

Splunk Enterprise Certified Architect Practice Test (Sample)

Study Guide



Everything you need from our exam experts!

This is a sample study guide. To access the full version with hundreds of questions,

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 from reliable sources accurate, complete, and timely information about this product.

SAMPLE

Table of Contents

Copyright	1
Table of Contents	2
Introduction	3
How to Use This Guide	4
Questions	6
Answers	10
Explanations	12
Next Steps	18

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. Don't worry about getting everything right, 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, and take breaks to retain information better.

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.

7. Use Other Tools

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!

SAMPLE

Questions

SAMPLE

- 1. What does the command 'splunk clean eventdata' do?**
 - A. Removes indexed data from Splunk.**
 - B. Clears the internal database.**
 - C. Deletes downloaded apps from a Splunk instance.**
 - D. Clears user session data.**

- 2. Which statement is true about Splunk indexer clustering?**
 - A. All peer nodes must run exactly the same Splunk version.**
 - B. The master node must run the same or a later Splunk version than search heads.**
 - C. The peer nodes must run the same or a later Splunk version than the master node.**
 - D. The search head must run the same or a later Splunk version than the peer nodes.**

- 3. How can you optimize Splunk's license usage?**
 - A. By increasing index size.**
 - B. By filtering out unnecessary data before ingestion.**
 - C. By ingesting more data without restrictions.**
 - D. By disabling alerts during high traffic.**

- 4. Which of the following is true regarding deploying Enterprise Security configurations?**
 - A. Only a single configuration file needs to be copied**
 - B. Configurations should be verified on a staging instance first**
 - C. Configurations can be modified directly on the search heads**
 - D. All configuration changes require a system reboot**

- 5. Why is RAID 10 recommended for Splunk indexing?**
 - A. It provides high performance and redundancy.**
 - B. It allows for cheaper storage options.**
 - C. It simplifies data recovery processes.**
 - D. It is the only option available.**

6. Which of the following best describes the role of a search head in Splunk?

- A. Stores indexed data.**
- B. Distributes search requests to indexers.**
- C. Serves as the primary data input source.**
- D. Acts as a backup for master nodes.**

7. How does a Heavy Forwarder differ from a Universal Forwarder in Splunk?

- A. A Heavy Forwarder can parse and index data before forwarding**
- B. A Universal Forwarder can only send data to other destinations**
- C. A Heavy Forwarder is used solely for data visualization**
- D. A Universal Forwarder analyzes data locally before sending**

8. What might prevent a colleague from seeing the src_ip field in their search results?

- A. The field was extracted as a private knowledge object.**
- B. The events are tagged as communicate, but are missing the network tag.**
- C. The Typing Queue is blocked.**
- D. The colleague did not explicitly use the field in the search and the search was set to Fast Mode.**

9. Which methods can be used to transfer captaincy in a search head clustering?

- A. Use the Monitoring Console.**
- B. Run the splunk transfer shcluster-captain command from the current captain.**
- C. Change settings from Splunk Web on any member.**
- D. Run the transfer command from the member you want to become captain.**

10. Which of the following is NOT considered a type of data that Splunk can ingest?

- A. Log files**
- B. Wireless network data**
- C. Only highly structured database records**
- D. Machine data**

SAMPLE

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. What does the command 'splunk clean eventdata' do?

- A. Removes indexed data from Splunk.**
- B. Clears the internal database.**
- C. Deletes downloaded apps from a Splunk instance.**
- D. Clears user session data.**

The command 'splunk clean eventdata' is specifically used to remove indexed data from a Splunk instance. When executed, it purges all events that have been indexed, effectively resetting the indexing process by deleting the data stored in the indexing directories. This command is crucial when there is a need to clear out data for testing or reconfiguration purposes — for instance, when re-indexing is required due to data corruption or changes in data structures. This command operates at a very low level within Splunk, and it is essential for users to understand that using it will result in the permanent loss of all indexed data at the specified index location. It is generally recommended to use this command with caution, particularly in production environments, because once the data is cleaned, it cannot be recovered. When considering the other options, they pertain to different functionalities: clearing the internal database involves different maintenance activities to ensure Splunk runs smoothly; deleting downloaded apps refers to the management of applications within Splunk; clearing user session data is relevant for managing user access and security but does not relate to indexed data removal. Each of these functions addresses specific components of the Splunk environment, making the command 'splunk clean eventdata' distinct in its purpose and application.

2. Which statement is true about Splunk indexer clustering?

- A. All peer nodes must run exactly the same Splunk version.**
- B. The master node must run the same or a later Splunk version than search heads.**
- C. The peer nodes must run the same or a later Splunk version than the master node.**
- D. The search head must run the same or a later Splunk version than the peer nodes.**

The statement that all peer nodes in a Splunk indexer cluster must run exactly the same Splunk version is true because consistency across peer nodes ensures that data is indexed in a uniform manner. When all peer nodes operate on the same version, it eliminates any potential discrepancies in functionality, features, or indexing processes that could arise from version differences. This uniformity is essential for maintaining data integrity and properly coordinating operations within the cluster. In an indexer cluster, if peer nodes were running different versions, it could lead to complications such as differing behaviors during data replication or potential errors during searches. Deploying the same version across all peers ensures that they can effectively communicate with one another without encountering version-specific limitations or bugs. This concept highlights the importance of version control and management within clusters, as ensuring that all components are on the same page can significantly contribute to the stability and reliability of the entire Splunk infrastructure.

3. How can you optimize Splunk's license usage?

- A. By increasing index size.
- B. By filtering out unnecessary data before ingestion.**
- C. By ingesting more data without restrictions.
- D. By disabling alerts during high traffic.

Optimizing Splunk's license usage is essential for managing costs and ensuring efficient data processing. Filtering out unnecessary data before ingestion is a key strategy because it directly reduces the volume of data that Splunk needs to index and store. When unnecessary or low-value data is excluded, it lowers your overall data ingestion volume, subsequently decreasing the number of license bytes consumed. This approach not only helps in staying within license limits but also enhances performance and allows for more meaningful data analysis. By being selective about what data to ingest, organizations can focus their Splunk resources on the most relevant and impactful information, improving both efficiency and insights generated from the data. Other strategies may not effectively address the objective of optimizing license usage, as merely increasing index size or ingesting more data can lead to exceeding license limits, while disabling alerts does not contribute to license savings but may instead impact system monitoring and operational effectiveness.

4. Which of the following is true regarding deploying Enterprise Security configurations?

- A. Only a single configuration file needs to be copied
- B. Configurations should be verified on a staging instance first**
- C. Configurations can be modified directly on the search heads
- D. All configuration changes require a system reboot

Verifying configurations on a staging instance first is a best practice in deploying Enterprise Security configurations. This approach allows for testing changes in a controlled environment, minimizing the risk of errors that could affect production environments. By validating the configurations in a staging area, administrators can ensure that any adjustments behave as expected, identify potential issues, and resolve them without impacting the live system. Additionally, this method promotes thorough testing of integrations and dependencies that may not be fully visible until the configurations are actively in use. It provides an opportunity to confirm that all components work harmoniously before deploying to production, thus enhancing overall system stability and reliability. In contrast, the other options do not align with best practices for configuration management. For example, simply copying a single configuration file may overlook the need for the interdependence of multiple files in the configuration process. Modifying configurations directly on search heads can lead to inconsistencies and challenges in maintaining version control, making it more difficult to track changes across deployments. Furthermore, requiring a system reboot for all configuration changes is typically not necessary in Splunk, as many changes can be applied dynamically, allowing for greater flexibility and uptime.

5. Why is RAID 10 recommended for Splunk indexing?

- A. It provides high performance and redundancy.**
- B. It allows for cheaper storage options.**
- C. It simplifies data recovery processes.**
- D. It is the only option available.**

RAID 10 is recommended for Splunk indexing primarily because it combines the benefits of high performance and redundancy. This configuration achieves both by mirroring data across multiple disks (which provides redundancy) and striping data across those mirrored sets (which enhances read and write performance). In the context of Splunk, where fast access to indexed data is crucial for search and analytics, the ability to rapidly read and write large volumes of data is essential. When data is mirrored, even if one disk fails, the data is still safe and accessible from the other mirrored disk, ensuring that the system remains reliable and minimizes downtime. Other choices do not align as closely with the specific requirements of Splunk indexers. While RAID configurations can potentially involve cost considerations or ease of recovery, these attributes don't directly address the core demands of indexing operations. RAID 10's combination of speed and fault tolerance makes it a robust choice for Splunk's data-heavy environment, thus firmly establishing it as the preferred choice for indexing tasks.

6. Which of the following best describes the role of a search head in Splunk?

- A. Stores indexed data.**
- B. Distributes search requests to indexers.**
- C. Serves as the primary data input source.**
- D. Acts as a backup for master nodes.**

The role of a search head in Splunk is fundamentally to manage and distribute search requests to indexers. The search head receives search queries from users and breaks these queries down into smaller tasks that can be executed in parallel across multiple indexers. This approach optimizes the search process by efficiently utilizing resources and speeding up results. Once the indexers complete their part of the search, the search head aggregates the results and presents them to the user. This role is crucial in ensuring that users can retrieve vast amounts of data quickly, effectively managing the resource-intensive nature of searching across potentially large sets of indexed data. In contrast, the other roles mentioned do not accurately describe the function of a search head. For instance, the storing of indexed data is a responsibility primarily held by the indexers in Splunk, as they are the components responsible for indexing and storing incoming data. While data input sources are managed by forwarders, which collect and send data to indexers, the search head is not involved in this process. Additionally, acting as a backup for master nodes pertains to cluster management and data integrity, which does not fall within the scope of a search head's responsibilities.

7. How does a Heavy Forwarder differ from a Universal Forwarder in Splunk?

- A. A Heavy Forwarder can parse and index data before forwarding**
- B. A Universal Forwarder can only send data to other destinations**
- C. A Heavy Forwarder is used solely for data visualization**
- D. A Universal Forwarder analyzes data locally before sending**

A Heavy Forwarder is designed to perform more advanced data processing than a Universal Forwarder. It can both parse and index data before sending it on to other destinations, such as an indexer or another forwarder. This capability makes the Heavy Forwarder useful in scenarios where data transformation or filtering needs to occur prior to forwarding. For example, it can apply additional parsing rules, enrich data with metadata, or do simple indexing operations, which can be beneficial in optimizing data flow and enhancing the efficiency of the ingestion process. In contrast, the Universal Forwarder is a lightweight agent primarily focused on one function: reliably transmitting log data to another Splunk instance without applying any substantial processing. Its design emphasizes minimal resource usage on the source machine, making it ideal for environments where you want to collect data without overwhelming the system's resources.

8. What might prevent a colleague from seeing the src_ip field in their search results?

- A. The field was extracted as a private knowledge object.**
- B. The events are tagged as communicate, but are missing the network tag.**
- C. The Typing Queue is blocked.**
- D. The colleague did not explicitly use the field in the search and the search was set to Fast Mode.**

The correct answer highlights an important aspect of how fields are managed in Splunk. When a field is extracted as a private knowledge object, it means that only the user who created it has access to that field in their searches. Therefore, if a colleague is trying to access the src_ip field but it was extracted privately, they would not see that field in their search results. In Splunk, knowledge objects such as fields can be defined at different levels of visibility. Private knowledge objects are only accessible to the user who created them, whereas public knowledge objects can be accessed by all users. This private setting can limit visibility and is critical for maintaining the appropriate access to sensitive information or fields. In contrast, the other options relate to different mechanisms of access or functionality that would not directly limit the visibility of a field solely based on its accessibility status. For instance, if events are tagged in a specific way or if a queue is blocked, these issues do not inherently control the visibility of an extracted field. Similarly, the Fast Mode setting being used for searches can affect the performance and speed of the searches but does not justifiably prevent access to certain fields based on how they were defined when extracted. Understanding the concept of public vs private knowledge objects is essential.

9. Which methods can be used to transfer captaincy in a search head clustering?

- A. Use the Monitoring Console.**
- B. Run the `splunk transfer shcluster-captain` command from the current captain.**
- C. Change settings from Splunk Web on any member.**
- D. Run the transfer command from the member you want to become captain.**

Transferring captaincy in a search head cluster involves designating a new captain among the cluster members. The correct option states that you should run the transfer command from the member you want to become the captain. This method is effective because it allows you to explicitly indicate which member should take on the role of captain, ensuring control over the cluster's leadership. In a search head cluster, the captain is responsible for coordinating actions like managing search requests and distributing tasks among the other members. By executing the command from the intended new captain, you are facilitating a seamless transition and allowing the cluster to maintain its operation without interruption. Other methods, while maybe relevant in different contexts, do not directly involve initiating the captaincy transfer from the intended member. For instance, using the Monitoring Console typically provides a high-level overview without the direct ability to change roles. Running commands from the current captain or changing settings from Splunk Web may not ensure the intended outcome effectively, as they do not focus on the specific process required to transfer the captaincy directly.

10. Which of the following is NOT considered a type of data that Splunk can ingest?

- A. Log files**
- B. Wireless network data**
- C. Only highly structured database records**
- D. Machine data**

Splunk is designed to ingest a wide variety of data types, making it versatile in handling different sources of data relevant for indexing and searching. Among the options provided, the inclusion of only highly structured database records as a category of data that Splunk ingests is inaccurate. Splunk is particularly effective with semi-structured and unstructured data, such as log files, machine data, and even wireless network data. These types of data often do not conform to strict schemas and can be collected directly from diverse sources, thus allowing Splunk to offer powerful analytical capabilities. The emphasis on highly structured database records suggests a limitation in the context of what Splunk can handle; however, Splunk extends its functionality to connect with structured data sources as well. This makes the assertion that Splunk only ingests highly structured records misleading. Instead, it is more accurate to say that Splunk excels at ingesting various less-structured and unstructured data types, which allows users to gain insights from a broader spectrum of information.

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://splunkcertifiedarchitect.examzify.com>

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

SAMPLE