

BGP

- BGP Overview
- BGP Operation
- BGP Neighbors

BGP Overview

- AS - Autonomous Systems
- Multihoming
- IGP vs. EGP
- When to use BGP?

BGP Overview

- **AS - Autonomous Systems**
- Multihoming
- IGP vs. EGP
- When to use BGP?

An autonomous system is a collection of networks under a single technical administration

AS numbers are 16 bits, 1 – 65535

Numbers from 64512 to 65535 are reserved for private use, much like RFC 1918 IP addresses

IANA controls IP address allocation and AS numbers

RIRs (Regional Internet Registry)

- American Registry for Internet Numbers (ARIN)
- Réseaux IP Européens Network Coordination Center (RIPE NCC)
- The Asia Pacific Network Information Center (APNIC)
- Latin American and Caribbean Internet Addresses Registry (LACNIC)
- AfriNIC


BGP Overview

- AS - Autonomous Systems
- **Multihoming**
- IGP vs. EGP
- When to use BGP?

If an organization wants to perform multihoming with BGP, there are three common ways to do this.

BGP Overview


- AS - Autonomous Systems
- **Multihoming**
- IGP vs. EGP
- When to use BGP?



Default route from all providers

BGP Overview

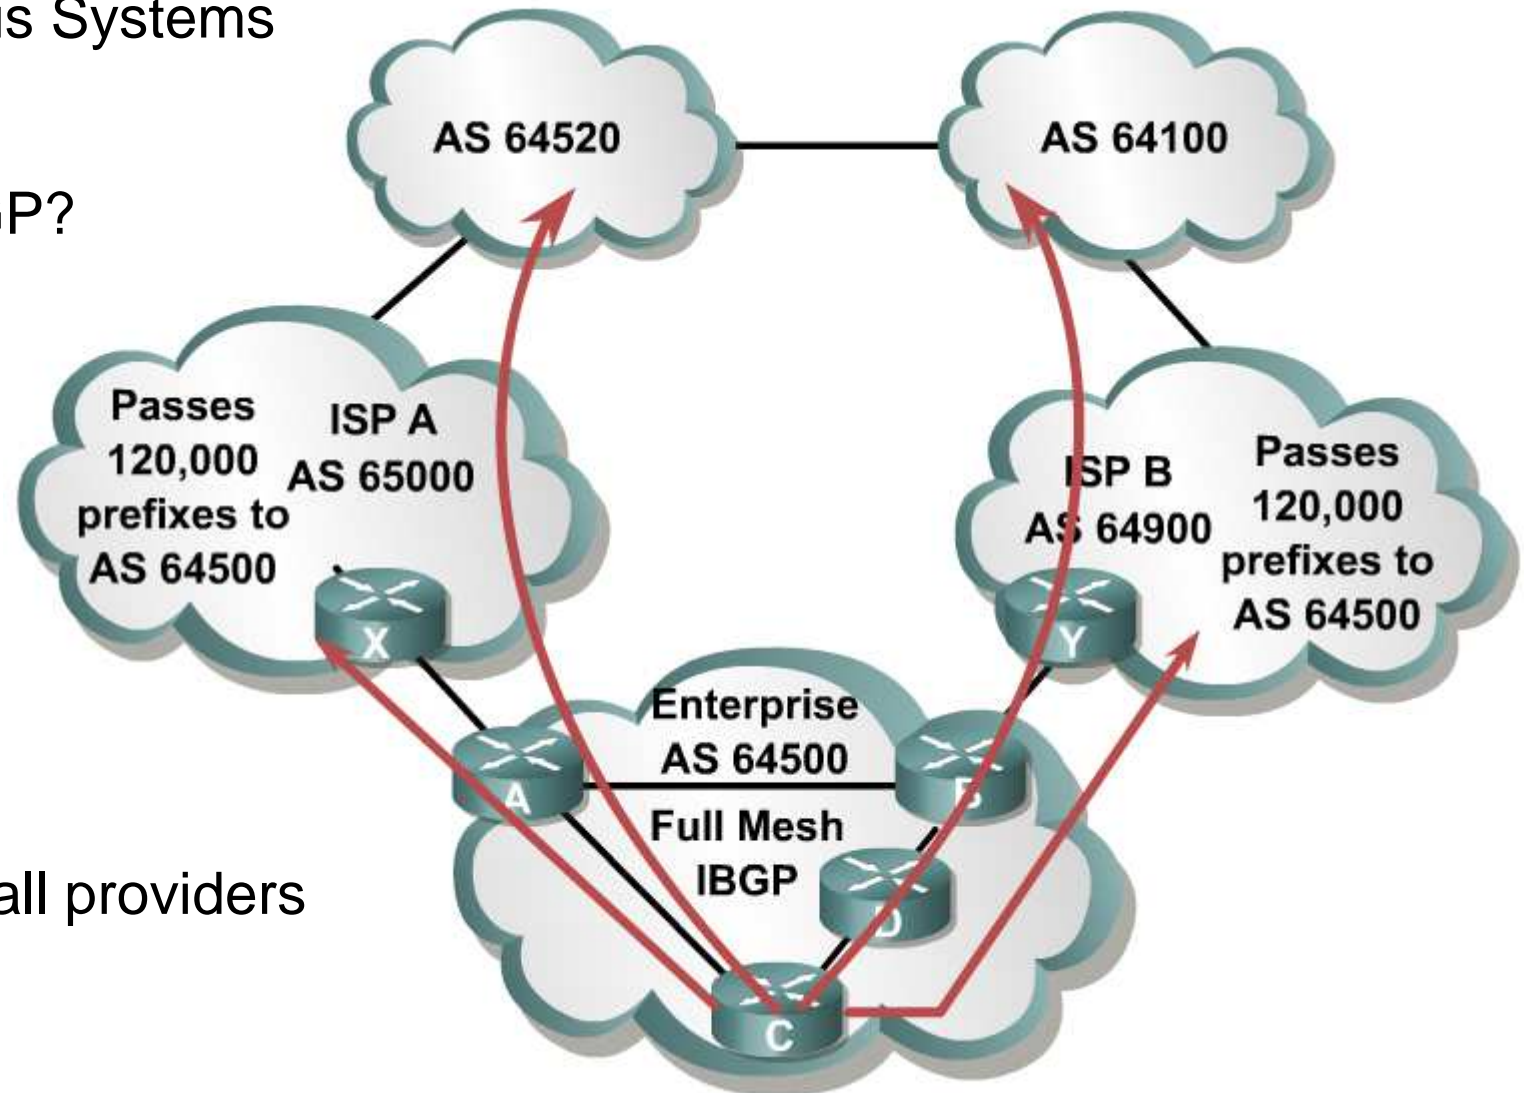
- AS - Autonomous Systems
- **Multihoming**
- IGP vs. EGP
- When to use BGP?



Default route and provider-owned partial routes from all providers

BGP Overview

- AS - Autonomous Systems
- **Multihoming**
- IGP vs. EGP
- When to use BGP?



Full routes from all providers

BGP Overview


- AS - Autonomous Systems
- Multihoming
- **IGP vs. EGP**
- When to use BGP?

IGP – Interior Gateway Protocol

Exchanges routing information within an autonomous system. RIP, IGRP, OSPF, IS-IS, and EIGRP are IGPs

EGP – Exterior Gateway Protocol

Exchanges routing information between different autonomous systems. BGP is an EGP.



BGP Overview

- AS - Autonomous Systems
- Multihoming
- IGP vs. EGP
- **When to use BGP?**

BGP is most appropriate when at least one of the following conditions exists:


- An AS allows packets to transit through it to reach other autonomous systems (for example, it is a service provider).
- An AS has multiple connections to other autonomous systems.
- Routing policy and route selection for traffic entering and leaving the AS must be manipulated.

BGP is not always appropriate. You do not have to use BGP if you have one of the following conditions:

- Limited understanding of route filtering and BGP path-selection process
- A single connection to the Internet or another AS
- Lack of memory or processor power to handle constant updates on BGP routers

BGP Operation

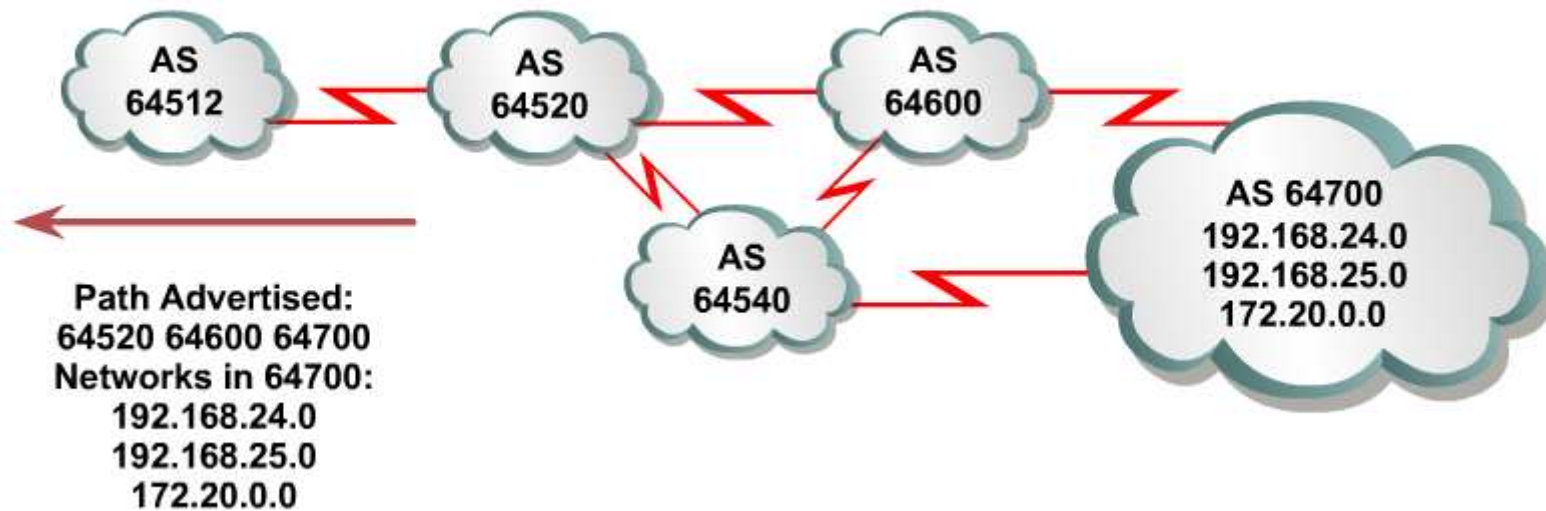
- **BGP is a policy-based routing protocol**
- Path-Vector routing protocol
- BGP Databases
- BGP message types



Improper filtering of BGP updates can potentially allow an outside autonomous system to affect the traffic flow to your autonomous system

BGP Operation

- BGP is a policy-based routing protocol
- **Path-Vector routing protocol**
- BGP Databases
- BGP message types



Avoid routing loops by as_path
Sends only the best paths to neighbors

BGP Operation

- BGP is a policy-based routing protocol
- Path-Vector routing protocol
- **BGP Databases**
- BGP message types

Neighbor database	BGP Table/BGP FIB	IP routing table/IP FIB
<ul style="list-style-type: none">• List of BGP neighbors	<ul style="list-style-type: none">• List of networks learned from each neighbor• Can contain multiple paths to destination networks• Contains BGP attributes for each path	<ul style="list-style-type: none">• List of best paths to destination networks

- Uses TCP port 179
- Uses TCP sliding windows to transfer huge amounts of information
- Sends TCP keepalives every 60 seconds by default

BGP Operation

- BGP is a policy-based routing protocol
- Path-Vector routing protocol
- BGP Databases
- **BGP message types**

BGP defines the following message types:


- Open
 - Includes holdtime and BGP router ID
- Keepalive
- Update
 - Information for one path only (could be to multiple networks)
 - Includes path attributes and networks
- Notification
 - When error is detected
 - BGP connection is closed as being sent

BGP Neighbors

- iBGP & eBGP
- BGP synchronization
- BGP Finite-state Machine

BGP Neighbors

- **iBGP & eBGP**
- BGP synchronization
- BGP Finite-state Machine



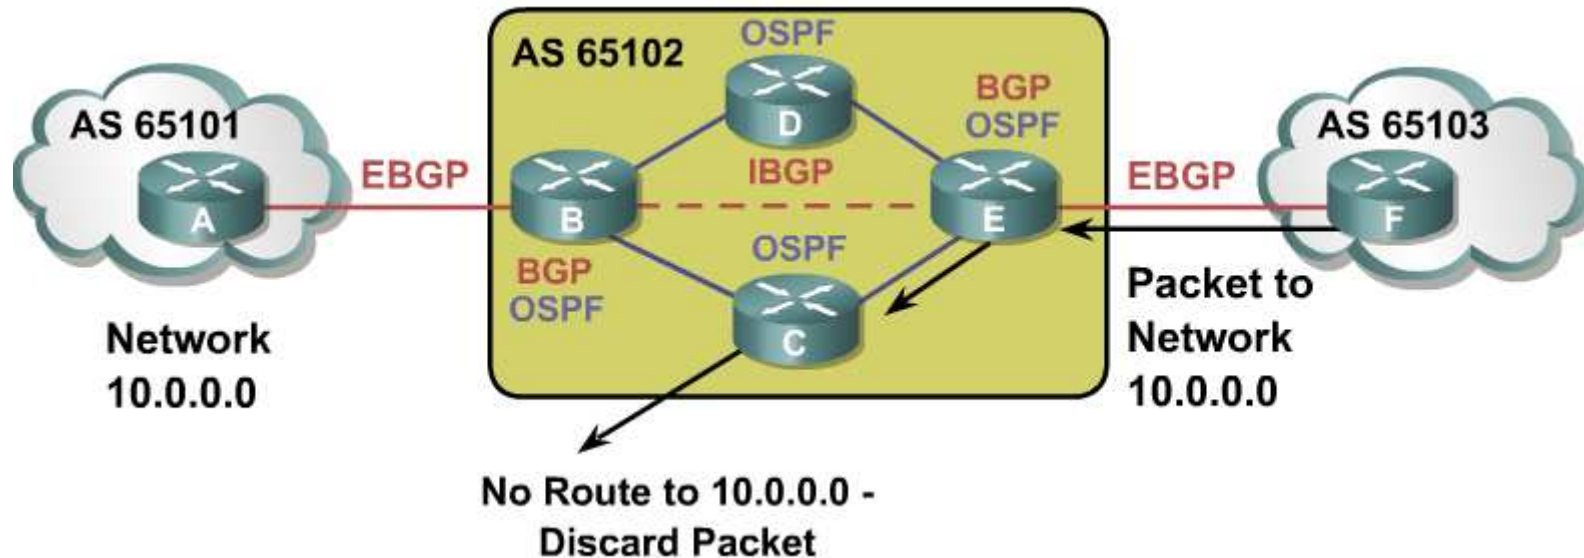
When BGP is running between peers in different autonomous systems, it is called eBGP. Neighbors need to be directly connected.

When BGP is running between routers in ... same AS, it is called iBGP. Neighbors do not have to be directly connected.

2 connected BGP speakers is called peers/neighbors

BGP Neighbors

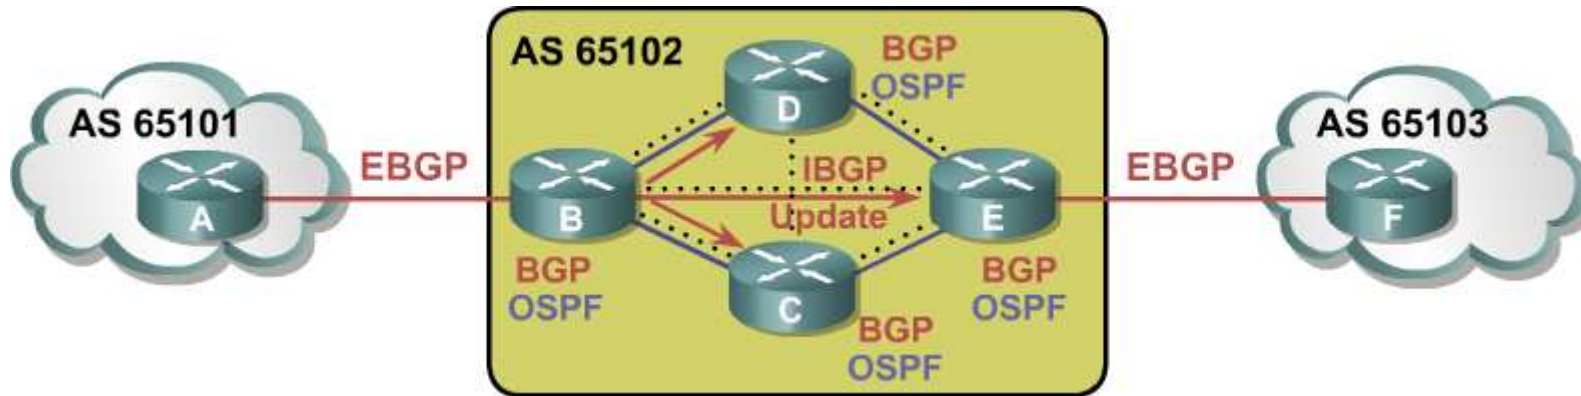
- iBGP & eBGP
- **BGP synchronization**
- BGP Finite-state Machine



- The BGP synchronization rule states that a BGP router should not use, or advertise to an external neighbor, a route that is learned from iBGP unless that route is local or the router learns it from the IGP
- Disabled by default on IOS 12.2(8)T and above

BGP Neighbors

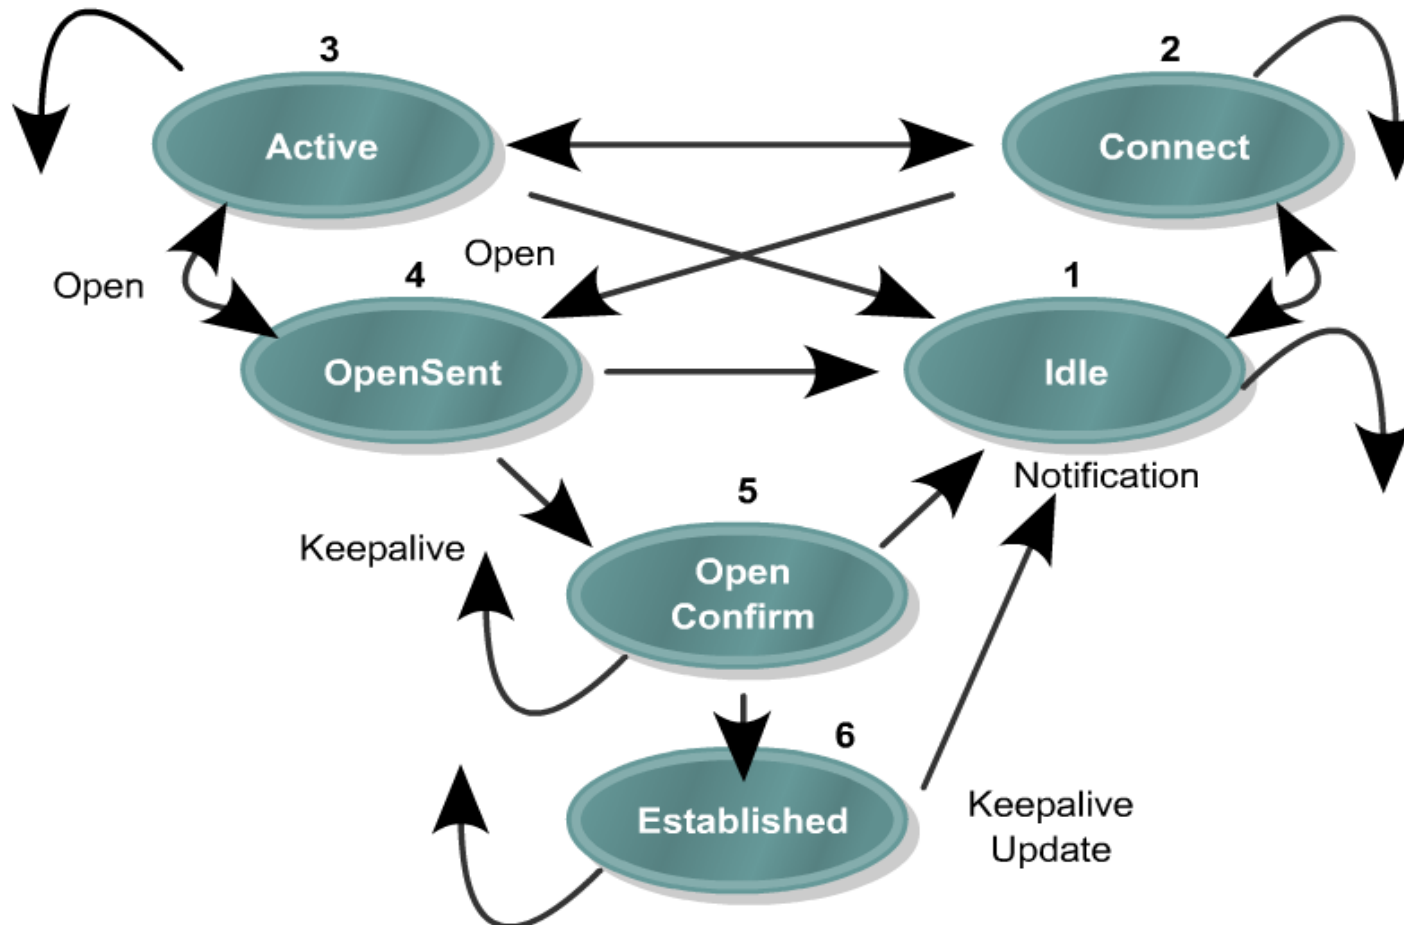
- iBGP & eBGP
- **BGP synchronization**
- BGP Finite-state Machine



- The BGP synchronization rule states that a BGP router should not use, or advertise to an external neighbor, a route that is learned from iBGP unless that route is local or the router learns it from the IGP
- Disabled by default on IOS 12.2(8)T and above

BGP Neighbors

- iBGP & eBGP
- BGP synchronization
- **BGP Finite-state Machine**



BGP Configuration

```
Router(config)#router bgp as-number
```

```
Router(router-config)#neighbor ip-number remote-as as-number
```

```
Router(router-config)#neighbor ip-number update-source interface-type interface-number
```

```
Router(router-config)#neighbor ip-number ebgp-multihop [t11]
```


```
Router(router-config)#network ip-number
```

```
Router(router-config)#no synchronization
```

BGP Configuration

`neighbor ip-number update-source interface-type interface-number`

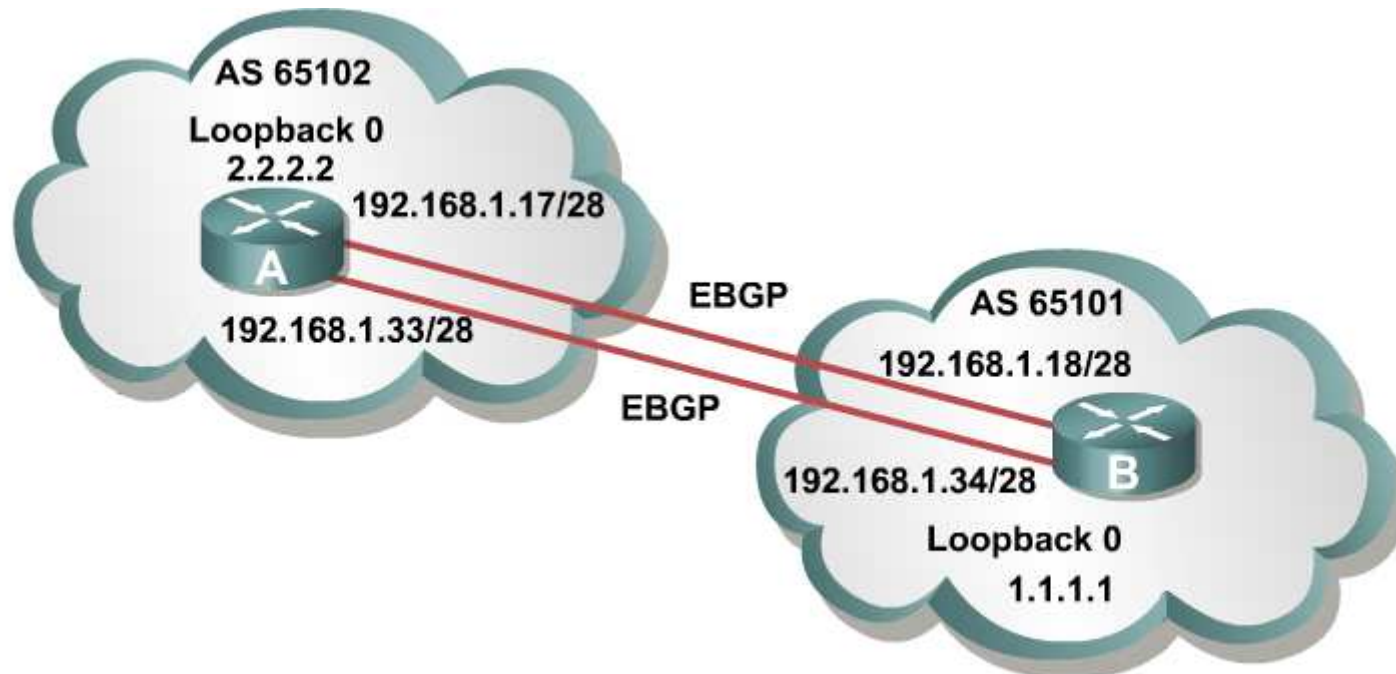
```
Router(config)#router bgp 65101
Router(router-config)#neighbor 3.3.3.3 remote-as 65101
Router(router-config)#neighbor 3.3.3.3 update-source Loopback0
```



BGP Configuration

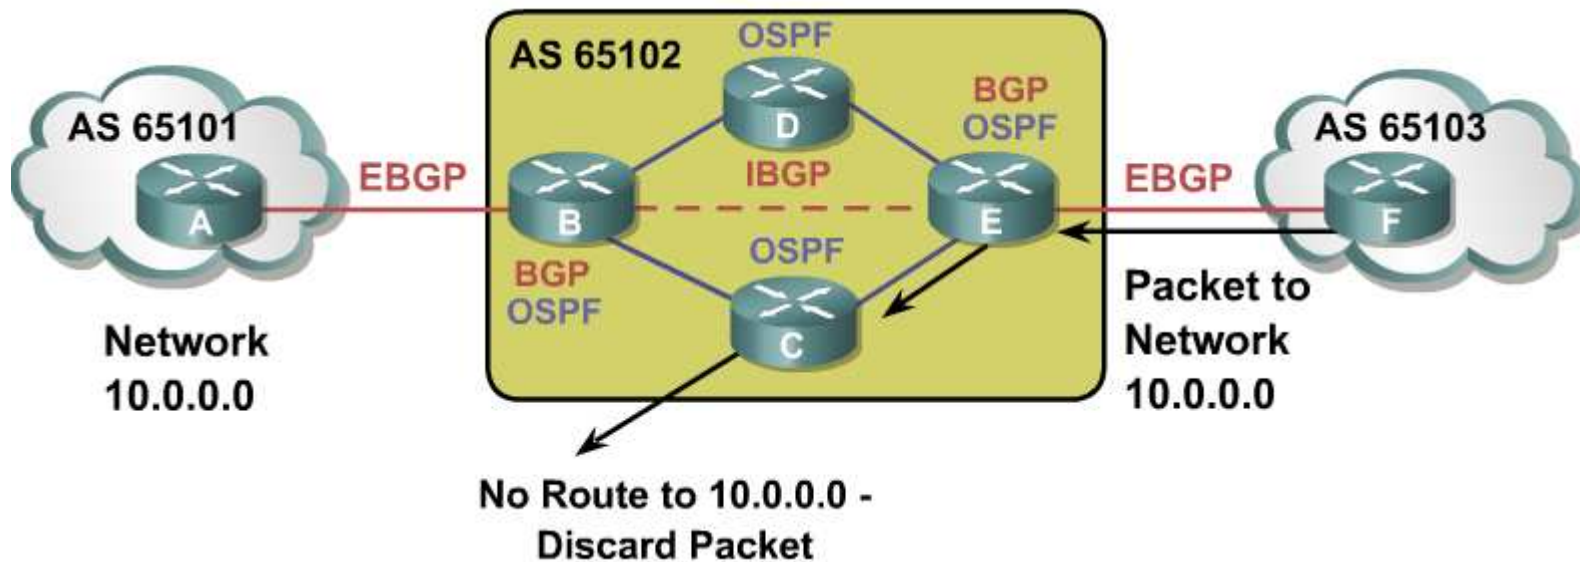
```
neighbor ip-number ebgp-multihop [ttl]
```

```
Router(config)#router bgp 65101  
Router(router-config)#neighbor 2.2.2.2 remote-as 65102  
Router(router-config)#neighbor 2.2.2.2 update-source Loopback0  
Router(router-config)#neighbor 2.2.2.2 ebgp-multihop 2
```



BGP Configuration

no synchronization



BGP Configuration

`Network network-number [mask network-mask]`

The network command is used to advertise networks

- Network command determines which network the router advertises
- BGP will only advertise a network if it is present in the routing table
- The network command does not start BGP on interfaces
- If network-mask is specified an exact match must exist in the routing table

```
Router(router-config)#network 192.168.1.1 mask 255.255.255.0
```

- Route 192.168.1.0/24 will not be advertised, if exists.
- Route 192.168.1.1/32 will not be advertised, if exists.

BGP Configuration


```
Network network-number [mask network-mask] [route-map map-tag]
```

Route-maps are used to filter advertised networks

- If no route-map is specified, all network are advertised
- If a route-map is specified but does not exist, no networks are advertised

BGP Configuration

Neighbor word peer-group
Neighbor word remote-as 65000



Router C Without a Peer Group

```
router bgp 65100
neighbor 192.168.24.1 remote-as 65100
neighbor 192.168.24.1 update-source Loopback 0
neighbor 192.168.24.1 next-hop-self
neighbor 192.168.24.1 distribute-list 20 out
neighbor 192.168.25.1 remote-as 65100
neighbor 192.168.25.1 update-source Loopback 0
neighbor 192.168.25.1 next-hop-self
neighbor 192.168.25.1 distribute-list 20 out
neighbor 192.168.26.1 remote-as 65100
neighbor 192.168.26.1 update-source Loopback 0
neighbor 192.168.26.1 next-hop-self
neighbor 192.168.26.1 distribute-list 20 out
```

Router C Using a Peer Group

```
router bgp 65100
neighbor internal peer-group
neighbor internal remote-as 65100
neighbor internal update-source Loopback 0
neighbor internal next-hop-self
neighbor internal distribute-list 20 out
neighbor 192.168.24.1 peer-group internal
neighbor 192.168.25.1 peer-group internal
neighbor 192.168.26.1 peer-group internal
```

BGP Configuration

`Clear ip bgp *`

- Resets all BGP connections with this router
- Entire BGP forwarding table is discarded
- BGP sessions go to idle state; everything must be relearned

`Clear ip bgp neighbor-address`

- Same as above but only with one neighbor

BGP Configuration

```
Clear ip bgp {* | neighbor-address} [soft out]
```

- Routes learned from this neighbor are not lost
- This router resends all BGP information to the neighbor
- The connections remains established

```
Neighbor neighbor-address soft-reconfiguration inbound
```

- The router stores all updates from the neighbor
- The command is memory-intensive

```
Clear ip bgp {* | neighbor-address} [soft in]
```

- Routes advertised to this neighbor are not withdrawn
- Revalidates learned routes
- The connections remains established

BGP Path Selection

- BGP attributes
- Common attributes
- Route selection decision process

BGP Path Selection

- **BGP attributes**
- Common attributes
- Route selection decision process

- Well-known attributes:
 - Must be recognized by all compliant BGP implementations
 - Are propagated to other neighbors
- Well-known mandatory attributes:
 - Must be present in all update messages
- Well-known discretionary attributes:
 - May be present in update messages

- Optional attributes:
 - Recognized by some implementations (could be private); expected not to be recognized by all BGP routers.
 - Recognized optional attributes are propagated to other neighbors based on their meaning.
- Optional transitive attributes:
 - If not recognized, are marked as partial and propagated to other neighbors
- Optional nontransitive attributes:
 - Discarded if not recognized

BGP Path Selection

- BGP attributes
- **Common attributes**
- Route selection decision process

Well-known mandatory:

- AS Path
- Next Hop
- Origin

Well-known discretionary:

- Local preference
- Atomic aggregate

Optional transitive

- Aggregator
- Community

Optional nontransitive

- Multi-exit discriminator
- Originator-ID
- Cluster-list

BGP Path Selection

- BGP attributes
- Common attributes
- **Route selection decision process**

Consider only (synchronized) routes with no AS loops and a valid next hop, and then:

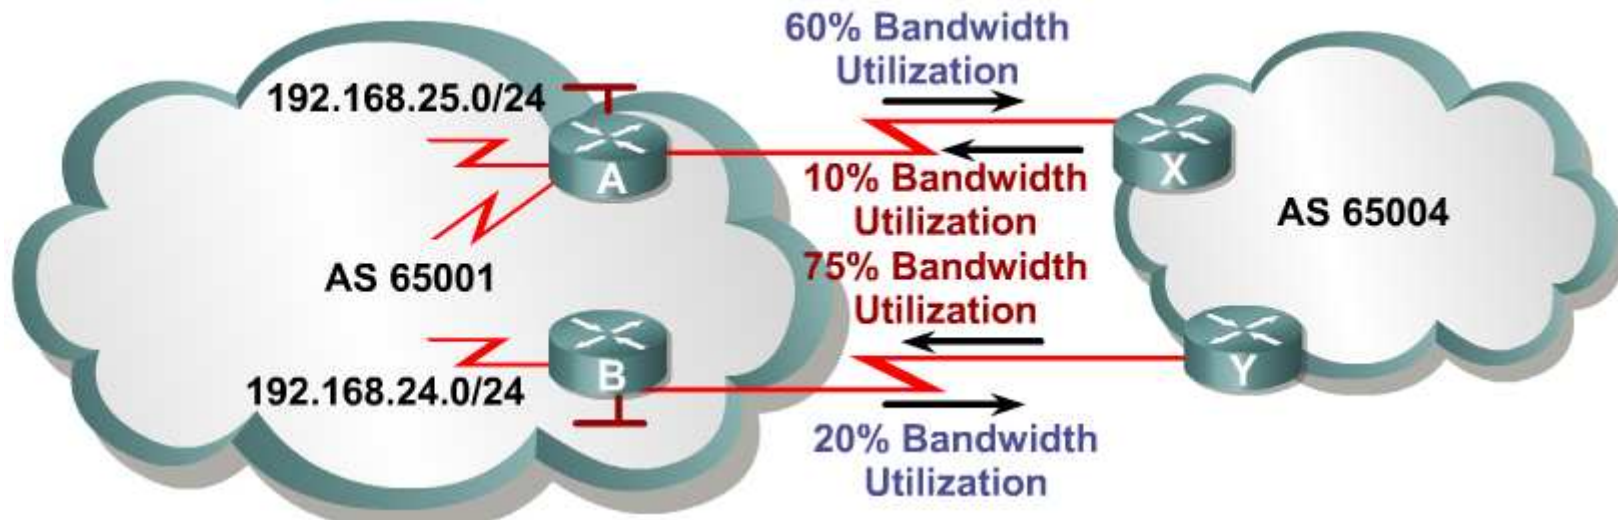
1. Prefer highest weight (local to router).
2. Prefer highest local preference (global within AS).
3. Prefer route originated by the local router (next hop = 0.0.0.0).
4. Prefer shortest AS path.
5. Prefer lowest origin code (IGP < EGP < incomplete).
6. Prefer lowest MED (from other AS).
7. Prefer EBGP path over IBGP path.
8. Prefer the path through the closest IGP neighbor.
9. Prefer oldest route for EBGP paths.
10. Prefer the path with the lowest neighbor BGP router ID.
11. Prefer the path with the lowest neighbor IP address.

BGP Manipulation

- Policy routing
- Route-maps
- Prefix-list

BGP Manipulation

- Policy routing
- Route-maps
- Prefix-list



BGP Manipulation

- Policy routing
- **Route-maps**
- Prefix-list


```
Router bgp 65001
  Neighbor 2.2.2.2 remote-as 65002
  Neighbor 2.2.2.2 route-map local_pref in
  !
Access-list 65 permit 172.16.0.0 0.0.0.255.255
  !
Route-map local_pref permit 10
  Match ip address 65
  Set local-preference 400
  !
Route-map local-pref permit 20
```

BGP Manipulation

- Policy routing
- Route-maps
- **Prefix-list**

```
Router bgp 65001
  Neighbor 2.2.2.2 remote-as 65002
  Neighbor 2.2.2.2 prefix-list prefix_in in
!
Ip prefix-list prefix-in seq 5 permit 172.16.0.0/16 le 32
Ip prefix-list prefix-in seq 10 deny 0.0.0.0/0
```

HOUSE OF
TECHNOLOGY  - en del af **mercantec**⁺



CISCO