ROUTING AND BGP
Lecture 23

COMPSCI 726
Network Defence and Countermeasures

Nalin Asanka Gamagedara Arachchilage

Slides from Muhammad **Rizwan** Asghar

September 22, 2021

Source of some slides: CMU and University of Virginia

HOW CAN WE ROUTE PACKETS?

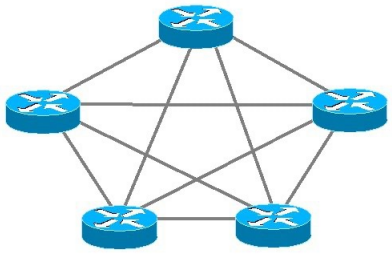

- Source routing
  - Storage: Each node stores routes to every destination
  - Update: Convergence time increases
  - Communication: Total message count increases
  - Issue: It does not scale well!
- Key observation
  - Need less information with increasing distance to destination
- Solution: Dynamic routing

DYNAMIC ROUTING

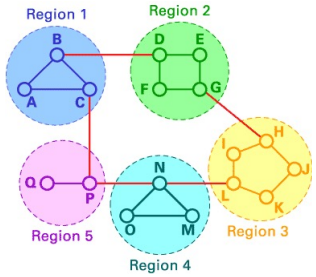

- Choose the best path from a source to the destination
- Best could be based on
  - Smallest number of hops
  - Shortest time delay
  - Least congested
  - Cheapest
  - Administratively allowed
  - Easiest to discover
  - Any combination of the above
- The solution must be quicker and guaranteed to avoid loops and deadlocks

FOR DYNAMIC ROUTING

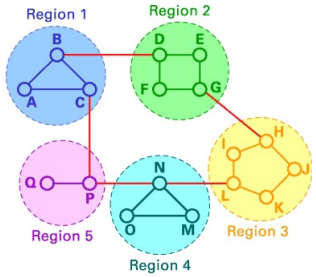

- Divide network into areas
  - Areas can have nested sub-areas
- Hierarchically address nodes in a network
  - Sequentially number top-level areas
  - Sub-areas of area are labelled relative to that area
  - Nodes are numbered relative to the smallest containing area

ROUTING HIERARCHY

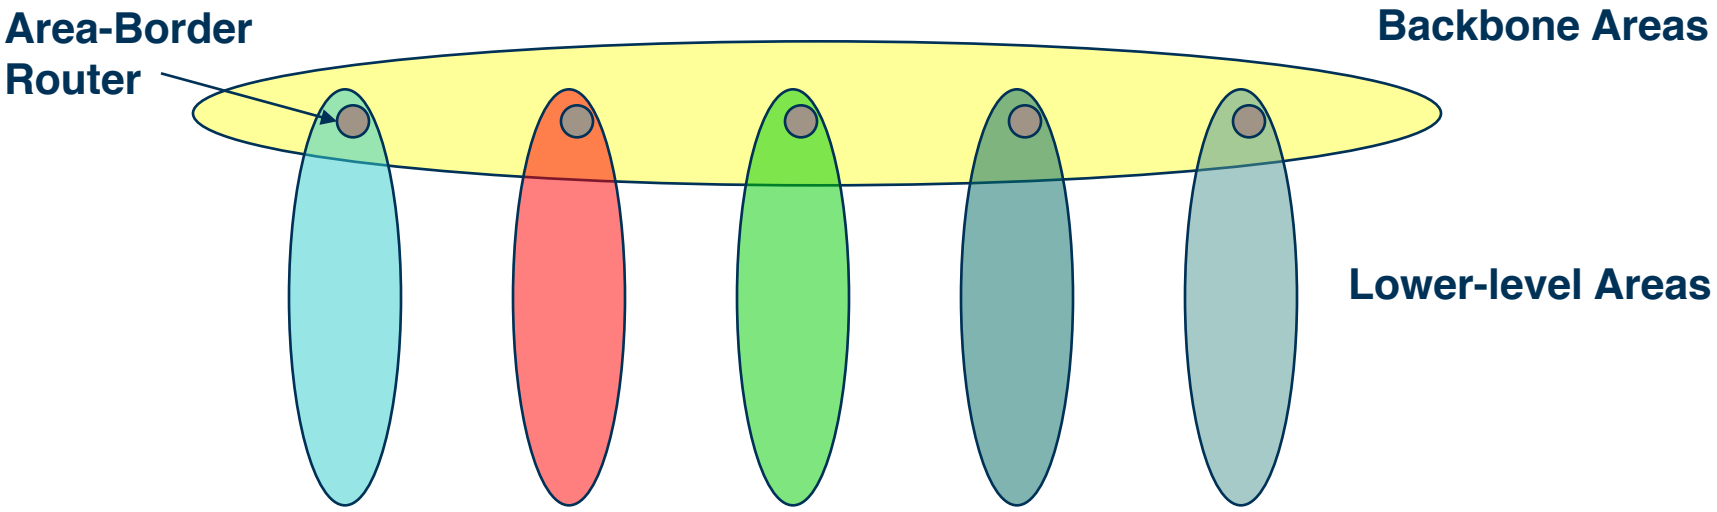

- Partition network into areas
  - Within area
    - Each node has route to every other node
  - Outside area
    - Each node has routes for other top-level areas only
    - Inter-area packets are routed to nearest appropriate border router

AREA HIERARCHY ADDRESSING

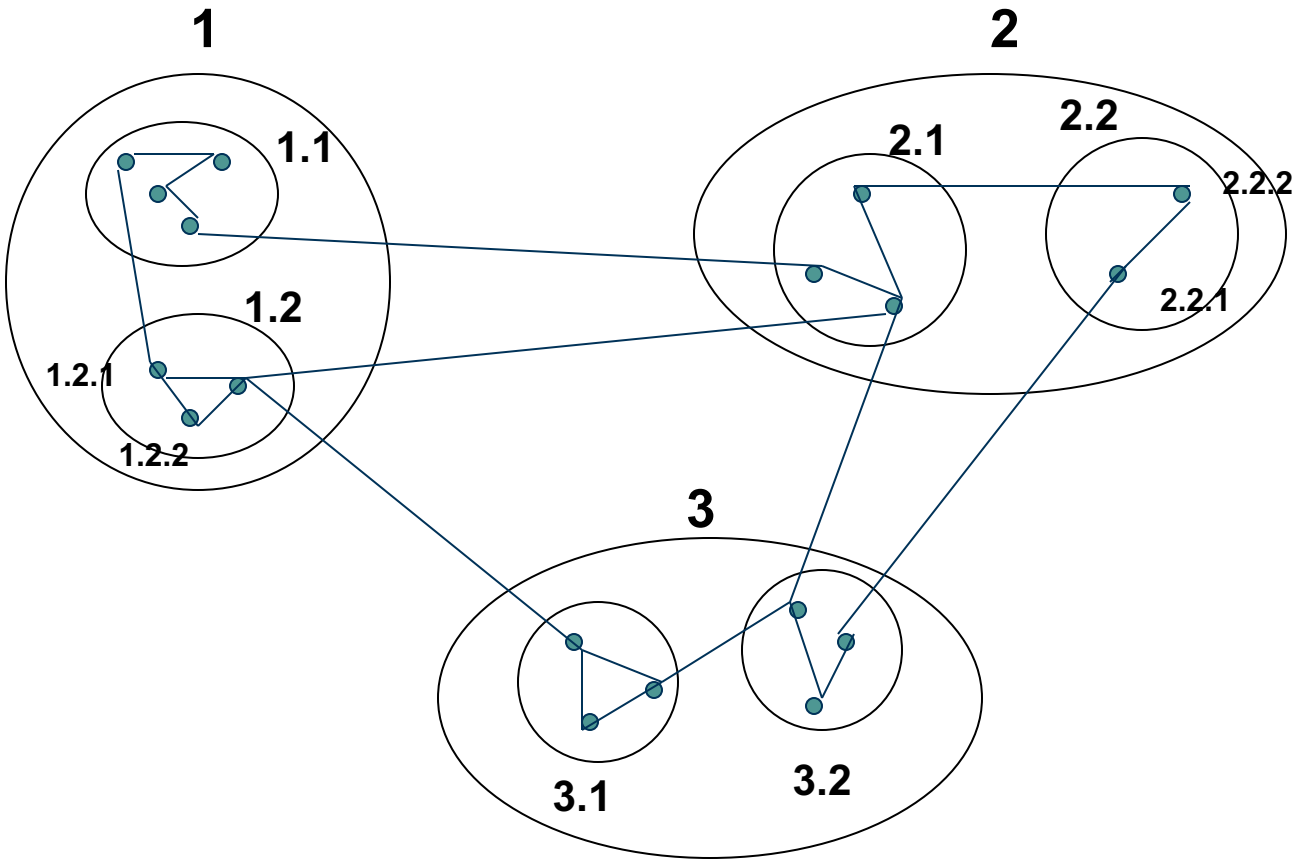

PATH SUB-OPTIMALITY

- Can result in sub-optimal paths

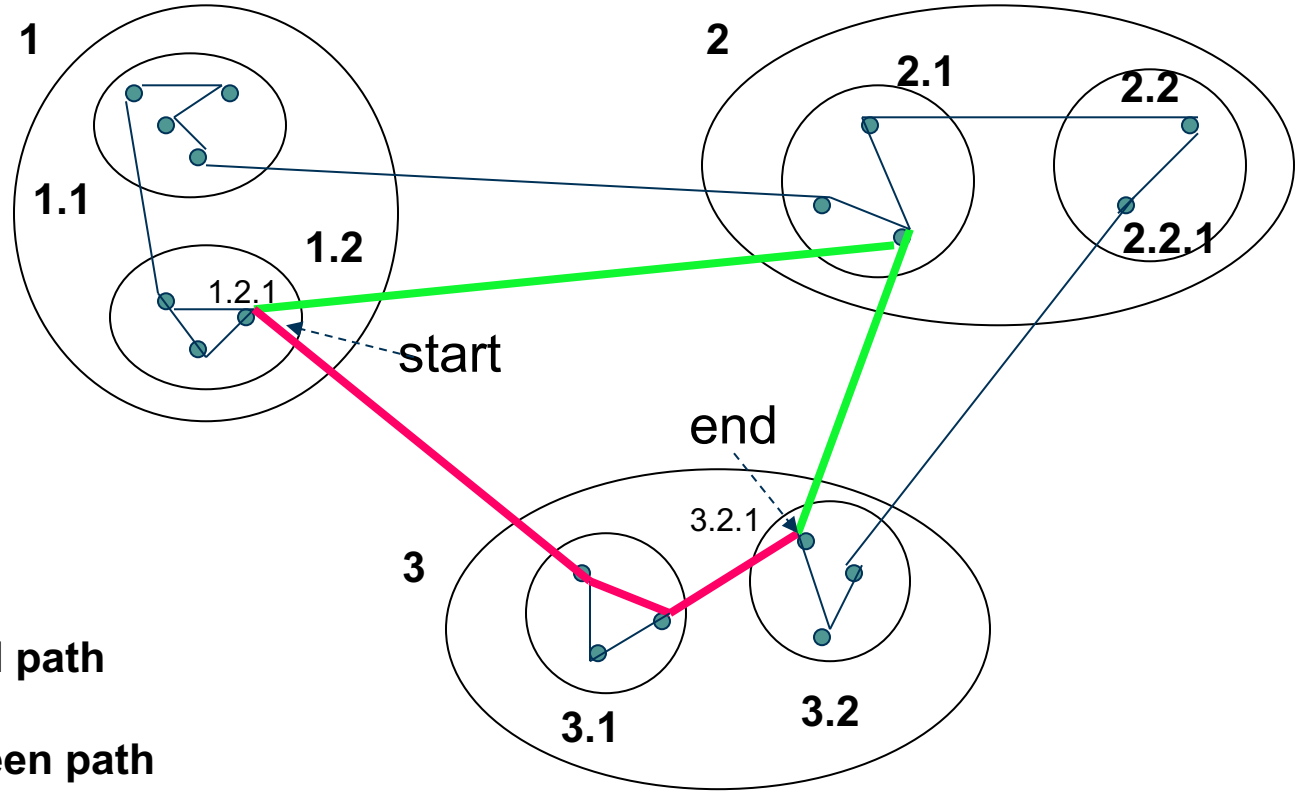

3-hop red path
vs.
2-hop green path

INTERNET'S AREA HIERARCHY

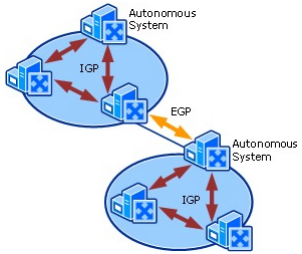

- Based on IP prefix
  - E.g., 192.168.1.5/24
  - E.g., 128.15.10.6/16
- What is an Autonomous System (AS)?
  - A set of routers under a single technical administration
- Interior Gateway Protocol (IGP) and common metrics are used to route packets within the AS
- Exterior Gateway Protocol (EGP) is used to route packets to other ASes
- Each AS is assigned a unique ID

AS NUMBER (ASN)

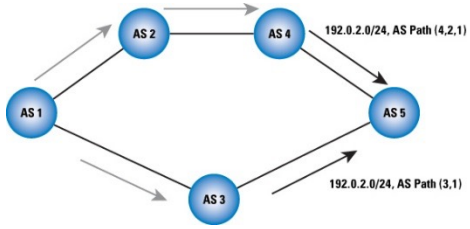

- ASNs represent units of routing
- ASNs are 16-bit values
- Max: 64K possibilities
- 64512 through 65535 are “private”
- Examples
  - Genuity: 1
  - MIT: 3
  - JANET: 786
  - UC San Diego: 7377
  - AT&T: 7018, 6341, 5074, …
  - UUNET: 701, 702, 284, 12199, …
  - Sprint: 1239, 1240, 6211, 6242, …

CURRENT COUNT OF ASNs

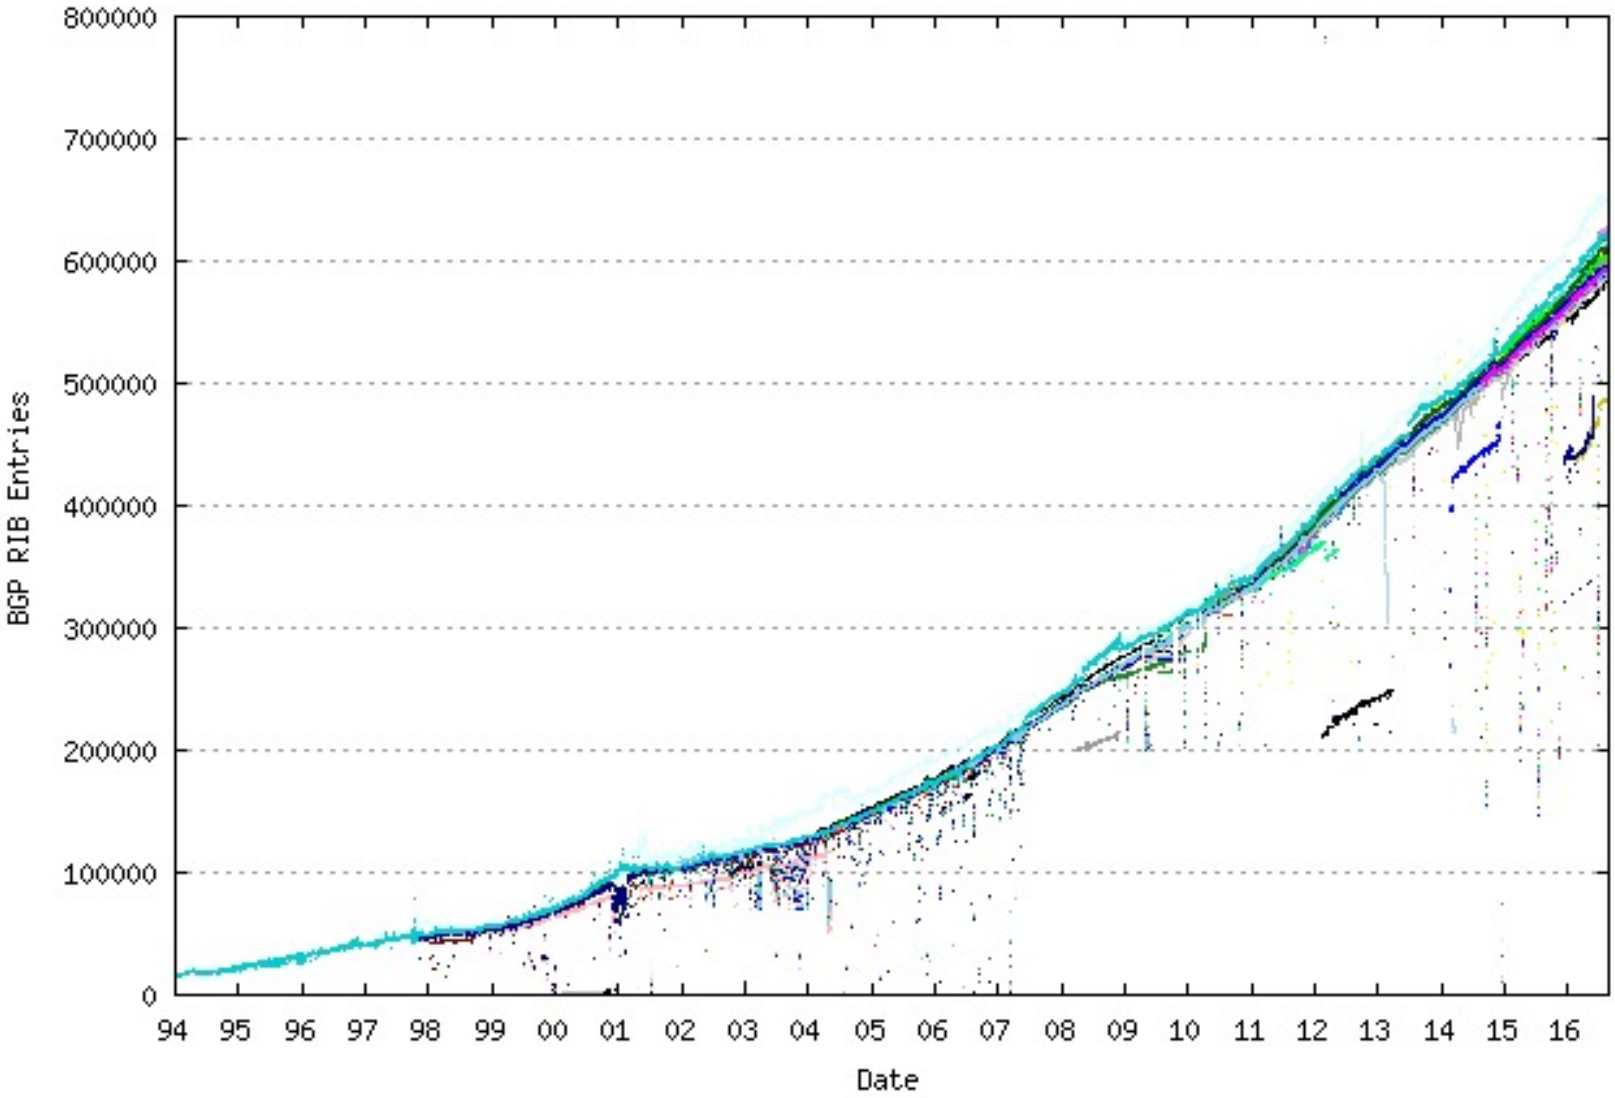

Source: http://bgp.potaroo.net/ (As of August 9, 2016)

IGP VS. EGP

- For routers to communicate within an AS
- Relies on IP address to construct paths
- Routing Information Protocol (RIP)
- Open Shortest Path First (OSPF)

- For routers to communicate among different ASes
- Relies on AS numbers to construct AS paths
- *Border Gateway Protocol (BGP)*

EXAMPLE

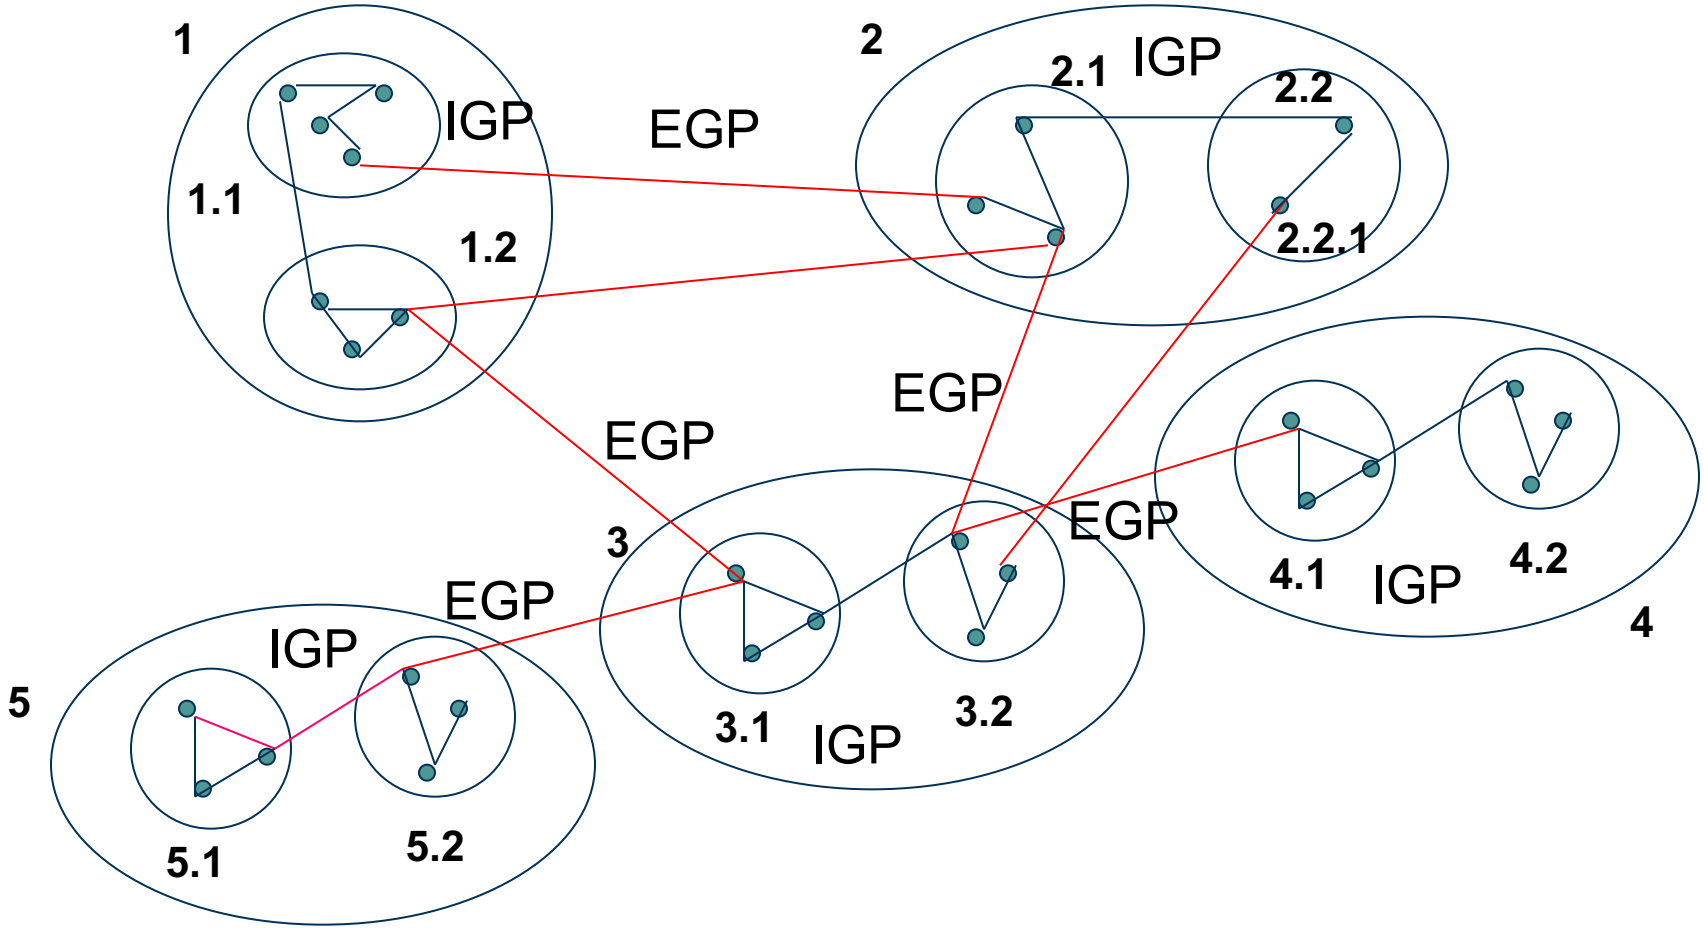
1
IGP
1.1
1.2
EGP
2
2.1
IGP
2.2
2.2.1
EGP
EGP
EGP
3
3.1
IGP
3.2
4
4.1
IGP
4.2
5
IGP
5.1
5.2
EGP

BGP OVERVIEW

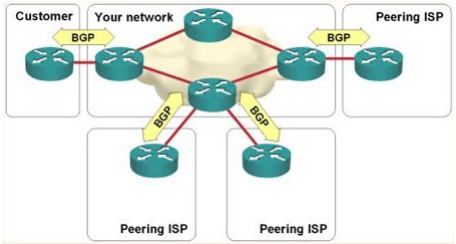

- BGP current version is 4
- Inter-AS routing protocol for exchanging network reachability information among BGP routers
- BGP uses TCP port 179 to send routing messages
- Routing messages in BGP contain complete routes
- Network administrators can specify routing policies

BGP OPERATIONS

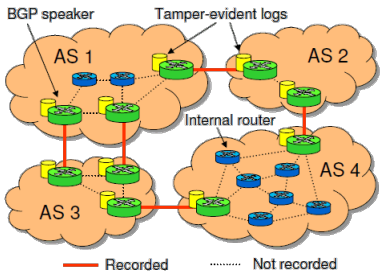

- Two BGP routers exchanging information on a connection are called peers
  - Initially, BGP peers exchange the entire BGP routing table
  - A BGP router retains the current version of the entire BGP routing tables of all of its peers for the duration of the connection
  - Subsequently, only incremental **updates** are sent as the routing tables change
  - **Keepalive** messages are sent periodically to ensure that the connection between the BGP peers is alive
  - **Notification** messages are sent in response to errors or special conditions
  - Routes are stored in the **Routing Information Base** (RIB)

THE BGP DECISION ALGORITHM

- After receiving updates, BGP will have to decide which paths to choose in order to reach a specific destination
- BGP will choose only a single path to reach a specific destination
- The decision process is based on different attributes, such as next hop, local preference, the route origin, and so on
- BGP will always propagate the best path to its neighbours

BGP THREATS / COUNTERMEASURE

- Route manipulation
  - When a malicious device alters the routing table
- Route hijacking
  - When a rogue BGP peer maliciously announces a victim's prefixes in an effort to reroute traffic
- Denial of Service (DoS)
  - When a malicious host sends unexpected or undesirable BGP traffic to a victim
- Solution
  - Use Message Authentication Code (MAC)
  - The secret shared key can be configured manually

RESOURCES

- **Protecting Border Gateway Protocol for the Enterprise**, http://www.cisco.com/c/en/us/about/security-center/protecting-border-gateway-protocol.html

- [Cohen-SIGCOMM16] Cohen, Avichai, Yossi Gilad, Amir Herzberg, and Michael Schapira, **Jumpstarting BGP Security with Path-End Validation**, ACM Special Interest Group on Data Communications (SIGCOMM) 2016

Questions?

Thanks for your attention!