Internet Protocol and Radio Frequency Networks: Creating Robust Military Networks

Executive Summary

Military organizations need reliable, versatile tactical networks for data, voice, and video communications. However, that goal presents considerable obstacles:

- New software-defined radios allow a single radio to operate with multiple waveforms to provide a wide range of capabilities, depending on frequency, waveform characteristics, and bandwidth, which makes it difficult to establish and maintain IP connectivity.
- Radio-based communications can be unreliable, and if routers are not aware of the current condition of each radio, they cannot make effective routing decisions.
- Different radios used in military networks use dissimilar connection methods, making it time-consuming to create a network and complex to add new radios to it.
- Radios and routers must be able to form ad hoc networks with minimal configuration or changes.

BENEFITS OF RADIO AWARE ROUTING

- Provides improved communications, coordination, and overall command and control capabilities
- Enables network-based applications and information to be delivered reliably and quickly over radio links
- Delivers faster convergence and optimal route selection to help ensure that sensitive traffic such as voice and video is not disrupted

To address these challenges, Cisco has introduced the concept of Radio Aware Routing (RAR). Radio Aware Routing optimizes IP routing over diverse radio networks to give users real-time access to critical information while on the move. It provides a standardized way of connecting routers and radios so that using multiple types of radios in a network can potentially be as simple and effective as using off-the-shelf components for wired networks.

By abstracting the physical and logical interface of a radio into a form applicable for virtually any type of radio, RAR offers a modular building-block approach for using IP routers over radio networks. Such an approach overcomes many of the underlying problems encountered when using radio-based technologies with IP.

This paper describes certain underlying aspects of the Internet's infrastructure, and how the linkage of many diverse networks evolved into what today we know as the Internet. It also discusses advances in mobile networking technologies that build on those fundamentals to make tactical military networks more modular, flexible, and robust.

Internet Fundamentals

Military data networks are a specialized application of today's enterprise networks. To understand the best way to create a radio-based military network, it is important to review the basics of the Internet and how it was built from many different individual networks. However, there are also important differences between enterprise networks and tactical military networks, which will be discussed later in this paper.

Though it did not come into prominence until the 1990s, the Internet has actually existed since the 1970s, arising from DARPA's packet-switched ARPANET military network. From the beginning, these networks were based on a few simple yet powerful concepts that are still imperative today.

IP Addresses

The most fundamental concept within the IP suite is the simple fact that everything that connects to the Internet gets an IP address. It is a simple idea, but one whose power cannot be overstated. Today, the Internet has become so ubiquitous that virtually any device—from cell phones and PDAs to high-tech refrigerators—can be assigned an IP address and then connect to it.

The concept of endpoint addresses did not originate with the Internet. Such addresses, using a variety of formats, are found in many of the underlying networks that form the Internet, such as SONET, ATM, Ethernet, X.25, and so on. Each of those technologies was designed to address a specific communications need, so they vary in speed, physical connection methods, and how two endpoints communicate with each other. But the concept of endpoints (computers or other devices) having addresses is consistent across these technologies.

Because the address format used by any of these network technologies is unique to that specific technology, complications arose when you wanted to connect multiple types of networks. For example, how could a computer using an Ethernet address share information with one using an ATM address? The disparate addressing schemes of those standalone network technologies therefore limited their communications ability.

This was the underlying problem that the Internet Protocol was designed to solve. It provides a universal address format, along with a method for allocating addresses to all endpoints on all networks that are connected together using this protocol. The overlay of IP addresses allows these diverse networks to participate in the global Internet.

The Internet Protocol ensures that all addresses on a given network are related—an essential aspect of how the Internet operates. For example, imagine two network devices with IP addresses of 10.1.249.10 and 10.1.249.13. Of the four octets of their 32-bit addresses, the first three are identical. From that, we can tell that these two devices are on the same subnet.

This structure of IP addresses is a critical point, because it tells us to which individual network—among the many hundreds of thousands that form the Internet—a certain device is connected, simply by looking at the beginning of its IP address. This architecture is analogous to getting an idea of where a telephone number is located by reading its area code.

Internet Routers

So far, we have seen how the Internet Protocol established the ability to form large internetworks from smaller diverse networks by assigning a unique address to any device on any of the component networks. This mechanism formed an overlay of IP addresses, and IP became the protocol that connected these diverse networks. But how do devices actually communicate across an internetwork?

The fundamental building block of the "network of networks" now known as the Internet is a device called an IP router. Essentially, a router connects two or more individual networks so that they can exchange data packets using the Internet Protocol. Multiply this basic architecture hundreds of thousands of times, and you get the Internet as it exists today. Routers are devices that support many different network protocols. They can therefore be used to connect different types of networks, using IP and a common language.

The basic concept of IP routing is depicted in Figure 1. Computers or other endpoint devices connected to one of the individual networks communicate with other endpoints by transmitting and receiving packets of data formatted according to the Internet Protocol. Each packet of data contains the IP addresses of both its source and destination, so the packet can find its way independently through the connected networks to reach its destination. This is a tremendously powerful concept with great benefits for radio-based networks, with their unpredictable connectivity.

Figure 1. Internet Structure

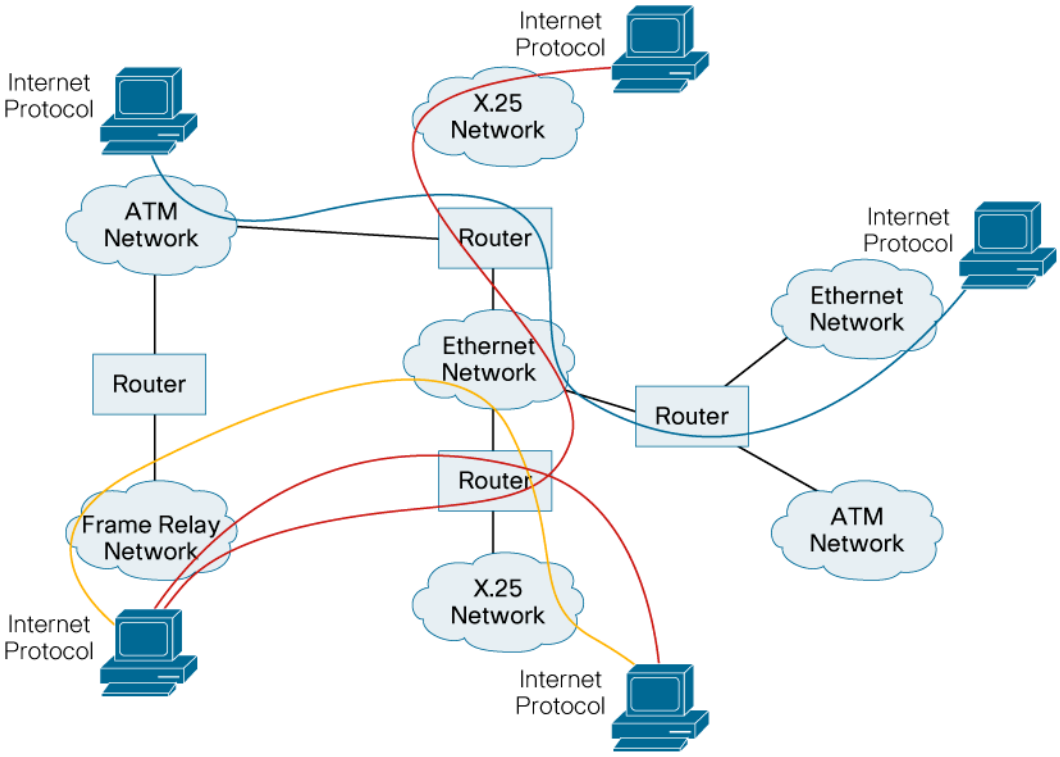

Of course, **how** a packet finds its way is not a simple matter. Most aspects of its journey are relatively straightforward, but the number of different factors that must work properly, and the sheer number of devices involved, make the total picture more complicated. Suffice it to say that the Internet operates because routers know how to route IP packets from their source to their destination across all of the intervening individual networks. The key technology and hence the heart of the Internet is IP routing.

At this point, it is helpful to review what happens in an endpoint that is attempting to communicate using the Internet Protocol. First, the computer (or other device) has to establish its own IP address. As mentioned earlier, an IP address is a 32-bit binary value; it is typically written in “dotted decimal” notation (such as 172.1.1.1) for easier understanding by humans. Remember that one benefit of the IP addressing scheme is the ability to determine on which specific network (among the hundreds of thousands that exist) any given IP address is located. This is accomplished by giving all the computers connected to a single physical network the same IP address prefix. It becomes the function of the routers to understand where all the prefixes are located within the structure of the internetwork.

Figure 2 shows the simple example of a single physical Ethernet network with five computers and a router. These days, even many homes have similar networks, with a couple of home computers, a network-attached printer, and a cable modem that acts as a router to the rest of the Internet. In our

example, everything connected to this small network has a prefix of 171.1.1.x, with the fourth value in the address being the unique identifier for each endpoint device.

Figure 2. A Simple Individual IP Network

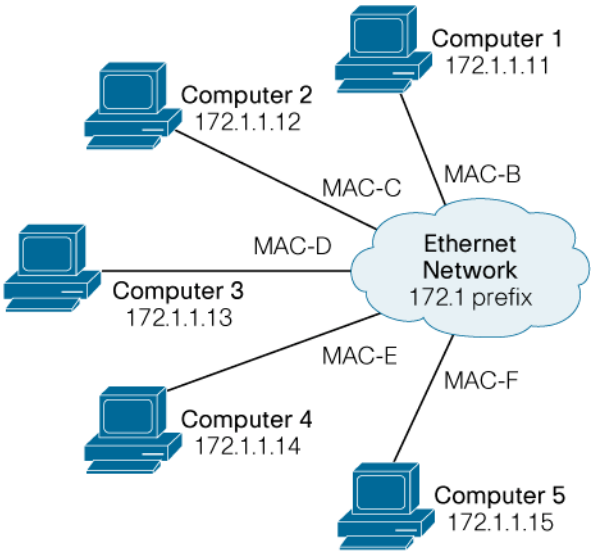

In order for everything to work properly, there are three pieces of information that absolutely **must** be configured into each computer connected to this network: its own IP address, a number known as a network mask, and the IP address of the "default router" that connects this network to other networks.

Every IP packet has a source address and a destination address that allows each packet of information to be routed individually from the source to the destination. Recall that for any given type of physical network (such as Ethernet, ATM, or X.25), there is some form of address that is usually very different from an IP address. Figure 2 shows a small Ethernet as an example, so every attached computer has a 48-bit Ethernet MAC address. MAC addresses are quite lengthy, so for simplicity, they are shown in Figure 2 as simply MAC-A through MAC-F.

The local addressing mechanism is an attribute of each type of network, so the form and specifics are different across different network technologies. Therefore, there needs to be a way of correlating a computer's IP address to its local network address. Each type of network has a different way of doing this. In the case of Ethernet, each computer keeps an internal table that lists both the MAC address and the associated IP address. In the simple example shown in Figure 2, Computer 1 would have a table something like this:

Table 1.

Computer	IP Addres	Local Network Address
Computer 2	171.1.1.12	MAC-C
Computer 3	171.1.1.13	MAC-D
Computer 4	171.1.1.14	MAC-E
Computer 5	171.1.1.15	MAC-F
Router	171.1.1.1	MAC-A

With that information, Computer 1 can communicate directly with any other computer on its local Ethernet network. However, it must communicate with computers on other networks only by sending packets through the router.

This is where the network mask and default router configurations come into play. The network mask is a number that helps network devices determine whether a destination IP address is on the local network, or whether it is only reachable via a router. If the latter, then the device must know the address of its default router. With these pieces of information, network devices can properly send their traffic either directly to other local devices, or to the appropriate router for remote delivery.

Network Layers

The concepts of mapping addresses (IP to local network) and routing between networks are as fundamental to running the Internet as electricity is to running the appliances in your home. Yet mapping and routing are invisible when you use the Internet, just as the flow of electricity is invisible when you plug in your alarm clock. Nonetheless, those underlying details have vital implications for how devices interconnect to construct an internetwork, which applications can run on those devices, and how well those applications perform.

The interface between a physical network device (such as a computer) and a network consists of special software protocols. The protocols that carry the actual data are implemented on top of the physical interface protocols required for the specific type of network being used. A formal breakdown of these virtual layers can be seen in a protocol reference model, such as the OSI Reference Model shown in Figure 3.

Figure 3. The OSI Reference Model

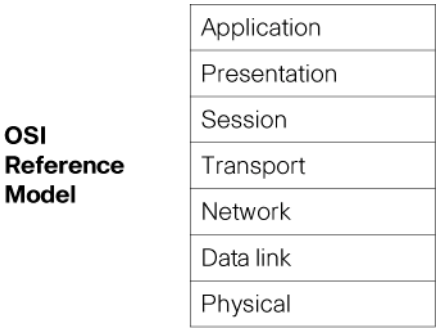

The ARPANET Reference Model (ARM), sometimes called the DoD Reference Model, is the forerunner of the OSI Model—and the protocol reference model that bears directly on the formation of the Internet. ARM is described in Figure 4.

Figure 4. The ARPANET Reference Model

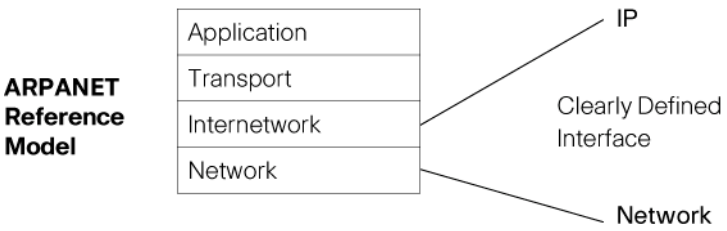

The OSI Reference Model did not contain a separate layer for internetworking, which led to the confusion of using the term “network” for both standalone and interconnected networks. When looking at the ARM, “network” and “internetwork” are clearly two separate layers, and therefore there must be a defined relationship between them. The internetwork layer was, in fact, meant to be the Internet Protocol. So, for example, an Internet-connected Ethernet network would map IP (internetwork) addresses to MAC (network) addresses.

There are other details of any underlying network technology that can affect the way things need to happen at the IP layer. For instance, some network technologies (including Ethernet) can send a single packet and have every connected computer receive it; this is called “broadcast” or “multicast” networking. Other networks, such as those using X.25 protocols, might require a packet to be sent individually to every attached computer. These underlying attributes for any given network are important when it comes to configuring routers to talk with each other—and keeping routers working is essential to keeping the Internet running. In radio-based networks, the underlying characteristics and capabilities of the radio network can have a tremendous impact on how the Internet Protocol should operate.

Radio-Based IP Networks

What do those essential aspects of the Internet—diverse underlying networks, unique IP addresses, routers, and network layers—have to do with tactical military networks?

In the ARPANET Reference Model, the details of the relationship between IP and the underlying network are defined in what are termed “Requests for Comment.” RFCs are essentially the designated standards for how things must operate and behave when using the Internet protocols. Today, these Internet standards are so well defined that you can purchase practically any commercial, off-the-shelf router and use it to connect almost any combination of networks.

All of the different implementations of IP, as well as many of the underlying routing protocols, have been vigorously tested throughout the decades of the Internet’s existence. We now take it for granted that routers and other devices can be purchased from a number of companies, connect to the same type of network the same way, and exchange information correctly. This compatibility was certainly not the case in the early days of the Internet build-out, and it is still a challenge in creating radio-based networks today.

The protocols used by Internet routers to exchange information about the connections between networks are also well defined. Routers understand the specific attributes of each network, such as its bandwidth, its delay times, and whether it is multicast/broadcast capable. They take that information into consideration when deciding how to route each packet to its destination, because being able to accurately and efficiently get packets to their destinations is what has made the Internet a staple of everyday life.

Now, we want to extend that proven architecture to a more challenging environment: tactical military networks. That means we must deal with the vagaries of wireless (radio-based) networks. Range, signal strength, type of antenna, and other attributes of radio systems can vary widely. To make the situation even more difficult, military radios are constantly changing their location. They may go in and out of a network because they are in a vehicle that goes behind a building, or in an aircraft that changes heading. All of these factors add complexity at the internetwork layer, because the routers must be able to account for the problems introduced by the underlying radio networks.

Even so, all of the previous concepts still apply. A network can be built around a common set of radios using the same waveform. There can be many simultaneous radio networks of different

types that must be interconnected. The Internet Protocol is still the unifying force being used to build an internetwork from the underlying radio networks. In effect, tactical networks form an Internet built over radio networks in much the same way the original Internetwork was built over wired network technologies. The key differences are in the underlying operation of the radio networks, which must be taken in account within IP routing functions.

MANETs

One of the most common terms in the realm of radio-based and mobile networking is “mobile ad hoc network,” or MANET. Recall that the term “network” can refer to a single instance of a network, or to a broader internetwork. In mobile radio networking, it has both meanings.

A single MANET consists of radios that can exchange data over a certain geographic area. They could be relatively short-range radios, or BLOS (beyond line-of-sight) radios with ranges of hundreds of miles. They could even use satellites to relay their signals. The details of each type of radio network may vary, but many of the underlying concepts are the same.

The underlying radio network must support an ad hoc capability in which radios can change in relation to each other, and can enter or leave the radio network. This is, in fact, a form of MANET that is operating at the radio network layer. This is a separate function from the idea of ad hoc IP routing, which also must occur. In fact, once a given radio comes into a radio network and stabilizes, the router using the radio must become part of the interconnected set of IP routers. Both functions will have to operate in an ad hoc manner.

The focus of Radio Aware Routing is to provide a mechanism for routers to connect and exchange information over radio paths and be able to obtain information about the radio links. Such detailed information about the links between radios will allow the IP routing to select the optimal paths, and enable more rapid convergence of IP routing due to any changes in the underlying radio network.

Radio Aware Routing is a broad concept that includes a number of areas. Foremost, is the development of the capability for routers to exchange crucial information with radios. This includes having standards for both the physical interface between a router and a radio, as well as the information to exchange and how it should be exchanged. It also includes a capability for a router to understand at least some minimal information about the underlying radio links even if a particular type of radio doesn’t cooperate in the exchange of RAR information. From a modularity standpoint, it is beneficial to have as few physical interface types as possible. It is also desirable to have a set of information-exchange protocols that enable a router to discover information about the underlying radio network. By having such standardized mechanisms in place, it becomes possible to take almost any type of radio, connect it to a router, and have the router optimize the relationship between the radio network and the IP routing layer.

The addressing schemes used in different radio networks will vary, just as they do across different wired networking technologies such as X.25, ATM, and Ethernet. For example, a radio with a broadcast (omnidirectional) antenna may simply use an existing addressing method, such as Ethernet MAC addresses. Other types of radio networks may use something as simple as an 8-bit identifier for each radio terminal. One can also layer addressing schemes on top of each other, which already happens with most networks that transport IP traffic.

An example will help illustrate these concepts. If each radio gets a unique 8-bit address, you could have up to 255 radios on a single network. Let’s assume that the radios are omnidirectional and can therefore broadcast information to all other radios in range. A simple representation of this notional radio network is shown in Figure 5.

Figure 5. A Simple Radio Network

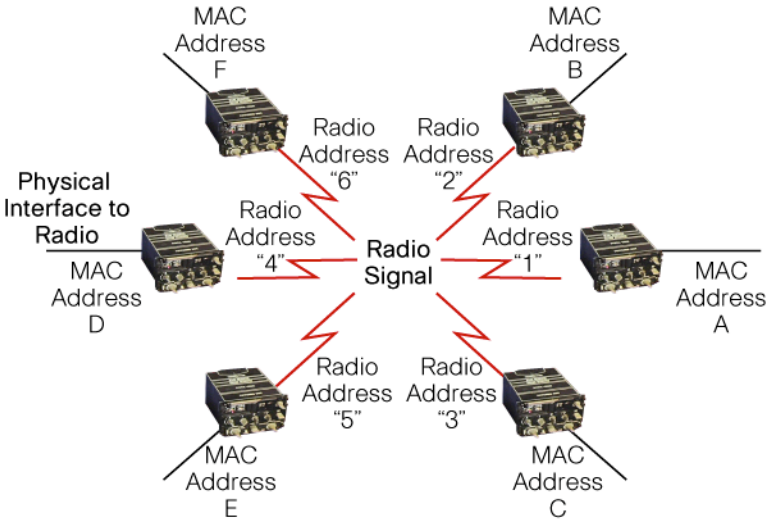

In this example, every radio has an individual address, numbered 1 through 6. The diagram also shows that there must be some type of physical interface to the computers (or other endpoint devices) that will use the radios to exchange data. If a new radio comes into range of this group, it must be able to join the network and begin exchanging data with the other radios. Having this occur without manual intervention or reconfiguration of the radios is a form of MANET operating at the network layer of the ARPANET Reference Model.

In this example, the physical interface is Ethernet, and as we know, every Ethernet device must have a 48-bit MAC address. In essence, this example makes the radio network look like a single Ethernet segment. The radios would keep track of the MAC addresses that are reachable through each 8-bit radio address.

To add the Internet Protocol to this example, let's connect a router to each radio, as shown in Figure 6.

Figure 6. Radio Network with Internet Routers

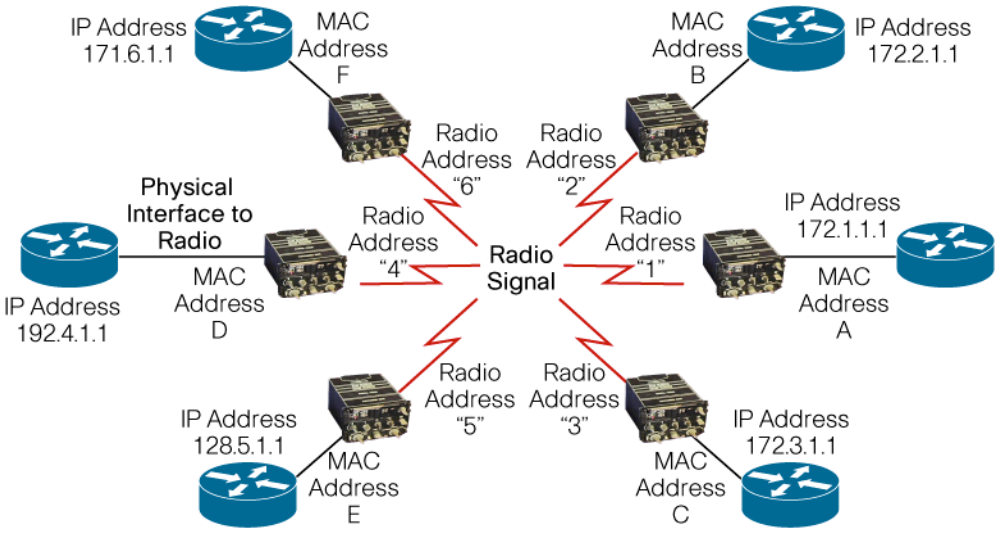

Radio Aware Routing

Now that we have a router behind each radio, we can expand this network to a MANET that uses the internetwork layer of the ARPANET Reference Model. In the example in Figure 6, only routers are connected to the radio network; individual computer systems are on other networks behind the routers. All of the end systems connecting through the routers operate just as they would over any IP network; they are not aware that their communications are being carried over radios. However, the routers need to take into consideration the aspects of radio networks that are different from wired network technologies. In other words, we need to implement Radio Aware Routing.

We mentioned earlier that all IP devices on a single network typically have the same IP address prefix. That organization allows any device on the internetwork to determine whether it is directly connected to a specific destination device, or whether it must use a router to reach that destination. However, the IP addresses assigned to the routers in Figure 6 do not have a common prefix. This is an attribute necessary for having ad hoc networking capability at the internetwork layer ¬– the ability for radio-connected devices to join and leave networks without manual intervention.

Therefore, routers have to operate differently when they are on a radio network than when they are on a traditional network. The specifics of the underlying network will determine exactly what accommodations the routers must make, but there will also be commonalities among radio networks that allow for some degree of standardization.

The diagram in Figure 6 shows several layers of address mapping that must occur, as well as two layers of MANET. Each radio must recognize other radios that have MAC addresses. When a new radio enters the network or leaves the network (perhaps because its antenna has become obscured), it is beneficial for the radio to notify the router. This allows the IP-layer MANET (or routing protocol in use by the router) to operate on nearly the same timescale as the radio-layer MANET being used in the underlying radio network.

This optimized information exchange along with standardized interfaces is the key functionality provided by RAR. By abstracting the details of the underlying radio networks to a standard set of interfaces, it makes it possible to change the routing and radio functions independently of each other.

There are many other issues to resolve in building radio networks: transiting multiple radios between endpoints, multicast distribution, disruption-tolerant networking paradigms, and long-delay paths, to name a few. But the importance of RAR is in laying the foundation to integrate the two key building blocks: radios and IP routers.

The trend in radio networking is to actually embed the router within the radio. Whether the router is implemented in software or via a card inside the radio, the result is the same as far as external devices are concerned. The details of how the router interacts with the radio network are completely hidden to them. However, difficulties can arise if communications between the router and the radio are proprietary or nonstandard, as we will discuss shortly.

End Systems and Routers

As mentioned earlier, each end device on an IP network requires three pieces of information to function: its own IP address, a network mask, and a default router. In Figure 7, all of the computers connect to each other and to the radio via Ethernet. This physical interface and network type is very well understood, is inexpensive, and supports data transfer speeds of a gigabit per second and beyond, all of which makes Ethernet a popular choice.

Figure 7. Multiple End Devices Sharing a Radio/Router

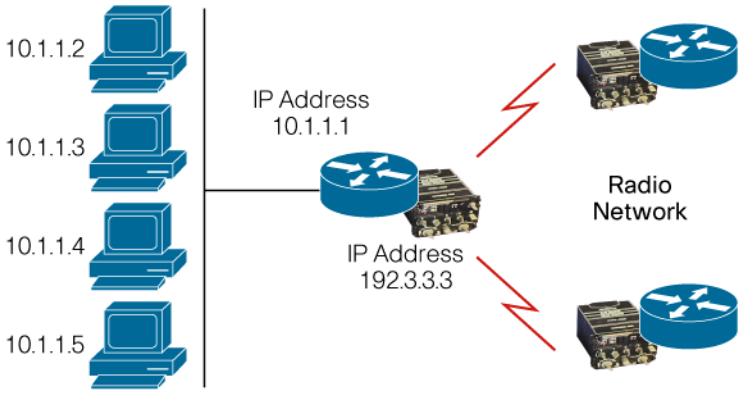

In Figure 7, all of the computers use the router embedded in the radio as their default router. As long as the routers using the radio network implement the necessary ad hoc routing, this example can work just fine. However, in a typical tactical scenario, such a simple topology with a single type of radio will rarely exist. More commonly, a tactical network would require the use of multiple types of radios, as demonstrated in Figure 8.

Figure 8. End Systems with Multiple Radio Types

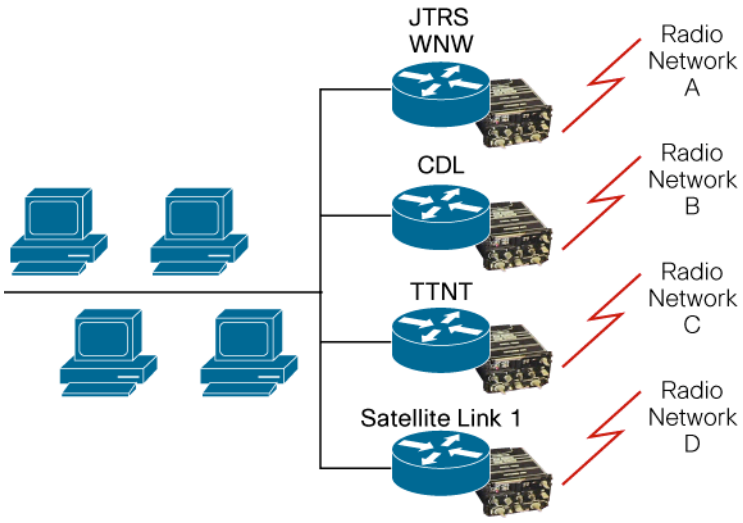

To add to the complexity of this scenario, a vehicle, aircraft, or ship would typically have several different radio paths, all differing in bandwidth, delay, range, and the attributes they present to connected routers. In such a dynamic situation, how do attached devices determine which of the available routers should be their default path to the internetwork? And how do the routers—which may be from different manufacturers and employ different types of both network and internetwork MANETs on their radio networks—communicate with each other across their Ethernet connection?

Both issues must be addressed to create an IP-based tactical internetwork. The most straightforward method currently used is to add another router just for the endpoint devices, as shown in Figure 9.

Figure 9. End Systems with Dedicated Aggregation Router

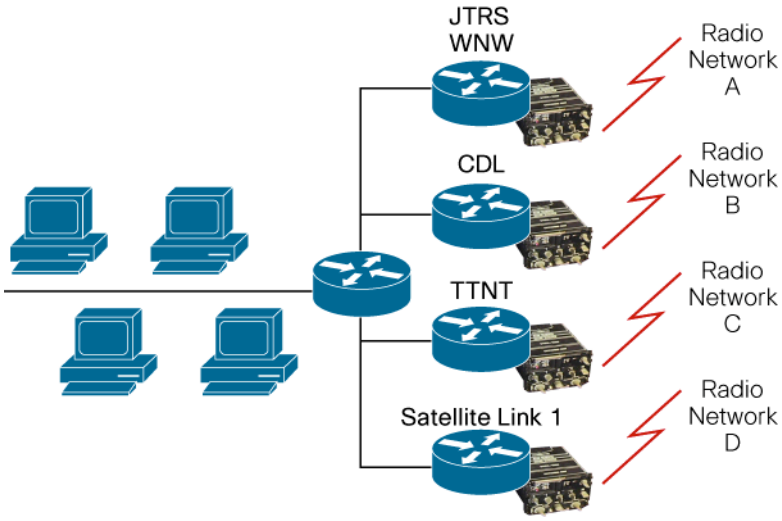

Adding this one router simplifies the configuration of the network attached computers. They now use the aggregation router's IP address as their default router, and send all non-local traffic to it for forwarding. This router can also be configured to exchange routing information with each of the radio routers, creating a single overall view of the tactical network's topology. However, we still need to establish a way for the aggregation router and the radio routers to exchange routing information.

This exchange of routing information is dependent on the capabilities and features in each connected radio and router. A specific configuration is usually required for each device. This design also isolates the aggregation router—which is being used by all of the attached devices—from any direct information about the underlying radio networks. This approach causes delay between the loss of a router in an underlying radio network and the ability of the aggregation router to reroute the IP traffic. Additionally, there will be delays in the aggregation router discovering changes in bandwidth or other radio network attributes.

The Promise of Radio Aware Routing

Whether routers are external to or embedded in radios, they still need some method of interacting with the underlying radio network. With current technologies, the method used depends on the kind of router and radio being used. Suppose, however, that we were to define a standard method for the radios and routers to interact. Information from the radio network (such as bandwidth, delay, and even whether a new radio has joined the network) could then be communicated to the router using a protocol designed for that exact purpose. Due to the very wide range of radios available and their underlying capabilities, there is, in fact, a need for a family or least a range of protocols for such information exchange.

Having this set of standards, the network pictured in Figure 9 would be transformed into the one shown in Figure 10.

Figure 10. Standard Router to Radio Interfaces

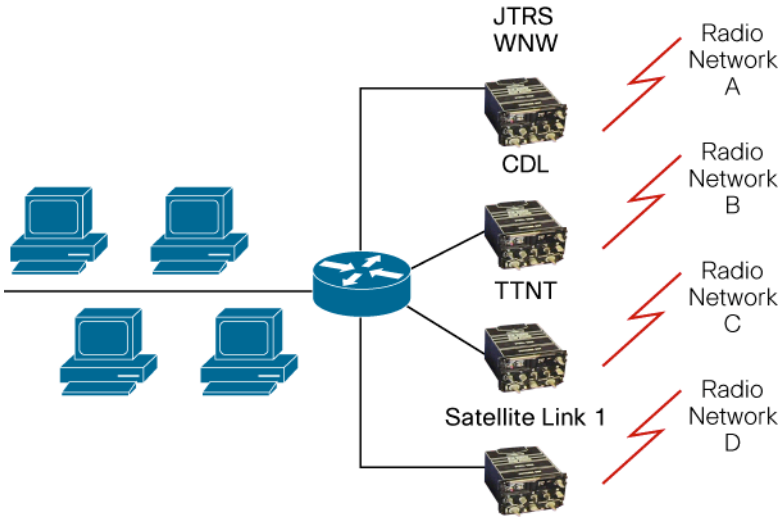

At first glance, this may seem like a simple change, but its effects would be profound. All attached computers would operate just as they do on the Internet today, unaware that some of their traffic is being carried over a radio network. On the radio side, it would become easy to add additional radio networks. As a result, such a system would offer both the robustness **and** versatility that tactical military networks demand.

To make such an integrated system feasible, new types of radios would have to support the selected method of exchanging information between the radio network and the router. Any new radio could then immediately connect to any router, without the need for developing yet more proprietary communications software.

Reaching a point where building tactical internetworks based on military or commercially available data radio systems becomes as straightforward as creating new segments of the Internet will ultimately depend on two factors:

- Standardizing on a single physical interface (or at least, a limited set of interfaces) between routers and radios
- Having a standardized protocol for exchanging information between routers and radios

What sorts of information would need to be included in those standards? Consider Ethernet being chosen as the physical interconnection. Ethernet is prevalent, it is inexpensive, and it can cover almost any range of bandwidth. But if the router connects to a radio via a 100-Mbps Ethernet connection, and the radio is only capable of transmitting at 3 Mbps, then there needs to be a way for the radio to inform the router that the actual bandwidth of the link is only 3 Mbps.

Another example occurs in the case of a radio network where each radio can keep track of every other radio on the network, such as with TDMA networks. When a new radio-based router joins the network, it's vital to have a method of notifying every other router about its availability, so that they can immediately begin exchanging routing information with it. This one simple step of proactive notification can have a tremendous impact on how quickly end devices are able to take advantage of the new router.

Conclusion

Only by choosing which underlying mechanisms to standardize and taking action on that decision can we create a level of connectivity within tactical military internetworks that we take for granted in the public Internet. Standardization is the core of the Radio Aware Routing concept. It will make it possible to build radio-based networks modularly, just as the Internet is built from a standard set of wired devices and protocols. Any approach that does not include such standardization could actually push the complexity of building a tactical network beyond what is attainable.

All of the underlying principles discussed in this paper are being used at this very moment on the Internet. The goal for mobile military networks must be to find the best way to leverage those existing technologies, while adding the necessary radio and satellite wireless capabilities.

To that end, Cisco has published RFC 4938, which sets out an extension to the Point-to-Point over Ethernet (PPPoE) protocol that will improve the performance of PPPoE over media that have variable bandwidth and limited buffering, such as mobile radio links. Essentially, the extension allows a radio to tell a router what it sees in the RF environment around it, allowing the router to make more intelligent decisions about how to keep network traffic flowing quickly and efficiently. In addition, changes to IP routing protocols, such as Open Shortest Path First Version 3 (OSPFv3), are being made to use information about underlying radio networks more effectively, and to operate in ad-hoc fashion, thus forming an IP MANET over the radio networks.

The advent of software-defined radios marks a key milestone in the evolution of radio networks. A single radio hardware implementation can, by virtue of software, function as many different types of radios. It may operate in some situations where there is only the one radio, or in other circumstances where there are many radios being used to build a complex radio-based internetwork. Implementing Radio Aware Routing as part of the software-defined radio enables a building-block approach with a well-defined set of standards for building IP-based networks over radio systems. With RAR, changing the functionality of the software-based radio will automatically modify the functionality of IP routing.

Cisco continues to research networking technologies that will make Radio Aware Routing a reality. Working with vendors who are willing to implement the necessary hardware interfaces and software protocols, we foresee solutions that will benefit not only military users, but any group that needs reliable, versatile mobile networks.

For More Information

Learn more about Radio Aware Routing and other Cisco® solutions for military networks by visiting the following websites:

- Cisco Government Solutions
- Cisco Defense Solutions
- Cisco Mobile Government Solutions
- Radio Aware Routing demonstration video
- The Future of Ad Hoc Mobility

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C11-504998-00 01/09