

ISTQB Certified Tester Performance Testing

ISTQB CT-PT

Version Demo

Total Demo Questions: 6

Total Premium Questions: 60

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Basic Concepts of Performance Testing	15
Topic 2, Performance Test Fundamentals	17
Topic 3, Performance Testing Activities	11
Topic 4, Performance Testing Tasks	5
Topic 5, Tools	12
Total	60

QUESTION NO: 1

A claims application is tested with a workload in which many virtual users submit updates to customer records. During the test, response times increase sharply and database monitoring shows substantial lock waits on a small set of customer records. The test data uses the same five customer accounts for every virtual user.

What should be done before deciding that the application has a production locking bottleneck?

- A. Increase the number of virtual users until the database server reaches full processor utilization.
- B. Disable database monitoring to prevent lock-wait measurements influencing the result.
- C. Parameterize the customer data to reflect the expected distribution of account updates.
- D. Replace all update transactions with read-only customer searches.

ANSWER: C

Explanation:

The test data should be reviewed and parameterized so that virtual users use representative, appropriately distributed customer accounts. Reusing a very small set of records can create artificial contention that does not reflect the operational profile.

Lock waits are useful evidence, but they must be interpreted in the context of the data distribution and transaction behaviour. Increasing the number of virtual users or disabling database monitoring would not determine whether the observed contention is realistic.

QUESTION NO: 2

A team plans to generate 80,000 requests per second from a single load generator. A preliminary test against a simple endpoint shows that the load generator CPU is saturated, while the system under test has low resource utilization and the achieved request rate is below target.

What should the team do before drawing conclusions about the capacity of the system under test?

- A. Increase the capacity of the load-generation infrastructure and validate the achieved workload
- B. Conclude that the system under test has sufficient capacity because its utilization is low
- C. Reduce monitoring to make more processor capacity available on the system under test
- D. Classify the result as a system-under-test performance failure

ANSWER: A

Explanation:

The load-generation infrastructure is the limiting component, so the intended workload has not been applied to the system under test. The team should validate and increase load-generation capacity, for example by distributing the workload across additional generators, before assessing system capacity.

Tuning the system under test or treating its low utilization as evidence of spare capacity would be inappropriate because the injector saturation prevents a valid test.

QUESTION NO: 3

A business stakeholder states that an online trading system must be fast for 10,000 concurrent users.

Which of the following is the best way to refine this statement into a performance requirement that can be used for test evaluation?

- A. The system shall provide an average response time of less than two seconds.
- B. The system shall be fast enough to satisfy active traders during busy periods.
- C. The system shall keep processor utilization below 70% during all tests.
- D. At the anticipated peak workload of 10,000 concurrent users and the defined transaction mix, 95% of order-submission transactions shall complete within two seconds.

ANSWER: D

Explanation:

A performance requirement needs measurable criteria and defined operating conditions. Specifying the transaction, percentile response-time target, concurrent workload, and operational mix makes the requirement testable and provides an objective basis for acceptance.

An average response time can hide unacceptable slow transactions, while a statement that the system should be fast remains subjective. A target based only on server utilization does not express the user-facing performance required.

QUESTION NO: 4

What challenge must be considered when using crowds to emulate load generation?

- A. The load generation will be difficult to reproduce.
- B. The load generation method is less sensitive to changes in the system under test.
- C. This type of load generation is more suitable for mainframe applications.
- D. This technique is more precise than other methods of load generation.

ANSWER: A

Explanation:

“The load generation will be difficult to reproduce.” is correct because crowd-based load generation depends on a distributed population of real people using their own devices, networks, locations, browsers, and operating conditions. Unlike a controlled load-generator environment, these conditions cannot be held constant from one execution to the next. The number of available participants, their timing, connection quality, device capability, and interaction patterns may all vary. Consequently, a repeated test may create a materially different workload even when the intended scenario is the same.

This lack of repeatability is an important consideration when interpreting measurements and comparing test runs. Performance testing requires test conditions and workload profiles to be sufficiently controlled so that observed response-time, throughput, resource-utilization, or error-rate differences can be attributed meaningfully to a system change rather than to variation in the load source itself. Crowd-based approaches can provide useful realism, especially for geographically distributed usage, but the test team must account for reduced control and reproducibility when planning the test and evaluating results. This aligns with the performance-testing concepts and guidance covered by the [ISTQB Certified Tester Performance Testing](#) certification and the terminology maintained in the [ISTQB Glossary](#).

QUESTION NO: 5

A test is planned to assess an application at 5,000 concurrent users for one hour. During execution, authentication failures prevent 1,200 virtual users from completing the login process. The remaining virtual users continue to execute the intended transaction mix.

What should the performance test report state about the test result?

- A. The system met the objective because the active virtual users completed the intended transaction mix.
- B. The result should be reported as evidence for 5,000 concurrent users because the test duration was one hour.

C. The report should record the achieved workload and authentication failures, and state that the 5,000-user objective was not assessed.

D. The response-time results should be extrapolated to estimate behaviour for the missing virtual users.

ANSWER: C

Explanation:

The report should state that the planned workload was not achieved, document the authentication failures and the actual achieved workload, and avoid claiming that the 5,000-user objective was met. The result may still provide diagnostic evidence, but it is not valid acceptance evidence for the intended workload.

Reporting only the measurements from the active users would conceal an important deviation from the test conditions. Estimating results for the missing users is not a substitute for executing a controlled test at the required workload.

QUESTION NO: 6

A test script records a customer identifier returned when a virtual user creates an order. The script then uses the identifier in a later request to retrieve the order status. When many virtual users execute the script, several requests retrieve the wrong order or fail.

Which action is most appropriate before using the script for load generation?

A. Correlate the dynamic identifier between the related requests

B. Increase think time between order creation and order retrieval

C. Remove the order-status request from the script

D. Run the script with more virtual users to confirm the failures

ANSWER: A

Explanation:

The customer or order identifier is dynamic and must be captured separately for each virtual user and reused in that user's subsequent request. This is correlation. Without it, the script may replay a recorded value or use an incorrect value, causing failures that invalidate the workload and the resulting measurements.

Increasing think time or virtual-user numbers does not correct the dependency. Replacing result verification would conceal rather than resolve the script defect.