

PART XIV


ROUTING: EXTERIOR GATEWAY PROTOCOLS AND AUTONOMOUS SYSTEMS (BGP)

General Principle

Although it is desirable for routers to exchange routing information, it is impractical for all routers in an arbitrarily large internet to participate in a single routing update protocol.

EXTERIOR GATEWAY PROTOCOLS (EGP)

- Consequence: routers must be divided into groups




A Practical Limit On Group Size

It is safe to allow up to a dozen routers to participate in a single routing information protocol across a wide area network; approximately five times as many can safely participate across a set of local area networks.

Router Outside A Group

- Does not participate directly in group's routing information propagation algorithm
- Will not choose optimal routes if it uses a member of the group for general delivery

The Extra Hop Problem



- Non-participating router picks one participating router to use (e.g., R_2)
- Non-participating router routes all packets to R_2 across backbone
- Router R_2 routes some packets back across backbone to R_1

Statement Of The Problem

Treating a group of routers that participate in a routing update protocol as a default delivery system can introduce an extra hop for datagram traffic; a mechanism is needed that allows nonparticipating routers to learn routes from participating routers so they can choose optimal routes.


Solving The Extra Hop Problem

- Not all routers can participate in a single routing exchange protocol (does not scale)
- Even nonparticipating routers should make routing decisions
- Need mechanism that allows nonparticipating routers to obtain correct routing information automatically (without the overhead of participating fully in a routing exchange protocol)

Hidden Networks

- Each site has complex topology
- Nonparticipating router (from another site) cannot attach to all networks

Illustration Of Hidden Networks



- Propagation of route information is independent of datagram routing
- Group must learn routes from nonparticipating routers
- Example: owner of networks 1 and 3 must tell group that there is a route to network 4

A Requirement For Reverse Information Flow

Because an individual organization can have an arbitrarily complex set of networks interconnected by routers, no router from another organization can attach directly to all networks. A mechanism is needed that allows nonparticipating routers to inform the other group about hidden networks.

Autonomous System Concept (AS)

- Group of networks under one administrative authority
- Free to choose internal routing update mechanism
- Connects to one or more other autonomous systems

Modern Internet Architecture

A large TCP/IP internet has additional structure to accommodate administrative boundaries: each collection of networks and routers managed by one administrative authority is considered to be a single autonomous system that is free to choose an internal routing architecture and protocols.

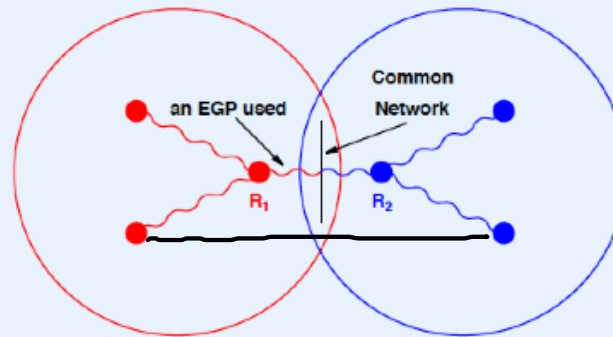
EGPs: Exterior Gateway Protocols

- Originally a single protocol for communicating routes between two autonomous systems
- Now refers to any exterior routing protocol
- Solves two problems
 - Allows router outside a group to advertise networks hidden in another autonomous system
 - Allows router outside a group to learn destinations in the group

Border Gateway Protocol

- The most popular (virtually the only) EGP in use in the Internet
- Current version is BGP-4
- Allows two autonomous systems to communicate routing information
- Supports CIDR (mask accompanies each route)
- Each AS designates a *border router* to speak on its behalf
- Two border routers become *BGP peers*

Illustration Of An EGP (Typically BGP)




Key Characteristics Of BGP

- Provides inter-autonomous system communication
- Propagates reachability information
- Follows next-hop paradigm
- Provides support for policies
- Sends path information
- Permits incremental updates
- Allows route aggregation
- Allows authentication

Additional BGP Facts

- Uses reliable transport (i.e., TCP)
 - Unusual: most routing update protocols use connectionless transport (e.g., UDP)
- Sends *keepalive* messages so other end knows connection is valid (even if no new routing information is needed)



Four BGP Message Types

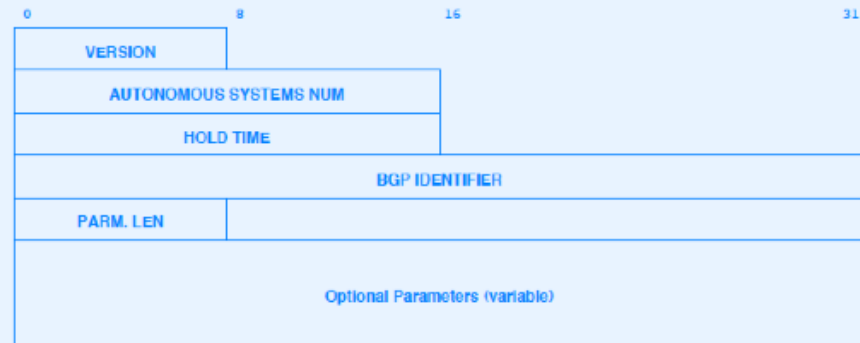
Type Code	Message Type	Description
1	OPEN	Initialize communication
2	UPDATE	Advertise or withdraw routes
3	NOTIFICATION	Response to an incorrect message
4	KEEPALIVE	Actively test peer connectivity

BGP Message Header

0	16	24	31
00	FF	11	EE
22	DD	33	BB
44	AA	55	99
66	88	77	77
LENGTH		TYPE	

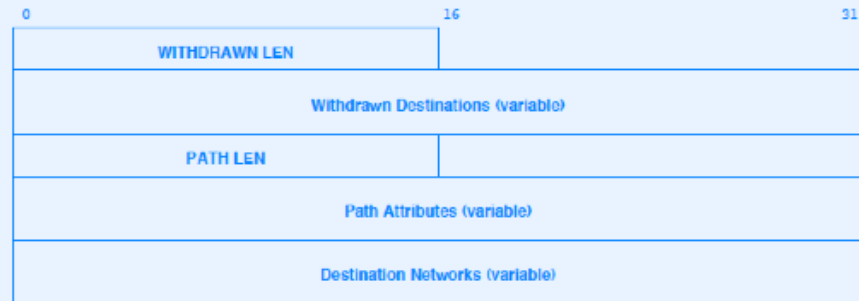
- Each BGP message starts with this header

BGP Open Message



- Used to start a connection
- **HOLD TIME** specifies max time that can elapse between BGP messages

BGP Update Message




- Sender can advertise new routes or withdraw old routes

Compressed Address Entries

- Each route entry consists of address and mask
- Entry can be compressed to eliminate zero bytes

Format Of BGP Address Entry That Permits Compression

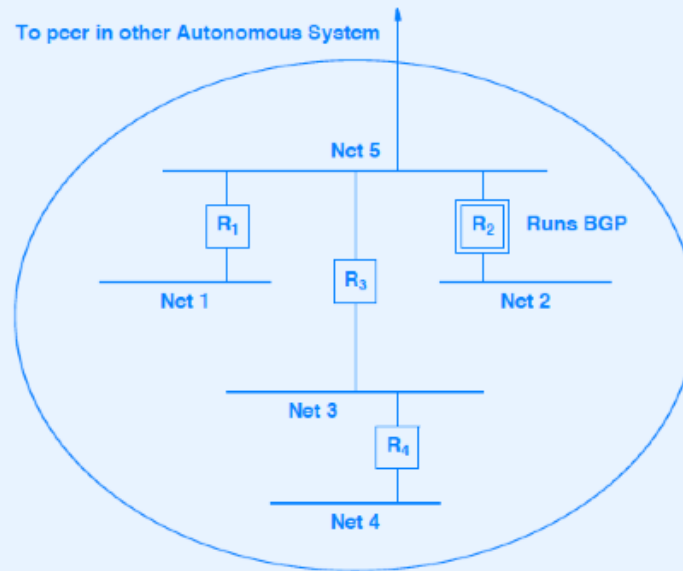


- LEN field specifies size of address that follows

Third-Party Routing Information

- Many routing protocols extract information from the local routing table
- BGP must send information “from the receiver’s perspective”

Example Of Architecture In Which BGP Must Consider Receiver's Perspective



Metric Interpretation

- Each AS can use its own routing protocol
- Metrics differ
 - Hop count
 - Delay
 - Policy-based values
- EGP communicates between two separate autonomous systems

Key Restriction On An EGP

An exterior gateway protocol does not communicate or interpret distance metrics, even if metrics are available.

- Interpretation: “my autonomous system provides a path to this network”

The Point About EGPs

Because an Exterior Gateway Protocol like BGP only propagates reachability information, a receiver can implement policy constraints, but cannot choose a least cost route. A sender must only advertise paths that traffic should follow.

Summary

- Internet is too large for all routers to participate in one routing update protocol
- Group of networks and routers under one administrative authority is called *Autonomous System (AS)*
- Each AS chooses its own interior routing update protocol
- Exterior Gateway Protocol (EGP) is used to communicate routing information between two autonomous systems
- Current exterior protocol is Border Gateway Protocol version 4, BGP-4
- An EGP provides reachability information, but does not associate metrics with each route