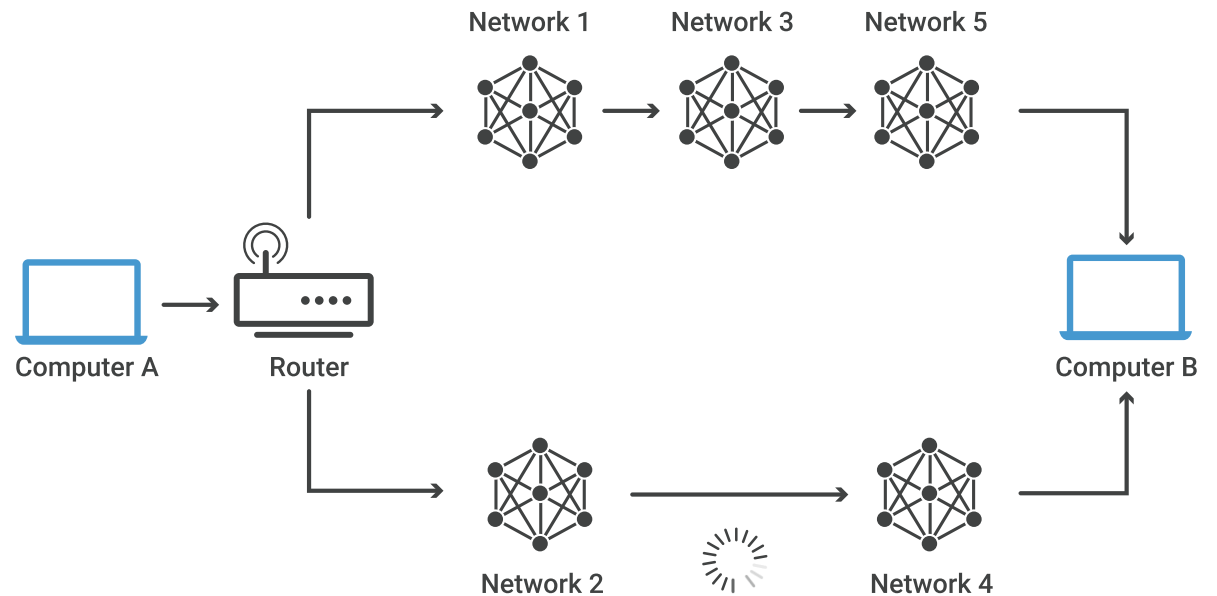


# Routing protocols

2023/2024

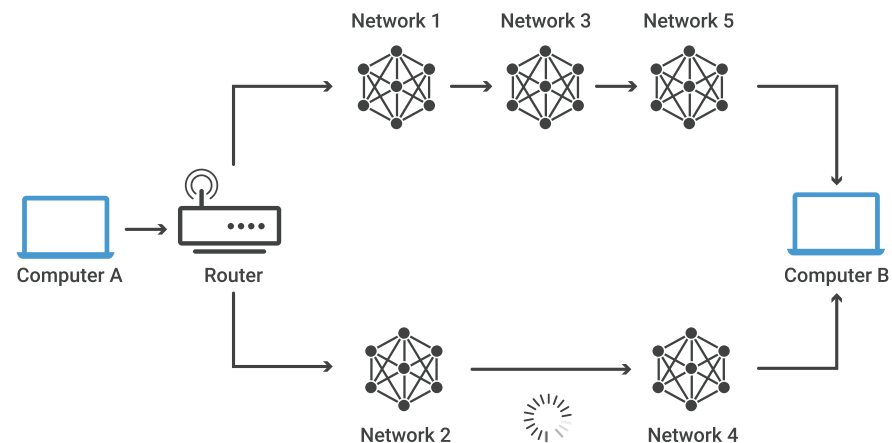
# Routing

- Routing is the process of selecting a path for traffic in a network or across multiple networks.
- Routing works by choosing the best path for data packets to take across a network when traveling from one computer to another.



# Routing

- Making sensible routing decisions is crucial to ensure reliable internet connections.
- These Internet routing decisions are made by specialized pieces of network hardware called **routers**.



# Routing

## What is a router?

- The router is a physical or virtual internetworking device that is designed to receive, analyze, and forward data packets between computer networks.
- A router examines a destination IP address of a given data packet, and it uses the headers and forwarding tables to decide the best way to transfer the packets.



# Routing

What is a router?

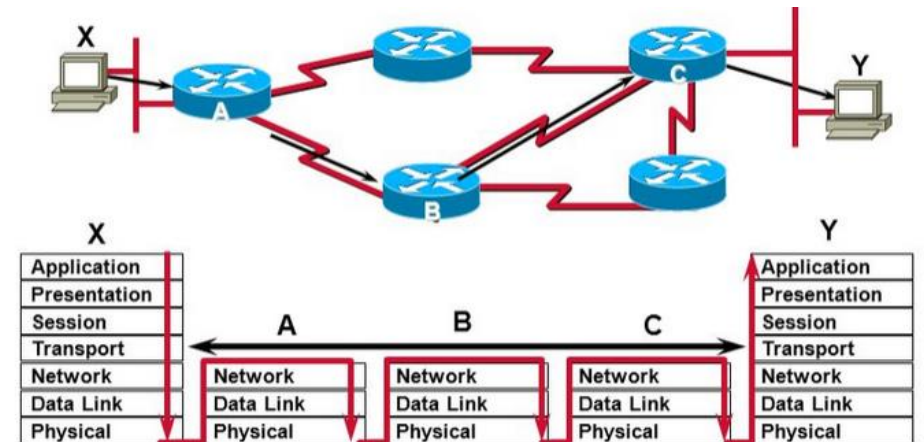

- The most familiar type of IP routers are home and small office routers that simply forward IP packets between the home computers and the Internet.



# Routing

## What is a router?


- When sending data to a remote destination, a host passes datagrams to a local router. The router forwards the datagrams toward the final destination.
- They travel from one router to another until they reach a router connected to the destination's LAN segment.



# Routing

## Routing table

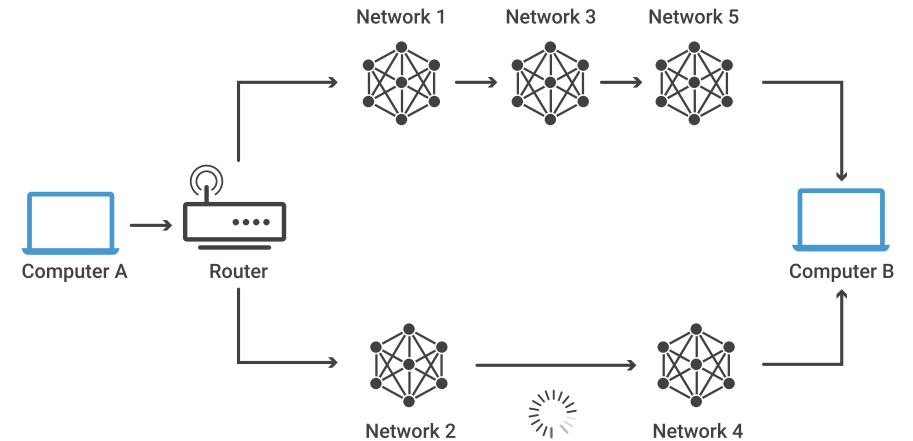
- is a data table stored in a router or a network host that lists the routes to particular network destinations, and in some cases, metrics (distances) associated with those routes. The routing table contains information about the topology of the network immediately around it.
- The construction of routing tables is the primary goal of **routing protocols**.



# Routing

## Routing table

- A routing table is a data file in RAM that is used to store route information about directly connected and remote networks.
- Nodes can also share the contents of their routing table with other nodes.
- With **hop-by-hop** routing, each routing table lists, for all reachable destinations, the address of the next device along the path to that destination: the **next hop**.

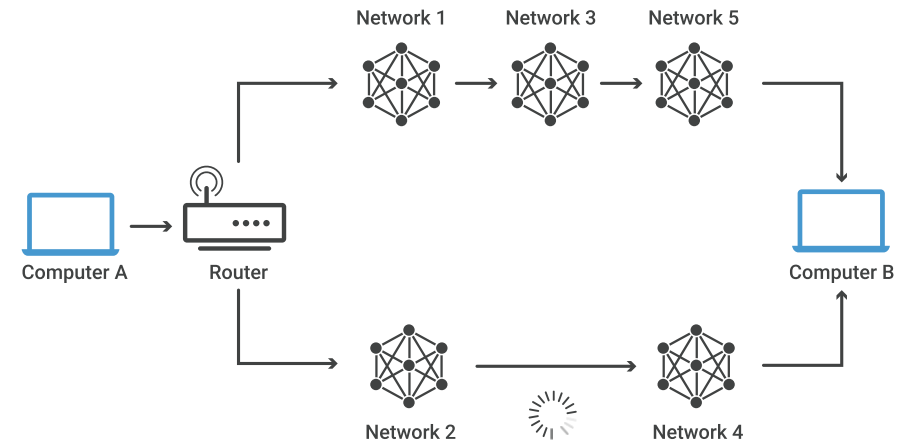




# Routing

## Routing table


- When a router interface is configured with an IP address and subnet mask, the interface becomes a host on that attached network.
- A directly connected network is a network that is directly attached to one of the router interfaces.
- The network address and subnet mask of the interface, along with the interface type and number, are entered into the routing table as a directly connected network.



# Routing

## Routing table

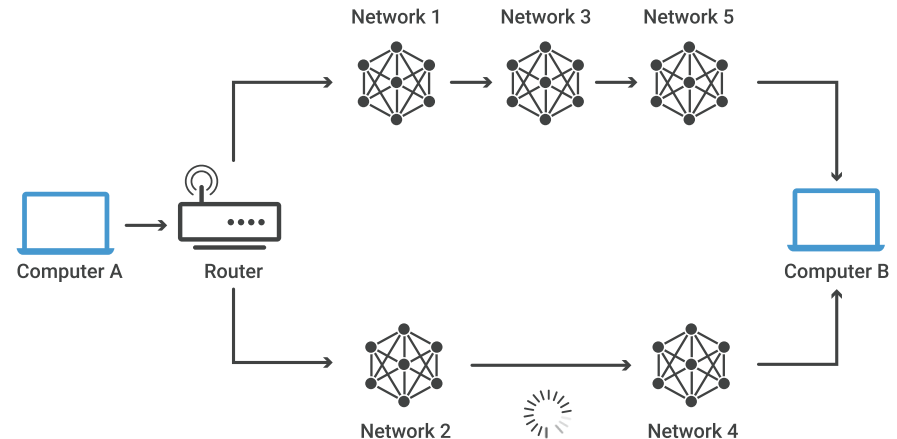
- A remote network is a network that can only be reached by sending the packet to another router.
- Routing table entries to remote networks may be either **dynamic** or **static**.
- Dynamic routes are routes to remote networks that were learned automatically by the router through a **dynamic routing protocol**.
- Static routes are routes that a **network administrator** manually configured.



# Routing

## Routing table

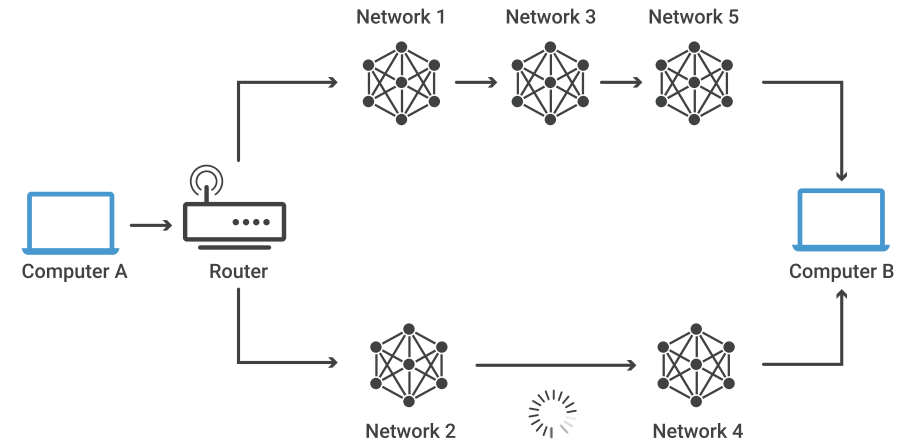
- The routing table consists of at least **three** information fields:
  1. **network identifier:** The destination subnet and netmask
  2. **metric:** The routing metric of the path through which the packet is to be sent. The route will go in the direction of the gateway with the lowest metric.
  3. **next hop:** The next hop, or gateway, is the address of the next station to which the packet is to be sent on the way to its final destination



# Routing

## Routing table

- Router metrics are configuration values used by a router to make routing decisions.
- Router metrics help the router choose the best route among multiple feasible routes to a destination. The route will go in the direction of the gateway with the **lowest** metric.
- A router metric is typically based on information such as path length, bandwidth, load, delay, etc.



# Routing

## Routing table

- Destination 192.168.0.0 and netmask 255.255.255.0 can be written as 192.168.0.0/24.
- The Gateway column contains the same information as the **Next hop**, i.e. it points to the gateway through which the network can be reached.
- The Interface indicates what locally available interface is responsible for reaching the gateway. In this example, gateway 192.168.0.1 (the internet router) can be reached through the local network card with address 192.168.0.100.

Network destination	Netmask	Gateway	Interface	Metric
0.0.0.0	0.0.0.0	192.168.0.1	192.168.0.100	10
127.0.0.0	255.0.0.0	127.0.0.1	127.0.0.1	1
192.168.0.0	255.255.255.0	192.168.0.100	192.168.0.100	10
192.168.0.100	255.255.255.255	127.0.0.1	127.0.0.1	10
192.168.0.1	255.255.255.255	192.168.0.100	192.168.0.100	10

# Routing

## Routing table

- Destination 192.168.0.0 and netmask 255.255.255.0 can be written as 192.168.0.0/24.
- The Gateway column contains the same information as the **Next hop**, i.e. it points to the gateway through which the network can be reached.
- The Interface indicates what locally available interface is responsible for reaching the gateway. In this example, gateway 192.168.0.1 (the internet router) can be reached through the local network card with address 192.168.0.100.

Network destination	Netmask	Gateway	Interface	Metric
0.0.0.0	0.0.0.0	192.168.0.1	192.168.0.100	10
127.0.0.0	255.0.0.0	127.0.0.1	127.0.0.1	1
192.168.0.0	255.255.255.0	192.168.0.100	192.168.0.100	10

# Routing

## Routing protocol

- A routing protocol specifies how routers communicate with each other to distribute information that enables them to select paths between nodes on a computer network.
- Data packets are forwarded through the networks of the internet from router to router until they reach their destination computer.
- Each router has a prior knowledge only of networks attached to it directly.
- A routing protocol shares this information first among immediate neighbors, and then throughout the network.

Network destination	Netmask	Gateway	Interface	Metric
0.0.0.0	0.0.0.0	192.168.0.1	192.168.0.100	10
127.0.0.0	255.0.0.0	127.0.0.1	127.0.0.1	1
192.168.0.0	255.255.255.0	192.168.0.100	192.168.0.100	10

# Routing

## Routing protocol

- This way, routers gain knowledge of the topology of the network.
- The specific characteristics of routing protocols include the manner in which they avoid routing loops, the manner in which they select preferred routes, using information about hop costs, etc.

Network destination	Netmask	Gateway	Interface	Metric
0.0.0.0	0.0.0.0	192.168.0.1	192.168.0.100	10
127.0.0.0	255.0.0.0	127.0.0.1	127.0.0.1	1
192.168.0.0	255.255.255.0	192.168.0.100	192.168.0.100	10



# Routing

## Routing protocol

- Two major classes are in widespread use on IP networks:
  - **Interior Gateway Protocols (IGPs):** Interior Gateway Protocols allow routers to exchange information within an autonomous system (AS). Examples of these protocols are Open Short Path First (OSPF) and Routing Information Protocol (RIP).
  - **Exterior Gateway Protocols (EGPs):** Exterior Gateway Protocols allow the exchange of summary information between autonomous systems. An example of this type of routing protocol is Border Gateway Protocol (BGP).

# Routing

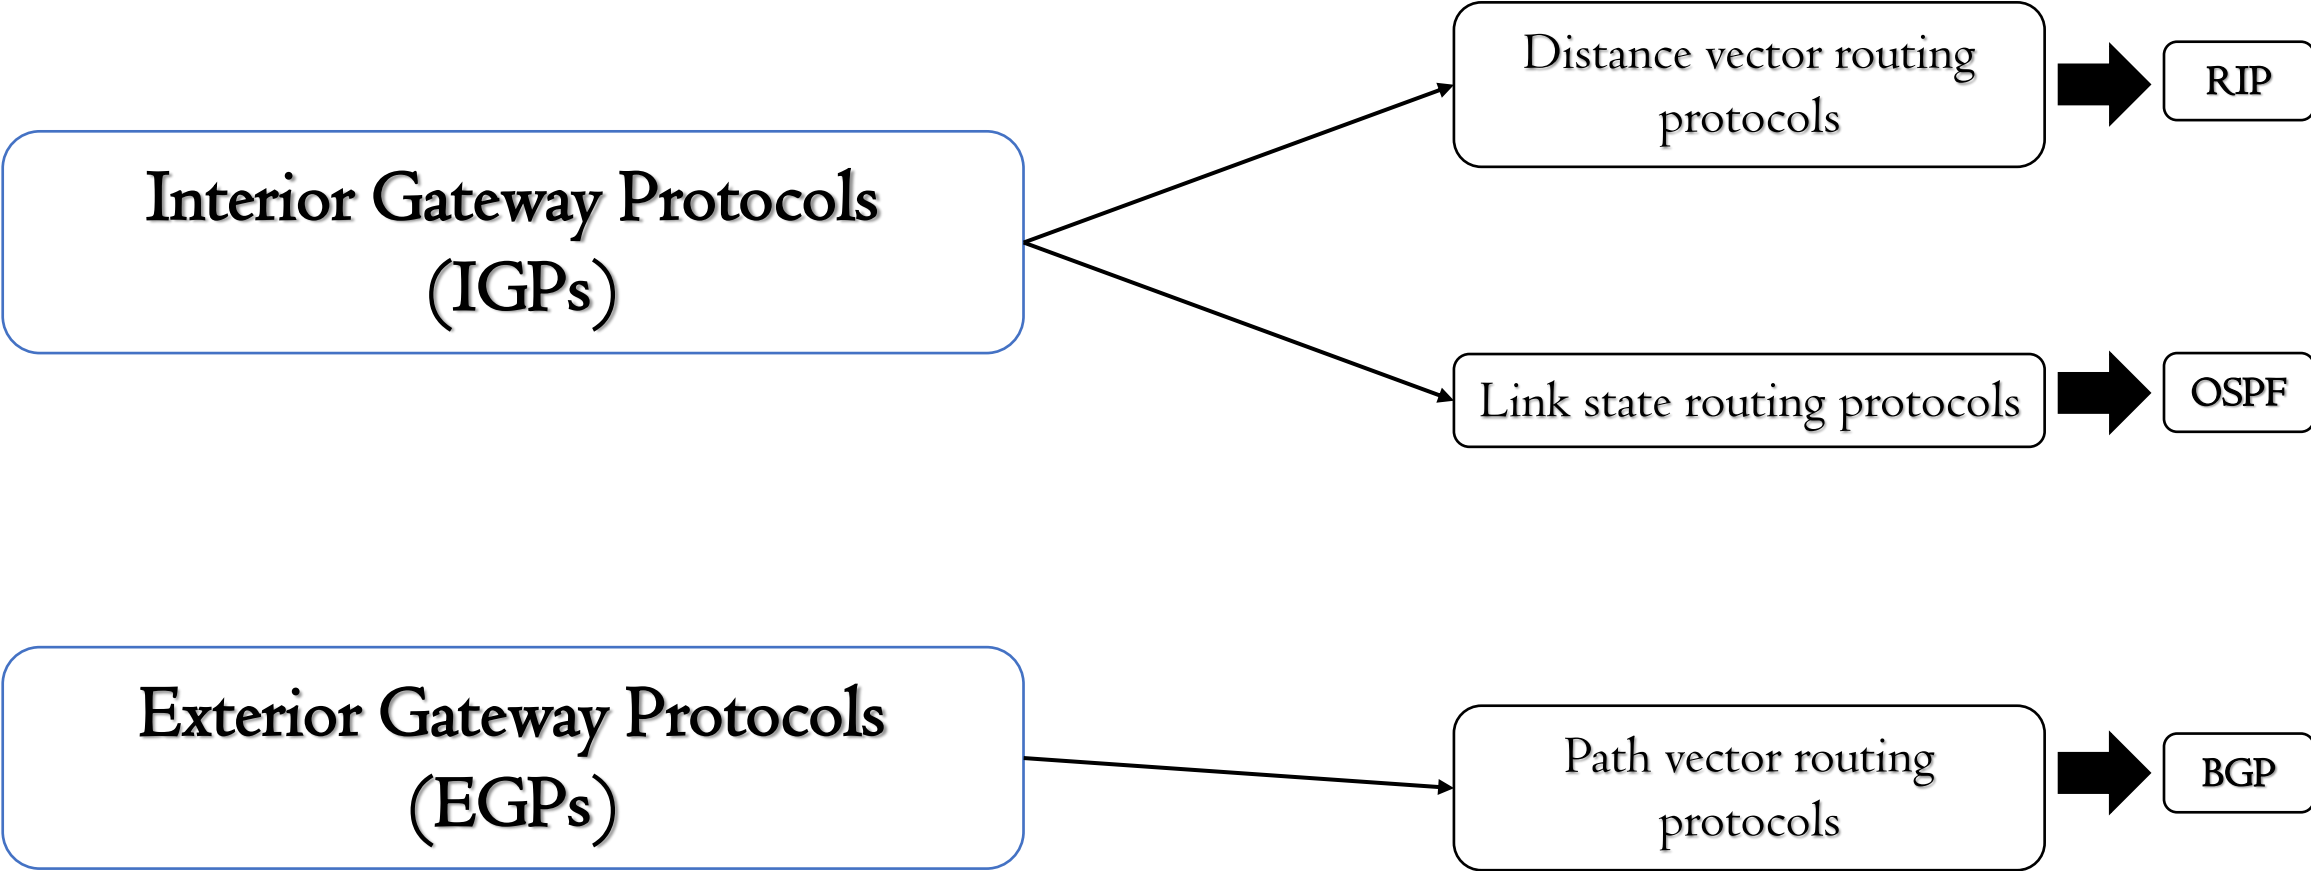
## Routing protocol

- Two major classes are in widespread use on IP networks:
  - **Interior Gateway Protocols (IGPs):** Interior Gateway Protocols allow routers to exchange information within an autonomous system (AS). Examples of these protocols are Open Short Path First (OSPF) and Routing Information Protocol (RIP).
  - **Exterior Gateway Protocols (EGPs):** Exterior Gateway Protocols allow the exchange of summary information between autonomous systems. An example of this type of routing protocol is Border Gateway Protocol (BGP).



An **AS** normally consists of an internetwork within an organization. It is administered by a single management authority (typically an Internet service provider (**ISP**)).

# Routing



# Routing

## Distance vector routing protocols

- Distance vector algorithms are examples of dynamic routing protocols. These algorithms allow each device in the network to automatically build and maintain a local IP routing table.
- The principle behind distance vector routing is simple. Each router in the internetwork maintains the distance or cost from itself to every known destination.
- The path represented by the smallest cost becomes the preferred path to reach the destination.

# Routing

## Distance vector routing protocols

- This information is maintained in a distance vector table. The table is periodically advertised to each neighboring router. Each router processes these advertisements to determine the best paths through the network.
- The main advantage of distance vector algorithms is that they are typically easy to implement and debug, especially in small networks,
- RIP is a popular example of a distance vector routing protocol

# Routing

## Routing Information Protocol (RIP)

- RIP is an example of an interior gateway protocol designed for use within small autonomous systems.
- The RIP protocol specifies two packet types. These packets can be sent by any device running the RIP protocol:
  - **Request packets:** A request packet queries neighboring RIP devices to obtain their distance vector table.
  - **Response packets:** A response packet is sent by a device to advertise the information maintained in its local distance vector table.

# Routing

## Routing Information Protocol (RIP)

- The table is automatically sent every 30 seconds.
- RIP sends and receives datagrams using UDP port 520.
- RIP datagrams have a maximum size of 512 octets.
- In LAN environments, RIP datagrams are sent using the MAC all-stations broadcast address and an IP network broadcast address.

# Routing

## Routing Information Protocol (RIP)

### Calculating distance vectors

- The distance vector table describes each destination network. The entries in this table contain the following information:
  - The destination network (vector) described by this entry in the table.
  - The associated cost (distance) of the most attractive path to reach this destination. This provides the ability to differentiate between multiple paths to a destination.
  - The IP address of the next-hop device used to reach the destination network.



# Routing

## Routing Information Protocol (RIP)

### Calculating distance vectors


- Each time a routing table advertisement is received by a device, it is processed to determine if any destination can be reached by a lower cost path.
- This is done using the RIP distance vector algorithm, which can be summarized as:
  - At router initialization, each device contains a distance vector table listing each directly attached networks and configured cost. Typically, each network is assigned a cost of 1. This represents a single hop through the network. The total number of hops in a route is equal to the total cost of the route. However, cost can be changed to reflect other measurements such as utilization, speed, or reliability.

# Routing

## Routing Information Protocol (RIP)

### Calculating distance vectors

- The IP routing table is updated to reflect the least-cost path to each destination



# Routing

## Routing Information Protocol (RIP)

### Calculating distance vectors

Router R2 Distance Vector Table

Net	Next hop	Metric
N1	R1	2
N2	Direct	1
N3	Direct	1
N4	R3	2
N5	R3	3
N6	R3	4

Router R3 Distance Vector Table

Net	Next hop	Metric
N1	R2	3
N2	R2	2
N3	Direct	1
N4	Direct	1
N5	R4	2
N6	R4	3


Router R4 Distance Vector Table

Net	Next hop	Metric
N1	R3	4
N2	R3	3
N3	R3	2
N4	Direct	1
N5	Direct	1
N6	R5	2

# Routing

## Routing Information Protocol (RIP)


In this network we have 4 routers A, B, C and D:



# Routing

## Routing Information Protocol (RIP)


- We mark the current time (or iteration) in the algorithm with  $T$ , and begin (at time 0, or  $T=0$ ) by creating distance matrices for each router to its immediate neighbours.
- As we build the routing tables below, the shortest path is highlighted in green, and a new shortest path is highlighted in yellow.
- Grey columns indicate nodes that are not neighbors of the current node, and are therefore not considered as a valid direction in its table. Red indicates invalid entries in the table since they refer to distances from a node to itself, or via itself.



# Routing

## Routing Information Protocol (RIP)

- We mark the current time (or iteration) in the algorithm with  $T$ , and begin (at time 0, or  $T=0$ ) by creating distance matrices for each router to its immediate neighbours.
- As we build the routing tables below, the shortest path is highlighted in green, and a new shortest path is highlighted in yellow.
- Grey columns indicate nodes that are not neighbors of the current node, and are therefore not considered as a valid direction in its table. Red indicates invalid entries in the table since they refer to distances from a node to itself, or via itself.



# Routing

## Routing Information Protocol (RIP)

T=0

	from A	via A	via B	via C	via D	from B	via A	via B	via C	via D	from C	via A	via B	via C	via D	from D	via A	via B	via C	via D
to A						3					23									
to B			3									2								
to C				23					2										5	
to D														5						

- At this point, all the routers (A,B,C,D) have new "shortest-paths" for their DV (the list of distances that are from them to another router via a neighbor). They each broadcast this new DV to all their neighbors: A to B and C, B to C and A, C to A, B, and D, and D to C. As each of these neighbors receives this information, they now recalculate the shortest path using it. For example: A receives a DV from C that tells A there is a path via C to D, with a distance (or cost) of 5. Since the current "shortest-path" to C is 23, then A knows it has a path to D that costs  $23+5=28$ . As there are no other shorter paths that A knows about, it puts this as its current estimate for the shortest-path from itself (A) to D, via C.

# Routing

## Routing Information Protocol (RIP)

	from	via	via	via	via	from	via	via	via	via	from	via	via	via	via	from	via	via	via	via
	A	A	B	C	D	B	A	B	C	D	C	A	B	C	D	D	A	B	C	D
T=1	to A					to A	3		25		to A	23	5			to A			28	
	to B		3	25		to B					to B	26	2			to B			7	
	to C		5	23		to C	26		2		to C					to C			5	
	to D			28		to D			7		to D				5	to D				

- Again, all the routers have gained in the last iteration (at T=1) new "shortest-paths", so they all broadcast their DVs to their neighbors; This prompts each neighbor to re-calculate their shortest distances again.
- For instance: A receives a DV from B that tells A there is a path via B to D, with a distance (or cost) of 7. Since the current "shortest-path" to B is 3, then A knows it has a path to D that costs  $7+3=10$ . This path to D of length 10 (via B) is shorter than the existing "shortest-path" to D of length 28 (via C), so it becomes the new "shortest-path" to D.



# Routing

## Routing Information Protocol (RIP)

T=2

from	via	via	via	via	from	via	via	via	via	from	via	via	via	via	from	via	via	via	via
A	A	B	C	D	B	A	B	C	D	C	A	B	C	D	D	A	B	C	D
to A					to A	3		7		to A	23	5		33	to A			10	
to B		3	25		to B					to B	26	2		12	to B			7	
to C		5	23		to C	8		2		to C					to C			5	
to D		10	28		to D	31		7		to D	51	9		5	to D				

- This time, only routers A and D have new shortest-paths for their DVs. So they broadcast their new DVs to their neighbors: A broadcasts to B and C, and D broadcasts to C. This causes each of the neighbors receiving the new DVs to re-calculate their shortest paths. However, since the information from the DVs doesn't yield any shorter paths than they already have in their routing tables, then there are no changes to the routing tables.

# Routing

## Routing Information Protocol (RIP)

T=3

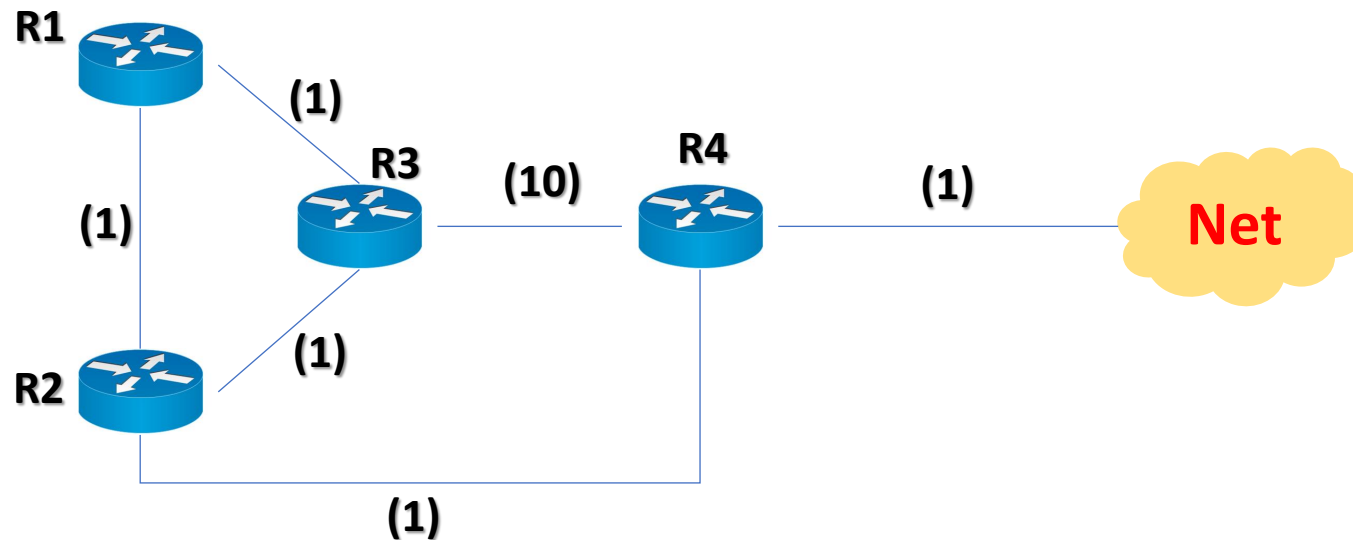
from	via	via	via	via	from	via	via	via	via	from	via	via	via	via	from	via	via	via	via
A	A	B	C	D	B	A	B	C	D	C	A	B	C	D	D	A	B	C	D
to A					to A	3		7		to A	23	5		15	to A			10	
to B		3	25		to B					to B	26	2		12	to B			7	
to C		5	23		to C	8		2		to C					to C			5	
to D		10	28		to D	13		7		to D	33	9		5	to D				

- None of the routers have any new shortest-paths to broadcast. Therefore, none of the routers receive any new information that might change their routing tables. The algorithm comes to a stop.

# Routing Information Protocol (RIP)

## Convergence and counting to infinity


- Given sufficient time, this algorithm will correctly calculate the distance vector table on each device.
- However, during this convergence time, erroneous routes may propagate through the network.



# Routing Information Protocol (RIP)

## Convergence and counting to infinity

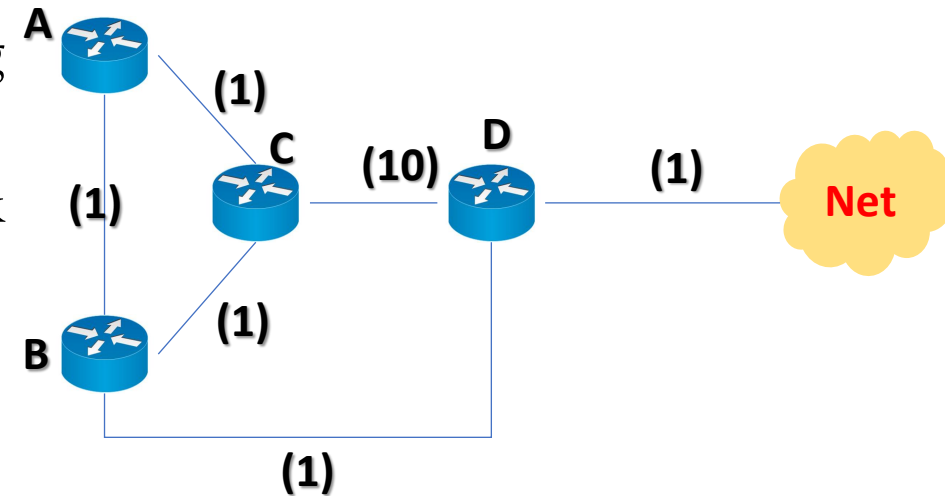
- After the network has converged, each device has routing information describing all networks.
- For example, to reach the target network, the routers have the following information:
  - Router D to the Net: Directly connected network. Metric is 1.
  - Router B to the Net: Next hop is router D. Metric is 2.
  - Router C to the Net: Next hop is router B. Metric is 3.
  - Router A to the Net: Next hop is router B. Metric is 3.



# Routing Information Protocol (RIP)

## Convergence and counting to infinity

- Consider an adverse condition where the link connecting router B and router D fails.
- After the network has reconverged, all routes use the link connecting router C and router D to reach the target network.
- However, this reconvergence time can be considerable.




# Routing Information Protocol (RIP)

## Convergence and counting to infinity

Time → →

D:	Direct	1	Direct	1	Direct	1	Direct	1	--	Direct	1	Direct	1
B:	Unreachable		C	4	C	5	C	6		C	11	C	12
C:	B	3	A	4	A	5	A	6		A	11	D	11
A:	B	3	C	4	C	5	C	6	--	C	11	C	12




- The figure illustrates how the routes to the target network are updated throughout the reconvergence period.

# Routing Information Protocol (RIP)

## Convergence and counting to infinity

Time → →

D:	Direct	1	Direct	1	Direct	1	Direct	1	--	Direct	1	Direct	1
B:	Unreachable		C	4	C	5	C	6		C	11	C	12
C:	B	3	A	4	A	5	A	6		A	11	D	11
A:	B	3	C	4	C	5	C	6	--	C	11	C	12




- Reconvergence begins when router B notices that the route to router D is unavailable.
- Router B is able to **immediately** remove the failed route because the link has timed out.
- **However**, a considerable amount of time passes before the other routers remove their references to the failed route.

# Routing Information Protocol (RIP)

## Convergence and counting to infinity

Time → →

D:	Direct	1	Direct	1	Direct	1	Direct	1	--	Direct	1	Direct	1
B:	Unreachable		C	4	C	5	C	6		C	11	C	12
C:	B	3	A	4	A	5	A	6		A	11	D	11
A:	B	3	C	4	C	5	C	6	--	C	11	C	12




Sequence of updates shown in the figure :

1. Prior to the adverse condition occurring, router A and router C have a route to the target network through router B.
2. The adverse condition occurs when the link connecting router D and router B fails. Router B recognizes that its preferred path to the target network is now invalid.

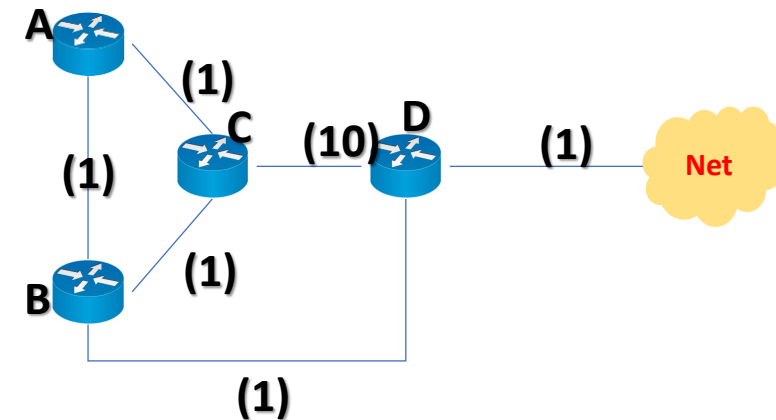


# Routing Information Protocol (RIP)

## Convergence and counting to infinity

Time 

D:	Direct	1	Direct	1	Direct	1	Direct	1	--	Direct	1	Direct	1
B:	Unreachable		C	4	C	5	C	6		C	11	C	12
C:	B	3	A	4	A	5	A	6		A	11	D	11
A:	B	3	C	4	C	5	C	6	--	C	11	C	12



Sequence of updates shown in the figure :

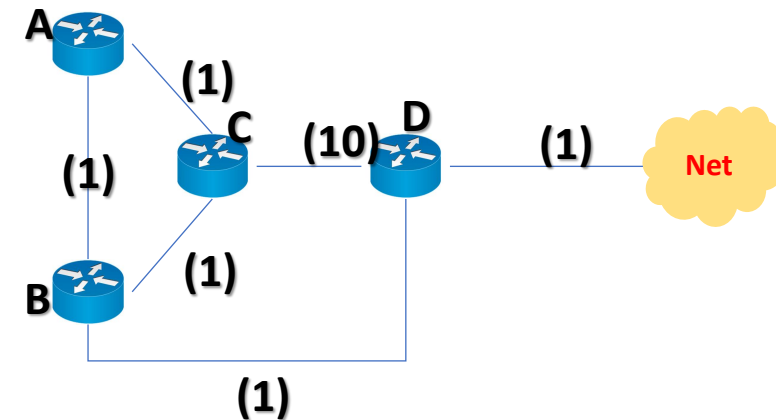
- Router A and router C continue to send updates reflecting the route through router B. This route is actually invalid because the link connecting router D and router B has failed.

# Routing Information Protocol (RIP)

## Convergence and counting to infinity

Time → →

D:	Direct	1	Direct	1	Direct	1	Direct	1	--	Direct	1	Direct	1
B:	Unreachable		C	4	C	5	C	6		C	11	C	12
C:	B	3	A	4	A	5	A	6		A	11	D	11
A:	B	3	C	4	C	5	C	6	--	C	11	C	12




Sequence of updates shown in the figure :


- Router B receives the updates from router A and router C. Router B believes it should now route traffic to the target network through either router A or router C. In reality, this is not a valid route, because the routes in router A and router C are vestiges of the previous route through router B.

# Routing Information Protocol (RIP)

## Convergence and counting to infinity

Time 

D:	Direct	1	Direct	1	Direct	1	Direct	1	--	Direct	1	Direct	1
B:	Unreachable		C	4	C	5	C	6		C	11	C	12
C:	B	3	A	4	A	5	A	6		A	11	D	11
A:	B	3	C	4	C	5	C	6	--	C	11	C	12




Sequence of updates shown in the figure :


- Using the routing advertisement sent by router B, router A and router C are able to determine that the route through router B has failed. However, router A and router C now believe the preferred route exists through the partner.

# Routing Information Protocol (RIP)

## Convergence and counting to infinity

Time 

D:	Direct	1	Direct	1	Direct	1	Direct	1	--	Direct	1	Direct	1
B:	Unreachable		C	4	C	5	C	6		C	11	C	12
C:	B	3	A	4	A	5	A	6		A	11	D	11
A:	B	3	C	4	C	5	C	6	--	C	11	C	12



Sequence of updates shown in the figure :


- Using the routing advertisement sent by router B, router A and router C are able to determine that the route through router B has failed. However, router A and router C now believe the preferred route exists through the partner.

# Routing Information Protocol (RIP)

## Convergence and counting to infinity

Time → →

D:	Direct	1	Direct	1	Direct	1	Direct	1	--	Direct	1	Direct	1
B:	Unreachable		C	4	C	5	C	6		C	11	C	12
C:	B	3	A	4	A	5	A	6		A	11	D	11
A:	B	3	C	4	C	5	C	6	--	C	11	C	12




- Network convergence continues as router A and router C engage in an extended period of **mutual deception**. Each device claims to be able to reach the target network through the partner device. The path to reach the target network now contains **a routing loop**.

# Routing Information Protocol (RIP)

## Convergence and counting to infinity

Time → →


D:	Direct	1	Direct	1	Direct	1	Direct	1	--	Direct	1	Direct	1
B:	Unreachable		C	4	C	5	C	6		C	11	C	12
C:	B	3	A	4	A	5	A	6		A	11	D	11
A:	B	3	C	4	C	5	C	6	--	C	11	C	12



- The manner in which the costs in the distance vector table increment gives rise to the term **counting to infinity**.
- The costs continues to increment, theoretically to infinity.

# Routing Information Protocol (RIP)

## Convergence and counting to infinity



- Imagine a subnet connected like A–B–C–D–E–F, and let the metric between the routers be "number of jumps".
- Now suppose that A is taken offline.
- In the vector-update-process B notices that the route to A, which was distance 1, is down – B does not receive the vector update from A.
- **The problem is**, B also gets an update from C, and C is **still not aware** of the fact that A is down – so it tells B that A is only two jumps from C (C to B to A).
- Since B doesn't know that the path from C to A is through itself (B), it updates its table with the new value "B to A = 2 + 1".
- Later on, B forwards the update to C and due to the fact that A is reachable through B (From C's point of view), C decides to update its table to "C to A = 3 + 1".
- This **slowly** propagates through the network until it becomes **infinity**

# Routing Information Protocol (RIP)

## Convergence and counting to infinity

- To minimize this exposure, whenever a network is unavailable, the incrementing of metrics through routing updates must be halted as soon as it is practical to do so.
- In a RIP environment, costs continue to increment until they reach **a maximum value of 16**.
- In a RIP environment, any path exceeding 15 hops is considered invalid. The routing algorithm will discard these paths.
- **A side effect** of the metric limit is that it also limits the number of hops a packet can traverse from source network to destination network.
- The excessive convergence time caused by counting to infinity can be reduced with the use of **split horizon** or **poison reverse**.




# Routing Information Protocol (RIP)

## Split horizon

- This rule dictates that routing information is **prevented** from exiting the router on an interface through which the information was received.
- Never send routing information back in the direction from which it was received.

### Example:

- The following figure shows an example of three routers used to forward packets between networks. In this simple design, the R3 router broadcasts routing information about the 10.0.0.0/16 network to the R2 router.




# Routing Information Protocol (RIP)

## Split horizon

### Example:

- The R2 router receives this information, updates its routing table and broadcasts the information to the R1 router. When the R1 router receives this information, it updates its own routing table.
- The R1 router can use the updated routing information to send packets to the 10.0.0.0/16 network via the R2 and R3 routers.




# Routing Information Protocol (RIP)

## Split horizon

### Example:

- If split horizon is **enabled**, the R1 router **will not** be able to advertise this network route **back** to the R2 router.




# Routing Information Protocol (RIP)

## Split horizon

### Example:

- If split horizon is **enabled**, the R1 router **will not** be able to advertise this network route **back** to the R2 router.
- If no split horizon is configured on the router, the R1 router will broadcast the route back to the R2 router, which will update its routing table **to include** the network route available through the R1 router.




# Routing Information Protocol (RIP)

## Split horizon

### Example:

- Under normal operations, the inclusion of the R1 route in the R2 routing table is not an issue because it's clearly a much costlier route than a direct R2-to-R3 connection.
- However, if the R2-to-R3 link fails and the R2 router receives a packet from R1 that's destined for the 10.0.0.0/16 network, the R2 router will send the packet back to R1 because the router advertised a viable network path.
- But the R1 router will simply return the packet to the R2 router based on its own routing information, resulting in a **routing loop** that continues until the packet expires.
- With split horizon in place, the R1 router will not advertise the network route to the R2 router, preventing the routing loop from occurring.



# Routing Information Protocol (RIP)


## Split horizon

Time → →

D:	Direct	1	Direct	1	Direct	1	Direct	1	--	Direct	1	Direct	1
B:	Unreachable		C	4	C	5	C	6		C	11	C	12
C:	B	3	A	4	A	5	A	6		A	11	D	11
A:	B	3	C	4	C	5	C	6	--	C	11	C	12

Time → →


D	Direct		1	Direct		1	Direct		1	Direct		1
B	Unreachable			Unreachable			Unreachable			C		12
C	B		3	A		4	D		11	D		11
A	B		3	C		4	Unreachable			C		12



# Routing Information Protocol (RIP)

## Poison reverse

- This is the equivalent of route poisoning all possible reverse paths -- that is, informing all routers that the path back to the originating node for a particular packet has an infinite metric.
- For example, in the figure below, the route will be advertised to R2 **but** with the distance marked as infinite, indicating that the network is unreachable via this route.
- For RIP, the metric value is set to 16, which is equivalent to infinity because the maximum RIP network hop count is 15.



# Link State routing protocols

## Poison reverse

- This is the equivalent of route poisoning all possible reverse paths -- that is, informing all routers that the path back to the originating node for a particular packet has an infinite metric.
- For example, in the figure below, the route will be advertised to R2 **but** with the distance marked as infinite, indicating that the network is unreachable via this route.
- For RIP, the metric value is set to 16, which is equivalent to infinity because the maximum RIP network hop count is 15.

