

MEF SOFTWARE- DEFINED WIDE AREA NETWORK (SD-WAN) PRACTICE EXAM (SAMPLE) STUDY GUIDE



SAMPLE STUDY GUIDE WITH LIMITED QUESTIONS

VISIT US ONLINE FOR
THE FULL VERSION STUDY GUIDE

<https://mefsdwan.examzify.com>



Copyright © 2026 by Examzify - A Kaluba Technologies Inc. product.

ALL RIGHTS RESERVED.

No part of this book may be reproduced or transferred in any form or by any means, graphic, electronic, or mechanical, including photocopying, recording, web distribution, taping, or by any information storage retrieval system, without the written permission of the author.

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

SAMPLE

Table of Contents

Copyright	1
Table of Contents	2
Introduction	3
How to Use This Guide	4
Questions	5
Answers	8
Explanations	10
Next Steps	15

SAMPLE

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

SAMPLE

- 1. The IP packet traveling from the subscriber toward the service provider is called the:**
 - A. Ingress IP Packet
 - B. Egress IP Packet
 - C. Underlay Connectivity Services
 - D. SD-WAN Edge

- 2. How does MEF SD-WAN support multi-vendor environments?**
 - A. By using open APIs, standardized interfaces, and an abstraction layer that allows different vendors' edge devices and controllers to interoperate.
 - B. By requiring all vendors to implement MEF-approved hardware only.
 - C. By isolating each vendor's domain with no cross-communication.
 - D. By using a single vendor's controller for all devices.

- 3. What term describes a pause between frames, potentially used for clock syncing or other purposes?**
 - A. Inter frame gap
 - B. Ingress IP Packet
 - C. Egress IP Packet
 - D. Internet Breakout / Local Internet Breakout

- 4. Which statement describes a feature of IPv6 link-local addressing?**
 - A. a Layer 2 service that can provide lower cost UCS with guaranteed minimum performance and availability (Carrier Ethernet)
 - B. It requires global routable addresses
 - C. Cannot function without routers
 - D. Is used for IPv4 only

- 5. What is the purpose of a Centralized Orchestrator in SD-WAN security?**
 - A. To provide device telemetry only.
 - B. To enforce consistent security policies, coordinate updates, collect telemetry, and ensure compliance across all edge devices and tenants.
 - C. To manage user access policy only.
 - D. To monitor bandwidth usage exclusively.

- 6. What is the function of a service graph in SD-WAN orchestration?**
- A. The sequence of policy actions applied to traffic
 - B. A log of traffic events
 - C. The policy enforcement engine at the edge
 - D. A device inventory
- 7. What do cloud egress/ingress metrics measure in SD-WAN?**
- A. Cloud egress/ingress metrics refer to traffic exiting or entering the cloud service, impacting performance and egress costs.
 - B. Only internal LAN throughput.
 - C. DNS query latency only.
 - D. Color of traffic.
- 8. Which metrics should be tracked to measure SD-WAN performance for cloud access?**
- A. Only bandwidth usage on the WAN link.
 - B. The color of the cabling at the edge.
 - C. The number of connected devices regardless of cloud access.
 - D. Application performance to cloud apps, end-to-end latency to cloud regions, jitter, packet loss, path reliability, and cloud egress/ingress metrics.
- 9. Which protocol operates at OSI Layer 3 and enables routing between Layer 2 networks?**
- A. IPv4
 - B. BFD
 - C. BGP
 - D. RIP
- 10. In SD-WAN, what is application proxying?**
- A. Routing traffic directly to the internet without optimization.
 - B. Routing traffic through a local or centralized proxy/gateway to optimize performance, apply security controls, or enable direct cloud access when appropriate.
 - C. Disabling proxies and using only direct host routes.
 - D. Installing proxies only for guest networks.

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. The IP packet traveling from the subscriber toward the service provider is called the:

- A. Ingress IP Packet**
- B. Egress IP Packet
- C. Underlay Connectivity Services
- D. SD-WAN Edge

Direction relative to a network boundary defines ingress versus egress. When a packet moves from the subscriber toward the service provider, it is entering the service provider's network from its perspective, so it is an ingress IP packet. Egress would be traffic leaving the service provider's network toward the subscriber. The other terms don't describe this directional labeling: underlay connectivity services refer to the transport system, and SD-WAN edge is the device at the edge, not the direction of a packet.

2. How does MEF SD-WAN support multi-vendor environments?

- A. By using open APIs, standardized interfaces, and an abstraction layer that allows different vendors' edge devices and controllers to interoperate.**
- B. By requiring all vendors to implement MEF-approved hardware only.
- C. By isolating each vendor's domain with no cross-communication.
- D. By using a single vendor's controller for all devices.

Interoperability across vendors is achieved by decoupling management from vendor-specific devices through open APIs, standardized interfaces, and an abstraction layer that presents a uniform control model to the orchestrator. In MEF SD-WAN, edge devices from different vendors expose common southbound interfaces and data models, and the orchestrator uses these open APIs to provision and manage services without caring about underlying vendor details. This enables a single control plane to coordinate devices from multiple vendors and deliver consistent services across the whole network. The other approaches would either lock you into one vendor or isolate domains, which would prevent true multi-vendor operation.

3. What term describes a pause between frames, potentially used for clock syncing or other purposes?

- A. Inter frame gap**
- B. Ingress IP Packet
- C. Egress IP Packet
- D. Internet Breakout / Local Internet Breakout

The pause between Ethernet frames is called the interframe gap. It's a mandated idle period defined by the Ethernet standards, measured in bit times, that separates consecutive frames. This gap lets devices finish processing the previous frame, prepare and time the next one, and can even aid receivers with timing or clock alignment in some designs. The other terms describe data traffic or routing concepts (IP packets moving in or out of a device, or local Internet breakout) rather than the timing gap between frames. So the interframe gap is the correct term for a pause between frames.

4. Which statement describes a feature of IPv6 link-local addressing?

- A. a Layer 2 service that can provide lower cost UCS with guaranteed minimum performance and availability (Carrier Ethernet)
- B. It requires global routable addresses
- C. Cannot function without routers
- D. Is used for IPv4 only

IPv6 link-local addresses are for on-link communications only, meaning they work within a single network segment and are not meant to be routed across the Internet. They're automatically configured on every IPv6-enabled interface and use the FE80::/10 prefix, typically forming a unique address without any manual setup or global routing configuration. Because they exist solely on the local link, these addresses don't require global routable addresses, and routers aren't needed for basic on-link communication. They're also IPv6-specific, not IPv4. The statements given don't describe this behavior: one refers to a Layer 2 service unrelated to addressing, another incorrectly claims global routable addresses are required, another suggests routers are essential, and another implies IPv4 usage only. The correct concept is that link-local IPv6 addresses are automatically configured, are local to the link, and are not globally routable.

5. What is the purpose of a Centralized Orchestrator in SD-WAN security?

- A. To provide device telemetry only.
- B. To enforce consistent security policies, coordinate updates, collect telemetry, and ensure compliance across all edge devices and tenants.
- C. To manage user access policy only.
- D. To monitor bandwidth usage exclusively.

In SD-WAN security, a Centralized Orchestrator provides a single control plane that manages and governs security across the entire network. It enforces consistent security policies across all edge devices and tenants, ensuring every site follows the same baseline rules and protections. It coordinates updates so policy changes, firmware, and software updates are pushed reliably to every device, avoiding drift and partial deployments. It collects telemetry from edge devices to give visibility into security events, health, and performance, enabling proactive monitoring and rapid incident response. It also ensures compliance by validating configurations against established security policies and regulatory requirements across the whole fleet and across multiple tenants. This holistic approach is what makes the centralized orchestrator the best answer. Relying on telemetry alone provides visibility but not enforcement. Focusing only on user access policy misses the broader security posture and multi-tenant governance. Monitoring bandwidth usage addresses throughput, not security policy orchestration or compliance.

6. What is the function of a service graph in SD-WAN orchestration?

- A. The sequence of policy actions applied to traffic
- B. A log of traffic events
- C. The policy enforcement engine at the edge
- D. A device inventory

In SD-WAN orchestration, a service graph expresses the sequence and combination of services that traffic should pass through as it moves through the network. It tells the controller the order in which policy actions and service functions (like firewall, DPI, NAT, QoS, encryption) are applied to different traffic flows, and how these services are chained across the path chosen for that traffic. This makes it possible to consistently deploy and enforce complex service chains across all edge devices. It's not a log of traffic events (which would be monitoring data), nor the policy enforcement engine at the edge (a runtime component), nor a device inventory (an asset listing). The service graph is the modeled plan of how services are applied to traffic.

7. What do cloud egress/ingress metrics measure in SD-WAN?

- A. Cloud egress/ingress metrics refer to traffic exiting or entering the cloud service, impacting performance and egress costs.
- B. Only internal LAN throughput.
- C. DNS query latency only.
- D. Color of traffic.

In SD-WAN, cloud egress and ingress metrics are about the traffic that leaves your network to reach cloud services and the traffic that returns from them. They track how much data is exchanged with cloud resources and how that traffic performs—latency, jitter, packet loss, and throughput—for both directions. This matters because cloud apps are accessed over the Internet or direct cloud connections, and cloud providers typically charge for data that leaves their networks. By monitoring these metrics, the SD-WAN can route cloud traffic along paths that deliver better performance and manage costs by avoiding costly or inefficient routes. This concept focuses on traffic to and from cloud resources, not on internal LAN throughput or DNS latency, and not on any meaningless attribute like color of traffic.

8. Which metrics should be tracked to measure SD-WAN performance for cloud access?

- A. Only bandwidth usage on the WAN link.
- B. The color of the cabling at the edge.
- C. The number of connected devices regardless of cloud access.
- D. Application performance to cloud apps, end-to-end latency to cloud regions, jitter, packet loss, path reliability, and cloud egress/ingress metrics.

Measuring SD-WAN performance for cloud access focuses on how cloud applications actually feel to users. The best set of metrics includes application performance to cloud apps, end-to-end latency to cloud regions, jitter, packet loss, path reliability, and cloud egress/ingress metrics. This combination captures user experience (latency, jitter, and loss), the health and reliability of the paths used to reach cloud services, and the specific behavior and cost implications of data entering and leaving cloud providers' networks. Since cloud traffic often travels across multiple networks and paths, monitoring end-to-end performance to the relevant cloud regions ensures you're measuring the actual experience rather than just a single interface metric. Focusing only on WAN bandwidth misses how well applications perform, and the color of cabling or the number of devices connected doesn't reflect cloud access quality or user experience.

9. Which protocol operates at OSI Layer 3 and enables routing between Layer 2 networks?

- A. IPv4**
- B. BFD
- C. BGP
- D. RIP

At Layer 3, routing packets across multiple networks relies on IP addressing. IPv4 provides the actual network-layer protocol used to forward packets between different Layer 2 domains (such as VLANs or separate Ethernet segments). Routers read the destination IP address and determine the next hop, enabling inter-network communication across diverse Layer 2 networks. Bidirectional Forwarding Detection isn't about routing decisions; it's a quick-failure-detection mechanism. Border Gateway Protocol is a routing protocol that runs over TCP and is used mainly for inter-domain routing between autonomous systems, not the fundamental layer-3 mechanism that moves IP packets across networks. Routing Information Protocol is a Layer 3 routing protocol, but it's one of several ways to learn routes; IPv4 is the core protocol that enables the actual routing of IP traffic between Layer 2 networks.

10. In SD-WAN, what is application proxying?

- A. Routing traffic directly to the internet without optimization.
- B. Routing traffic through a local or centralized proxy/gateway to optimize performance, apply security controls, or enable direct cloud access when appropriate.**
- C. Disabling proxies and using only direct host routes.
- D. Installing proxies only for guest networks.

Application proxying in SD-WAN means steering application traffic through a local or centralized proxy or gateway so it can be optimized, inspected, and governed before it reaches its destination. The proxying point can perform WAN optimization (acceleration and caching), apply security controls (firewalls, threat protection, TLS inspection, access policies), and still allow direct, efficient access to cloud applications when policy says it's appropriate. This centralized handling enables consistent security and performance for different apps, including SaaS and cloud services, while giving the network the flexibility to choose the best path or proxy based on the application and policy. This is why routing traffic through a proxy/gateway to optimize performance, enforce security controls, or enable direct cloud access when suitable is the correct concept. Routing directly to the internet without optimization bypasses these benefits, disabling proxies removes the control point, and proxies only for guest networks don't capture the broader, policy-driven approach used in SD-WAN.

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://mefsdwan.examzify.com>

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

SAMPLE