# Routing Protocols

This document is intended as an introduction to routing, routing protocols and a comparison of their use in enterprises

**What is a Router?**

A router is an electronic device and/or software that connects at least two networks and forwards packets among them according to the information in the packet headers and routing tables. Routers are fundamental to the operation of the Internet and other complex networks (such as enterprise-wide networks).
A network consists of two or more computers, and typically other devices as well (such as printers and external hard drives), that are linked together so that they can communicate with each other and thereby share files and the devices. Examples of the networks connected by a router can be two LANs (local area networks) or WANs (wide area networks) or a LAN and its ISP's (Internet service provider's) network.
A packet is the fundamental unit of information transport in all modern computer networks, and increasingly in other communications networks as well. A packet header is the portion of a data packet that precedes the body (i.e., a portion of the message being transmitted) and which contains source and destination IP addresses as well as control and timing information required for successful transmission.

**What is Routing?**

Routing is the process of moving packets across a network from one host to another. It is usually performed by dedicated devices called routers.
Packets are the fundamental unit of information transport in all modern computer networks, and increasingly in other communications networks as well. They are transmitted over packet switched networks, which are networks on which each message (i.e., data that is transmitted) is cut up into a set of small segments prior to transmission. Each packet is then transmitted individually and can follow the same path or a different path to the common destination. Once all of the packets have arrived at the destination, they are automatically reassembled to recreate the original message.
Routing is a key feature of the Internet and it, together with a great deal of deliberate redundancy of high capacity transmission lines (e.g., optical fiber cable and microwave), is a key factor in the robustness (i.e., resistance to equipment failure) of the Internet. Each intermediary router performs routing by passing along the message to the next router. Part of this process involves analysing self-configuring routing tables to determine the best (i.e., optimal) path.
Routing is sometimes confused with bridging, which performs a somewhat similar function. The main difference is that the latter occurs at a lower level of the OSI (open systems

interconnect) model and is thus more of a hardware function; the former occurs at a higher level where the software component is more important, and thus it can perform more complex analysis to determine the optimal path for each packet.
Routing is also used by circuit switched networks, in which a dedicated circuit is established for the duration of the transmission of each message. The dominant circuit switched network is the public switched telephone network (PSTN), which is the worldwide collection of interconnected public telephone networks that was designed primarily for voice traffic.

**Routing protocol**

A generic term that refers to a set of rules, or protocol, used by a router to determine the appropriate path over which data packets is transmitted. The routing protocol also specifies how routers in a network share information with each other and report changes. The routing protocol enables a network to make dynamic adjustments to its conditions (e.g., changes in network topology and traffic patterns), so routing decisions do not have to be predetermined and static.
Packet routing on the Internet is divided into two main groups: interior routing and exterior routing. The former occurs inside of independent networks, referred to as autonomous systems. The latter are used between the autonomous systems.
There are two main types of algorithms for IP (Internet protocol) routing: distance vector and link state. The former determine best path on the basis of how far the destination is, typically in terms of the number of hops to the destination. The latter use more sophisticated methods taking into consideration link variables such as bandwidth, delay, reliability and load.
Examples of distance vector routing protocols include RIPv2 (routing information protocol) and IGRP (interior gateway routing protocol). Examples of link state protocols are BGP (border gateway protocol), OSPF (open shortest path first), NLSP (NetWare link services protocol) and IS-IS (intermediate system to intermediate system).

**Static Routing**

Manually configuring routes on your router can be both beneficial and disadvantageous. Static routing has the following benefits:
No extra processing and added resources as in the case of dynamic routing protocols
No extra bandwidth requirement caused by the transmission of excessive packets for the routing table update process
Extra security by manually admitting or rejecting routing to certain networks
Disadvantages of static routing include the following:

Network Administrators need to know the complete network topology very well in order to configure routes correctly
Topology changes need manual adjustment to all routers something which is very time consuming

**RIPv2 (Routing Information Protocol Version2)**

The Routing Information Protocol (RIP) defines a way for routers, which connect networks using the Internet Protocol (IP), to share information about how to route traffic among networks. RIP is classified by the Internet Engineering Task Force (IETF) as an Interior Gateway Protocol (IGP), one of several protocols for routers moving traffic around within a larger autonomous system network -- e.g., a single enterprise's network that may be comprised of many separate local area networks (LANs) linked through routers.

Although once the most widely used IGP, Open Shortest Path First (OSPF) routing has largely replaced RIP in corporate networks.

Each RIP router maintains a routing table, which is a list of all the destinations (networks) it knows how to reach, along with the distance to that destination. RIP uses a distance vector algorithm to decide which path to put a packet on to get to its destination. It stores in its routing table the distance for each network it knows how to reach, along with the address of the "next hop" router -- another router that is on one of the same networks -- through which a packet has to travel to get to that destination. If it receives an update on a route, and the new path is shorter, it will update its table entry with the length and next-hop address of the shorter path; if the new path is longer, it will wait through a "hold-down" period to see if later updates reflect the higher value as well, and only update the table entry if the new, longer path is stable.

Using RIP, each router sends its entire routing table to its closest neighbours every 30 seconds. (The neighbours are the other routers to which this router is connected directly -- that is, the other routers on the same network segments this router is on.) The neighbours in turn will pass the information on to their nearest neighbours, and so on, until all RIP hosts within the network have the same knowledge of routing paths, a state known as convergence.

If a router crashes or a network connection is severed, the network discovers this because that router stops sending updates to its neighbours, or stops sending and receiving updates along the severed connection. If a given route in the routing table isn't updated across six successive update cycles (that is, for 180 seconds) a RIP router will drop that route, letting the rest of the network know via its own updates about the problem and begin the process of recon verging on a new network topology.

RIP uses a modified hop count as a way to determine network distance. (Modified reflects the fact that network engineers can assign paths a higher cost.) By default, if a router's neighbour owns a destination network (i.e., if it can deliver packets directly to the destination network without using any other routers), that route has one hop, described as a cost of 1. RIP allows only 15 hops in a path. If a packet can't reach a destination in 15 hops, the destination is considered unreachable. Paths can be assigned a higher cost (as if they involved extra hops) if the enterprise wants to limit or discourage their use. For example, a satellite backup link might be assigned a cost of 10 to force traffic follow other routes when available.

RIP has been supplanted mainly due to its simplicity and its inability to scale to very large and complex networks. Other routing protocols push less information of their own onto the network, while RIP pushes its whole routing table every 30 seconds. As a result, other protocols can converge more quickly, use more sophisticated routing algorithms, include latency, packet loss, actual monetary cost and other link characteristics, as well as hop count with arbitrary weighting.

Due to the deficiencies of RIPv1, RIP version 2 (RIPv2) was developed sometime in 1993. It’s equipped with the ability to support subnet information and supports Classless Inter-Domain Routing (CIDR). A router that receives routing updates from multiple routers advertising the same classful summary route cannot determine which subnets belong to which summary route. This inability leads to unexpected results including misrouted packets.

However, with RIPv2 automatic summarization can be disabled with the no auto-summary command. Automatic summarization must be disabled to support discontinuous networks.

RIPv2 still maintains the hop count limit of 15 and incorporated a password authentication mechanism. However, passwords were transmitted in clear-text format, which were found insufficient for secure communications on the Internet.

The default version of RIP is version 1. The command version 2 is used to modify RIPv1 to RIPv2.

Use the show ip protocols command to view that RIP is now sending and receiving version 2 updates and whether or not automatic summarization is in effect.

RIPv2 is actually an enhancement of RIPv1's features and extensions rather than an entirely new protocol. Some of these enhanced features include:

Next-hop addresses included in the routing updates
Use of multicast addresses in sending updates
Authentication option available

Like RIPv1, RIPv2 is a distance vector routing protocol. Both versions of RIP share the following features and limitations:

Use of holddown and other timers to help prevent routing loops.
Use of split horizon or split horizon with poison reverse to also help prevent routing loops.
Use of triggered updates when there is a change in the topology for faster convergence.

Maximum hop count limit of 15 hops, with the hop count of 16 signifying an unreachable network.

**OSPF (Open Shortest Path First)**

An interior gateway routing protocol developed for IP networks based on the shortest path first or link-state algorithm.
Routers use link-state algorithms to send routing information to all nodes in an internetwork by calculating the shortest path to each node based on a topography of the Internet constructed by each node. Each router sends that portion of the routing table (keeps track of routes to particular network destinations) that describes the state of its own links, and it also sends the complete routing structure (topography).
The advantage of shortest path first algorithms is that they results in smaller more frequent updates everywhere. They converge quickly, thus preventing such problems as routing loops and Count-to-Infinity (when routers continuously increment the hop count to a particular network). This makes for a stable network.
The disadvantage of shortest path first algorithms is that they require a lot of CPU power and memory. In the end, the advantages outweigh the disadvantages.
OSPF Version 2 is defined in RFC1583. It is rapidly replacing RIP on the Internet.

**EIGRP**
Enhanced Interior Gateway Protocol (EIGRP) is a proprietary hybrid routing protocol developed by Cisco Systems. EIGRP uses the same distance vector algorithm and distance information as IGRP. However, as its name implies, EIGRP has been enhanced in convergence properties and operating efficiency over IGRP. Principally, EIGRP has been enhanced to use more advanced features to avoid routing loops and to speed convergence time. In addition, EIGRP transmits the subnet mask for each routing entry, enabling EIGRP to support features such as VLSM and route summarization.

EIGRP Features
EIGRP provides advanced features over its predecessors IGRP and RIP:

Increased network width— With IP RIP, the largest possible width of your network is 15 hops. When IP EIGRP is enabled, the largest possible width is 224 hops.
Fast convergence— EIGRP uses an algorithm called the Diffusing Update Algorithm (DUAL). This algorithm guarantees loop-free operation at every instant throughout a route computation and allows all routers involved in a topology change to synchronize at the same time. Routers that are not affected by topology changes are not involved in recomputations. DUAL provides a system for routers to not only calculate the best current route to each

subnet, but also to calculate alternative routes that could be used if the current route fails. The alternate route, called the feasible successor route, is guaranteed to be loop-free, so convergence can happen quickly. Because of DUAL, the convergence time of EIGRP rivals that of other existing routing protocols.

Partial updates— EIGRP sends incremental updates when the state of a destination changes, instead of sending the entire contents of the routing table. This feature reduces the bandwidth required for EIGRP packets and also reduces CPU processing.

Neighbour-discovery mechanism— this is a simple hello mechanism used to learn about neighbouring routers. It is protocol-independent.

VLSM and route summarization— EIGRP supports variable-length subnet masks and route summarization.

Automatic redistribution— Because IGRP and EIGRP share the same metrics, IP IGRP routes can be automatically redistributed into EIGRP, and IP EIGRP routes can be automatically redistributed into IGRP. If desired, you can turn off redistribution. Redistribution is covered in more detail in Chapter 11, "Route Redistribution."

EIGRP Components

EIGRP has four basic components:

Neighbour discovery/recovery— Routers use this process to dynamically learn of other routers on their directly attached networks. Routers also must discover when their neighbours become unreachable or inoperative. Neighbour discovery/recovery is achieved with low overhead by periodically sending small hello packets. As long as hello packets are received, a router can determine that a neighbour is alive and functioning. When this status is determined, the neighbouring routers can exchange routing information.

Reliable transport protocol— this protocol is responsible for guaranteed, ordered delivery of EIGRP packets to all neighbours. It supports intermixed transmission of multicast and unicast packets. Some EIGRP packets must be transmitted reliably and others need not be. For efficiency, reliability is provided only when necessary. For example, on a multi-access network that has multicast capabilities, such as Ethernet, it is not necessary to send hellos reliably to all neighbours individually. Therefore, EIGRP sends a single multicast hello with an indication in the packet informing the receivers that the packet need not be acknowledged. Other types of packets, such as updates, require acknowledgment, and this is indicated in the packet. The reliable transport has a provision to send multicast packets quickly when there are unacknowledged packets pending. Doing so helps ensure that convergence time remains low in the presence of varying speed links.

DUAL finite state machine— this embodies the decision process for all route computations. It tracks all routes advertised by all neighbours. DUAL uses the distance information, known as a metric, to select efficient, loop-free paths. DUAL selects routes to be inserted into a routing table based on feasible successors. A successor is a neighbouring router used for packet forwarding that has a least-cost path to a destination that is guaranteed not to be part of a routing loop. When there are no feasible successors but there are neighbours advertising the destination, a recomputation must occur. This is called "going active." A router asks all of its neighbours if they have a feasible successor to the destination. If none replies, the neighbours go active and the process repeats. This is the process whereby a new successor is determined. The amount of time that it takes to recompute a route affects the convergence time. Even though the recomputation is not processor-intensive, it is advantageous to avoid recomputation if it is not necessary. When a topology change occurs, DUAL tests for changes to feasible successors. If feasible successors exist, it uses any that it finds, to avoid unnecessary recomputation.

## Predictions for the Future of the Internet of Things

Wired magazine described a new era where "the most mundane items in our lives can talk wirelessly among themselves, performing tasks on command, giving us data we've never had before." The Internet of Things (IoT) is a world where up to 50 billion things (or devices) will be connected to the Internet by 2020; or, the equivalent of 6 devices for every person on the planet.

We are already starting to see the emergence of smart cities, connected utilities, connected railways, connected factories, connected cars, and even connected mines, to name but a few. The Internet of Things will fundamentally transform businesses, generate enormous economic wealth and create immeasurable social value.

## Bibliography

https://supportforums.cisco.com/document/127851/routing-protocol-selection-guide-igrp-eigrp-ospf-bgp

http://www.cisco.com/c/en/us/support/docs/ip/routing-information-protocol-rip/13769-5.html#irp

http://docwiki.cisco.com/wiki/Routing_Information_Protocol

http://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-1.html

http://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/16406-eigrp-toc.html

http://www.wired.com/tag/internet-of-things/