

RIP protocol



Successfully pass the free certification exam at IW Academy and become an Infinet Certified Engineer.

[To the certification exam](#)

Table of contents

- [Table of contents](#)
- [Description](#)
- [RIP's operational algorithm](#)
 - [Start RIP](#)
 - [Routing information distribution](#)
 - [Adding routes to the RIB](#)
 - [Adding routes to the FIB](#)
 - [Timers control](#)
 - [Network changes control](#)
 - [Adding a new link](#)
 - [A link is out of order](#)
 - [A router is out of order](#)
- [False routes](#)
- [RIP features](#)
- [Additional materials](#)
 - [Other](#)

Description

RIP (Routing Information Protocol) is a dynamic routing protocol based on the Bellman-Ford algorithm. The protocol has the following features:

- the first RIP version was developed in 1969, as described in [RFC 1058](#);
- the second RIP version was developed in 1994, as described in [RFC 2454](#). This version is the main one used in IPv4 networks. There is no backward compatibility between the first and the second versions;
- there is a RIP version developed for IPv6 networks. This version is called RIPng and it is described in [RFC 2080](#);
- RIP is an internal distance vector routing protocol;
- the number of hops is used as a metric, i.e. the number of routers on the path to the destination network. The maximum metric value is 16 and it limits the network size in which RIP can be used;
- the multicast address 224.0.0.9 is reserved for RIP version 2. The first version of the protocol uses the broadcast address 255.255.255.255;
- UDP datagrams are used for service information transmission and port 520 is assigned to the protocol;
- the first RIP version only supports routes to classful networks, the second to classless networks;
- the distance value of 120 is used for RIP;
- RIP supports authentication: routing information will only be accepted from a router having the same key value.

RIP's operational algorithm

Let's look at an example on how is the routing information distributed over a network using RIP. The example follows the scheme in (Figure 1):

- the network consists of four routers R1, R2, R3 and R4, connected in a ring. These routers form a RIP domain. A RIP domain is a set of routers that exchange routing information using RIP;
- router R1 has an external link to the WAN-1 network;
- an independent IP network class is assigned for each link.

RIP must be configured on all the routers. WAN-1 is the external network, i.e. the route to this network must be added to the RIP domain as external.

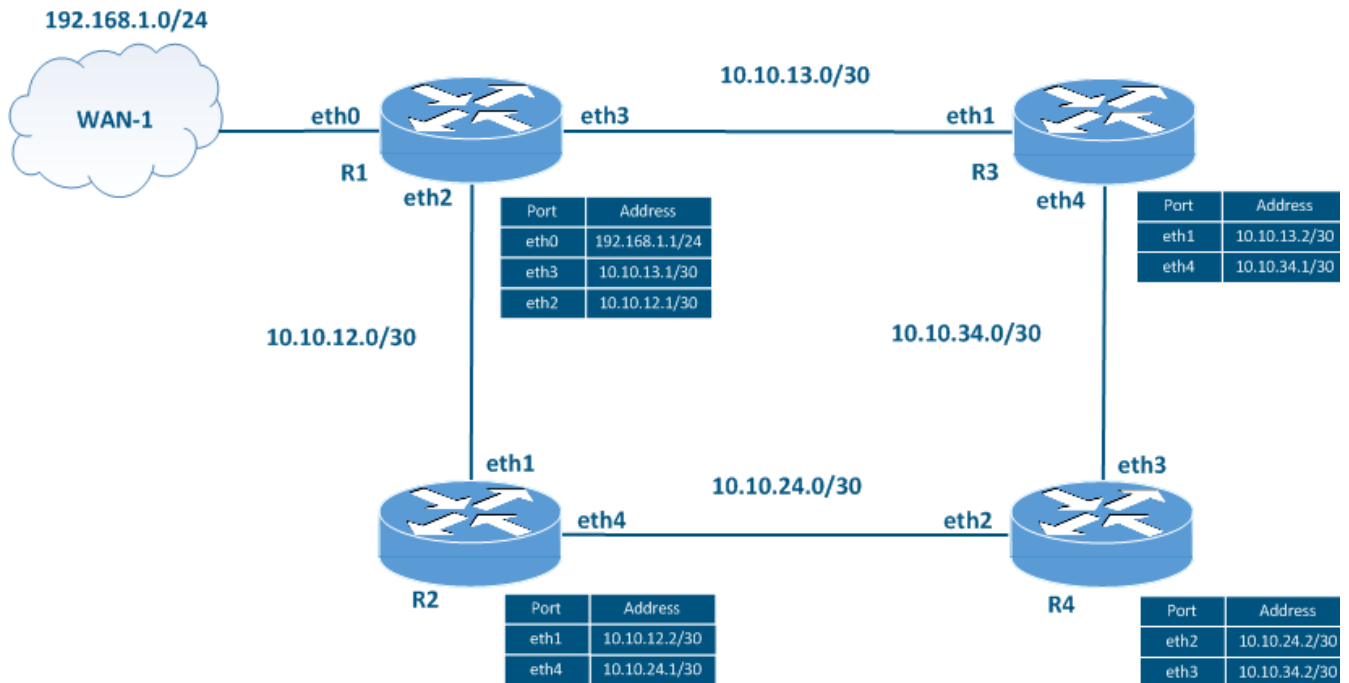


Figure 1 - Network scheme used to explain the RIP's operational principles

In this example, only the distribution of the routing information from router R1 will be described, because the routes from the other devices will be distributed in the same way.

The RIP operational algorithm is the following:

1. RIP starts.
2. The routing information distributed.
3. Routes are added to the RIB.
4. Routes are added to the FIB.
5. Timers control.
6. Network changes control.

Start RIP

At this stage, RIP must be enabled on the routers and the list of interfaces that will participate in the RIP process must be defined.

The network list definition is performed using the "network A.B.C.D/M" command. All network interfaces whose IP addresses belong to the specified range participate in RIP. This means that the routing information will be distributed from these interfaces, and the information about the networks associated with these interfaces will be transmitted to other routers.

If an interface participates in the RIP process, but routing information should not be distributed over that interface, then the interface can be configured as passive.

Let's assume that the commands "network 10.10.13.0/30" and "network 10.10.12.0/30" have been executed on router R1. In this case, the router will assume that RIP is running on eth2 and eth3.

Routing information distribution

The routing information distribution about the networks connected to router R1 is performed in the following sequence:

- **Step 1:** router R1 generates service messages with routing information. The service messages contain information about networks 10.10.13.0/30 and 10.10.12.0/30, since the interfaces associated with these networks are included in RIP when the [protocol is started](#). The service messages include also information about the 192.168.1.0/24 network, since router R1 must redistribute the external routes. The metric value for each network is set to 1.
- **Step 2:** the router sends out the service messages generated at step 1 (Figure 2a). The distribution is performed through the eth2 and the eth3 interfaces included in RIP when the protocol is started.




Figure 2a - R1's routing information distribution

- **Step 3:** routers R2 and R3 generate service messages to transmit the routing information to router R4:
 - **R2:** the router includes in the service message the 192.168.1.0/24 and the 10.10.13.0/30 networks received from router R1, incrementing the metric value. The information about the 10.10.12.0/30 network received from R1 is ignored because R2 is directly connected to this network. Instead, router R2 includes the 10.10.12.0/30 network with the metric 1 in the service message, announcing the network by itself. Also, the service message includes information about the network 10.10.24.0/30.
 - **R3:** the router includes in the service message the 192.168.1.0/24 and the 10.10.12.0/30 networks received from router R1, incrementing the metric value. The information about the 10.10.13.0/30 network received from R1 is ignored because R3 is directly connected to this network. Instead, router R3 includes the 10.10.13.0/30 network with the metric 1 in the service message, announcing the network by itself. Also, the service message includes information about the network 10.10.34.0/30.
- **Step 4:** routers R2 and R3 send the service messages generated in the previous step to router R4 (Figure 2b). Note that R2 and R3 send routing information to R1, but this process will not be described in this example.

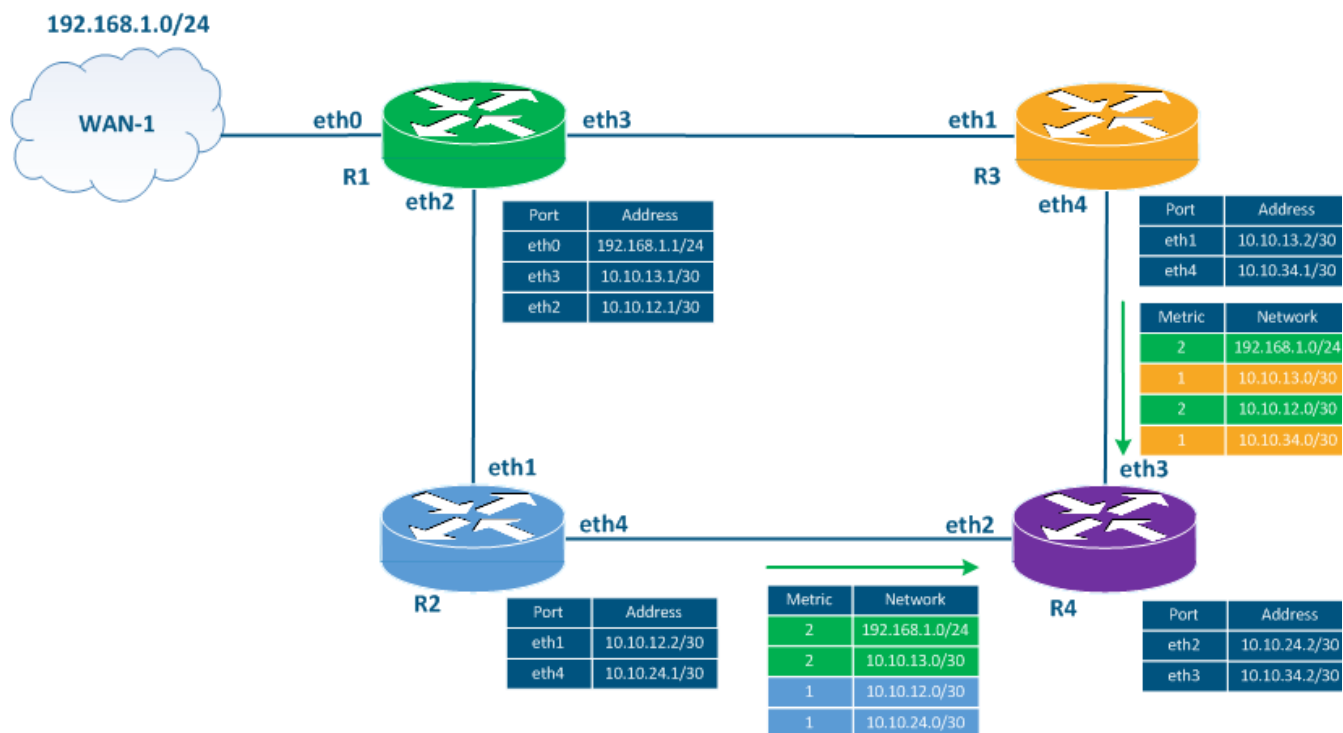


Figure 2b - Distribution of the routing information by routers R2 and R3

Adding routes to the RIB

After exchanging the service messages, the routers add the received routes to the RIB. In this case, some routes are filtered out. To filter the routing information, the devices are guided by the following principles:

- only one route to the destination network can be added to the RIB;
- if there are two routes towards the same network, the route with a lower metric value is added to the RIB;
- a route with a higher metric value can be added to the RIB if it is received from the same source;
- if there are two routes towards the same destination network with the same metric values, then the route received first will be added to the RIB.

Let's assume that router R4 received a service message from router R2 earlier than from R3. Then, according to the described principles, the routers add the following routing information to the RIB:

R2 router		
Destination network	Metric	Gateway
10.10.12.0/30	1	-
10.10.24.0/30	1	-
10.10.13.0/30	2	10.10.12.1
192.168.1.0 /24	2	10.10.12.1
R3 router		
Destination network	Metric	Gateway
10.10.13.0/30	1	-
10.10.34.0/30	1	-

10.10.12.0/30	2	10.10.13.1
192.168.1.0/24	2	10.10.13.1
R4 router		
Destination network	Metric	Gateway
10.10.24.0/30	1	-
10.10.34.0/30	1	-
10.10.12.0/30	2	10.10.24.1
10.10.13.0/30	2	10.10.34.1
192.168.1.0/24	3	10.10.24.1

Adding routes to the FIB

The export of the routes from the RIB to the FIB is associated with the analysis of the distance value assigned to each route source. Distance value 120 is assigned to RIP, so some of the routes added to the RIB in the previous step will be filtered out. For example, all the routes with the metric 1 will be discarded because these are routes to directly connected networks and have a 0 distance value.

Timers control

The distribution of the routing information described [above](#) is cyclically repeated. The repetition period is defined by the "update timer's" value. By default, the update timer is 30 seconds. Thus, all the routing information in the network is distributed every 30 seconds.

Note that the routers do not establish neighboring relations and do not synchronize the routing information exchange time, therefore, the time for sending service messages is distributed in the 30 seconds interval for every router, avoiding this way the service traffic bursts in the network.

After a route that was learned via RIP is placed in the RIB, the timeout timer is started, its default value being 180 seconds. If the router does not receive an update within 180 seconds, the route will be marked as unavailable, i.e. the metric value of such a route is set to 16. The router cannot use this route for data transmission.

The RIP service data exchange is not guaranteed, so the timeout timer value must be greater than the update timer. Otherwise, in case of delays or packet loss a route can be falsely marked as unavailable.

After adding a route to the RIB, in addition to the timeout timer, the garbage timer is also started with the default value of 240 seconds. If the router does not receive an update within 240 seconds, then such a route will be removed from the routing table.

When routing information updates are received, the timeout timer and the garbage timer are reset to their original values.

Network changes control

RIP is a dynamic routing protocol and must adapt to the changes in the network. There are three main scenarios that describe the network changes:

- a new link has been added;
- a link is out of order;
- a router is out of order.

Adding a new link

Let's add a new link between router R1 and WAN-2 (Figure 3a). In this case, the WAN-2 network will be external for RIP, so the route to this network will be distributed similarly to the route towards the WAN-1 network: router R1 will advertise this network with the metric 1, R2 and R3 will increment the metric's value and transmit the route to R4. In case of new routers included into the RIP domain, the process will be similar.

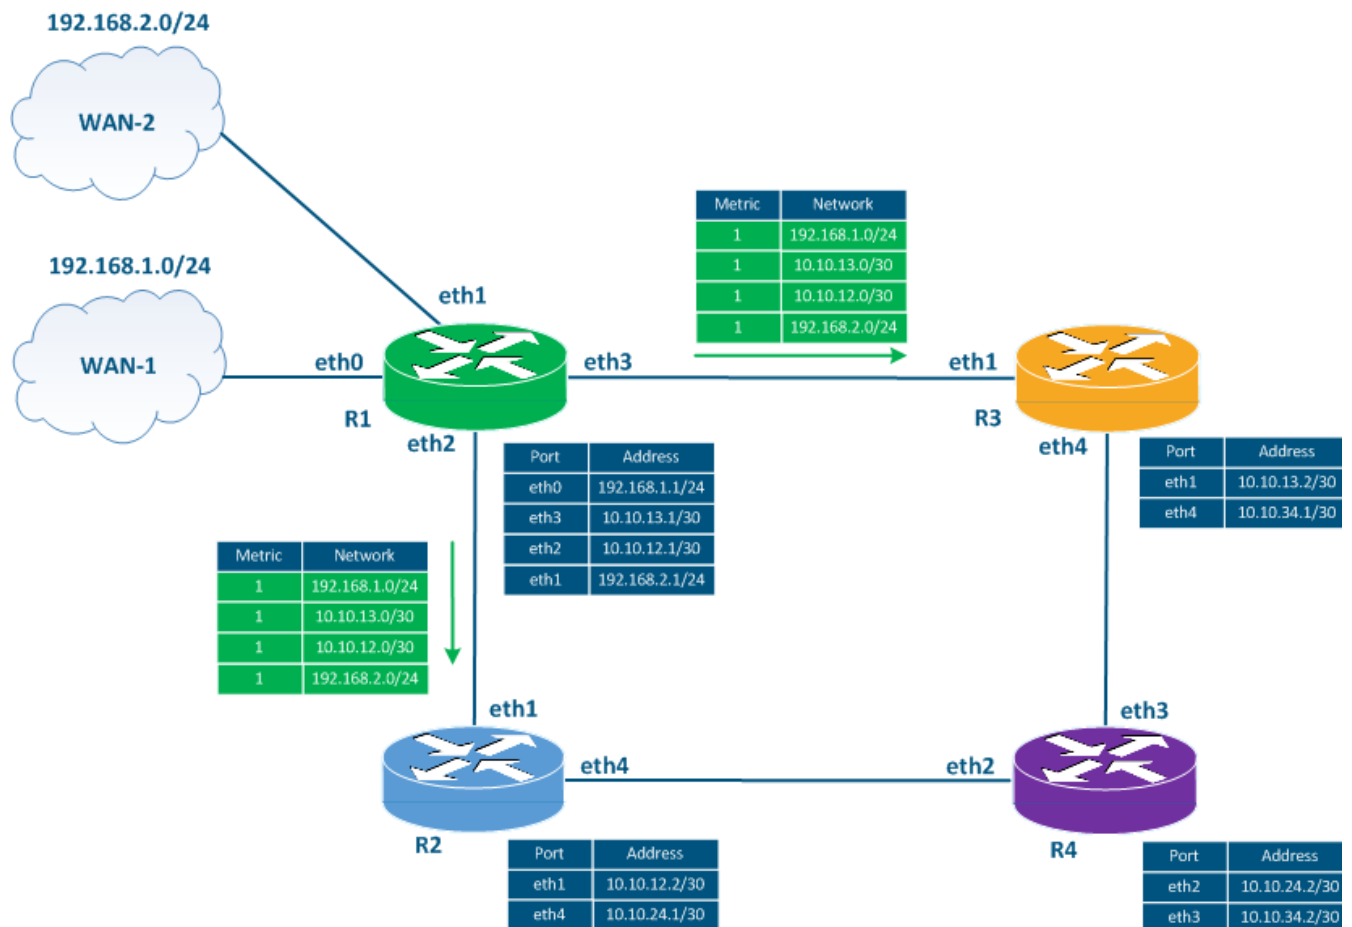


Figure 3a - Transmission of the service messages when a new link has been added

A link is out of order

Let's assume that the link between routers R1 and R3 has failed (Figure 3b):

- **Step 1:** the interfaces eth3 of R1 and eth1 of R3 will become down.
- **Step 2:** R1 and R3 will set the metric's value to 16 in the RIB for the route towards network 10.10.13.0/30.
- **Step 3:** according to the update timer's expiration, routers R1 and R3 will generate a service message including the known routing information. This message includes the route to the unreachable network with the metric 15.
- **Step 4:** routers R2 and R4 receive the service messages from R1 and R3, increment the metric value, and add the route to the RIB. The RIB of the routers R2 and R4 already has a route to the 10.10.13.0/30 network with the metric 2, but since the sources of the new routes with the worst metric are the same routers, then the route in the RIB will be replaced with a route with the metric 16. Thus, all the routers in the network will receive information about the link's unavailability.

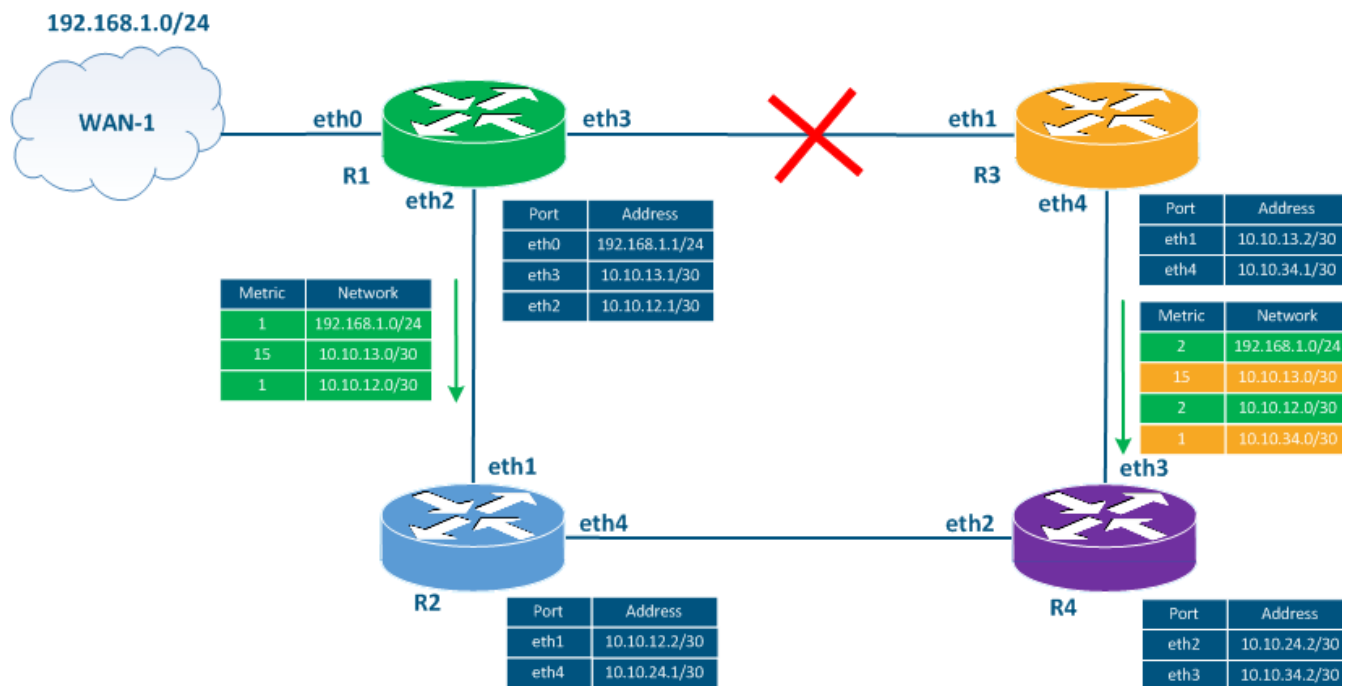


Figure 3b - Transmission of the service messages in case of link failure

The time during which the updated information is distributed depends on the size of the network and on the update timer's value. As the maximum network size is 16 hops and the update timer by default is 30 seconds, then in the worst scenario, the information about the link's unavailability will be transmitted in $15 * 30 = 450$ seconds.

A router is out of order

Let's add two switches SW1 and SW2 to the scheme in (Figure 3c) and assume that router R1 has failed.

If router R1 fails, it does not have time to send the routing information update about its status to routers R2 and R3. In addition, R2 and R3 won't be informed that R1 is unavailable because they are directly connected to switches SW1 and SW2 which don't support RIP. In this case, the other routers will assume that R1 is available during the timeout timer and will use the routes associated with R1. Particularly, the routes to networks 10.10.12.0/30, 10.10.13.0/30 and 192.168.1.0/24 will be valid in the routing tables of all the routers.




Figure 3c - Transmission of the service messages in case of router failure

False routes

Let's simplify the scheme, leaving two routers R1 and R3 (Figure 4a). We will assume that the routers have exchanged routing information and the device's routing tables are up to date.




Figure 4a - Network scheme for the false route scenario

R1 router		
Destination network	Metric	Gateway
10.10.13.0/30	1	-
192.168.1.0/24	1	-

R3 router		
-----------	--	--

Destination network	Metric	Gateway
10.10.13.0/30	1	-
192.168.1.0/24	2	10.10.13.1

The link between router R1 and the external WAN-1 network fails (Figure 4b). Router R1 sets the metric equal to 16 for such a route, but does not generate a service message with updated routing information, since the update timer value is 25 seconds. At the same time, the update timer on router R3 is 3 seconds.

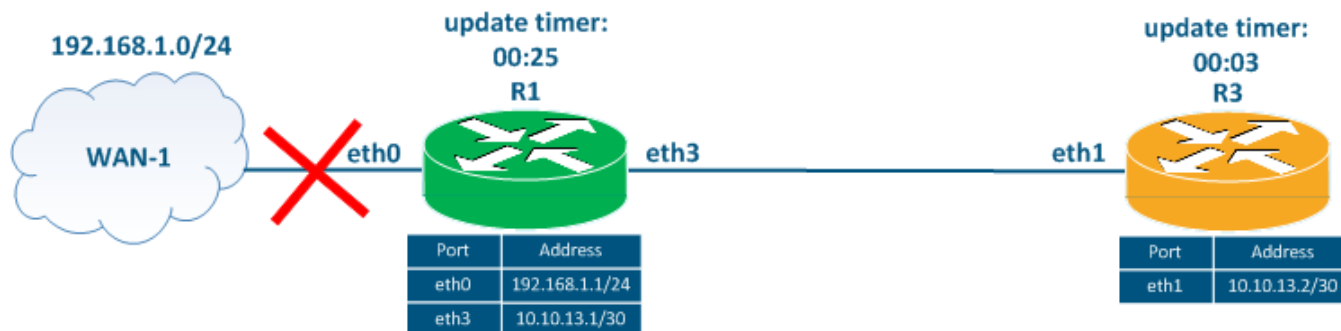


Figure 4b - The R1 - WAN-1 connection fails

R1 router		
Destination network	Metric	Gateway
10.10.13.0/30	1	-
192.168.1.0/24	16	-

R3 router		
Destination network	Metric	Gateway
10.10.13.0/30	1	-
192.168.1.0/24	2	10.10.13.1

Three seconds later after the R1 - WAN-1 link has failed, R3's update timer expires, so R3 sends a routing update to R1 (Figure 4c). This update contains information about the 10.10.13.0/30 and 192.168.1.0/24 networks. Router R1 adds a route towards the 192.168.1.0/24 network to the RIB, because its metric is better than the route that has already been added to the RIB.

A routing loop is formed, each router, passing traffic to the 192.168.1.0/24 network. Thus, the traffic will be transmitted between the routers until the TTL expires. In addition to the routing loop generation, both routes are false as the connection with the network 192.168.1.0/24 is lost.

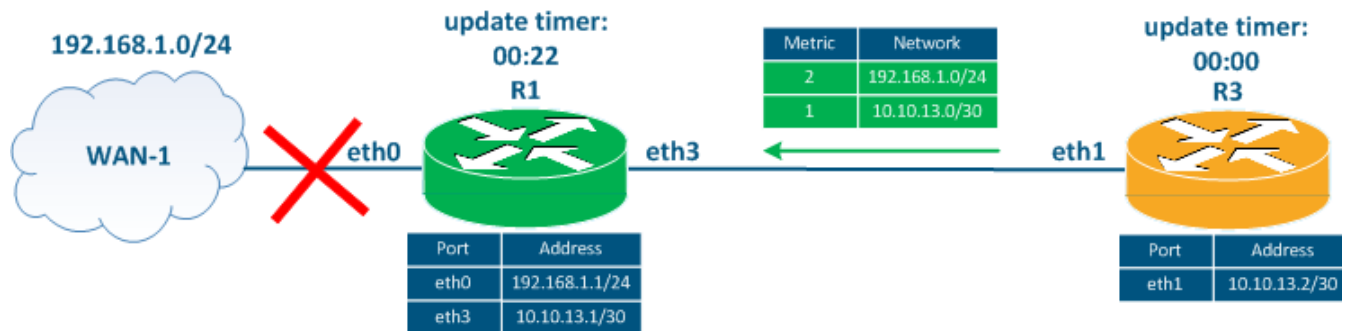


Figure 4c - Distribution of the service messages by R3

R1 router		
Destination network	Metric	Gateway
10.10.13.0/30	1	-
192.168.1.0/24	3	10.10.13.2

R3 router		
Destination network	Metric	Gateway
10.10.13.0/30	1	-
192.168.1.0/24	2	10.10.13.1

The false routes will continue until the timeout timer expires on R3. After that, R3 will set the metric value for this route to 16. In the next routing information distribution, router R1 will send routing information to R3 about the 10.10.13.0/30 and 192.168.1.0/24 networks (Figure 4d). Since the metric for the route to 192.168.1.0/24 is 3, and less than 16, R3 will add a false route to the RIB, incrementing the metric value.

The false route exchange will continue until the metric value reaches the invalid value of 16. With the default settings, this will happen 36 minutes later after the occurrence of the link failure between R1 and WAN-1.




Figure 4d - Service message distribution with routing information by router R1

R1 router		
Destination network	Metric	Gateway
10.10.13.0/30	1	-
192.168.1.0/24	3	10.10.13.2

R3 router		
Destination network	Metric	Gateway
10.10.13.0/30	1	-
192.168.1.0/24	4	10.10.13.1

To avoid the appearance of the false routes, the following mechanisms are used:

- **split-horizon:** the routers do not distribute any routing information updates to those routers which are the sources of the information. In the example above (Figure 4d), router R3 will not include information about the route to the 192.168.1.0/24 network in the message for router R1, since R1 is the source of this information. This will allow to avoid the described situation, but will not work effectively in a scheme with a large number of routers. In the scheme with four routers (Figure 1) when the link with WAN-1 is out of order, routers R2 and R3 will not transmit information about the route to the 192.168.1.0/24 network for R1. However, this route is also in the R4's RIB and it will be included in the periodic distribution of the routing information, which will lead to the appearance of a false route.
- **poison-reverse:** this mechanism implies the immediate distribution of the updates about unavailable routes without waiting for the update timer expiration. An unreachable route is sent with a metric of 15, so each router will increment the metric and lead to an unavailable route in the RIB. The mechanism has a disadvantage, since the routing information distribution about an unavailable route requires a certain time, during which a message with old routing information can be sent and include the unavailable route marked as available.
- **garbage-timer:** the device does not accept updates for routes marked as unavailable before the garbage-timer expires.

The listed mechanisms make it possible to adapt RIP for the modern networks, since the existence of the false routes for 36 minutes or a protocol convergence time of 450 seconds is unacceptable for many services.

RIP features

RIP has the following features:

- the operational algorithm used by RIP is easy to understand and to configure;
- RIP is a standard protocol and can be implemented using devices from different manufacturers;
- false routes may appear in the protocol: this disadvantage was eliminated by adding new functions to the protocol, but this makes the protocol's operation and implementation more complicated;
- there is a network size limitation of 16 hops;
- it takes a long time for the network to adapt to changes.

Additional materials

Other

1. [RFC 1058](#)
2. [RFC 2454](#)
3. [RFC 2080](#)
4. [ARDA \(Aqua Router Daemon\)](#)
5. [arip command](#)
6. [rip command](#)