

# Exam DevOps-SRE questions and answers



Our specialists check whether the contents of DevOps-SRE real exam are updated every day. If there are newer versions, they will be sent to users in time to ensure that users can enjoy the latest resources in the first time. In such a way, our DevOps-SRE guide materials can have such a fast update rate that is taking into account the needs of users. Users using our DevOps-SRE Study Materials must be the first group of people who come into contact with new resources. When you receive an update reminder from DevOps-SRE practice questions, you can update the version in time and you will never miss a key message.

The PeopleCert DevOps-SRE Certification Exam is ideal for individuals who seek to advance their careers in DevOps and SRE. It is also suitable for IT professionals who want to enhance their skills in managing IT infrastructure, applications, and services. PeopleCert DevOps Site Reliability Engineer (SRE) certification is recognized globally and is highly regarded by organizations seeking to hire professionals in this field.

**>> Exam DevOps-SRE Format <<**

## CHOOSE THE BEST PLATFORM FOR ACING THE Peoplecert DevOps-SRE EXAM

You may be given the Peoplecert DevOps-SRE practice exam results as soon as they have been saved in the software. NewPassLeader modified Peoplecert DevOps-SRE exam dumps allow students to learn effectively about the real Peoplecert DevOps-SRE Certification Exam. Peoplecert DevOps-SRE practice exam software allows students to review and refine skills in a preceding test setting.

## Peoplecert PeopleCert DevOps Site Reliability Engineer (SRE) Sample Questions (Q23-Q28):

### NEW QUESTION # 23

What metrics will embracing failure help to improve?

- A. Mean time to detect and mean time between system incidents
- B. Empirical test data and mean time to recover service
- C. Change lead time and change failure rate
- D. Mean time to detect and mean time to recover

**Answer: D**

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

Embracing failure-through practices such as blameless postmortems, chaos engineering, and proactive detection-enables organizations to improve their incident response performance. This directly improves:

- \* MTTD (Mean Time to Detect)
- \* MTTR (Mean Time to Recover)

The Site Reliability Engineering Book, chapter "Postmortem Culture," states:

"By examining failures without blame and learning from them, organizations improve their ability to detect issues faster and recover more quickly." Similarly, in the SRE Workbook, section on incident response:

"Learning from incidents is essential to reducing time to detection and time to mitigation." Why the other options are incorrect:

\* A MTBSI (Mean Time Between System Incidents) is influenced by architecture and testing, not directly by embracing failure.

\* B These are DORA metrics - important, but not primarily tied to failure-embracing practices.

\* C Too vague and not a standard SRE metric pair.

Thus, D is the correct answer.

References:

Site Reliability Engineering Book, "Postmortem Culture"

SRE Workbook, "Incident Response"

#### NEW QUESTION # 24

When outages are repetitive and similar they become a form of toil.

Which of the following describes the MOST compelling reason to adopt advanced technologies and artificial intelligence (AI)?

- A. To increase the mean time to restore services (MTRS)
- B. To increase reliability and achieve perfect MTRS
- C. To increase reliability by reducing MTTR and MTRS
- D. To increase the mean time to repair services (MTTR)

Answer: C

#### NEW QUESTION # 25

Which of these approaches can alleviate linear scaling toil?

- A. Outsourcing development
- B. Manual scaling of services
- C. Using auto-scaling capabilities
- D. Switching cloud providers

Answer: C

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

Linear-scaling toil refers to work whose effort increases proportionally to service growth, such as manually provisioning servers or handling capacity expansion. The Google SRE Book, Chapter "Eliminating Toil," explains:

"Toil is work that scales linearly with the size of your service. A core strategy for reducing toil is to introduce automation that breaks the linear relationship." Auto-scaling capabilities directly address linear-scaling toil by automating resource allocation based on load or demand. This prevents engineers from repeatedly and manually adjusting infrastructure as usage grows.

The SRE Workbook also emphasizes:

"Infrastructure automation such as auto-scaling removes a major source of linear scaling toil by ensuring that capacity adjusts automatically as services grow." Why the other options are incorrect:

\* A Manual scaling is linear-scaling toil, not a solution.

\* C Outsourcing development does not reduce operational toil.

\* D Switching cloud providers alone does not solve toil unless automation is introduced.

Thus, B is the correct answer.

References:

Site Reliability Engineering Book, "Eliminating Toil"

SRE Workbook, "Toil Reduction Strategies"

#### NEW QUESTION # 26

Which of the following is the BEST description of a Customer Reliability Engineer (CRE)?

- A. They use deep engineering expertise to improve the cloud provider's services
- B. They work with the cloud provider's SRE team to ship and build new features
- C. They integrate with the customer's operations team to share responsibilities
- D. They take a software engineering approach to redesign all cloud services

Answer: C

Explanation:

#### Comprehensive and Detailed Explanation From Exact Extract:

Customer Reliability Engineering (CRE) is described in Google's SRE literature as an extension of SRE practices outward to customers who run workloads on cloud platforms. The SRE Book and the SRE Workbook state: "CRE is the practice of sharing SRE principles with customers, working closely with their operations teams, and establishing shared responsibility for reliability." (SRE Workbook - Chapter: Customer Reliability Engineering). A CRE team collaborates directly with customer engineering and operations teams to identify reliability risks, review architectures, and co-manage SLOs, but does not redesign cloud services or build new features.

Option D matches the exact intention: CRE integrates with the customer's operations team to share reliability responsibilities, applying SRE methods to customer systems and ensuring both customer and provider work jointly on reliability goals.

Option A is incorrect-CRE does not redesign cloud services.

Option B misinterprets CRE as improving the provider's internal systems.

Option C incorrectly focuses on feature shipping: CRE is about reliability guidance, not feature delivery.

Thus, D is the correct and SRE-authentic answer.

References:

The Site Reliability Workbook, Chapter: "Customer Reliability Engineering" Google Cloud documentation on CRE practices.

#### NEW QUESTION # 27

What is the benefit of strategically burning the Error Budget to zero every month?

- A. It can be revised every month as necessary
- B. It creates a dialog between strategic partners
- **C. It allows a balance between velocity and stability**
- D. It allows for the measurement of capacity and reliability

**Answer: C**

Explanation:

#### Comprehensive and Detailed Explanation From Exact Extract:

Burning the error budget to zero - strategically, not accidentally - helps ensure the correct balance between release velocity and system stability, which is the fundamental purpose of error budgets. Error budgets exist to encourage a healthy level of risk-taking up to the point where user experience is not impacted.

From the Site Reliability Engineering Book, SLO chapter:

"Error budgets provide a mechanism for balancing innovation and reliability by allowing measured risk-taking while ensuring user expectations are met." The SRE Workbook adds:

"Teams should aim to use their full error budget. Not using it implies missed opportunities to deliver features or improvements." This means that strategically burning the error budget to zero ensures:

- \* Teams are shipping value at maximum safe velocity
- \* Reliability goals are still respected
- \* Risk is managed and intentional

Why other options are incorrect:

- \* B Capacity measurement is unrelated to error budget consumption.
- \* C Error budgets should not be continually revised unless business needs change.
- \* D Conversations with partners may occur, but this is not the primary benefit.

Thus, the correct answer is A.

References:

Site Reliability Engineering Book, "Service Level Objectives"

SRE Workbook, "SLO Engineering"

#### NEW QUESTION # 28

.....

The Peoplecert DevOps-SRE is so flexible that you can easily change the timings, types of questions, and topics for each mock exam. Peoplecert DevOps-SRE practice test contains all the important questions that will appear in the actual DevOps-SRE Exam. NewPassLeader offers updates for Peoplecert DevOps-SRE Exam questions up to 365 days after purchase, to match the changes in the latest DevOps-SRE exam syllabus.

**Simulated DevOps-SRE Test:** <https://www.newpassleader.com/Peoplecert/DevOps-SRE-exam-preparation-materials.html>

- Top DevOps-SRE Questions  DevOps-SRE Exam Sample Questions  Valid Exam DevOps-SRE Book  Search

on [www.prepawayexam.com](http://www.prepawayexam.com) for **✓** DevOps-SRE **✓✓** to obtain exam materials for free download **✓** DevOps-SRE Minimum Pass Score