IP Routing

Routing is clearly one of the most important functions of IP since, otherwise, packets would never actually get anywhere. Simply, it is the process of deciding along which path to send packets. The first routing decision is made on the host when it prepares to send a packet. Additional routing decisions are made on any routers (either dedicated hardware devices or hosts that deal with routing as well as any other function) on the way to the destination.

The aim of all routing decisions is to get any given packet to its destination in the best way possible. How "best" is defined may vary depending on the requirements of the application that is actually generating the traffic. For many requirements, the fastest route is required; but for other applications it may be desirable to use the cheapest route, or the most reliable.

The host's decision process to allow it to choose how to send a packet is effectively:

1. Is the destination on the same subnet or on a different subnet or network?
2. If not on the same subnet, check the internal routing table for a route.
3. If there is no route specified, use the default gateway.

Assuming the packet is not destined for a machine on the local network and that no other route is defined, it will pass on to a router, which will check its own internal routing table. If that doesn't contain a suitable route, the packet will then be sent to the router's default gateway. This process is repeated for each part of the packets journey or **hop,** until it finally reaches a router that knows how to get to the host. The last router in the chain will have an interface on the same subnet as the destination host, and so the packet will finally be delivered to the destination host itself.

The routing tables referred to, and the routes they contain, are created and maintained using a combination of two routing methods: **static** and **dynamic**. Static routes must be maintained by the System Administrator. Dynamic routing is maintained by one or a combination of protocols, depending on the size and type of network and to what else the network is connected.

Static Routing

The functions required to use static routing are built into the implementation of IP on a host. When it needs to send out a packet, the IP layer consults a table, which is held in memory, containing details about itself, the default gateway, any additional routes which have been configured using the route utility and any routes which have been learned as a result of receiving an ICMP redirect message (where a router has offered a different route to a host).

Static routes are not shared between routers (or hosts acting as routers in the case of a *Windows NT* system); each router must be configured manually.

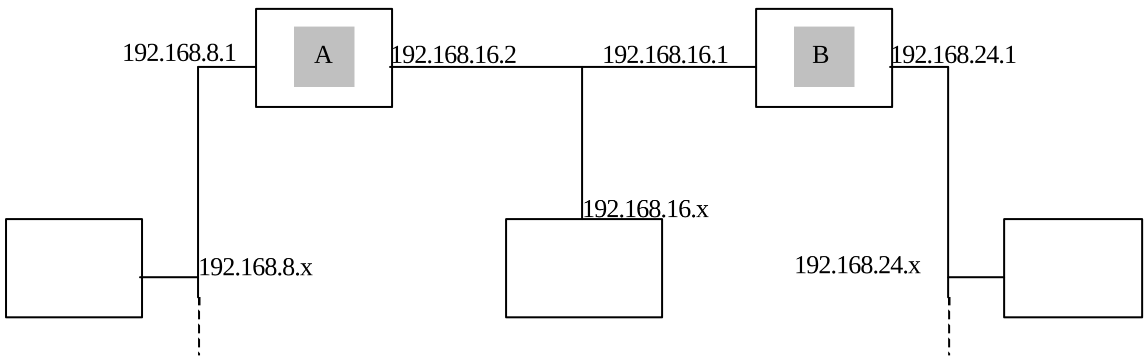

In the example above, both A and B are hosts that are also configured to act as routers. Since they have network interfaces connecting to more than one network, they are referred to as being multihomed.

A is only connected to the 192.168.8.0 and 192.168.16.0 networks; therefore, without additional configuration, is not able to contact the 192.168.24.0 network. Similarly, B is connected to 192.168.16.0 and 192.168.24.0, and cannot contact 192.168.8.0.

In order to allow the hosts on all the networks to communicate with each other, the following steps must be taken:

1. A static route is created on A, creating a route to the 192.168.24.0 network using the interface 192.168.16.1

(route add 192.168.24.0 mask 255.255.255.0 192.168.16.1)

2. A static route is created on B, creating a route to the 192.168.7.0 network using the interface 192.168.16.2

(route add 192.168.7.0 mask 255.255.255.0 192.168.16.2)

As was shown earlier with the description of IP subnetting, it is possible to implement a simple routing system by using the default gateway entries for the router. However, this is only really effective for networks with only one or two routers. On *Windows NT*, static persistent routes are stored in the Registry. Non-persistent routes are only stored in memory while they are active, and are lost when the system is restarted. route (the utility) is described in more detail in the TCP/IP utilities section.

Dynamic Routing

Dynamic routing is the process by which routers exchange information about the routes they know to other networks, by using one of the routing protocols. *Windows NT* supports RIP, although this is somewhat outdated. More modern systems, especially those providing routing services on larger networks or the Internet, use OSPF (Open Shortest Path First) or BGP (Border Gateway Protocol).

A dynamic routing facility becomes almost essential with even a very small number of routers on an internet, due to the administration overhead implicit in using static routing. With dynamic routing, minimal Administrator intervention is required. For the hosts on the network, all that is required is that their default gateway is set to the address of the local interface of the router, which has access to the rest of the internet.

Although there are a number of routing protocols available, choosing which one to use is usually fairly easy. In the case of a network where the routers are *Windows NT* systems, you are limited to RIP, since that is the only **interior gateway protocol** (IGP) that *Windows NT* supports. If dedicated hardware routers are in use, then the newer OSPF becomes an option.

When connecting networks together, there often isn't much choice of routing protocol. If, of the two **autonomous systems** (AS), the other is already communicating with other networks, then they will have already chosen the protocol to be used and any new networks connecting will have to comply. This is particularly true of the Internet, where BGP is usually a requirement.

Static routers do not exchange routing information with dynamic routers. In order to route from a static router through a dynamic router, you need to add static routes to both the static and dynamic routers. An additional complication, potentially, is that dynamic routers (some implementations of RIP in particular) may not propagate these static routes. In this case, all the routers on the internet must have the appropriate static routes added and subsequently maintained.

RIP

The Routing Information Protocol (RIP) is a simple and well-supported protocol, which is especially suited to smaller networks. Each router maintains a complete list of all the networks it can reach. Routers send this table out as a broadcast update every 30 seconds (by default) on all the networks to which they are connected.

When a RIP update is received by a router, either because it was just started and requested an update or because it was a periodic update sent by another router, the new routes received are put into the routing table (or not) using the following criteria:

- If the update contains a route that is not already in the local table, add it, unless it would be too far away (more than 15 hops).
- If the update contains a route which is already in the table, ignore it unless it is nearer (fewer hops) than the existing route.
- If the update contains a route which has more than 15 hops, delete it, unless another route is known.

RIP uses the hop count - the number of routers that must be gone through - to get to a given destination, as its metric. A route with a cost of *n* has *n* hops, or goes through *n* routers to get to its destination. The maximum number of hops that RIP will allow is 15. Any network further away than that is considered to be unreachable.

If an update has not been received from a router within 3 minutes (6 updates missed) all the routes through that gateway are set to a metric of 16 (which means infinity or unreachable in RIP terms), and marked for deletion. Actual deletion of the route is then delayed for another 60 seconds to ensure that the deletion or invalidation of that route is propagated. A router that receives RIP updates, but which does not send out any itself, is known as a silent RIP router.

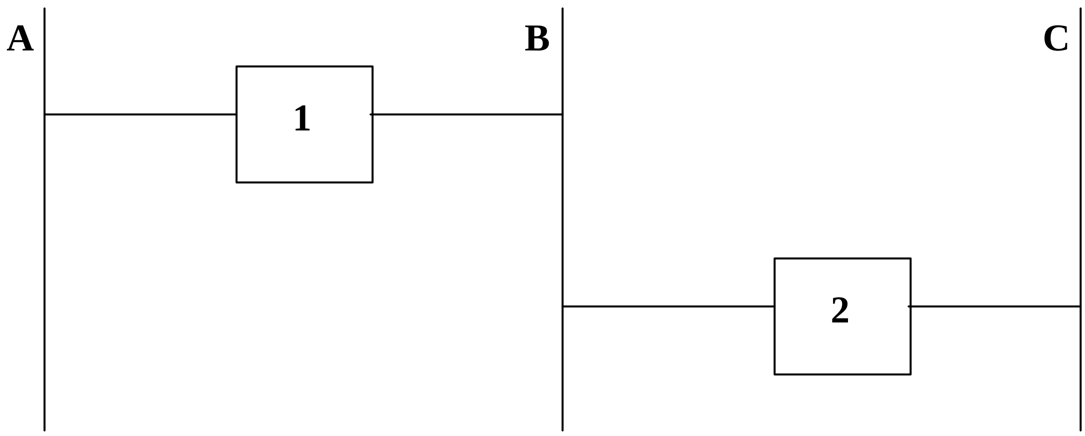

In the above diagram, there is one hop, via Router 1, between networks A and B, and two hops between networks A and C. Router 1 would have to have received a routing update from Router 2 to know about network C.

If there are multiple routes to a network, a router will normally choose the one with the lowest metric. If this is a slow or busy route, it may be necessary to adjust the hop count in favour of a nominally "longer" route, which may actually get packets to their destination faster than the "shorter" route.

RIP is only really suitable for smaller internets, owing to its initial LAN-based design. Since each router must maintain a complete list of all the routes on the network, a network with a large number of routes would obviously require more memory in each router. Then, because routers broadcast their routing table periodically, there can be quite a lot of network traffic due to RIP updates alone. The maximum size of a RIP message is 512 bytes, which may contain up to 25 routes (actually each route takes 20 bytes so 20 bytes x 25 routes + 4 bytes RIP header = 504 bytes). If a router has more routes than this to advertise, multiple packets are used.

Since the longest route supported is 15 hops, no part of the internet may be more than that far away from any other part if connectivity is to be maintained. In some cases, it could be necessary to add extra routers to bring networks within reach of each other, with a corresponding increase in RIP broadcasts.

A RIP-based internet may take a long time, often of the order of several minutes, to become stable (all routes valid) in the case of the failure of a link or router. While the network routing is settling again, routing loops may occur, and some parts of a network may become temporarily unreachable. The process of re-stabilising the network is called **convergence.** Once all the routes are up to date again, the network is said to have converged.

Finally, RIP has no support for subnets.

There are currently two versions of RIP. Version 2 (RIP-2 or RIP-II as it's known) keeps the same basic protocol, but adds some extra information in fields that RIP sets to zero. This includes support for subnets and exterior gateway protocols. RIP and RIP-2 may inter-operate as long as these fields are ignored. A simple authentication scheme is added, which includes a password. Multicasting is supported as well as broadcast, which may reduce the load on hosts that are not listening for broadcast traffic; although not reducing the overall amount of network traffic generated.

RIP messages are sent as UDP datagrams, port number 520.

OSPF

The Open Shortest Path First protocol is a newer alternative to RIP, overcoming its limitations. In contrast to RIP, for which the distance to another network is the most important criterion (otherwise described as being a **distance-vector** protocol), OSPF keeps track of the state of the links between routers (a **link-state** protocol). Link-state protocols are able to converge more quickly than distance vector protocols, in the event of a problem somewhere.

In operation, each router periodically checks the status of the links to each neighbouring router. This information is then sent to all the other neighbouring routers which also propagate it on to their neighbours until the information from any given router has worked its way thought the entire internet. In this way, each router is able to build a complete routing table.

Other OSPF features:

- Uses IP directly, not UDP or TCP.
- Simple password-based authentication scheme (like RIP-2).
- Subnets are supported.
- Uses multicast (as is an option for RIP-2).
- Can calculate a different route for each IP type-of-service.
- Each interface and type-of-service may be arbitrarily assigned its own cost, based on distance, link capacity, reliability or some other criteria.
- If there are a number of routes to the same destination each with the same cost, OSPF will **load balance** the traffic, sending roughly equal amounts by each route.
- Links that are solely between routers do not need an IP address at both ends. These are called **unnumbered networks.**

Detailed knowledge of OSPF is not required for the certification exam.

BGP

Border Gateway Protocol (BGP), currently up to version 4, is an exterior gateway protocol used to communicate between routers in different autonomous systems, imposing no constraints on the underlying network topology.

The key feature of BGP is that the path availability information, which routers exchange, includes the complete list of autonomous systems that need to be travelled to get to a particular network. This allows for the easy suppression of looping routing information and the building of a complete picture of connectivity. Routing policies may also be implemented based on this information. BGP is neither purely a distance vector (e.g. RIP) nor link state (e.g. OSPF) protocol.

After the initial complete exchange of routing information when a router first goes online, BGP subsequently uses incremental updates, which exchange only the changes in routing information. This requires a reliable transport between routers, therefore TCP is used.

Attaching to an existing internet such as the Internet will require the issuing of an Autonomous Systems Number (ASN) so that the new network can be included in the routing tables.

> Knowledge of BGP is not required for the certification exam.

Routing on Windows NT

Windows NT Server is capable of both static and dynamic routing however Windows NT Workstation only supports static routes. Static routes are implemented the same way on both platforms, using the `route` utility.

Implementing Dynamic Routing

To use a *Windows NT* server as a dynamic router, follow these steps.

1. Install all the required network interface cards, and configure with appropriate IP addresses and subnet masks.

2. Open the `Network` applet in `Control Panel`. On the `Protocols` tab select TCP/IP Protocol and click `Properties`. On the `Routing` tab check the `Enable IP Forwarding` option is selected then click `OK`.

3. On the `Services` tab of the `Network` applet, add the RIP for Internet Protocol service.

Dead Gateway Detection

If multiple gateways have been defined (under `Advanced IP Addressing` in the TCP/IP properties), then *Windows NT* can update its routing table and select the next available gateway if the default gateway appears to fail. Failure is recognised by the gateway not responding within half of the value specified for `TcpMaxDataRetransmissions`.

Summary

- Routing: How to get the packet to its destination in the best way possible.
- Routing information may be static, dynamic or a combination.
- Static routes must be maintained by Administrator.
- Dynamic routes shared automatically between routers.
- *Windows NT Server* supports RIP. More modern protocols are RIP-2 and OSPF.
- *Windows NT Workstation* supports only static routing.
- BGP used between service providers and other major networks on the Internet.