



Download Free PeopleCert DevOps-SRE Exam PDF | PrepBolt

Don't miss out! Download the latest free PeopleCert Site Reliability Engineering Foundation v1.2 exam PDF questions. Access real DevOps-SRE dumps with verified answers and boost your chances to pass your certification on the first try with [PrepBolt](https://prepbolt.com/DevOps-SRE.html) DevOps-SRE exam pdf questions and answers.

Thank you for Downloading DevOps-SRE exam PDF Demo

<https://prepbolt.com/DevOps-SRE.html>

QUESTIONS & ANSWERS
DEMO VERSION
(LIMITED CONTENT)

Question 1

Question Type: MultipleChoice

In a safety culture, engineers are allowed to do more with the production environment without fear of repercussions.

What else do engineers need to do?

Options:

- A- Share production incidents on social media
- B- Be accountable for their actions
- C- Skip all blameless post-mortems
- D- Avoid being on-call

Answer:

B

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

In a safety culture, SRE emphasizes psychological safety so engineers can work effectively in production without fear of blame. However, safety never removes accountability. Engineers must take responsibility for their actions, decisions, and assumptions, particularly during incidents.

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

"Blamelessness does not eliminate accountability. Individuals must still explain the context, assumptions, and reasoning behind their decisions so that the organization can learn."

Google stresses that:

Engineers must feel safe to act and report issues

Engineers must remain responsible and accountable

Accountability enables learning, not punishment

Why other options are incorrect:

A Sharing incidents on social media violates confidentiality

C Blameless postmortems are required, not skipped

D Avoiding on-call is contrary to SRE responsibilities

Thus, B is correct.

Site Reliability Engineering Book, "Postmortem Culture"

SRE Workbook, "Learning from Incidents"

Question 2

Question Type: MultipleChoice

In a blameless post-mortem, those involved report

Options:

- A- Without fear of retribution
- B- Assumptions they had made
- C- Both A and B
- D- Using testing data

Answer:

C

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

A blameless post-mortem is a foundational SRE practice that encourages truthful, detailed reporting after an incident. The purpose is to learn, not punish. Google SRE emphasizes that engineers must feel psychologically safe to report what they did, what they assumed, and why they made those decisions.

From the Site Reliability Engineering Book, Chapter "Postmortem Culture":

"Blameless postmortems encourage engineers to share the full details of their actions and assumptions without fear of punishment, enabling learning and preventing repeated failures."

The book further states:

"Understanding the assumptions made during an incident is critical to uncovering systemic issues."

Thus:

Engineers must report without fear of retribution

They must report assumptions and decisions made during the incident

Therefore, the correct answer is C. Both A and B.

Why the other options are insufficient:

A Only partially correct

B Only partially correct

D Testing data may be included, but it is not the defining feature of blameless postmortems

Site Reliability Engineering Book, "Postmortem Culture"

SRE Workbook, "Learning from Incidents"

Question 3

Question Type: MultipleChoice

If SREs own some sections of a service, but not others, then this organizational approach is known as

Options:

A- Consultant

B- Full SRE

C- Slice and dice

D- Platform

Answer:

C

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

The Slice-and-Dice model is an SRE adoption pattern where the SRE team owns specific portions of a service---typically the most critical, complex, or high-risk components---while development teams own the rest.

From the SRE Workbook, Organizational Models section:

"In the slice-and-dice model, SREs take responsibility for particular portions of a service or system rather than owning the entire thing. This works well when parts of the system require stronger reliability engineering than others."

This model is used when:

Services are large or complex

Only certain components need SRE-level reliability

Full SRE ownership is not feasible

Why the other options are incorrect:

A Consultant SREs advise; they do not own components

B Full SRE fully owns the entire service

D Platform SRE builds shared reliability tooling, not owning service slices

Thus, C. Slice and dice is the correct answer.

SRE Workbook, "SRE Organizational Patterns"

Site Reliability Engineering Book, "Engagement Models"

Question 4

Question Type: MultipleChoice

In which of the following SRE adoption models is reliability a 'first class citizen'?

Options:

A- Consulting

B- Embedded

C- Platform

D- Full

Answer:

D

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

In the Full SRE model, reliability becomes a first-class citizen because SREs own the complete operational responsibility for the service and apply SRE principles end-to-end. The Google SRE Book

describes several adoption models (Consulting, Embedded, and Full), and only the Full SRE model has SREs fully accountable for reliability outcomes.

From the Site Reliability Engineering Book, Chapter "SRE Engagement Models":

"In the Full SRE model, the SRE team is responsible for end-to-end reliability. Reliability becomes a first-class objective, supported through SLOs, error budgets, and systematic reduction of toil."

The Full model includes:

Full ownership of reliability

Enforcement of SLOs

Error budget policies

Engineering-driven improvement

Other models:

Consulting SRE gives guidance but doesn't own reliability

Embedded temporary embedding to train teams, not full ownership

Platform focuses on shared tooling, not service ownership

Thus, D. Full is correct.

Site Reliability Engineering Book, "Engagement Models"

SRE Workbook, "Adopting SRE in Your Organization"

Question 5

Question Type: MultipleChoice

What does the Westrum model study?

Options:

A- The flow of technology in order to deploy faster and more frequently

B- The flow of information and how it affects organizational culture

C- The flow of critical resources necessary to recover in the event of a major incident

D- The flow of communication and how it affects organizational design

Answer:

B

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

The Westrum Organizational Culture Model categorizes organizations based on how effectively they allow information to flow, particularly in high-risk and high-performance environments. This model is widely referenced in SRE and DevOps practices and is cited by Google in their SRE culture guidance.

The Westrum model identifies three culture types:

Pathological (power-oriented)

Bureaucratic (rule-oriented)

Generative (performance-oriented)

From the SRE Workbook and Google's DevOps Research (DORA):

"The Westrum model is a taxonomy of organizational cultures based on how information flows. Generative cultures---with high cooperation and open information flow---correlate with higher reliability and better outcomes."

Why the other options are incorrect:

A Deployment frequency is a DORA metric, not Westrum's focus.

C Resource flow is unrelated to Westrum.

D While communication affects design, Westrum specifically studies information flow as a cultural health indicator, not structural design.

Thus, the correct answer is B.

SRE Workbook, "Cultural Change"

Google DORA Research, "Accelerate State of DevOps Report"

Ron Westrum, "A Typology of Organizational Cultures"

Question 6

Question Type: MultipleChoice

Which of the following is the LEAST useful metric when working to improve antifragility?

Options:

- A- Mean Time To Detect
- B- Service Level Objective
- C- Deployment frequency
- D- Recovery Point Objective

Answer:

C

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

Anti-fragility focuses on an organization's ability to respond, adapt, learn, and recover from incidents. The most useful metrics relate to incident detection, response, reliability, and recovery. Deployment frequency, while important in DevOps and DORA metrics, does not directly measure anti-fragility.

From the SRE Workbook, Incident Response section:

"Improving antifragility requires better detection, better recovery mechanisms, and clear reliability goals."

Key metrics relevant to anti-fragility:

MTTD (Mean Time To Detect) --- quicker detection improves resilience

MTTR/RPO --- recoverability measures

SLOs --- define acceptable reliability thresholds and guide learning

Deployment frequency primarily measures delivery velocity, not resilience.

The Site Reliability Engineering Book emphasizes:

"Antifragility is improved by learning from incidents and strengthening recovery mechanisms rather than by increasing release cadence."

Why other options are correct for anti-fragility:

- A . Mean Time To Detect --- critical for detecting failures quickly
- B . SLOs --- define boundaries for reliability and failure tolerance
- D . Recovery Point Objective --- measures potential loss during failures

Thus, C is the least useful metric for improving antifragility.

SRE Workbook, "Incident Response"

Site Reliability Engineering Book, "Postmortem Culture"

Question 7

Question Type: MultipleChoice

How does chaos engineering as an anti-fragility strategy improve Mean Time to Recover Service?

Options:

- A- It helps to identify weaknesses and dependencies pinpointing areas where more resilience may be required
- B- It optimizes monitoring tools making it more likely we will detect real incidents
- C- It creates automation for auto-recovery
- D- Caching data in the case of a database outage instead could mean the SLO is met

Answer:

A

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

Chaos engineering is an SRE-aligned practice where systems are intentionally subjected to controlled failure scenarios so teams can observe how the system responds. This practice supports anti-fragility, meaning the system becomes stronger through exposure to failure.

The SRE Workbook, Chapter "Handling Overload" and Chaos Engineering sections, explains:

"Injecting failure in a controlled environment exposes the hidden dependencies, weaknesses, and systemic risks that only appear under stress."

The Site Reliability Engineering Book reinforces this concept:

"By understanding how systems behave during partial failures, teams can make targeted improvements that reduce recovery time during real incidents."

Improving Mean Time to Recover (MTTR) happens because:

Weak points and bottlenecks are identified early

Engineers gain familiarity with failure modes

Systems are hardened ahead of actual outages

Dependencies that cause cascading failures are revealed

Why other options are incorrect:

B Monitoring optimization is helpful but not the core mechanism of chaos engineering.

C Chaos engineering does not create auto-recovery automation; it reveals where it is required.

D Caching is an architectural resilience strategy, not an outcome of chaos engineering itself.

Thus, A is the correct answer.

SRE Workbook, "Chaos Engineering"

Site Reliability Engineering Book, "Managing Critical State"

Question 8

Question Type: MultipleChoice

What metrics will embracing failure help to improve?

Options:

A- Mean time to detect and mean time between system incidents

B- Change lead time and change failure rate

C- Empirical test data and mean time to recover service

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.

Site Reliability Engineering Book, "Postmortem Culture"

SRE Workbook, "Incident Response"

Question 9

Question Type: MultipleChoice

If an organization wants to promote changes automatically and reduce dependency errors, what steps should they take?

Options:

- A- Ensure that the artifacts used to deploy system components are tested and visible externally
- B- Ensure that they use only verified and signed artifacts to deploy system components
- C- Ensure that Error Budgets are agreed with oversight and policies
- D- Establish Service Level Objectives that define how artifacts are signed and verified

Answer:

B

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

Using verified and signed artifacts is essential for safe automation, ensuring that deployments are consistent and free of dependency or supply chain errors. This is a fundamental principle in Google's release engineering and SRE practices.

The Site Reliability Engineering Book, chapter "Release Engineering," states:

"Releases should be built once, tested, signed, and stored in a secure repository. Only signed and verified artifacts should be promoted to production to prevent configuration drift and dependency inconsistencies."

The SRE Workbook echoes this:

"Automated promotions depend on the integrity and immutability of artifacts. Signed artifacts ensure consistency and prevent errors related to mismatched dependencies."

Why the other options are incorrect:

A External visibility is irrelevant and may create security risks.

C Error budgets relate to reliability, not artifact promotion.

D SLOs do not define artifact signing; this is handled by release engineering processes.

Thus, the correct answer is B.

Site Reliability Engineering Book, "Release Engineering"

SRE Workbook, "Automation and Safe Releases"

Thank You for trying DevOps-SRE PDF Demo

To try our DevOps-SRE practice exam software
visit link below

<https://prepbolt.com/DevOps-SRE.html>

Start Your DevOps-SRE Preparation

Use Coupon **"SAVE50"** for extra 50% discount on the purchase of Practice Test Software. Test your DevOps-SRE preparation with actual exam questions.