

Mesh Networking and AREDN

Thirty years ago, packet networking was the new “in” thing. Many of us marveled at our ability to communicate digitally over hundreds of miles using little more than an HT, computer terminal, and TNC by traversing the network. Net/ROM, TheNET, TexNet, and ROSE were the choices available for firmware, and each had its advantages. Later, such systems as Germany’s FlexNet increased the utility of these networks, and hams by the thousands were drawn into the digital age. Commercial competition was in the form of dial-up bulletin board systems and pay services such as CompuServe, America Online, and others.

Around that time, the internet formed, and the rapid rise of reliable and inexpensive broadband access to the World Wide Web proved too tempting. Tunnels through the Internet were used to expand packet networks (to their eventual detriment), and some decided that 128 kb/s DSL access was far superior to 1200 baud. And so, the radio was leaving packet radio networking, as were its users.

But packet never really died. It proved valuable for emergency communications (EmComm), even at 300-baud speeds on HF bands, because it was far faster than anything else available. The APRS (Automatic Packet Reporting System) phenomenon had many hams repurposing their terminal node controllers (TNCs) to mobile and tactical uses, including EmComm. And WinLink 2000 combined internet, VHF, and HF radio to build a

worldwide email system that, while slow, worked reliably and is credited with saving many lives in bad situations.

HSMM

Later, HSMM (high-speed multimedia) networking became a “thing,” and folks started using low-cost and readily-available consumer-grade 802.11 (Wi-Fi) equipment to traverse short distances at very high speeds. Several standard 2.4-GHz Wi-Fi channels are on frequencies shared with hams, so the modems could be modified to operate under FCC Part 97 rules rather than the more restrictive Part 15 rules. The idea was that for temporary situations, hams could offer excellent compatibility with existing applications and equipment, delivering “last mile” connectivity in a seamless and transparent manner. Larger networks could be built, since deploying many devices wasn’t a financial issue, unlike packet.

The problem, as it always has been, is finding enough sites to make the network usable. Sure, there are a few hams in almost every town, but not everyone is interested in participating. Consumer-grade Wi-Fi networking equipment also has a limited range not because of power (tens of milliwatts), but due to built-in 802.11 networking firmware that requires the receiving station to acknowledge receiving data within a certain timeframe. If the ACK doesn’t come soon enough, the sender assumes it was lost and re-sends. That’s good, unless the ACK time is shorter than the speed-of-light delay we’ll encounter over distances of, say, 2 miles, in which case the link cannot work.

* c/o CQ magazine
Email: <N2IRZ@cq-amateur-radio.com>



Figure 1. A Ubiquiti Rocket M3 serving as an AREDN™ node. On these frequencies, line of sight is important, so long-haul backbones tend to be at good sites. EmComm users could then deploy portable nodes to deliver “last mile” connectivity into an area, providing intranet connections between an EOC or other operational command center and its deployed resources. (Courtesy of AREDN)

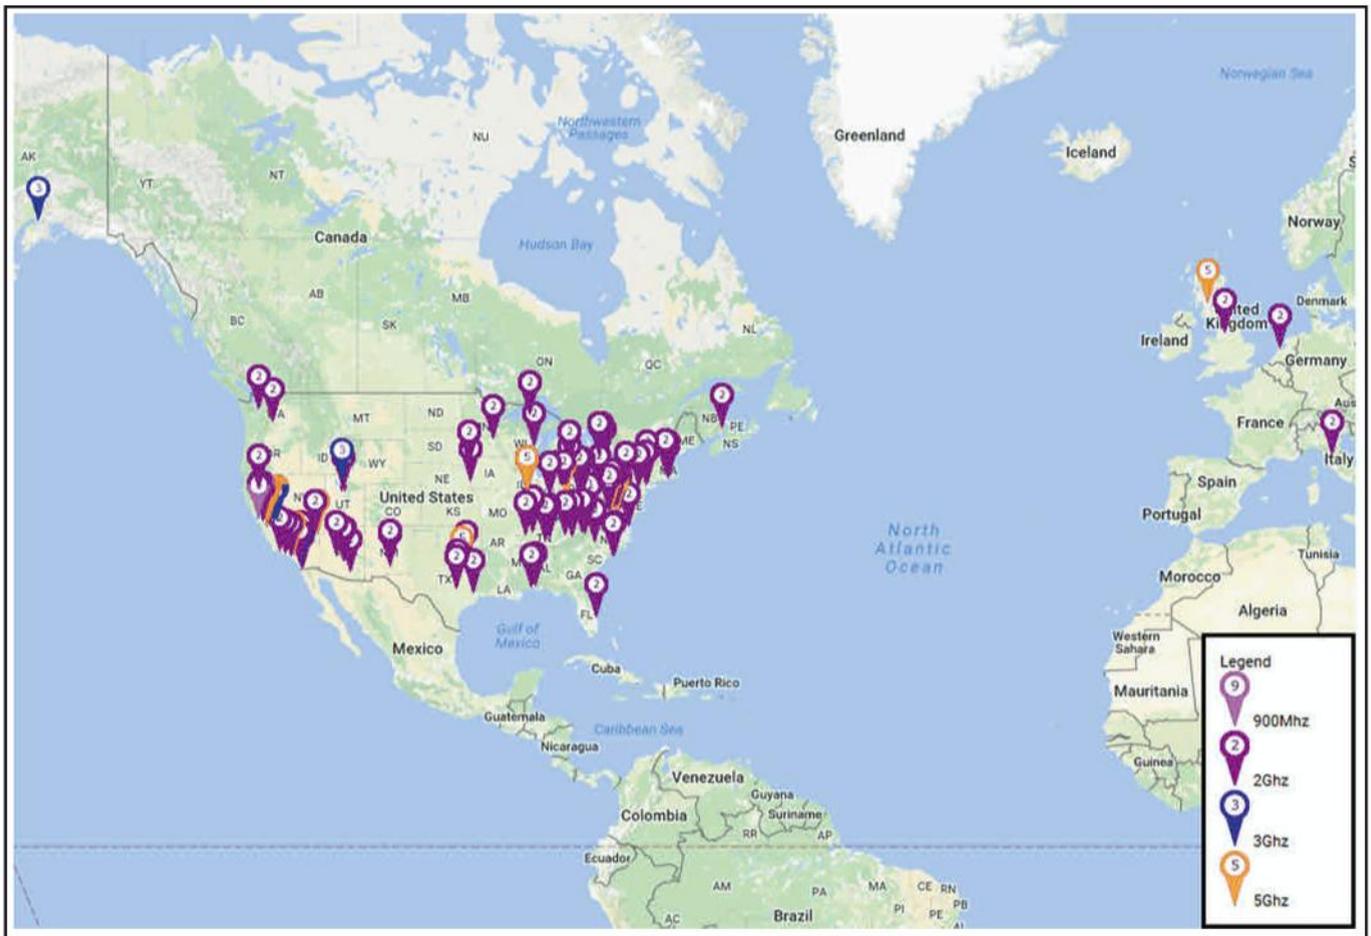


Figure 2. A screen shot from the AREDN website, showing some of the AREDN nodes. The map data comes from sites that have voluntarily posted to the map, so it represents only a percentage of active sites. Visit <http://www.aredn.org/content/aredn-map-0> for the latest version.

To get past this inherent distance limit (unrelated to transmit power), some employed Wi-Max equipment (802.16), which is built to have links of several miles. But two issues led to a low acceptance: These are Part 15 devices, and so there were limitations to be accounted for, and the cost, which could be several hundred dollars per device. Commercial users can afford to spend a lot of dollars to save even more dollars, but hams are pretty “price sensitive.”

And all this bypasses the idea that these 802.11 devices were primarily intended to connect a house full of users to a Wide-Area Network (WAN), such as the internet, and not meant to form their own network.

Broadband-Hamnet™

This problem was easily addressed by entrepreneurial hams at HSMM-Mesh™ (now known as Broadband-Hamnet™ <http://www.broadband-hamnet.org/>), who leveraged the Linux operating system of some popular 802.11 gear. Just as TheNET and ROSE were networks written for the Z-80 processors that were found in TNCs, hams wrote firmware for the Linux kernel in the 802.11 gear and repurposed it to form robust networks.

This didn't happen too long ago, but in case you hadn't noticed, we're not generally awash with HSMM networks running on 802.11 gear, are we? Sigh. The problems with consumer-grade gear — low power and poor antennas (both for Part 15 compliance) along with indoor-only packaging — still

prevent them from being a good choice for anything greater than a few hundred yards, or what I call “shouting distance.” And if you can shout, who needs a network?

AREDN™

Enter AREDN™, the *Amateur Radio Emergency Data Network*. As you can see from its name, its purpose is not building a network for fun, or to deliver alternate access to the internet, but primarily for EmComm. AREDN's approach is much like what came before — indeed, some were members of the Broadband Hamnet team — except they have focused on a newer generation of outdoor-mounted, high-power, and low-cost commercial-grade Wi-Fi gear. As before, new firmware is loaded into the device, changing it from a Wi-Fi access point to something else: A mesh network node.

Mesh Networks

A mesh network is somewhat different from an AX.25 packet network, in that it is self-configuring and self-healing (fault tolerant). A packet network was built by a bunch of hams who had to configure its interoperability with other nearby nodes. The configuration was critical: Mess it up and it didn't work. This wasn't a limitation of the hardware or protocol, but the networking firmware itself.

A good mesh network basically configures itself after it is powered up. Sure, you'll want to set up specific links using dish or Yagi antennas for good performance, but if a dozen

hams were to converge on New York City's Central Park, each with a powered mesh device, a robust network would form out of thin air as the devices learned about each other and configured themselves based on their view of reality. A near-ideal situation if some kind of emergency was in progress, wouldn't you agree?

A mesh network is also self-redundant to the extent possible. The idea is to deploy many low-cost devices so there are several connections and paths to other nodes available. That means that if a particular node stops operating, the network immediately sends data along a different path, seamlessly. Indeed, the network is always looking for the lowest "cost" path (meaning the best performance), and switches to it as soon as it is detected.

I encourage you to do your own research, but if you're in the EmComm community, this is the next big thing. While technically able to support an internet link, EmComm users know to not count on one being available. Instead, your AREDN mesh network is intended to restore an intranet network connection between operational command centers (EOC, Red Cross Chapter HQ, etc.) and their deployed resources. An unmodified Wi-Fi access point or router (like the one in your house) is deployed for those users, offering broadband connectivity both wireless and wired. If a location is underserved, someone can grab an AREDN device and put it where it's needed, with no technical expertise needed: Place it and turn on the power.

So what's the process for creating an AREDN node? There are a few simple steps, which we'll examine in more detail:

1. Buy a compatible device
2. Upload new firmware to it
3. Configure the node with information such as callsign and password
4. Mount the device and apply power

Compatible Devices

You can find the current list of supported devices on AREDN's website <www.aredn.org>. The list keeps growing, so anything I published here would be out of date by the time you read this. Most of these are in the Ubiquiti AirMAX "M" product line. Many can be had new for under \$80 on some bands, and on the used market for even less. These are available for different bands, so choosing a band for your backbone link is an important first decision.

The lower the frequency, the better the signal's penetration of obstructions like trees. However, the 902-928-MHz amateur allocation is relatively noisy from Part 15 devices and some other unlicensed users. If we tune the device down into exclusive amateur frequencies, and limit the bandwidth to 5 MHz, noise is much less of a problem.

Moving up to 2.4 GHz, the noise level is even higher, but again there are some frequencies that are outside Part 15 but within Part 97 amateur allocations. With this capability,

Figure 3. The Basic Setup screen from one of Gordon, W2TTT's, AREDN nodes. While setup is simple (see the text), each node needs a unique name, and the distance to the farthest neighboring node needs to be set (6 kilometers in this example). The user interface has several screens that are used for node configuration and network monitoring.

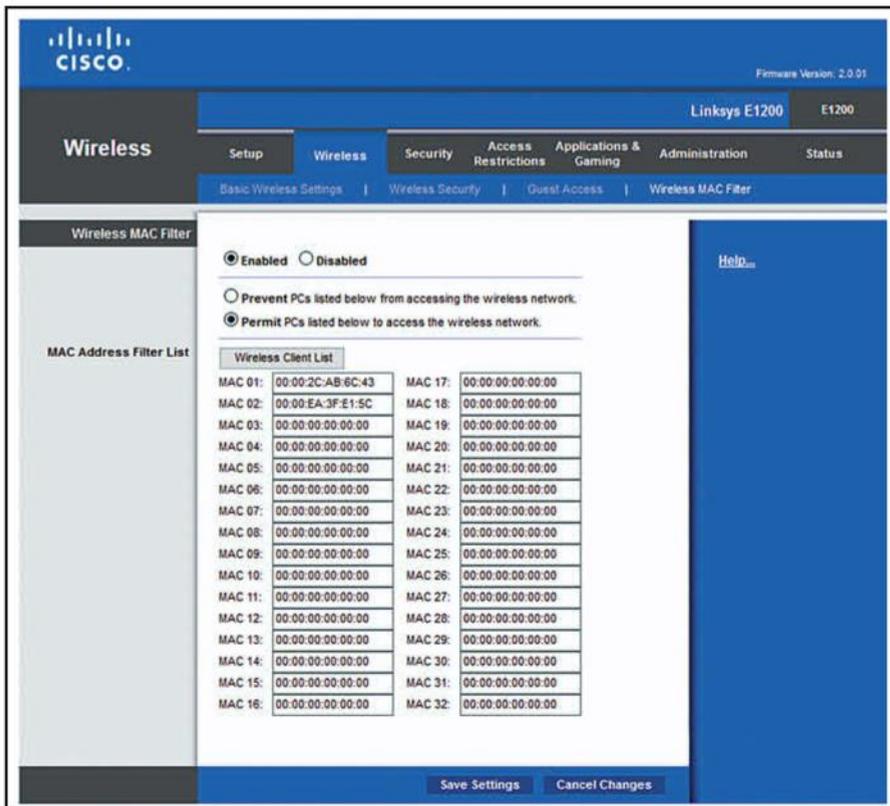


Figure 4. This is a screen shot of my Linksys E1200 router showing the MAC filtering page. Most consumer-grade routers have this feature. The data is fake, since I have encryption on and don't broadcast the SSID, but this is an example of how you can prevent unauthorized users from associating with a Wi-Fi router or access point.

nobody using AREDN would even think of using the Part 15 channels — the noise is just too high. On the plus side, these devices tend to be the least expensive compared to the other bands.

The allocation at 3.4 GHz is by far the least congested band, but in some locations the primary military users make this band unsuitable. Tight antenna patterns and a good awareness of other users can make this very useful for long-haul backbone links of 50 miles or more.

On 5.8 GHz there is a lot of room, so co-located gear has a better chance of being non-interfering because there are several ham-exclusive channels available. The smaller wavelength means that the antenna for a given size has more gain, but this is essentially offset by the reduced propagation, which behaves more like light and less like radio, meaning line-of-sight is essential. Large obstructions are out, and even trees are problematic.

Where I live, the trees and terrain severely limit higher frequencies, except for those lucky few on top of a hill or equipped with a tall tower. I was hoping to establish a 900-MHz link to a site

about 9 miles from here, but a hill in the way nixed that idea.

How did I know?

I went to Radio Mobile <<http://www.cplus.org/rmw/english1.html>>, a site run by Roger, VE2DBE, which lets you calculate the expected propagation between two points on a map. It accounts for terrain, power, antennas and several other factors. You need to create a (free) account to use the site, but that takes only a few minutes. After entering the details, I got the bad news that it just wouldn't work. Sigh, maybe I can find someone up on a hill.

Uploading Firmware

The AREDN website has detailed, illustrated instructions for uploading new firmware into your Ubiquiti device, including some warnings about compatibility. Look for the "How-To" link and follow that. The basic idea is to connect to the device through its web interface, take down some information, and optionally copy the existing firmware in case we ever want to use it for its original purpose. Then, just upload the AREDN firmware file, let the device reboot, and



Authorized Distributor for



Linear Amplifiers for the Americas

ExpertLinears.com
281-259-7877

**Fully Automatic
Solid State Amplifiers!**

**EXPERT 1K-FA ★ EXPERT 2K-FA
EXPERT 1.3K-FA**

**THE MOST
TECHNOLOGICALLY
ADVANCED IN THE WORLD!**



Built-in Power supply & Fully Automatic
Antenna Tuner – Fully Remoteable

**ELA Offers
SALES & EXPERT SERVICE**

- ★ SPE Factory-trained with over 5 years repairing Expert Amps
- ★ Fast Turnaround
- ★ Over 50 years RF Experience
- ★ Over 60 years ham radio repair experience



**EXPERT LINEARS
AMERICA, LLC**

PO Box 1224
Magnolia, TX 77353
Contact: Bob Hardie W5UQ
Bob@ExpertLinears.com

reconnect to it at its new IP address. There's also some information on what to do if anything goes wrong.

The next step is to configure your node. Start by entering your callsign as the "Node Name," along with a unique identifier so each node you create has a (slightly) different name (See *Figure 3* for example). Set a password — remember, you can reconfigure things over the air — and set the "Distance" parameter to the distance to the farthest node you expect to communicate with. Then save, reboot, and reset your wired network (by simply pulling the cable and reconnecting it). Verify that you can connect to the status screen when you're done, the node is ready to use.

Installation

Now the "hard" part: Installing your node. Basically, you mount it to a pole outdoors, point it in the right direction, and it works. If only real life went so easily.

Mounting the node — many have a built-in antenna — consists of tightening the clamps to a pole, and running an Ethernet (Cat5) cable. These devices use Power over Ethernet (PoE) which uses some of the Cat5 wires to deliver power to the device. This saves running a second cable for power.

When mounting any high-gain antenna, pointing it in the right direction is a little bit of a science, and a little bit of an art. If you can actually see the other node — perhaps with binoculars — then pointing is fairly obvious. The beamwidth of these antennas vary from a few degrees to 120°, so unless you need to squeeze every last decibel out of the link margin, close enough is usually close enough. If you can't see the other node — very likely unless you have eagle eyes — Google maps is your friend. Find both points on the map, identify the line that connects both sites, then zoom in to orient yourself with respect to local features — roads, houses, even trees are visible from the aerial view. (I've found that Google maps blurs out trees, while Bing maps doesn't, so if you're using a tree as a reference, try using Bing.)

If you're installing your node at home, maintenance is easy, since you're right there. AREDN provides ways of upgrading the software and performing a reset remotely. The Ubiquiti boxes are quite stable and reliable, and you have the option of installing remotely-controlled switches if something really goes wrong. Do you really want to drive an hour just to flick a switch?

In a real emergency, people are focused on survival, so hackers are less

likely to be up to mischief. But in a tactical event like supporting a marathon, hackers are just itching to get onto your network and see what makes it tick. To help avoid problems, even if innocent, use MAC address filtering to prevent unauthorized Wi-Fi connections: You make a list of the MAC address of each device authorized to connect, and those are the only ones who can connect.

I'd like to thank Andre Hansen, K6AH, of AREDN for his kind assistance in preparing this month's column, and my old friend Gordon Beattie, W2TTT, for his enthusiastic explanations of AREDN.

So we have a new year now, 2017. Some have big plans for this year, others are content to let it slide by without notice. I'm never one to be content with things the way they are. No, I don't have

any big plans to shake up the world, but I'm always looking for opportunities. Here in Atlanta, we don't get many ice storms, so not many chances for antenna work, but I'll be keeping busy in the workshop puttering around and making sawdust. This is a good time of year for planning out your summer, so maybe find a project and dig in.

If you liked this month's column, send me a note. Also if you'd like to see me write about something. As with most of my columns, this one was sparked by a reader's comments at the Huntsville Hamfest. I already have my hotel room at Dayton (Xenia, really) this year, so if you want to tell me where to go in person, you know where to find me. Hope to see you there!

— *Until next time, 73 de N2IRZ*

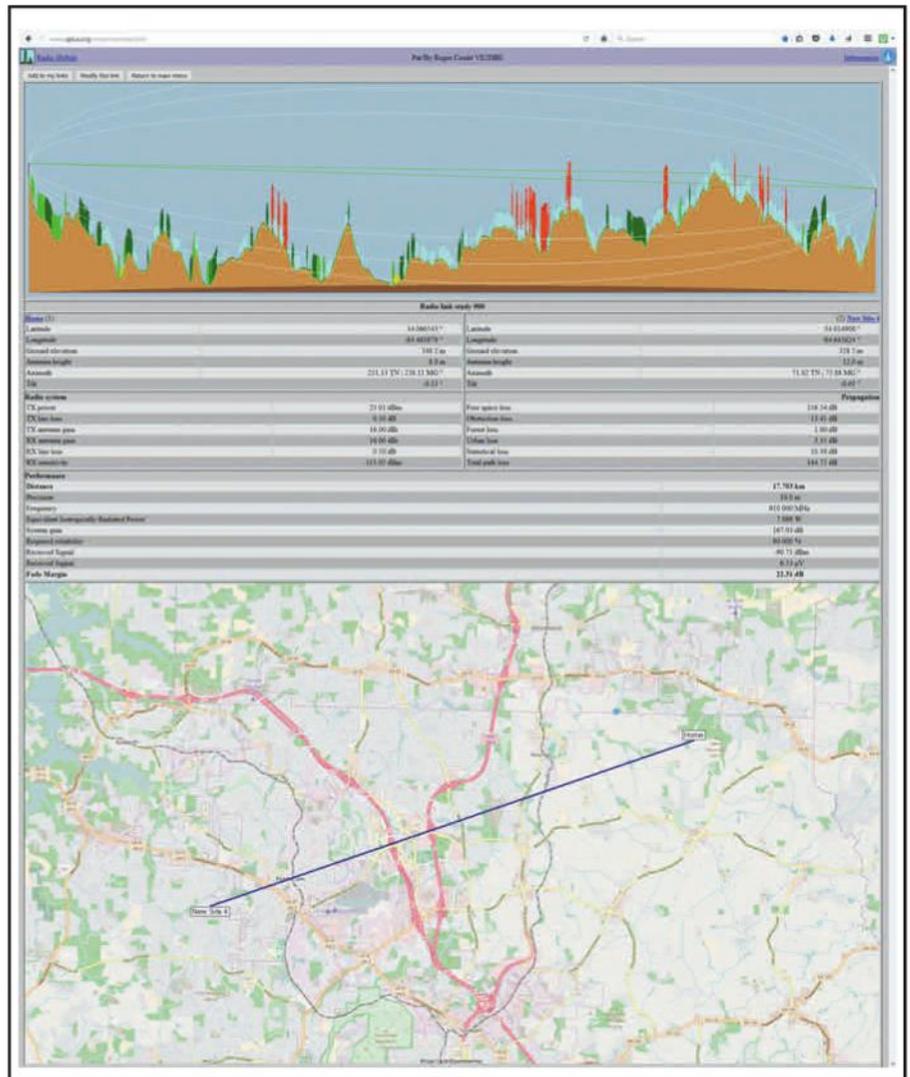


Figure 5. This is an output page from the Radio Mobile website, run by Roger, VE3DPE, showing a proposed radio path and the associated calculations. I was hoping for a 900-MHz path between my home and a nearby friend, but alas, that hill near my house (on the right) dashed those plans. See the text for the web address and other details.