

Unit 5

Intradomain and Interdomain Routing Protocols

Overview

Description

The basic function of the router in a network is to determine the best route for a packet to travel across a network from its source to its destination. Routers accomplish this task by storing information in tables and referring to these tables to determine the best path. Various routing protocols have been developed to assist the router in mapping the network and create routing tables.

Some routing protocols operate only within domains, a collection of hosts and routers that function as a group, and are referred to as intradomain routing protocols. Other routing protocols operate within and between domains and are referred to as interdomain routing protocols.

- Lesson one will discuss an intradomain routing protocol, the Routing Information Protocol (RIP), and the Bellman-Ford algorithm underlying RIP.
- Lesson two will discuss the Open Shortest Path First (OSPF) protocol, which is also an intradomain routing protocol.
- Lesson three covers the interdomain routing protocol, the Border Gateway Protocol (BGP).

Unit Table of Contents

This unit contains the following three lessons:

Lessons	Pages	Length
Lesson 5-1: Routing Information Protocol	184-202	5 hours
Lesson 5-2: Open Shortest Path First Protocol	203-224	5 hours
Lesson 5-3: Border Gateway Protocol	225-236	5 hours

Lesson 5-1: Routing Information Protocol

At a Glance



The Routing Information Protocol (RIP) is an intradomain routing protocol that uses the Bellman-Ford Distance Vector Algorithm to build a routing table of the entire network. A router that uses RIP to build its routing table:

- Requests routing information from neighboring routers.
- Responds to requests from neighboring routers.
- Advertise across the network to inform other routers of its configuration.

What You Will Learn

After completing this lesson, you will be able to do the following:

- Identify the steps of the Bellman-Ford Distance Vector Algorithm
- Diagram the counting to infinity problem.
- Describe how RIP resolves the counting to infinity problem
- Explain how split horizon attempts to prevent the counting to infinity problem.

Tech Talk



- **Adjacent Routers**—Refers to routers that are next to each other sharing the same wire in a network.
- **Advertising**—Routing process where routers send updates at specific intervals so that other routers on the network can maintain current routing tables.
- **Domain**—A collection of hosts and routers that function as a group.
- **Drop (a route)**—To remove a route from a routing table.
- **Hop**—The transmission of a packet through a node, e.g. router.
- **Hop Count**—The routing metric used by RIP to measure the distance between a source and a destination on a network.
- **Host**—A computer system on a network. It does not include routers.
- **Intradomain Routing Protocols**—Routing protocols that operate only within domains.
- **Metric**—A value, for example, hop count or bandwidth, used by a routing algorithm to determine the best route for transmitting a packet across the network.
- **Node**—Another name for a network device, for example, workstations and routers.
- **Routing Table**—A table stored in a router that is used to keep track of specific routes between nodes.
- **Unicasting**—One-to-one transmissions.

Distance-Vector Routing


Distance-vector routing, also known as the Fulkerson, or Bellman-Ford routing algorithm, after the researchers that first proposed it, is the oldest and simplest class of routing protocols. The most common implementation of a distance-vector protocol is RIP, the Routing Information Protocol, which may still be the most common dynamic routing protocol of any kind. Like TCP/IP its popularity stems not so much from technical superiority as from availability. It was freely available in the research institutions that formed the early internetworked community. RIP was actually used for many years before it was ever documented or standardized.

Bellman-Ford Algorithm

The chief advantage of distance-vector protocols, such as RIP, is their simplicity and ease of implementation. To understand how distance-vector routing protocols work, consider the fictitious company, Acme Delivery Service as an example.

Acme offices operate vans that deliver packages. An office is only allowed to run its vans on streets that actually touch the office parking lot. The office in Ripville, for instance, delivers packages only to businesses along 1st, 2nd, and 3rd Avenues.

Acme Delivery Service: Ripville




The clever manager of the Ripville office realizes, one day, that she could expand her business by collaborating with the office in Ford City. Fortunately, that office is right on 1st Street, in Ford City, so the company rules allow her to deliver packages there. She sends a memo to the manager at the Ford City office, which reads as follows:

"I can ship packages to businesses on 1st Street, and 2nd and 3rd Avenues. If you have any such packages, just ship them to me and I will deliver the packages for you."

The manager of the Ford City office immediately recognizes this as a great idea. He begins to advertise to his customers that he can now ship, not only to 1st Street, and 4th and 5th Avenues, but to 2nd and 3rd Avenues as well. He does not need the help of the Ripville manager in reaching 1st Street, because he can already deliver there, according to the company rules.

Ripville and Ford City Offices Add Each Other's Route Information



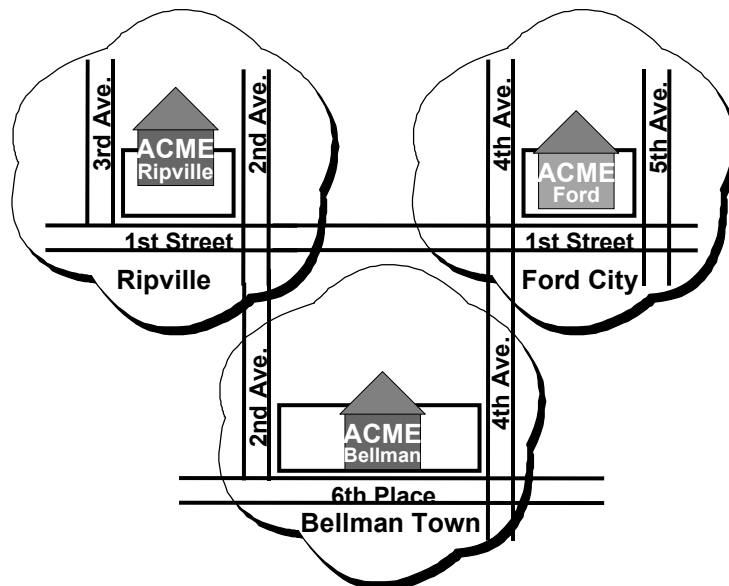
The Ford City manager decides to continue the expansion. He replies, not only to his colleague in Ripville, but to another colleague in Bellmantown.

He decides that he can send the same message to both other managers. The message looks like this:

"I can ship packages to businesses on 1st Street, and 4th and 5th Avenues. I can also ship packages to 2nd and 3rd Streets by sending them to my colleague in Ripville. If you have any such packages, just ship them to me."

When the Bellmantown manager receives this message from Ford City, she updates her own delivery information. She already can ship to 2nd Avenue, 4th Avenue, and 6th Place, without help. The rest of the routes, to 1st Street and 3rd and 5th Avenues are new, and she advertises them to her customers.

Bellmantown Office Adds Ripville and Ford City Route Information

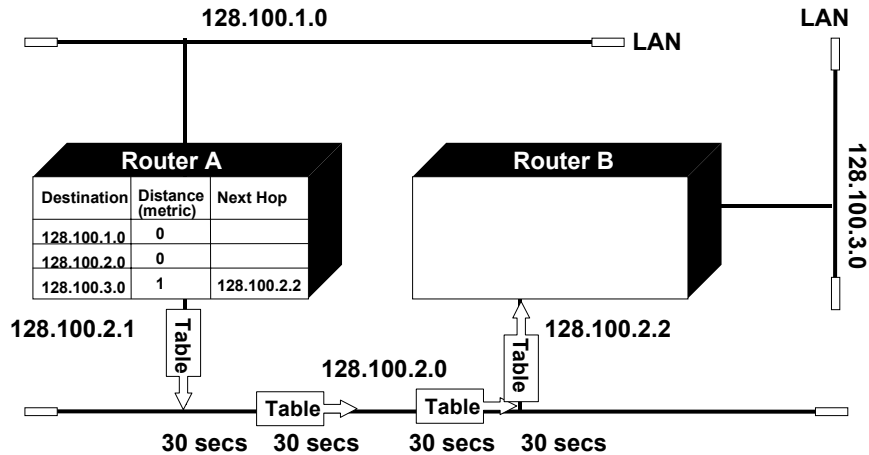


Soon all of the offices are exchanging messages, and every office knows how to get to any road that any other office knows about.

The previous example illustrates the Bellman-Ford Distance-vector algorithm. To use this algorithm, each router keeps a table, called the routing table. Each entry in the table contains three fields: a destination, a distance or hop count, and a next hop. The hop count is the routing metric which tracks how many hops a packet passes through from its source to its destination.

In addition to maintaining their own information, routers also exchange it with each other. The exchanged information is just the first two columns of the routing table. A new router table update is advertised or transmitted over the network every 30 seconds.


Router A Broadcasts Its Router Table



Each time a router receives new information from a neighbor, it updates its own routing table using the following algorithm:

- Get the next row from the table sent by the neighbor router.
- Add 1 to the distance field in the neighbor's row.
- Find a row in the local routing table with the same destination.
- If no such row exists, add the neighbor's row, and record the neighbor who sent the information as the next hop.
- If the row exists, compare the distance in the corresponding row in the local table to the distance from the row in the neighbor's table.
- If the neighbor's distance is less than the local distance replace the local row with the neighbor's row, and record the neighbor who sent the information as the next hop. Go on to the next row.
- If the neighbor's distance is greater or equal to the local distance ignore it, and go to the next row.

Router B Updates Its Router Table



The addition of the distance field is essential. It allows a router to choose between destinations it hears from two neighbors. In the Acme Delivery example, the Bellmantown manager heard about 3rd Avenue from the manager in Ford City. If all of the managers exchanged delivery information, though, she would hear about it from the Ripville manager as well. If she were concerned about getting her packages delivered as quickly as possible, she would send them to Ripville, instead of the long way around, through Ford City.

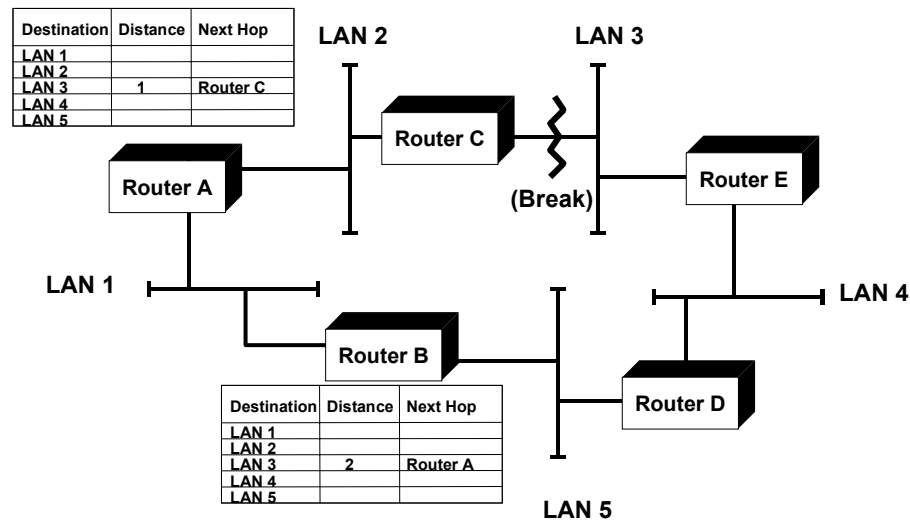
Check Your Understanding

- ◆ Describe as simply as possible how the Bellman-Ford algorithm works in the Routing Information Protocol.

Discovering Link Failure

The Bellman-Ford algorithm allows each router to discover the entire network. What happens, though, if a connection breaks? Suppose that, in the network pictured below, all of the routers have discovered the whole network, when suddenly the link between C and E breaks.

Link Failure Discovered



Router B periodically sends information that LAN 3 is 2 hops away, for packets headed to Router E, with Router A as the next hop. Router A's routing table displays LAN 3 as only 1 hop away and has C as the next hop for packets headed to Router E. The Bellman-Ford rules don't allow A to actively discover the working route. It must wait for updates from neighboring routers.

There are several solutions to this problem. Within RIP, routers set time limits and periodically drop or expire routes. RIP routers advertise their routes once every 30 seconds. When a neighbor hears a route advertised, it sets a timer. Unless the same neighbor advertises the exact route again before the timer reaches 180, it is removed from the routing table. In the example above, then, when Router C sends its sixth update with no mention of Router E, Router A will drop the route with the distance of 1, and include the route it now hears from B (with the distance of 2).

The developers of the algorithm, as a result of their repeated experiences, chose 30 seconds between announcements, and 180 seconds for dropping a route as the timeout intervals. If routing tables are exchanged more often, especially in a large network, they may start to interfere with real data. A route timeout of 180 seconds means that a few packets may be lost here

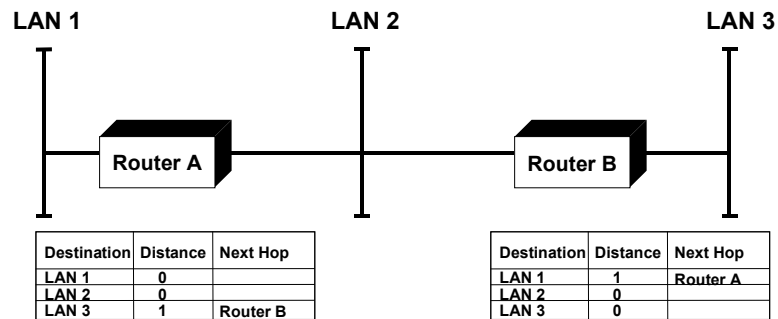
and there, without causing a perfectly useful link to vanish. Still, with hundreds, even thousands of packets a second traversing some network connections, 180 seconds is a long time to wait for a broken network to recover.

Counting to Infinity

A problem inherent in distance-vector protocols can make this situation even worse. It is called the counting to infinity problem.

In the network below, all the routers have exchanged and updated their routing tables to store the paths to each LAN. When the network is running, Router B hears from Router A that LAN 1 is one hop away.

Router A and B have Exchanged Routing Tables

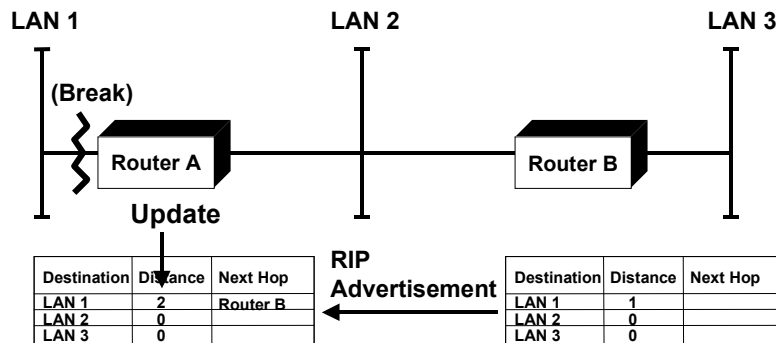


Suppose that the link from Router A to LAN 1 fails.

Router A immediately removes the route for LAN 1 from its routing table, when the hardware indicates a failure. Router A must then recalculate the distance to LAN 1. Unfortunately, in less than 30 seconds, Router B, who still has the route in its routing table, with a distance of 1, will advertise the route to Router A.

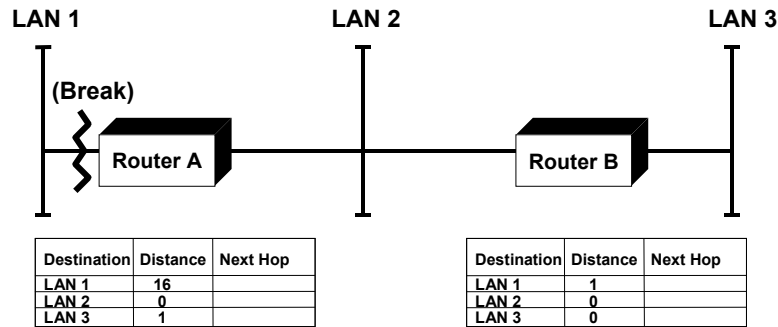
According to the rules, Router A will add LAN 1 to its routing table with a distance of 2, with Router B as the next hop.

Router A Updates its Routing Table



Soon, Router B will time out the route it heard from Router A, with the distance 1. As soon as it does so, though, it will hear Router A's updated advertisement with a distance 2. Router B will then add a new entry to its table with a distance of 3. This is called counting to infinity since the routers are caught in a loop that increases the distance by 1 each time a router updates its routing table.

This process could continue for a very long time. To prevent that, RIP chooses the number 16 to mean “unreachable.” Any route with a distance of 16 or more is dropped from the routing table, and not advertised.



Unreachable Routes Are Dropped

In the situation above, as soon as Router A or Router B discovers that the distance has climbed to 16, it will drop the route, not advertise it, and the cycle will stop. However, this increases the time it takes for the network to recover from a failure.

Split Horizon

There have been many attempts to fix the problems in distance-vector routing protocols. Split-horizon is an attempt to fix the counting-to-infinity problem. Split horizon is a modification of the original protocol rules. It adds a rule which requires that a router never send information learned through an interface, back out through that same interface.

In the case above, since Router B learned about LAN 1 through its interface on LAN 2, it can not re-advertise it there. Router A will never get the reflection, and the problem will never occur. Although this does prevent counting-to-infinity in the particular case described above, it does not eliminate the problem.

Poison Reverse

Poison reverse is an attempt to decrease the amount of time it takes RIP, to recover from a broken link. It gives a special meaning to advertising a route with an infinite (16 for RIP) distance. Such an advertisement indicates to the receiver that it should immediately remove the indicated route from its table, without waiting for it to timeout. In the example above, Router A might send such a message to Router B.

RIP 2

Routing Information Protocol 2 is the latest enhancement to the original RIP. Both protocols use hop counts as their metric for creating routing tables and both advertise the routing tables every 30 seconds.

RIP 2 differs from RIP 1 in that RIP 2 allows more information to be included in the message packets and an authentication mechanism was added.

Check Your Understanding

- ◆ Explain what is meant by “counting to infinity” and how the problem is resolved.

Try It Out

Design a RIP Game



Materials Needed:

- Windows 95 PC (optional)
- Pen/Pencil and Paper

Devise a team action game for 6-8 students based on the Routing Information Protocol. Document a simple set of rules and describe their relationship to RIP.

Rubric: Suggested Evaluation Criteria and Weightings

Criteria	%	Your Score
Simple, but clear set of rules.	25	
Creative and enticing game.	50	
Accurate relationship to RIP concepts.	25	
TOTAL	100	

Stretch Yourself

Explore Multicasting



This lesson has explored distance-vector routing protocols. Sending one-to-one transmissions is unicasting and has been the traditional method used by the Internet in past years. However, with the exponential increase in Internet traffic, unicasting is inefficient. Multicasting has been developed to efficiently transmit text, audio, and video over the Internet to multiple nodes at one time.

Materials Needed:

- Windows 95 PC (optional)
 - Internet Connection (optional)
 - Any Word Processor (e.g., MS Word)
 - Pen/Pencil and Paper
 - MS PowerPoint (optional)
1. Using whatever resources you find helpful, explore the concept of multicasting.
 2. Include in your research the Distance-Vector Multicasting Routing Protocol (DVMRP) and compare it to RIP. What are the basic differences? What advantages does one protocol have over the other?

3. Present your research in any one of the following methods:
 - a. Write a research report on your findings. Include in your report if you recommend using one protocol over the other and why. Document your resources.
 - b. Create a PowerPoint presentation demonstrating the similarities and differences in these concepts and protocols. Include in your presentation if you recommend using one protocol over the other and why. Document your resources.
 - c. Create a collage that demonstrates the similarities and differences in these concepts and protocols. If you recommend using one protocol over the other, include in your collage a conceptual representation of this recommendation. Document your resources.

Rubric: Suggested Evaluation Criteria and Weightings

Criteria	%	Your Score
Quality project suitable for presentation	30	
Analysis and synthesis of research	30	
Well thought out comparison	30	
Resources documented	10	
TOTAL	100	

Network Wizards

Counting to Infinity and Split Horizon



The discussion of Distance Vector routing brought up the "Counting to Infinity" problem that is a key problem with RIP. The lesson also suggested that there have been additions to the protocol, Split-Horizon, and Poison-Reverse among them, which attempt to address the problem. Recall that Split-Horizon is simply an additional rule in the RIP protocol, which states that a RIP router may not advertise a route that it hears through a given interface, back out to that same interface.

Materials Needed:

- Windows 95 PC (optional)
 - Any Word Processor (e.g., MS Word) (optional)
 - Pen/Pencil and Paper
1. Be certain that you understand the problem, and how Split-Horizon fixes it in the example. When you are sure of yourself, try to think of a network in which "Counting to Infinity" is still a problem, even when all RIP routers are using Split-Horizon. Think of a way that a router might hear its own advertisement through a different interface than the one it sent out.
 2. Diagram your network and explain why you think Split Horizon will not solve the problem. Exchange diagrams and conclusions with a classmate. Together, examine each other's diagrams and determine if you agree that your diagrams defy the Split Horizon solution.
 3. Contact a network manager in your area and ask if he/she would critique your diagrams and explanations as well. It is acceptable to make this exchange by mail, but you are encouraged to make an appointment to discuss your diagrams face-to-face.
 4. Report both your conclusions and the conclusions of the network professional.

Rubric: Suggested Evaluation Criteria and Weightings

Criteria	%	Your Score
Thoughtful analysis of the problem	50	
Quality critique of partner's diagram	25	
Thorough report on the conclusions of the network professional's critique	25	
TOTAL	100	

Summary

In this lesson, you learned the following:

- The steps of the Bellman-Ford Distance Vector Algorithm
- How to diagram the counting to infinity problem.
- How RIP resolves the counting to infinity problem
- How split horizon attempts to prevent the counting to infinity problem.

Review Questions**Name** _____**Lesson 5-1: Routing Information Protocol****Part A**

Number the steps of the Bellman-Ford algorithm in order of occurrence.

Step Number	Step Description
	Find a row in the local routing table with the same destination.
	Add 1 to the distance field in the neighbor's row.
	If the neighbor's distance is less than the local distance replace the local row with the neighbor's row and record the neighbor who sent the information as the next hop. Go on to the next row.
	If no such row exists, add the neighbor's row and record the neighbor who sent the information as the next hop.
	If the neighbor's distance is greater or equal to the local distance, ignore it, and go to the next row.
	Get the next row from the table sent by the neighbor's router.
	If the row exists, compare the distance in the corresponding row in the local table to the distance from the row in the neighbor's table.

Part B

1. Diagram the count to infinity problem and write a short explanation of your diagram and how RIP solves the problem.

Part C

1. Describe how split horizon resolves the counting to infinity problem. Include a brief example in your description.

Scoring

Rubric: Suggested Evaluation Criteria and Weightings

Criteria	%	Your Score
Part A: Identify the steps of the Bellman-Ford Distance Vector Algorithm	25	
Part B: Diagram the count to infinity problem. Describe how RIP resolves the counting to infinity problem	50	
Part C: Explain how split horizon attempts to prevent the count to infinity problem.	25	
TOTAL	100	
Try It Out	100	
Stretch Yourself	100	
Network Wizards	100	
FINAL TOTAL	100	

Resources

Comer, D. E. (1995). Internetworking With TCP/IP : Principles, Protocols, and Architecture 3rd edition, *Prentice Hall*, Upper Saddle River, New Jersey.

Perlman, R. (1992). Interconnections: Bridges and Routers, *Addison-Wesley Publishing Co.*, Reading, MA.

Stevens, W. R. (1994). TCP/IP Illustrated, Volume 1: The Protocols, *Addison-Wesley Publishing Co.*, Reading, MA.

Tannenbaum, A.S. (1996). Computer Networks, *Prentice Hall*, Upper Saddle River, New Jersey.