


SYSTEM SOFTWARE, ARCHITECTURE, MEMORY AND STORAGE PRACTICE TEST (SAMPLE) STUDY GUIDE



SAMPLE STUDY GUIDE WITH LIMITED QUESTIONS

VISIT US ONLINE FOR
THE FULL VERSION STUDY GUIDE

<https://syssoftwarearchimemorystorage.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. Which statement correctly defines clock speed in a CPU?**
 - A. The number of cores in a CPU
 - B. The total processing time of a program
 - C. The number of instructions a single processor core can carry out per second (Hz)
 - D. The amount of cache per core

- 2. What can cause a context switch after an interrupt service routine completes?**
 - A. The scheduler may switch to a higher-priority ready task after the ISR.
 - B. The system always returns to the interrupted task without scheduling.
 - C. The interrupt preempts the current thread until it completes.
 - D. No context switch occurs after an ISR.

- 3. What is primary storage?**
 - A. Non-volatile external storage where OS and applications are stored when not in use like HDD, SSD, CDs and SD cards.
 - B. External cloud storage.
 - C. Non-volatile memory used for long-term backups.
 - D. Mostly volatile memory areas that the CPU can access quickly like registers, RAM and cache.

- 4. During the boot process, which of the following is commonly performed by BIOS?**
 - A. It initializes the disk drives and loads user programs
 - B. It runs the OS kernel before loading BIOS settings
 - C. It copies the operating system into RAM to start it
 - D. It formats the hard drive

- 5. What does x16 denote in PCIe lane configurations?**
 - A. A single data channel.
 - B. 16 lanes providing higher bandwidth.
 - C. A 16-bit data path.
 - D. Sixteen physical lanes used together for increased bandwidth.

- 6. In memory management using paging, which attributes are conveyed by a page table entry?**
- A. Present, Read/Write, User/Supervisor, Dirty, Accessed
 - B. Present, Read-Only, User/Kernel, Dirty, Accessed
 - C. Present, Read/Write, User/Supervisor, Dirty, Accessed
 - D. Absent, Read/Write, User/Supervisor, Dirty, Accessed
- 7. Which unit is used to control the fetch-decode-execute cycle and data flow inside the CPU?**
- A. Arithmetic Logic Unit
 - B. Cache
 - C. Control Unit
 - D. Memory Address Register
- 8. What is thrashing and how is it related to the working set of a process?**
- A. Thrashing happens when the working set is larger than available memory, causing high page fault rates and paging CPU usage.
 - B. Thrashing happens when there is no memory pressure and no paging occurs.
 - C. Thrashing occurs when the OS runs out of disk space.
 - D. Thrashing happens when all processes share the same memory.
- 9. Which statement accurately describes the interrupt vector table and how the OS manages interrupt priorities (including priority levels, APICs, and nested/masked interrupts)?**
- A. All interrupts are handled with the same priority; there is no nesting or masking.
 - B. The vector table stores process IDs for scheduling.
 - C. The OS may use interrupt priority levels or APIC; nested or masked interrupts manage priority.
 - D. The vector table maps memory addresses for I/O devices.
- 10. What are the disadvantages of Hard Disk Drives (HDDs)?**
- A. Damaged by large impacts like being dropped.
 - B. Very fast against SSDs.
 - C. No power consumption.
 - D. Non-mechanical drive.

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. Which statement correctly defines clock speed in a CPU?

- A. The number of cores in a CPU
- B. The total processing time of a program
- C. The number of instructions a single processor core can carry out per second (Hz)**
- D. The amount of cache per core

Clock speed is the rate at which the CPU's clock ticks, measured in cycles per second (Hz). That ticking sets the pace for fetching, decoding, and executing instructions. A faster clock means more cycles occur each second, which generally allows more work to be done per second, although the exact instructions-per-second throughput also depends on the architecture and memory behavior. The statement that clock speed equals the number of instructions a core can carry out per second (in Hz) captures this per-second rate idea, which is why it's considered correct here. The other options describe separate things—how many cores you have, the total elapsed time to run a program, or how much cache is available per core—not the clock's speed.

2. What can cause a context switch after an interrupt service routine completes?

- A. The scheduler may switch to a higher-priority ready task after the ISR.**
- B. The system always returns to the interrupted task without scheduling.
- C. The interrupt preempts the current thread until it completes.
- D. No context switch occurs after an ISR.

After an interrupt service routine finishes, the system often runs the scheduler to decide what to execute next. In a preemptive system, an ISR can wake up or unblock a higher-priority task. When the ISR returns, the scheduler sees that a higher-priority task is now ready and performs a context switch to that task, saving the state of the current task and loading the state of the higher-priority one. This is why a context switch after an ISR is possible and often expected: the interrupt may reveal more urgent work that should run immediately. This isn't guaranteed to always resume the interrupted task, because if a higher-priority task became ready during the ISR, the scheduler will switch to that task. It's also not accurate to say there's no context switch after an ISR in a system with preemption, since the exit from the ISR is precisely when the scheduler decides whether a switch is needed. The idea that the interrupt preempts the current thread until it completes is a misunderstanding of how scheduling and ISR handling interact: the ISR runs, then the scheduler decides if a higher-priority task should run next, which can involve a context switch.

3. What is primary storage?

- A. Non-volatile external storage where OS and applications are stored when not in use like HDD, SSD, CDs and SD cards.
- B. External cloud storage.
- C. Non-volatile memory used for long-term backups.
- D. Mostly volatile memory areas that the CPU can access quickly like registers, RAM and cache.**

Primary storage is the fast, directly usable memory the CPU uses while executing programs. It is typically volatile, meaning it loses data when power is removed. This category includes the CPU's registers, the caches (L1/L2/L3), and RAM. These memory types are designed for very low latency access to the instructions and data the processor is actively working with. In contrast, non-volatile storage like HDDs, SSDs, optical disks, external cloud storage, or tape is used for long-term persistence and backup, not for the immediate, high-speed work of the CPU. So the correct idea is that primary storage consists of volatile, quickly accessible memory areas such as registers, RAM, and cache.

4. During the boot process, which of the following is commonly performed by BIOS?

- A. It initializes the disk drives and loads user programs
- B. It runs the OS kernel before loading BIOS settings
- C. It copies the operating system into RAM to start it**
- D. It formats the hard drive

During boot, firmware on the motherboard (the BIOS) initializes hardware, performs a short self-test, and then locates a bootable device. It loads the initial bootstrap code from that device into RAM and transfers control to it. That bootstrap loader is responsible for loading the operating system into memory so it can run. In this sense, the BIOS participates in bringing the OS into RAM to start it by moving the boot code into memory and handing off execution to it, which ultimately leads to the OS kernel being loaded into RAM. The other statements describe actions that aren't what BIOS does during the boot process: it doesn't load user programs directly at boot, it doesn't run the OS kernel before configuring BIOS settings, and it doesn't format the hard drive.

5. What does x16 denote in PCIe lane configurations?

- A. A single data channel.
- B. 16 lanes providing higher bandwidth.
- C. A 16-bit data path.
- D. Sixteen physical lanes used together for increased bandwidth.**

PCIe uses a scalable, high-speed serial interface. The xN notation shows how many physical lanes form one link. Each lane is a separate serial channel, and the lanes operate in parallel to boost total bandwidth. So x16 means sixteen lanes are bonded to work together as a single PCIe link, giving much more bandwidth than a single lane. The actual speed per lane depends on the PCIe generation (Gen 3, Gen 4, Gen 5), but the core idea is that sixteen lanes are used together to carry data. This is different from a 16-bit data path, since PCIe lanes carry serial data, not a single wide 16-bit parallel bus. A single data channel describes a single-lane link, which is not what x16 represents.

6. In memory management using paging, which attributes are conveyed by a page table entry?

- A. Present, Read/Write, User/Supervisor, Dirty, Accessed
- B. Present, Read-Only, User/Kernel, Dirty, Accessed
- C. Present, Read/Write, User/Supervisor, Dirty, Accessed**
- D. Absent, Read/Write, User/Supervisor, Dirty, Accessed

In paging, a page table entry carries both where the data lives and how the page can be used. The Present bit signals whether the page's frame is currently in physical memory; if not, the system handles a page fault to bring it in. The Read/Write bit sets the access permission: with it set, the page can be written to; if cleared, writes are disallowed. The User/Supervisor bit (often called User/Kernel) enforces protection between user-mode processes and the kernel, controlling who may access the page. The Dirty bit marks whether the page has been modified since it was loaded, guiding whether the page must be written back to storage when evicted. The Accessed (Referenced) bit indicates that the page has been read or written, supporting page-replacement decisions. The other options either mix up these permissions, use a different naming convention for privilege, or replace Present with Absent, which is not an independent attribute. So the combination Present, Read/Write, User/Supervisor, Dirty, and Accessed is the correct set conveyed by a page table entry.

7. Which unit is used to control the fetch-decode-execute cycle and data flow inside the CPU?

- A. Arithmetic Logic Unit
- B. Cache
- C. Control Unit**
- D. Memory Address Register

Coordinating the fetch-decode-execute cycle and directing data flow inside the CPU is the job of the Control Unit. It orchestrates pulling an instruction from memory, decoding it to understand the required operation, and generating the precise control signals that move data between registers, the Arithmetic Logic Unit, and memory so each step happens in the correct order and timing. The control unit itself doesn't perform arithmetic or logic; it directs those operations by telling other parts of the CPU what to do. The Arithmetic Logic Unit handles computations, the Cache speeds access to recently used data, and the Memory Address Register holds addresses for memory operations. Because only the Control Unit manages the sequencing and signaling for the entire instruction cycle, it is the correct choice.

8. What is thrashing and how is it related to the working set of a process?

- A. Thrashing happens when the working set is larger than available memory, causing high page fault rates and paging CPU usage.**
- B. Thrashing happens when there is no memory pressure and no paging occurs.
- C. Thrashing occurs when the OS runs out of disk space.
- D. Thrashing happens when all processes share the same memory.

Thrashing happens when the process's working set—the set of pages it actively uses over a recent window of time—cannot fit into physical memory. When the working set is larger than available RAM, the system spends most of its time swapping pages in and out, causing a flood of page faults and a lot of CPU time devoted to paging rather than useful work. This intense paging behavior is exactly what thrashing is: the CPU is busy managing memory instead of executing the processes. If there's no memory pressure and no paging, thrashing won't occur; running out of disk space or all processes sharing the same memory doesn't define thrashing in the same way.

9. Which statement accurately describes the interrupt vector table and how the OS manages interrupt priorities (including priority levels, APICs, and nested/masked interrupts)?

- A. All interrupts are handled with the same priority; there is no nesting or masking.
- B. The vector table stores process IDs for scheduling.
- C. The OS may use interrupt priority levels or APIC; nested or masked interrupts manage priority.**
- D. The vector table maps memory addresses for I/O devices.

The concept being tested is how interrupts are organized and serviced: there is a dedicated table that maps each interrupt to the code that should run, and the system uses priorities and masking to control which interrupts can interrupt others. The interrupt vector table (often realized as an IDT on x86 or a similar structure in other architectures) stores the addresses of the interrupt service routines. When a device signals an interrupt, the CPU uses the interrupt number to jump to the corresponding handler, so this table is all about routing to the right piece of code. Beyond routing, the OS can assign priority levels to interrupts and use a Programmable Interrupt Controller (APIC) to enforce them. Higher-priority interrupts can preempt lower-priority ones, which enables nesting: a high-priority interrupt can interrupt a handler that was started for a lower-priority interrupt. Masking is the mechanism that temporarily blocks certain interrupts (by software or hardware settings) to protect critical sections or prevent interference while a high-priority task is being completed. In modern systems, Local APICs and I/O APICs coordinate to deliver and prioritize interrupts across multiple cores, supporting both nesting and masking to keep the system responsive and correct under load. Remember, the vector table stores addresses of interrupt handlers, not process IDs or device memory maps, and interrupts do not all share the same priority.

10. What are the disadvantages of Hard Disk Drives (HDDs)?

- A. Damaged by large impacts like being dropped.**
- B. Very fast against SSDs.
- C. No power consumption.
- D. Non-mechanical drive.

Hard disk drives rely on moving parts—spinning platters and a movable read/write head. That mechanical design makes them particularly vulnerable to physical shocks, such as drops. A strong jolt can cause the head to crash into the platter or misalign, leading to data loss or drive failure. This is the primary downside when comparing HDDs to solid-state drives. The other statements don't describe disadvantages correctly: HDDs are not faster than SSDs, they do consume power (they spin the disks and move the head), and they are mechanical, not non-mechanical.

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://syssoftwarearchimemorystorage.examzify.com>

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

SAMPLE