RIPE Border Gateway Protocol (BGP) Security Practice Test (Sample)

Study Guide

BY EXAMZIFY

Everything you need from our exam experts!

Notice: Examzify makes every reasonable effort to obtain accurate, complete, and timely information about this product from reliable sources.

Table of Contents

Introduction

Preparing for a certification exam can feel overwhelming, but with the right tools, it becomes an opportunity to build confidence, sharpen your skills, and move one step closer to your goals. At Examzify, we believe that effective exam preparation isn't just about memorization, it's about understanding the material, identifying knowledge gaps, and building the test-taking strategies that lead to success.

This guide was designed to help you do exactly that.

Whether you're preparing for a licensing exam, professional certification, or entry-level qualification, this book offers structured practice to reinforce key concepts. You'll find a wide range of multiple-choice questions, each followed by clear explanations to help you understand not just the right answer, but why it's correct.

The content in this guide is based on real-world exam objectives and aligned with the types of questions and topics commonly found on official tests. It's ideal for learners who want to:

- **Practice answering questions under realistic conditions,**
- **Improve accuracy and speed,**
- **Review explanations to strengthen weak areas, and**
- **Approach the exam with greater confidence.**

We recommend using this book not as a stand-alone study tool, but alongside other resources like flashcards, textbooks, or hands-on training. For best results, we recommend working through each question, reflecting on the explanation provided, and revisiting the topics that challenge you most.

Remember: successful test preparation isn't about getting every question right the first time, it's about learning from your mistakes and improving over time. Stay focused, trust the process, and know that every page you turn brings you closer to success.

Let's begin.

How to Use This Guide

This guide is designed to help you study more effectively and approach your exam with confidence. Whether you're reviewing for the first time or doing a final refresh, here's how to get the most out of your Examzify study guide:

1. Start with a Diagnostic Review

Skim through the questions to get a sense of what you know and what you need to focus on. Your goal is to identify knowledge gaps early.

2. Study in Short, Focused Sessions

Break your study time into manageable blocks (e.g. 30 – 45 minutes). Review a handful of questions, reflect on the explanations.

3. Learn from the Explanations

After answering a question, always read the explanation, even if you got it right. It reinforces key points, corrects misunderstandings, and teaches subtle distinctions between similar answers.

4. Track Your Progress

Use bookmarks or notes (if reading digitally) to mark difficult questions. Revisit these regularly and track improvements over time.

5. Simulate the Real Exam

Once you're comfortable, try taking a full set of questions without pausing. Set a timer and simulate test-day conditions to build confidence and time management skills.

6. Repeat and Review

Don't just study once, repetition builds retention. Re-attempt questions after a few days and revisit explanations to reinforce learning. Pair this guide with other Examzify tools like flashcards, and digital practice tests to strengthen your preparation across formats.

There's no single right way to study, but consistent, thoughtful effort always wins. Use this guide flexibly, adapt the tips above to fit your pace and learning style. You've got this!

Questions

1. **Why is BGP a highly applicable routing protocol for the Internet?**
   - A. The Internet uses only the hop-by-hop routing paradigm and BGP can support any policy that conforms to that paradigm
   - B. It uses distance-vector routing
   - C. It floods all networks
   - D. It ignores policies

2. **What are the four BGP message types?**
   - A. Open, Keepalive, Update and Notification
   - B. Hello, Update, Withdraw, Tear-Down
   - C. Start, Stop, Update and Notify
   - D. Connect, Open, Keepalive and Shutdown

3. **True or False: A BGP router forms a neighbor relationship with a large number of BGP routers**
   - A. True
   - B. Depends on the router
   - C. BGP supports only a single neighbor
   - D. False

4. **In BGP, what characterizes the open sent state?**
   - A. The router received agreement on the parameters for establishing a session.
   - B. An open message was sent with the parameters for the BGP session.
   - C. The router found a route to the neighbor and has completed the three-way handshake.
   - D. Peering is established and routing begins.

5. **What is the next-hop attribute?**
   - A. The IP address used to reach the next AS
   - B. The origin code
   - C. The AS path length
   - D. The MED value

6. Name the three well-known mandatory attributes.

- **A. AS path, Next Hop, Origin**
- **B. Local Preference, Atomic Aggregate, Community**
- **C. Aggregator, Community, Next Hop**
- **D. Origin, Local Preference, AS Path**

7. Which command enters BGP configuration mode?

- **A. Router bgp autonomous-system in global config**
- **B. Router bgp in global config**
- **C. bgp config enter**
- **D. Enable BGP in interface mode**

8. Which command displays the BGP state?

- **A. Show IP Route Summary**
- **B. Show BGP Neighbors**
- **C. Show Run BGP**
- **D. Show IP BGP Summary**

9. How is the BGP router ID chosen?

- **A. Highest Loopback Interface**
- **B. Highest Active Interface**
- **C. Automatically Chosen Randomly**
- **D. Configured Using the BGP Router-ID Command**

10. Which statement about the BGP next-hop behavior for eBGP is true?

- **A. The next-hop is the IP address to reach the next AS**
- **B. The next-hop is always the neighbor's interface address**
- **C. The next-hop is always the next router's IP**
- **D. There is no next-hop in BGP**

Answers

1. **A**
2. **A**
3. **B**
4. **B**
5. **A**
6. **A**
7. **A**
8. **D**
9. **D**
10. **A**

Explanations

1. Why is BGP a highly applicable routing protocol for the Internet?

A. The Internet uses only the hop-by-hop routing paradigm and BGP can support any policy that conforms to that paradigm

B. It uses distance-vector routing

C. It floods all networks

D. It ignores policies

The mechanism that makes BGP highly applicable is its policy-based approach to interdomain routing. BGP is designed to exchange reachability between autonomous systems and to carry policy information that operators use to control which paths traffic should take. As a path-vector protocol, it tracks the route's journey through ASes (AS_PATH) and uses attributes like LOCAL_PREF, MED, and communities to influence path selection. This lets networks encode business relationships (who is a peer, who is a transit provider) and perform traffic engineering, all while scaling to the Internet's size. This fits the hop-by-hop style of Internet routing, with decisions made at each border between ASes. The other options don't fit because BGP isn't distance-vector, it doesn't flood the network, and it is designed to enforce rather than ignore policies.

2. What are the four BGP message types?

A. Open, Keepalive, Update and Notification

B. Hello, Update, Withdraw, Tear-Down

C. Start, Stop, Update and Notify

D. Connect, Open, Keepalive and Shutdown

BGP operates with four message types that drive session setup, maintenance, and routing updates. Open is used to establish a peering session and negotiate parameters between peers. Keepalive exchanges maintain the session and confirm that both sides are still reachable. Update messages carry the actual routing information, including new paths and withdrawals of previously advertised routes. Notification signals errors or issues and can terminate the session when something goes wrong. This set—Open, Keepalive, Update, and Notification—is specific to BGP. Other terms like Hello belong to other routing protocols (for example, IGPs), and Start, Stop, Tear-Down, Connect, or Shutdown are not BGP message types.

3. True or False: A BGP router forms a neighbor relationship with a large number of BGP routers

A. True

B. Depends on the router

C. BGP supports only a single neighbor

D. False

BGP sessions are established per neighbor, and how many you can maintain is not fixed by the protocol itself. The number of neighbor relationships a router can form depends on the device's resources and software capabilities. High-end routing platforms used by large operators are designed to support many hundreds or even thousands of BGP sessions, while smaller or less capable devices may be limited to just a handful. This is where practical scalability comes in: the limits come from memory for routing tables, CPU for processing updates, and how the device implements BGP state management. In addition, architectural choices like using route reflectors or confederations in an iBGP environment can affect how many peers you actually need to reach the same routing goals. So the statement is not universally true or false; it depends on the router.

4. In BGP, what characterizes the open sent state?

A. The router received agreement on the parameters for establishing a session.

B. An open message was sent with the parameters for the BGP session.

C. The router found a route to the neighbor and has completed the three-way handshake.

D. Peering is established and routing begins.

OpenSent represents the moment BGP starts negotiating the session parameters. After the TCP connection is up, the speaker sends an OPEN message that carries the session details: BGP version (usually 4), its own AS number, the hold time, the router ID, and any optional capabilities. This is the step where the parameters for the BGP session are proposed. The router then waits for the neighbor's OPEN and, once validated, moves toward OpenConfirm and finally Established. The other statements describe later stages or different events (receiving the neighbor's parameter agreement, TCP handshake, or actually establishing the peering and starting to route), which occur after OpenSent.

5. What is the next-hop attribute?

A. The IP address used to reach the next AS

B. The origin code

C. The AS path length

D. The MED value

In BGP, the next-hop attribute tells you the IP address to send packets to in order to reach the destination prefix—the next router along the path toward the destination. This address is what downstream routers use to forward traffic toward the destination, and it is typically the IP of the neighbor that advertised the route (for eBGP). For routes learned via iBGP, the next-hop is usually preserved from the origin unless you explicitly modify it (for example, with next-hop-self). This attribute is distinct from the origin code, the AS path length, or the MED value, which describe how the route was learned, how many ASes it traverses, and inbound traffic engineering, respectively.

6. Name the three well-known mandatory attributes.

A. AS path, Next Hop, Origin

B. Local Preference, Atomic Aggregate, Community

C. Aggregator, Community, Next Hop

D. Origin, Local Preference, AS Path

In BGP, three attributes are required in every route update because they provide essential, universally-needed information about how the route arrived and how to reach it. The AS_PATH lists the sequence of autonomous systems the route has traversed. This helps prevent routing loops and supports policy decisions because you can see the path the update took across the Internet. The NEXT_HOP tells you the IP address of the next router to reach the destination; without this, a router wouldn't know where to forward the packets. The ORIGIN attribute indicates how the route was learned—whether from an internal IGP, an external EGP, or if the origin is incomplete—so a router can make informed decisions when comparing multiple paths. The other attributes mentioned—Local Preference, Atomic Aggregate, Community, and Aggregator—are useful for shaping routing policies and route handling, but they're not required to be present on every update. Local Preference is a local policy knob within an AS, Atomic Aggregate and Aggregator relate to how routes are summarized or who performed aggregation, and Community is a tagging mechanism for policy purposes.

7. Which command enters BGP configuration mode?

A. Router bgp autonomous-system in global config

B. Router bgp in global config

C. bgp config enter

D. Enable BGP in interface mode

In BGP on most IOS devices, you enter BGP configuration mode from global configuration mode by issuing the command router bgp <AS-number>. This places you into the BGP router configuration context where you define neighbors, networks, route maps, and policies. The option that matches this practice uses the router bgp syntax with the autonomous system number specified inside global config, which is why it's the best choice. The other options aren't valid: there is no bgp config enter command, you don't enable BGP from interface mode, and simply typing router bgp without the AS number isn't the correct way to enter BGP configuration in most environments.

8. Which command displays the BGP state?

A. Show IP Route Summary

B. Show BGP Neighbors

C. Show Run BGP

D. Show IP BGP Summary

Understanding the current status of BGP peering is essential to know whether routes can be exchanged. The BGP session state shows how ready the TCP session with each neighbor is (Idle, Connect, Active, Established). The most direct way to view this for all peers at once is the Show IP BGP Summary command. It presents a compact table listing each neighbor and its session state, plus a quick glance at how many routes have been learned, so you can immediately see if peering is up and stable. Other commands pull in different information: Show IP Route Summary focuses on the routing table rather than the peering status; Show Run BGP reveals how BGP is configured but not the real-time state; Show BGP Neighbors may provide per-neighbor details but often doesn't give a single, concise view of the overall BGP session states across all peers.

9. How is the BGP router ID chosen?

A. Highest Loopback Interface

B. Highest Active Interface

C. Automatically Chosen Randomly

D. Configured Using the BGP Router-ID Command

The main point here is that BGP uses a unique identifier for each router in a BGP session, called the router ID. This value is what BGP uses to distinguish speakers in the network. Configuring the router ID with the BGP Router-ID command gives you a fixed, stable identifier that remains the same across reboots and interface changes. This stability is crucial because auto-selected IDs can change if interfaces go up or down or if IP addresses change, potentially disrupting existing neighbor sessions or causing unexpected updates in the BGP domain. By explicitly setting a router ID, you ensure a predictable identity for the router, which helps maintain consistent neighbor relationships and avoids ambiguity in the routing system. It's common to choose a loopback address for the router ID because loopbacks tend to be always up and provide a persistent, reachable-looking address, though the key is simply that the value is stable and unique within the routing environment. Relying on automatic selection is allowed, but it's more error-prone in practice: changes in interface state can lead to a different router ID being chosen, which can upset established BGP sessions or lead to confusion in network management. That's why explicitly configuring the router ID is considered the best practice.

10. Which statement about the BGP next-hop behavior for eBGP is true?

A. The next-hop is the IP address to reach the next AS

B. The next-hop is always the neighbor's interface address

C. The next-hop is always the next router's IP

D. There is no next-hop in BGP

In BGP, the next-hop attribute indicates the immediate router to forward to in order to reach the destination network. For routes learned via external BGP, the next-hop is set to the IP address of the advertising neighbor’s interface on the directly connected link. This means the next-hop is the IP address you would use to reach the next AS along the path—the address of the neighbor that announced the route. This behavior ensures you have a valid, directly reachable next hop to forward traffic toward the destination AS. It isn’t a generic “final destination” address, and it isn’t a global constant; it’s specifically the neighbor’s address on the link between the peers.

Next Steps

Congratulations on reaching the final section of this guide. You've taken a meaningful step toward passing your certification exam and advancing your career.

As you continue preparing, remember that consistent practice, review, and self-reflection are key to success. Make time to revisit difficult topics, simulate exam conditions, and track your progress along the way.

If you need help, have suggestions, or want to share feedback, we'd love to hear from you. Reach out to our team at hello@examzify.com.

Or visit your dedicated course page for more study tools and resources:

https://ripebgpsecurity.examzify.com

We wish you the very best on your exam journey. You've got this!