Andre Hansen, K6AH
The AREDN Project (AREDN.org)
2113 Via Monserate
Fallbrook, CA 92028
K6AH@ARRL.net

Abstract

Mesh technology has been around for over ten years. Over the past two years developers on the AREDN™ team have advanced the art by porting Broadband-Hamnet's extremely popular mesh firmware to the Ubiquiti airMAX line of commercial Wireless ISP routers. This has literally changed the complexion of mesh implementations from an experimental, hobby-oriented, novelty into a viable alternative network suitable for restoring some degree of Inter/intra-net connectivity "when all else fails."

More recently, the developers of this software have kicked-off a new project, AREDN, focused on taking this technology to the next level in EMCOMM communications.

This paper begins with an introduction to the AREDN Project and mesh networking and concludes with a roadmap for the Project's future. It dives into implementation techniques and considerations as well as avoidable pitfalls.

Keywords: AREDN, EMCOMM, mesh, BBHN

Introduction

The typical Emcomm message-passing scenario today involves the sender conveying the message to a ham, who transcribes it onto an ICS-213 form. Then the message is spoken over VHF/UHF radio to another ham who writes it down on another ICS-213 form. The form is then delivered to the recipient, who reads it and signs it. The acknowledgement is then conveyed back over the radio to the sending ham who confirms the receipt to the originator.

Emcomm "Customer" expectations aren't being met

Customer expectations differ wildly than this. They expect the continued use of tools with which they are accustomed: email, phone service, chat, and other web-based tools specific to their roles within the organization.

Over $4B in ham-compatible radios is sold to non-hams each year and most hams wouldn't recognize them to be ham radios. These devices follow the 802.11 standard and operate in several of our microwave bands. They are all around us, and coupled with the privileges our license offers, we should be using this technology to deliver on these customer expectations.

So what is AREDN?

AREDN is an RF network mesh of radio/routers operating under the FCC rules, Part 97 in the ham microwave bands, controlled by hams with a Tech license or higher. It is a high-speed data network with rates of up to 54 Mbps designed to provide a TCP/IP medium when other network infrastructure has failed. While technically capable, it is not intended to be a general Internet access alternative.

AREDN is written for Linux-based WIFI and WISP devices by the AREDN Development Team which also authored the BBHN (Broadband-Hamnet) releases from v1.0.1 to 3.0.1.

AREDN replaces the manufacturer's operating system with the following major components:

1. OpenWRT, an OpenSource wireless routing framework onto which custom applications can be built
2. OLSR (Optimized Link State Routing Protocol), an IP routing protocol optimized for dynamic ad hoc networks
3. Web-based GUI for node configuration
4. Automatic device-specific TCP/IP network configuration based on the device MAC address

The primary objectives of the project are to empower the typical ham to become a deployable part of the network by simply installing the firmware, entering the station's call-sign and an administrative password, and then pointing the node's antenna toward an existing network node.

The secondary objectives are to provide a means to monitor & manage the network and to specify a set of operational standards & services for Emcomm's utilization of the technology.

To date this technology has attracted 3 very different user types:

1. Look at the cool things you can do! They're intrigued by the autonomous nature of the network and quickly setup neighborhood networks for gaming, VoIP, etc. They tend to attract other computer types who may be enticed into Ham Radio as a result.
2. Applying it toward a need. These guys weren't looking for it, but see the value in it and apply it toward a specific need, such as Field-Day logging, race support, surveillance cameras, etc.
3. Those who have longed for it… To these guys, the technology is game-changing. They are in the process of exploiting it by building infrastructure around it.

This last type include the Emcomm guys. They are the primary target of this technology and the focus of the AREDN™ Project.

How it works

It is easier to understand this technology if we start with how standard WIFI works.

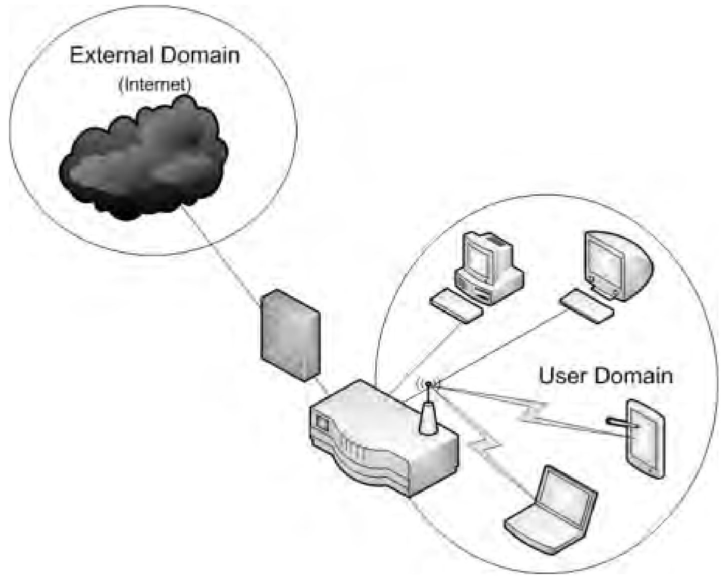

In the diagram above, we see two distinct “domains.” The User domain includes both wired and wireless devices. These are all in the same address space and nothing distinguishes them aside from how they connect to the WIFI router.

The second domain is for the Internet. You will note a firewall protects the User domain from unauthorized access and other threats from the external domain.

How AREDN “repurposes” the device

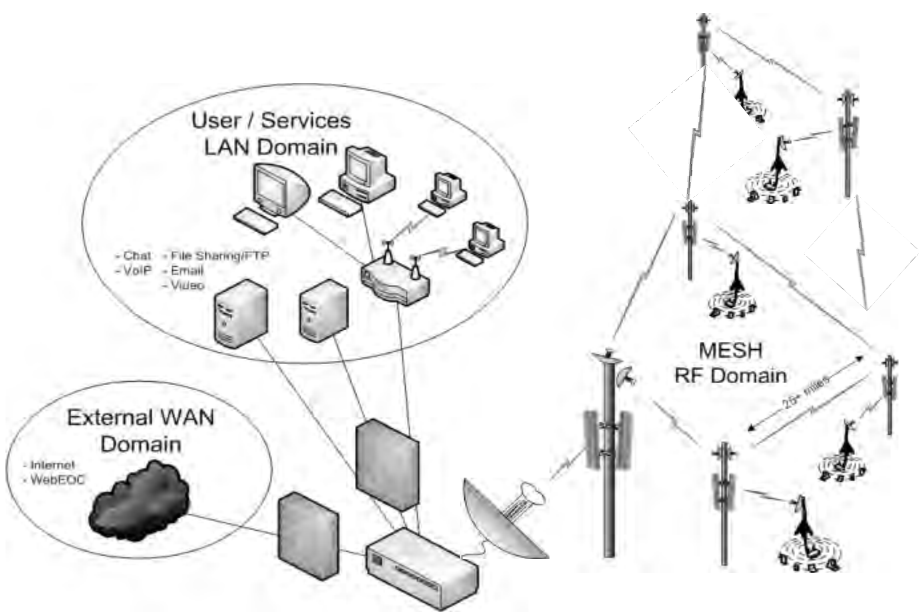

What were two domains in our Standard WIFI diagram have now become three. We see the familiar external and user domains... although the user domain now contains a WIFI router and new computers which, in this case, deliver services such as email, FTP, VoIP, chat, etc.

The new domain here is an RF mesh network which forms the business end of the AREDN technology.

The Hardware

I’ll use this term “mesh” to describe the interconnection of devices, and “nodes” to describe the devices. Nodes are typically comprised of a Linux computer, a software-defined radio operating within a predefined microwave band, and a strip-line amplifier. Some utilize an on-board transverter as a means of reusing an existing device in another band. The computer has at least one Ethernet port, although some have two with an internal hub or switch. All of these components are contained within the single node. The radios contain either internal antenna or one or two antenna ports. These radio receivers are hot… with sensitivities in the range of -95dBm. Power output is in the range of 23 to 28 dBm (200 to 600 mW).

Historically, Ham Radio mesh networks have been built on Linksys WRT54 devices intended for home and office use. These lower-powered (19dBm, 79mW) units have required environmental protection when used outside, and as a result have been difficult to utilize in meshes extending beyond the neighborhood. Over the past 2 years AREDN developers have extended existing mesh technology to environmentally robust, commercially available, Ubiquiti, hardware. The devices supported are currently manufactured by Ubiquiti under their product group “airMAX” and the TP-LINK “Pharos” series. Support of these devices has literally changed the complexion of mesh implementations from an experimental, hobby-oriented, novelty into a viable alternative network suitable for restoring some level of Inter/intra-net connectivity when “all else fails.”

The ARDEN software supports these devices on the 900 MHz, 2.4 GHz, 3.4 GHz and 5.8 GHz Amateur Radio bands.

How OLSR Works

Optimized Link State Routing determines the best path for data transmission through the network.

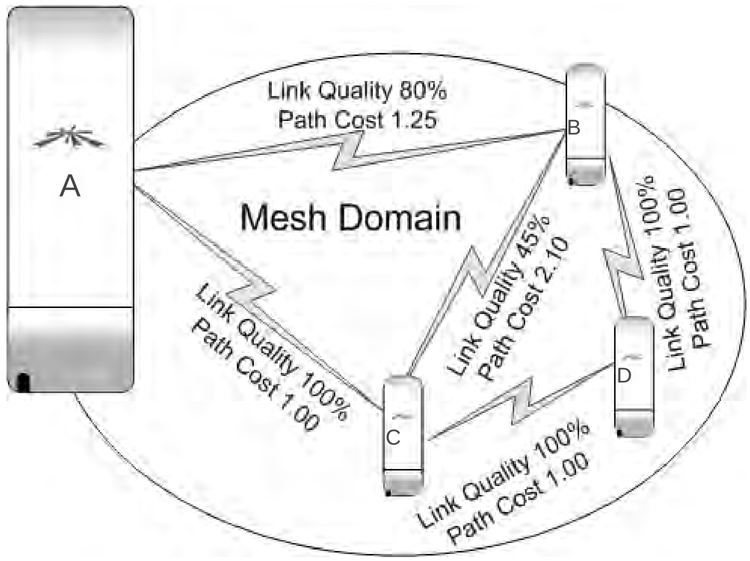

The four devices illustrated above, all Ubiquiti NanoStations, have formed a "mesh." The route data will take through this network is dependent on the quality of the links between them. Note the link between Node A and Node B. From historical broadcast reception, Node A knows that 80% of the data from Node B is received without error. In addition, Node B knows from historical reception, that data from Node A is received without error 100% of the time… it is said to be 80% reliable based on the following formula, where LQ = Link Quality and PL = % Packet Loss

$$LQ = PL_{A\rightarrow B} \times PL_{B\rightarrow A}$$

$$LQ = .8 \times 1$$

$$LQ = 80\%$$

Based on this, it assigns a "cost" (or ETX, estimated transmission) to the A→B link inversely proportional to the Link Quality:

$$\text{Cost}_{A\rightarrow B} \text{ (aka ETX)} = 1/LQ = 1/.8 = \$1.25$$

All traffic from A to B will take this path, because it is the least expensive route available. To determine the path through multiple nodes, the individual link costs are simply summed.

If the link between A and B where to fail, then Node A would quickly calculate an alternate path. Two are available:

- Path A-C-B at a cost of $1.00 + $2.10=$3.10
- Path A-C-D-B at a cost of $1.00 + $1.00 +$1.00= $3.00

Node A's routing table is updated with the new optimal path of A-C-D-B. These updates are performed multiple times each second, with routing table propagations through the mesh taking some amount of additional time depending on the size of the mesh.

What the GUI does

The Graphical User Interface is generated using HTTP by the AREDN software running on the embedded Linux computer. The GUI provides access to a variety of administrative and operational functions, such as:

- Checking for traffic / congestion on the channel
- Reporting the current node status and other nodes both directly or indirectly connected
- The basic node configuration settings
- More advanced administrative setting for
  - Port Forwarding / DHCP / Advertised Services
  - Network Address Translations for complex network environments
  - Updating the AREDN firmware
  - Installing useful service packages on the node

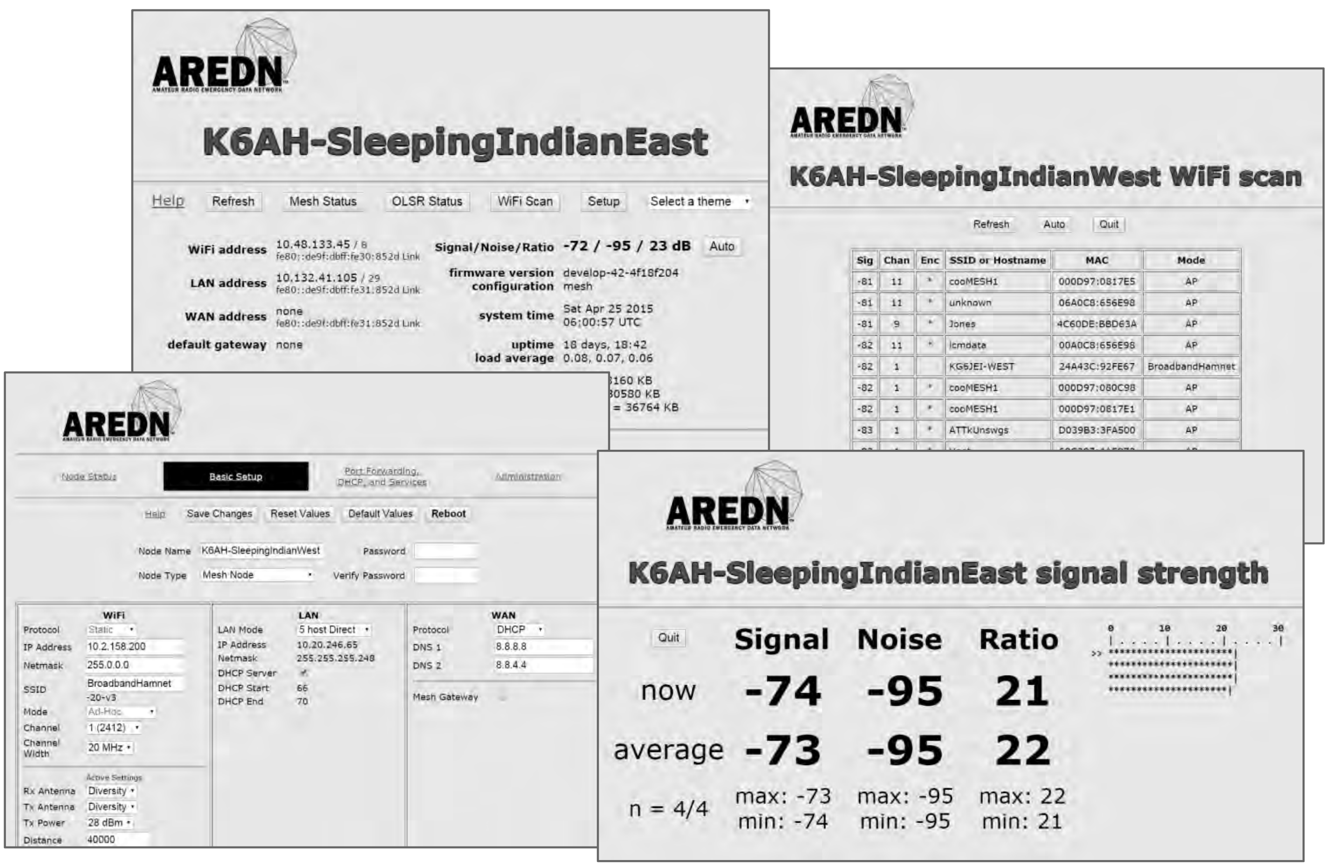

How Do I Build an AREDN Network?

Building an AREDN network is not difficult:

1. We encourage one to have a specific objective before you begin. That may be simply to understand the technology, but such an objective should not creep into a production implementation without restarting with that new objective.
2. The next step is to plan and deploy core nodes onto which mesh is formed. These nodes form an initial, primary path for network traffic. Note that they may or may not retain that distinction as the mesh grows.
3. Setup purposeful services that align with your objective and user requirements.
4. Utilize the mesh routinely to ensure it is operational and will be so with needed.

Supported Device Details

Ubiquiti airMAX M-series wireless routers share the following general characteristics:

- They are tower mountable and environmentally robust:
  - Temperature: -40° to +176°F,
  - Humidity: 5 - 95% Condensing
- Many utilize a combination of horizontal and vertically polarized antenna to minimize unwanted interference from on-channel or adjacent WIFI noise. When conditions allow, they also support the combining of these polarizations for increased data throughput—called MIMO (Multi-In, Multi-Out).

Here are a few representative devices and characteristic selection criteria:

- Rocket – A two RF port MIMO node (500mW) that is “plug and play” with a variety of Ubiquiti antenna systems: 90° and 120° Sector antenna, Dual-polarization Verticals, and 30-34dBi Dish antenna. Note that MIMO nodes split the power between the vertical and horizontal domains.
- Bullet – A single RF port high-power (600mW), non-MIMO node with an N-Type female connector suitable for direct connection to many 3rd party antennas. With the right antenna this could easily achieve ranges of 50+ km.
- NanoStation – A fully contained node with an internal 11dBi patch antenna and a 45° coverage pattern.
- airGrid – A larger node available in several size/gain grid-reflector antenna configurations. Designed to be a highly directive, it performs at a range of from 10 to 30 km

All of these devices obtain their operating power delivered over the CAT5 cabling (PoE or Power over Ethernet). They have a broad DC input voltage specification to accommodate a wide range of cable lengths which may be required for tower applications: 10.5 to 24VDC at the CAT5 connector.

In planning to deploy the core nodes it is advisable to use propagation prediction software such as Radio Mobile to avoid the hassle and expense of experimentation. I will diverge for a bit for the benefit of those unfamiliar with this type or software.

Radio Mobile is a free propagation simulation system. It utilizes data from the Space Shuttle Radar Terrain Mapping Mission (SRTM) resulting in elevation contours which have then been overlaid with satellite imaging and road maps. It further utilizes the Longley-Rice radio propagation prediction method, which computes the attenuation of radio signals using an "irregular terrain model"... a technique that has been successfully used in commercial radio coverage planning since the 1960's.

It is available both as a software download and a Web-based tool. While the download results in a more flexible tool, its installation is not for the faint-of-heart. I would advise that, if you do not consider yourself a computer expert, then the Web-based tool will more than adequately meet your needs. The English language portal is at: http://www.cplus.org/rmw/english1.html

Sufficient use of this tool to explore the variables of band, node-model receiver sensitivity, and node-model power output, will result in the required antenna gain in either point-to-point (PtP) or point-to-multi-point (PtMP) topologies.

Here are a screenshot from a typical Radio Mobile analysis:

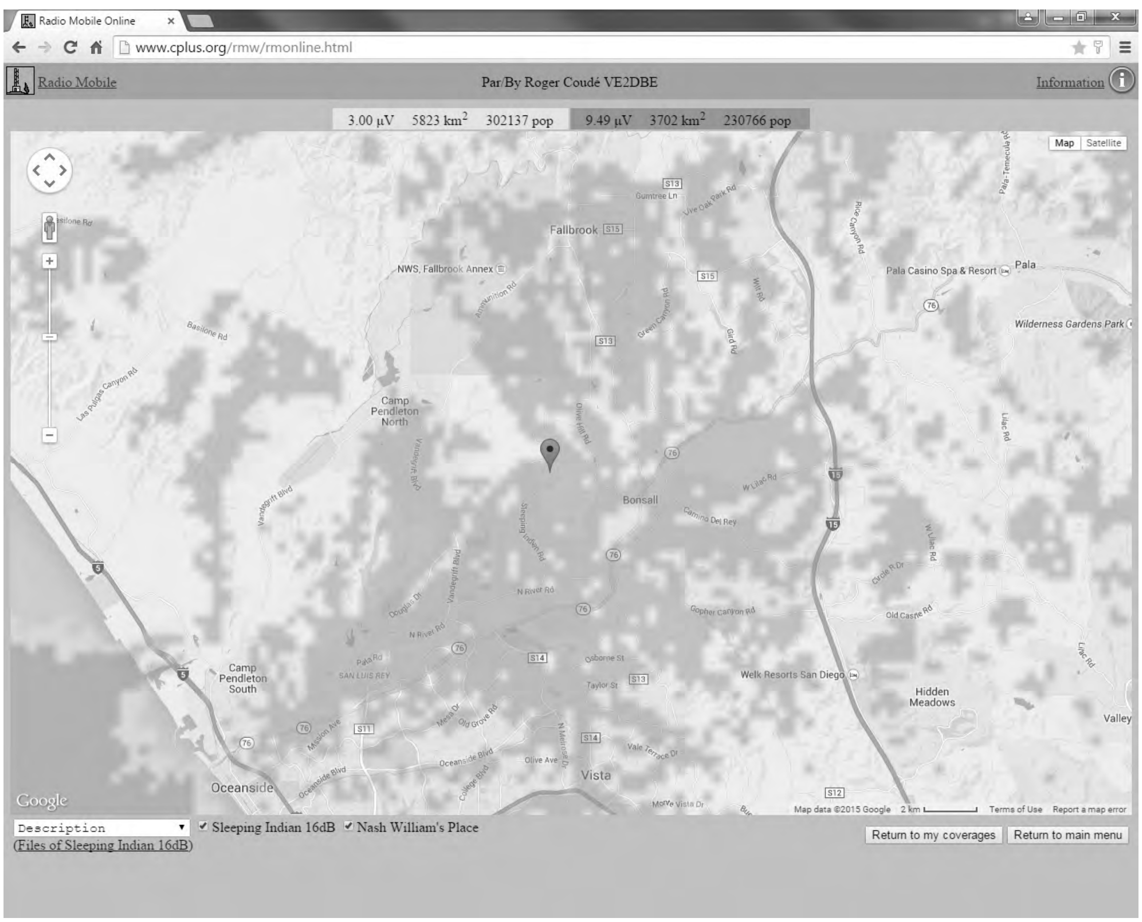

In my planning for a mesh in northern San Diego County, I secured a location in the Bonsall area. The marker locates the potential node(s) under review… in this case a pair of 2.4 GHz Ubiquiti Rocket nodes with 120° sector antenna, one pointed in the general southwest direction and other to the northeast. Based on the specific parameters I entered for these nodes the green-shaded areas will hear the node under review at approximately 9.5µV (-87.5dBm, 7.5dB above the receive sensitivity) and the yellow area at about 3µV (-97dBm, 2dB below the limit of the receiver sensitivity).

I was interested in connecting this cluster of nodes to a high-ground “backbone” node recently placed near Mt Palomar, a 5000’+ peak directly to the east.

The analysis proves out the viability of the prospective link with a more than ample fade margin of 26dB. The author normally considers a margin of 15dB or more as adequate in most instances. Whether this is achievable is highly dependent on the RF environment the node is being placed. Under ideal circumstances you want to take advantage of the full receiver sensitivity (approx. -95dBm for most Ubiquiti devices). But this may not be possible if there is in-band competing activity from other nodes, or other RF sources. For any given situation, it is the signal-to-noise ratio (SNR) that is important, not the signal strength above the receiver’s sensitivity. SNR is also expressed in dB. For example, if the noise floor at a given location is -85dB, then you would need a received signal strength of -70dBm in order to achieve an SNR of 15dB. Remember that Radio Mobile, or any other predictive software is not knowledgeable of the noise environment you will encounter. So uncertainty will remain until you visit the prospective site and confirm this variable.

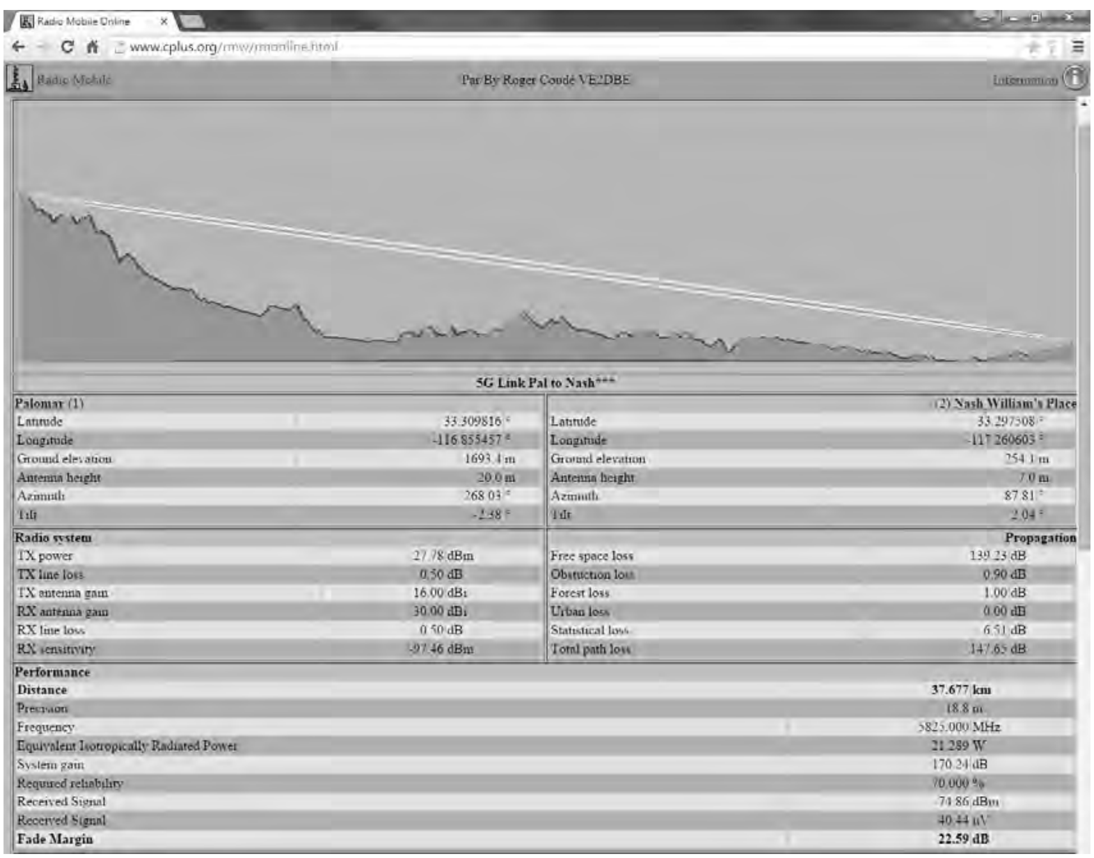

Power Density

As hams we generally don't concern ourselves with too much with power density. Common modulation techniques such as CW, SSB, AM, FM are relatively narrow, but that's not the case with 802.11b which uses direct sequence spread spectrum (complementary coded keying - CCK) or 802.11a/g which utilize orthogonal frequency division multiplexing (OFDM). The AREDN software allows user-selectable bandwidths of 5, 10, and 20 MHz.

So why would a user want to reduce the bandwidth… and the corresponding maximum throughput of the link? The answer is, to increase the signal-to-noise ratio. Each halving of the bandwidth will improve the SNR by 3 dB. Therefore, going from 20MHz to 5MHz has a 6dB improvement. This can easily represent the difference between a reliable link and problematic one.

Fresnel Zones

One more technical concept you need to be cognizant of is Fresnel (pronounced "fren-l") Zones. They are imaginary oblong lines which emanate from one node to another resembling an elongated cigar. The primary cigar represents an area within which interfering objects will cause destructive signal reflections at the receiving end. This is also referred to as "multi-pathing." The formula for computing this area is closely approximated by the following simplified formula:

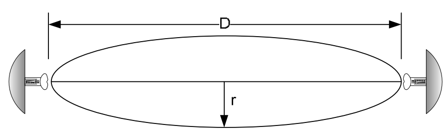

$$r\ (in\ feet) = 36\sqrt{\frac{D\ (in\ miles)}{4f\ (in\ GHz)}}$$

This zone should be clear of all obstructions, including vegetation. If your modeling exercise is not playing out in the real world, the likely culprit is an obstruction in the Fresnel Zone. Notably, the 900 MHz band is tolerant of some vegetation in this zone.

For those less mathematically adept, here are some examples:

900 MHz at 10 miles = 60'	900 MHz at 20 miles = 85'
2.4 GHz at 10 miles = 37'	2.4 GHz at 20 miles = 52'
3.4 GHz at 10 miles = 31'	3.4 GHz at 20 miles = 44'
5.8 GHz at 10 miles = 24'	5.8 GHz at 20 miles = 33'

Conceptual AREDN Network Implementation

Below is a conceptual layout of an AREDN network's core nodes:

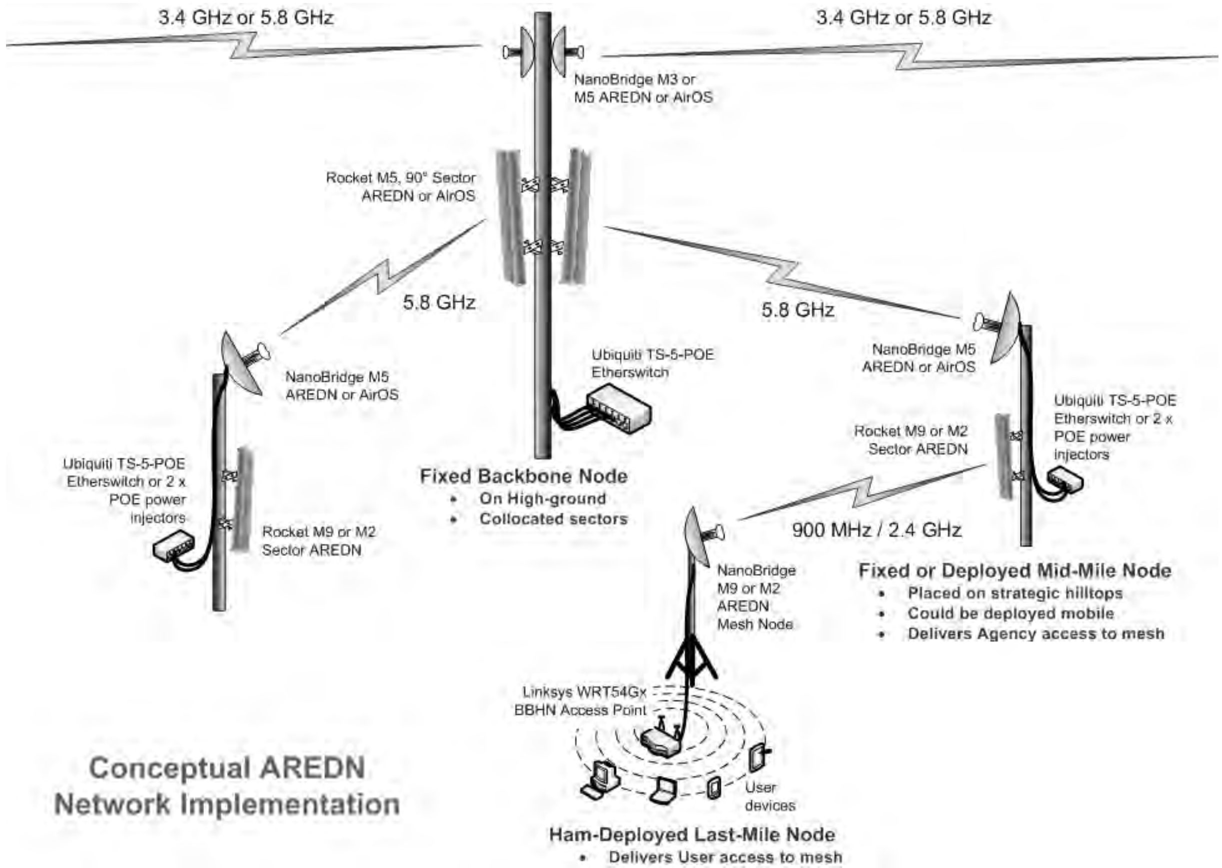

While the term "mesh networking" implies a peer relationship between nodes, a structured core of nodes is defined to establish an initial, predictable path for the network. Three node types are illustrated.

- Ham-Deployed Last Mile Node: As the name indicates, this node is carried by the Ham deployed to a served-agency location requiring the pre-established data services. The "Go-Kit" is comprised of a 2.4 GHz node, or for longer, outlying areas a 900 MHz node with a high-gain antenna pointed at a Mid-Mile Node. The kit also contains a WIFI router to provide network access for the local devices on site. Batteries or a generator powers these devices.

- Mid-Mile Node: This node is either fixed or vehicle deployed as necessary and per the specific network's design. It forms a "reachable" collection point for surrounding Last Mile Nodes and has the required higher-gain antenna to reach a higher-ground based Backbone Node. Sector antenna(s) allow broad downstream accessibility.

- Fixed Backbone Node: These nodes are permanent installations that extend the mesh to the extreme corners of the planned coverage area. For reliability they operate on the least congested band and are optimized for data throughput. Again here, sector antenna(s) are utilized to maximize their downstream accessibility to Mid-Mile Nodes.

Deployment Challenges

One of the most challenging aspects of a mesh implementation is inter-connecting “mesh islands” that have formed in the more-easily meshed areas. Without a competed mesh network, it is difficult to justify the expense/effort of building out network services (email, Voice-over-IP (VoIP) telephony, web-based utilities, etc.) which are needed to demonstrate the network to prospective EMCOMM clients. It may also be difficult to justify the expense of acquiring strategic high-ground properties necessary to connect the mesh-islands.

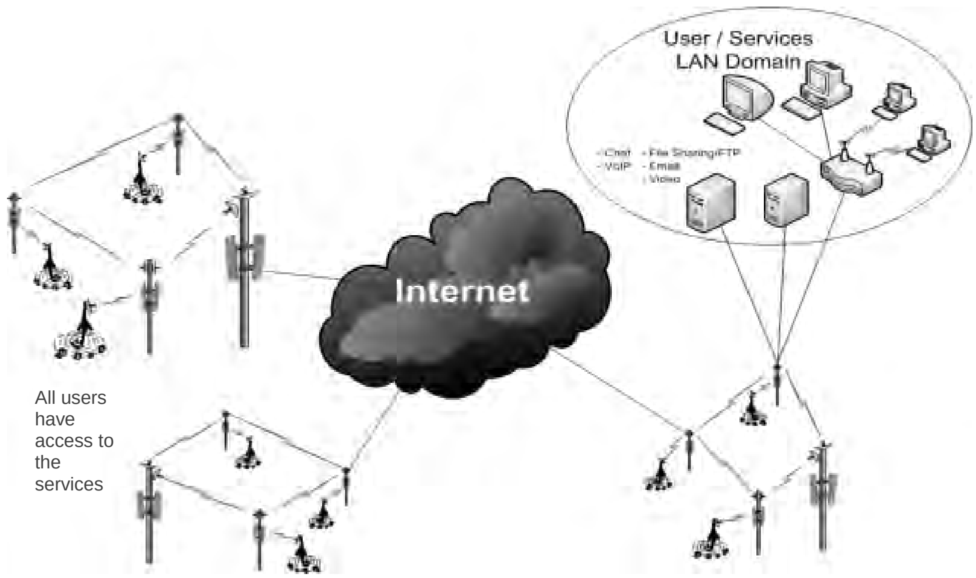

The interim solution AREDN has provided is based on Internet tunneling. This involves setting up an encrypted tunnel between one tunnel server-node and one node in each of the other mesh-islands. This has the effect of connecting all participating mesh-islands together in the same network. In doing so, you can gain the benefits of having completed the network and, at the same time justify the build-out of IP-based services for the users and demonstrate the utility to prospective customers.

While tunneling is an effective way to gain that critical mass, it is a poor strategy for EMCOMM deployment and should only be used as a temporary means of achieving a specific goal. Tunnels will likely not be functional in a real disaster.

A Roadmap for the Near Term

The AREDN Development Team recognizes that the following important improvements are warranted and is commitment to the advancement of this technology.

A More Advanced GUI – More advanced users have taken to other network devices in deploying more sophisticated network topologies. We are committed to providing an optional, more advanced GUI to allow for the configuration of these progressive designs. The effort will preserve the auto-configuration features for the non-network savvy Hams.

Network Management with SNMP – Utilizing the Simple Network Management Protocol, the AREDN network should be manageable using any number of standard tools in use today. Customized MIBs (Management Information Bases) will be developed specific to this technology.

These tools will allow visualization of the mesh against a map background, see the network throughput at a node-interface level, locate choke points within the network that would benefit from improved RF quality and bandwidth, and be alerted to segment outages.

Quality of Service (QoS) – Critical traffic needs priority. EMCOMM data during a disaster must not be hampered by casual traffic. Therefore, some form of Quality of Service is warranted. We are not certain how this would best be implemented, but will likely require some advanced configuration or certification at the user-level.

Operations outside of the ISM bands – The available Ham band is broader than the Ubiquiti and TP Link devices support in their manufactured form. There would be great value in moving the device off the ISM-portion of the band in terms of reduced interference and higher resultant SNR. We are exploring this and believe this will be possible.

Preserving the 3.4 GHz band – Having recently released software supporting devices in the 3.4 GHz band, the team is just coming to realize the huge asset this band represents. This band has no commercial allocation in the US and is considered to be in jeopardy of loss to commercial interests. With the exception of some military radar, there is little interference, making it the quietest of alternatives. A Ham presence, particularly in the EMCOMM space, could prove useful in retaining it for Ham purposes.

AREDN Team Developed Deployments – The number of network deployments utilizing this technology are too numerous to list, however it is fair to say that groups in most major US cities are, at a minimum, exploring this technology as well as major cities in Canada and Europe.

Where 2 years ago I was begging for speaking opportunities, they now approach me. I am certain we are past the tipping point.

How One Gets Involved – Most local deployments start at a grass-roots level. It is an excellent way to span the gap between radio and computer technologies… and, if approached correctly, can attract a new generation to Ham Radio. Larger implementations deserve a network specialist, so I encourage you to find one early. Hams tend to become overwhelmed with the networking elements, so having someone to offload that worry lets Hams worry more about the radio aspects of this technology than the data.

There is a fabulous getting-started primer entitled "Wireless Networking in the Developing World" which I encourage anyone interested in this technology to read. It is available for free as a PDF download at: www.wndw.net.

Conclusion

There are a variety of mesh network systems today. AREDN is unique in that it operates under Part 97 under the authorizations inherent in our Amateur license grant. It is easy to configure and is deployable by typical hams to served agencies without any knowledge of data networking or the design of the mesh to which a node is being connected. It can be used to provide a variety of IP-based services or to restore failed intranet-based agency services.

The AREDN Project team provides support via its website at www.aredn.org to Emcomm groups wishing to deploy this technology.
