

# Advanced IP Networking

## Part 1: Dynamic Routing Protocols

---

Rik Wade

[rik@rikwade.com](mailto:rik@rikwade.com)

WYLUG 14/June/2004



# What We'll Cover

---

- IP Routing Recap
- Introduction to Dynamic Protocols
- A Basic Dynamic Protocol – RIP
- Link-state Protocols – OSPF (briefly)
- Zebra / Quagga



# IP Routing Recap

---

- Routing table entries automatically created when interfaces configured
- Generally populated with static and connected routes
- Routing table can be dynamically updated in real-time using.... dynamic routing protocols
  - Large networks
  - Redundant routing
  - Multiple upstream networks



# Dynamic Protocols

---

- Update the routing table in real-time
- Conversation between systems running the same routing protocol (routers)
- Routers swap information about their known routes (connected or otherwise)
- Timers and more advanced methods of detecting loss of a neighbor or the withdrawal of routes from the database
- Specific software and configuration required




# Dynamic Protocols

---

- Nothing clever or magic about them
  - Run as a process (daemon) on a router
  - Listen and talk on a given IP interface and Port
- Maintain a database of networks
- Algorithm selects networks from the database and installs in the routing table
- Clever Stuff(tm) is in the data that is stored, the way in which it is collected and the route selection algorithm

# Dynamic Protocols



- How can Host A talk to Host B?
- RIP carries information of 192.168.3.0/24 to Router A and 192.168.2.0/24 to Router B



# A Basic Protocol - RIP


---

- Routing Information Protocol - RIP (distance vector protocol). See RFC 1058.
- Broadcasts all or part of routing table every 30 seconds (on Ethernet). DST UDP port 520.
- If a route is not refreshed for 3 minutes, it is marked as unreachable and removed.
- Metric for RIP is "Hop Count".
- Hop count is limited to 15 to avoid loops. 16 is used to signify unreachable networks.



# Basic RIP

---



- Router B advertises 192.168.3.0 to Router A with a Metric of 1 Hop.
- Router A installs 192.168.3.0 in its table with a Metric of 1. Advertises with a Metric of 2.



# Before RIP – Router A

---

```
Router-A#sh ip route
```

```
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP  
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area  
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2  
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP  
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate def  
U - per-user static route, o - ODR
```

```
Gateway of last resort is not set
```

```
C    192.168.1.0/24 is directly connected, Serial0
```

```
C    192.168.2.0/24 is directly connected, Ethernet0
```



# Before RIP – Router B

---

```
Router-B#sh ip route
```

```
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP  
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area  
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2  
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP  
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate def  
U - per-user static route, o - ODR
```

```
Gateway of last resort is not set
```

```
C    192.168.1.0/24 is directly connected, Serial0  
C    192.168.3.0/24 is directly connected, Ethernet0
```



# With RIP – Router A

---

```
Router-A#sh ip route
```

```
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP  
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area  
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2  
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP  
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate def  
U - per-user static route, o - ODR
```

```
Gateway of last resort is not set
```

```
C    192.168.1.0/24 is directly connected, Serial0  
C    192.168.2.0/24 is directly connected, Ethernet0  
R    192.168.3.0/24 [120/1] via 192.168.1.2, 00:00:15, Serial0
```



# With RIP – Router B

---

```
Router-B#sh ip route
```

```
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
```

```
[edited for brevity]
```

```
Gateway of last resort is not set
```

```
C    192.168.1.0/24 is directly connected, Serial0
```

```
R    192.168.2.0/24 [120/1] via 192.168.1.1, 00:00:12, Serial0
```

```
C    192.168.3.0/24 is directly connected, Ethernet0
```

## RIP Debug Output:

```
00:40:48: RIP: sending v1 update to 255.255.255.255 via Ethernet0  
    (192.168.3.1)
```

```
00:40:48:      network 192.168.1.0, metric 1
```

```
00:40:48:      network 192.168.2.0, metric 2
```

```
00:40:48: RIP: Update contains 2 routes
```



# RIP Network Failure

---

```
Router-A#sh ip route
```

```
Codes: [edited for brevity]
```

```
Gateway of last resort is not set
```

```
C    192.168.1.0/24 is directly connected, Serial0
```

```
C    192.168.2.0/24 is directly connected, Ethernet0
```

```
R    192.168.3.0/24 is possibly down, routing via 192.168.1.2, Serial0
```

## RIP Debug output:

```
00:44:41: RIP: sending v1 update to 255.255.255.255 via Ethernet0  
(192.168.2.1)
```

```
00:44:41:      network 192.168.1.0, metric 1
```

```
00:44:41:      network 192.168.3.0, metric 16
```



# Other Routing Protocols

---

- RIP is:
  - Old, noisy, slow to converge, **classful**
  - Fixed somewhat in RIPv2
- OSPF and others are:
  - New, efficient, quick to converge, **classless**
  - Link-state protocols are “in”
  - Exchange what is require, when required
  - Can converge within 1 second, even in large networks



# Zebra and Quagga

---

- Zebra (<http://www.zebra.org>)
  - A GNU form of a commercial package
  - Certain protocols and features are only available in the commercial package
  - The community forked this project
- Quagga (<http://www.quagga.net>)
  - Same architecture as Zebra
  - Many more features
  - Community-driven

The logo graphic consists of a vertical black line intersecting a horizontal black line. To the left of the intersection, there are three overlapping squares: a yellow one at the top, a red one in the middle, and a blue one at the bottom. The word "Quagga" is written in a blue, sans-serif font to the right of the vertical line.

# Quagga

---

- Packages available for:
  - RedHat & Fedora Core
  - Debian
  - Gentoo
- Quagga command line is very similar to Cisco's IOS (<http://www.cisco.com>)
- Used in many ISPs and can out-perform dedicated routers from Cisco and others



# Quagga

---

- Quagga consists of
  - Zebra daemon: collects routing information and interfaces with the system's tables
  - OSPFd: OSPF v2
  - OSPF6d: OSPF v3 (for IPv6)
  - RIPd: RIP v1 and v2
  - RIPngd: RIP v3 (for IPv6)
  - BGPd: BGP v4+
  - VTYsh: a command-line interface



# Quagga

---

- Excellent documentation at:
  - <http://www.quagga.net/docs.php>
- Many advanced features in addition to basic routing protocol support
  - Access lists – traffic manipulation
  - Route maps – route manipulation
  - Route redistribution – e.g. From RIP to BGP



# FIN

---

- TCP/IP Illustrated Volume 1 (The Protocols)
  - W. Richard Stevens, ISBN 0-201-63346-9
- Routing TCP/IP Volume 1
  - Jeff Doyle, ISBN 1-57870-041-8
- Quagga
  - <http://www.quagga.net>
- Cisco References
  - [http://www.cisco.com/en/US/tech/tk365/tech\\_protocol\\_families.html](http://www.cisco.com/en/US/tech/tk365/tech_protocol_families.html)
  - Excellent reference for all protocols