



Introduction to BGP and Deployment Prerequisites

This chapter introduces the key concepts of BGP, explains how it differs from other routing protocols, and outlines the technical and operational prerequisites needed for a successful BGP implementation.

- [Border Gateway Protocol, on page 1](#)
- [How BGP works, on page 2](#)
- [BGP router identifier, on page 3](#)
- [Prerequisites for implementing BGP, on page 4](#)

Border Gateway Protocol

Border Gateway Protocol (BGP) is an exterior gateway protocol that

- enables loop-free interdomain routing between autonomous systems
- exchanges routing information between different networks operated by distinct organizations, and
- supports policy-based routing decisions for scalable and reliable Internet operation.

Autonomous system

An autonomous system (AS) is a group of routers managed under a single technical administration that operates as a unified routing entity. Each AS consists of one or more IP networks under the control of a single organization and is assigned a unique Autonomous System Number (ASN).

Comparison of routing protocol types

To understand BGP, it's essential to distinguish between two types of routing protocols:

- **Interior Gateway Protocols (IGPs):** These protocols such as OSPF, IS-IS, and EIGRP, are used within a single AS. They determine the best paths for routing data within the internal network of that AS.
- **Exterior Gateway Protocols (EGPs):** These protocols, primarily BGP, are used for routing between different ASes. They enable communication and the exchange of routing information across different networks.

Feature	Exterior Gateway Protocols (EGPs)	Interior Gateway Protocols (IGPs)
Scope	Inter-AS (between organizations)	Intra-AS (within an organization)
Typical protocols	BGP	OSPF, IS-IS, EIGRP, RIP
Routing policy support	Policy-based (highly flexible)	Generally least-cost or shortest path
Scalability	Highly scalable for large networks	Suitable for limited, internal domains

- An Internet Service Provider uses BGP to exchange routing information with other service providers, ensuring that data can travel efficiently between different networks on the global Internet.
- Enterprises with multi-homed Internet connections deploy BGP to manage traffic policies and redundancy across their upstream providers.

How BGP works

Summary

The key components involved in BGP operation include:

- BGP: A protocol that uses TCP as its transport protocol to exchange routing information between routers.
- Peer routers (BGP peers/neighbors): Two routers that establish a TCP connection to exchange BGP routing information.
- BGP routing table: A table that contains the routing information exchanged between peers.
- Keepalive packets: Packets that are periodically exchanged between peers to verify that the connection is still active.

BGP operates by establishing a TCP connection between routers (peers) to exchange network reachability information, which includes paths represented by Autonomous System Numbers (ASNs). This information is used to construct a loop-free routing graph, enabling efficient routing and policy enforcement. Once connected, peers exchange complete routing tables initially, followed by incremental updates for changes such as new routes, modifications, or withdrawals. To maintain the connection, periodic keepalive packets are exchanged, and notification packets are sent for errors or special conditions, ensuring reliable inter-domain routing.

Workflow

These stages describe how BGP works.

1. Establishing a TCP connection:
 - Two BGP routers, referred to as BGP peers or neighbors, establish a TCP connection.
 - These routers exchange messages to open the connection and confirm the connection parameters.
2. Exchanging network reachability information:

- After the connection is established, BGP routers exchange network reachability information.
 - This information includes the complete paths represented by BGP autonomous system numbers (ASNs) that a route should take to reach a destination network.
 - The exchanged information helps construct a graph of autonomous systems. This graph shows loop-free paths and identifies where routing policies can be applied to enforce restrictions on routing behavior.
3. Constructing a loop-free routing graph:
 - BGP uses the exchanged network reachability information to construct a graph.
 - This graph identifies loop-free autonomous systems and highlights where routing policies can be applied to enforce restrictions on routing behavior.
 4. Exchanging routing tables:

When the BGP connection is established, the peers exchange their complete BGP routing tables.
 5. Sending incremental updates:

After the initial exchange, BGP routers send only incremental updates as changes occur in the routing table. These updates may include new routes, modified routes, or route withdrawals.
 6. Tracking routing table versioning:

BGP keeps a version number for the routing table. This version number is identical across all BGP peers and is updated whenever changes are made to the table due to routing information updates.
 7. Maintaining the connection:
 - BGP routers send periodic **keepalive** packets between peers to ensure that the connection remains active.
 - BGP routers send **notification** packets when errors or special conditions occur.

Result

This process ensures that BGP operates efficiently, maintains reliable communication, and adapts to changes in routing information dynamically.

BGP router identifier

Each BGP router needs a unique router identifier (ID) to establish BGP sessions. This ID is sent to BGP peers in the OPEN message when a BGP session starts.

Methods for determining BGP router ID

BGP tries to obtain a router ID in this order of preference:

1. **Configured router ID:** BGP uses the router ID set explicitly with the **bgp router-id** command in router configuration mode.

2. **Highest IPv4 address on loopback interface:** If no router ID is configured, BGP uses the highest IPv4 address from a loopback interface, provided the router was booted with a saved loopback address configuration.
3. **Primary IPv4 address of the first configured loopback:** If no loopback address exists in the saved configuration, BGP uses the primary IPv4 address of the first loopback interface that is configured.

What happens if BGP fails to obtain a router ID

If none of these methods for obtaining a router ID succeeds, BGP does not have a router ID and cannot establish any peering sessions with BGP neighbors.

In such an instance:

- The system logs an error message.
- The **show bgp summary** command displays a router ID of 0.0.0.0.

Behavior after a router ID is assigned

Once BGP obtains a router ID, it uses that ID continuously, even if a better option becomes available based on the rules. This behavior prevents unnecessary flapping of BGP sessions.

Exception

If the current router ID becomes invalid, for instance due to a down interface or configuration change, BGP selects a new one according to the previously outlined rules. However, this change resets all established peering sessions.

Best practice

We strongly recommend configuring the router ID using the **bgp router-id** command. This helps prevent unexpected changes to the router ID and avoids potential BGP session flapping.

Prerequisites for implementing BGP

User group and task ID requirements

- You must be part of a user group associated with a task group that includes the required task IDs.
- Refer to the command reference guides to find the task IDs necessary for each command.
- Contact your AAA administrator for assistance if you suspect that your user group assignment is restricting access to use a command.

BGP table memory requirements

- The current Internet BGP table contains approximately 1.1 million IPv4 routes and 200,000 IPv6 routes. Ensure your router has at least 5.5 GB of RAM to manage the full Internet BGP table, with an average of two paths per route.
- You must plan for higher memory requirements as the IPv6 Internet routing table continues to expand.

Hardware requirements

Use the Service Edge (SE) version of Route Processor (RP) cards, Route Switch Processor (RSP) cards, or fixed chassis routers for devices that need to maintain a full BGP table.

