



Global Knowledge®

Expert Reference Series of White Papers

# Interdomain Routing with BGP

# Interdomain Routing with BGP

Carol Kavalla, Global Knowledge Instructor

---

## Introduction

Border Gateway Protocol (BGP) is a very complex routing protocol where the administrator can influence the route choice in more ways than with an interior gateway protocol. This white paper will provide you with a high level overview of interdomain routing with BGP, including: interior vs. exterior gateway protocol; basic characteristics of BGP; attributes of BGP; how relationships between BGP neighbors are defined; and how BGP chooses the best path. But before we dive into BGP, let's review the function of a router and how a router installs routes in its routing table.

## Routing Basics

The function of a router is to find a path to a destination network and forward packets. The path a packet takes is ultimately derived from information received from a routing protocol neighbor.

If more than one routing protocol is enabled on a router, it will install the route with the lower administrative distance for identical routes learned. If there is only one routing protocol enabled on a router, the router will install the route with the lowest metric (for example: hop count; cost; or bandwidth and delay).

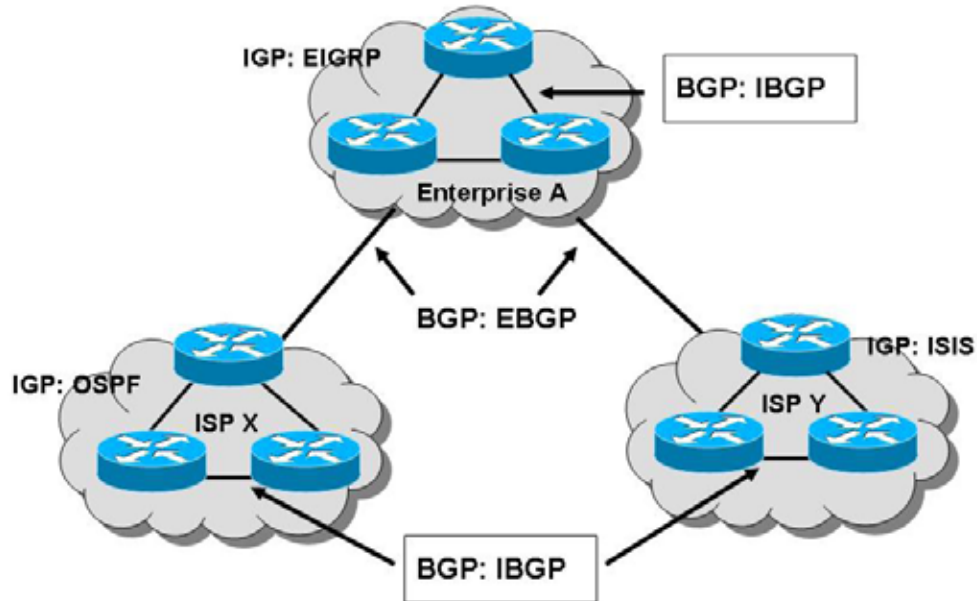
Several alternate paths to a destination network may exist, but only the best next hop is stored in the routing (or forwarding) table. Depending on the routing protocol, a router can store several equal cost paths in its routing table. When the topology of a network changes, a router will attempt to install an alternate path and will then continue to forward packets along that alternate path.

The IP route lookup is done on the destination IP address in the IP packet and is based on "longest match" routing. A more specific prefix is preferred over a less specific prefix. For example: an IP packet with a destination address of 192.168.100.33/27 is sent to the router announcing the 192.168.100.32/27 rather than the one announcing the 192.168.100.0/24 network.

After determining the longest match for a particular route, the router then determines to which interface the IP packet is to be sent and switches out the packet to that interface.

Routing decisions in BGP are based on routing policies and metrics, more accurately referred to as attributes (more on this later).

# Example: Interior Gateway Protocol & Exterior Gateway Protocol



## Interior Gateway Protocol vs. Exterior Gateway Protocol

A routing protocol can route packets within an autonomous system (AS) or between autonomous systems. An AS is a network or group of networks under a common administration and with common routing policies.

BGP is used to exchange routing information in the Internet between Internet Service Providers (ISP)—autonomous systems. Enterprise networks make use of Interior Gateway Protocols (IGP) such as Enhanced Interior Gateway Protocol (EIGRP) or Open Shortest Path First (OSPF) for the exchange of routing information within their Autonomous Systems. Larger users connected to ISP's, and ISPs use BGP to exchange customer and Internet routes. When BGP is used between autonomous systems, the protocol is referred to as External BGP (EBGP). When BGP is used to exchange routes within an AS, then the protocol is referred to as Interior BGP (IBGP).

### Interior Gateway Protocol

- Automatic neighbor discovery
- Generally trust your IGP routers
- Binds routers in one AS together

### Exterior Gateway Protocol

- Specifically configured peers
- Connecting with outside networks
- Sets administrative boundaries

## Basic Characteristics

Basic characteristics of BGP include:

- BGP is a distance vector protocol with enhancements:
  - Reliable updates
  - Triggered updates only
  - Rich metrics (called path attributes)
- Designed to scale to huge internetworks


BGP uses TCP, port 179, as its transport protocol. Because TCP is a reliable transport protocol, it makes it unnecessary for BGP to utilize acknowledgement packets for reliability.

When two BGP neighbors first exchange routing information, they exchange their entire BGP Routing Information Bases (RIB)—best paths only. After the initial exchange, they only send triggered-updates. Cisco's implementation of BGP batches the updates and sends them periodically. The default batch timer is every 5 seconds for Internal BGP peers and every 30 seconds for External BGP peers. If there are no triggered updates, BGP peers send keepalives every 60 seconds (default).

BGP was designed to be implemented in an extremely large network, such as the Internet. The developers of BGP decided that scalability was more important than fast convergence of topology changes, because it was expected that a very large network would have changing routes. Consequently, today BGP has been able to scale to over 290,000 routes, and it converges slowly compared to an internal gateway protocol like OSPF.

## Multihomed Customers

- Customer connecting to more than one service provider



BGP is used when customers are multihomed or are connected to two different service providers. Internet Service Providers also employ BGP, both within their autonomous system (IBGP) and between autonomous systems (EBGP).

## Attributes of BGP

BGP attributes have a similar function to IGP metrics. Routes learned via BGP have associated properties or attributes that are used to determine the best path to a destination network when multiple paths exist. These properties are called path attributes.

The path attributes belong in the following categories:

- **Well Known Mandatory** – must be present in all BGP update messages and must be understood by all BGP routers
- **Well Known Discretionary** – must be recognized by all BGP routers, but do not have to be present in all BGP updates
- **Optional** – do not have to be present in BGP updates and do not have to be recognized by all BGP routers

Well Known Mandatory BGP attributes include:

- Origin
  - The origin of a BGP route
- AS-path
  - Sequence of AS numbers through which the network is accessible
- Next-hop
  - IP address of the next-hop router

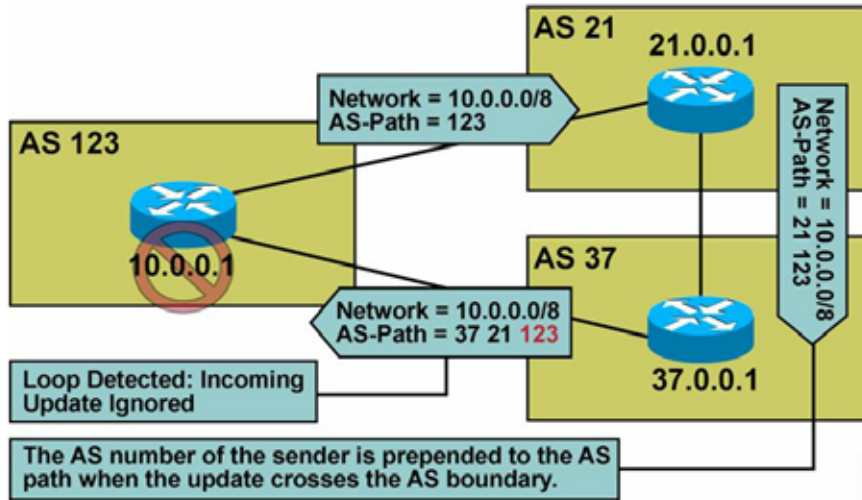
The origin attribute indicates how BGP learned about a particular route. The origin attribute can have one of three values:

- i route originated in an IGP
- e route originated in EGP
- ? route was redistributed into BGP

The AS-path attribute is a sequence of AS numbers through which the network is accessible and is null when a local route is inserted in the BGP table. The AS-path value is added when the routing update crosses an AS boundary. When an AS receives routing information with its own AS number in the AS path, it will ignore the information. This is used for loop avoidance.

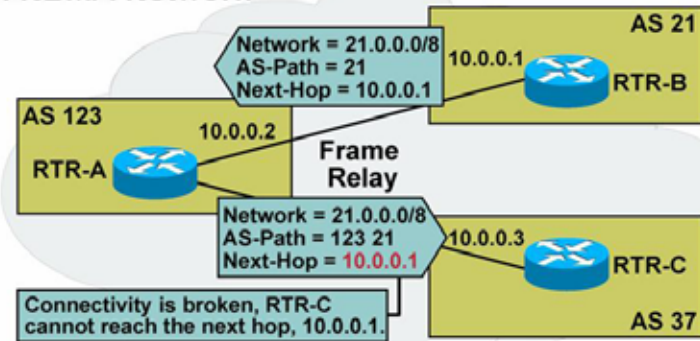
The next hop is used to reach the advertising router in order to get to the destination network. In BGP the next hop is the next autonomous system hop, not always the adjacent router.

# AS-Path Attribute Example



# Next-Hop Attribute Example

## Next-Hop Processing on NBMA Network



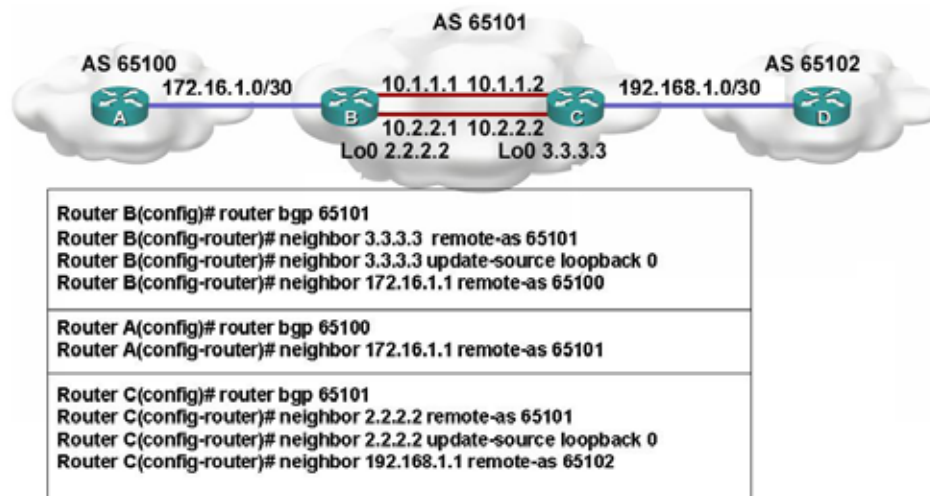
- BGP next-hop processing can break connectivity with improper network designs over partially meshed WAN networks.

## Becoming BGP Neighbors

BGP was designed to be more secure than interior gateway routing protocols in part because a BGP user may not be in control of the neighbor router.

With IGP's the routing protocol is enabled, a network statement (or statements) is configured, and a router will begin to exchange routing information with its neighbors. BGP is different. Both sides of the neighbor relationship need to define the neighbor relationship. If one side is not configured or is misconfigured, routing information will not be exchanged. Below is an example of a minimal configuration on Routers A, B, and C.

## Configuring BGP Neighbors



Routers A and B will become neighbors and exchange BGP routes. Also, routers B and C will become neighbors. Router D is not configured for BGP, so routers C and D will not become neighbors and will not exchange BGP routes. In addition, if additional security is needed, route authentication can be enabled on BGP peers.

## Choosing the Best Path

BGP will only choose one best path. It does not load share. To determine the best path it uses BGP attributes. Although all of the attributes listed below are beyond the scope of this white paper, the list indicates how a BGP router will pick the best path when it learns about a route from more than one neighbor.

### BGP Route Selection Criteria

- Exclude routes with inaccessible next hop
- Prefer highest weight (local to router)

- Prefer highest local preference (global within AS)
- Prefer routes that the router originated
- Prefer shortest AS path (only length is compared)
- Prefer lowest origin code (IGP < EGP < Incomplete)
- Prefer lowest MED
- Prefer external (EBGP) paths over internal (IBGP)
- For IBGP paths, prefer path through closest IGP neighbor
- For EBGP paths, prefer oldest (most stable) path
- Prefer paths from router with the lowest BGP router-ID

This is a top down process. As soon as the router finds a differentiator, such as shortest AS path, path selection is complete.

## Learn More

Learn more about how you can improve productivity, enhance efficiency, and sharpen your competitive edge. Check out the following Global Knowledge courses.

[BGP – Configuring BGP on Cisco Routers v3.2](#)

[CCNA Boot Camp v2.0](#)

[BSCI – Building Scalable Cisco Internetworks v3.0](#)

Our courses and enhanced, hands-on labs offer practical skills and tips that you can immediately put to use. Our expert instructors draw upon their experiences to help you understand key concepts and how to apply them to your specific work situation. Choose from our more than 1,200 courses, delivered through Classrooms, e-Learning, and On-site sessions, to meet your IT and business training needs.

For more information or to register, visit [www.globalknowledge.com](http://www.globalknowledge.com) or call 1-800-COURSES to speak with a sales representative.

## About the Author

Carol Kavalla's background includes teaching at Rockland Community College in New York, managing networks and being a consultant for the NYS small business development center. For the last eight-and-a-half years, Carol has taught for Global Knowledge and is certified to teach nine Cisco Courses: ICND1; ICND2; CCDA; BSCI; BCMSN; TCN; ICMI; BGP; and ARCH. She also has a consulting firm in Charleston, South Carolina, where she works with small companies (100-200 nodes) installing, configuring routers and switches, and troubleshooting network problems.