

ITGSS Certified DevOps Engineer Practice Test (Sample)

Study Guide



Everything you need from our exam experts!

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	16

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. What is a Service Level Agreement (SLA)?**
 - A. A guideline for system architecture**
 - B. A contract that defines the expected level of service between a provider and a customer**
 - C. A document outlining the troubleshooting process**
 - D. A competitive analysis of service providers**

- 2. What is the benefit of using containers in DevOps?**
 - A. They replace the need for cloud storage**
 - B. They automate coding tasks**
 - C. They enable consistent application deployment across various environments**
 - D. They guarantee zero downtime for applications**

- 3. Which tool is commonly used for automation in DevOps?**
 - A. Photoshop**
 - B. Excel**
 - C. Ansible**
 - D. PowerPoint**

- 4. The architecture of Kubernetes is primarily designed for what purpose?**
 - A. To run single container applications**
 - B. To manage distributed applications at scale**
 - C. To simplify networking**
 - D. To manage database interactions**

- 5. Which statement best defines containerization?**
 - A. A process for creating virtual machines**
 - B. A technique for packaging applications and their dependencies into isolated units**
 - C. A method for encrypting data during transmission**
 - D. A way to scale applications horizontally without redundancy**

6. How does Continuous Delivery (CD) differentiate from Continuous Deployment?

- A. Continuous Delivery requires manual approval for deployment, while Continuous Deployment automates the deployment process**
- B. Continuous Delivery is only applicable to web applications, whereas Continuous Deployment is for mobile apps**
- C. Continuous Delivery focuses on testing environments, while Continuous Deployment focuses on production environments**
- D. Continuous Delivery involves fewer integrations than Continuous Deployment**

7. What is the primary function of the Service Mesh?

- A. Storing application logs**
- B. Coordinating supporting services for microservices applications**
- C. Deploying new microservices**
- D. Managing resources in Kubernetes**

8. What is a common tool used for Configuration Management?

- A. Ansible, Chef, or Puppet**
- B. Jenkins, Travis, or CircleCI**
- C. React, Angular, or Vue.js**
- D. MySQL, MongoDB, or PostgreSQL**

9. What feature of the Linux kernel allows processes to see different sets of resources?

- A. Control groups**
- B. Namespaces**
- C. Modules**
- D. Allocators**

10. Which security controls should not conflict in a Kubernetes environment?

- A. Network segmentation and log management**
- B. Resource allocation and service discovery**
- C. Security policies and orchestrator functions**
- D. User permissions and roles**

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. What is a Service Level Agreement (SLA)?

- A. A guideline for system architecture
- B. A contract that defines the expected level of service between a provider and a customer**
- C. A document outlining the troubleshooting process
- D. A competitive analysis of service providers

A Service Level Agreement (SLA) is fundamentally a contract that specifies the expected level of service between a service provider and a customer. It defines measurable performance metrics, such as uptime, response times, and support availability, which both parties agree to uphold. The purpose of an SLA is to establish clear expectations and responsibilities, thereby facilitating better communication and understanding between the provider and the customer. By setting these clear guidelines, an SLA helps in holding the provider accountable for delivering the promised level of service, while also allowing the customer to understand what to expect in terms of service quality and performance. This is crucial for maintaining a good business relationship, as both parties have a reference point for service obligations. Other options relate to different aspects of service management or documentation but do not capture the essence of what an SLA is meant to achieve in terms of service provision and customer expectations. Ultimately, an SLA plays a vital role in ensuring that both the service provider and customer share a mutual understanding of what constitutes proper service delivery.

2. What is the benefit of using containers in DevOps?

- A. They replace the need for cloud storage
- B. They automate coding tasks
- C. They enable consistent application deployment across various environments**
- D. They guarantee zero downtime for applications

The benefit of using containers in DevOps primarily revolves around enabling consistent application deployment across various environments. Containers encapsulate an application and its dependencies, ensuring that it runs uniformly regardless of where it is deployed—be it on local machines, virtual machines, or in the cloud. This consistency reduces the common issue of “it works on my machine” syndrome, where applications function properly in one environment but fail in another due to differences in configurations, libraries, or system settings. By using containers, teams can package their applications along with everything they need to run, including code, runtime, libraries, and environment variables. This results in a streamlined process for moving applications from development to testing to production without the usual pitfalls of environmental mismatch, leading to faster deployments and more reliable application performance. In contrast, the other options do not accurately capture the role of containers: - While cloud storage is critical in many scenarios, containers do not replace the need for it; they operate independently and can utilize it as needed. - Containers do not specifically automate coding tasks; instead, they facilitate the deployment and running of applications. Automation is often a separate concern handled by other tools and practices in the DevOps pipeline. - Although containers can help reduce downtime during deployments through strategies like blue-green

3. Which tool is commonly used for automation in DevOps?

- A. Photoshop**
- B. Excel**
- C. Ansible**
- D. PowerPoint**

Ansible is a widely used tool in the DevOps ecosystem for automation. It is designed to automate software provisioning, configuration management, and application deployment. Ansible uses a simple and human-readable language (YAML) to define the automation processes, which makes it accessible to both developers and operations teams. The primary strength of Ansible lies in its agentless architecture and its ability to manage systems over SSH or WinRM, eliminating the need for additional software agents to be installed on managed servers. This is particularly beneficial in DevOps for achieving greater efficiency and consistency in deployments across different environments. Furthermore, Ansible plays a pivotal role in continuous integration and continuous deployment (CI/CD) by enabling infrastructure as code (IaC). This allows teams to version control their infrastructure, ensuring that changes are trackable and reversible. Beyond its ease of use, Ansible's wide adoption in the industry means that it has strong community support and a plethora of modules that extend its capabilities, making it suitable for using in various stages of the software development lifecycle. In contrast, the other tools listed are primarily designed for different purposes: Photoshop is for graphic design, Excel for data analysis and visualization, and PowerPoint for presentations. None of them provide the automation capabilities essential for

4. The architecture of Kubernetes is primarily designed for what purpose?

- A. To run single container applications**
- B. To manage distributed applications at scale**
- C. To simplify networking**
- D. To manage database interactions**

The architecture of Kubernetes is primarily designed to manage distributed applications at scale. This platform provides a robust framework for automating the deployment, scaling, and management of containerized applications across clusters of hosts. Kubernetes excels in orchestrating containers, allowing developers to efficiently manage workloads and simplify the complexities involved in running applications distributed over a large number of servers. It handles potentially hundreds or thousands of containers, ensuring that the applications remain highly available, resilient, and able to automatically recover from failures. Features such as rolling updates and scaling capabilities are fundamental to its architecture, enabling seamless adjustments in resource allocation depending on demand. The other options do not capture the essence of Kubernetes' design intent. Running single container applications does not leverage the full capabilities of Kubernetes, which is designed to handle multiple containers and services. Simplifying networking and managing database interactions are certainly aspects of Kubernetes, but they are not the primary purposes of its architecture. Instead, Kubernetes encompasses a broader scope focused on distributed application management, making it a pivotal tool in modern DevOps practices.

5. Which statement best defines containerization?

- A. A process for creating virtual machines
- B. A technique for packaging applications and their dependencies into isolated units**
- C. A method for encrypting data during transmission
- D. A way to scale applications horizontally without redundancy

Containerization is best defined as a technique for packaging applications and their dependencies into isolated units. This approach allows developers to create, deploy, and run applications consistently across different computing environments. Each container includes everything needed for the application to run, such as the code, runtime, libraries, and system tools, ensuring that it operates the same way regardless of where it is deployed. This isolation is one of the key benefits of containerization, as it minimizes conflicts between software versions and dependencies by providing a self-contained environment. This is particularly valuable in DevOps practices, where rapid deployment and scalability are essential. Containers can be easily distributed and managed, making them ideal for microservices architectures and cloud-native applications. In contrast, other choices discuss concepts that do not encapsulate the essence of containerization. For instance, the idea of creating virtual machines refers to a different technology that involves full operating system virtualization, while encrypting data relates to securing information during transmission rather than packaging software. Lastly, scaling applications horizontally without redundancy is more about infrastructure design rather than how applications are packaged and managed.

6. How does Continuous Delivery (CD) differentiate from Continuous Deployment?

- A. Continuous Delivery requires manual approval for deployment, while Continuous Deployment automates the deployment process**
- B. Continuous Delivery is only applicable to web applications, whereas Continuous Deployment is for mobile apps
- C. Continuous Delivery focuses on testing environments, while Continuous Deployment focuses on production environments
- D. Continuous Delivery involves fewer integrations than Continuous Deployment

Continuous Delivery (CD) is a software development practice that ensures code changes are automatically prepared for a production release. One of its primary characteristics is that, although the deployment process is automated up to the point of production, a manual approval step is often required before the code is actually released into production. This allows teams to verify and validate changes before they go live, ensuring that everything is functioning as intended and meets business goals or compliance requirements. In contrast, Continuous Deployment takes this a step further by completely automating the release process without requiring any manual intervention. Anytime code passes the automated testing phase, it is immediately deployed to production. This swift approach allows teams to release updates and features more rapidly and frequently, reducing the time between development and user feedback. The other choices do not accurately capture the distinction between Continuous Delivery and Continuous Deployment. Continuous Delivery is not limited to specific types of applications, nor does it primarily focus on testing environments; it encompasses a wide range of applications and environments. Moreover, the frequency of integrations is not a defining factor between the two practices, as both Continuous Delivery and Continuous Deployment rely on regular integration of code changes.

7. What is the primary function of the Service Mesh?

- A. Storing application logs**
- B. Coordinating supporting services for microservices applications**
- C. Deploying new microservices**
- D. Managing resources in Kubernetes**

The primary function of a Service Mesh revolves around facilitating the communication between microservices in a microservices architecture. It provides a dedicated infrastructure layer that allows different services to talk to each other, managing how they interact, connect, and secure these communications. This layer often handles concerns such as service discovery, traffic management, load balancing, failure recovery, metrics, and monitoring, which are crucial in a microservices environment. By coordinating supporting services, the Service Mesh ensures that requests between services are efficiently routed, making it easier to manage complex inter-service communications. This enables developers to focus more on the business logic of their applications rather than on the intricacies of network communication and service interactions. In contrast, storing application logs is a distinct function typically handled by logging systems or tools instead of a Service Mesh. Deploying new microservices is generally part of a CI/CD (Continuous Integration/Continuous Deployment) pipeline and is not a core responsibility of the Service Mesh. Managing resources in Kubernetes pertains to orchestration and infrastructure management rather than the specialized communication functionalities that a Service Mesh provides.

8. What is a common tool used for Configuration Management?

- A. Ansible, Chef, or Puppet**
- B. Jenkins, Travis, or CircleCI**
- C. React, Angular, or Vue.js**
- D. MySQL, MongoDB, or PostgreSQL**

The choice of Ansible, Chef, or Puppet as a common tool for Configuration Management is accurate because these tools are specifically designed for managing and automating the configuration of systems and software across various environments. Ansible is known for its simplicity and use of YAML for defining configurations, enabling users to automate software provisioning, configuration management, and application deployment with ease. Chef uses a more code-intensive approach but provides strong capabilities for configuration with its Ruby-based DSL (domain-specific language). Puppet, on the other hand, is widely recognized for its declarative language, helping system administrators define the desired state of system configurations, which it can enforce across multiple servers. These tools help ensure consistency in system configurations, enabling teams to maintain system integrity, automate real-time deployment, and reduce manual intervention errors, which is a cornerstone of effective DevOps practices. Other options consist of tools that serve different purposes. Jenkins, Travis, and CircleCI are primarily continuous integration and continuous deployment (CI/CD) tools; they focus on automating the process of building, testing, and deploying applications rather than managing system configurations. React, Angular, and Vue.js are frameworks used for building user interfaces, while MySQL, MongoDB, and PostgreSQL are database management systems, thus making

9. What feature of the Linux kernel allows processes to see different sets of resources?

- A. Control groups
- B. Namespaces**
- C. Modules
- D. Allocators

Namespaces are a fundamental feature of the Linux kernel that enable the isolation of system resources for different processes. By utilizing namespaces, each process can operate with its own view of system resources such as process IDs, user IDs, file systems, network interfaces, and more. This allows applications to run in a confined environment without interference from others, thereby enhancing security and resource management. For instance, in a containerized environment, different containers can use the same process IDs or file system paths without conflicting with one another because each container is utilizing its own namespace. This encapsulation is crucial for resource allocation and process management in modern operating systems, especially in cloud computing and DevOps practices where multiple applications may share the same underlying infrastructure but need to function independently of each other. Other options, while important components of the Linux kernel, do not provide the same level of resource isolation. Control groups focus on resource management and limiting resource usage, but they do not enable distinct views of resources for processes. Modules refer to loadable kernel modules that extend the kernel's capabilities, and allocators manage memory allocation but do not contribute to resource visibility or isolation as namespaces do. Thus, namespaces are the specific feature that allows processes to perceive different sets of resources.

10. Which security controls should not conflict in a Kubernetes environment?

- A. Network segmentation and log management
- B. Resource allocation and service discovery
- C. Security policies and orchestrator functions**
- D. User permissions and roles

In a Kubernetes environment, security policies and orchestrator functions must operate in harmony to ensure optimal security without disrupting essential orchestration activities. Security policies define the rules and guidelines that control access and permissions within the Kubernetes cluster, ensuring that only authorized users and processes can access or interact with resources. On the other hand, orchestrator functions, such as deployment and scaling of applications, rely on these security policies to manage workloads effectively without compromising the integrity or security of the environment. When security policies and orchestrator functions are aligned, it ensures that policies are enforced without hindering the automation and management capabilities of the orchestrator. For example, a well-implemented security policy would prevent unauthorized access while still allowing the orchestrator to function normally by incorporating role-based access controls or network policies that do not interfere with deployment processes. Overall, the correct answer highlights the necessity for seamless interaction between security measures and the operational mechanisms of Kubernetes. In contrast, the other options represent areas where conflicts may arise. For instance, network segmentation and log management can introduce complexities in monitoring and analyzing traffic patterns if overly restrictive network policies are in place. Resource allocation and service discovery are fundamentally operational concerns that need to balance performance with security but can produce conflicts when resource constraints inhibit service discovery processes. Lastly

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://itgss-certifieddevopsengineer.examzify.com>

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

SAMPLE