EXPRESS_23269 | Advances_in_PNM_with_a_leader_in_Comcasts_Next_Gen_Operations_Technology_group

All right, happy Friday, everyone. Today is September 9th, and this is episode 23, in all things DOCSIS I'm Brady Volpe, founder of the Volpe Firm, and Nimble This. Today we have a fantastic episode for you on PNM, 3.0, 2016 advances in technology, and what to expect coming in 2017. But first, let me introduce our panelists for today's show. With us we have the Diviner of DOCSIS John Downey. CMTS technical leader at Cisco Systems. John, great to have you with us today.

It's good to be back, as usual.

Yes, and fantastic audio, as always. Also with us, for the first time on the show, we have the Ayatollah of rock and roll for proactive network maintenance, the one and only, Larry Wilcott, Comcast fellow in operations technology. Larry, so good to have you with us today.

What an introduction, of [? the ayat-noise ?] of rock and roll-a Ayatollah.

Ayatollah of rock and roll-a.

I love you, Brady. You've got a [INAUDIBLE]. OK. That's great. So it's great to be here. I've been a big fan of you guys for a long time. Obviously my heart still is in this PNM stuff. A cable junkie from the get go. And it's an honor to be here. Really excited to have some fun with you guys, and talk about this great stuff.

Yes. So fantastic. Obviously the discussion topic today is going to be around a proactive network maintenance, A.K.A. PNM. I would like to start, and open up the discussion just to talk about how far we've come from the early days of PNM. You know, from where we were just looking at pre-equalization values from the cable modem itself, to now where we are today, where we have full band capture, where we're looking at downstream spectrum. We're doing upstream spectrum analysis. We're looking at intermit modem capabilities to find noise. So I mean, we're almost changing the term of PNM to, you know, maybe some new name like, PNM plus, or advanced PNM, because we're doing so much more with PNM. And I think we're getting just so many more capabilities. And so this is some of the things I want us to cover in today's meeting. Do you guys agree?

Absolutely. Great stuff, Brady.

John? Good? All right. So, you know, Larry, maybe we can start with you. I know you're doing some pretty good, cool things, and stuff. Is there any particular item you want to open up with?

Yeah, actually it's funny, if you look at the board behind me here, I'm actually squatting in a conference room at Cable. So all that cool formulas, and jazz back there is not mine. I can't take credit for it. But it's like the perfect scenario for doing this kind of thing, and so appropriate. That's looks like a 2.1, or something. Let's roll back into the PNM time machine to 2000. Well, let's go to 1999. And just for a quick history lesson on the stuff, adaptive equalization coefficients are where this all started. That's, can you see an echo tunnel between the cable modem, and the CMTS on the transmission line? And then what, right?

And that's where this whole thing started, in '99, with DOCSIS 1.1. And I remember distinctly when this came out, kind of looking at it. I had no idea what it really meant, which is funny. And then I kind of shook my head. And the problem with PNM, or with adaptive equalization, at the time, was they this 8 tap equalization structure. And the resolution was horrible. And that was really the first technical constraint. Then really, more so, the problem with it was that early implementations of equalization were just really bad. We turned it on, and customers wouldn't come online. And just lots of bad stuff happened. And equalization just started off with a bad name. And it took years to overcome the stigma that was created by the early implementations. And here we are today. If you look at number one, just equalization as a concept, we can't function 64 client reliably without it. And if you look at the pre and post constallation and MER that we see today, in many, many areas of our network, it's a must. We just couldn't do what we can do, deliver the quality, and the amount of service that we do without it. And the evolution from 1999, till today, and all of the folly in between in terms of industry adoption, and operational adoption, and learning, and evolution, is remarkable.

And I remember very fondly when I first saw how this could be operationalized, in 2007 when I was visited by Alberto Compos, and Eduardo Cardona. They came and visited, and showed me the spreadsheet. Brady, have you ever seen Albert's spreadsheet?

I don't know that I saw one that he had on PNM. I did see Alberto's early presentations at Cable Tech Expo on PNM. And I think it was a similar thing, as you said, where I had no idea what he was talking about.

Exactly. And the crazy thing is, history repeats itself. And now that we've operationalized it, it's in the tools, alarm's pod. People are dispatched to go fix stuff. And still to this day, I got to be careful where I say this, the interpretation of information is very difficult. And that's kind of the opportunity. And this is a great segue into the evolution of where we're at, and where we're going in the future in machine learning, all this school stuff we're going to talk about. About how do you interpret the data? And what do you do with it? And actionability? And how do you translate that into a business? And how you operate your cable system? It's such a fantastic topic.

Let me throw in some historical data too, or my two cetns. And I know I've talked about it before, with you, and Brady. Cisco is one of the only ones who used a TI Intel CMTS chip. Everyone was [? Drawcom. ?] And we ran into some issues with adaptive [INAUDIBLE] or actually, [INAUDIBLE]. It was an [? attach ?] equalizer. And what we found out was that when you're different firmware, you might have impulse noise. That pre-EQ could get out of whack, meaning-- I'll give you an example. Say MER ruptures. MER says 20dB. I turn pre-EQ, it goes to 25. A week later it drops down to 17. We're like, wait a minute. What's going on? Turns out the pre-EQ is like pre-equalizing the wrong way, or out of time, getting the wrong coefficient.

So Cisco, a while back, made a decision that we'd have like a catch all. On the CMTS, by default we have this on. Once you turn on pre-EQ, because pre-EQ's not on by default for our CMTS. Once its run pre-EQ, if the MER of a single modem drops by here is 3dB, we kind of attribute it to pre-EQ problems. So we send a directly to the modem, basically saying, clean yourself out, and start over again. So it actually works well as kind of a catch all, Band-Aid if you will. I don't know if [? dot aires ?] has the same thing. Maybe they do. But it's actually helped save me, quite a few times.

That's good to know, Dan. Many years of doing this, and I've seen the behavior. I didn't know exactly why. You'd just explained it. I should come to these more often. All right. Cool. Then, to basically just reign that in a little bit, that was some of the history. We started with the first scout flux implementation. That's our first crack at it. Here's equalizers, here's-- we didn't really know what the heck it was going to tell us, or what we're going to do about it. But it was kind of the first output of that coefficient analysis. We just kind of quietly put a little cool looking equalizer button and the tools. And just let people stumble on it with no fanfare, no meaning. I mean, we did a couple of decks that talked about what equalizers did, and how to interpret it. We didn't really know.

And then just kind of out of grass roots, from the field, which is really, I think the secret sauce to our success for the program, the technicians started picking up, that cared enough about it, and starting engaging us by email, and calling us, and creating relationships in the field. And that was kind of an early implementation of this dev ops model, by happy accident, if you will. I still have fantastic relationships with a lot of folks out in the field, and a very cohesive feedback, and development lifecycle. So that if something happens, or a feature, or function, is just magic when you get close to the field, and respond quickly to that.

So we turned it on in 2009. 10 years, ironically, after the very first equalization implementation. 10 years after what, after having it out there, running, we figured out how to operationalize it. And then along with that came all of the trials, and tribulations, of now what's? And it's the now what where things get interesting. Brady, let me hand it back to you. Do you deal with the now what every day of your life with your Nimble products, and out there with the operators? Anything to add?

Well yeah. So we got involved, I think, it was around 2011. As I said, I didn't understand it, with Alberto's original presentations, because they were very technical. But I think it was when Cable Web's came out with their reference implementation, around 2011, that we got involved. And the reference implementation really helped understand how you put this long EQ string into their reference implementation, you get out the taps, you get out the in channel frequency response. And that actually makes sense. The challenge with reference implementation is it wasn't operational. So then we started Nimble This, and we turned the reference implementation into something that had database, behind it, you could start importing data, SNMP polling modems, and CMTSs. And then we started working with some small cable operators to really say, OK, you know, we see these problems. Let's see if we can actually find where the problem is, in the field. And that's where it started to become really exciting for us. But the challenge that it became after that was, there was so many modems that were sort of non-compliant, where we would pull the modem, and we would get back-- and Larry, I think you coined the term-- a bunk equalizers string.

And I really liked that. So you kind of had to build up this library over time with this modem, then that. So you can start dispelling it. You need to start training cable operators that, know how to turn on pre-EQ, how to put the CMTS into TDMA mode, because of it's TDMA mode only, then the modem's are operating as a DOCSIS 1.1 modem. They're not taking advantage of their 2.0, or capabilities. So there was, you know, a learning curve for us, as a commercial PNM provider, and also the cable operators, as some of the things that they were doing weren't taking full advantage of the pre-equalization in the cable modem.

So you know, it's one thing to just understand the fundamental technologies of pre-equalization. Then it's another thing to understand the limitations, because none of this was part of the DOCSIS specification. So there was so many issues in cable modems not fully supporting, it that we later fixed. You know, I think you, guys drove a lot of this, in fixing issues in the firmware, and the modems, and CMTSs, and us also working with vendors. And then also training cable operators on how to optimize their systems, and taking advantage of this.

Again, this is still classic PNM. So it was really exciting then also to say, now in the DOCSIS 3.1 spec, all this is going to be standard. So a huge learning curve for both us, and operators. And as you said, getting this operationalized within cable operators. Now cable operators really understand how to be reactive. When a modem's offline we know we got to fix that right away. But it's completely different to say, be proactive. Where we'd really like not to have the cable modems go offline. We'd like to fix these before the subscriber's impacted. So that's a very that's a big shift within organizations. And getting there takes a lot of work internally, working with operators.

Yup. Absolutely. John, talk a little bit about your early introduction to PNM, and equalization, and what was your reaction at the time?

You know, I was going to add on to Brady's PNM react of RNM. You know, everyone does RNM, reactive network maintenance. So when PNM came out--

[INAUDIBLE]

In the cable industry was C Cor Electronics, and then I went to Wavetek. And so I had a lot of background in upstream, downstream sweeping. And we used to do this way back in the day. We called it, basically, FDR. Frequency domain reflexometry. Basically, you're doing an upstream sweep, with really tight granularity on your sweep points, avoiding overlapping with actual channels. But if I made my sweep points tight enough, I could get better granularity in my sweep response.

And looking at an upstream sweep trace, I could see standing waves. I could use the well known formula 492 times [? blossom ?] propagation, divided by the delta in the frequency, standing wave, and just find out the distance to the fault. So I was well aware of troubleshooting that way, but now that we have so many modems in the field doing this already, we could utilize, and exploit that functionality, right? So the more modems we had, the more test equipment, the more visibility we have out in the field.

So what I started seeing, and I was sick at SET Expo years ago, where Alberto was showing this stuff. And I was like, this is pretty cool. And I then I think there was winter or summer conference table that's Jiggle Apps, outside of Denver. Keystone? Maybe it was at Keystone. And he showing me some of the functionality. And I was like, this is this pretty powerful stuff.

It's just a matter of operationalizing it, like Brady was talking about. You can put a paper out in the web, and people look at it like it's a Greek, right? The smaller operators didn't know what to do with it. You work for Comcast, so you're a bigger operator, you have SNP expertise, a little bit more resources that can go into it. So you guys really ran with it, and I was pleasantly surprised. Now I got everyone else asking, well how can I get that functionality as well? How can I get visibility into that into what's happening out in the field? And it's not just, you know, or two channels.

Now it might 5 and 42 megahertz, or 5 to 60. Now we're looking at systems going 5 to 85 megahertz. You know, the spectrum's getting full in the upstream. So now we don't have much spectrum that we can just throw away. We're using close to the [INAUDIBLE]. So we definitely need to review.

I always thought lower below 20 megahertz is a lot of impulse noise. And I asked this question to you guys during that panel last year. I said, is there any detriment to run a pre-EQ. Like, should I not run pre-EQ below 20 megahertz, if impulse noise could negatively affect the pre-equalizaiton. Then I thought, well, if it's lower than 20 megahertz, it's real far far away from the [INAUDIBLE]. But that's not our problem now. It's mostly micro reflections, not group delay. Group delay might be a problem near filter roll offs, and stuff like that, but in the lower frequencies, we still have a lot of micro-reflections, no matter what. And I remember the one guy, I think Keith from Cable Ad said, no I would run pre-EQ everywhere. Do you feel that same way?

Absolutely.

Yeah. So you don't think the impulse noise detriment is enough to turn off pre-EQ?

No. So early on-- and that's a good point John, and credit where due-- the early implementations of pre-EQ-- in fact, you can actually see them in like the post equalization coefficients-- where there was constantly attempting to just, like a non-linear affect of an impulse noise that was present. And I think it's all just kind of worked itself out where the burst receiver, and the vendors have all come up with the appropriate filters, so they don't whack out like they used to. So, in general I'd say I would always recommend using adaptive equalization.

Yeah.

Good question though. Yeah.

I've also had great results with dynamic upstream interleaving in the mod profile when I use frequencies below 20 megahertz. That usually creates the uncorrectable factor, go to correctable factor, in some regards. So it actually works out pretty well. It's something that's not on by default, but I usually give a recommended mod profile, if people run frequencies below 20 megahertz.

Yeah, that kind of reminds me of, if it's raining, or pouring outside, run around really fast, and you'll won't get hit as much. That's a joke.

Unless it's a hurricane.

Of course. Yeah. OK. Great segue I think your comments about how your early days of doing return sweeps, and that kind of evolved, and it connected the dots where we were able to use CPE. So effectively, do return sweeps for us. And that really changed the game for return set-ups, and return sweeps, in my opinion. It certainly did in my world, and I can't remember the last time I've heard of anybody having to do a return sweep, just to be honest with you.

Now, is in both the forward, and the return, Larry? Because I mean the--

[INAUDIBLE].

Yeah, the return sweep is something that I think is really going away. Forward sweep still happens a lot.

Oh, absolutely. Yes. But the point was, you know, with the intelligence of the CPU, and the ubiquitous deployment of all this stuff, why would you drive around doing things that your CPU could tell you? Kind of the segue into the next part of the conversation.

Yeah.

That came up with a local SETO I was at, and I said, you know, with full bandwidth capture on the downstream-- and we can use that term, right? FBC,

Mm-hmm.

Full bandwidth capture. It's not [INAUDIBLE] and copyrighted?

We're good.

I mean, remember saying generically they were going to call it a wide band capture, I don't know. And I felt like, with so much spectrum being used in a downstream full bandwidth capture, you'd almost have a downstream sweep everywhere. But if we go to 1.218 gigahertz upgrades, you're probably not going to have carriers there so, you might still need to inject sweep points. Then it comes down to, where do you put the equipment to inject? I mean, Brady, and I have talked about this, when we go eventually to remote [? fi, ?] now it's all digital out to the node. So you have to inject the sweep points at the node, somehow. You're not going to have test equipment hanging off the polls, you know, and next to the node. So do you incorporate some type of injecting, sweeping carrier, in the chipset itself? Or do you just use the full bandwidth capture, and look at the noise floor, like a sweep trace? I mean, noise does get affected by gain and loss, pretty much the same way, right?

Great point.

if you normalize it, then you can kind of get sort of a-- it's not really apples to apples, but you get the idea. A delta.

Yeah, absolutely. And that's a fact. But the problem with that approach is that-- that's actually a great one. We should talk on the side about getting a patent on that. But the problem with that is the filter response across the entire spectrum is very difficult to get a flat, filtered response across the entire spectrum, all the way up to 1.2 gig, or whatever. So there's a lot of stuff in there you have to normalize out too, which would be, again, the filer itself.

So it's worth looking at. But I'll tell you, we call it a sacrificial crumb. You just put it out on in the roll off, and you know, you can basically-- and this is a great conversation for, how do I use full spectrum 3.0 cable modems? And we're doing this today. It's fantastic. To characterize the roll off of the spectrum where we want a cheat and use 2 net 1 or OFDM channels. If I'm going to put a 96 megahertz wide OFDM channel in there, what does it look like before I even deploy it?

Of course. You said test? The frequency you plan to use, with the type of signal you plan to run.

Exactly. So you couldn't just turn it on, and use it the 3.0 spectrum analyzers to tell what the frequency response is, even without participating in the 3.1 conversation, if you will, but you can learn a lot about the frequency response in the fidelity of plant, at whatever frequency time used. But a lot of systems, I'll be honest, they're just sticking a quam out. Calling it a sacrificial quam out in the high end, in a roll off area. And it tells you a lot. If you look at this--

Here's a good question for you. I brought that same question up of, I'm going to turn on more quams. The CMTSs now are C-Cap, right? They have full spectrum capability. So do I turn on more RF? Do I turn on more RF? Is that you phone, or somebody else's.

Sorry, that was me.

So if I turn on more RF, can I do it without being charged a license? So think about that one. That's probably higher above your pay grade. I don't know. But I said the same thing to my own guys. I'm like, I should, be as an RF guy, should able to turn on RF that doesn't carry any content, and not have to pay a license to turn it on, because I might want to just use it as a test signal.

Interesting.

You didn't hear that from me, and this is not being recorded, right?

Just for the whole world. Just that the whole world.

I play the Devil's Advocate a little bit.

We'll ask everyone to keep it quiet.

OK. My official response is, no comment. That makes assumptions about your environment, right? You could still be running an old CMTS. You could have a combiner in there. You could have a quam. You you could pull a [? Tom ?] [? Melanes ?] and run a [? revidium ?] stabilized, you know, tong out there.

You obviously could inject pretty much anything you wanted.

Exactly.

With a generator, you know?

It could just be a CW from an analog generator, which doesn't cost you a quam license.

Exactly. But that's no fun. It's fun to see [INAUDIBLE].

Yes. That looks much better.

No, but you're--

Or an OFTM channel all the way out and 1.218 gig.

Yeah. That's money. And we do that. Here We'll stand up an OFTM channel wherever it is you're running it, and use our DOCSIS population of spectrum analyzers, which right now is-- just 50% of every DOCSIS device on Comcast for footprint has a spectrum analyzer, minimum one gig to spectrum visibility to it. So the cool thing about this-- and I'll go geeky on you guys-- you know what nyquist theorem was all about. You can instruct the FFT parameters on a spectrum capture engine to hit that thing at exactly the right place at 25 kilohertz resolution bandwidth, and you can sample the OFDM sub-carriers with excruciating detail, and accuracy. So if there's a roll off in there, you can literally see the pilots. You can see the PLCs. You can discern a lot about the performance, before you even deploy a single 3.1 cable monitor. So there's a lot of intrinsic value in your system population now.

And you're just talking about the full bandwidth capture capabilities of deployed DOCSIS 3.0 modems, Larry, is that correct?

Exactly. Absolutely. Yes. So now, I'm not saying we don't sweep, and we don't balance. But what we don't do is go sweep them out where we don't have to, right? So think about the operational efficiency. If I have 10,000 route miles in a-- we'll talk a little about that a little later in the show, Brady-- 10,000 route miles, and we use this technique. And you can basically weed out 80% of the network that's not of interest. This is perfect from here. So that now you can focus you very limited and high value resources out in the field to go fix the problems. So spend time fixing problems, and not time finding problems.

Sure. So what you're doing is you're weeding it out. You're saying, because we don't see any suck outs, we don't see any peaking, we don't see noise, roll off, et cetera. On all of these full band capture modems that are deployed in this x-amount of plant, there's no sense in going out and sweeping this x amount of plant, because the modems aren't reporting any problems back to us. Is that an accurate assessment?

It would be a foolish waste of time to go sweep perfectly good plant.

Yeah. That's awesome.

You know what used to crack me up as we'd always squabble over a dB here and there. And then originally people were balancing based on CW carriers at the high and low end. I'm like, how do two carriers represent your entire spectrum? If you have roll off on both ends, you could be balancing, changing your pad based on 3 dB low.

So now it's all humped up in the middle, which creates even more power, and more distortions. [INAUDIBLE] being based on a sweep trace, and getting the median, or the average power where you need it to be, or normalized. That really, to me, makes more sense. If it rolls off, it rolls off. You don't pick the whole thing up just so there's two end points at zero reference point. You know, now you're creating a hump in the middle. I've been doing it for a long time, and I always thought that looking at the sweep trace as a balancing technique was much better than just relying on a couple of CW carriers.

Or you're missing lots of suck outs in-between. And I think what Larry is talking about is exactly the huge benefits that we're seeing in advancements in PNM, where we can use full band capture modems to see huge areas of the plant. We have visibility all the way to the subscriber's home, not just to a sweep point that might be off of an amplifier, or tap. We're seeing right into the subscriber's home, any impairment that may exist between the head end, and many, many, many individual subscribers homes. So it's huge.

It brings up a good point too. Originally, I thought this whole talk, you know, was about PNM. We think PNM is upstream pre-EQ. But PNM is everything now, right? It's down stream, upstream, upstream spectrum analysis, full bandwidth capture, pre-equalization. I think the 3:01 modems can generate an upstream sweep trace, practically, like a sweeping carrier. I even used that to a lab one time to prove that the modem itself, its internal filters was bad, because they'd shut off upstream bonding, and upstream three and four were sucked out. And it was the internal filter had a 4285 filter that was kind of overlapping itself. It was sucking up right where the crossover was

Yeah, crossover.

Yeah. So I mean, what was kind of interesting is this is the type of functionality we have in the modems now, and it's, how do we utilize that? What else do we got in the 3.1 modems. I mean where are we going with this hang out, Brady? I'm trying to segue us.

So well, I mean, basically we keep on getting more and more advanced features. What do we want to talk about next? Or what do we want to reveal next?

Well let's talk about the upstream. I mean, this is the secret sauce. It's the Achilles heel of cable. I'll share some numbers with you, just off the cuff here. You know, from the maintenance technician perspective, we're talking bucket trucks, not vans. You know, it's upwards of, conservatively, depending on the area, 70% upwards of 85, 90% of the time that a network maintenance technician is spending in the field in the pursuit of return analysis, and mitigation, of course. And by mitigation, that may or may not be fixing it. Let's just be real frank about it. A lot of times we can't get access, a maintenance tech may not have fulfillment tech. A maintenance tech might not want to go in a house. The house is most of where Ingress comes from. That is the Achilles heel of cable. And there's a lot of inertia on that. And how we chase it in, and what we do about it, is involved quickly. So let's all let's shift gears to upstream in general.

OK. So I think we're doing a lot of things with upstream. We do upstream spectrum analysis. We use the CMTS to do that. We've just pushed through a new ESO at Cable Labs as part of the DOCSIS 3.1 specification, that will allow us to have a standard mib, for all CMTS to do triggered upstream spectrum analysis, with all DOCSIS 3.1 CMTSs. So that's nice. Everyone has a common mib a common software development platform across all CMTSs. Now, there will be some time before vendors can implement this. So there's some excitement there. We'll have that capability in it.

And remote fi, as well.

What's that?

I said remote fi, as well.

So that is a really, really good point to bring up, John. It's one thing to have a head end, where all signals are culminating at, and you can put whatever type of equipment you want. But you talk about a remote fi, where you have-- and if you don't have an understanding of remote fi, you basically are taking the RF portion, the back portion of your CMTS and putting it out in a node, right? And you're taking your RF signals at that node, that the fiber node. You're converting them from RF, back to data, and then transporting that data optically, no longer as analog, but as true digital data, back to your CMTS core. So at this point, if you want to do any return-path analysis of, you know, where's my noise coming from, it has to be done at that fiber node. So now, to your point earlier John, we can't be hanging a bunch of test equipment off of a fiber node. We have to be communicating--

Not going to put an [INAUDIBLE] out there.

Yeah. It's an expensive place to put it. So now, we want to be able to communicate with that remote fi. If we know we have a noise problem out there, you know, maybe we have some really high, uncorrectable code word errors, and we believe we have noise in grass, we want to be able to communicate without remote fi, either SNMP, or IPDR, or maybe some other communication method, and say, you know, if my return is 5 to 42, 5 to 85, or even 5 to 200 megahertz, where's the noise at? What type of noise is it? Is it impulse noise? Is it HPNA? Is it some other type of noise? I want to be able to see it, just like using regular spectrum analyzer, or return-path monitoring system, and that's what this basically ECO that went through, gives us the ability to monitor that remote fi, without having to hang heavy and expensive test equipment on with the fiber node itself. So this is awesome capability.

I mean, it's good to hear, but I mean I have experience with the CMTS spectrum analysis, and I told you it was never meant to replace a real spectrum analyzer, because of how quickly you could capture, say impulse noise, because with the spectrum analysis, it was SNMP based, so how quickly are you getting SNMP responses? So getting that SNMP, and how quickly is it-- what is the resolution bandwidth, video bandwidth? I couldn't give zero span mode, so I don't know if this ECO addresses how quickly it can capture it. Can it really capture impulse noise? That's my point. I never could really capture impulse noise before.

Yeah, if you're lucky.

Get in Larry.

In full service, and discloser, John, can you talk just a minute on your early involvement with-- I believe it was CBI-- no, sorry, what was the two-- CBT, yes.

Cisco broadband troubleshooter. So basically, it was Cisco's proprietary mib, but we've given it out to a lot of people. And it really was the information from the [INAUDIBLE] [? mark ?] TI, into an upstream chip, and then making a dewey to pull that information. And I've even, to get a faster response, I could set my SNMP to pollings zero. Meaning, do SNMP as fast as I could. But it was still NFFT. It was still grabbing information. And it was great for seeing the spectrum. It was great for tagging a specific Mac address, and because a CMTS does all the timing, you could also say, show me the upstream with no bursts. So that was a great functionality, because you know, PathTrak couldn't do that. PathTrak had no idea about timing of many slots, but a CMTS does.

So you could say, you know what? Show me the Ingress end of the carrier. I've even used the CBT to say, let me see the upstream port with no traffic, and if I saw traffic, then I knew it was bad timing, or it was actually upstream that was meant for different port, and poor isolation in the head end bleeding over to the wrong port. So I was able to kind of make some, you know, deductions from seeing, and knowing how it worked. But I was never able to capture impulse noise. Maybe it would once in a while. And the one thing I was always asking for, that I never could do, and maybe Brady can mention it if we are going to do this. I don't know if the chipsets will do this. I want to go below 5 megahertz. I need to see AM radio noise, that's causing laser clipping. So I want to go down to 0, or 0.5, or 500 kilohertz, or something like that, and then all the way up to, say, 85 now. I mean, we can do up to 85 now. And anything on the chip sets we have installing now will go to 204 megahertz, because they're DOCSIS 4.1. So I don't know if that kind of covers what you're after.

No, so it's a great history about, and here's kind of a segue to talking about using the CMTS burst receiver to the spectrum analysis the capability. Has been there for a long, long time. What year would you say, John, you started it CB2, or first a came to your attention?

Shit. I've been there at Cisco 15 years now. And that was there, probably, 12, 10 or 12 years.

Wow. So what we're getting is a-- not intentionally plugging the product-- but Nimble does exactly that, right? We expose the CMTS burst receiver mips to offer kind of a low cost, in terms of spectrum analysis, using the burst receiver.

Yes. It's exactly correct. ANd some of the things that John mentioned, like being able to trigger on a sid, or a Mac address. Or a time when no modems are transmitting, so you can see just noise floor, all that functionality is there. To John's point, right now, with the existing CMTSs, you can't trigger on burst noise. One of the things that we've done in the new mib, DOCSIS 3.1 is that you have the ability set thresholds. So ideally we can capture that burst voice, that we can't do in DOCSIS 2.0, or 3.0 CMTSs.

And the reason I ask is, so we do this on Comcast too. It's obviously very intriguing to be able to use the burst receiver because then, in particular, the facility's real estate is really getting into a crunch, now that we're kind of going through architecture migrations, and these things, our head end's our, basically, the golden space now. So rack space is important. And we really like having the CMTS bust receiver's knowledge of many slot timing, for the reasons that we mentioned. It's pretty cool to be able to see under the carrier, truly, at exactly the moment when there should be an open transfer. So you don't have to [? min ?] [? hole ?] [? trucks, ?] and all this other stuff that we've done over the years, you know, to try and perceive you know, things about the spectrum.

But I'll be honest, that has limitations too. Burst noise is one of them. You know, we dynamic range problems, and depending on how people set up their CMTSs, what the nominal level of the receiver level is set up for. You people padding, or no padding. It's just all over the place when you have such diverse environments, at the burst receiver. So, you know, we're having challenges with that too. But very much looking forward to the 3.1 hooks that are coming. It's just going to be extraordinary. Very much game changing, from the upstream perspective. One other issue, and this I think is probably a segue to the other stuff. What else is out there, Brady? Using the CMTS is fantastic, but that doesn't solve for like fiber deep noise, where now we're starting to do pretty heavy combining of the CMTS ports, instead of having to drive out to one or two nodes, basically to isolate the Ingress point, or the leg of the node. Jeez, now you got to go out to, minimum, four nodes, before you can even begin to troubleshoot out in the field.

Right. So I'm going to share just a couple of slides here. We're going to be doing a press release at Cable Tech Expo. So I don't want to go into too many details here on this, but I want to share this, because it's some really cool stuff. We'll just call this, for now, PNM remote monitor. So can you guys see what I have on my screen?

Got it.

Yeah, I got it.

Awesome. So we have this little silver box here, and it's a cable modem. But as you see, it has additional ports on it. And the little cable modem you would put into an amplifier. So one of the ports, this main cable modem port here, is going to go into a test point on the amplifier, where it can see both downstream, and upstream. The other ports are going to be monitoring test points in the amplifier itself.

So you can see the upstream and also, more importantly, the downstream. And so what this would be, that actual device installed in an amplifier. So again, this is just a cable modem. It's a full band capture cable modem. Port one here would be the cable modem port itself. So here, what I'm showing here, is a screen capture on it. I didn't want to do a live demo, because that always creates problems, and takes a while. And so I thought I'd do a screen capture. We can see, we have preload noise floor on this port, down at minus 50 dB, MV. So we have a pretty decent dynamic range. You can see the forward channel here. So this goes all the way up to a gig.

So to your point, John, we can see whatever the forward channels are, and the returned channels. And then we can we can also see the modem temperature. So this is in degrees Celsius. If you click on it, it turns to degrees Fahrenheit. It's nice to see the basically in. So if this is installed in a fiber node, or maybe a trunk amplifier, bridge amplifier, you can monitor the temperature while you're going. And now, we want to look at one of the other ports. And, you know, we we've been told there's a bunch of modems out on port four. So the next slide it's going to show you the difference between port one, and poor four, on this amplifier, and we're going to see that the noise floor goes up pretty significantly.

Wow.

This is this is port four, on this same exact amplifier. So like, to what Larry was saying, you know normally, if we wanted to compare the noise difference on two ports of the same amplifier you would have to go out and pull the pad on this amplifier, and see what the difference is, back at the CMTS. But now if you install this little tiny silver box. It's a very tiny modem. And this modem is powered off the power pack in the amplifier itself, you can just toggle back between port one, and see, hey, port one, really low noise floor. No problems, if I could also see the forward pass signals. Port four, this is why my modems are offline on port four. We've got a really high noise floor. So--

I mean how are you plugged in? You're plugged in to the test points?

Yeah. We're plugged into the test points in the amplifier.

Bidirectional, are you outside of the [INAUDIBLE]? I guess, it almost seems like you'd have to be plugged in to an upstream reading, and a downstream output test point.

So it's going to depend on the amplifier. You know, some amplifiers have bidirectional test points. You're going to see the [INAUDIBLE] forward and return. Some amplifiers, you're going to plug into the return-path test points. And so then you're only going to see the return on i.

I actually wonder if like, if you had that case, if you just plugged in to the upstream reading test point, and then you kind of eye-balled, you assumed a certain amount of isolation of the [INAUDIBLE], you would still see the downstream, right?

Yeah. You're still going to get some downstream level, but it's going to be, you know, on a 20 dB down test point. So you're going to still see some. And this is just an example of that same port four. You can see that not only is there a lot of Ingress noise, but there's something there's something serious going on with the forward on this as well. So you know, it's impacting both the return, letting a lot of Ingress in, and it's also seriously impacting the forward, because you can see there's a good bit of ringing going on in here.

That's just brilliant.

Tremendous, tremendous diagnostic capability, without ever having to open up the amplifier. I mean, we know there's a problem on this one port of the amplifier, you know the other ports are OK, so right away, we have the ability to make a diagnosis know, which amplifier, and which leg of the amplifier is having the problem, all from the comfort of your desk. And then send the right person to the right place.

And to be clear, what's really intriguing by this example you pose here, Brady, is that the age old Achilles's heel of cable is, that port four, right now, you have a port level resolution of where the noise is. Before this kind of a solution you would have a node level resolution, or at the combining level you'd see the noise getting in. But you would have no idea which leg is getting in. So that, from the fuel perspective, it's a trip out to that active of-- well, actually you start at the node, and you begin breaking it down to the leg. You're disrupting your customers by pulling pads, and you're driving trucks to places, dividing and conquering, and blocking tackling, and then eventually get there. This short circuits all of that stuff.

Yeah. So we're super excited they'll be on a press release about it, right? At Expo. So you'll find us you'll find us on the show floor, demo-ing this product at Expo.

Very exciting. I'll be there.

Yeah, I mean it's interesting because we've tried to push status monitoring forever. And now I think, maybe with the technology, and lower Cascades, and the importance of everything, and the cable modem, and the inexpensive way of doing it, this makes a lot more sense now.

Yeah. It's is super cheap. It's commodity technology.

Yeah. Timing is everything. And it's funny. So I'll be just honest with everybody here. I've seen this before. I'm in love with it. It's fantastic. And my understanding is we have seen folks with this thing running as low as 1.5 watts, the heat signature is physically nonexistent, if you have it mounted properly to the shell of the amplifier node, or whatever you're going to put it in. The cost is commodity based. It's mind boggling. But it reminds me of kind of when the iPad took off, right? The iPad was not the first tablet ever invented. There had been tablets since the 80s. But the technology wasn't right. It took the stars aligning for touchscreens, and speed, and doing away with the stylus. And it took everything to get right before it basically transformed the way we compute. And I see this as being one of those opportunities in cable. Just the time, and the technology is finally right. It's very exciting.

Hey Brady, was there a certain resolution that was setting on that? Or is it unset?

This is a cable modem, so you know, you can go-- right now we have a program to go as low as 30 kilohertz. The default is 100 kilohertz. Or you can send it at 300 kilohertzs. Those are our three default settings. But, you know, you can change the bin size on it, using the standard DOCSIS mibs to change to other resolution kilohertz bandwidths if you would like to.

Nice. Very nice. So let's take a devil's advocate view here. I know we're pushing up on close to an hour, but you know, I see trouble with this. Like John, eluded, you status monitoring, and transponders have been around for a long time. And arguably, some places, and systems, I've let them lapse, and haven't used them. And just because you've got it there doesn't save the day, right? You have to integrate it with your operations, and it has to become a thing. What makes us think that this is really right? What makes us think that we're going to start sticking things out in our plan, activating them, and now have yet another point of maintenance, and power consumption, and thermal considerations, and that aren't going to someday break, or become out of date, that need to be removed or maintained? What makes us think we're going to actually deliver the promise this time?

So I think one of the things that we have going for us is that we already have some proof of concepts out there in status monitoring modems, that are in power supplies. So we have proof that this works, because we have a lot of power supplies with cable modems in them that are being used as full band capture devices, today. And we have a lot of cable modems out there. I mean, people are buying cable modems, and putting these next to power supplies. Commercial cable modems, because they don't want to pay a little bit more money for the status monitoring, you know, sort of ruggedized versions that come with the power supply. So

That kind of gives me a really good feeling that if that works, and we're buying DOCSIS 3 status monitoring devices, putting in power supplies, and those are working really well. We're getting good data off of those. Then this is the next evolution, where we're taking a cable modem, an industrial great cable modem, putting a four port switch on it, and now putting it in an amplifier, this is just taking it one step further.

It's genius.

It's going to be, I believe, solid. Will we have some growing pains? Probably. Will we overcome those growing pains? Absolutely.

Fantastic. So to that note, I will say, I've seen that at Comcast too. We have really fantastic status monitoring, we have transponders, systems. But one of the motivators that I've seen, individual technicians and systems, taking that put a commodity DOCSIS box out there as a measuring point, isn't necessarily the cost, or the trouble, or expense. It's really because it seamlessly integrates with all the tools we've got already, right? If you drop one of those out into a power supply, you know, it now magically participates in all of our Comcast tools. Spectrum up blocks, last tower. And it becomes kind of a nonevent to extend the capabilities of the network monitoring, without having to swivel to yet another application, just because it's plain old DOCSIS.

Yeah, to your point, I mean we've had many things in the past, like you said tablets, and stuff, that didn't work, because they weren't part of a standard. Apple's iPad works so well because it was part of the Apple ecosystem. It worked with iTunes, and everything else Apple. This works so well because it's part of DOCSIS. Is part of the DOCSIS ecosystem. It works with SNMP. It works with all the tools that you currently have. That's what makes it so elegant.

Exactly. I love it. Thanks for that.

I don't know the whole functionality of it Brady, but I'm thinking, you know, as a cable modem, I'd like a cable modem to see on the downstream between minus 5, plus 5, and I'd like it to transmit about 45. So if you're just plugging it into a node, through a 20dB test point, and you don't put extra padding, it could be transmitting as a 25. If it goes off the rail, at some time, it can go up to 55, and then cause laser clipping on itself, you know? So I always like to keep my modems transmitting close to 45, so there's not much range left if it does go kind of crazy. You know, it starts babbling. So 45, 48 would be good.

But I'm thinking also, you can't control which MSO, where they put their DOCSIS channel. On the low end, or the high end. If you come up off downstream test point, you're going to have tilt. I think your-- I'm going to call it your black box silver box. Your silver black box. Your silver black box, I almost think you should put a cable simulator in there, because you're always going to see tilt from the downstream test point. So you could drive put 10, 12, 15 dB of cable simulation to flatten it out, and then padding in your black box, so that you know-- I mean, basically all nodes are you going to come out, you know, distribution levels, nowadays, because there's really no trunk line. You could figure it out, right? So I don't know what you're doing about that. If you're going to have potential pads, or something inside that black box. Maybe go back to dibble sticks. Bring back the dibble sticks.

No, no, no, no.

[INTERPOSING VOICES]

I've got an idea. Let me start throw my two cents in. Let's go with the software. That way you can normalize it. You can [INAUDIBLE] it with software, if that's really a requirement. Too much data is a good thing.

Yeah, but you see, you got my point, right? I mean, hitting a decent downstream level, between minus five, plus five. I don't want this device, it's a cable modem, but I've seen where cable modems are not set up properly, and they cause problems. So I don't want this to cause a problem.

Yeah, I understand what you're saying with transmitting to hot. That's something that can be done with proper installation guidance.

But as Larry just mentioned, maybe if you have software control, had EQ internally, then you do adjust whatever you want.

OK. So I love this. Let me recap this for you, make sure I interpret this properly. Basically, a module. And it looks like you can power it, and connect it up as a DOCSIS. Just like Honey, I Shrunk The Cable Modem Into A Nice Tight Little Box. What precludes you from sticking that into a power supply countable?

So it goes into a power supply cabinet. It could be, just like you're saying, a transponder. Monitor your downstream, your upstream. You can put this anywhere you like, as long as you have 24 volt DC powering, a wall ward, even.

OK, so if there--

We have no [INAUDIBLE]. There's really no RF at the power supply cabinet, right? I mean, we're trying to look at RF, aren't we?

You can run a coax drop.

Exactly I mean from a tap, maybe, or something like that, right?

Well, we do that for transponders today.

OK. I see it more as a node functionality, best bang for your buck, being in the node.

OK. How else? You could put it in an amplifier. That would give you-- let's say, hypothetically a use case. Five m plus four or five cascade, and I'm busy with a really noisy node, and it's in a high density residential area, with

MDUs. And loose connectors are killing me all the time, I'm rolling trucks constantly, pulling pads, trying to isolate legs on this m plus five, with all these through legs, and such. And there's tons of splints. And it's just a total pain to maintain this, and constantly mitigate noise. Hypothetically, could you embed one in every amplifier, and the node, if you really wanted to, in the cascade?

Yeah, absolutely. And I think you're bringing up another point I hadn't considered before, were MDBU environments, where you have big clusters, big closets, full of connectors, and splinters. Actually, an MDU might be a good place to put these for monitoring, where you cluster many of these together in a shelf. So these become like awesome monitoring systems.

They make MDUs that have maybe three or four risers, right? You might get able to segregate--

Absolutely. Yeah. You know, MDUs use are notoriously difficult to troubleshoot. Especially, like with a PNM system, where it's difficult, because your stands, where you have echo cavities, are very, very short. So having something where you can start monitoring different floors, and different legs very quickly, and troubleshooting problems, and in high density buildings, multi-dwelling units, this becomes something that's because it's low cost, it doesn't cost a lot to deploy many modules in this. So it becomes a very, very-- the potentials of this becomes limitless, for deploying it.

Yeah, thank you generic module. We love it. So let's think about what he said. MDU is a fantastic use case. I mean, I just kind of posed it, but didn't think about it that way. But you're right. We could put it on a rack in the cabin, at the MDU complex, and just automatically start dividing the problem into spans and legs, without having just to disrupt the entire complex.

And while you're rack mounting it, you may also want to put it in your head end, or your hub site. And so there's many places that-- if it's rack mounted, you could do a lot with this device.

Yeah. That's fantastic. So and not just upstream, but also downstream, depending, obviously, on the capability of the test points. But man, can you imagine a head end instrument in full test point upstream, and downstream, visibility in real-time?

[INTERPOSING VOICES]

Up to a gigahertz. It does forward path, and return-path monitoring, it all in one device.

Unbelievable.

It might not work in line extensions, flat lids, right? Because there's no space. You're talking more like amplifiers that have a lid that has spacing.

I don't know. Brady, how thick is that thing?

It's shy of an inch. I think it's maybe 3/4 of an inch, or less. It's maybe even less.

The one I saw was a quarter inch thick.

Oh, is it a quarter inch? OK. I haven't measured it. So it's very thin, actually. So it probably will fit in the line extender lid. But I don't know.

All right, I'm being coy. I've seen it in a line extender?

You have?

Not in my line extenders, but some it will work in. It's pretty cool.

Very nice. So, Joe, as you mentioned, being devil's advocate. You know, you talk about it's a rack mount. As extra test points, and all that. But what's going keep someone from saying, well why would I put this in the head end, when I'll just put a modem in the head end? It's a modem, I have a modem. So why would I put this in? What does it give me that's more so than just a modem?

One thing that I noticed about the, at least three of those ports on the RF switch, is they're uninhabited [? Dioplex ?] shelter, that you don't get from a cable modem.

Fair. That's what I was leading to.

That's the secret sauce, you don't have the [? Dioplex ?] rejection.

It's also density and power, because you've got multiple ports. You'd have to have multiple modems. You don't have the [? Dioplex ?] filter there, and they're very low power.

Yes.

Now, if you do a head end, or a hub-site, is it going to be 110 volt? Can you plug it in with 110 volt?

You would have a power supply to feed DC to it.

That's right.

That might be tricky.

Where'd you go and put the power supply?

Oh, jeez that's funny. You guys are great. I got to be honest. I know we're coming up on an hour. I still have a little time. But it's such a thrill for me to be here with you guys. I've known you for a long time. We're all part of birds of a feather with this PNM business. I can always count on seeing all of you at Expo, and the Cable [INAUDIBLE] conferences, and we can talk for hours, and hours on this stuff, because it's all of our passion. So it's really a pleasure to be here with you. I've known you guys for a long time. I'm huge fans. Thanks a million.

And Larry, you're doing a great paper at Expo this year, right? It's--

Thanks. I didn't want to do the cheap plug, but because you brought it up, yes. Thursday, the last day. Thanks STCE council for doing this to me, but they always-- I think they like to do a book end kind of thing. It's the very last day, the very last thing on Thursday. But it's cool. It's kind of an institution. It's the PNM, you know, conversation. This year I'm going to be doing a pretty comprehensive case study, which we'll talk about a lot of the stuff we've talked about here. A lot of gory detail about the financial aspects, and the things that have been traditionally hard to ensure, we've put a lot of energy into that, about cost versus benefit, PNO, and areas of adoption, and all of these things. So I would encourage anybody to come say hi. If you listen to this podcast, mention it. I'd love to hear that. I know there's thousands of people that do tune in occasionally to your great show, Brady.

All right, I recommend everyone catch Larry's show, or Larry's presentation, at Expo. Stay for the last day. John, I think you and I are both out this year. We're not presenting, right?

No, I didn't submit anything.

So we're taking this one off. We'll make great, [? cashiers, ?] Larry.

You just come up on stage with me, because you deserve as much credit, cause you guys are pioneers.

Larry, are you going to cover DOCSIS 3.1 stuff as well, for PNM.

It's like a couple of pages. But there is there's some great stuff. Another plug for my colleagues at Comcast. An old, great friend of mine, Rob Thompson, will be going over, kind of a very novel, not DOCSIS 3.1 specific act. But that's on the first day. So they-- I said bookends. Don't come in late, and don't leave early, because you're going to miss all my favorite stuff. Rob Thompson's going to be doing some stuff. But Maurice Garcia is doing some really fantastic work at Comcast, and in the same session with Rob Thompson, on Monday. So don't come in late. Go see Rob and Maurice, and I'd stay all the way through, and come see me. And by all means, anybody out there, we're all friends here. Look me up. I'd love to go have a beer, and a pizza with anybody who loves PNM as much as we all do here.

It's supposed to be a pretzel, and a Philly cheese steak.

Pretty much.

All right, guys. Hey, thanks so much for your time. Great information. Have a great weekend. We'll catch you later.

Oh, PNM.

Thanks guys.

Thanks a