Routing Introduction

Direct vs. Indirect Delivery
Static vs. Dynamic Routing
Distance Vector vs. Link State

"The most simple way to accelerate a Router is at 9.8 m/sec/sec."

Seen on Usenet

Routing Basics

- **Routing Introduction**
  - **Direct Delivery**
  - **Indirect Delivery**
  - **Static Routing**
  - **Default Routing**
- **Dynamic Routing**
  - **Distance Vector Routing**
  - **Link State Routing**

3

In this chapter we talk about routing basics. About the 2 ways to deliver a packet: direct and indirect, and then about the 3 kinds of routing: the static routing, the default routing and then about the most important routing today, the dynamic routing.

What is routing?

- *Finding a path to a destination address*
- **Direct delivery** performed by host
  - Destination network = local network
- **Indirect delivery** performed by router
  - Destination network ≠ local network
  - Packet is forwarded to default gateway

There are 2 ways to delivery a packet. The **direct** delivery and the **indirect** delivery. When there is a direct delivery (destination network = local network) the host makes for example an ARP-request (Ethernet) and then delivery the packet to the right host. If there is a indirect delivery (destination network ≠ local network) the router forwards the packet to his default gateway.

Direct Delivery

- **IP host checks if packet's destination network is identical with local network**
  - ◆ **By applying the configured subnet mask of the host's interface**
- **If destination network = local network then the L2 address of the destination is discovered using ARP**
  - ◆ **Not necessary on point-to-point connections**

Before the IP host sends out his packet, he checks if the destination address of the packet is identical with the local network (subnet mask). If the destination network = local network the IP host needs a layer 2 address to deliver the packet correct. To do this he sends out an ARP-request. With the information the host receives he can send the packet to the right host in his local network.

IP Host Facts

- **Also IP hosts have routing tables !**
  - **But typically only a static route to the default gateway is entered**
- **ARP cache aging timer: 20 minutes**

Note that also simple workstations and PCs maintain routing tables—but not for routing passthrough packets, rather locally originated packets should be routed to the most reasonable next hop. Typically, the routing table consists only of a single entry, which is the default gateway for this host.

But also additional entries can be made, indicating other gateways for some dedicated routes.

Additionally, an ARP cache must be maintained by a host. The ARP cache stores layer-2 MAC addresses and associated IP addresses of interfaces to which communication had occured recently. Any ARP result is stored in this cache, thus subsequent packets to the same destination do not invoke the ARP each time. Per default the ARP cache is flushed after 20 minutes. Of course this value can be configured individually—even by DHCP.

Indirect Delivery

- **Default gateway delivers packet in behalf of its host using a routing table**
- **Routing table components**
  - **Destination network (+ subnet mask)**
  - **Next hop (+ outgoing interface)**
  - **Metric (+ Administrative Distance)**

Every router has his own routing table. This table contains many information such as destination network + subnet mask, next hop, metric, etc. If the destination address of a packet ≠ local network the host sends the packet to the router. The router compares this address with his routing table and make a forward decision. Most of the time router in small networks have a so called **default gateway.** For example this gateway is used to forward a packet to a router who is connected to the internet.

Router

- **Initially Unix workstations with several network interface cards**
- **Today specialized hardware**
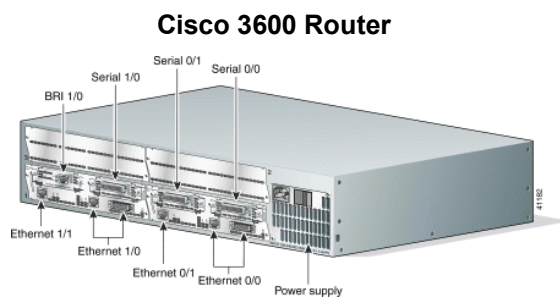
Cisco 3600 Router

The picture above shows one of the most used routers today, the Cisco 3600 platform, employing various Ethernet and Serial interfaces.

Update: Today (2008) the most commonly used Cisco router series are the "Integrated Services Routers" 800 (SOHO), 1800, 2800, and 3800, while in larger networks the 7200 or 7600 series routers are found.
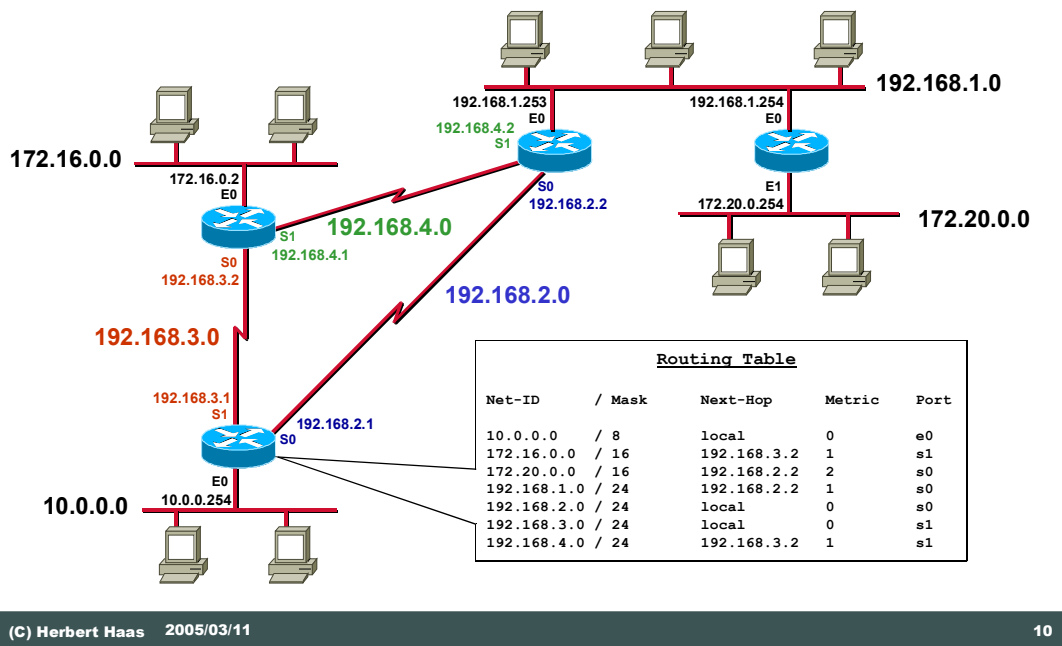
Routing Table Example

```
Gateway of last resort is 175.18.1.2 to network 0.0.0.0

     10.0.0.0 255.255.0.0 is subnetted, 4 subnets
C       10.1.0.0 is directly connected, Ethernet1
R       10.2.0.0  [120/1] via 10.4.0.1, 00:00:05, Ethernet0
R       10.3.0.0  [120/5] via 10.4.0.1, 00:00:05, Ethernet0
C       10.4.0.0 is directly connected, Ethernet0
R     192.168.12.0  [120/3] via 10.1.0.5, 00:00:08, Ethernet1
S     194.30.222.0 [1/0] via 10.4.0.1
S     194.30.223.0 [1/0] via 10.1.0.5
C     175.18.1.0 255.255.255.0 is directly connected, Serial0
S*    0.0.0.0 0.0.0.0 [1/0] via 175.18.1.2
```

In the picture above there is example of a routing table. 0.0.0.0 is used for default gateway. The single letters at the beginning of each entry indicates how the routes were learned, for example "C" corresponds to "Directly Connected", "R" means "learned by RIP", "S" means "static route", and so on. The numbers in the brackets denote the administrative distance and the metric. For example [120/5] means AD=120, metric=5.

IP Routing Basics

172.16.0.0

192.168.1.253
192.168.4.2
E0
S1

192.168.1.254
E0

192.168.1.0

172.16.0.2
E0

S0
192.168.2.2

S1
192.168.4.1

192.168.4.0

E1
172.20.0.254

172.20.0.0

S0
192.168.3.2

192.168.2.0

192.168.3.0

192.168.3.1
S1

192.168.2.1
S0

E0
10.0.0.254

10.0.0.0

```
                        Routing Table

Net-ID      / Mask        Next-Hop      Metric     Port

10.0.0.0    / 8           local         0          e0
172.16.0.0  / 16          192.168.3.2   1          s1
172.20.0.0  / 16          192.168.2.2   2          s0
192.168.1.0 / 24          192.168.2.2   1          s0
192.168.2.0 / 24          local         0          s0
192.168.3.0 / 24          local         0          s1
192.168.4.0 / 24          192.168.3.2   1          s1
```

In the picture above there is small network, and a good example of a routing table. For example a host in network 10 want to send a packet to a user in network 192.168.1.

The destination address ≠ local address so the router must do a forward decision. The router compare the destination address with his routing table and found the right match (192.168.1.0/ 24 192.168.2.2 1 s0). Now he sends out the packet via port s0 to the next hop, the router with the IP-Address of 192.168.2.2. This router is direct connected to the network 192.168.1. After an ARP-request the packet deliver to the right user.

Static or Dynamic

- **Static routing entries are configured manually**
  - ♦ **Override routes learned via dynamic routing**
  - ♦ ***Can* be set as permanent (will not be removed if interface goes down)**
  - ♦ **Only way for certain technologies (DDR)**
- **Dynamic routing entries are learned by routing protocols**
  - ♦ **Adapts to topology changes**
  - ♦ **But additional routing-traffic overhead**

The only different between static and dynamic routing is, that static routing entries configured manually, and dynamic routing entries are learned by routing protocols. Static routes can be set as permanent, this means that such kind of entries will not be removed when the interface goes down or cannot overwrite by routing protocols.

Reasons for Static Routing

- **Very low bandwidth links (e. g. dialup links)**
- **Administrator needs control over the link**
- **Backup links**
- **Link is the only path to a stub network**
- **Router has very limited resources and cannot run a routing protocol**

ip route *prefix mask* **{***ip-address* **|** *interface-type interface-number***}** **[***distance***] [tag** *tag***] [permanent]**

Tag value that can be used as a "match" value for controlling redistribution via route maps

Specifies that the route will not be removed, even if the interface shuts down

(If you don't understand the tag keyword then please wait for the follow-up lectures).

Routing Paradigm

- **Destination Based Routing**
  - **Source address is not taken into account for the forward decision**
- **Hop by Hop Routing**
  - **IP datagram's follow the signposts given by routing table entries**
  - **Network's routing state must be loop-free and consistent**
- **Least Cost Routing**
  - **Typically only the best path is entered into routing table**

The IP routing paradigm is fundamental in IP routing. Firstly, IP routing is "destination based routing", that means the source IP address is never examined during the routing process. Secondly, IP routing is "hop-by-hop", which emphasizes the difference to virtual circuit principles. The routing table in every router within the autonomous system must be both accurate and up to date so that datagrams can be directed across the network to their destination.
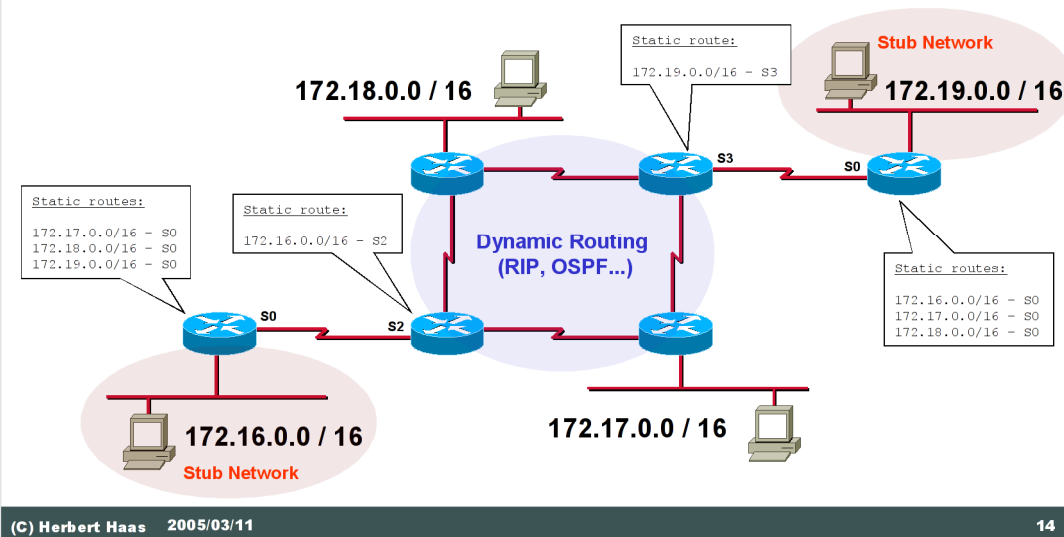
In IP the path of a packet is not pre-defined and not connection oriented, rather each single router performs a routing decision for each packet. Thirdly, IP routing is "least cost" in that only that path with the lowest metric is selected in case of multiple redundant paths to the same destination.

Note that several vendors extend these rules by providing additional features, but the routing paradigm generally holds for most of the routers in the Internet, at least for the basic routing processes.

Static Routing (1)

- **Static routes to and from stub networks**

172.18.0.0 / 16

Static route:
172.19.0.0/16 – S3

Stub Network
172.19.0.0 / 16

S3   S0

Static routes:
172.17.0.0/16 – S0
172.18.0.0/16 – S0
172.19.0.0/16 – S0

Static route:
172.16.0.0/16 – S2

**Dynamic Routing
(RIP, OSPF...)**

Static routes:
172.16.0.0/16 – S0
172.17.0.0/16 – S0
172.18.0.0/16 – S0

S0   S2

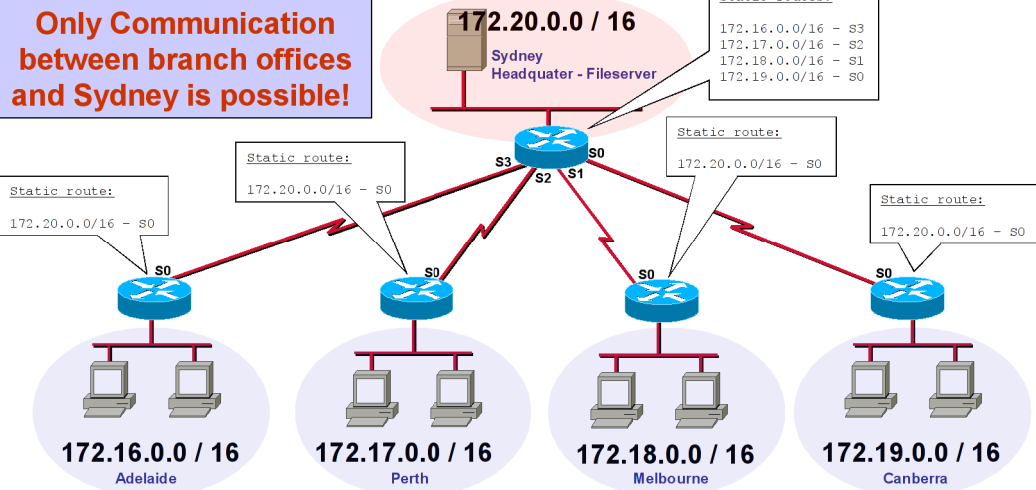172.16.0.0 / 16
Stub Network

172.17.0.0 / 16

In the picture above we see a good example of static routes. Every router in the stub networks is configured manually, because there is only one way the packet can go. Most of the time this entries are set to permanent, so the routes need not be configured again when the interface goes down.

Static Routing (2)

- ## Static routes in "Hub and Spoke" topologies

Only Communication between branch offices and Sydney is possible!

172.20.0.0 / 16

Sydney
Headquater - Fileserver

Static routes:

```
172.16.0.0/16 - S3
172.17.0.0/16 - S2
172.18.0.0/16 - S1
172.19.0.0/16 - S0
```

S3   S0
S2   S1

Static route:

```
172.20.0.0/16 - S0
```

Static route:

```
172.20.0.0/16 - S0
```

Static route:

```
172.20.0.0/16 - S0
```

Static route:

```
172.20.0.0/16 - S0
```

S0   S0   S0   S0

172.16.0.0 / 16
Adelaide

172.17.0.0 / 16
Perth

172.18.0.0 / 16
Melbourne

172.19.0.0 / 16
Canberra

Here you see a other example of static routing. Every branch office is connected over static routes with the Sydney headquarter.
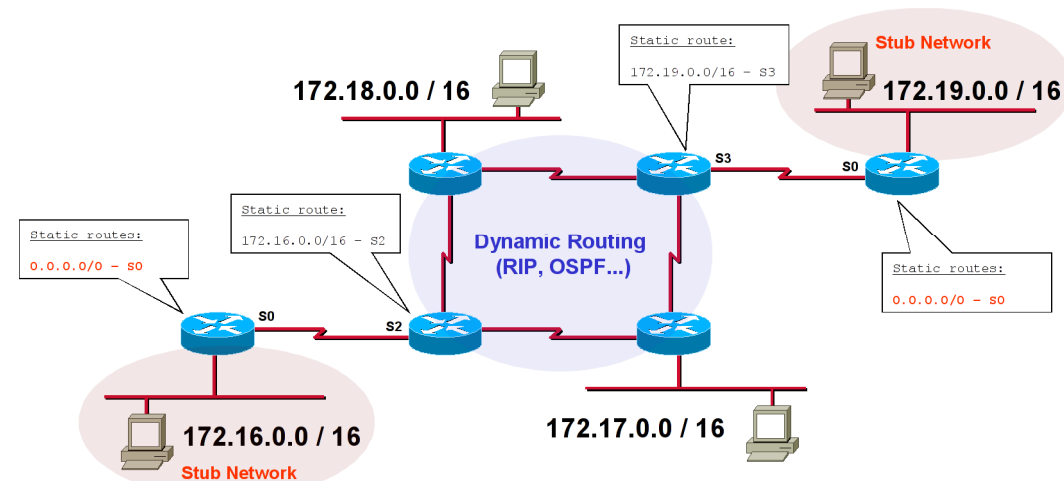
Default Routing

- **Special static route**
  - **Traffic to unknown destinations are forwarded to default router ("Gateway of Last Resort")**
- **Routing table entry "0.0.0.0 0.0.0.0"**
- **Hopefully, default gateway knows more destination networks**
- **Advantage: Smaller routing tables!**

To get smaller routing tables there is the so called default gateway. When a router receives a packet, and when the router couldn't find the destination address of the packet in his routing table he is forward this packet over his default gateway, hopefully the next router knows more.

Default Routing (1)

- **Default Routes from stub networks**

Static route:
172.19.0.0/16 – S3

Stub Network

172.18.0.0 / 16

172.19.0.0 / 16

S3

S0

Static routes:

0.0.0.0/0 – S0

Static route:
172.16.0.0/16 – S2

**Dynamic Routing
(RIP, OSPF...)**

Static routes:

0.0.0.0/0 – S0

S0

S2

172.16.0.0 / 16

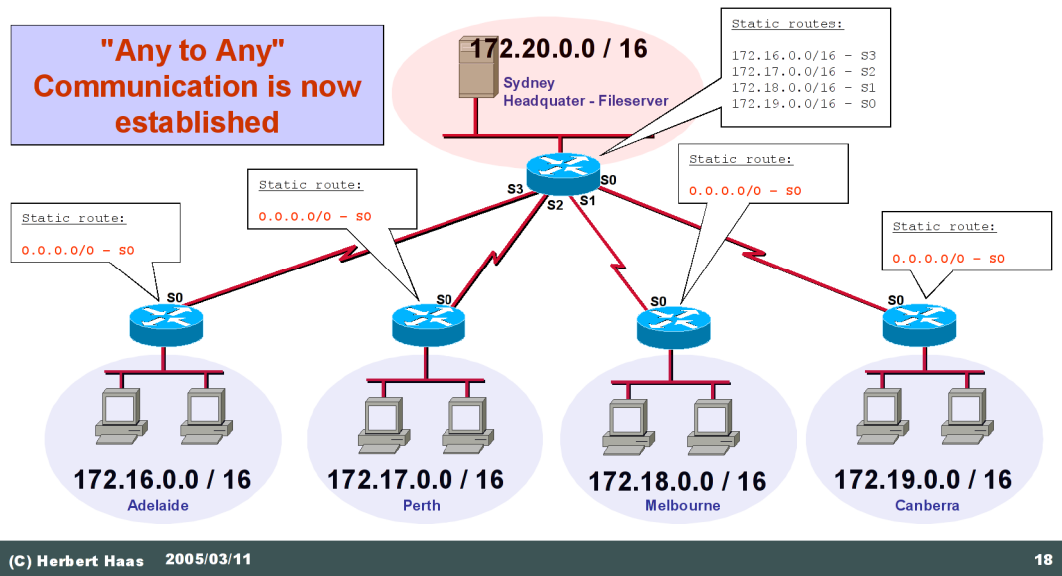**Stub Network**

172.17.0.0 / 16

In this example you see the work of default routing. Every packet a router of a stub network receives will be forward to the next router, doesn't matter what destination address.

Default Routing (2)

- **Default routes in "Hub and Spoke" topologies**

"Any to Any" Communication is now established

172.20.0.0 / 16
Sydney
Headquater - Fileserver

Static routes:
```
172.16.0.0/16 - S3
172.17.0.0/16 - S2
172.18.0.0/16 - S1
172.19.0.0/16 - S0
```

S3 S0
S2 S1

Static route:
0.0.0.0/0 - S0

Static route:
0.0.0.0/0 - S0

Static route:
0.0.0.0/0 - S0

Static route:
0.0.0.0/0 - S0

S0 S0 S0 S0

172.16.0.0 / 16
Adelaide

172.17.0.0 / 16
Perth

172.18.0.0 / 16
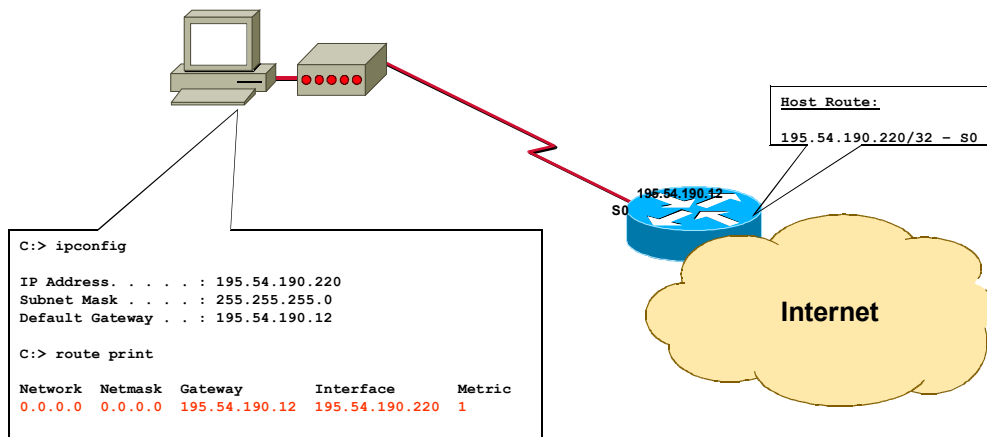Melbourne

172.19.0.0 / 16
Canberra

With default routing it is now possible that every branch office can talk with each other, and not only with the headquater.

Default Routing (3)

- **Default Routes to the Internet**

```
C:> ipconfig

IP Address. . . . . : 195.54.190.220
Subnet Mask . . . . : 255.255.255.0
Default Gateway . . : 195.54.190.12

C:> route print

Network  Netmask  Gateway        Interface       Metric
0.0.0.0  0.0.0.0  195.54.190.12  195.54.190.220  1
```

Host Route:
195.54.190.220/32 – S0

195.54.190.12
S0

Internet

Also your home pc uses the default gateway.

Router IP Address 195.54.190.12

Once the host dials in, the router assigns an IP-Address (195.54.190.220) and a default gateway (195.54.190.12) to that host and also creates a "Host Route" (dynamic) that points to that host. The host takes that default gateway information and creates a default route pointing to its local interface

On Demand Routing (ODR)

- **Efficient for hub-and-spoke topologies**
  - Same configuration at each router
- **Uses CDP to send the prefixes of attached networks from the spokes, or stub networks, to the hub or core router**
  - CDP does this automatically (!)
- **The hub router sends its interface address of the shared link as the default route for the stub router**
- **Note:**
  - Don't enable routing protocols on spoke routers
  - CDP must be enabled (don't forget e. g. ATM interfaces)
  - Every 60 sec a CDP message is sent per default (change with "cdp timer" command)

```
(config)# router odr                                    ! Only on hub router
```

ODR has the advantage of sending minimal information, such as the prefix and mask and the metric of one, every 60 seconds by default. This information populates the routing table of the hub router and can be redistributed into a routing protocol. Because the mask is sent in the update, VLSM can be used.

Dynamic Routing

- **Each router can run one or more routing protocols**
- **Routing protocols are information sources to create routing table**
- **Routing protocols differ in convergence time, loop avoidance, network size, complexity**

In contrast to static routing where every route must be configured manually, dynamic routing works with one or more routing protocols. These protocols inform the router and create the routing table automatically. Widely used in the Internet.

Routing Protocol Comparison

Routing Protocol	Complexity	Max. Size	Convergence Time	Reliability	Protocol Traffic
RIP	very simple	16 Hops	Up to 480 secs	Not absolutely loop-safe	High
RIPv2	very simple	16 Hops	Up to 480 secs	Not absolutely loop-safe	High
IGRP	simple	x	x	medium	medium
EIGRP	complex	x	x	x	x
OSPF	very complex	Thousands of Routers	Fast	High	low/ depends
IS-IS	complex	Thousands of Routers	Fast	High	x
BGP-4	complex	more than 100,000 networks	Fast	Very High	x

The table above gives a rough comparison of the most important routing protocols used today. Note that some values can not easily determined and are left blank for this reason.

Metric

- **Routing protocols typically find out more than one route to the destination**
- **Metrics help to decide which path to use**
  - **Hop count**
  - **Cost (reciprocal value of bandwidth)**
  - **Load, Reliability, Delay, MTU**

Often router find more than one path to forward a packet to a given destination. The metric helps router find the "best" way. Note that there are several types of metrics used in modern routing protocols. Typically they cannot be compared with each other. For example a simple hop-count is no measure for speed (bandwidth).

Administrative Distance

- **Several routing protocols independently find out different routes to same destination**
  - **Which one to choose?**
- **"Administrative Distance" is a trustiness-value associated to each routing protocol**
  - **The lower the better**
  - **Can be changed**

If several different routing protocols suggest different paths to the same destination at the same time, the router makes a trustiness decision based on the "Administrative Distance", which is a Cisco feature. Each routing protocol has assigned a static AD value indicating the "trustiness" – the lower the better. Of course these values can be manipulated for special purposes.

Administrative Distances Chart

Unknown	255
I-BGP	200
E-EIGRP	170
EGP	140
RIP	120
IS-IS	115
OSPF	110
IGRP	100
I-EIGRP	90
E-BGP	20
EIGRP Summary Route	5
Static route to next hop	1
Static route through interface	0
Directly Connected	0

Note the difference between static routes, if the next hop either points to an interface (AD=1) or if the route is configured as directly connected (AD=0)

AD also tells the router that E-BGP updates are more trustworthy than I-BGP messages.

Remember

1) Using the METRIC *one* routing protocol determines the best path to a destination.

2) A router running multiple routing protocols might be told about multiple possible paths to one destination.

3) Here the METRIC cannot help for decisions because *different type of METRICS* cannot be compared with each other.

4) A router chooses the route which is proposed by the routing protocol with the *lowest* ADMINISTRATIVE DISTANCE

AD with Static Routes

- **Each static route can be given a different administrative distance**
- **This way fall-back routes can be configured**

AD = 20

AD = 5

AD = 10

Dialup ISDN
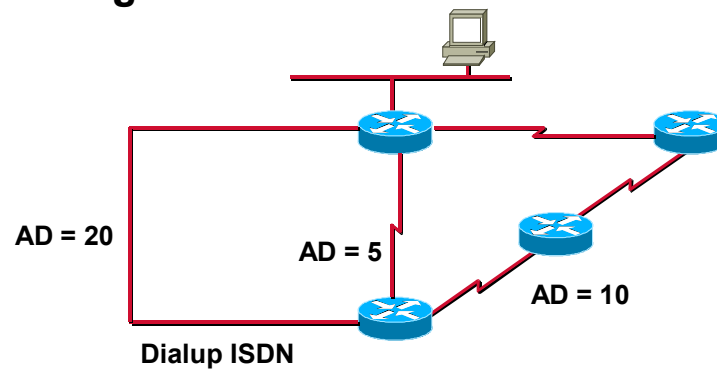
In the example above, there are several static routes to same destination. There are three paths with different quality (more or less hops, BW, ...). So every path has assign a different AD. If there are problems with the main path (AD 5) the router automatically change to the next path (AD 10) and so on.

Classification

- **Depending on age:**
  - **Classful** (no subnet masks)
  - **Classless** (VLSM/CIDR supported)
- **Depending on scope:**
  - **IGP** (Inside an Autonomous System)
  - **EGP** (Between Autonomous Systems)
- **Depending on algorithm:**
  - **Distance Vector** (Signpost principle)
  - **Link State** (Roadmap principle)

All routing protocols can be classified three-fold. If routing protocols are able to carry a subnet mask for each route we call them "classless", otherwise "classful". Today, most modern routing protocols are classless and therefore support VLSM and CIDR. If routing protocols are used inside an autonomous system we call it "Interior Gateway Protocol (IGP)", while only "Exterior Gateway Protocols (EGPs)" are used between autonomous systems. Technically, all routing protocols use one of two possible algorithms: "Distance Vector" protocols rely on the signpost principle, while "Link State" protocols maintain a road-map for the whole network.

Classful route match (1/2)

- **1) If the classful network number is NOT listed in the routing table, use the default route if available (otherwise discard the packet)**
- **2) If the classful network number is listed in the routing table:**
  - **If the listed network number is NOT subnetted and matches the IP-packet's destination address then use this route**
  - **If this network is subnetted, then lookup the corresponding subnet; if no subnet matches then *discard the packet* (even if a default route exists!)**

Example:

```
IP Packet
DA = 10.35.72.26
SA = ...
```

```
              Routing Table:

10.0.0.0/8 is subnetted, 4 subnets:
10.22.0.0/16   via 172.17.7.19
10.31.0.0/16   via 172.17.8.31
10.34.0.0/16   via 172.18.1.254
10.35.0.0/16   via 192.186.176.254
0.0.0.0/0      via 172.19.41.254
```

Classful route match (2/2)

- 1) If the classful network number is NOT listed in the routing table, use the default route if available (otherwise discard the packet)
- 2) If the classful network number is listed in the routing table:
  - If the listed network number is NOT subnetted and matches the IP-packet's destination address then use this route
  - **If this network is subnetted, then lookup the corresponding subnet; if no subnet matches then *discard the packet* (even if a default route exists!)**

Example:

IP Packet
DA = 10.35.72.26
SA = …

```
                Routing Table:

10.0.0.0/8 is subnetted, 4 subnets:
10.22.0.0/16    via 172.17.7.19
10.31.0.0/16    via 172.17.8.31
10.34.0.0/16    via 172.18.1.254
0.0.0.0/0       via 172.19.41.254
```

DISCARD THE PACKET (!)

Classless routing: Longest match

- **The subnet mask of each route entry tells how many bits must be compared with the IP-packet's destination address**
- **The router takes the route with the longest match**

Example:

IP Packet
DA = 10.35.72.26
SA = …

```
           Routing Table:

10.0.0.0/8      via 172.16.1.1
10.22.0.0/16    via 172.17.7.19
10.31.0.0/16    via 172.17.8.31
10.34.0.0/16    via 172.18.1.254
10.35.0.0/16    via 192.186.176.254
10.35.64.0/19   via 192.186.177.254
10.35.192.0/19  via 172.19.54.1
```

Distance Vector (1)

- **After powering-up each router only knows about directly attached networks**
- **Routing table is sent periodically to all neighbor-routers**
- **Received updates are examined, changes are adopted in own routing table**
- **Metric information (originally) is number of hops**
- **"Bellman-Ford" algorithm**

Distance vector protocols works with the Signpost principle. A Part of the own routing table is sent periodically to all neighbor routers (e.g.: RIP: every 30 seconds).

A signpost carries the Destination network, the Hop Count (metric, "distance") and the Next Hop.

After a router receives a update, he extracts new information's. Known routes with worse metric are ignored.

Distance Vector (2)

- **Next hop is always originating router**
  - **Topology behind next hop unknown**
  - **Signpost principle**
- **Loops can occur!**
  Additional mechanisms needed:
  - **Maximum hop count**
  - **Split horizon (with poison reverse)**
  - **Triggered update**
  - **Hold down**
- **Examples: RIP, RIPv2, IGRP (Cisco)**

Routing loops are big problems with distance vector protocols. Because of the simple principle of Distance Vector protocols, we cannot prevent rooting loops. Access Lists, Disconnection and connections, Router malfunction, etc can always lead to it, there is no 100% solution.

Link State (1)

- **Each two neighbored routers establish adjacency**
- **Routers learn real topology information**
  - ◆ **Through "Link State Advertisements"**
  - ◆ **Stored in database (Roadmap principle)**
- **Updates only upon topology changes**
  - ◆ **Propagated by *flooding*
    (very fast convergence)**

Link-state routing protocol were designed for large networks. This kind of protocols are more reliable and convergence fast.

The smallest topological unit is simply the information: ROUTER-LINK-ROUTER

Link State (2)

- **Routing table entries are calculated by applying the Shortest Path First (SPF) algorithm on the database**
  - **Loop-safe**
  - **Alternative paths immediately known**
  - **CPU and memory greedy**
- **Large networks can be split into areas**
- **Examples: OSPF, Integrated IS-IS**

Applying the SPF algorithm on the link state database, each router can create routing table entries by its own.

Summary

- **Routing is the "art" of finding the best way to a given destination**
- **Can be static or dynamic**
  - ◆ **Static means: YOU are defining the way packets are going**
  - ◆ **Dynamic means: A routing protocol is "trying" to find the best way to a given destination**
- **In today's routers the route with the longest match is used**
- **Routing protocols either implement the principle *Distance Vector* or *Link State***