

Routing for IP/Packet

7-July-06

Routing for IP/Packet

- With IP packet, once configured, “normal” programs can be used, reducing training needs
- Routing is a major advantage of IP
- Routing is the one “technical” issue that may need to be addressed under fire
- More than a few people must be knowledgeable about routing

7-July-2006 IP Routing 2

A huge advantage of IP for packet is that once configured, normal network programs can be used, greatly reducing training. Since configuration is done once, only a few people really need to be highly skilled.

Automatic routing is also a major advantage.

However, when an incident occurs, changes may be needed to routing. It is probably the only configuration item that might need to be adjusted under pressure of an incident.

As a result, nearly everyone must understand routing.

Routing command syntax

- The syntax for establishing routes is different between JNOS and command line, and slightly different between Linux and Windows command line
- Routing depends on how you are connected – with a laptop, or if some infrastructure is disabled, may need to change routing

7-July-2006 IP Routing 3

Although the problem is the same for the various environments, the syntax is different. JNOS is different from operating system routing, and Windows is slightly different than Linux.

Routing does depend on the situation, so changes may be needed from time to time as conditions dictate.

Definitions

- When setting up routes, some definitions are needed:
 - IP address – the assigned address for a network function, like 192.168.1.2
 - MAC address – the serial number associated with a particular piece of hardware, like 00:02:f0:f1:2f:7e
 - Network – a related group of IP addresses
 - Interface – an identifier for a piece of hardware in your computer

7-July-2006 IP Routing 4

A few definitions are needed before we can talk.

An **IP address** is the four numbers separated by dots that you need to participate in a network. Packet IP addresses start with 44 and are assigned through the mi-drg web page. LAN addresses usually start with 192

A **MAC address** is usually essentially a serial number for a piece of hardware. Every Ethernet card ever manufactured was assigned a MAC address by its manufacturer. For packet, our call is used as the MAC address.

For this discussion, a **network** is a group of related IP addresses. For routing purposes, they must be a contiguous block of numbers.

An **interface** is the name of a devices used for accessing a network.

Describing Networks

- Networks are described in two different ways, depending on context:
 - Number of bits:
 - 192.168.1/24 (i.e. 192.168.1.???)
 - 192.168.1.1/24
 - 44/8 (i.e. 44.????.????.???)
 - With a netmask:
 - 192.168.1.0 MASK 255.255.255.0

7-July-2006 IP Routing 5

IP addresses are really 32 bit numbers. They are written as four, eight bit (0-255) numbers to make them easier for humans to deal with.

A network is a group of addresses for which the higher order bits are all the same. The size of the network determines how many bits must be the same.

We can describe a network by listing it's higher order bits and how many bits must be the same. So for example, a typical home LAN has an address space of 256 addresses, and is defined as 192.168.1.0/24. That is to say, the 24 bits that are the same for that network have values 192.168.1.x.

The ampr.org network has only 8 bits the same, so any address beginning with 44 belongs to the amateur radio packet network.

Another way to describe a network is to list essentially two addresses. The first is that same as in the previous example, and the second has all the bits turned in that must be the same for this network.

Note that addresses 8, 16, and 24 bits wide are easy; the leftmost 1, 2, or 3 octets must be the same. It is somewhat less obvious if the number of bits is different. For example, Midland county could be described as 44.102.200/22, which includes 44.102.200.1, but also 44.102.202.13.

Carving up the network

- Some addresses may be only reachable through some interfaces.

7-July-2006 IP Routing 6

If a computer has more than one interface, we need to carve up our view of the network. Somehow, the computer must be told what addresses to find on what interface. An obvious example is a PC connected to our home LAN and also to a TNC. We have to tell the computer to find packet stations on the TNC, but the rest of our home LAN is still attached to the Ethernet interface.

Naming Interfaces

- Windows:
 - Interfaces are named sequentially
 - 0x1, 0x2, 0x3
- Linux
 - Interfaces are named sequentially within a hardware type:
 - ttyS4, eth1, tun0
- JNOS
 - Within JNOS, Internet interfaces take Linux-style names, packet interfaces are assigned names

7-July-2006 IP Routing 7

The various devices connected to the various networks are called interfaces, and they are named somewhat differently depending on the context.

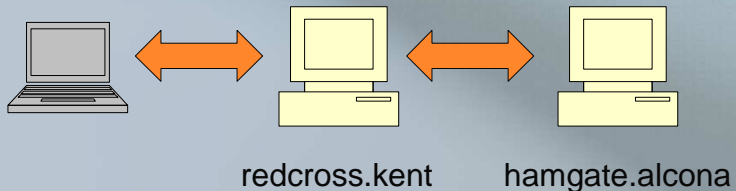
In Windows, interfaces are named sequentially as they are installed. They are creatively named 1, 2, 3, etc.

On Linux, the interface sequence numbers are preceded with a few cryptic characters that describe the type of interface. Thus, serial port 4 would be called ttyS3, a NIC might be described as eth1, and a virtual network device (such as is used for JNOS) as tun0.

Within JNOS, packet devices are assigned names in autoexec.nos. In Michigan, we have a convention that these names represent the frequency. Other devices are assigned Linux-style names.

Gateways

- Sometimes parts of the network are only reachable with some other computer's help:



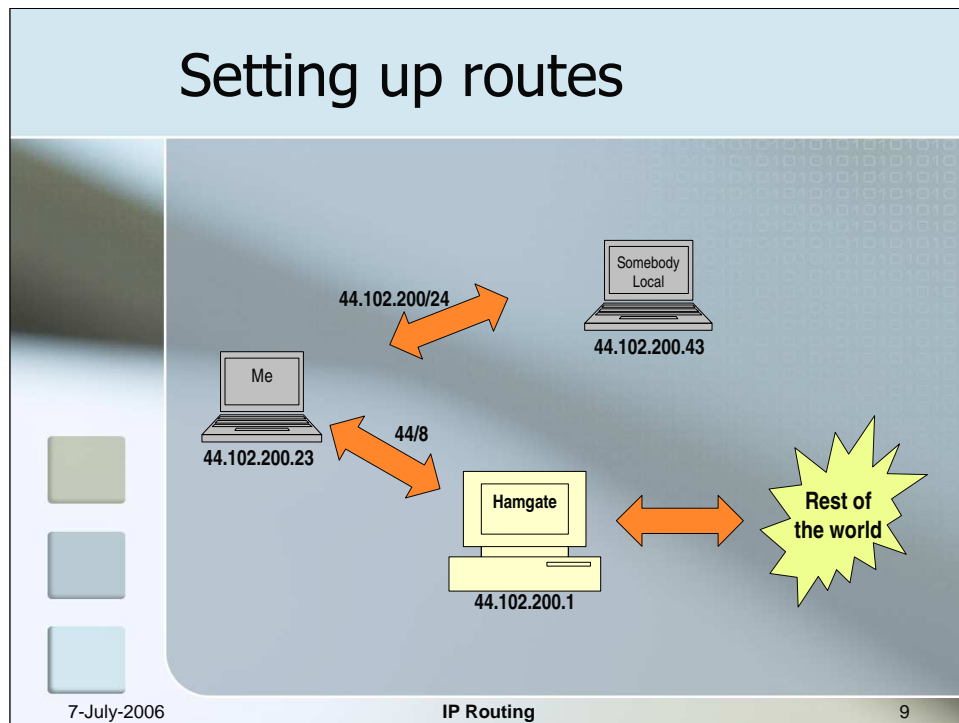
7-July-2006 IP Routing 8

Of course, we rarely have a direct connection to all the other stations we want to communicate with. Computers tasked with connecting us to stations we can't reach directly are called gateways.

Within our routing instructions we need to describe when we need to use these gateways.

For example, if we wanted to communicate from the Red Cross to, say, hamgate.alcona, we would want redcross.kent to perform the gateway function for us.

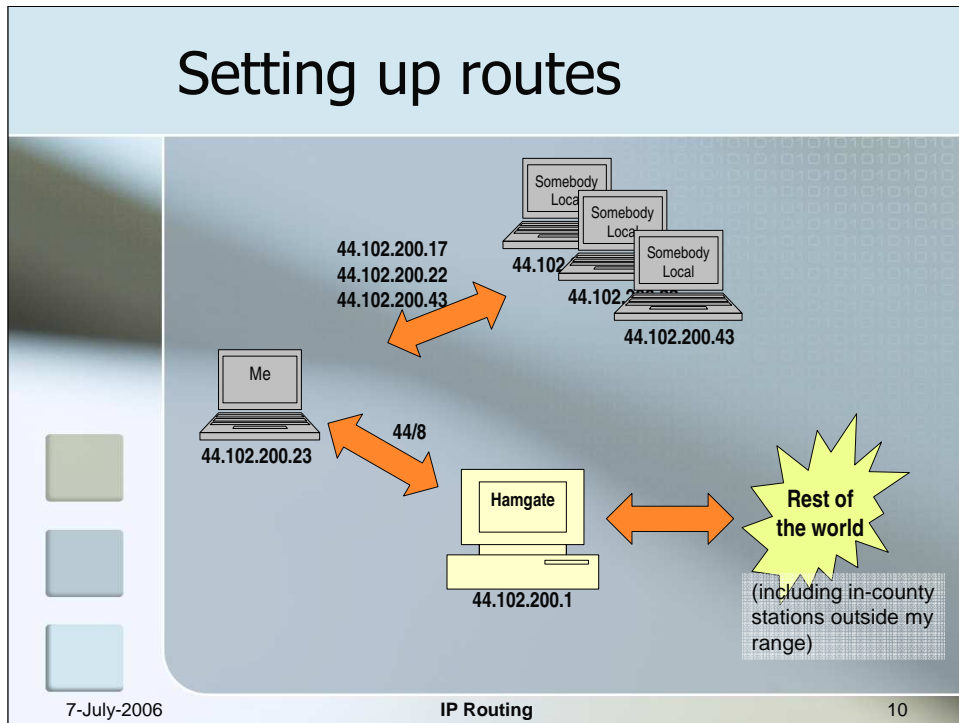
Now, redcross.kent might have several interfaces it is concerned with, but we don't have to care. We only need to describe how to get to the next gateway. Beyond that, it is the gateway's problem.



Let's look at a simple case. Imagine we are in the field with a lone laptop, say running AGW's IP stack. We have enough of an antenna that we can reach all the stations in our county, but beyond that, we need the help of a gateway.

Within ampr.org, each county has one or more 256 address subnets. So in this case, we would set up 2 routes. One that says to use our RF interface for our county (assume for this example, 44.102.200/24), and for any other 44 address, to use the RF interface to get to the nearest HamGate (in this case, 44.102.200.1) and let it worry about anything outside the county.

Notice that the HamGate might use the Internet to get to distant counties, or it may use RF to get to nearby counties. It may even have several RF interfaces, and choose one depending on the particular target. When we are setting up our local routes, however, we don't care. Getting to that other county is the HamGate's problem, not ours.



Perhaps we have a lower-profile station, and we can only reach a few other stations. Instead of setting up a route for the entire county, we might set up routes for a few stations, and allow the HamGate to choose the path to the rest.

Defining routes in Windows

■ Step 1: Identify the interface

```
C:\>ROUTE PRINT
=====
Interface List
0x1 ..... MS TCP Loopback interface
0x2 ...00 01 6c ca 24 28 ..... Broadcom NetXtreme Gigabit Ethernet - Packet Sch
eduler Miniport
0x3 ...00 12 f0 f1 2f 7e ..... Intel(R) PRO/Wireless 2200BG Network Connection
- Packet Scheduler Miniport
0x4 ...ae 84 24 70 34 70 ..... SV2AGW TCPIP Over Radio NDIS Driver - Packet Sch
eduler Miniport
```

ROUTE PRINT command displays interfaces and current routes

Note interface number for interface of interest

7-July-2006 IP Routing 11

Before we define any routes, we need to know the name of our interface. On Windows, the ROUTE PRINT command will list all our current routes, but they will be preceded by a list of interfaces.

In the case show, there is a loopback interface which will always be present, and in this particular case, there happens to be both an Ethernet and a wireless interface. The fourth interface on the list is the AGW IP stack. It's address, surprisingly enough, is 4.

Defining Routes in Windows

■ Step 2: Add route(s)

```

:
: Local ampr.org IP address and Interface
:
SET MYIP=44.102.200.20
SET IF=4
:
: Hamgate - change to 44.102.200.1 when hamgate.midland up
:
SET HAMGATE=44.102.200.17
:
: Routes to Midland County
:
ROUTE -p ADD 44.102.200.0 MASK 255.255.255.0 %MYIP% IF %IF%
ROUTE -p ADD 44.102.201.0 MASK 255.255.255.0 %MYIP% IF %IF%
:
: Default all ampr.org through hamgate
:
ROUTE -p ADD 44.0.0.0 MASK 255.0.0.0 %HAMGATE% IF %IF%

```

Gateway

Note: in Windows, all bits = 0 in mask must also be 0 in address

7-July-2006 IP Routing 12

Once we know the interface, and the IP addresses we care about, we can then add routes with the ROUTE ADD command. The -p switch means to make these routes persistent; that is, don't forget them when we reboot.

Here we have shown a batch file which we might want so that it is easy to make changes without losing track of what we have done.

We can delete routes with the ROUTE DELETE command.

Defining routes in JNOS

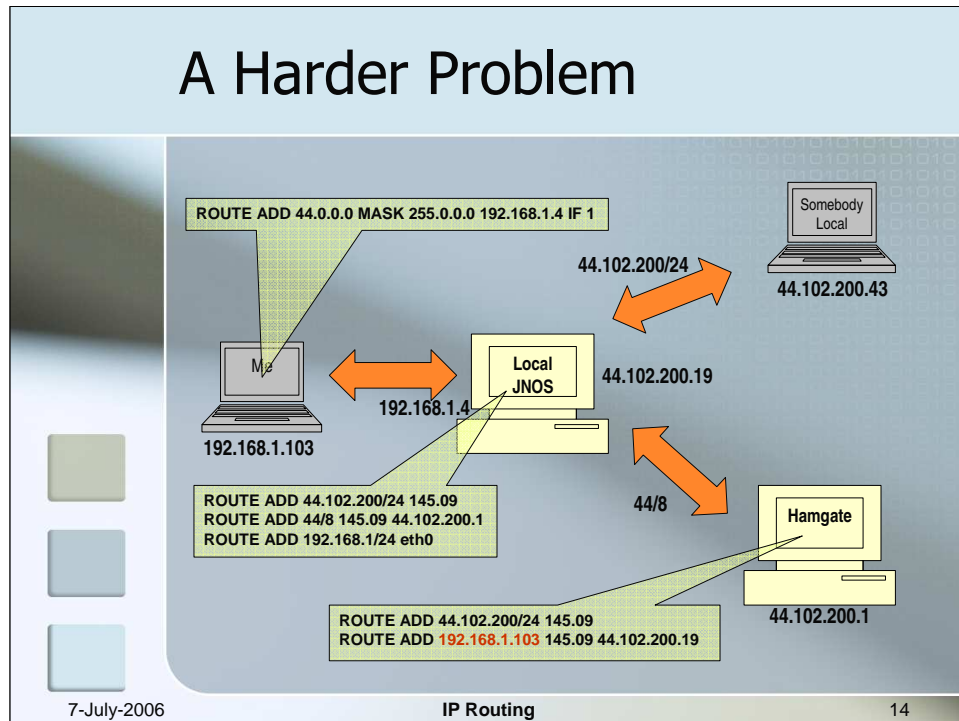
- Use /bits notation, otherwise similar

```
#  
# Midland County subnet  
route add 44.102.200/24 145.09  
route add 44.102.201/24 145.09
```

Michigan convention: RF
interfaces named by frequency

7-July-2006 IP Routing 13

In JNOS, the syntax is similar. However we don't have to use the (clumsy) MASK syntax, and the interface name can be almost anything, although here in Michigan, it is the frequency.



There are messier situations we might encounter. What if we have a JNOS box (or similar router) on our LAN, and we want to use a LAN-connected computer to query the local HamGate for our messages using POP3.

For the client computer, it is pretty simple. We tell it to use the Ethernet card for all 44-dot addresses (since we don't have a TNC), but to use our local JNOS as a gateway.

The local JNOS, however, has two interfaces, and hence, two IP addresses. We need to tell it to communicate with anything on our local LAN (192.168.1/24) over its Ethernet interface, to other stations in the county over the RF interface, and to other ampr.net addresses over the RF interface through the local HamGate.

The HamGate knows to get to the local stations over the RF port. It also has a (long) list of all the other HamGates and how to reach them over the Internet. But an interesting thing happens. When a packet arrives from our client computer, it's FROM address will be from our local LAN. It won't know how to get a response back unless it has a routing entry telling it that *our* local JNOS is a gateway to that computer.

Wrapup

- In general, we can configure our IP/packet system once and forget it
- For routing, however, we may need to adjust
 - Missing infrastructure requires different routes
 - Temporary location requires different routes
- Routing is something almost everyone needs to understand
- We need to consider routes in both directions
- Depending on services we wish to provide, we may need to coordinate LAN addresses

7-July-2006 IP Routing 15

So while we can do most of our configuration once, for routing, we might need to be able to make changes on the fly.

As a consequence, almost everyone should know how to deal with routing.

Routing isn't hard if we limit ourselves to the single problem at hand.

We always need to keep in mind that we need paths both ways.