Internet Routing

Dr. Miled M. Tezeghdanti

May 8, 2012

Introduction

- Networks may be interconnected using
  - Repeaters
    - Work at the Physical layer
  - Bridges
    - Work at the Data-Link layer
    - Bridging uses physical addressing
    - Spanning Tree Protocol and Algorithm "STP"
  - Routers
    - Work at the Network layer
    - Routing uses logical addressing
    - Routing computes routes for available networks and stores them in a routing table
    - Static Routing
    - Dynamic Routing

Static Routing

- Manual Process
- Network Administrator
  - route command
  - route add 10.1.2.0/24 10.1.1.2
  - route delete 10.1.2.0/24 10.1.1.2
- Not tolerant to failures
  - Repair the failure
  - Change the route
- No processing overhead
- No traffic overhead
- Good for default routes and stub networks

Dynamic Routing

- Automated Process
- Routing Protocol
    - IGP vs EGP (Interior Gateway Protocol vs Exterior Gateway Protocol)
    - Metric-based vs Policy-based
    - Distance Vector vs Link State
- Tolerant to failures
- Processing overhead
- Traffic overhead
- Good for large networks

Routing Protocol Characteristics

- Convergence time
- Scalability
- Resource usage (CPU and Bandwidth)
- Implementation and maintenance

Distance Vector Routing

- Distance or metric of a destination
- Vector or direction to that destination
- Each router knows only the distance or metric to each destination and the direction (next-hop) to take to get there.
- Routing information is exchanged between directly connected neighbors
- RIP "Routing Information Protocol"

RIP

- Each router can reach directly connected networks with a distance 1
- Each router sends periodically to its neighbors the content of its routing table
- Each router learns new routes from its neighbors
- Each router stores for each destination the best route with the minimum distance

RIP Example

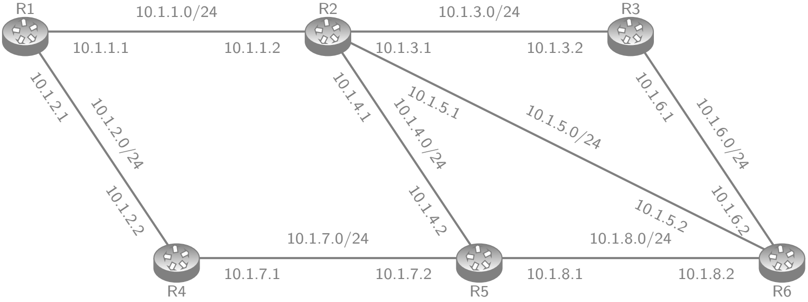
R1
10.1.1.0/24
R2
10.1.3.0/24
R3
10.1.1.1
10.1.1.2
10.1.3.1
10.1.3.2
10.1.2.1
10.1.2.0/24
10.1.2.2
10.1.4.1
10.1.5.1
10.1.4.0/24
10.1.5.0/24
10.1.4.2
10.1.6.1
10.1.6.0/24
10.1.6.2
10.1.5.2
10.1.7.0/24
10.1.8.0/24
10.1.7.1
10.1.7.2
10.1.8.1
10.1.8.2
R4
R5
R6

RIP Example

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.2.0/24	C	1

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.3.0/24	C	1
10.1.4.0/24	C	1
10.1.5.0/24	C	1

Destination	Gateway	Cost
10.1.3.0/24	C	1
10.1.6.0/24	C	1

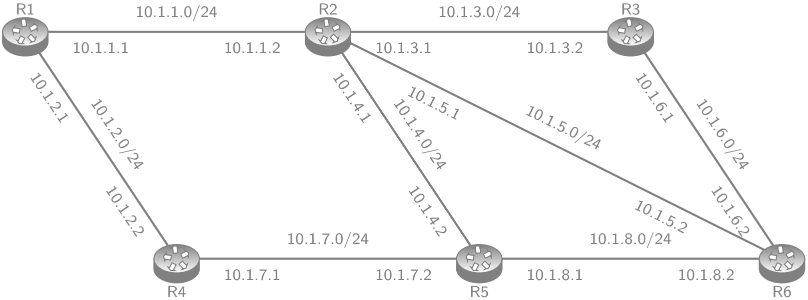

Destination	Gateway	Cost
10.1.2.0/24	C	1
10.1.7.0/24	C	1

Destination	Gateway	Cost
10.1.4.0/24	C	1
10.1.7.0/24	C	1
10.1.8.0/24	C	1

Destination	Gateway	Cost
10.1.5.0/24	C	1
10.1.6.0/24	C	1
10.1.8.0/24	C	1

RIP Example

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.2.0/24	C	1

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.3.0/24	C	1
10.1.4.0/24	C	1
10.1.5.0/24	C	1
10.1.2.0/24	10.1.1.1	2

Destination	Gateway	Cost
10.1.3.0/24	C	1
10.1.6.0/24	C	1

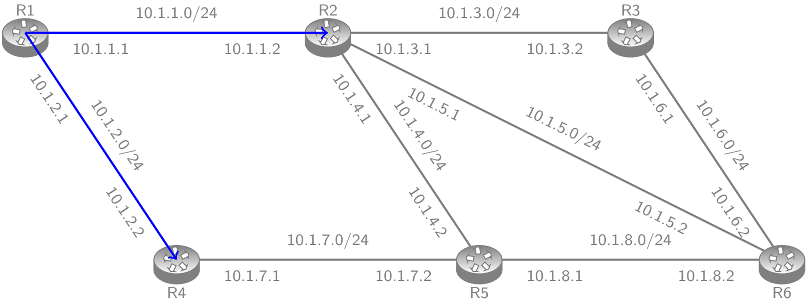

Destination	Gateway	Cost
10.1.2.0/24	C	1
10.1.7.0/24	C	1
10.1.1.0/24	10.1.2.1	2

Destination	Gateway	Cost
10.1.4.0/24	C	1
10.1.7.0/24	C	1
10.1.8.0/24	C	1

Destination	Gateway	Cost
10.1.5.0/24	C	1
10.1.6.0/24	C	1
10.1.8.0/24	C	1

RIP Example

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.2.0/24	C	1
10.1.3.0/24	10.1.1.2	2
10.1.4.0/24	10.1.1.2	2
10.1.5.0/24	10.1.1.2	2

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.3.0/24	C	1
10.1.4.0/24	C	1
10.1.5.0/24	C	1
10.1.2.0/24	10.1.1.1	2

Destination	Gateway	Cost
10.1.3.0/24	C	1
10.1.6.0/24	C	1
10.1.1.0/24	10.1.3.1	2
10.1.4.0/24	10.1.3.1	2
10.1.5.0/24	10.1.3.1	2
10.1.2.0/24	10.1.3.1	3

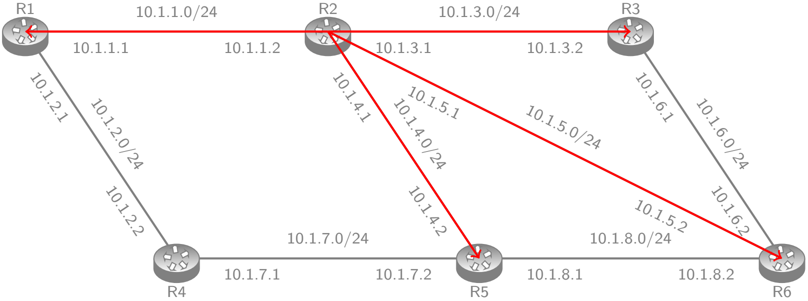

Destination	Gateway	Cost
10.1.2.0/24	C	1
10.1.7.0/24	C	1
10.1.1.0/24	10.1.2.1	2

Destination	Gateway	Cost
10.1.4.0/24	C	1
10.1.7.0/24	C	1
10.1.8.0/24	C	1
10.1.1.0/24	10.1.4.1	2
10.1.3.0/24	10.1.4.1	2
10.1.5.0/24	10.1.4.1	2
10.1.2.0/24	10.1.4.1	3

Destination	Gateway	Cost
10.1.5.0/24	C	1
10.1.6.0/24	C	1
10.1.8.0/24	C	1
10.1.1.0/24	10.1.5.1	2
10.1.3.0/24	10.1.5.1	2
10.1.4.0/24	10.1.5.1	2
10.1.2.0/24	10.1.5.1	3

RIP Example

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.2.0/24	C	1
10.1.3.0/24	10.1.1.2	2
10.1.4.0/24	10.1.1.2	2
10.1.5.0/24	10.1.1.2	2

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.3.0/24	C	1
10.1.4.0/24	C	1
10.1.5.0/24	C	1
10.1.2.0/24	10.1.1.1	2
10.1.6.0/24	10.1.3.2	2

Destination	Gateway	Cost
10.1.3.0/24	C	1
10.1.6.0/24	C	1
10.1.1.0/24	10.1.3.1	2
10.1.4.0/24	10.1.3.1	2
10.1.5.0/24	10.1.3.1	2
10.1.2.0/24	10.1.3.1	3

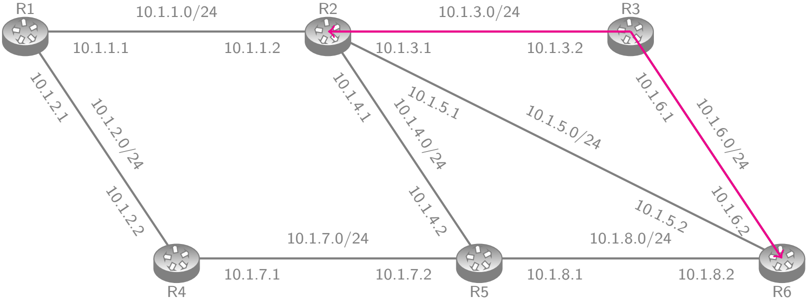

Destination	Gateway	Cost
10.1.2.0/24	C	1
10.1.7.0/24	C	1
10.1.1.0/24	10.1.2.1	2

Destination	Gateway	Cost
10.1.4.0/24	C	1
10.1.7.0/24	C	1
10.1.8.0/24	C	1
10.1.1.0/24	10.1.4.1	2
10.1.3.0/24	10.1.4.1	2
10.1.5.0/24	10.1.4.1	2
10.1.2.0/24	10.1.4.1	3

Destination	Gateway	Cost
10.1.5.0/24	C	1
10.1.6.0/24	C	1
10.1.8.0/24	C	1
10.1.1.0/24	10.1.5.1	2
10.1.3.0/24	10.1.5.1	2
10.1.4.0/24	10.1.5.1	2
10.1.2.0/24	10.1.5.1	3

RIP Example

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.2.0/24	C	1
10.1.3.0/24	10.1.1.2	2
10.1.4.0/24	10.1.1.2	2
10.1.5.0/24	10.1.1.2	2
10.1.7.0/24	10.1.2.2	2

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.3.0/24	C	1
10.1.4.0/24	C	1
10.1.5.0/24	C	1
10.1.2.0/24	10.1.1.1	2
10.1.6.0/24	10.1.3.2	2

Destination	Gateway	Cost
10.1.3.0/24	C	1
10.1.6.0/24	C	1
10.1.1.0/24	10.1.3.1	2
10.1.4.0/24	10.1.3.1	2
10.1.5.0/24	10.1.3.1	2
10.1.2.0/24	10.1.3.1	3

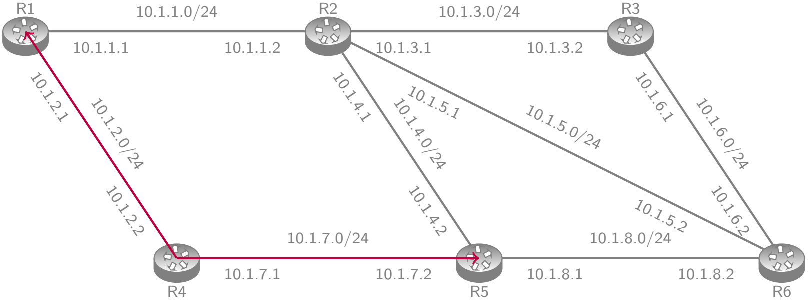

Destination	Gateway	Cost
10.1.2.0/24	C	1
10.1.7.0/24	C	1
10.1.1.0/24	10.1.2.1	2

Destination	Gateway	Cost
10.1.4.0/24	C	1
10.1.7.0/24	C	1
10.1.8.0/24	C	1
10.1.1.0/24	10.1.4.1	2
10.1.3.0/24	10.1.4.1	2
10.1.5.0/24	10.1.4.1	2
10.1.2.0/24	10.1.7.1	2

Destination	Gateway	Cost
10.1.5.0/24	C	1
10.1.6.0/24	C	1
10.1.8.0/24	C	1
10.1.1.0/24	10.1.5.1	2
10.1.3.0/24	10.1.5.1	2
10.1.4.0/24	10.1.5.1	2
10.1.2.0/24	10.1.5.1	3

RIP Example

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.2.0/24	C	1
10.1.3.0/24	10.1.1.2	2
10.1.4.0/24	10.1.1.2	2
10.1.5.0/24	10.1.1.2	2
10.1.7.0/24	10.1.2.2	2

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.3.0/24	C	1
10.1.4.0/24	C	1
10.1.5.0/24	C	1
10.1.2.0/24	10.1.1.1	2
10.1.6.0/24	10.1.3.2	2
10.1.7.0/24	10.1.4.2	2
10.1.8.0/24	10.1.4.2	2

Destination	Gateway	Cost
10.1.3.0/24	C	1
10.1.6.0/24	C	1
10.1.1.0/24	10.1.3.1	2
10.1.4.0/24	10.1.3.1	2
10.1.5.0/24	10.1.3.1	2
10.1.2.0/24	10.1.3.1	3

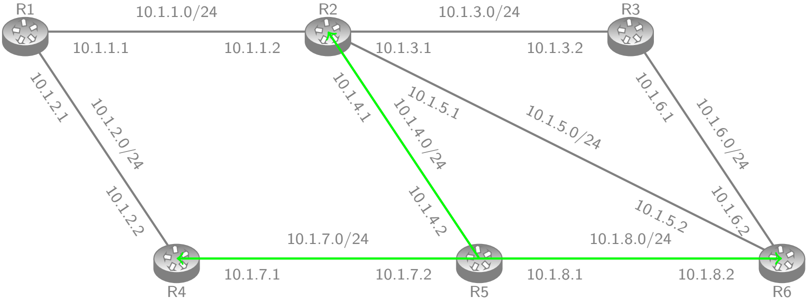

Destination	Gateway	Cost
10.1.2.0/24	C	1
10.1.7.0/24	C	1
10.1.1.0/24	10.1.2.1	2
10.1.4.0/24	10.1.7.2	2
10.1.8.0/24	10.1.7.2	2
10.1.3.0/24	10.1.7.2	3
10.1.5.0/24	10.1.7.2	3

Destination	Gateway	Cost
10.1.4.0/24	C	1
10.1.7.0/24	C	1
10.1.8.0/24	C	1
10.1.1.0/24	10.1.4.1	2
10.1.3.0/24	10.1.4.1	2
10.1.5.0/24	10.1.4.1	2
10.1.2.0/24	10.1.7.1	2

Destination	Gateway	Cost
10.1.5.0/24	C	1
10.1.6.0/24	C	1
10.1.8.0/24	C	1
10.1.1.0/24	10.1.5.1	2
10.1.3.0/24	10.1.5.1	2
10.1.4.0/24	10.1.5.1	2
10.1.2.0/24	10.1.5.1	3
10.1.7.0/24	10.1.8.1	2

RIP Example

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.2.0/24	C	1
10.1.3.0/24	10.1.1.2	2
10.1.4.0/24	10.1.1.2	2
10.1.5.0/24	10.1.1.2	2
10.1.7.0/24	10.1.2.2	2

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.3.0/24	C	1
10.1.4.0/24	C	1
10.1.5.0/24	C	1
10.1.2.0/24	10.1.1.1	2
10.1.6.0/24	10.1.3.2	2
10.1.7.0/24	10.1.4.2	2
10.1.8.0/24	10.1.4.2	2

Destination	Gateway	Cost
10.1.3.0/24	C	1
10.1.6.0/24	C	1
10.1.1.0/24	10.1.3.1	2
10.1.4.0/24	10.1.3.1	2
10.1.5.0/24	10.1.3.1	2
10.1.2.0/24	10.1.3.1	3
10.1.8.0/24	10.1.6.2	2
10.1.7.0/24	10.1.6.2	3

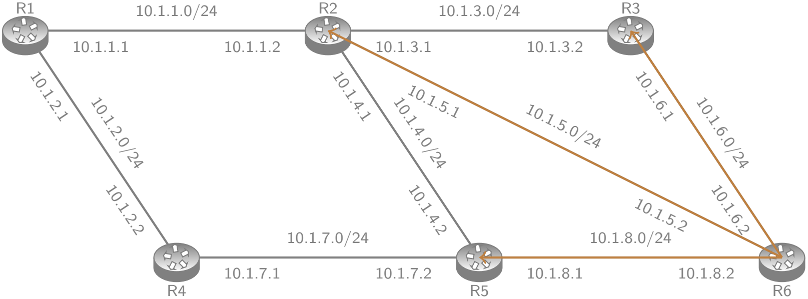

Destination	Gateway	Cost
10.1.2.0/24	C	1
10.1.7.0/24	C	1
10.1.1.0/24	10.1.2.1	2
10.1.4.0/24	10.1.7.2	2
10.1.8.0/24	10.1.7.2	2
10.1.3.0/24	10.1.7.2	3
10.1.5.0/24	10.1.7.2	3

Destination	Gateway	Cost
10.1.4.0/24	C	1
10.1.7.0/24	C	1
10.1.8.0/24	C	1
10.1.1.0/24	10.1.4.1	2
10.1.3.0/24	10.1.4.1	2
10.1.5.0/24	10.1.4.1	2
10.1.2.0/24	10.1.7.1	2
10.1.6.0/24	10.1.8.2	2

Destination	Gateway	Cost
10.1.5.0/24	C	1
10.1.6.0/24	C	1
10.1.8.0/24	C	1
10.1.1.0/24	10.1.5.1	2
10.1.3.0/24	10.1.5.1	2
10.1.4.0/24	10.1.5.1	2
10.1.2.0/24	10.1.5.1	3
10.1.7.0/24	10.1.8.1	2

RIP Example

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.2.0/24	C	1
10.1.3.0/24	10.1.1.2	2
10.1.4.0/24	10.1.1.2	2
10.1.5.0/24	10.1.1.2	2
10.1.7.0/24	10.1.2.2	2

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.3.0/24	C	1
10.1.4.0/24	C	1
10.1.5.0/24	C	1
10.1.2.0/24	10.1.1.1	2
10.1.6.0/24	10.1.3.2	2
10.1.7.0/24	10.1.4.2	2
10.1.8.0/24	10.1.4.2	2

Destination	Gateway	Cost
10.1.3.0/24	C	1
10.1.6.0/24	C	1
10.1.1.0/24	10.1.3.1	2
10.1.4.0/24	10.1.3.1	2
10.1.5.0/24	10.1.3.1	2
10.1.2.0/24	10.1.3.1	3
10.1.8.0/24	10.1.6.2	2
10.1.7.0/24	10.1.6.2	3

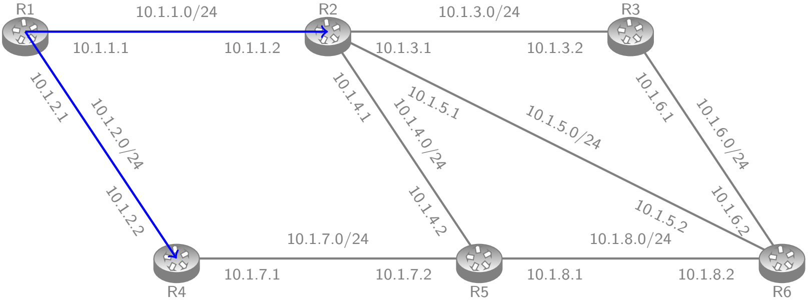

Destination	Gateway	Cost
10.1.2.0/24	C	1
10.1.7.0/24	C	1
10.1.1.0/24	10.1.2.1	2
10.1.4.0/24	10.1.7.2	2
10.1.8.0/24	10.1.7.2	2
10.1.3.0/24	10.1.7.2	3
10.1.5.0/24	10.1.7.2	3

Destination	Gateway	Cost
10.1.4.0/24	C	1
10.1.7.0/24	C	1
10.1.8.0/24	C	1
10.1.1.0/24	10.1.4.1	2
10.1.3.0/24	10.1.4.1	2
10.1.5.0/24	10.1.4.1	2
10.1.2.0/24	10.1.7.1	2
10.1.6.0/24	10.1.8.2	2

Destination	Gateway	Cost
10.1.5.0/24	C	1
10.1.6.0/24	C	1
10.1.8.0/24	C	1
10.1.1.0/24	10.1.5.1	2
10.1.3.0/24	10.1.5.1	2
10.1.4.0/24	10.1.5.1	2
10.1.2.0/24	10.1.5.1	3
10.1.7.0/24	10.1.8.1	2

RIP Example

R1:

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.2.0/24	C	1
10.1.3.0/24	10.1.1.2	2
10.1.4.0/24	10.1.1.2	2
10.1.5.0/24	10.1.1.2	2
10.1.7.0/24	10.1.2.2	2
10.1.6.0/24	10.1.1.2	3
10.1.8.0/24	10.1.1.2	3

R2:

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.3.0/24	C	1
10.1.4.0/24	C	1
10.1.5.0/24	C	1
10.1.2.0/24	10.1.1.1	2
10.1.6.0/24	10.1.3.2	2
10.1.7.0/24	10.1.4.2	2
10.1.8.0/24	10.1.4.2	2

R3:

Destination	Gateway	Cost
10.1.3.0/24	C	1
10.1.6.0/24	C	1
10.1.1.0/24	10.1.3.1	2
10.1.4.0/24	10.1.3.1	2
10.1.5.0/24	10.1.3.1	2
10.1.2.0/24	10.1.3.1	3
10.1.8.0/24	10.1.6.2	2
10.1.7.0/24	10.1.6.2	3

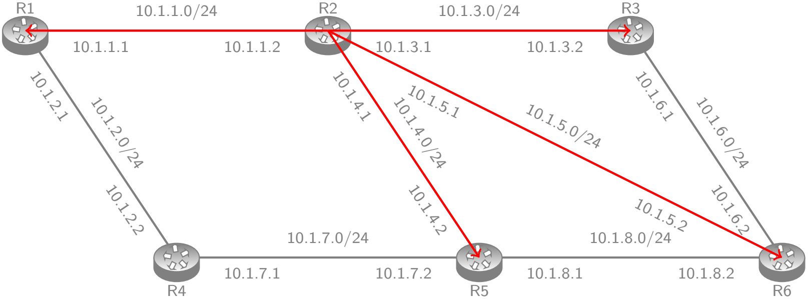

R4:

Destination	Gateway	Cost
10.1.2.0/24	C	1
10.1.7.0/24	C	1
10.1.1.0/24	10.1.2.1	2
10.1.4.0/24	10.1.7.2	2
10.1.8.0/24	10.1.7.2	2
10.1.3.0/24	10.1.7.2	3
10.1.5.0/24	10.1.7.2	3

R5:

Destination	Gateway	Cost
10.1.4.0/24	C	1
10.1.7.0/24	C	1
10.1.8.0/24	C	1
10.1.1.0/24	10.1.4.1	2
10.1.3.0/24	10.1.4.1	2
10.1.5.0/24	10.1.4.1	2
10.1.2.0/24	10.1.7.1	2
10.1.6.0/24	10.1.8.2	2

R6:

Destination	Gateway	Cost
10.1.5.0/24	C	1
10.1.6.0/24	C	1
10.1.8.0/24	C	1
10.1.1.0/24	10.1.5.1	2
10.1.3.0/24	10.1.5.1	2
10.1.4.0/24	10.1.5.1	2
10.1.2.0/24	10.1.5.1	3
10.1.7.0/24	10.1.8.1	2

RIP Example

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.2.0/24	C	1
10.1.3.0/24	10.1.1.2	2
10.1.4.0/24	10.1.1.2	2
10.1.5.0/24	10.1.1.2	2
10.1.7.0/24	10.1.2.2	2
10.1.6.0/24	10.1.1.2	3
10.1.8.0/24	10.1.1.2	3

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.3.0/24	C	1
10.1.4.0/24	C	1
10.1.5.0/24	C	1
10.1.2.0/24	10.1.1.1	2
10.1.6.0/24	10.1.3.2	2
10.1.7.0/24	10.1.4.2	2
10.1.8.0/24	10.1.4.2	2

Destination	Gateway	Cost
10.1.3.0/24	C	1
10.1.6.0/24	C	1
10.1.1.0/24	10.1.3.1	2
10.1.4.0/24	10.1.3.1	2
10.1.5.0/24	10.1.3.1	2
10.1.2.0/24	10.1.3.1	3
10.1.8.0/24	10.1.6.2	2
10.1.7.0/24	10.1.6.2	3

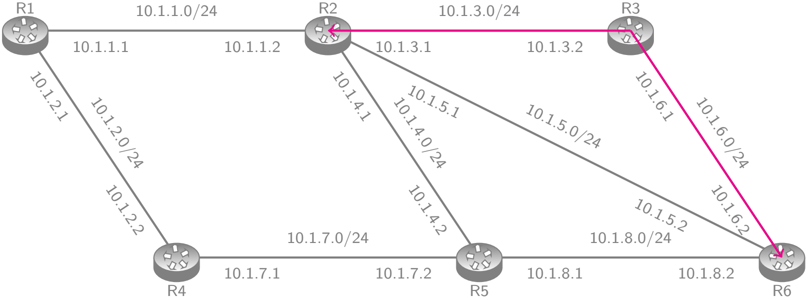

Destination	Gateway	Cost
10.1.2.0/24	C	1
10.1.7.0/24	C	1
10.1.1.0/24	10.1.2.1	2
10.1.4.0/24	10.1.7.2	2
10.1.8.0/24	10.1.7.2	2
10.1.3.0/24	10.1.7.2	3
10.1.5.0/24	10.1.7.2	3

Destination	Gateway	Cost
10.1.4.0/24	C	1
10.1.7.0/24	C	1
10.1.8.0/24	C	1
10.1.1.0/24	10.1.4.1	2
10.1.3.0/24	10.1.4.1	2
10.1.5.0/24	10.1.4.1	2
10.1.2.0/24	10.1.7.1	2
10.1.6.0/24	10.1.8.2	2

Destination	Gateway	Cost
10.1.5.0/24	C	1
10.1.6.0/24	C	1
10.1.8.0/24	C	1
10.1.1.0/24	10.1.5.1	2
10.1.3.0/24	10.1.5.1	2
10.1.4.0/24	10.1.5.1	2
10.1.2.0/24	10.1.5.1	3
10.1.7.0/24	10.1.8.1	2

RIP Example

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.2.0/24	C	1
10.1.3.0/24	10.1.1.2	2
10.1.4.0/24	10.1.1.2	2
10.1.5.0/24	10.1.1.2	2
10.1.7.0/24	10.1.2.2	2
10.1.6.0/24	10.1.1.2	3
10.1.8.0/24	10.1.1.2	3

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.3.0/24	C	1
10.1.4.0/24	C	1
10.1.5.0/24	C	1
10.1.2.0/24	10.1.1.1	2
10.1.6.0/24	10.1.3.2	2
10.1.7.0/24	10.1.4.2	2
10.1.8.0/24	10.1.4.2	2

Destination	Gateway	Cost
10.1.3.0/24	C	1
10.1.6.0/24	C	1
10.1.1.0/24	10.1.3.1	2
10.1.4.0/24	10.1.3.1	2
10.1.5.0/24	10.1.3.1	2
10.1.2.0/24	10.1.3.1	3
10.1.8.0/24	10.1.6.2	2
10.1.7.0/24	10.1.6.2	3

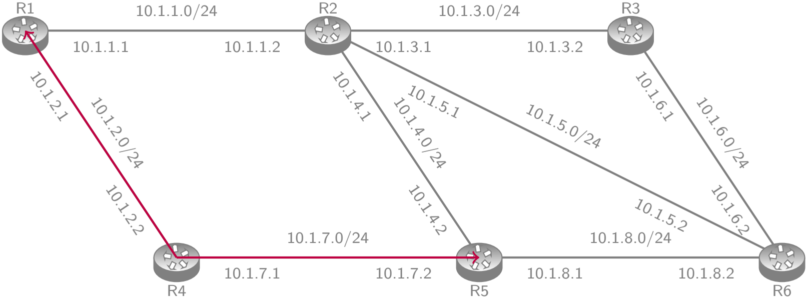

Destination	Gateway	Cost
10.1.2.0/24	C	1
10.1.7.0/24	C	1
10.1.1.0/24	10.1.2.1	2
10.1.4.0/24	10.1.7.2	2
10.1.8.0/24	10.1.7.2	2
10.1.3.0/24	10.1.7.2	3
10.1.5.0/24	10.1.7.2	3

Destination	Gateway	Cost
10.1.4.0/24	C	1
10.1.7.0/24	C	1
10.1.8.0/24	C	1
10.1.1.0/24	10.1.4.1	2
10.1.3.0/24	10.1.4.1	2
10.1.5.0/24	10.1.4.1	2
10.1.2.0/24	10.1.7.1	2
10.1.6.0/24	10.1.8.2	2

Destination	Gateway	Cost
10.1.5.0/24	C	1
10.1.6.0/24	C	1
10.1.8.0/24	C	1
10.1.1.0/24	10.1.5.1	2
10.1.3.0/24	10.1.5.1	2
10.1.4.0/24	10.1.5.1	2
10.1.2.0/24	10.1.5.1	3
10.1.7.0/24	10.1.8.1	2

RIP Example

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.2.0/24	C	1
10.1.3.0/24	10.1.1.2	2
10.1.4.0/24	10.1.1.2	2
10.1.5.0/24	10.1.1.2	2
10.1.7.0/24	10.1.2.2	2
10.1.6.0/24	10.1.1.2	3
10.1.8.0/24	10.1.1.2	3

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.3.0/24	C	1
10.1.4.0/24	C	1
10.1.5.0/24	C	1
10.1.2.0/24	10.1.1.1	2
10.1.6.0/24	10.1.3.2	2
10.1.7.0/24	10.1.4.2	2
10.1.8.0/24	10.1.4.2	2

Destination	Gateway	Cost
10.1.3.0/24	C	1
10.1.6.0/24	C	1
10.1.1.0/24	10.1.3.1	2
10.1.4.0/24	10.1.3.1	2
10.1.5.0/24	10.1.3.1	2
10.1.2.0/24	10.1.3.1	3
10.1.8.0/24	10.1.6.2	2
10.1.7.0/24	10.1.6.2	3

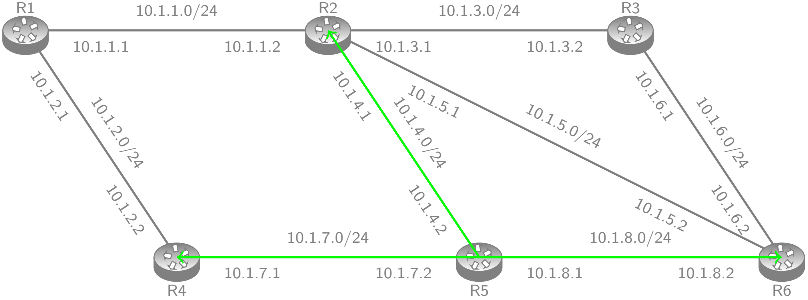

Destination	Gateway	Cost
10.1.2.0/24	C	1
10.1.7.0/24	C	1
10.1.1.0/24	10.1.2.1	2
10.1.4.0/24	10.1.7.2	2
10.1.8.0/24	10.1.7.2	2
10.1.3.0/24	10.1.7.2	3
10.1.5.0/24	10.1.7.2	3
10.1.6.0/24	10.1.7.2	3

Destination	Gateway	Cost
10.1.4.0/24	C	1
10.1.7.0/24	C	1
10.1.8.0/24	C	1
10.1.1.0/24	10.1.4.1	2
10.1.3.0/24	10.1.4.1	2
10.1.5.0/24	10.1.4.1	2
10.1.2.0/24	10.1.7.1	2
10.1.6.0/24	10.1.8.2	2

Destination	Gateway	Cost
10.1.5.0/24	C	1
10.1.6.0/24	C	1
10.1.8.0/24	C	1
10.1.1.0/24	10.1.5.1	2
10.1.3.0/24	10.1.5.1	2
10.1.4.0/24	10.1.5.1	2
10.1.2.0/24	10.1.5.1	3
10.1.7.0/24	10.1.8.1	2

RIP Example

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.2.0/24	C	1
10.1.3.0/24	10.1.1.2	2
10.1.4.0/24	10.1.1.2	2
10.1.5.0/24	10.1.1.2	2
10.1.7.0/24	10.1.2.2	2
10.1.6.0/24	10.1.1.2	3
10.1.8.0/24	10.1.1.2	3

Destination	Gateway	Cost
10.1.1.0/24	C	1
10.1.3.0/24	C	1
10.1.4.0/24	C	1
10.1.5.0/24	C	1
10.1.2.0/24	10.1.1.1	2
10.1.6.0/24	10.1.3.2	2
10.1.7.0/24	10.1.4.2	2
10.1.8.0/24	10.1.4.2	2

Destination	Gateway	Cost
10.1.3.0/24	C	1
10.1.6.0/24	C	1
10.1.1.0/24	10.1.3.1	2
10.1.4.0/24	10.1.3.1	2
10.1.5.0/24	10.1.3.1	2
10.1.2.0/24	10.1.3.1	3
10.1.8.0/24	10.1.6.2	2
10.1.7.0/24	10.1.6.2	3

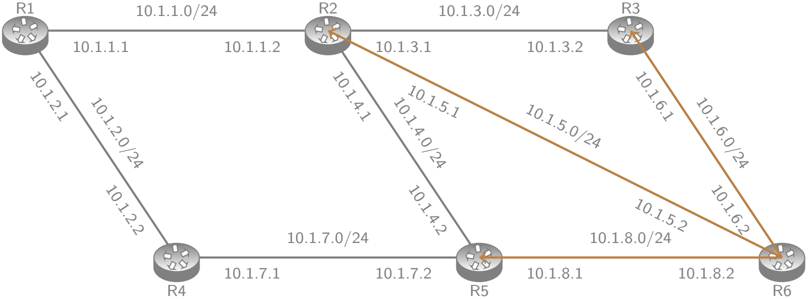

Destination	Gateway	Cost
10.1.2.0/24	C	1
10.1.7.0/24	C	1
10.1.1.0/24	10.1.2.1	2
10.1.4.0/24	10.1.7.2	2
10.1.8.0/24	10.1.7.2	2
10.1.3.0/24	10.1.7.2	3
10.1.5.0/24	10.1.7.2	3
10.1.6.0/24	10.1.7.2	3

Destination	Gateway	Cost
10.1.4.0/24	C	1
10.1.7.0/24	C	1
10.1.8.0/24	C	1
10.1.1.0/24	10.1.4.1	2
10.1.3.0/24	10.1.4.1	2
10.1.5.0/24	10.1.4.1	2
10.1.2.0/24	10.1.7.1	2
10.1.6.0/24	10.1.8.2	2

Destination	Gateway	Cost
10.1.5.0/24	C	1
10.1.6.0/24	C	1
10.1.8.0/24	C	1
10.1.1.0/24	10.1.5.1	2
10.1.3.0/24	10.1.5.1	2
10.1.4.0/24	10.1.5.1	2
10.1.2.0/24	10.1.5.1	3
10.1.7.0/24	10.1.8.1	2

Maximizing Hop Count

- Routing information is broadcast between RIP neighbors every 30 seconds.
- Routing information requires 8 minutes to reach 16 hops away routers (worst case).
- Any network that is more than 15 hops away is treated as unreachable and assigned a hop count equal to infinity
- Infinity has the value 16
- This maximum hop count is called the network diameter.

Routing Loops

- A condition in which a packet is continuously transmitted within a series of routers without ever reaching its destination.
- IP TTL was designed to drop looping packets
- Mechanisms used to prevent against routing loops
  - Split Horizon
  - Split Horizon with Poison Reverse

Split Horizon

- A router should not advertise a network through the interface from which the update came

Split Horizon with Poison Reverse

- The rule states that once a router learns of an unreachable route through an interface, advertise it as unreachable back through the same interface
- It is an exception to the Split Horizon rule

RIP Updates

- Periodic Updates
  - Entire routing table is sent
  - Update timer
  - RIP Jitter
  - The update timer is offset by a small random time (+/- 0 to 5 seconds) to prevent synchronized updates
- Triggered Updates
  - Selected entries of the routing table are sent
  - Interface changes state
  - Route becomes unreachable
  - Route is placed in routing table

RIP Request Messages

- Request entire routing table
  - If there is exactly one entry in the request, and it has an address family identifier of zero and a metric of infinity (i.e., 16), then this is a request to send the entire routing table.
- Request a part of routing table
  - Only requested entries are sent back
  - If a network is inaccessible, it is sent back with a metric 16

RIP Timers

- Update Timer
  - Default value 30 seconds
  - RIP router sends an unsolicited Response message containing the complete routing table to every neighboring router.
  - The update timer is offset by a small random time (+/- 0 to 5 seconds) each time it is set to avoid synchronized updates.
- Triggered Timer
  - Triggered updates can cause excessive load on networks
  - After a triggered update is sent, a timer should be set for a random interval between 1 and 5 seconds.
  - If other changes that would trigger updates occur before the timer expires, a single update is triggered when the timer expires.
  - A triggered update should be suppressed if a regular update is due by the time the triggered update would be sent.

RIP Route Timers

- Route Timeout Timer
  - Also called Invalid Timer
  - Associated with each route
  - Initialized when a route is established, and any time an update message is received for the route.
  - Default value 180 seconds
  - Upon expiration of the timeout, the route is no longer valid
  - It is retained in the routing table for a short time so that neighbors can be notified that the route has been dropped.
- Route Garbage Collection Timer
  - Also called Flush Timer
  - Associated with each route
  - Default value 120 seconds
  - Initialized when a route is no longer valid
  - Upon expiration of the garbage-collection time, the route is removed from the routing table

RIPv1

- RFC 1058
- UDP Port 520
- Based on the program "routed", distributed with the 4.3 Berkeley Software Distribution.
- RIPv1 is a classful routing protocol
- Uses broadcast address 255.255.255.255 to send updates to neighbors.
- The maximum datagram size is 512 octets (RIP-PDU).

RIPv2

- RFC 2453
- UDP Port 520
- RIPv2 is a classless routing protocol
- Uses multicast address 224.0.0.9 to send updates to neighbors.
- The maximum datagram size is 512 octets (RIP-PDU).
- RIPv2 supports authentication

RIP Packets

32 bits (RIPv1 Packet)

command	version	must be zero
address family identifier		must be zero
IPv4 address		
must be zero		
must be zero		
metric		

32 bits (RIPv2 Packet)

command	version	must be zero
Address Family Identifier		Route Tag
IPv4 address		
Subnet Mask		
Next Hop		
Metric		

RIP Packet Fields

- Command
  - 1: RIP Request Packet
  - 2: RIP Response Packet
- Version
  - 1: RIPv1
  - 2: RIPv2

RIPv2 Entry Fields

- Address Family Identifier
  - 0: Request entire routing table
    - Single RIP entry
    - Metric is set to Infinity (16)
  - 2: AF_INET (IPv4 addresses)
- IPv4 Address
  - Address of the network/subnetwork/host
- Subnet Mask
  - Subnet mask the network/subnetwork/host
- Next Hop
  - Should be put to 0.0.0.0 if it is the advertising router
- Metric
  - Distance to reach the network/subnetwork/host

RIP-2 Entry

- Each RIP packet contains between 1 and 25 (inclusive) RIP entries
- If Authentication is used, it should be the first RIP entry.

Address Family Identifier	Route Tag
IPv4 address	
Subnet Mask	
Next Hop	
Metric	

0XFFFF	Authentication Type
Authentication	

RIP

- Interior Routing Protocol
  - Used inside Autonomous Systems
- Distance Vector Routing Protocol
  - Sends its routing table periodically
- Limited to small networks
  - Cannot work in networks where the diameter is greater than 16 routers.

Link State Routing

- Each router advertises the states of its local network links to its neighbors.
- These link state advertisements are then distributed to all other routers.
- The end result is that all routers obtain the same database of these link state advertisements.
- These link state advertisements describe the current map of the network.
- From the network map, each router then runs the Dijkstra algorithm to determine the shortest path to each destination.
- OSPF "Open Shortest Path First"

OSPF

- RFC 2328
- Version 2
- Open stands for non-proprietary publicly available specification
- Developed by the OSPF working group of the Internet Engineering Task Force.
- Carried in IP (Protocol: 89)
- Uses Multicast address: 224.0.0.5

OSPF Operation

- Neighbor Discovery
- Link State Database (LSDB) Synchronization
- Link State Requests
- Link State Updates
- Link State Acknowldgements

OSPF Packet Header

32 bits

Version	Type	Packet length	
Router ID			
Area ID			
Checksum		AuType	
Authentication			
Authentication			

OSPF Neighbor Discovery

- A router discovers neighbors by periodically sending OSPF Hello packets out all of its interfaces.
- By default, a router sends Hello message out an interface every 10 seconds (HelloInterval).
- If no Hello message was recieved from a neighbor for a RouterDeadInterval, the neighbor is considered as down
- Default value for HelloInterval: 10 seconds
- Default value for RouterDeadInterval: 40 seconds

OSPF Hello Packet

32 bits

Version	Type	Packet length	
Router ID			
Area ID			
Checksum		AuType	
Authentication			
Authentication			
Network Mask			
HelloInterval		Options	Rtr Pri
RouterDeadInterval			
Designated Router			
Backup Designated Router			
Neighbor			

Database Synchronization

- If a new neighbor is discovered, a synchronization between the two topology databases is performed
- Only keys are exchanged between the two neighbors
- Master/Slave relationship is established between the two neighbors
- The router with the high router-id will be the master
- The master sends a packet with a random sequence number X containing a list of keys
- The packet size should be less than the maximum transmission unit.
- The slave replies with a packet conatining a list of its keys with the same sequence number X
- When the Slave reply is recieved, the Master sends a new list of its keys and increments the sequence number
- When the Slave reply is not received, the Master retransmits the last sent packet again after the expiration of RxmtTimer
- If the slave receives a duplicate packet, it must retransmit its last sent packet to the master.

OSPF Database Description packet

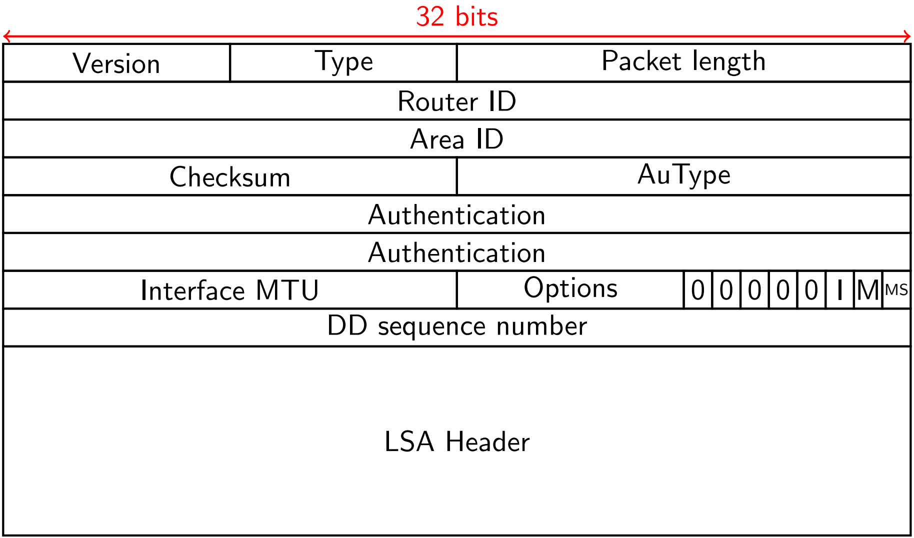

Link State Request

- When a router receives a new or a more recent item from its neighbor, it will send a request to the neighbor to ask for that LSA
- If a reply to that request is not received until the expiration of RxmtTimer, the request is retransmitted again to the neighbor

OSPF Link State Request Packet

32 bits

Version	Type	Packet length
Router ID		
Area ID		
Checksum		AuType
Authentication		
Authentication		
LS type		
Link State ID		
Advertising Router		

Link State Update

- Responses to Link State Requests are sent using Link State Updates
- Link State Updates sent as a reply to Link State Requests are not acknoweldged
- Triggered Link State Updates are sent to adjacent neighbors when there is a change in the topology database
- Triggered Link State Updates are retransmitted after the expiration of RxmtTimer

Link State Update Packet

32 bits

<table>
<tr><td>Version</td><td>Type</td><td colspan="2">Packet length</td></tr>
<tr><td colspan="4">Router ID</td></tr>
<tr><td colspan="4">Area ID</td></tr>
<tr><td colspan="2">Checksum</td><td colspan="2">AuType</td></tr>
<tr><td colspan="4">Authentication</td></tr>
<tr><td colspan="4">Authentication</td></tr>
<tr><td colspan="4">Number of LSAs</td></tr>
<tr><td colspan="4">LSA</td></tr>
</table>

Link State Ack

- Acknowldgements are sent as a reply to triggered Link State Updates

OSPF Link State Acknowledgment packet

32 bits

Version	Type	Packet length	
Router ID			
Area ID			
Checksum		AuType	
Authentication			
Authentication			
LSA Header			

Network Types

- Point-to-Point Networks
- Point-to-Multipoint Networks
- Broadcast Networks
- Non-Broadcast Multiple Access Networks

Inter-Area Routing

- The autonomous system is divided in many areas
- Each area runs a copy of OSPF
- Backbone Area (Area 0) connects all areas together
- Backbone Area is a contiguous area
- A virtual link may be used to connect a distant area
- Inter-Area routing uses distance vector approach

Router Types

- Internal Router: belongs to a given area
- Area Border Router: connects an area to the Backbone area
- Autonomous System Border Router: connects the AS to another AS
- Backbone Router: belongs to the backbone area

Link State Advertisements (LSAs)

- Each router in the Autonomous System originates one or more link state advertisements (LSAs).
- The collection of LSAs forms the link-state database.
- Each separate type of LSA has a separate function.
  - Router-LSA
  - Network-LSA
  - Type 3 Summary-LSA
  - Type 4 Summary-LSA
  - AS-External-LSA
- Router-LSAs and network-LSAs describe how an area's routers and networks are interconnected.
- Summary-LSAs provide a way of condensing an area's routing information.
- AS-external-LSAs provide a way of transparently advertising externally-derived routing information throughout the Autonomous System.

LSA Header

- Each LSA begins with a standard 20-byte header.
- The LSA header contains the LS type, Link State ID and Advertising Router fields.
- The combination of these three fields uniquely identifies the LSA.

LSA Header

32 bits

LS Age	Options	LS Type
Link State ID		
Advertising Router		
LS sequence number		
LS checksum	length	

Router-LSA

- Originated by all routers.
- This LSA describes the collected states of the router's interfaces to an area.
- Flooded throughout a single area only.

Router-LSA

32 bits

<table>
<tr><td colspan="9">LS Age</td><td>Options</td><td>LS Type</td></tr>
<tr><td colspan="11">Link State ID</td></tr>
<tr><td colspan="11">Advertising Router</td></tr>
<tr><td colspan="11">LS sequence number</td></tr>
<tr><td colspan="9">LS checksum</td><td colspan="2">length</td></tr>
<tr><td>0</td><td>0</td><td>0</td><td>0</td><td>0</td><td>V</td><td>E</td><td>B</td><td>0</td><td colspan="2">number of links</td></tr>
<tr><td colspan="11">Link ID</td></tr>
<tr><td colspan="11">Link Data</td></tr>
<tr><td colspan="8">Type</td><td>number of TOS</td><td colspan="2">metric</td></tr>
<tr><td colspan="8">TOS</td><td>0</td><td colspan="2">TOS metric</td></tr>
</table>

Network-LSA

- Originated for broadcast and NBMA (Non-Broadcast Multiple Access) networks by the Designated Router.
- This LSA contains the list of routers connected to the network.
- Flooded throughout a single area only.

Network-LSA

32 bits

<table>
<tr><td colspan="2">LS Age</td><td>Options</td><td>LS Type</td></tr>
<tr><td colspan="4">Link State ID</td></tr>
<tr><td colspan="4">Advertising Router</td></tr>
<tr><td colspan="4">LS sequence number</td></tr>
<tr><td colspan="2">LS checksum</td><td colspan="2">length</td></tr>
<tr><td colspan="4">Network Mask</td></tr>
<tr><td colspan="4">Attached Router</td></tr>
<tr><td colspan="4">...</td></tr>
</table>

Type 3 Summary-LSA

- Originated by area border routers, and flooded throughout the LSA's associated area.
- Each summary-LSA describes a route to a destination outside the area, yet still inside the AS (i.e., an inter-area route).
- Type 3 summary-LSAs describe routes to networks.

Type 3 Summary-LSA

32 bits

LS Age		Options	LS Type
Link State ID			
Advertising Router			
LS sequence number			
LS checksum		length	
Network Mask			
0	metric		
TOS	metric		
...			

Type 4 Summary-LSA

- Originated by area border routers, and flooded throughout the LSA's associated area.
- Each summary-LSA describes a route to a destination outside the area, yet still inside the AS (i.e., an inter-area route).
- Type 4 summary-LSAs describe routes to AS boundary routers.

Type 4 Summary-LSA

32 bits

LS Age		Options	LS Type
Link State ID			
Advertising Router			
LS sequence number			
LS checksum		length	
Network Mask			
0	metric		
TOS	metric		
...			

AS-External-LSA

- Originated by AS boundary routers, and flooded throughout the AS.
- Each AS-external-LSA describes a route to a destination in another Autonomous System.
- Default routes for the AS can also be described by AS-external-LSAs.

AS-External-LSA

32 bits

LS Age	Options	LS Type						
Link State ID								
Advertising Router								
LS sequence number								
LS checksum	length							
Network Mask								
E	0	0	0	0	0	0	0	metric
Forwarding address								
External Route Tag								
E	TOS	TOS metric						
Forwarding address								
...								

Exterior Gateway Protocol

- Routing between Autonomous Systems
- EGP: Exterior Gateway Protocol
- BGP: Border Gateway Protocol

BGP

- RFC 4271
- Inter-Autonomous System routing protocol
- Version 4
- Support for Classless Interdomain Routing (CIDR)
- Is a Policy-Based routing protocol
- iBGP: internal BGP (between peers in the same AS)
- eBGP: external BGP (between peers in different ASs)

BGP

- BGP uses TCP as its transport protocol
- Port 179
- No need to implement the following functions
  - fragmentation
  - retransmission
  - acknowledgement
  - sequencing

BGP Messages

- OPEN Message: Establish a peering session.
- UPDATE Message: Announcing new routes or withdrawing previously announced routes.
- KEEPALIVE Message: Handshake at regular intervals.
- NOTIFICATION Message: Shuts down a peering session.

Routing Information Base

- Adj-RIBs-In
  - The Adj-RIBs-In stores routing information learned from BGP peers
  - Received in UPDATE messages
- Loc-RIB
  - Local Routing Information Base
  - Selected by applying local policies to the routing information contained in the Adj-RIBs-In
- Adj-RIBs-Out
  - The Adj-RIBs-Out stores routing information selected for advertisement to BGP peers
  - Sent in UPDATE messages

BGP Header

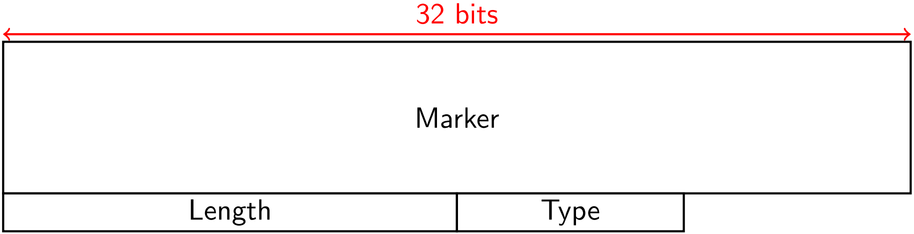

- Marker: (16 bytes) it is included for compatibility; it MUST be set to all 1s.
- Length: (2 bytes) it indicates the total length of the message, including the header in bytes.
  - $19 \leq \text{Length} \leq 4096$
- Type: (1 byte) it indicates the type code of the message.
  - 1: OPEN
  - 2: UPDATE
  - 3: NOTIFICATION
  - 4: KEEPALIVE

BGP OPEN Message

- After a TCP connection is established, the first message sent by each side is an OPEN message.
- If the OPEN message is acceptable, a KEEPALIVE message confirming the OPEN is sent back.

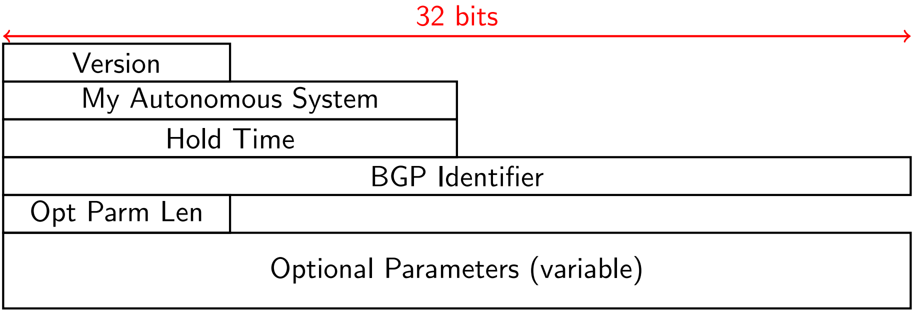

BGP UPDATE Message

- UPDATE messages are used to transfer routing information between BGP peers.
- The information in the UPDATE message can be used to construct a graph that describes the relationships of the various ASs.
- An UPDATE message is used to advertise feasible routes that share common path attributes to a peer, or to withdraw multiple unfeasible routes from service.
- An UPDATE message MAY simultaneously advertise a feasible route and withdraw multiple unfeasible routes from service.

32 bits

Withdrawn Routes Length	
Withdrawn Routes (variable)	
Total Path Attribute Length	
Path Attributes (variable)	
Network Layer Reachability Information (variable)	

BGP UPDATE Message

- Withdrawn Routes (variable)
    - Length
        - 1 byte
        - Indicates the length in bits of the IP address prefix
        - A length of zero indicates a prefix that matches all IP addresses with prefix of zero bytes
    - Prefix
        - Contains an IP address prefix
        - Followed by the minimum number of trailing bits needed to make the end of the field fall on a byte boundary

Length (1 byte)
Prefix (variable)

BGP UPDATE Message

- Path Attributes (variable)
  - A variable-length sequence of path attributes
  - TLV (Type/Length/Value) Encoding
  - Type
    - 2 bytes
    - Attr. Flags
    - Optional (bit 0): It defines whether the attribute is optional or well-known
    - Transitive (bit 1): It defines whether an optional attribute is transitive or non-transitive (always 1 for well-known)
    - Partial (bit 2): It defines whether the information contained in the optional transitive attribute is partial or complete (always 0 for other cases).
    - Extended (bit 3): It defines whether the Attribute Length is one byte or two bytes
    - Unsed (bits 4-7)
    - Attr. Type Code

BGP UPDATE Message

- Path Attributes (variable)
  - Length
    - 1 byte if the Extended Flag is 0
    - 2 bytes if the Extended Flag is 1
    - Indicates the length of the value field in bytes

BGP Path Attributes

- ORIGIN
  - Type Code: 1
  - Well-known mandatory attribute
  - Value (1 byte)
    - 0: IGP (NLRI is interior to the originating AS)
    - 1: EGP (NLRI learned via the EGP protocol)
    - 2: INCOMPLETE (NLRI learned by some other means)
- AS-PATH
  - Type Code: 2
  - Well-known mandatory attribute
  - Path segment type
    - 1 byte
    - AS-SET: unordered set of ASes traversed by the route
    - AS-SEQUENCE: ordered set of ASes traversed by the route
  - Path segment length
    - 1 byte
    - Contains the number of ASes in the Path segment field
  - Path segment value
    - Contains one or more AS numbers each encoded as a 2-byte field

BGP Path Attributes

- NEXT-HOP
  - Type Code: 3
  - Well-known mandatory attribute
  - Defines the unicast IP address of the router that should be used as the next hop for the NLRI
- MULTI-EXIT-DISC
  - Type Code: 4
  - Optional non-transitive attribute
  - 4 bytes (unsigned integer)
  - May be used by a BGP router to discriminate between multiple entry points to a neighboring AS
- LOCAL-PREF (iBGP)
  - Type Code: 5
  - Well-known attribute
  - 4 bytes (unsigned integer)
  - Used to inform other iBGP peers of the advertising router degree of preference for an advertised route

BGP Path Attributes

- ATOMIC-AGGREGATE
  - Type Code: 6
  - Well-known discretionary attribute
  - length 0
- AGGREGATOR
  - Type Code: 7
  - Optional transitive attribute
  - Length: 6 bytes
  - 2 bytes (AS that formed the aggregate route)
  - 4 bytes (BGP router that formed the aggregate route)

BGP UPDATE Message

- Network Layer Reachability Information (variable)
  - Length
    - 1 byte
    - Indicates the length in bits of the IP address prefix
    - A length of zero indicates a prefix that matches all IP addresses with prefix of zero bytes
  - Prefix
    - Contains an IP address prefix
    - Followed by the minimum number of trailing bits needed to make the end of the field fall on a byte boundary

Length (1 byte)
Prefix (variable)

BGP KEEPALIVE Message

- BGP does not use any TCP-based, keep-alive mechanism to determine if peers are reachable.
- Instead, KEEPALIVE messages are exchanged between peers often enough not to cause the Hold Timer to expire.
- A reasonable maximum time between KEEPALIVE messages would be one third of the Hold Time interval.
- KEEPALIVE messages MUST NOT be sent more frequently than one per second.
- If the negotiated Hold Time interval is zero, then periodic KEEPALIVE messages MUST NOT be sent.
- A KEEPALIVE message consists of only the message header and has a length of 19 octets.

BGP NOTIFICATION Message

- A NOTIFICATION message is sent when an error condition is detected.
- The BGP connection is closed immediately after it is sent.

32 bits

Error code	Error subcode	Data (variable)

BGP Path Attributes

- ORIGIN (mandatory)
- AS-PATH (AS-SETs,AS-SEQUENCEs) (mandatory)
- NEXT-HOP (mandatory)
- MULTI-EXIT-DISC
- LOCAL-PREF (iBGP)
- ATOMIC-AGGREGATE
- AGGREGATOR