BGP

Overview

- BGP overview
  - When to use BGP?
- BGP terminology
- BGP operation
- Configuring BGP
- Lab scenario: BGP peering

Objectives

- To understand BGP protocol functions
- Describe how to connect an AS
- To be able to configure a router running BGP protocol
- To able to verify the operation of BGP inside the network

BGP overview

Autonomous systems

- An autonomous system (AS) is a collection of networks controlled by a common or single administrator
- Autonomous systems operate using:
  - Interior Gateway Protocol (IGP)
    - RIPv2, EIGRP, OSPF, ISIS
  - Exterior Gateway Protocol (EGP)
    - BGP verision 4 (RFC 1771)

Autonomous systems

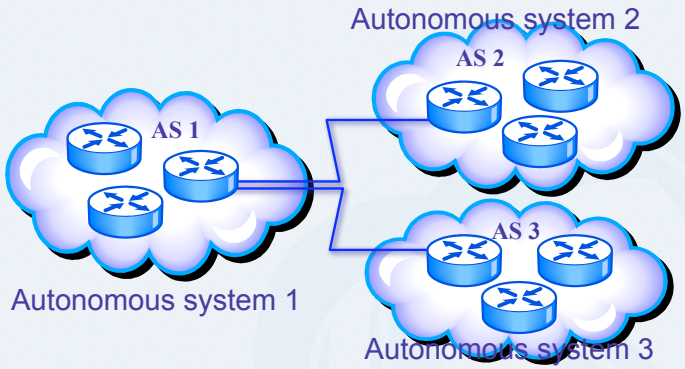

The connection protocol used within an AS (IGP) can be RIPv2, EIGRP, OSPF, or ISIS

BGP is used for the interconnection of different ASes (EGP)

Border Gateway Protocol (BGP)

- Inter-domain routing protocol
- BGP Version 4 (BGP-4)
- RFC 1771
- Used to connect different organisations using an Autonomous System Number (ASN)
- There are two types of ASN
  - Private
    - 64512 - 65535
  - Public
    - Issued by the internet registries
      - APNIC, ARIN, RIPE NCC, LACNIC, AFriNIC

BGP in between Autonomous systems

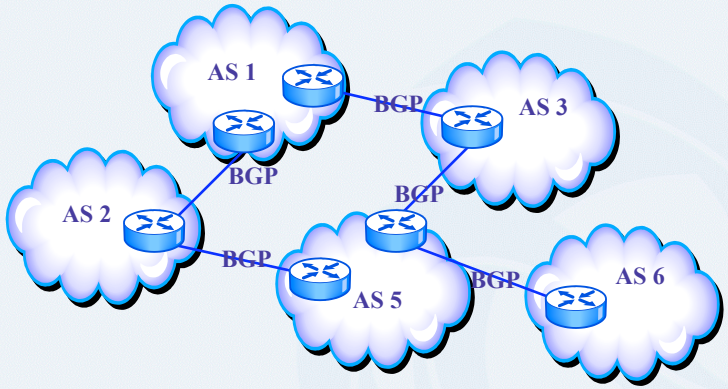

BGP is used to interconnect ASes
This guarantees loop-free routing information

When to use BGP

- BGP is the most appropriate application is for the following conditions:
  - An AS has multiple connections to different ASes
  - Packets are transmitted (transit) between third party ASes (as in an ISP scenario)
  - Decision is needed to control the traffic flow entering and leaving an AS
  - Route summarisation and aggregation of announcement exchanges from ASes

When BGP is not needed?

- BGP is not appropriate for the following conditions:
  - The AS only has single connection to the Internet or a different AS (upstream provider)
  - Routing policy is not the main concern to control the traffic flow to/from an AS

With other conditions:

- Router has limited memory and processing power
  - Can run BGP with partial or default route configuration
  - If full routes are required and transit is provided, may require higher capacity router to accommodate full routes or customer transit traffic
- Low bandwidth between AS
  - Can run even with 64Kpbs link but with minimum configuration only (partial or default route only)

Make use of static routes ☺

BGP terminology

Characteristics of BGP

- BGP is a path vector protocol
  - TCP port 179
  - PVP is UDP protocol 17
  - Incremental and triggered updates only
  - TCP connectivity is verified using periodic keepalives
  - Designed for large scale networks

BGP packets

- Since BGP run on top of TCP “port 179” it relies on TCP protocol for the reliability of the session

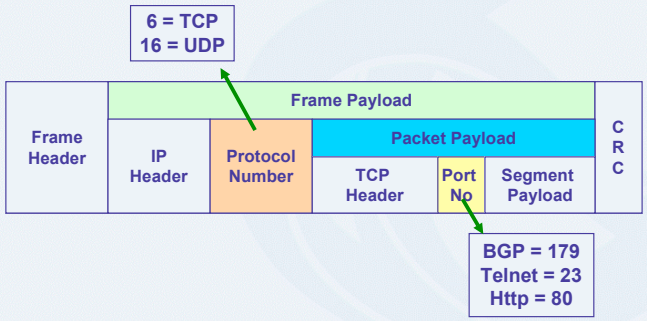

The tables

- If BGP is configured and running is creates its own table (BGP routing table) in addition to the existing IP routing table of the router (static route, IGP routes)
- However, both sets of information can be exchanged between the two tables (IP and BGP table)


```
P2R2#sh ip route

Gateway of last resort is not set

     192.168.2.0/28 is subnetted, 4 subnets
C       192.168.2.16 is directly connected, Loopback0
C       192.168.2.32 is directly connected, Serial0/0
C       192.168.2.48 is directly connected, FastEthernet0/0
```

```
P2R2#sh ip bgp

BGP table version is 8, local router ID is 192.168.2.49
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 192.168.2.0/28   192.168.2.33             1         0 i
*  192.168.2.32/28  192.168.2.33             1         0 i
*> 10.100.100.0/24  192.168.2.33             1         0 i
*> 192.168.2.16/28  0.0.0.0                  0     32768 i
*> 192.168.2.32/28  0.0.0.0                  0     32768 i
*> 192.168.2.48/28  0.0.0.0                  0     32768 i
```

Peers = neighbors

- Two or more ASes exchanging BGP information are called peers or neighbors

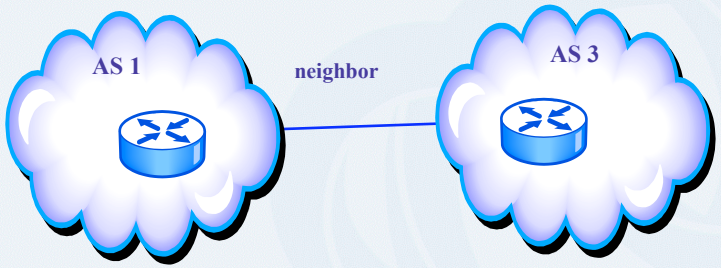

Internal BGP (iBGP)

- neighbors that belong to the same AS can use internal BGP (iBGP)
- Note that these neighbors don’t need to be directly connected

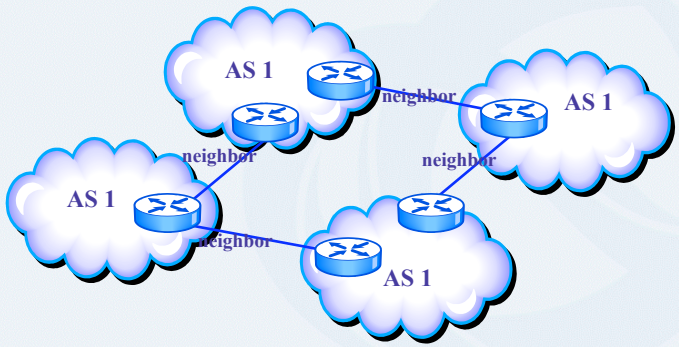

External BGP (eBGP)

- neighbors that belong to different AS use external BGP (eBGP)
- Note that these neighbors need to be directly connected

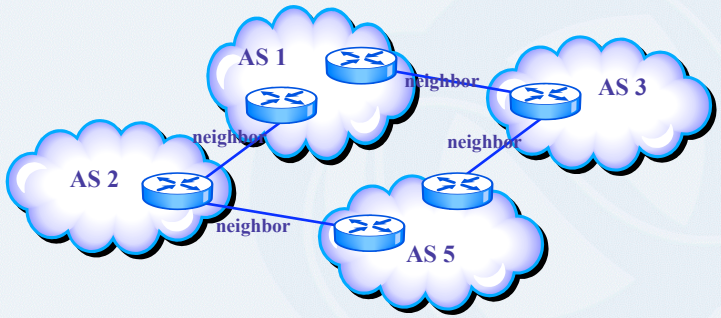

Policy routing with BGP

- BGP supports the definition of policies or rules to manipulate the flow of data through the AS
  - rules is based on hop-by-hop routing
- However, some policies which are not supported by hop-to-hop may require using different techniques
  - For example source routing

BGP attributes

- The metrics used by BGP are called path attributes
- Two types of attributes
  - Well-know
  - Optional

Well-know	Optional
Mandatory Discretionary	Transitive Nontransitive

BGP attributes (cont.)

- Well-know mandatory
  - AS-Path
  - Next-hop
  - Origin

Well-know discretionary
Local preference

Optional transitive
Community

Optional non-transitive
Multi-exit-discriminator (MED)

AS-path attribute

- Presents the list of ASes that a route has traversed in order to reach its destination

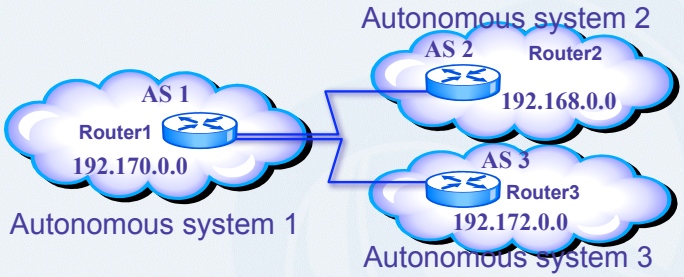

Router2 to needs to go to network 192.172.0.0
The AS-path passes through AS1 to AS3
AS-path to the network destination is AS1 and AS3 from AS2

Next-hop attribute

- Indicates the next-hop IP address used to reach the destination

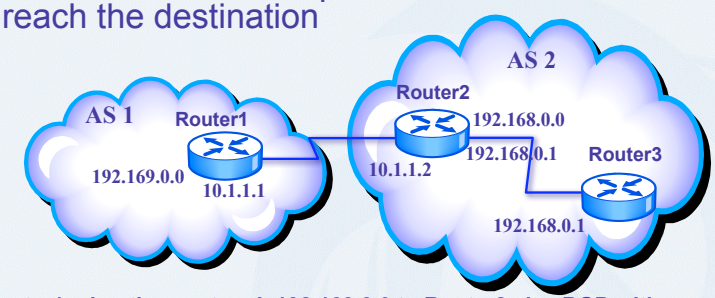

Router1 advertises network 192.169.0.0 to Router2 via eBGP with next-hop 10.1.1.1 (Router 2 serial address)

Router2 then advertises 192.169.0.0 via iBGP to Router3, and keeps the next-hop address 10.1.1.1 as the next-hop for the network 192.169.0.0

Next-hop attribute on multiaccess network

- Appropriate next-hop IP address is required to avoid insertion of additional hop into the network

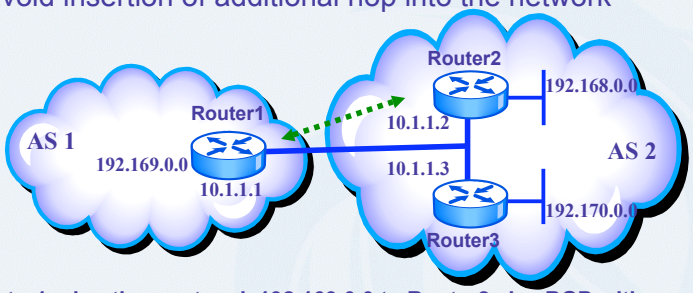

Router1 advertises network 192.169.0.0 to Router2 via eBGP with next-hop 10.1.1.2 not 10.1.1.1

The purpose is to avoid an unnecessary hop

Origin attribute

- A well-known mandatory attribute that defines the path origin
  - The (i) for IGP if achieved by the *network* command in BGP
  - The (e) EGP which was coming from the redistribution made from EGP
  - The (?) is the incomplete mark for redistributed network from IGP or static

Local preference attribute

- Provides indication to router which AS path is preferred to exit the AS
  - Highest value is preferred once configured with routers running BGP
  - Allowed only for routers within the same AS

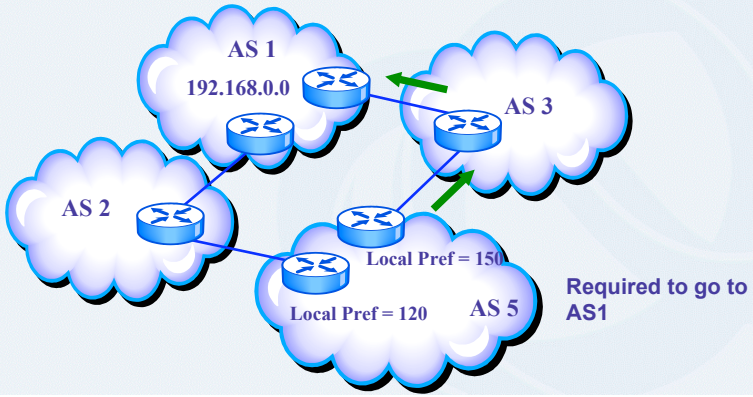

Community attribute

- Is an optional transitive attributes used for tagging of routes to ensure consistency on filtering and route-selection policy
- Tagging of routes can be made for the incoming and outgoing routing updates in the following purposes
  - Filtering of incoming routes
  - Outgoing routes updates from internal network or customer networks being announced
- Communities are dropped by default if the router does not understand

MED attribute

- Lowest value is preferred if configured with BGP
- Used with routers connecting to external BGP peers only

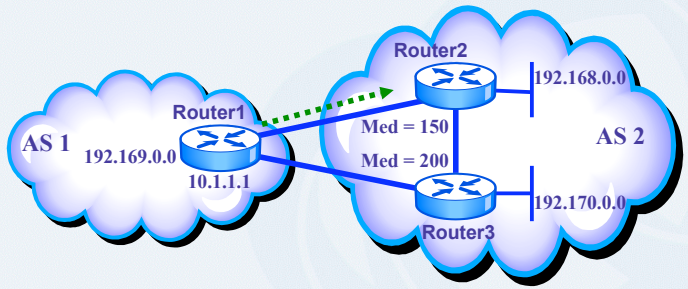

BGP synchronisation rules

- Routers cannot use or advertise any routes learned via iBGP to an external neighbor, until a route match is learned via IGP.
  - Ensuring route consistency throughout the AS but safer to turn off because it can cause problem sometimes

BGP synchronisation example

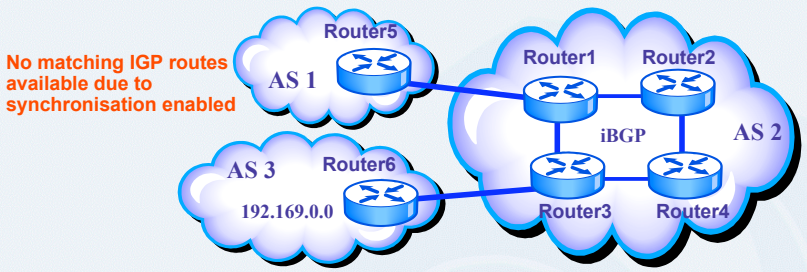

Example network with BGP synchronisation ON (default)

Router1, Router2, and Router4 would not use or advertise the route 192.169.0.0 until they receive the matching route via IGP which will keep Router 5 from not hearing anything about the network due to non-availability of routes in the IGP.

BGP synchronisation example

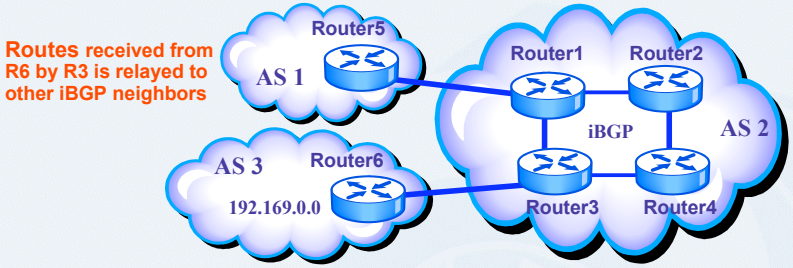

Example network with BGP synchronisation OFF

Router1, Router2, and Router4 would use and advertise the route they receive via iBGP from Router 3 and will allow announcement to Router5 so that Router 5 can hear about 192.169.0.0

BGP operation

BGP messages

- BGP messages types are very important to understand to make sure that BGP is perfectly running
- Understanding the messages types will make it more easier to troubleshoot BGP problem
  - The “Open” message type contains the hold timer for BGP including the BGP router ID
  - The “Keepalive” is used for hold timer expiration
  - The “Update” handles the information for BGP updates but single path only
  - The “Notification” is for error detection to triggers the BGP protocol to close immediately if needed

Route selection decision

- The stages for the process selection of route decisions below is based on the assumption that routes are synchronised and no AS loops and valid next-hop:
  - Prefer highest weight (local to router)
  - Prefer highest local preference (within the AS)
  - Prefer routes originated by the local router
  - Prefer shortest AS-path
  - Prefer lowest origin code (IGP < EGP < incomplete)
  - Prefer lowest MED (from other AS)
  - Prefer eBGP path over iBGP path
  - Prefer the path through the closest IGP neighbor
  - Prefer oldest route for eBGP paths
  - Prefer the path with the lowest neighbor BGP routes ID
  - Prefer the path with the lowest neighbor IP address

Aggregated address

- Routes can be aggregated when sending announcement to ASes

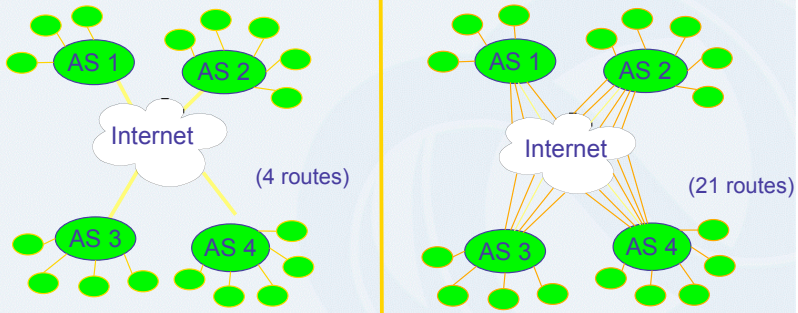

(Aggregated announcement) (No aggregation)

Written exercise 6 ☺

Please answer the attached worksheet material in your student manual

- Objective
  - To practice what you have learned

Written exercise 6

- BGP terminology and operations ☺

Configuring BGP

Note: all example commands are Zebra base

BGP configuration commands

- Starting the BGP routing process

  router bgp *autonomous-system-number*

- Defining the network to advertise

  network *network-number* ***mask*** *network-mask*

BGP configuration commands

- Setting the neighbour individually

  neighbor *ip-addresss* ***remote-as*** *autonomous-system-number*

- Setting the neighbours and defining peer groups

  neighbor *ip-addresss | peer-group-name* ***remote-as*** *autonomous-system-number*

BGP configuration commands

- Forcing the next-hop address

  neighbor *ip-addresss | peer-group* ***next-hop-self***

- Disabling synchronisation

  no synchronisation

- Summarising or aggregating routes

  aggregate-address *ip-addresss mask* ***[summary-only] [as-set]***

BGP configuration sample

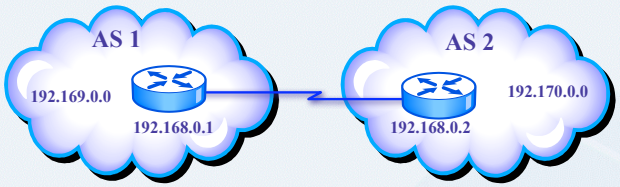

AS 1 (Router 1)

```
Router1(config) router bgp 1
Router1(config-router) neighbor 192.168.0.2 remote-as 2
Router1(config-router) network 192.169.0.0
```

AS 2 (Router 2)

```
Router2(config) router bgp 2
Router2(config-router) neighbor 192.168.0.1 remote-as 1
Router2(config-router) network 192.170.0.0
```

BGP configuration sample

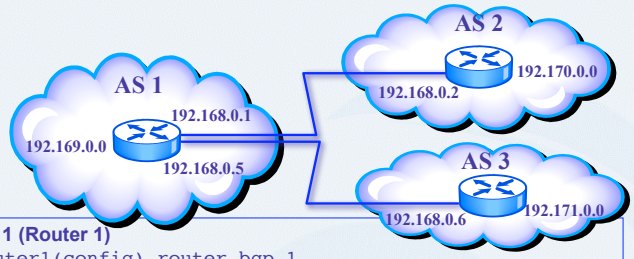

AS 1 (Router 1)

```
Router1(config) router bgp 1
Router1(config-router) neighbor 192.168.0.2 remote-as 2
Router1(config-router) neighbor 192.168.0.6 remote-as 3
Router1(config-router) network 192.169.0.0
```

AS 2 (Router 2)

```
Router2(config) router bgp 2
Router2(config-router) neighbor 192.168.0.1 remote-as 1
Router2(config-router) network 192.170.0.0
```

AS 3 (Router 3)

```
Router3(config) router bgp 3
Router3(config-router) neighbor 192.168.0.5 remote-as 1
Router3(config-router) network 192.171.0.0
```

Managing and verifying BGP

- To be able to manage and verify the BGP session running the following commands can be used:
  - Reset or route refresh for the BGP session to a neighbour

    clear ip bgp {* | *ip-address}* *[soft* ***[in*** | ***out]]***
  - Commands to view the BGP sessions informative

```
show ip bgp
show ip bgp paths
show ip bgp summary
show ip bgp neighbors
```

Lab exercise 4 ☺

Please perform the required configuration in the worksheet material of your student manual

- Objective
  - To practice what you have learned

Lab exercise 4

- Configuring BGP peering ☺

Questions?