allen set screws and they refused to turn. The Allen screws were also very much seized. As a temporary

measure, I propped up one of the rings with a piece of wood jammed between the concrete and the conduit. That "temporary" fix stayed in place for a very long time.

The reworked flange and spacer greatly reduced the stress on the flange connection.

While working on the repair of the transformers, I revisited the problem of the droopy primary windings. The attachment of the conduit to the bottom of the tower base insulator looked like a terrible design, placing tremendous bending moment on the bolts. I couldn't understand how an engineer could design such a poor attachment scheme. Then it hit me – the installers must have put the flange in upside down! The next day, I confirmed that the flanges for each primary coil were indeed installed upside down.

Correcting the placement of those flanges would be a significant job. I also discovered hairline cracks in the casting of the base insulators where the bolts go in at two of the towers. I needed a good permanent solution that would center the rings and reduce stress on the base insulator casting without major work.

My original "temporary" fix using a chunk of wood wasn't such a bad idea, but it needed improvement. I lifted the primary coil into the correct position with a ratchet strap attached to the tower. Using clearance measuring blocks, I determined the exact distance between the conduit and the concrete. I then cut some heavy-wall aluminum tubing on a lathe to the exact dimension needed. After placing this in position, I released the ratchet strap and repeated the process on the other two towers.

This solution removed the bending moment from the flange connection on the base insulator and replaced it with a relatively small amount of shear stress on the bolts. I believe this reduced the stress on the flange connection by a couple orders of magnitude. For the first time, all the rings are centered nicely, and the repair is solid and permanent.

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

Zetta

February has been a whirlwind of a month. Not only is the month short by a few days, it was just jam packed with Zetta.

It appears most of the weirdness has been worked out. Scott, the Field Technical Support guy who came out to do the installation, fixed a lot of stuff while we were working on other things. The audio in the studios seems to be holding solid. NexGen, Zetta and Audition can all be running with no issues.

People at the office in Denver are loving Zetta. We have one lady who works from home when the weather is bad, and she is loving Zetta2Go. She has a good Shure MV7 microphone and can do all her voice tracks easily from the comfort of her home.

I haven't heard anything bad from our other two remote hosts, Tron and Julie. Once we got them set up with Zetta2Go, we haven't heard a peep. We know it works though because

I ordered Mike Kernen in Detroit this Epson labeler. Now he can label wires like a pro!

the voice tracks for their shows are in the logs and the VTs sound great!

We've learned how to schedule macros in the Music Master logs/clocks for KLVZ. Since that is mostly a music station, we went with doing the scheduling through Music Master rather than having a separate clock in Zetta. After a few weekends of live shows having issues because the macros got blown out of the log to turn the console on, we figured out how to do it properly, in Music Master or in the mini logs that are used for various weekend shows.

The one thing I am working on now is how to automate FTP downloads. Todd Dixon spent countless hours figuring this out to work with NexGen. I had problems due to some naming issues with our programs, so I never got it set up with NexGen. I would imagine that those who have it set up with NexGen won't have an issue with Zetta.

I have plans to get with one of our people in California to see their setup and see how it works and really investigate it with NexGen in hopes I can find a way to make it work with Zetta.

I have had Scott working on it for days with no luck. I have called some other RCS support people, all with the same result – nothing working yet. I am going to keep working on it, and hopefully I can get something figured out before Birmingham gets their Zetta setup installed.

Worst case, I don't, and Todd has more work to do. But at least there would be two of us trying to figure it out. I think in Denver, for one station we have twelve daily shows to download from our own FTP server. I don't even want to know how many more shows we have to download from other places. Our goal is to automate this to free up people's time to get other work done.

Overall, we are all very happy with Zetta. It wasn't without its growing pains and issues, but once we worked through it, it's been smooth sailing.

KLDC Connectivity Issues

We continue to have connectivity issues at the KLDC transmitter site. For a while, the Cambium link's Ethernet connection would go down. I would still be able to see the radio and it was passing data, but the network on the other end wasn't there. It would sometimes come back after a reboot; other times, it was a waiting game.

In the process of troubleshooting all this, we found the backup internet would go down, or at least the Barix would. We went to the site, did a full power cycle on the Cambium and moved the network cable from one port to another on our network switch. This cleared it up for several weeks, but the issue has returned a couple times since.

I also mentioned the backup internet. We have CenturyLink internet at the site and a Barix Exstreamer 500 for backup audio. What I have witnessed is that the Barix will lose its link. I can get in remotely to the laptop on the CenturyLink side, so I know the port forwarding is working. The transmitter site Barix is very slow to respond if it does at all. I sometimes get a timed-out response when trying to get on, other times it will just load slowly. I have to go in and reboot the NetGear router, which we have replaced to try and remedy the issue. Since this isn't an everyday issue, troubleshooting has been slow.

LXE OLED

As noted elsewhere in this issue, we noticed in January that the OLED displays on the KLZ LXE console were pretty dim. With the help of Mike Kernen, we were able to find a suitable replacement cheaper than going through Wheatstone. I bought enough to change all the OLEDs, I think, if I remember correctly, there are 39 on the consoles. With the help of my dad, we were able to get these replaced with few issues.

We did learn some things, like sometimes, things don't come back up properly and you need to power cycle the console. Another thing we found was that the EXT1 and 2 buttons, which are programmed for our Pre-Delay feed and Air feed, came unprogrammed somehow. It took us a good long while of messing around before we realized how to fix it.

The new OLEDs look bright and amazing and I have already received comments on how great the console looks now that they can see the displays.

Coming Up

Spring is nearly here, and I am very excited. I have a bunch of work to do at the transmitter sites. At KLTT, once the ground thaws, I have some copper strap repairs to deal with. All the sites could use a good cleaning. Thankfully, on a recent trip to KLZ to have an electric meter replaced, I was able to take time and clean the place up some. I look forward to later in the spring and mowing. We'll see how the growth is, but my hope is to spend even more time at the sites getting things mowed down.

I think that about covers it for this edition. I pray you all stay safe and well. And to those markets of ours getting Zetta, "May the odds be ever in your favor."

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